



ExtremeXOS Command Reference Guide for Release 15.4

Copyright © 2001-2013 Extreme Networks

AccessAdapt, Alpine, Altitude, BlackDiamond, Direct Attach, EPICenter, ExtremeWorks Essentials, Ethernet Everywhere, Extreme Enabled, Extreme Ethernet Everywhere, Extreme Networks, Extreme Standby Router Protocol, Extreme Turbodrives, Extreme Velocity, ExtremeWare, ExtremeWorks, ExtremeXOS, Go Purple Extreme Solution, ExtremeXOS ScreenPlay, ReachNXT, Ridgeline, Sentrant, ServiceWatch, Summit, SummitStack, Triumph, Unified Access Architecture, Unified Access RF Manager, UniStack, XNV, the Extreme Networks logo, the Alpine logo, the BlackDiamond logo, the Extreme Turbodrives logo, the Summit logos, and the Powered by ExtremeXOS logo are trademarks or registered trademarks of Extreme Networks, Inc. or its subsidiaries in the United States and/or other countries.

sFlow is the property of InMon Corporation.

iBooks is property of Apple, Inc.

Specifications are subject to change without notice.

All other registered trademarks, trademarks, and service marks are property of their respective owners.

For additional information on Extreme Networks trademarks, please see:
www.extremenetworks.com/about-extreme/trademarks.aspx.

120886-00 Rev 1.0

Table of Contents

Chapter 1: Introduction to the ExtremeXOS Command Reference Guide 57

- Conventions 57
- Related Publications 59
- Providing Feedback to Us 59

Chapter 2: Command Reference Overview 60

- Structure of this Guide 60
- Platforms and Required Software Versions 61
- Software Required 62
- Understanding the Command Syntax 64
- Port Numbering 67
- Line-Editing Keys 68
- Command History 68

Chapter 3: Commands for Accessing the Switch 70

- clear account lockout 72
- clear session 73
- configure account 74
- configure account encrypted 75
- configure account password-policy char-validation 77
- configure account password-policy history 78
- configure account password-policy lockout-on-login-failures 79
- configure account password-policy max-age 80
- configure account password-policy min-length 81
- configure banner 83
- configure cli max-sessions 84
- configure cli max-failed-logins 85
- configure dns-client add 86
- configure dns-client default-domain 87
- configure dns-client delete 88
- configure failsafe-account 89
- configure idletimeout 91
- configure safe-default-script 91
- configure time 93
- configure timezone 94
- create account 98
- delete account 100
- disable cli prompting 101
- disable cli space-completion 102
- disable clipaging 103
- disable idletimeout 104
- enable cli prompting 105
- enable cli space-completion 106
- enable clipaging 106
- enable idletimeout 107
- history 108
- ping 109
- reboot 111



show accounts 113
show accounts password-policy 115
show banner 116
show banner 117
show failsafe-account 119
show switch 120
traceroute 123
unconfigure banner 124

Chapter 4: Commands for Managing the Switch 126

NTP 130
PTP 130
SNMP 130
Telnet 131
TFTP 131
System Redundancy with Dual Management Modules Installed, Modular Switches Only 132
Power Supply Management 132
SNTP 132
clear cdp counters 133
clear cdp neighbor 133
clear network-clock ptp counters 134
configure cdp device-id 135
configure cdp frequency 136
configure cdp hold-time 137
configure node priority 138
configure ntp local-clock none 139
configure ntp local-clock stratum 140
configure ntp restrict-list 140
configure ntp server/peer add 141
configure ntp server/peer delete 142
configure power led motion-detector 143
configure power monitor 144
configure power supply 145
configure snmp access-profile 148
configure snmp add community 150
configure snmp add trapreceiver 152
configure snmp delete community 153
configure snmp delete trapreceiver 155
configure snmp sysContact 156
configure snmp sysLocation 157
configure snmp sysName 158
configure snmpv3 add access 159
configure snmpv3 add community 160
configure snmpv3 add filter 162
configure snmpv3 add filter-profile 164
configure snmpv3 add group user 165
configure snmpv3 add mib-view 166
configure snmpv3 add notify 167
configure snmpv3 add target-addr 169
configure snmpv3 add target-params 171



configure snmpv3 add user	172
configure snmpv3 add user clone-from	174
configure snmpv3 delete access	176
configure snmpv3 delete community	177
configure snmpv3 delete filter	178
configure snmpv3 delete filter-profile	179
configure snmpv3 delete group user	180
configure snmpv3 delete mib-view	182
configure snmpv3 delete notify	183
configure snmpv3 delete target-addr	184
configure snmpv3 delete target-params	185
configure snmpv3 delete user	186
configure snmpv3 engine-boots	187
configure snmpv3 engine-id	188
configure snmpv3 target-addr retry	189
configure snmpv3 target-addr timeout	190
configure snmp-client	191
configure snmp-client update-interval	192
configure ssh2 access-profile	192
configure telnet access-profile	195
configure telnet port	197
configure telnet vr	198
configure web http access-profile	199
create network-clock ptp	201
create ntp key	203
create snmp trap	204
delete network-clock ptp	205
delete ntp key	206
disable auto-provision	207
disable cdp ports	208
disable dhcp vlan	208
disable network-clock ptp	209
disable network-clock ptp boundary unicast-negotiation	210
disable ntp broadcast-server	211
disable snmp access	212
disable snmp access vr	213
disable snmp community	214
disable snmp traps	215
disable snmpv3	216
disable snmp-client	217
disable telnet	217
disable watchdog	218
enable auto-provision	219
disable ntp	220
disable ntp authentication	221
disable ntp broadcast-client	222
disable ntp broadcast-server	223
disable ntp vlan	224
enable cdp ports	224



enable dhcp vlan 225
enable network-clock ptp 226
enable network-clock ptp unicast-negotiation 227
enable network-clock ptp end-to-end transparent 228
enable snmp access 229
enable snmp access vr 230
enable snmp community 231
enable snmp traps 232
enable snmpv3 233
enable snmp-client 234
enable telnet 235
enable watchdog 236
enable ntp 237
enable ntp authentication 238
enable ntp broadcast-client 238
enable ntp broadcast-server 239
enable ntp vlan 240
exit 241
logout 242
quit 243
show access-list counters process 244
show auto-provision 245
show checkpoint-data 246
show dhcp-client state 248
show management 249
show network-clock ptp 252
show network-clock ptp (datasets) 253
show network-clock ptp (interface) 256
show network-clock ptp end-to-end-transparent ports 257
show network-clock ptp boundary unicast-master 259
show network-clock ptp boundary unicast-slave 260
show network-clock ptp counters 261
show node 263
show ntp 265
show ntp association 266
show ntp association statistics 267
show ntp key 268
show ntp restrict-list 269
show ntp server 270
show ntp sys-info 271
show ntp vlan 272
show odometers 273
show power 275
show power budget 281
show power controller 284
show power led motion-detector 286
show session 287
show snmp 289
show snmp vr_name 290



- show snmpv3 access 291
- show snmpv3 community 293
- show snmpv3 context 295
- show snmpv3 counters 295
- show snmpv3 engine-info 297
- show snmpv3 extreme-target-addr-ext 298
- show snmpv3 filter 299
- show snmpv3 filter-profile 300
- show snmpv3 group 301
- show snmpv3 mib-view 303
- show snmpv3 notify 305
- show snmpv3 target-addr 307
- show snmpv3 target-params 308
- show snmpv3 user 309
- show snmp-client 311
- telnet 312
- telnet msm 314
- telnet slot 315
- tftp 316
- tftp get 321
- tftp put 323

Chapter 5: Commands for Managing the ExtremeXOS Software 325

- cp 326
- disable xml-mode 329
- enable xml-mode 330
- ls 331
- mv 334
- restart process 337
- rm 339
- show heartbeat process 342
- show memory 344
- show memory process 349
- show process 351
- start process 358
- terminate process 360

Chapter 6: SummitStack Feature Commands 363

- configure stacking alternate-ip-address 364
- configure stacking easy-setup 365
- configure stacking license-level 368
- configure stacking mac-address 370
- configure stacking master-capability 371
- configure stacking priority 373
- configure stacking protocol 374
- configure stacking redundancy 375
- configure stacking slot-number 377
- configure stacking-support stack-ports 378
- disable stacking 382
- disable stacking-support 383



- enable stacking 384
- enable stacking-support 386
- show power (Stack Nodes Only) 387
- show stacking 388
- show stacking configuration 391
- show stacking detail 393
- show stacking stack-ports 395
- show stacking-support 397
- synchronize stacking 399
- unconfigure stacking 400
- unconfigure stacking alternate-ip-address 402
- unconfigure stacking license-level 403
- unconfigure stacking-support 404

Chapter 7: Commands for Configuring Slots and Ports on a Switch 406

- clear counters ports 413
- clear lacp counters 414
- clear counters edp 415
- clear slot 415
- configure ces add peer ipaddress 416
- configure ces add peer ipaddress fec-id-type pseudo-wire 418
- configure ces add peer mac-address 419
- configure ces add service 420
- configure ces delete peer 421
- configure ces delete service 422
- configure ces delete peer 423
- configure ces filler-pattern 423
- configure ces jitter-buffer 424
- configure ces lops-threshold 425
- configure ces payload-size 426
- configure ces peer ipaddress static-pw 427
- configure ces qosprofile 428
- configure ces ttl 429
- configure edp advertisement-interval 429
- configure forwarding external-tables 430
- configure forwarding switching-mode 432
- configure ip-fix domain 433
- configure ip-fix flow-key ipv4 434
- configure ip-fix flow-key ipv6 435
- configure ip-fix flow-key nonip 437
- configure ip-fix ip-address 438
- configure ip-fix ports 439
- configure ip-fix ports flow-key ipv4 mask ipaddress 440
- configure ip-fix ports flow-key ipv6 mask ipaddress 441
- configure ip-fix ports record 441
- configure ip-fix source ip-address 442
- configure ip-mtu vlan 443
- configure jumbo-frame-size 445
- configure lacp member-port priority 446
- configure mirror add ports anomaly 447



configure mirror add	447
configure mirror delete	449
configure mirror description	450
configure mirror name	451
configure mirror to port	452
configure mlag peer interval	453
configure mlag peer ipaddress	454
configure MLAG peer lacp-mac	455
configure mlag ports convergence-control	456
configure network-clock clock-source input	457
configure network-clock clock-source output	458
configure network-clock ptp (priority)	459
configure network-clock ptp announce interval	461
configure network-clock ptp announce timeout	462
configure network-clock ptp boundary add vlan	463
configure network-clock ptp boundary add unicast-slave	465
configure network-clock ptp boundary delete unicast-slave	466
configure network-clock ptp delete	467
configure network-clock ptp delay-request-interval	468
configure network-clock ptp end-to-end transparent	469
configure network-clock ptp ordinary add	470
configure network-clock ptp sync-interval	471
configure network-clock ptp add unicast-master	473
configure network-clock ptp delete unicast-master	474
configure network-clock sync-e	475
configure network-clock sync-e clock-source	476
configure port description-string	477
configure ports auto off	478
configure ports auto on	480
configure ports auto-polarity	482
configure ports display-string	483
configure ports dwdm channel	484
configure ports dwdm channel none	486
configure ports eee enable	487
configure ports far-end-fault-indication	489
configure ports isolation	490
configure ports link-scan interval	490
configure ports mode	492
configure ports partition	492
configure ports preferred-medium	494
configure ports redundant	495
configure ports tdm cable-length	496
configure ports tdm clock-source	498
configure ports tdm display-string	499
configure ports tdm framing	500
configure ports tdm idle-code	501
configure ports tdm line-coding	502
configure ports tdm recovered-clock	503
configure ports tdm signaling	504



configure ports tdm trunk-conditioning 505
configure ports wan-phy clocking 506
configure ports wan-phy framing 506
configure ports wan-phy loopback 507
configure ports wan-phy trace-path 508
configure ports wan-phy trace-section 509
configure sharing add ports 510
configure sharing address-based custom 512
configure sharing delete ports 513
configure sharing health-check member-port add tcp-tracking 514
configure sharing health-check member-port delete tcp-tracking 516
configure sharing health-check member-port tcp-tracking 516
configure sharing lacp activity-mode 517
configure sharing lacp defaulted-state-action 518
configure sharing lacp system-priority 520
configure sharing lacp timeout 521
configure sharing port-based key 522
configure slot module 523
configure slot restart-limit 524
configure tdm hierarchy 525
configure tdm service circuit add port 526
configure tdm service circuit delete port 528
configure tdm service circuit seized-code 529
create ces psn 530
create mirror to port 530
create mlag peer 531
create tdm service circuit 532
delete ces 533
delete mirror name 534
delete mlag peer 535
delete tdm service circuit 536
disable ces 537
disable edp ports 537
disable flow-control ports 538
disable ip-fix ports 540
disable jumbo-frame ports 541
disable learning port 542
disable mirror 543
disable mlag port 544
disable network-clock ptp end-to-end-transparent ports 545
disable network-clock sync-e 545
disable port 546
disable ports tdm 547
configure slot module 548
disable slot 549
disable smartredundancy 551
disable snmp traps port-up-down ports 552
enable ces 553
enable | disable ces peer ipaddress 553



- enable edp ports 554
- enable flow-control ports 556
- enable | disable ip-fix 558
- enable ip-fix ports 559
- enable jumbo-frame ports 560
- enable learning port 561
- enable mirror 562
- enable mlag port peer id 562
- enable network-clock ptp 564
- enable network-clock ptp end-to-end-transparent ports 565
- enable network-clock sync-e 565
- enable port 567
- enable ports tdm 567
- enable ports tdm loopback 568
- enable sharing grouping 569
- enable slot 574
- enable smartredundancy 576
- enable snmp traps port-up-down ports 576
- restart ports 577
- run failover 578
- run msm-failover 579
- show cdp 580
- show cdp counters 581
- show cdp neighbor 582
- show cdp ports 583
- show ces 584
- show ces clock-recovery 586
- show ces errors 587
- show ces peer 590
- show dwdm channel-map 591
- show edp 592
- show ip-fix 594
- show lacp 595
- show lacp counters 597
- show lacp lag 598
- show lacp member-port 603
- show mirror 605
- show mirroring 607
- show mlag peer 609
- show mlag ports 611
- show network-clock clock-source 613
- show network-clock sync-e ports 614
- show port eee 616
- show ports 618
- show ports anomaly 621
- show ports buffer 622
- show ports collisions 624
- show ports configuration 626
- show port information 628



- show ports ip-fix 636
- show ports link-scan 637
- show ports packet 638
- show ports redundant 640
- show ports sharing 641
- show ports tdm alarms 643
- show ports tdm configuration 644
- show ports tdm errors 645
- show ports tdm information 647
- show ports tdm no-refresh 650
- show ports transceiver information 651
- show ports transceiver information detail 653
- show ports utilization 657
- show ports wan-phy configuration 660
- show ports wan-phy errors 662
- show ports wan-phy events 663
- show ports wan-phy overhead 665
- show sharing distribution port-based 666
- show sharing health-check 667
- show sharing port-based keys 668
- show slot 670
- show tdm hierarchy 677
- show tdm service 678
- unconfigure ip-fix 679
- unconfigure ip-fix flow-key 680
- unconfigure ip-fix ip-address 681
- unconfigure ip-fix ports 682
- unconfigure ip-fix ports flow-key mask 682
- unconfigure ip-fix source ip-address 683
- unconfigure mlag peer interval 684
- unconfigure mlag peer ipaddress 685
- unconfigure network-clock sync-e 686
- unconfigure network-clock sync-e clock-source 687
- unconfigure port description-string 687
- unconfigure ports display string 688
- unconfigure ports redundant 689
- unconfigure ports tdm display string 690
- unconfigure ports tdm recovered-clock 690
- unconfigure ports wan-phy 691
- unconfigure slot 692

Chapter 8: Universal Port Management Commands 694

- configure log target upm filter 695
- configure log target upm match 696
- configure upm event 697
- configure upm profile maximum execution-time 698
- configure upm timer after 698
- configure upm timer at 700
- configure upm timer profile 701
- create log target upm 701



- create upm profile 702
- create upm timer 704
- delete log target upm 704
- delete upm profile 705
- delete upm timer 706
- disable log target upm 707
- disable upm profile 708
- edit upm profile 708
- enable log target upm 709
- enable upm profile 710
- run upm profile 711
- show log configuration target upm 711
- show upm events 712
- show upm history 713
- show upm history exec-id 714
- show upm profile 715
- show upm timers 717
- unconfigure upm event 718
- unconfigure upm timer 719

Chapter 9: CLI Scripting Commands 721

- configure cli mode 721
- configure cli mode scripting 723
- configure cli script timeout 723
- delete var 725
- delete var key 726
- disable cli scripting 727
- disable cli scripting output 727
- ELSE 728
- enable cli scripting 729
- enable cli scripting output 730
- ENDIF 731
- ENDWHILE 732
- IF ... THEN 734
- load var key 735
- return 736
- save var key 737
- set var 738
- show var 739
- WHILE ... DO 740

Chapter 10: Commands for Configuring LLDP 742

- clear lldp neighbors 743
- configure lldp med fast-start repeat-count 744
- configure lldp ports management-address 745
- configure lldp ports port-description 746
- configure lldp ports system-capabilities 746
- configure lldp ports system-description 747
- configure lldp ports system-name 748
- configure lldp ports vendor-specific avaya-extreme call-server 749



configure lldp ports vendor-specific avaya-extreme dot1q-framing 750
 configure lldp ports vendor-specific avaya-extreme file-server 752
 configure lldp ports vendor-specific avaya-extreme poe-conservation-request 753
 configure lldp ports vendor-specific dot1 port-vlan-ID 754
 configure lldp ports vendor-specific dot1 port-protocol-vlan-ID 755
 configure lldp ports vendor-specific dot1 vlan-name 756
 configure lldp ports vendor-specific dot3 link-aggregation 758
 configure lldp ports vendor-specific dot3 mac-phy 758
 configure lldp ports vendor-specific dot3 max-frame-size 759
 configure lldp ports vendor-specific dot3 power-via-mdi 760
 configure lldp ports vendor-specific med capabilities 762
 configure lldp ports vendor-specific med location-identification 763
 configure lldp ports vendor-specific med policy application 765
 configure lldp ports vendor-specific med power-via-mdi 767
 configure lldp reinitialize-delay 769
 configure lldp snmp-notification-interval 769
 configure lldp transmit-delay 770
 configure lldp transmit-hold 771
 configure lldp transmit-interval 772
 disable lldp ports 773
 disable snmp traps lldp 774
 disable snmp traps lldp-med 774
 enable lldp ports 775
 enable snmp traps lldp 776
 enable snmp traps lldp-med 777
 show lldp 778
 show lldp neighbors 780
 show lldp statistics 783
 unconfigure lldp 785

Chapter 11: Commands for OAM 787

clear counters bfd 789
 clear counters cfm segment <segment_name> 790
 clear counters cfm segment all 792
 clear counters cfm segment all frame-delay 795
 clear counters cfm segment all frame-loss 797
 clear counters cfm segment frame-delay 800
 clear counters cfm segment frame-loss 802
 clear counters cfm segment frame-loss mep 804
 configure bfd vlan 805
 configure bfd vlan authentication 806
 configure cfm domain add association integer 807
 configure cfm domain add association string 808
 configure cfm domain add association vlan-id 809
 configure cfm domain add association vpn-id oui index 810
 configure cfm domain association add 811
 configure cfm domain association add remote-mep 813
 configure cfm domain association delete 814
 configure cfm domain association delete remote-mep 815
 configure cfm domain association destination-mac-type 816



configure cfm domain association end-point add group 817
configure cfm domain association ports end-point ccm 818
configure cfm domain association end-point delete group 819
configure cfm domain association ports end-point mepid 819
configure cfm domain association ports end-point sender-id-ipaddress 821
configure cfm domain association end-point transmit-interval 822
configure cfm domain association ports end-point 823
configure cfm domain association remote-mep mac-address 824
configure cfm domain delete association 825
configure cfm domain md-level 825
configure cfm group add rmep 826
configure cfm group delete rmep 827
configure cfm segment add domain association 828
configure cfm segment delete domain association 829
configure cfm segment dot1p 829
configure cfm segment frame-delay dot1p 830
configure cfm segment frame-delay/frame-loss transmit interval 831
configure cfm segment frame-delay window 832
configure cfm segment frame-loss dot1p 833
configure cfm segment frame-loss window 834
configure cfm segment frame-loss mep 835
configure cfm segment frame-loss consecutive 836
configure cfm segment frame-loss ses-threshold 836
configure cfm segment threshold 837
configure cfm segment timeout 838
configure cfm segment transmit-interval 839
configure cfm segment window 840
create cfm domain dns md-level 841
create cfm domain mac md-level 842
create cfm domain string md-level 843
create cfm segment destination 844
delete cfm domain 845
delete cfm segment 846
disable cfm segment frame-delay measurement 847
disable cfm segment frame-loss measurement mep 848
enable/disable bfd vlan 849
enable cfm segment frame-delay measurement 849
enable cfm segment frame-loss measurement mep 851
ping mac port 852
show bfd 853
show bfd counters 854
show bfd session client 855
show bfd session counters vr all 856
show bfd session detail vr all 857
show bfd session vr all 859
show bfd vlan 860
show bfd vlan counters 861
show cfm 862
show cfm detail 864



show cfm groups 866
show cfm segment 867
show cfm segment frame-delay 869
show cfm segment frame-delay/frame-loss mep id 869
show cfm segment frame-delay statistics 872
show cfm segment frame-loss 873
show cfm segment frame-loss statistics 874
show cfm segment mep 876
traceroute mac port 878
unconfigure bfd vlan 880
unconfigure cfm domain association end-point transmit-interval 881

Chapter 12: PoE Commands 883

Extreme Networks PoE Devices 884
Summary of PoE Software Features 884
clear inline-power stats ports 885
configure inline-power budget 886
configure inline-power disconnect-precedence 887
configure inline-power label ports 889
configure inline-power operator-limit ports 890
configure inline-power priority ports 891
configure inline-power usage-threshold 892
disable inline-power 893
disable inline-power legacy 894
disable inline-power legacy slot 895
disable inline-power ports 896
disable inline-power slot 897
enable inline-power 898
enable inline-power legacy 899
enable inline-power legacy slot 900
enable inline-power ports 902
enable inline-power slot 903
reset inline-power ports 904
show inline-power 905
show inline-power configuration ports 906
show inline-power info ports 907
show inline-power slot 910
show inline-power stats 912
show inline-power stats ports 913
show inline-power stats slot 915
unconfigure inline-power budget slot 916
unconfigure inline-power disconnect-precedence 917
unconfigure inline-power operator-limit ports 918
unconfigure inline-power priority ports 918
unconfigure inline-power usage-threshold 919

Chapter 13: Commands for Status Monitoring and Statistics 921

Event Management System 923
Extreme Link Status Monitoring 924
sFlow Statistics 925



RMON 925

- clear counters 926
- clear counters xml-notification 927
- clear cpu-monitoring 928
- clear elsm ports auto-restart 929
- clear elsm ports counters 930
- clear log 931
- clear log counters 932
- clear sys-recovery-level 933
- configure elsm ports hellotime 935
- configure elsm ports hold-threshold 936
- configure elsm ports uptimer-threshold 937
- configure log display 938
- configure log filter events 939
- configure log filter events match 942
- configure log target filter 946
- configure log target format 949
- configure log target match 954
- configure log target severity 956
- configure log target syslog 958
- configure log target xml-notification filter 959
- configure ports monitor vlan 960
- configure sflow agent ipaddress 962
- configure sflow collector ipaddress 963
- configure sflow max-cpu-sample-limit 964
- configure sflow poll-interval 965
- configure sflow ports sample-rate 966
- configure sflow sample-rate 967
- configure sys-health-check all level 968
- configure sys-health-check interval 971
- configure sys-recovery-level 972
- configure sys-recovery-level slot 974
- configure sys-recovery-level switch 979
- configure syslog add 981
- configure syslog delete 983
- configure xml-notification target 984
- configure xml-notification target add/delete 985
- create log filter 986
- create log target xml-notification 987
- create xml-notification target url 988
- delete log filter 989
- delete log target xml-notification 990
- delete xml-notification target 991
- disable cli-config-logging 991
- disable cpu-monitoring 992
- disable elsm ports 993
- disable elsm ports auto-restart 994
- disable log display 996
- disable log target 997



disable log target xml-notification	998
disable rmon	999
disable sflow	1000
disable sflow ports	1001
disable sys-health-check	1002
disable syslog	1003
enable cli-config-logging	1004
enable cpu-monitoring	1005
enable elsm ports	1006
enable elsm ports auto-restart	1008
enable log display	1009
enable log target	1010
enable log target xml-notification	1012
enable rmon	1013
enable sflow	1015
enable sflow ports	1016
enable sys-health-check	1017
enable syslog	1019
enable/disable xml-notification	1019
show configuration "xmlc"	1020
show cpu-monitoring	1022
show elsm	1026
show elsm ports	1028
show fans	1032
show log	1036
show log components	1040
show log configuration	1044
show log configuration filter	1046
show log configuration target	1047
show log configuration target xml-notification	1049
show log counters	1050
show log events	1053
show ports rxerrors	1054
show ports statistics	1057
show ports txerrors	1061
show ports vlan statistics	1064
show rmon memory	1065
show sflow configuration	1068
show sflow statistics	1070
show temperature	1071
show version	1074
show vlan statistics	1079
show xml-notification configuration	1080
show xml-notification statistics	1082
unconfigure log filter	1083
unconfigure log target format	1084
unconfigure ports monitor vlan	1085
unconfigure sflow	1086
unconfigure sflow agent	1087



unconfigure sflow collector 1088
unconfigure sflow ports 1089
unconfigure xml-notification 1090
upload log 1091

Chapter 14: VLAN Commands 1094

configure private-vlan add network 1095
configure private-vlan add subscriber 1095
configure private-vlan delete 1096
configure protocol add 1097
configure protocol delete 1099
configure vlan add ports 1100
configure vlan add ports private-vlan translated 1102
configure vlan delete ports 1102
configure vlan description 1104
configure vlan ipaddress 1104
configure vlan name 1106
configure vlan protocol 1107
configure vlan tag 1108
configure vlan-translation add loopback-port 1109
configure vlan-translation add member-vlan 1110
configure vlan-translation delete loopback-port 1111
configure vlan-translation delete member-vlan 1112
create private-vlan 1113
create protocol 1114
create vlan 1115
delete private-vlan 1117
delete protocol 1118
delete vlan 1119
disable loopback-mode vlan 1119
disable vlan 1120
enable loopback-mode vlan 1121
enable vlan 1122
show private-vlan 1123
show private-vlan <name> 1125
show protocol 1126
show vlan 1127
show vlan description 1134
unconfigure vlan description 1135
unconfigure vlan ipaddress 1135

Chapter 15: VMAN (PBN) Commands 1137

configure port ethertype 1138
configure vman add ports 1138
configure vman add ports cep 1141
configure vman delete ports 1143
configure vman ethertype 1143
configure vman ports add cvid 1144
configure vman ports delete cvid 1146
configure vman tag 1147



- create vman 1148
- delete vman 1149
- disable dot1p examination inner-tag ports 1150
- disable vman cep egress filtering ports 1151
- enable dot1p examination inner-tag port 1152
- enable vman cep egress filtering ports 1153
- show vman 1154
- show vman eaps 1157
- show vman ethertype 1157
- unconfigure vman ethertype 1159

Chapter 16: FDB Commands 1161

- clear fdb 1161
- configure fdb mac-tracking ports 1162
- configure fdb static-mac-move packets 1163
- create fdbentry vlan ports 1164
- delete fdb mac-tracking entry 1166
- delete fdbentry 1167
- disable fdb static-mac-move 1168
- disable flooding ports 1169
- disable learning iparp sender-mac 1171
- disable learning port 1171
- disable snmp traps fdb mac-tracking 1172
- enable fdb static-mac-move 1173
- enable flooding ports 1174
- enable learning iparp sender-mac 1176
- enable learning port 1176
- enable snmp traps fdb mac-tracking 1177
- show fdb 1178
- show fdb mac-tracking configuration 1181
- show fdb mac-tracking statistics 1181
- show fdb static-mac-move configuration 1183
- show fdb stats 1184

Chapter 17: Data Center Solution Commands 1187

- show vm-tracking repository 1188
- configure fip snooping add vlan 1189
- configure fip snooping add fcf 1190
- configure fip snooping delete vlan 1192
- configure fip snooping delete fcf 1193
- configure fip snooping fcf-update 1194
- configure fip snooping fcmmap 1195
- configure fip snooping port location 1197
- configure lldp ports dcbx add application 1198
- configure lldp ports dcbx delete application 1199
- configure lldp ports vendor-specific dcbx 1201
- configure port reflective-relay 1202
- configure vlan dynamic-vlan uplink-ports 1202
- configure vm-tracking authentication database-order 1203
- configure vm-tracking blackhole 1204



configure vm-tracking local-vm 1205
configure vm-tracking nms 1207
configure vm-tracking nms timeout 1208
configure vm-tracking repository 1208
configure vm-tracking timers 1209
configure vm-tracking vpp add 1210
configure vm-tracking vpp delete 1211
configure vm-tracking vpp vlan-tag 1212
configure vm-tracking vpp counters 1213
create vm-tracking local-vm 1214
create vm-tracking vpp 1216
delete vm-tracking local-vm 1216
delete vm-tracking vpp 1217
disable fip snooping 1218
disable vm-tracking 1219
disable vm-tracking dynamic-vlan ports 1220
disable vm-tracking ports 1221
enable fip snooping 1221
enable vm-tracking 1223
enable vm-tracking dynamic-vlan ports 1223
enable vm-tracking ports 1224
run vm-tracking repository 1225
show fip snooping access-list 1226
show fip snooping counters 1228
show fip snooping enode 1229
show fip snooping fcf 1231
show fip snooping virtual-link 1232
show fip snooping vlan 1233
show lldp dcbx 1234
show vlan dynamic-vlan 1240
show vm-tracking 1241
show vm-tracking local-vm 1243
show vm-tracking network-vm 1244
show vm-tracking nms 1245
show vm-tracking port 1246
show vm-tracking repository 1247
show vm-tracking vpp 1248
unconfigure vm-tracking local-vm 1249
unconfigure vm-tracking repository 1250
unconfigure vm-tracking vpp vlan-tag 1251
unconfigure vm-tracking vpp 1252
unconfigure vm-tracking nms 1252

Chapter 18: AVB Commands 1254

clear msrp counters 1255
clear mvrp counters 1256
clear network-clock gptp counters 1257
configure mrp ports timers 1258
configure msrp latency-max-frame-size 1259
configure msrp ports sr-pvid 1260



configure msrp ports traffic-class delta-bandwidth	1261
configure msrp timers first-value-change-recovery	1262
configure mvrp stpd	1263
configure mvrp tag ports registration	1264
configure mvrp tag ports transmit	1265
configure mvrp vlan auto-creation	1266
configure mvrp vlan registration	1267
configure network-clock gtp default-set	1268
configure network-clock gtp ports announce	1269
configure network-clock gtp ports peer-delay	1270
configure network-clock gtp ports sync	1272
disable avb	1273
disable avb ports	1274
disable msrp	1275
disable msrp ports	1276
disable mvrp	1276
disable mvrp ports	1277
disable network-clock gtp	1278
disable network-clock gtp ports	1279
enable avb	1280
enable avb ports	1280
enable msrp	1281
enable msrp ports	1282
enable mvrp	1283
enable mvrp ports	1284
enable network-clock gtp	1285
enable network-clock gtp ports	1285
show avb	1286
show mrp ports	1287
show msrp	1288
show msrp listeners	1289
show msrp ports	1291
show msrp ports bandwidth	1293
show msrp ports counters	1294
show msrp streams	1296
show msrp talkers	1298
show mvrp	1300
show mvrp ports counters	1301
show mvrp tag	1302
show network-clock gtp	1304
show network-clock gtp ports	1306
unconfigure avb	1309
unconfigure mrp ports timers	1310
unconfigure msrp	1311
unconfigure mvrp	1312
unconfigure mvrp stpd	1313
unconfigure mvrp tag	1314
unconfigure network-clock gtp ports	1314

Chapter 19: TRILL Commands 1316



clear trill counters 1316
configure trill add access tag 1317
configure trill add network vlan 1318
configure trill appointed-forwarder 1320
configure trill delete access tag 1321
configure trill delete network vlan 1323
configure trill designated-vlan 1324
configure trill inhibit-time 1325
configure trill mtu probe 1326
configure trill mtu probe fail-count 1327
configure trill mtu size 1328
configure trill nickname 1329
configure trill ports 1330
configure trill ports protocol 1331
configure trill pseudonode 1332
configure trill system-id 1333
configure trill timers csnp 1334
configure trill timers hello 1335
configure trill timers lsp 1336
configure trill timers spf 1337
configure trill timers spf backoff-delay 1338
configure trill tree prune vlan 1339
create trill nickname 1340
delete trill nickname 1342
disable trill 1343
enable trill 1344
show trill 1345
show trill distribution-tree 1347
show trill lsdb 1348
show trill neighbor 1349
show trill ports 1350
show trill rbridges 1352
traceroute trill 1353

Chapter 20: Commands for Virtual Routers 1355

clear counters vr 1355
configure vr add ports 1356
configure vr add protocol 1357
configure vr delete ports 1359
configure vr delete protocol 1360
configure vr description 1361
configure vr rd 1362
configure vr route-target 1364
configure vr vpn-id 1365
create virtual-router 1366
delete virtual-router 1368
disable snmp trap l3vpn 1369
disable virtual-router 1370
enable snmp trap l3vpn 1371
enable virtual-router 1372



show counters vr 1373
show virtual-router 1374
unconfigure vr description 1380
unconfigure vr rd 1381
unconfigure vr vpn-id 1382
virtual-router 1382

Chapter 21: Policy Manager Commands 1385

check policy 1385
check policy attribute 1386
edit policy 1388
refresh policy 1389
show policy 1391

Chapter 22: ACL Commands 1393

clear access-list counter 1394
clear access-list meter 1395
configure access-list 1396
configure access-list add 1398
configure access-list delete 1400
configure access-list network-zone 1401
configure access-list rule-compression port-counters 1402
configure access-list vlan-acl-precedence 1403
configure access-list width 1404
configure access-list zone 1405
configure flow-redirect add nexthop 1407
configure flow-redirect delete nexthop 1408
configure flow-redirect health-check 1409
configure flow-redirect nexthop 1410
configure flow-redirect no-active 1411
configure flow-redirect vr 1411
create access-list 1412
create access-list zone 1414
create access-list network-zone 1415
create flow-redirect 1416
delete access-list 1417
delete access-list network-zone 1418
delete access-list zone 1419
delete flow-redirect 1419
disable access-list permit to-cpu 1420
disable access-list refresh blackhole 1421
enable access-list permit to-cpu 1422
enable access-list refresh blackhole 1423
refresh access-list network-zone 1424
show access-list 1425
show access-list configuration 1427
show access-list counter 1428
show access-list dynamic 1429
show access-list dynamic counter 1430
show access-list dynamic rule 1431



show access-list interface 1433
show access-list network-zone 1435
show access-list usage acl-mask port 1437
show access-list usage acl-range port 1438
show access-list usage acl-rule port 1439
show access-list usage acl-slice port 1441
show access-list width 1444
show flow-redirect 1446
unconfigure access-list 1447

Chapter 23: QoS Commands 1449

clear counters wred 1450
configure diffserv examination code-point qosprofile 1451
configure diffserv replacement code-point 1452
configure dot1p type 1453
configure meter 1455
configure port shared-packet-buffer 1457
configure ports qosprofile 1459
configure ports rate-limit egress 1460
configure qosprofile 1461
configure qosprofile qp8 weight 1465
configure qosprofile wred 1466
configure qosscheduler weighted-deficit-round-robin 1468
create meter 1470
create qosprofile 1471
delete meter 1472
delete qosprofile 1472
disable diffserv examination ports 1473
disable diffserv replacement ports 1474
disable dot1p examination ports 1475
disable dot1p replacement ports 1476
enable diffserv examination ports 1477
enable diffserv replacement ports 1478
enable dot1p examination ports 1478
enable dot1p replacement ports 1479
show access-list meter 1481
show diffserv examination 1482
show diffserv replacement 1483
show dot1p 1484
show meter 1485
show ports congestion 1486
show ports qosmonitor 1488
show ports qosmonitor {congestion} 1490
show ports wred 1492
show qosprofile 1494
show wredprofile 1496
unconfigure diffserv examination 1497
unconfigure diffserv replacement 1498
unconfigure qosprofile 1498
unconfigure qosprofile wred 1500



Chapter 24: Network Login Commands 1502

- clear netlogin state 1504
- configure netlogin add mac-list 1504
- configure netlogin add proxy-port 1506
- configure netlogin agingtime 1507
- configure netlogin allowed-refresh-failures 1507
- configure netlogin authentication database-order 1508
- configure netlogin authentication failure vlan 1509
- configure netlogin authentication service-unavailable vlan 1510
- configure netlogin banner 1512
- configure netlogin base-url 1513
- configure netlogin delete mac-list 1514
- configure netlogin delete proxy-port 1515
- configure netlogin dot1x eapol-transmit-version 1515
- configure netlogin dot1x guest-vlan 1516
- configure netlogin dot1x timers 1518
- configure netlogin dynamic-vlan 1520
- configure netlogin dynamic-vlan uplink-ports 1522
- configure netlogin local-user 1523
- configure netlogin local-user security-profile 1526
- configure netlogin mac timers reauth-period 1526
- configure netlogin move-fail-action 1527
- configure netlogin port allow egress-traffic 1528
- configure netlogin ports mode 1529
- configure netlogin ports no-restart 1532
- configure netlogin ports restart 1533
- configure netlogin redirect-page 1534
- configure netlogin session-refresh 1535
- configure netlogin vlan 1536
- configure vlan netlogin-lease-timer 1537
- create netlogin local-user 1538
- delete netlogin local-user 1541
- disable netlogin 1541
- disable netlogin authentication failure vlan ports 1542
- disable netlogin authentication service-unavailable vlan ports 1543
- disable netlogin dot1x guest-vlan ports 1544
- disable netlogin logout-privilege 1545
- disable netlogin ports 1545
- disable netlogin reauthenticate-on-refresh 1546
- disable netlogin redirect-page 1547
- disable netlogin session-refresh 1548
- enable netlogin 1549
- enable netlogin authentication failure vlan ports 1549
- enable netlogin authentication service-unavailable vlan ports 1550
- enable netlogin dot1x guest-vlan ports 1551
- enable netlogin logout-privilege 1552
- enable netlogin ports 1553
- enable netlogin reauthentication-on-refresh 1554
- enable netlogin redirect-page 1555



enable netlogin session-refresh 1555
show banner netlogin 1556
show netlogin 1557
show netlogin authentication failure vlan 1562
show netlogin authentication service-unavailable vlan 1563
show netlogin banner 1564
show netlogin guest-vlan 1565
show netlogin local-users 1566
show netlogin mac-list 1567
unconfigure netlogin allowed-refresh-failures 1568
unconfigure netlogin authentication database-order 1569
unconfigure netlogin authentication failure vlan 1570
unconfigure netlogin authentication service-unavailable vlan 1571
unconfigure netlogin banner 1571
unconfigure netlogin dot1x guest-vlan 1572
unconfigure netlogin local-user security-profile 1573
unconfigure netlogin session-refresh 1574
unconfigure netlogin vlan 1574

Chapter 25: Commands for Identity Management 1576

clear counters identity-management 1577
configure identity-management access-list 1578
configure identity-management blacklist 1579
configure identity-management database memory-size 1581
configure identity-management detection 1582
configure identity-management greylist 1584
configure identity-management kerberos snooping aging time 1586
configure identity-management kerberos snooping force-aging time 1586
configure identity-management kerberos snooping forwarding 1587
configure identity-management kerberos snooping server 1589
configure identity-management list-precedence 1590
configure identity-management ports 1591
configure identity-management role 1592
configure identity-management role add child-role 1593
configure identity-management role add dynamic-rule 1594
configure identity-management role add policy 1595
configure identity-management role-based-vlan 1596
configure identity-management role delete child-role 1597
configure identity-management role delete dynamic-rule 1598
configure identity-management role delete policy 1599
configure identity-management role match-criteria inheritance 1600
configure identity-management role priority 1601
configure identity-management stale-entry aging-time 1602
configure identity-management whitelist 1605
configure ldap domain 1607
configure ldap domain add server 1607
configure ldap domain base-dn 1609
configure ldap domain bind-user 1611
configure ldap domain delete server 1612
configure ldap domain netlogin 1613



create identity-management role 1614
create ldap domain 1617
configure ldap hierarchical-search-oid 1618
delete identity-management role 1619
delete ldap domain 1620
disable identity-management 1621
disable snmp traps identity-management 1622
enable identity-management 1623
enable snmp traps identity-management 1624
refresh identity-management role 1625
show identity-management 1626
show identity-management blacklist 1627
show identity-management entries 1628
show identity-management greylist 1632
show identity-management list-precedence 1633
show identity-management role 1634
show identity-management statistics 1636
show identity-management whitelist 1637
show ldap domain 1638
show ldap statistics 1641
unconfigure identity-management 1643
unconfigure identity-management list-precedence 1644
unconfigure ldap domains 1645

Chapter 26: Security Commands 1646

SSH 1650
User Authentication 1650
Denial of Service 1650
clear ip-security anomaly-protection notify cache 1650
clear ip-security arp validation violations 1651
clear ip-security dhcp-snooping entries 1652
clear ip-security source-ip-lockdown entries ports 1652
clear vlan dhcp-address-allocation 1653
configure dos-protect acl-expire 1654
configure dos-protect interval 1655
configure dos-protect trusted ports 1656
configure dos-protect type l3-protect alert-threshold 1657
configure dos-protect type l3-protect notify-threshold 1658
configure ip-security anomaly-protection icmp ipv4-max-size 1658
configure ip-security anomaly-protection icmp ipv6-max-size 1659
configure ip-security anomaly-protection notify cache 1660
configure ip-security anomaly-protection notify rate limit 1661
configure ip-security anomaly-protection notify rate window 1661
configure ip-security anomaly-protection notify trigger off 1662
configure ip-security anomaly-protection notify trigger on 1663
configure ip-security anomaly-protection tcp 1664
configure ip-security dhcp-snooping information check 1665
configure ip-security dhcp-snooping information circuit-id port-information port 1666
configure ip-security dhcp-snooping information circuit-id vlan-information 1667
configure ip-security dhcp-snooping information option 1667



configure ip-security dhcp-snooping information policy	1668
configure ip-security dhcp-bindings add	1669
configure ip-security dhcp-bindings delete	1670
configure ip-security dhcp-bindings storage	1671
configure ip-security dhcp-bindings storage filename	1672
configure ip-security dhcp-bindings storage location	1673
configure mac-lockdown-timeout ports aging-time	1674
configure ports rate-limit flood	1675
configure ports vlan	1676
configure radius server client-ip	1679
configure radius shared-secret	1681
configure radius timeout	1682
configure radius-accounting server client-ip	1683
configure radius-accounting shared-secret	1685
configure radius-accounting timeout	1686
configure ssh2 key	1687
configure sshd2 user-key add user	1689
configure sshd2 user-key delete user	1690
configure ssl certificate pregenerated	1691
configure ssl certificate privkeylen	1692
configure ssl privkey pregenerated	1693
configure tacacs server client-ip	1694
configure tacacs shared-secret	1696
configure tacacs timeout	1697
configure tacacs-accounting server	1698
configure tacacs-accounting shared-secret	1699
configure tacacs-accounting timeout	1700
configure trusted-ports trust-for dhcp-server	1701
configure trusted-servers add server	1702
configure trusted-servers delete server	1703
configure vlan dhcp-address-range	1704
configure vlan dhcp-lease-timer	1705
configure vlan dhcp-options	1706
create sshd2 key-file	1707
create sshd2 user-key	1708
delete sshd2 user-key	1709
disable dhcp ports vlan	1710
disable dos-protect	1711
disable iparp gratuitous protect vlan	1711
disable ip-security anomaly-protection	1712
disable ip-security anomaly-protection ip	1713
disable ip-security anomaly-protection l4port	1714
disable ip-security anomaly-protection tcp flags	1715
disable ip-security anomaly-protection tcp fragment	1715
disable ip-security anomaly-protection icmp	1716
disable ip-security anomaly-protection notify	1717
disable ip-security arp gratuitous-protection	1718
disable ip-security arp learning learn-from-arp	1719
disable ip-security arp learning learn-from-dhcp	1720



disable ip-security arp validation 1721
disable ip-security dhcp-bindings restoration 1722
disable ip-security dhcp-snooping 1723
disable ip-security source-ip-lockdown ports 1724
disable mac-lockdown-timeout ports 1725
disable radius 1726
disable radius-accounting 1727
disable ssh2 1728
disable tacacs 1728
disable tacacs-accounting 1729
disable tacacs-authorization 1730
disable web http 1731
disable web https 1732
download ssl certificate 1732
download ssl privkey 1734
enable dhcp ports vlan 1736
enable dos-protect 1737
enable dos-protect simulated 1738
enable iparp gratuitous protect 1738
enable ip-option loose-source-route 1739
enable ip-security anomaly-protection 1740
enable ip-security anomaly-protection icmp 1741
enable ip-security anomaly-protection ip 1742
enable ip-security anomaly-protection l4port 1743
enable ip-security anomaly-protection notify 1743
enable ip-security anomaly-protection tcp flags 1744
enable ip-security anomaly-protection tcp fragment 1745
enable ip-security arp gratuitous-protection 1746
enable ip-security arp learning learn-from-arp 1747
enable ip-security arp learning learn-from-dhcp 1749
enable ip-security arp validation violation-action 1750
enable ip-security dhcp-bindings restoration 1752
enable ip-security dhcp-snooping 1753
enable ip-security source-ip-lockdown ports 1754
enable mac-lockdown-timeout ports 1756
enable radius 1756
enable radius-accounting 1758
enable ssh2 1759
enable tacacs 1761
enable tacacs-accounting 1762
enable tacacs-authorization 1763
enable web http 1764
enable web https 1764
scp2 1765
show dhcp-server 1768
show dos-protect 1769
show ip-security anomaly-protection notify cache ports 1770
show ip-security arp gratuitous-protection 1771
show ip-security arp learning 1772



show ip-security arp validation	1773
show ip-security arp validation violations	1774
show ip-security dhcp-snooping entries	1775
show ip-security dhcp-snooping information-option	1776
show ip-security dhcp-snooping information circuit-id port-information	1777
show ip-security dhcp-snooping information-option circuit-id vlan-information	1779
show ip-security dhcp-snooping	1780
show ip-security dhcp-snooping violations	1781
show ip-security source-ip-lockdown	1782
show mac-lockdown-timeout fdb ports	1783
show mac-lockdown-timeout ports	1784
show ports rate-limit flood	1785
show radius	1789
show radius-accounting	1791
show ssh2 private-key	1792
show sshd2 user-key	1793
show ssl	1794
show tacacs	1795
show tacacs-accounting	1797
show vlan dhcp-address-allocation	1798
show vlan dhcp-config	1799
show vlan security	1800
ssh2	1801
unconfigure ip-security dhcp-snooping information check	1804
unconfigure ip-security dhcp-snooping information circuit-id port-information ports	1804
unconfigure ip-security dhcp-snooping information circuit-id vlan-information	1805
unconfigure ip-security dhcp-snooping information option	1806
unconfigure ip-security dhcp-snooping information policy	1806
unconfigure radius	1807
unconfigure radius-accounting	1808
unconfigure tacacs	1809
unconfigure tacacs-accounting	1810
unconfigure trusted-ports trust-for dhcp-server	1811
unconfigure vlan dhcp	1812
unconfigure vlan dhcp-address-range	1813
unconfigure vlan dhcp-options	1814
upload dhcp-bindings	1815
Chapter 27: CLEAR-Flow Commands	1816
disable clear-flow	1817
enable clear-flow	1817
show clear-flow	1818
show clear-flow acl-modified	1819
show clear-flow rule	1820
show clear-flow rule-all	1822
show clear-flow rule-triggered	1823
Chapter 28: EAPS Commands	1825
clear eaps counters	1826
configure eaps add control vlan	1827



configure eaps add protected vlan 1828
configure eaps cfm 1829
configure eaps config-warnings off 1830
configure eaps config-warnings on 1831
configure eaps delete control vlan 1832
configure eaps delete protected vlan 1833
configure eaps failtime 1835
configure eaps failtime expiry-action 1836
configure eaps fast-convergence 1837
configure eaps hello-pdu-egress 1838
configure eaps hellotime 1839
configure eaps mode 1840
configure eaps multicast add-ring-ports 1842
configure eaps multicast send-igmp-query 1843
configure eaps multicast temporary-flooding 1844
configure eaps multicast temporary-flooding duration 1845
configure eaps name 1846
configure eaps port 1847
configure eaps priority 1849
configure eaps shared-port common-path-timers 1850
configure eaps shared-port link-id 1851
configure eaps shared-port mode 1852
configure eaps shared-port segment-timers expiry-action 1853
configure eaps shared-port segment-timers health-interval 1854
configure eaps shared-port segment-timers timeout 1855
configure forwarding L2-protocol fast-convergence 1856
configure ip-arp fast-convergence 1857
create eaps 1859
create eaps shared-port 1859
delete eaps 1860
delete eaps shared-port 1861
disable eaps 1862
enable eaps 1863
show eaps 1865
show eaps cfm groups 1870
show eaps counters 1871
show eaps counters shared-port 1876
show eaps shared-port 1880
show eaps shared-port neighbor-info 1884
show vlan eaps 1885
unconfigure eaps shared-port link-id 1887
unconfigure eaps shared-port mode 1888
unconfigure eaps port 1888

Chapter 29: ERPS Commands 1891

clear counters erps 1892
configure erps dynamic-state clear 1893
configure erps add control vlan 1894
configure erps add protected vlan 1894
configure erps cfm md-level 1895



configure erps cfm port ccm-interval	1896
configure erps cfm port group	1897
configure erps cfm port mepid	1898
configure erps delete control vlan	1899
configure erps delete protected vlan	1900
configure erps name	1901
configure erps neighbor port	1902
configure erps notify-topology-change	1902
configure erps protection-port	1903
configure erps revert	1904
configure erps ring-ports east west	1905
configure erps subring-mode	1906
configure erps timer guard	1907
configure erps timer hold-off	1907
configure erps timer periodic	1908
configure erps timer wait-to-block	1909
configure erps timer wait-to-restore	1910
configure erps topology-change	1911
create erps ring	1912
debug erps	1912
debug erps show	1913
delete erps	1914
disable erps	1915
disable erps block-vc-recovery	1915
disable erps ring-name	1916
disable erps topology-change	1917
enable erps	1918
enable erps block-vc-recovery	1919
enable erps ring-name	1919
enable erps topology-change	1920
run erps force-switch manual-switch	1921
show erps	1922
show erps ring-name	1923
show erps statistics	1924
unconfigure erps cfm	1925
unconfigure erps neighbor-port	1926
unconfigure erps notify-topology-change	1927
unconfigure erps protection-port	1928
unconfigure erps ring-ports west	1928
configure erps cfm protection group	1929

Chapter 30: STP Commands 1931

STP	1932
RSTP	1932
MSTP	1933
Spanning Tree Domains	1933
STP Rules and Restrictions	1936
clear counters stp	1937
configure mstp format	1938
configure mstp region	1939



configure mstp revision 1940
 configure stpd add vlan 1941
 configure stpd default-encapsulation 1944
 configure stpd delete vlan 1946
 configure stpd description 1947
 configure stpd flush-method 1948
 configure stpd forwarddelay 1949
 configure stpd hellotime 1950
 configure stpd maxage 1951
 configure stpd max-hop-count 1952
 configure stpd mode 1953
 configure stpd ports active-role disable 1955
 configure stpd ports active-role enable 1955
 configure stpd ports bpdu-restrict 1957
 configure stpd ports cost 1958
 configure stpd ports edge-safeguard disable 1959
 configure stpd ports edge-safeguard enable 1961
 configure stpd ports link-type 1963
 configure stpd ports mode 1966
 configure stpd ports port-priority 1967
 configure stpd ports priority 1968
 configure stpd ports restricted-role disable 1970
 configure stpd ports restricted-role enable 1971
 configure stpd priority 1972
 configure stpd tag 1973
 configure vlan add ports stpd 1974
 create stpd 1977
 delete stpd 1979
 disable stpd 1980
 disable stpd auto-bind 1981
 disable stpd ports 1982
 disable stpd rapid-root-failover 1983
 enable stpd 1984
 enable stpd auto-bind 1984
 enable stpd ports 1987
 enable stpd rapid-root-failover 1988
 show stpd 1989
 show stpd ports 1993
 show vlan stpd 1996
 unconfigure mstp region 1998
 unconfigure stpd 1999
 unconfigure stpd ports link-type 2000

Chapter 31: ESRP Commands 2002

clear esrp counters 2003
 clear esrp neighbor 2004
 clear esrp sticky 2005
 configure esrp add elrp-poll ports 2006
 configure esrp add master 2007
 configure esrp add member 2008



configure esrp add track-environment 2009
configure esrp add track-iproute 2010
configure esrp add track-ping 2011
configure esrp add track-vlan 2012
configure esrp aware add selective-forward-ports 2013
configure esrp aware delete selective-forward-ports 2014
configure esrp delete elrp-poll ports 2015
configure esrp delete master 2016
configure esrp delete member 2017
configure esrp delete track-environment 2018
configure esrp delete track-iproute 2019
configure esrp delete track-ping 2020
configure esrp delete track-vlan 2021
configure esrp domain-id 2021
configure esrp election-policy 2022
configure esrp elrp-master-poll disable 2026
configure esrp elrp-master-poll enable 2027
configure esrp elrp-premaster-poll disable 2028
configure esrp elrp-premaster-poll enable 2029
configure esrp group 2030
configure esrp mode 2032
configure esrp name 2033
configure esrp ports mode 2034
configure esrp ports no-restart 2035
configure esrp ports restart 2035
configure esrp ports weight 2036
configure esrp priority 2037
configure esrp timer hello 2038
configure esrp timer neighbor 2040
configure esrp timer neutral 2041
configure esrp timer premaster 2042
configure esrp timer restart 2043
create esrp 2044
delete esrp 2046
disable esrp 2046
enable esrp 2047
show esrp 2048
show esrp aware 2051
show esrp counters 2053

Chapter 32: VRRP Commands 2055

clear counters vrrp 2055
configure vrrp vlan vrid accept-mode 2057
configure vrrp vlan vrid add ipaddress 2058
configure vrrp vlan vrid add track-iproute 2060
configure vrrp vlan vrid add track-ping 2061
configure vrrp vlan vrid add virtual-link-local 2062
configure vrrp vlan vrid add track-vlan 2063
configure vrrp vlan vrid advertisement-interval 2064
configure vrrp vlan vrid authentication 2065



configure vrrp vlan vrid delete 2066
configure vrrp vlan vrid delete track-iproute 2067
configure vrrp vlan vrid delete track-ping 2068
configure vrrp vlan vrid delete track-vlan 2069
configure vrrp vlan vrid dont-preempt 2070
configure vrrp vlan vrid preempt 2071
configure vrrp vlan vrid priority 2072
configure vrrp vlan vrid track-mode 2073
configure vrrp vlan vrid version 2074
create vrrp vlan vrid 2075
delete vrrp vlan vrid 2077
disable vrrp vrid 2078
enable vrrp vrid 2078
show vrrp 2079
show vrrp vlan 2083

Chapter 33: MPLS Commands 2085

clear counters l2vpn 2089
clear counters mpls 2090
clear counters mpls ldp 2091
clear counters mpls rsvp-te 2092
clear counters mpls static lsp 2093
clear counters vpls 2094
clear fdb vpls 2095
configure forwarding switch-fabric protocol 2096
configure iproute add default 2097
configure iproute add lsp 2098
configure iproute delete 2100
configure iproute delete default 2101
configure l2vpn 2102
configure l2vpn add peer 2104
configure l2vpn add service 2106
configure l2vpn delete peer 2107
configure l2vpn delete service 2108
configure l2vpn health-check vccv 2109
configure l2vpn peer mpls lsp 2111
configure l2vpn peer 2112
configure l2vpn peer mpls lsp 2113
configure l2vpn vpls peer static-pw 2115
configure l2vpn vpls redundancy 2116
configure l2vpn vpls add peer ipaddress 2117
configure l2vpn vpws add peer ipaddress 2119
configure l2vpn vpws peer static-pw 2120
configure mpls add vlan 2121
configure mpls delete vlan 2122
configure mpls exp examination 2123
configure mpls exp replacement 2124
configure mpls labels max-static 2125
configure mpls ldp advertise 2127
configure mpls ldp loop-detection 2128



configure mpls ldp pseudo-wire	2129
configure mpls ldp timers	2130
configure mpls lsr-id	2132
configure mpls rsvp-te bandwidth committed-rate	2133
configure mpls rsvp-te lsp add path	2134
configure mpls rsvp-te lsp change	2136
configure mpls rsvp-te lsp delete path	2137
configure mpls rsvp-te lsp fast-reroute	2137
configure mpls rsvp-te lsp path use profile	2138
configure mpls rsvp-te lsp transport	2139
configure mpls rsvp-te metric	2141
configure mpls rsvp-te path add ero	2142
configure mpls rsvp-te path delete ero	2143
configure mpls rsvp-te profile	2144
configure mpls rsvp-te profile (fast-reroute)	2147
configure mpls rsvp-te timers lsp rapid-retry	2149
configure mpls rsvp-te timers lsp standard-retry	2150
configure mpls rsvp-te timers session	2152
configure mpls static lsp	2154
configure mpls static lsp transport	2155
configure vpls	2157
configure vpls add peer	2159
configure vpls add service	2160
configure vpls delete peer	2162
configure vpls delete service	2163
configure vpls health-check vccv	2164
configure vpls peer mpls lsp	2165
configure vpls peer	2167
configure vpls peer mpls lsp	2168
configure vpls snmp-vpn-identifier	2169
create l2vpn fec-id-type pseudo-wire	2170
create mpls rsvp-te lsp	2171
create mpls rsvp-te path	2172
create mpls rsvp-te profile	2173
create mpls rsvp-te profile fast-reroute	2174
create mpls static lsp	2175
create vpls fec-id-type pseudo-wire	2176
delete l2vpn	2177
delete mpls rsvp-te lsp	2179
delete mpls rsvp-te path	2179
delete mpls rsvp-te profile	2180
delete mpls static lsp	2181
delete vpls	2182
disable bgp mpls-next-hop	2183
disable iproute mpls-next-hop	2184
disable l2vpn	2185
disable l2vpn vpls fdb mac-withdrawal	2186
disable l2vpn health-check vccv	2187
disable l2vpn service	2188



disable l2vpn sharing	2189
disable mpls	2190
disable mpls bfd	2190
disable mpls exp examination	2192
disable mpls exp replacement	2193
disable mpls ldp	2193
disable mpls ldp bgp-routes	2194
disable mpls ldp loop-detection	2195
disable mpls php	2196
disable mpls protocol ldp	2197
disable mpls protocol rsvp-te	2198
disable mpls rsvp-te	2199
disable mpls rsvp-te bundle-message	2200
disable mpls rsvp-te fast-reroute	2201
disable mpls rsvp-te lsp	2202
disable mpls rsvp-te summary-refresh	2203
disable mpls static lsp	2204
disable mpls vlan	2204
disable ospf mpls-next-hop	2205
disable snmp traps l2vpn	2206
disable snmp traps mpls	2207
disable vpls	2208
disable vpls fdb mac-withdrawal	2209
disable vpls health-check vccv	2210
disable vpls service	2211
enable bgp mpls-next-hop	2212
enable iproute mpls-next-hop	2212
enable l2vpn	2213
enable l2vpn vpls fdb mac-withdrawal	2214
enable l2vpn health-check vccv	2215
enable l2vpn service	2216
enable l2vpn sharing	2217
enable mpls	2218
enable mpls bfd	2220
enable mpls exp examination	2221
enable mpls exp replacement	2222
enable mpls ldp	2222
enable mpls ldp bgp-routes	2223
enable mpls ldp loop-detection	2224
enable mpls php	2225
enable mpls protocol ldp	2226
enable mpls protocol rsvp-te	2227
enable mpls rsvp-te	2228
enable mpls rsvp-te bundle-message	2229
enable mpls rsvp-te fast-reroute	2230
enable mpls rsvp-te lsp	2231
enable mpls rsvp-te summary-refresh	2232
enable mpls static lsp	2233
enable mpls vlan	2234



enable ospf mpls-next-hop	2234
enable snmp traps l2vpn	2235
enable snmp traps mpls	2236
enable vpls	2237
enable vpls fdb mac-withdrawal	2238
enable vpls health-check vccv	2239
enable vpls service	2240
ping mpls lsp	2241
restart process mpls	2243
show bandwidth pool	2244
show ces	2245
show l2vpn	2247
show mpls	2253
show mpls bfd	2254
show mpls exp examination	2255
show mpls exp replacement	2256
show mpls interface	2257
show mpls label	2259
show mpls label usage	2262
show mpls ldp	2265
show mpls ldp interface	2267
show mpls ldp label	2268
show mpls ldp label advertised	2270
show mpls ldp label l2vpn	2271
show mpls ldp label l2vpn retained	2273
show mpls ldp label lsp retained	2274
show mpls ldp label retained	2275
show mpls ldp lsp	2276
show mpls ldp peer	2278
show mpls rsvp-te	2281
show mpls rsvp-te bandwidth	2282
show mpls rsvp-te interface	2284
show mpls rsvp-te lsp	2286
show mpls rsvp-te lsp [egress transit]	2289
show mpls rsvp-te lsp ingress	2291
show mpls rsvp-te neighbor	2293
show mpls rsvp-te path	2295
show mpls rsvp-te profile	2296
show mpls rsvp-te profile fast-reroute	2297
show mpls static lsp	2298
show mpls statistics l2vpn	2300
show vpls	2302
traceroute mpls lsp	2311
unconfigure l2vpn dot1q ethertype	2313
unconfigure l2vpn vpls redundancy	2314
unconfigure mpls	2315
unconfigure mpls exp examination	2316
unconfigure mpls exp replacement	2316
unconfigure mpls vlan	2317



unconfigure vpls dot1q ethertype 2318
unconfigure vpls snmp-vpn-identifier 2319

Chapter 34: IP Unicast Commands 2321

clear ip dad 2324
clear iparp 2325
configure bootprelay add 2326
configure bootprelay delete 2327
configure bootprelay dhcp-agent information check 2328
configure bootprelay dhcp-agent information circuit-id port-information 2329
configure bootprelay dhcp-agent information circuit-id vlan-information 2330
configure bootprelay dhcp-agent information option 2331
configure bootprelay dhcp-agent information policy 2332
configure bootprelay dhcp-agent information remote-id 2333
configure bootprelay include-secondary 2334
configure bootprelay vlan include-secondary 2335
configure forwarding sharing 2336
configure ip dad 2338
configure iparp add 2339
configure iparp add proxy 2339
configure iparp delete 2341
configure iparp delete proxy 2342
configure iparp distributed-mode 2343
configure iparp max_entries 2344
configure iparp max_pending_entries 2346
configure iparp max_proxy_entries 2346
configure iparp timeout 2347
configure ipforwarding originated-packets 2348
configure iproute add (IPv4) 2349
configure iproute add blackhole 2351
configure iproute add blackhole ipv4 default 2352
configure iproute add default 2353
configure iproute delete 2354
configure iproute delete blackhole 2355
configure iproute delete blackhole ipv4 default 2356
configure iproute delete default 2357
configure iproute priority 2358
configure iproute reserved-entries 2360
configure iproute sharing hash-algorithm crc 2363
configure iproute sharing max-gateways 2364
configure irdp 2366
configure vlan add secondary-ipaddress 2367
configure vlan delete secondary-ipaddress 2368
configure vlan subvlan 2369
configure vlan subvlan-address-range 2370
configure vlan delete secondary-ipaddress 2371
disable bootp vlan 2372
disable bootprelay 2373
disable icmp address-mask 2374
disable icmp parameter-problem 2375



disable icmp port-unreachables	2376
disable icmp redirects	2376
disable icmp time-exceeded	2377
disable icmp timestamp	2378
disable icmp unreachable	2379
disable icmp userredirects	2380
disable iparp checking	2381
disable iparp refresh	2382
disable ipforwarding	2382
disable ip-option loose-source-route	2383
disable ip-option record-route	2384
disable ip-option record-timestamp	2385
disable ip-option router-alert	2386
disable ip-option strict-source-route	2386
disable iproute bfd	2387
disable iproute compression	2388
disable iproute sharing	2389
disable irdp	2390
disable subvlan-proxy-arp vlan	2390
disable udp-echo-server	2391
enable bootp vlan	2392
enable bootprelay	2393
enable icmp address-mask	2394
enable icmp parameter-problem	2395
enable icmp port-unreachables	2396
enable icmp redirects	2397
enable icmp time-exceeded	2398
enable icmp timestamp	2399
enable icmp unreachable	2400
enable icmp userredirects	2401
enable iparp checking	2402
enable iparp refresh	2402
enable ipforwarding	2403
enable ip-option record-route	2405
enable ip-option record-timestamp	2406
enable ip-option strict-source-route	2406
enable ip-option router-alert	2407
enable iproute bfd	2408
enable iproute compression	2409
enable iproute sharing	2410
enable irdp	2411
enable subvlan-proxy-arp vlan	2412
enable udp-echo-server	2413
rtlookup	2414
run ip dad	2415
show bootprelay	2416
show bootprelay configuration	2418
show bootprelay dhcp-agent information circuit-id port-information	2419
show bootprelay dhcp-agent information circuit-id vlan-information	2420



show bootprelay configuration ipv4 2421
show bootprelay ipv6 2422
show ip dad 2424
show iparp 2425
show iparp distributed-mode statistics 2428
show iparp proxy 2431
show iparp security 2431
show iparp stats 2433
show ipconfig 2435
show iproute 2437
show iproute mpls 2438
show iproute mpls origin 2439
show iproute origin 2441
show iproute reserved-entries 2443
show iproute reserved-entries statistics 2444
show ipstats 2446
show udp-profile 2449
unconfigure bootprelay dhcp-agent information check 2450
unconfigure bootprelay dhcp-agent information circuit-id port-information 2451
unconfigure bootprelay dhcp-agent information circuit-id vlan-information 2452
unconfigure bootprelay dhcp-agent information option 2452
unconfigure bootprelay dhcp-agent information policy 2453
unconfigure bootprelay dhcp-agent information remote-id 2454
unconfigure bootprelay include-secondary 2455
unconfigure icmp 2456
unconfigure iparp 2456
unconfigure iproute priority 2457
unconfigure irdp 2459
unconfigure vlan subvlan-address-range 2460
unconfigure vlan udp-profile 2461

Chapter 35: IPv6 Unicast Commands 2462

clear ipv6 dad 2464
clear neighbor-discovery cache 2465
configure iproute add (IPv6) 2466
configure iproute add blackhole 2467
configure iproute add blackhole ipv6 default 2468
configure iproute add default 2469
configure iproute delete 2470
configure iproute delete blackhole 2471
configure iproute delete blackhole ipv6 default 2472
configure iproute delete default 2473
configure iproute ipv6 priority 2474
configure iproute sharing max-gateways 2476
configure ipv6 dad 2477
configure ipv6 hop-limit 2478
configure neighbor-discovery cache add 2479
configure neighbor-discovery cache delete 2480
configure neighbor-discovery cache max_entries 2481
configure neighbor-discovery cache max_pending_entries 2482



configure neighbor-discovery cache timeout 2482
configure vlan router-discovery add prefix 2483
configure vlan router-discovery delete prefix 2484
configure vlan router-discovery default-lifetime 2485
configure vlan router-discovery link-mtu 2486
configure vlan router-discovery managed-config-flag 2487
configure vlan router-discovery max-interval 2488
configure vlan router-discovery min-interval 2489
configure vlan router-discovery other-config-flag 2490
configure vlan router-discovery reachable-time 2491
configure vlan router-discovery retransmit-time 2491
configure vlan router-discovery set prefix 2492
configure tunnel ipaddress 2494
create tunnel 6to4 2495
create tunnel gre destination source 2496
create tunnel ipv6-in-ipv4 2497
delete tunnel 2498
disable icmp redirects ipv6 fast-path 2499
disable ipforwarding ipv6 2500
disable iproute ipv6 compression 2501
disable iproute ipv6 sharing 2502
disable neighbor-discovery refresh 2503
disable router-discovery 2503
disable tunnel 2504
enable icmp redirects ipv6 fast-path 2505
enable ipforwarding ipv6 2506
enable ipforwarding 2507
enable iproute ipv6 compression 2508
enable iproute ipv6 sharing 2508
enable neighbor-discovery refresh 2509
enable router-discovery 2510
enable tunnel 2511
rtlookup 2512
rtlookup rpf 2513
run ipv6 dad 2514
show ipconfig ipv6 2515
show iproute ipv6 2516
show iproute ipv6 origin 2518
show ipstats ipv6 2520
show ipv6 dad 2520
show neighbor-discovery cache ipv6 2522
show router-discovery 2523
show tunnel 2524
unconfigure iproute ipv6 priority 2525
unconfigure neighbor-discovery cache 2527
unconfigure vlan router-discovery 2528
unconfigure vlan router-discovery default-lifetime 2529
unconfigure vlan router-discovery hop-limit 2530
unconfigure vlan router-discovery link-mtu 2530



unconfigure vlan router-discovery managed-config-flag 2531
unconfigure vlan router-discovery max-interval 2532
unconfigure vlan router-discovery min-interval 2533
unconfigure vlan router-discovery other-config-flag 2534
unconfigure vlan router-discovery reachable-time 2534
unconfigure vlan router-discovery retransmit-time 2535
unconfigure tunnel 2536

Chapter 36: RIP Commands 2538

clear rip counters 2539
configure rip add vlan 2540
configure rip delete vlan 2541
configure rip garbage-time 2542
configure rip import-policy 2542
configure rip routetimeout 2543
configure rip updatetime 2544
configure rip vlan cost 2545
configure rip vlan route-policy 2546
configure rip vlan rxmode 2547
configure rip vlan trusted-gateway 2548
configure rip vlan txmode 2549
disable rip 2550
disable rip aggregation 2550
disable rip export 2551
disable rip poisonreverse 2553
disable rip splithorizon 2553
disable rip triggerupdates 2554
disable rip use-ip-router-alert 2555
enable rip 2556
enable rip aggregation 2557
enable rip export 2558
enable rip originate-default cost 2559
enable rip poisonreverse 2560
enable rip splithorizon 2561
enable rip triggerupdates 2562
enable rip use-ip-router-alert 2562
show rip 2563
show rip interface 2564
show rip interface vlan 2566
show rip memory 2567
show rip routes 2568
unconfigure rip 2569

Chapter 37: RIPng Commands 2571

clear ripng counters 2572
configure ripng add 2573
configure ripng cost 2574
configure ripng delete 2574
configure ripng garbage-time 2575
configure ripng import-policy 2576



configure ripng route-policy 2577
configure ripng routetimeout 2579
configure ripng trusted-gateway 2579
configure ripng updatetime 2581
disable ripng 2581
disable ripng export 2582
disable ripng originate-default 2583
disable ripng poisonreverse 2584
disable ripng splithorizon 2585
disable ripng triggerupdate 2586
enable ripng 2587
enable ripng export 2588
enable ripng originate-default 2590
enable ripng poisonreverse 2591
enable ripng splithorizon 2592
enable ripng triggerupdates 2593
show ripng 2594
show ripng interface 2595
show ripng routes 2597
unconfigure ripng 2599

Chapter 38: OSPF Commands 2600

Licensing 2602
OSPF Edge Mode 2602
clear ospf counters 2602
configure ospf add virtual-link 2603
configure ospf add vlan area 2604
configure ospf add vlan area link-type 2605
configure ospf area external-filter 2606
configure ospf area interarea-filter 2607
configure ospf area add range 2608
configure ospf area normal 2609
configure ospf area nssa stub-default-cost 2610
configure ospf area stub stub-default-cost 2611
configure ospf area timer 2612
configure ospf ase-limit 2613
configure ospf ase-summary add 2614
configure ospf ase-summary delete 2615
configure ospf authentication 2616
configure ospf bfd 2617
configure ospf cost 2618
configure ospf delete virtual-link 2619
configure ospf delete vlan 2620
configure ospf import-policy 2621
configure ospf lsa-batch-interval 2621
configure ospf metric-table 2622
configure ospf priority 2623
configure ospf restart 2624
configure ospf restart grace-period 2625
configure ospf restart-helper 2626



configure ospf routerid 2627
configure ospf spf-hold-time 2628
configure ospf virtual-link timer 2629
configure ospf vlan area 2630
configure ospf vlan neighbor add 2631
configure ospf vlan neighbor delete 2632
configure ospf vlan timer 2633
create ospf area 2634
delete ospf area 2635
disable ospf 2636
disable ospf capability opaque-lsa 2637
disable ospf export 2638
disable ospf originate-default 2639
disable ospf restart-helper-lsa-check 2640
disable ospf use-ip-router-alert 2641
disable snmp traps ospf 2641
enable ospf 2642
enable ospf capability opaque-lsa 2643
enable ospf export 2644
enable ospf originate-default 2645
enable ospf restart-helper-lsa-check 2646
enable ospf use-ip-router-alert 2647
enable snmp traps ospf 2648
show ospf 2649
show ospf area 2650
show ospf ase-summary 2651
show ospf interfaces 2652
show ospf interfaces detail 2654
show ospf lsdb 2655
show ospf memory 2656
show ospf neighbor 2657
show ospf virtual-link 2659
unconfigure ospf 2660

Chapter 39: OSPFv3 Commands 2662

Licensing 2663
OSPF Edge Mode 2664
clear ospfv3 counters 2664
configure ospfv3 add interface 2666
configure ospfv3 add interface all 2667
configure ospfv3 add virtual-link 2668
configure ospfv3 area add range 2669
configure ospfv3 area cost 2670
configure ospfv3 area delete range 2671
configure ospfv3 area external-filter 2672
configure ospfv3 area interarea-filter 2674
configure ospfv3 area normal 2675
configure ospfv3 area priority 2676
configure ospfv3 area stub 2677
configure ospfv3 area timer 2678



configure ospfv3 bfd 2680
configure ospfv3 delete interface 2681
configure ospfv3 delete virtual-link 2681
configure ospfv3 import-policy 2682
configure ospfv3 interface area 2684
configure ospfv3 interface cost 2685
configure ospfv3 interface priority 2686
configure ospfv3 interface timer 2687
configure ospfv3 metric-table 2688
configure ospfv3 routerid 2690
configure ospfv3 spf-hold-time 2691
configure ospfv3 virtual-link timer 2692
create ospfv3 area 2693
delete ospfv3 area 2694
disable ospfv3 2695
disable ospfv3 export 2696
enable ospfv3 2697
enable ospfv3 export 2698
show ospfv3 2700
show ospfv3 area 2702
show ospfv3 interfaces 2703
show ospfv3 lsdb 2706
show ospfv3 lsdb stats 2707
show ospfv3 memory 2709
show ospfv3 neighbor 2710
show ospfv3 virtual-link 2711
unconfigure ospfv3 2712

Chapter 40: IS-IS Commands 2714

clear isis counters 2715
clear isis counters area 2716
clear isis counters vlan 2717
configure isis add vlan 2718
configure isis area add area-address 2719
configure isis area add summary-address 2720
configure isis area area-password 2721
configure isis area delete area-address 2722
configure isis area delete summary-address 2723
configure isis area domain-password 2724
configure isis area interlevel-filter level 1-to-2 2725
configure isis area interlevel-filter level 2-to-1 2726
configure isis area is-type level 2727
configure isis area metric-style 2728
configure isis area overload-bit on-startup 2729
configure isis area system-id 2731
configure isis area timer lsp-gen-interval 2732
configure isis area timer lsp-refresh-interval 2733
configure isis area timer max-lsp-lifetime 2733
configure isis area timer restart 2734
configure isis area timer spf-interval 2735



configure isis area topology-mode	2736
configure isis circuit-type	2737
configure isis delete vlan	2738
configure isis hello-multiplier	2739
configure isis import-policy	2740
configure isis link-type	2741
configure isis mesh	2742
configure isis metric	2743
configure isis password vlan	2744
configure isis priority	2745
configure isis restart	2746
configure isis restart grace-period	2747
configure isis timer csnp-interval	2748
configure isis timer hello-interval	2749
configure isis timer lsp-interval	2750
configure isis timer restart-hello-interval	2751
configure isis timer retransmit-interval	2752
configure isis wide-metric	2752
create isis area	2753
delete isis area	2754
disable isis	2755
disable isis area adjacency-check	2756
disable isis area dynamic-hostname	2757
disable isis area export	2758
disable isis area export ipv6	2759
disable isis area originate-default	2760
disable isis area overload-bit	2761
disable isis hello-padding	2761
disable isis restart-helper	2762
enable isis	2763
enable isis area adjacency-check	2764
enable isis area dynamic-hostname	2765
enable isis area export	2766
enable isis area export ipv6	2767
enable isis area originate-default	2768
enable isis area overload-bit	2769
enable isis hello-padding	2770
enable isis restart-helper	2771
show isis	2772
show isis area	2772
show isis area summary-addresses	2773
show isis counters	2774
show isis lsdb	2775
show isis neighbors	2776
show isis topology	2777
show isis vlan	2778
unconfigure isis area	2779
unconfigure isis vlan	2780

Chapter 41: BGP Commands 2782



clear bgp flap-statistics	2784
clear bgp neighbor counters	2786
configure bgp add aggregate-address	2788
configure bgp add confederation-peer sub-AS-number	2789
configure bgp add network	2791
configure bgp as-display-format	2792
configure bgp as-number	2793
configure bgp cluster-id	2794
configure bgp confederation-id	2795
configure bgp delete aggregate-address	2797
configure bgp delete confederation-peer sub-AS-number	2798
configure bgp delete network	2799
configure bgp export shutdown-priority	2800
configure bgp import-policy	2801
configure bgp local-preference	2802
configure bgp maximum-paths	2804
configure bgp med	2805
configure bgp neighbor allowas-in	2806
configure bgp neighbor dampening	2808
configure bgp neighbor description	2810
configure bgp neighbor dont-allowas-in	2811
configure bgp neighbor maximum-prefix	2812
configure bgp neighbor next-hop-self	2814
configure bgp neighbor no-dampening	2816
configure bgp neighbor password	2818
configure bgp neighbor peer-group	2819
configure bgp neighbor route-policy	2821
configure bgp neighbor route-reflector-client	2822
configure bgp neighbor send-community	2824
configure bgp neighbor shutdown-priority	2826
configure bgp neighbor soft-reset	2827
configure bgp neighbor source-interface	2828
configure bgp neighbor timer	2829
configure bgp neighbor weight	2830
configure bgp peer-group allowas-in	2832
configure bgp peer-group dampening	2833
configure bgp peer-group dont-allowas-in	2835
configure bgp peer-group maximum-prefix	2837
configure bgp peer-group next-hop-self	2839
configure bgp peer-group no-dampening	2840
configure bgp peer-group password	2842
configure bgp peer-group remote-AS-number	2843
configure bgp peer-group route-policy	2844
configure bgp peer-group route-reflector-client	2846
configure bgp peer-group send-community	2847
configure bgp peer-group soft-reset	2848
configure bgp peer-group source-interface	2850
configure bgp peer-group timer	2851
configure bgp peer-group weight	2852



configure bgp restart 2853
configure bgp restart address-family 2854
configure bgp restart restart-time 2856
configure bgp restart stale-route-time 2857
configure bgp restart update-delay 2858
configure bgp routerid 2859
configure bgp soft-reconfiguration 2860
create bgp neighbor peer-group 2861
create bgp neighbor remote-AS-number 2863
create bgp peer-group 2865
delete bgp neighbor 2866
delete bgp peer-group 2867
disable bgp 2868
disable bgp adj-rib-out 2869
disable bgp advertise-inactive-route 2870
disable bgp aggregation 2871
disable bgp always-compare-med 2872
disable bgp community format 2873
disable bgp export 2874
disable bgp export vr 2876
disable bgp fast-external-fallover 2877
disable bgp neighbor 2878
disable bgp neighbor capability 2879
disable bgp neighbor originate-default 2881
disable bgp neighbor remove-private-AS-numbers 2882
disable bgp neighbor soft-in-reset 2883
disable bgp peer-group 2885
disable bgp peer-group capability 2886
disable bgp peer-group originate-default 2887
disable bgp peer-group remove-private-AS-numbers 2888
disable bgp peer-group soft-in-reset 2889
enable bgp 2890
enable bgp adj-rib-out 2891
enable bgp advertise-inactive-route 2892
enable bgp aggregation 2893
enable bgp always-compare-med 2894
enable bgp community format 2895
enable bgp export 2896
enable bgp export vr 2898
enable bgp fast-external-fallover 2899
enable bgp neighbor 2900
enable bgp neighbor capability 2901
enable bgp neighbor originate-default 2904
enable bgp neighbor remove-private-AS-numbers 2905
enable bgp neighbor soft-in-reset 2906
enable bgp peer-group 2908
enable bgp peer-group capability 2909
enable bgp peer-group originate-default 2911
enable bgp peer-group remove-private-AS-numbers 2912



- enable bgp peer-group soft-in-reset 2913
- show bgp 2914
- show bgp memory 2919
- show bgp neighbor 2921
- show bgp neighbor [flap-statistics | suppressed-routes] 2928
- show bgp peer-group 2931
- show bgp routes 2932
- show bgp routes summary 2937

Chapter 42: L3 VPN Commands 2939

- disable snmp traps l3vpn 2939
- enable snmp traps l3vpn 2940

Chapter 43: OpenFlow Commands 2941

- clear openflow counters 2941
- configure openflow controller 2942
- debug openflow 2944
- debug openflow show flows 2944
- disable openflow 2945
- disable openflow vlan 2946
- enable openflow 2947
- enable openflow vlan 2948
- show openflow 2949
- show openflow controller 2950
- show openflow flows 2951
- show openflow vlan 2952
- unconfigure openflow controller 2953

Chapter 44: IP Multicast Commands 2955

- clear igmp group 2957
- clear igmp snooping 2958
- clear pim cache 2959
- clear pim snooping 2960
- configure forwarding ipmc compression 2961
- configure forwarding ipmc lookup-key 2962
- configure igmp 2964
- configure igmp router-alert receive-required 2965
- configure igmp router-alert transmit 2966
- configure igmp snooping filters 2967
- configure igmp snooping flood-list 2969
- configure igmp snooping leave-timeout 2970
- configure igmp snooping timer 2971
- configure igmp snooping vlan ports add dynamic group 2973
- configure igmp snooping vlan ports add static group 2974
- configure igmp snooping vlan ports add static router 2975
- configure igmp snooping vlan ports delete static group 2976
- configure igmp snooping vlan ports delete static router 2977
- configure igmp snooping vlan ports filter 2978
- configure igmp snooping vlan ports set join-limit 2979
- configure igmp ssm-map add 2980
- configure igmp ssm-map delete 2981



configure ipmcforwarding	2982
configure ipmroute add	2983
configure ipmroute delete	2984
configure iproute add (Multicast)	2985
configure iproute delete	2987
configure mcast ipv4 cache timeout	2988
configure mvr add receiver	2989
configure mvr add vlan	2990
configure mvr delete receiver	2990
configure mvr delete vlan	2991
configure mvr mvr-address	2992
configure mvr static group	2993
configure pim add vlan	2995
configure pim border	2996
configure pim cbsr	2997
configure pim crp static	2998
configure pim crp timer	2999
configure pim crp vlan	3000
configure pim delete vlan	3002
configure pim dr-priority	3003
configure pim iproute sharing hash	3004
configure pim register-policy	3005
configure pim register-policy rp	3006
configure pim register-rate-limit-interval	3007
configure pim register-suppress-interval register-probe-interval	3008
configure pim register-checksum-to	3009
configure pim shutdown-priority	3010
configure pim spt-threshold	3011
configure pim ssm range	3012
configure pim state-refresh	3013
configure pim state-refresh timer origination-interval	3014
configure pim state-refresh timer source-active-timer	3015
configure pim state-refresh ttl	3016
configure pim timer vlan	3017
configure pim vlan trusted-gateway	3018
disable igmp	3019
disable igmp snooping	3020
disable igmp snooping vlan fast-leave	3021
disable igmp ssm-map	3022
disable ipmcforwarding	3023
disable mvr	3024
disable pim	3024
disable pim iproute sharing	3025
disable pim snooping	3026
disable pim ssm vlan	3027
enable igmp	3028
enable igmp snooping	3029
enable igmp snooping vlan fast-leave	3031
enable igmp snooping with-proxy	3032



enable igmp ssm-map	3033
enable ipmcforwarding	3034
enable mvr	3035
enable pim	3036
enable pim iproute sharing	3037
enable pim snooping	3038
enable pim ssm vlan	3038
mrinfo	3039
mtrace	3041
rtlookup	3044
rtlookup rpf	3045
show igmp	3046
show igmp group	3048
show igmp snooping	3049
show igmp snooping cache	3050
show igmp snooping vlan	3051
show igmp snooping vlan filter	3053
show igmp snooping vlan static	3053
show igmp ssm-map	3054
show ipmroute	3055
show iproute multicast	3057
show L2stats	3059
show mcast cache	3059
show mvr	3061
show mvr cache	3062
show pim	3063
show pim cache	3069
show pim snooping	3071
unconfigure igmp	3072
unconfigure igmp snooping vlan ports set join-limit	3072
unconfigure igmp ssm-map	3073
unconfigure pim	3074
unconfigure pim ssm range	3075
Chapter 45: IPv6 Multicast Commands	3077
clear mld counters	3078
clear mld group	3079
clear mld snooping	3080
configure mld	3081
configure mcast ipv6 cache timeout	3082
configure mld snooping fast-learning	3083
configure mld snooping filters	3084
configure mld snooping vlan ports add dynamic group	3085
configure mld snooping vlan ports add static group	3086
configure mld snooping vlan ports delete static group	3087
configure mld snooping vlan ports add static router	3088
configure mld snooping vlan ports delete static router	3089
configure mld snooping vlan ports filter	3090
configure mld snooping vlan ports join-limit	3092
configure mld snooping flood-list	3093



configure mld snooping leave-timeout 3095
configure mld snooping timer 3096
disable ipmcforwarding ipv6 3097
disable mld 3098
disable mld snooping 3099
enable ipmcforwarding ipv6 3100
enable mld 3100
enable mld snooping 3102
enable mld snooping with-proxy 3102
show mcast ipv6 cache 3103
show mld 3105
show mld counters 3106
show mld group 3107
show mld snooping 3108
show mld snooping vlan filter 3110
show mld snooping vlan static 3111
unconfigure mld 3112

Chapter 46: MSDP Commands 3114

clear msdp counters 3115
clear msdp sa-cache 3117
configure msdp as-display-format 3118
configure msdp max-rejected-cache 3118
configure msdp originator-id 3120
configure msdp peer default-peer 3121
configure msdp peer description 3122
configure msdp peer mesh-group 3123
configure msdp peer no-default-peer 3125
configure msdp peer password 3126
configure msdp peer sa-filter 3127
configure msdp peer sa-limit 3128
configure msdp peer source-interface 3129
configure msdp peer timer 3131
configure msdp peer ttl-threshold 3132
configure msdp sa-cache-server 3133
configure pim border 3134
create msdp mesh-group 3135
create msdp peer 3136
delete msdp mesh-group 3137
delete msdp peer 3138
disable msdp 3139
disable msdp data-encapsulation 3140
disable msdp export local-sa 3141
disable msdp peer 3141
disable msdp process-sa-request 3142
enable msdp 3144
enable msdp data-encapsulation 3144
enable msdp export local-sa 3145
enable msdp peer 3147
enable msdp process-sa-request 3148



show msdp 3149
show msdp memory 3150
show msdp mesh-group 3152
show msdp peer 3153
show msdp sa-cache 3155
unconfigure msdp sa-cache-server 3156
unconfigure pim border 3157

Chapter 47: Configuration and Image Commands 3159

clear license-info 3160
configure firmware 3161
download image 3163
enable license 3176
enable license file 3178
install bootrom 3179
install firmware 3181
install image 3187
load script 3190
run update 3192
save configuration 3193
save configuration as-script 3196
show configuration 3197
show licenses 3199
show memorycard 3200
show script output autoexec 3201
show script output default 3202
synchronize 3203
unconfigure switch 3206
uninstall image 3207
upload configuration 3209
use configuration 3213
use image 3215

Chapter 48: CNA Agent Commands 3218

clear cna-testplug counters 3218
configure cna-testplug scheduler ipaddress 3219
configure cna-testplug vlan 3220
disable cna-testplug 3221
enable cna-testplug 3222
show cna-testplug 3223

Appendix A: Troubleshooting Commands 3226

Extreme Loop Recovery Protocol 3227
disable log debug-mode 3227
clear elrp counters 3228
clear esvt traffic-test 3229
configure debug core-dumps 3230
configure elrp-client disable ports 3232
configure elrp-client one-shot 3233
configure elrp-client periodic 3234
configure forwarding fabric hash 3236



configure forwarding hash-algorithm	3238
configure forwarding hash-recursion-level	3241
configure tech-support add collector	3242
configure tech-support collector	3243
configure tech-support collector data-set	3244
configure tech-support collector frequency error-detected	3246
configure tech-support collector report	3247
configure tech-support delete collector	3248
disable elrp-client	3249
disable led locator	3250
disable log debug-mode	3251
disable tech-support collector	3252
eject memorycard	3252
enable elrp-client	3254
enable led locator	3255
enable log debug-mode	3256
enable tech-support collector	3257
nslookup	3258
run diagnostics	3259
run elrp	3265
run esvt traffic-test	3267
run tech-support report	3268
save debug tracefiles memorycard	3270
show debug	3271
show diagnostics	3272
show elrp	3276
show elrp disabled-ports	3278
show esvt traffic-test	3279
show forwarding configuration	3281
show tech support	3283
show tech-support collector	3285
stop esvt traffic-test	3287
top	3288
unconfigure elrp-client	3289
unconfigure elrp-client disable ports	3290
upload debug	3291

Index 3294

1 Introduction to the ExtremeXOS Command Reference Guide

Conventions

Related Publications

Providing Feedback to Us

This guide provides the complete syntax for all the commands available in the currently supported versions of the ExtremeXOS® software running on switches from Extreme Networks®. Included with each command is a description, the default if applicable, usage guidelines, an example, a history of the command and platform availability.

This guide is intended for use as a reference by network administrators who are responsible for installing and setting up network equipment. It assumes knowledge of Extreme Networks switch configuration. For conceptual information and guidance on configuring Extreme Networks switches, see the *ExtremeXOS Concepts Guide* for your version of the ExtremeXOS software.

Using ExtremeXOS Publications Online

You can access ExtremeXOS publications at the Extreme Networks website (www.extremenetworks.com). Publications are provided in HTML, ePub, and Adobe® PDF formats.

To navigate this guide online, use the table of contents found in the navigation bar on the left. You can also use the **prev** | **next** links at the top and bottom of the page.

To download the EXOS books in PDF or ePub format, click the links below:

[EXOS Concepts PDF](#)

[EXOS Concepts ePub](#)

[EXOS Commands PDF](#)

[EXOS Commands ePub](#)



Note

To enable cross-referenced linking between the concepts and command reference guides in the PDF, we recommend that you keep both files open on your computer desktop.

Conventions

This section discusses the conventions used in this guide.

Text Conventions

The following tables list conventions that are used throughout this guide.

Table 1: Notice Icons





Icon	Notice Type	Alerts you to...
	Note	Important features or instructions.
	Caution	Risk of personal injury, system damage, or loss of data.
	Warning	Risk of severe personal injury.
	New	This command or section is new for this release.

Table 2: Text Conventions

Convention	Description
Screen displays	This typeface indicates command syntax, or represents information as it appears on the screen.
The words enter and type	When you see the word “enter” in this guide, you must type something, and then press the Return or Enter key. Do not press the Return or Enter key when an instruction simply says “type.”
[Key] names	Key names are written with brackets, such as [Return] or [Esc]. If you must press two or more keys simultaneously, the key names are linked with a plus sign (+). Example: Press [Ctrl]+[Alt]+[Del]
<i>Words in italicized type</i>	Italics emphasize a point or denote new terms at the place where they are defined in the text. (Italics are also used when referring to publication titles.)

Platform-Dependent Conventions

Following are the platforms supported by ExtremeXOS software:

- BlackDiamond® X8 series switch.
- BlackDiamond 8800 series switches.
- Summit® family switches.
- E4G-200 Cell Site Router.



- E4G-400 Cell Site Aggregation Router.
- SummitStack™ .

Each command has a separate entry for platform availability that addresses which of these platforms support the entire feature.

In many cases, although the command is available on all platforms, each platform uses specific keywords. These keywords specific to each platform are shown in the Syntax Description and discussed in the Usage Guidelines.

Terminology

When features, functionality, or operation is specific to a switch family, the family name is used. Explanations about features and operations that are the same across all product families simply refer to the product as the “switch.”

Related Publications

The publications related to this one are:

- *ExtremeXOS Release Notes*
- *ExtremeXOS Hardware and Software Compatibility Matrix*
- *ExtremeXOS Concepts Guide*
- *BlackDiamond X8 Switch Hardware Installation Guide*
- *BlackDiamond 8800 Series Switches Hardware Installation Guide*
- *Summit Family Switches Hardware Installation Guide*
- *Extreme Networks Pluggable Interface Installation Guide*

Some ExtremeXOS software files have been licensed under certain open source licenses. Information is available at: www.extremenetworks.com/services/osl-exos.aspx.

Documentation for Extreme Networks products is available at: www.extremenetworks.com.

Providing Feedback to Us

We are always striving to improve our documentation and help you work better, so we want to hear from you! We welcome all feedback but especially want to know about:

- Content errors or confusing or conflicting information.
- Ideas for improvements to our documentation so you can find the information you need faster.
- Broken links or usability issues.

If you would like to provide feedback to the Extreme Networks Information Development team about this document, please contact us using our short online [Feedback form](#). You can also email us directly at internalinfodev@extremenetworks.com.



2 Command Reference Overview

Structure of this Guide

Platforms and Required Software Versions

Software Required

Understanding the Command Syntax

Port Numbering

Line-Editing Keys

Command History

This guide provides details of the command syntax for all ExtremeXOS commands in this ExtremeXOS version.

The guide does not provide feature descriptions, explanations of the technologies, or configuration examples. For information about the various features and technologies supported by Extreme Networks switches, see the *ExtremeXOS Concepts Guide*.

This chapter includes the following sections:

- Audience
- Structure of this Guide
- Platforms and Required Software Versions
- Understanding the Command Syntax
- Port Numbering
- Line-Editing Keys
- Command History

Structure of this Guide

This guide documents each ExtremeXOS command.

Related commands are grouped together and organized into chapters based on their most common usage. The chapters reflect the organization of the ExtremeXOS Concepts Guide. If a specific command is relevant to a wide variety of functions and could be included in a number of different chapters, we have attempted to place the command in the most logical chapter. Within each chapter, commands appear in alphabetical order.

For each command, the following information is provided:

- **Command Syntax**—The actual syntax of the command. The syntax conventions (the use of braces, for example) are defined in the section [Understanding the Command Syntax](#).
- **Description**—A brief (one sentence) summary of what the command does.
- **Syntax Description**—The definition of any keywords and options used in the command.

- **Default**—The defaults, if any, for this command. The default can be the default action of the command if optional arguments are not provided, or it can be the default state of the switch (such as for an enable/disable command).
- **Usage Guidelines**—Information to help you use the command. This may include prerequisites, prohibitions, and related commands, as well as other information.
- **Example**—Examples of the command usage, including output, if relevant.
- **History**—The version of ExtremeXOS in which the command was introduced, and version(s) where it was modified, if appropriate.
- **Platform Availability**—Platforms on which the command is available.

Platforms and Required Software Versions

The following table lists the Extreme Networks platforms that run ExtremeXOS software.

Table 3: ExtremeXOS Switches

Switch Series	Switches
BlackDiamond X Series	BlackDiamond X8.
BlackDiamond 8800 Series	BlackDiamond 8810 BlackDiamond 8806.
Mobile Backhaul Products	E4G-200 Cell Site Router E4G-400 Cell Site Aggregation Router.
Summit X430 Series	Summit X430-24t Summit X430-48t
Summit X440 Series	Summit X440-8t Summit X440-8p Summit X440-24t Summit X440-24p Summit X440-24t-10G Summit X440-24p-10G Summit X440-48t Summit X440-48p Summit X440-48t-10G Summit X440-48p-10G Summit X440-L2-24t Summit X440-L2-48t.
Summit X460 Series	Summit X460-24x Summit X460-24t Summit X460-24p Summit X460-48x Summit X460-48t Summit X460-48p.
Summit X480 Series	Summit X480-24x Summit X480-48x Summit X480-48t
Summit X670 Series	Summit X670-48x Summit X670V-48x Summit X670V-48t.
Summit X770 Series	Summit X770-32q
SummitStack	All Summit family switches, except the Summit X440-L2 model series.



Software Required

The tables in this section describe the software version required for each switch that runs ExtremeXOS software.



Note

The features available on each switch are determined by the installed feature license and optional feature packs. For more information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

The following table lists the BlackDiamond X8 series modules and the ExtremeXOS software version required to support each module.

Table 4: BlackDiamond X8 Series Modules and Required Software

Module Series Name	Modules	Minimum ExtremeXOS Software Version
BlackDiamond X8	BDX-MM1.	15.1.1
	BDXA-FM10T.	15.1.1
	BDXA-FM20T.	15.1.1
	BDXA-10G48X.	15.1.1
	BDXA-40G12X.	15.1.1
	BDXA-40G24X.	15.1.1

The following table lists the BlackDiamond 8000 series modules and the ExtremeXOS software version required to support each module.

Table 5: BlackDiamond 8000 Series Modules and Required Software

Module Series Name	Modules	Minimum ExtremeXOS Software Version
MSMs	8500-MSM24 MSM-48c 8900-MSM128.	ExtremeXOS 12.3 ExtremeXOS 12.1 ExtremeXOS 12.3.
c-series	G24Xc G48Xc 10G4Xc 10G8Xc G48Tc S-10G1Xc S-10G2Xc S-G8Xc.	ExtremeXOS 12.1 ExtremeXOS 12.1 ExtremeXOS 12.1 ExtremeXOS 12.1 ExtremeXOS 12.1 ExtremeXOS 12.1 ExtremeXOS 12.5.3 ExtremeXOS 12.1.
	8900-G96T-c 8900-10G24X-c.	ExtremeXOS 12.3.
e-series	8500-G24X-e 8500-G48T-e G48Te2.	ExtremeXOS 12.3 ExtremeXOS 12.3 ExtremeXOS 12.1.
xl-series	8900-G48X-xl 8900-G48T-xl 8900-10G8X-xl.	ExtremeXOS 12.4.
xm-series	8900-40G6X-xm.	ExtremeXOS 12.6.



The following guidelines provide additional information on the BlackDiamond 8000 series modules described in [Table 5: BlackDiamond 8000 Series Modules and Required Software](#) on page 62:

- The term BlackDiamond 8000 series modules refers to all BlackDiamond 8500, 8800, and 8900 series modules. Beginning with the ExtremeXOS 12.5 release, it does not include other modules formerly listed as original-series modules.
- Module names that are not preceded with 8500 or 8900 are BlackDiamond 8800 series modules.
- The c-series, e-series, xl-series, and xm-series names are used to distinguish between groups of modules that support different feature sets.

The following table lists the Summit family switches that run ExtremeXOS software and the minimum ExtremeXOS software version required.

Table 6: Summit Family Switches and Required Software

Switch Series	Switches	Minimum ExtremeXOS Software Version
Summit X430 Series	Summit X430-24T. Summit X430-48T.	ExtremeXOS 15.3.2.
Summit X440 Series	Summit X440-24t Summit X440-24p.	ExtremeXOS 15.1.1.
	Summit X440-8t Summit X440-8p Summit X440-24t-10G Summit X440-24p-10G Summit X440-48t Summit X440-48p Summit X440-48t-10G Summit X440-48p-10G.	ExtremeXOS 15.1.2.
	Summit X440-L2-24t Summit X440-L2-48t.	ExtremeXOS 15.2.1.
Summit X460 Series	Summit X460-24x Summit X460-24t Summit X460-24p Summit X460-48x Summit X460-48t Summit X460-48p.	ExtremeXOS 12.5.
Summit X480 Series	Summit X480-24x Summit X480-48x Summit X480-48t.	ExtremeXOS 12.4.
Summit X670	Summit X670-48x Summit X670V-48x Summit X670v-48t	ExtremeXOS 12.6.
Summit X770	Summit X770-32q	ExtremeXOS 15.4.
SummitStack	Summit family of switches.	ExtremeXOS 12.0.

Beginning with the ExtremeXOS 12.5 release, the term Summit Family Switches includes the Summit X440, X460, X480, X670, and X770 switches.

A SummitStack is a combination of up to eight Summit family switches (excluding the Summit X440-L2 series) that are connected together with stacking cables.

The mobile backhaul routers (E4G-200 and E4G-400) both require ExtremeXOS version 15.1.



Understanding the Command Syntax

This section covers the following topics:

- Access Levels
- Syntax Symbols
- Syntax Helper
- Object Names
- Command Shortcuts

Access Levels

When entering a command at the prompt, ensure that you have the appropriate privilege level.

Most configuration commands require you to have the administrator privilege level.

Syntax Symbols

You may see a variety of symbols shown as part of the command syntax.

These symbols explain how to enter the command, and you do not type them as part of the command itself. The following table summarizes command syntax symbols.



Note

ExtremeXOS software does not support the ampersand (&), left angle bracket (<), or right angle bracket (>) because they are reserved characters with special meaning in XML.

Table 7: Command Syntax Symbols

Symbol	Description
angle brackets < >	Enclose a variable or value. You must specify the variable or value. For example, in the syntax <code>configure vlan <i>vlan_name</i> ipaddress <i>ip_address</i></code> you must supply a VLAN name for <i>vlan_name</i> and an address for <i>ip_address</i> when entering the command. Do not type the angle brackets and do not include spaces within angle brackets.
square brackets []	Enclose a required value or list of required arguments. One or more values or arguments can be specified. For example, in the syntax <code>use image [<i>primary</i> <i>secondary</i>]</code> you must specify either the primary or secondary image when entering the command. Do not type the square brackets.



Table 7: Command Syntax Symbols (continued)

Symbol	Description
vertical bar	Separates mutually exclusive items in a list, one of which must be entered. For example, in the syntax <pre>configure snmp community [readonly readwrite] alphanumeric_string</pre> you must specify either the read or write community string in the command. Do not type the vertical bar.
braces { }	Enclose an optional value or a list of optional arguments. One or more values or arguments can be specified. For example, in the syntax <pre>reboot {time month day year hour min sec} {cancel} {msmslot_id} {slotslot-number node-addressnode-address stack-topology {as-standby} }</pre> you can specify either a particular date and time combination, or the keyword cancel to cancel a previously scheduled reboot. In this command, if you do not specify an argument, the command will prompt asking if you want to reboot the switch now. Do not type the braces.

Syntax Helper

The CLI has a built-in syntax helper. If you are unsure of the complete syntax for a particular command, enter as much of the command as possible and press [Tab]. The syntax helper provides a list of options for the remainder of the command, and places the cursor at the end of the command you have entered so far, ready for the next option.

If the command is one where the next option is a named component, such as a VLAN, access profile, or route map, the syntax helper also lists any currently configured names that might be used as the next option. In situations where this list might be very long, the syntax helper lists only one line of names, followed by an ellipses (...) to indicate that there are more names than can be displayed.

Some values (such as the <node-address> used in Summit stack) are lengthy, but limited in number. ExtremeXOS places these values into a “namespace.” This allows command completion on these values.

The syntax helper also provides assistance if you have entered an incorrect command.

Abbreviated Syntax

Abbreviated syntax is the shortest unambiguous allowable abbreviation of a command or parameter.

Typically, this is the first three letters of the command. If you do not enter enough letters to allow the switch to determine which command you mean, the syntax helper provides a list of the options based on the portion of the command you have entered.



Note

When using abbreviated syntax, you must enter enough characters to make the command unambiguous and distinguishable to the switch.



Object Names

All named components within a category of the switch configuration, such as VLAN, must be given a unique object name.

Object names must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but they cannot contain spaces. The maximum allowed length for a name is 32 characters.

Object names can be reused across categories (for example, STPD and VLAN names). If the software encounters any ambiguity in the components within your command, it generates a message requesting that you clarify the object you specified.



Note

If you use the same name across categories, Extreme Networks recommends that you specify the identifying keyword as well as the actual name. If you do not use the keyword, the system may return an error message.

Reserved Keywords

Keywords such as `vlan`, `stp`, and other 2nd level keywords, are determined to be reserved keywords and cannot be used as object names. This restriction applies to the specific word (`vlan`) only, while expanded versions (`vlan2`) can be used.

A complete list of the reserved keywords for ExtremeXOS 12.4.2 and later software is displayed in the *ExtremeXOS Concepts Guide*. Any keyword that is not on this list can be used as an object name. Prior to 12.4.2, all keywords were reserved, that is, none of them could be used for naming user-created objects such as VLANs.

Command Shortcuts

Components are typically named using the `create` command.

When you enter a command to configure a named component, you do not need to use the keyword of the component. For example, to create a VLAN, enter a VLAN name:

```
create vlan engineering
```

Once you have created the VLAN with a unique name, you can then eliminate the keyword `vlan` from all other commands that require the name to be entered (unless you used the same name for another category such as STPD or EAPS). For example, instead of entering the modular switch command:

```
configure vlan engineering delete port 1:3,4:6
```

you could enter the following shortcut:

```
configure engineering delete port 1:3,4:6
```



Port Numbering

Commands that require you to enter one or more port numbers use the parameter *port_list* in the syntax.

The available variables differ on a stand-alone switch, modular switch, mobile backhaul routers, and SummitStack.



Note

The keyword `all` acts on all possible ports; it continues on all ports even if one port in the sequence fails.

Stand-alone Switch Numerical Ranges

On Summit family switches, the port number is simply noted by the physical port number (such as 5).

Separate the port numbers by a dash to enter a range of contiguous numbers, and separate the numbers by a comma to enter a range of non-contiguous numbers:

- `x-y`—Specifies a contiguous series of ports on a stand-alone switch.
- `x,y`—Specifies a noncontiguous series of ports on a stand-alone switch.
- `x-y,a,d`—Specifies a contiguous series of ports and a noncontiguous series of ports on a stand-alone switch.

Modular Switch and SummitStack Numerical Ranges

On modular switches, such as the BlackDiamond 8800 series or a SummitStack, the port number is a combination of the slot number and the port number.

The nomenclature for the port number is as follows:

slot:port

Example

If an I/O module that has a total of four ports is installed in slot 2 of the chassis, the following ports are valid:

- 2:1
- 2:2
- 2:3
- 2:4

Wildcards and combinations

You can also use wildcard combinations (*) to specify multiple modular slot and port combinations. The following wildcard combinations are allowed:



- slot:*—Specifies all ports on a particular I/O module.
- slot:x-slot:y—Specifies a contiguous series of ports on a particular I/O module.
- slot:x-y—Specifies a contiguous series of ports on a particular I/O module.
- slota:x-slotb:y—Specifies a contiguous series of ports that begin on one I/O module or SummitStack node and end on another node.

Mobile Backhaul Routers

Commands operating on a *port_list* for mobile backhaul routers all use the keyword **tdm**. When the **tdm** keyword is present, the *port_list* is expanded to include only time division multiplexing (TDM) ports, omitting any Ethernet ports occurring within the *port_list* range.

Existing CLI commands without the **tdm** keyword continue to work as usual without any change, and these commands omit any TDM ports that may lie within the *port_list* range.

Line-Editing Keys

The following table describes the line-editing keys available using the CLI.

Table 8: Line-Editing Keys

Key(s)	Description
Left arrow or [Ctrl] + B	Moves the cursor one character to the left.
Right arrow or [Ctrl] + F	Moves the cursor one character to the right.
[Ctrl] + H or Backspace	Deletes character to left of cursor and shifts remainder of line to left.
Delete or [Ctrl] + D	Deletes character under cursor and shifts remainder of line to left.
[Ctrl] + K	Deletes characters from under cursor to end of line.
Insert	Toggles on and off. When toggled on, inserts text and shifts previous text to right.
[Ctrl] + A	Moves cursor to first character in line.
[Ctrl] + E	Moves cursor to last character in line.
[Ctrl] + L	Clears screen and moves cursor to beginning of line.
[Ctrl] + P or Up Arrow	Displays previous command in command history buffer and places cursor at end of command.
[Ctrl] + N or Down Arrow	Displays next command in command history buffer and places cursor at end of command.
[Ctrl] + U	Clears all characters typed from cursor to beginning of line.
[Ctrl] + W	Deletes previous word.
[Ctrl] + C	Interrupts the current CLI command execution.

Command History



ExtremeXOS “remembers” all the commands you enter.

You can display a list of these commands by using the following command:

`history`

If you use a command more than once, consecutively, the history will list only the first instance.



3 Commands for Accessing the Switch

```
clear account lockout
clear session
configure account
configure account encrypted
configure account password-policy char-validation
configure account password-policy history
configure account password-policy lockout-on-login-failures
configure account password-policy max-age
configure account password-policy min-length
configure banner
configure cli max-sessions
configure cli max-failed-logins
configure dns-client add
configure dns-client default-domain
configure dns-client delete
configure failsafe-account
configure idletimeout
configure safe-default-script
configure time
configure timezone
create account
delete account
disable cli prompting
disable cli space-completion
disable clipaging
disable idletimeout
enable cli prompting
enable cli space-completion
enable clipaging
enable idletimeout
history
ping
reboot
show accounts
show accounts password-policy
```

show banner
show banner
show failsafe-account
show switch
traceroute
unconfigure banner

This chapter describes commands used for:

- Accessing and configuring the switch including how to set up user accounts, passwords, date and time settings, and software licenses.
- Managing passwords.
- Configuring the Domain Name Service (DNS) client.
- Checking basic switch connectivity.
- Enabling and displaying licenses.
- Returning the switch to safe defaults mode.

ExtremeXOS supports the following two levels of management:

- User.
- Administrator.

A user-level account has viewing access to all manageable parameters, with the exception of:

- User account database.
- SNMP community strings.

A user-level account can change the password assigned to the account name and use the `ping` command to test device reachability.

An administrator-level account can view and change all switch parameters. It can also add and delete users and change the password associated with any account name. The administrator can disconnect a management session that has been established by way of a Telnet connection. If this happens, the user logged on by way of the Telnet connection is notified that the session has been terminated.

The DNS client in ExtremeXOS augments certain ExtremeXOS commands to accept either IP addresses or host names. For example, DNS can be used during a Telnet session when you are accessing a device or when using the `ping` command to check the connectivity of a device.

The switch offers the following commands for checking basic connectivity:

- `ping`
- `traceroute`

The `ping` command enables you to send Internet Control Message Protocol (ICMP) echo messages to a remote IP device. The `traceroute` command enables you to trace the routed path between the switch and a destination endstation.

This chapter describes commands for enabling and displaying software, security, and feature pack licenses.



clear account lockout

```
clear account [all | name] lockout
```

Description

Re-enables an account that has been locked out (disabled) for exceeding the permitted number failed login attempts. This was configured by using the `configure account password-policy lockout-on-login-failures` command.

Syntax Description

all	Specifies all users.
<i>name</i>	Specifies an account name.

Default

N/A.

Usage Guidelines

This command applies to sessions at the console port of the switch as well as all other sessions.

You can re-enable both user and administrative accounts, once they have been disabled for exceeding the 3 failed login attempts.



Note

The failsafe accounts are never locked out.

This command only clears the locked-out (or disabled) condition of the account. The action of locking out accounts following the failed login attempts remains until you turn it off by issuing the `configure account [all | name] password-policy lockout-on-login failures off` command.

Example

The following command re-enables the account `finance`, which had been locked out (disabled) for exceeding 3 consecutive failed login attempts:

```
clear account finance lockout
```

History

This command was first available in ExtremeXOS 11.2.



Platform Availability

This command is available on all platforms.

clear session

```
clear session [history | sessId | all]
```

Description

Terminates a Telnet and/or SSH2 sessions from the switch.

Syntax Description

history	Clears the chronology of sessions that were opened.
<i>sessId</i>	Specifies a session number from show session output to terminate.
all	Terminates all sessions.

Default

N/A.

Usage Guidelines

An administrator-level account can disconnect a management session that has been established by way of a Telnet connection.

You can determine the session number of the session you want to terminate by using the `show session` command. The `show session` output displays information about current Telnet and/or SSH2 sessions including:

- The session number.
- The login date and time.
- The user name.
- The type of Telnet session.
- Authentication information.

Depending on the software version running on your switch, additional session information may be displayed. The session number is the first number displayed in the `show session` output.

When invoked to clear the session history, the command clears the information about all the previous sessions that were logged. The information about the active sessions remains intact.



Example

The following command terminates session 4 from the system:

```
clear session 4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure account

```
configure account [all | name]
```

Description

Configures a password for the specified account, either user account or administrative account.

Syntax Description

all	Specifies all accounts (and future users).
<i>name</i>	Specifies an account name.

Default

N/A.

Usage Guidelines

You must create a user or administrative account before you can configure that account with a password.

Use the `create account` command to create a user account.

The system prompts you to specify a password after you enter this command. You must enter a password for this command; passwords cannot be null and cannot include the following characters: "<", ">", and "?".

Note



Once you issue this command, you cannot have a null password. However, if you want to have a null password (that is, no password on the specified account), use the `create account` command.



Passwords can have a minimum of 0 character and can have a maximum of 32 characters. Passwords are case-sensitive. User names are not case-sensitive.

**Note**

If the account is configured to require a specific password format, the minimum is 8 characters. See `configure account password-policy char-validation` for more information.

You must have administrator privileges to change passwords for accounts other than your own.

Example

The following command defines a new password green for the account marketing:

```
configure account marketing
```

The switch responds with a password prompt:

```
password: green
```

Your keystrokes will not be echoed as you enter the new password. After you enter the password, the switch will then prompt you to reenter it:

```
Reenter password: green
```

Assuming you enter it successfully a second time, the password is now changed.

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure account encrypted

```
configure account [all | name] encrypted e-password
```

Description

Encrypts the password that is entered in plain text for the specified account, either user account or administrative account.



Syntax Description

all	Specifies all accounts (and future users).
<i>name</i>	Specifies an account name.
<i>e-password</i>	Enter in plain text the string you for an encrypted password. See Usage Guidelines for more information.

Default

N/A.

Usage Guidelines

You must create a user or administrative account before you can configure that account with a password.

Use the `create account` account command to create a user account.

When you use this command, the following password that you specify in plain text is entered and displayed by the switch in an encrypted format. Administrators should enter the password in plain text. The encrypted password is then used by the switch once it encrypts the plain text password. The encrypted command should be used by the switch only to show, store, and load a system-generated encrypted password in configuration; this applies with the following commands: `save configuration`, `show configuration`, and `use configuration`.

Note



Once you issue this command, you cannot have a null password. However, if you want to have a null password (that is, no password on the specified account), use the `create account` command.

Passwords can have a minimum of 0 character and can have a maximum of 32 characters. Passwords are case-sensitive. User names are not case-sensitive.

Note



If the account is configured to require a specific password format, the minimum is 8 characters. See `configure account password-policy char-validation` for more information.

You must have administrator privileges to change passwords for accounts other than your own.

Example

The following command encrypts the password red for the account marketing:

```
configure account marketing encrypted red
```



History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

configure account password-policy char-validation

```
configure account [all | name] password-policy char-validation [none | all-char-groups]
```

Description

Requires that the user include an upper-case letter, a lower-case letter, a digit, and a symbol in the password.

Syntax Description

all	Specifies all users (and future users).
<i>name</i>	Specifies an account name.
none	Resets password to accept all formats.
all-char-groups	Specifies that the password must contain at least two characters from each of the four groups. NOTE: The password minimum length will be 8 characters if you specify this option.

Default

N/A.

Usage Guidelines

This feature is disabled by default.

Once you issue this command, each password must include at least two characters of each of the following four types:

- Upper-case A-Z.
- Lower-case a-z.
- 0-9.
- !, @, #, \$, %, ^, *, (,).

The minimum number of characters for these specifically formatted passwords is 8 characters and the maximum is 32 characters.

Use the none option to reset the password to accept all formats.



Example

The following command requires all users to use this specified format for all passwords:

```
configure account all password-policy char-validation all-char-groups
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure account password-policy history

```
configure account [all | name] password-policy history [num_passwords | none]
```

Description

Configures the switch to verify the specified number of previous passwords for the account. The user is prevented from changing the password on a user or administrative account to any of these previously saved passwords.

Syntax Description

all	Specifies all accounts (and future users).
<i>name</i>	Specifies an account name.
<i>num_passwords</i>	Specifies the number of previous passwords the system verifies for each account. The range is 1 to 10 passwords.
none	Resets the system to not remember any previous passwords.

Default

N/A.

Usage Guidelines

Use this command to instruct the system to verify new passwords against a list of all previously used passwords, once an account successfully changes a password.

The limit is the number of previous passwords that the system checks against in the record to verify the new password.

If this parameter is configured, the system returns an error message if a user attempts to change the password to one that is saved by the system (up to the configured limit) for that account; this applies



to both user and administrative accounts. This also applies to a configured password on the default admin account on the switch.

The limit of previous passwords that the system checks for previous use is configurable from 1 to 10. Using the none option disables previous password tracking and returns the system to the default state of no record of previous passwords.

Example

The following command instructs the system to verify that the new password has not been used as a password in the previous 5 passwords for the account engineering:

```
configure account engineering password-policy history 5
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure account password-policy lockout-on-login-failures

```
configure account [all | name] password-policy lockout-on-login-failures [on | off]
```

Description

Disables an account after the user has 3 consecutive failed login attempts.

Syntax Description

all	Specifies all users (and future users).
<i>name</i>	Specifies an account name.
on	Specifies an account name.
off	Resets the password to never lockout the user.

Default

N/A.



Usage Guidelines

If you are not working on SSH, you can configure the number of failed logins that trigger lockout, using the `configure cli max-failed-logins num-of-logins` command.

This command applies to sessions at the console port of the switch as well as all other sessions and to user-level and administrator-level accounts. This command locks out the user after 3 consecutive failed login attempts; the user's account must be specifically re-enabled by an administrator.

Using the `off` option resets the account to allow innumerable consecutive failed login attempts, which is the system default. The system default is that 3 failed consecutive login attempts terminate the particular session, but the user may launch another session; there is no lockout feature by default.



Note

The switch does not allow to lock out of at least one administrator account.

Example

The following command enables the account finance for lockout.

After 3 consecutive failed login attempts, the account is subsequently locked out:

```
configure account finance password-policy lockout-on-login-failures on
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure account password-policy max-age

```
configure account [all | name] password-policy max-age [num_days | none]
```

Description

Configures a time limit for the passwords for specified accounts. The passwords for the default admin account and the failsafe account do not age out.

Syntax Description

all	Specifies all accounts (and future users).
<i>name</i>	Specifies an account name.



<i>num_days</i>	Specifies the length of time that a password can be used. The range is 1 to 365 days.
none	Resets the password to never expire.

Default

N/A.

Usage Guidelines

The passwords for the default admin account and the failsafe account never expire.

The time limit is specified in days, from 1 to 365 days. Existing sessions are not closed when the time limit expires; it will not open the next time the user attempts to log in.

When a user logs into an account with an expired password, the system first verifies that the entered password had been valid prior to expiring and then prompts the user to change the password.



Note

This is the sole time that a user with a user-level (opposed to an administrator-level) account can make any changes to the user-level account.

Using the none option prevents the password for the specified account from ever expiring (it resets the password to the system default of no time limit).

Example

The following command sets a 3-month time limit for the password for the account marketing:

```
configure account marketing password-policy max-age 90
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure account password-policy min-length

```
configure account [all | name] password-policy min-length [num_characters | none]
```

Description

Requires a minimum number of characters for passwords.



Syntax Description

all	Specifies all accounts (and future users).
<i>name</i>	Specifies an account name.
<i>num_characters</i>	Specifies the minimum number of characters required for the password. The range is 1 to 32 characters. NOTE: If you configure the <code>configure account password-policy char-validation</code> parameter, the minimum length is 8 characters.
none	Resets password to accept a minimum of 0 characters. NOTE: If you configure the <code>configure account encrypted</code> parameter, the minimum length is 8 characters.

Default

N/A.

Usage Guidelines

Use this command to configure a minimum length restriction for all passwords for specified accounts.

This command affects the minimum allowed length for the next password; the current password is unaffected.

The minimum password length is configurable from 1 to 32 characters. Using the none option disables the requirement of minimum password length and returns the system to the default state (password minimum is 0 by default).

Note



If the account is configured to require a specific password format, the minimum is 8 characters. See `configure account password-policy char-validation` for more information.

Example

The following command requires a minimum of 8 letters for the password for the account management:

```
configure account management password-policy min-length 8
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.



configure banner

```
configure banner {after-login | { before-login } { acknowledge } | before-login
{acknowledge} save-to-configuration}
```

Description

Configures the banner string to be displayed for CLI screens.

Syntax Description

after-login	Specifies that a banner be displayed after login.
before-login	Specifies that a banner be displayed before login.
acknowledge	Require acknowledgement of the banner before login.
save-to-configuration	Save the before login banner to the configuration file as well as non-volatile memory.

Default

N/A.

Usage Guidelines

Use this command to configure two types of banners:

- A banner for a CLI session that displays before login.
- A banner for a CLI session that displays after login.

If no optional parameters are specified, the command defaults to configuring a banner that is displayed before the CLI session login prompt.

For each CLI session banner, you can enter up to 24 rows of 79-column text.

Press [Return] at the beginning of a line to terminate the command and apply the banner. To clear the banner, press [Return] at the beginning of the first line.



Note

The system does not wait for a keypress when you use SSH for access; this only applies to the serial console login sessions and Telnet sessions.

To disable the acknowledgement feature, use the `configure banner` command omitting the `acknowledge` parameter.

To display any configured banners, use the `show banner` command.

To unconfigure one or more configured banners, use the `unconfigure banner` command.



Example

```
# show banner
Before-Login banner:
****
**** configure banner before-login acknowledge
****
Acknowledge: Enabled
Save to      : Non-volatile memory only
# show banner
Before-Login banner:
****
**** configure banner before-login acknowledge save-to-configuration
****
Acknowledge: Enabled
Save to      : Configuration file and non-volatile memory
```

History

This command was first available in ExtremeXOS 10.1.

The acknowledge parameter was added in ExtremeXOS 11.5.

The after-login option was added in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

configure cli max-sessions

```
configure cli max-sessions num-of-sessions
```

Description

Limits number of simultaneous CLI sessions on the switch.

Syntax Description

<i>num-of-sessions</i>	Specifies the maximum number of concurrent sessions permitted. The range is 1 to 16.
------------------------	--

Default

The default is eight sessions.



Usage Guidelines

The value must be greater than 0; the range is 1 to 16.

Example

The following command limits the number of simultaneous CLI sessions to ten:

```
configure cli max-sessions 10
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure cli max-failed-logins

configure cli max-failed-logins *num-of-logins*

Description

Establishes the maximum number of failed logins permitted before the session is terminated.

Syntax Description

<i>num-of-logins</i>	Specifies the maximum number of failed logins permitted; the range is 1 to 10.
----------------------	--

Default

The default is three logins.

Usage Guidelines

The value must be greater than 0; the range is 1 to 10.

Example

The following command sets the maximum number of failed logins to five:

```
configure cli max-failed-logins 5
```



History

This command was first available in ExtremeXOS 10.1.


Platform Availability

This command is available on all platforms.

configure dns-client add

```
configure dns-client add [domain-suffix domain_name | name-server ip_address {vr vr_name}]
```

Syntax Description

domain-suffix	Specifies adding a domain suffix.
<i>domain_name</i>	Specifies a domain name.
name-server	Specifies adding a name server.
<i>ip_address</i>	Specifies an IP address for the name server.
vr	Specifies use of a virtual router. <div style="border: 1px solid #ccc; padding: 5px; margin: 5px 0;">  <p>Note User-created VRs are supported only on the platforms listed for this feature in Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i>.</p> </div>
<i>vr_name</i>	Specifies a virtual router.

Description

Adds a domain suffix to the domain suffix list or a name server to the available server list for the DNS client.

Default

N/A.

Usage Guidelines

The domain suffix list can include up to six items.

If the use of all previous names fails to resolve a name, the most recently added entry on the domain suffix list will be the last name used during name resolution. This command will not overwrite any exiting entries. If a null string is used as the last suffix in the list, and all other lookups fail, the name resolver will attempt to look up the name with no suffix.

Up to eight DNS name servers can be configured. The default value for the virtual router used by the DNS client option is VR-Default.



Example

The following command configures a domain name and adds it to the domain suffix list:

```
configure dns-client add domain-suffix xyz_inc.com
```

The following command specifies that the switch use the DNS server 10.1.2.1:

```
configure dns-client add name-server 10.1.2.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure dns-client default-domain

```
configure dns-client default-domain domain_name
```

Description

Configures the domain that the DNS client uses if a fully qualified domain name is not entered.

Syntax Description

<i>domain_name</i>	Specifies a default domain name.
--------------------	----------------------------------

Default

N/A.

Usage Guidelines

The default domain name will be used to create a fully qualified host name when a domain name is not specified.

For example, if the default domain name is set to “food.com” then when a command like “ping dog” is entered, the ping will actually be executed as “ping dog.food.com”.



Example

The following command configures the default domain name for the server:

```
configure dns-client default-domain xyz_inc.com
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure dns-client delete

```
configure dns-client delete [domain-suffix domain_name | name-server ip_address
{vr vr_name}]
```

Description

Deletes a domain suffix from the domain suffix list or a name server from the available server list for the DNS client.

Syntax Description

domain-suffix	Specifies deleting a domain suffix.
<i>domain_name</i>	Specifies a domain name.
name-server	Specifies deleting a name server.
<i>ip_address</i>	Specifies an IP address for the name server.
vr	Specifies deleting a virtual router. NOTE: User-created VRs are supported only on the platforms listed for this feature in Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i> .
<i>vr_name</i>	Specifies a virtual router.

Default

N/A.

Usage Guidelines

Specifying a domain suffix removes an entry from the domain suffix list.

If the deleted item was not the last entry in the list, all items that had been added later are moved up in the list. If no entries in the list match the domain name specified, an error message will be displayed.



The default value for the virtual router used by the DNS client option is VR-Default.

Example

The following command deletes a domain name from the domain suffix list:

```
configure dns-client delete domain-suffix xyz_inc.com
```

The following command removes a DNS server from the list:

```
configure dns-client delete name-server 10.1.2.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure failsafe-account

```
configure failsafe-account {[deny | permit] [all | control | serial | ssh {vr vr-name} | telnet {vr vr-name}]}
```

Description

Configures a name and password for the failsafe account, or restricts access to specified connection types.

Syntax Description

deny	Prohibits failsafe account usage over the specified connection type(s).
permit	Allows a failsafe account to be used over the specified connection type(s).
all	Specifies all connection types.
control	Specifies internal access between nodes in a SummitStack or between MSMs/MMs in a chassis.
serial	Specifies access over the switch console port.
ssh	Specifies access using SSH on specified or all virtual routers.
telnet	Specifies access using Telnet on specified or all virtual routers.



Default

The failsafe account is always configured.

The default connection types over which failsafe account access is permitted are the same as if **permit all** is configured.

Usage Guidelines

The failsafe account is the account of last resort to access your switch.

If you use the command with no parameters, you are prompted for the failsafe account name and prompted twice to specify the password for the account. The password does not appear on the display at any time. You are not required to know the current failsafe account and password in order to change it.

If you use the command with the permit or deny parameter, the permitted connection types are altered as specified.

The failsafe account or permitted connection types are immediately saved to NVRAM on all MSMs/MMs or active nodes in a SummitStack.

Note



The information that you use to configure the failsafe account cannot be recovered by Extreme Networks. Technical support cannot retrieve passwords or account names for this account. Protect this information carefully.

Once you enter the failsafe account name, you are prompted to enter the password. Once you successfully log in to the failsafe account, you are logged in to an admin-level account.

Example

The following example restricts usage of the failsafe account to the series console port and to access between MSMs:

```
BD-8810.1 # configure failsafe-account deny all
BD-8810.2 # configure failsafe-account permit serial
BD-8810.3 # configure failsafe-account permit control
BD-8810.4 #
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.



configure idletimeout

configure idletimeout *minutes*

Description

Configures the time-out for idle console, SSH2, and Telnet sessions.

Syntax Description

<i>minutes</i>	Specifies the time-out interval, in minutes. Range is 1 to 240 (1 minute to 4 hours).
----------------	---

Default

The default time-out is 20 minutes.

Usage Guidelines

This command configures the length of time the switch will wait before disconnecting idle console, SSH2, or Telnet sessions.

The idletimeout feature must be enabled for this command to have an effect (the idletimeout feature is enabled by default).

Example

The following command sets the time-out for idle login and console sessions to 10 minutes:

```
configure idletimeout 10
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure safe-default-script

configure safe-default-script



Description

Allows you to change management access to your device and to enhance security.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command runs an interactive script that prompts you to choose to enable or disable SNMP, Telnet, and enabled ports.

Refer to [Use Safe Defaults Mode](#) in the *Extreme Networks XOS Concepts Guide* for complete information on the safe default mode.

Once you issue this command, the system presents you with the following interactive script:

```
Telnet is enabled by default. Telnet is unencrypted and has been the
target of security exploits in the past. Would you like to disable
Telnet? [y/N]: SNMP access is enabled by default. SNMP uses no
encryption, SNMPv3 can be configured to eliminate this problem. Would you
like to disable SNMP? [y/N]: All ports are enabled by default. In some
secure applications, it maybe more desirable for the ports to be turned
off. Would you like unconfigured ports to be turned off by default?
[y/N]: Changing the default failsafe account username and password is
highly recommended. If you choose to do so, please remember the username
and password as this information cannot be recovered by Extreme Networks.
Would you like to change the failsafe account username and password now?
[y/N]: Would you like to permit failsafe account access via the
management port? [y/N]: Since you have chosen less secure management
methods, please remember to increase the security of your network by
taking the following actions: * change your admin password * change your
failsafe account username and password * change your SNMP public and
private strings * consider using SNMPv3 to secure network management
traffic
```

Example

The following command reruns the interactive script to configure management access:

```
configure safe-default-script
```



History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure time

configure time *month day year hour min sec*

Description

Configures the system date and time.

Syntax Description

<i>month</i>	Specifies the month. The range is 1-12.
<i>day</i>	Specifies the day of the month. The range is 1-31.
<i>year</i>	Specifies the year in the YYYY format. The range is 2003 to 2036.
<i>hour</i>	Specifies the hour of the day. The range is 0 (midnight) to 23 (11 pm).
<i>min</i>	Specifies the minute. The range is 0-59.
<i>sec</i>	Specifies the second. The range is 0-59.

Default

N/A.

Usage Guidelines

The format for the system date and time is as follows:

```
mm dd yyyy hh mm ss
```

The time uses a 24-hour clock format. You cannot set the year earlier than 2003 or past 2036. You have the choice of inputting the entire time/date string. If you provide one item at a time and press [Tab], the screen prompts you for the next item. Press [cr] to complete the input.



Example

The following command configures a system date of February 15, 2002 and a system time of 8:42 AM and 55 seconds:

```
configure time 02 15 2002 08 42 55
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure timezone

```
configure timezone {name tz_name} GMT_offset {autodst {name dst_timezone_ID}  
{dst_offset} {begins [every floatingday | on absoluteday] {at time_of_day} {ends  
[every floatingday | on absoluteday] {at time_of_day}}} | noautodst}
```

Description

Configures the Greenwich Mean Time (GMT) offset and Daylight Saving Time (DST) preference.

Syntax Description

<i>tz_name</i>	Specifies an optional name for this timezone specification. May be up to six alphabetic characters in length. The default is an empty string.
<i>GMT_offset</i>	Specifies a Greenwich Mean Time (GMT) offset, in + or - minutes.
autodst	Enables automatic Daylight Saving Time.
<i>dst_timezone-ID</i>	Specifies an optional name for this DST specification. May be up to six characters in length. The default is an empty string.
<i>dst_offset</i>	Specifies an offset from standard time, in minutes. Value is in the range of 1 to 60. Default is 60 minutes.
<i>floatingday</i>	Specifies the day, week, and month of the year to begin or end DST each year. Format is: <i>week day month</i> where: <i>week</i> is specified as [first second third fourth last] or 1-5. <i>day</i> is specified as [sunday monday tuesday wednesday thursday friday saturday] or 1-7 (where 1 is Sunday). <i>month</i> is specified as [january february march april may june july august september october november december] or 1-12. Default for beginning is second sunday march; default for ending is first sunday november.



<i>absoluteday</i>	Specifies a specific day of a specific year on which to begin or end DST. Format is: <i>month day year</i> where: <i>month</i> is specified as 1-12. <i>day</i> is specified as 1-31. <i>year</i> is specified as 2003-2035. The year must be the same for the begin and end dates.
<i>time_of_day</i>	Specifies the time of day to begin or end Daylight Saving Time. May be specified as an hour (0-23) or as hour:minutes. Default is 2:00.
noautodst	Disables automatic Daylight Saving Time.

Default

Autodst, beginning every second Sunday in March, and ending every first Sunday in November.

Usage Guidelines

Network Time Protocol (NTP) server updates are distributed using GMT time.

To properly display the local time in logs and other timestamp information, the switch should be configured with the appropriate offset to GMT based on geographic location.

The *GMT_offset* is specified in +/- minutes from the GMT time.

Automatic DST changes can be enabled or disabled. The default configuration, where DST begins on the second Sunday in March at 2:00 AM and ends the first Sunday in November at 2:00 AM, applies to most of North America (beginning in 2007), and can be configured with the following syntax:
`configure timezone GMT_offset autodst.`

The starting and ending date and time for DST may be specified, as these vary in time zones around the world.

- Use the **every** keyword to specify a year-after-year repeating set of dates (for example, the last Sunday in March every year).
- Use the **on** keyword to specify a non-repeating, specific date for the specified year. If you use this option, you will need to specify the command again every year.
- The **begins** specification defaults to every second Sunday in March.
- The **ends** specification defaults to every first Sunday in November.
- The ends date may occur earlier in the year than the begins date. This will be the case for countries in the Southern Hemisphere.
- If you specify only the starting or ending time (not both) the one you leave unspecified will be reset to its default.
- The *time_of_day* specification defaults to 2:00.
- The timezone IDs are optional. They are used only in the display of timezone configuration information in the `show switch` command.

To disable automatic DST changes, re-specify the GMT offset using the `noautodst` option: `configure timezone gmt_offset noautodst.`



Greenwich Mean Time offsets

NTP updates are distributed using GMT time. To properly display the local time in logs and other timestamp information, the switch should be configured with the appropriate offset to GMT based on geographical location. `configure timezone` on page 94 describes the GMT offsets.

GMT Offset in Hours	GMT Offset in Minutes	Common Time Zone References	Cities
+0:00	+0	GMT - Greenwich Mean UT or UTC - Universal (Coordinated) WET - Western European	London, England; Dublin, Ireland; Edinburgh, Scotland; Lisbon, Portugal; Reykjavik, Iceland; Casablanca, Morocco
-1:00	-60	WAT - West Africa	Cape Verde Islands
-2:00	-120	AT - Azores	Azores
-3:00	-180		Brasilia, Brazil; Buenos Aires, Argentina; Georgetown, Guyana;
-4:00	-240	AST - Atlantic Standard	Caracas; La Paz
-5:00	-300	EST - Eastern Standard	Bogota, Columbia; Lima, Peru; New York, NY, Trevor City, MI USA
-6:00	-360	CST - Central Standard	Mexico City, Mexico
-7:00	-420	MST - Mountain Standard	Saskatchewan, Canada
-8:00	-480	PST - Pacific Standard	Los Angeles, CA, Cupertino, CA, Seattle, WA USA
-9:00	-540	YST - Yukon Standard	
-10:00	-600	AHST - Alaska-Hawaii Standard CAT - Central Alaska HST - Hawaii Standard	
-11:00	-660	NT - Nome	
-12:00	-720	IDLW - International Date Line West	
+1:00	+60	CET - Central European FWT - French Winter MET - Middle European MEWT - Middle European Winter SWT - Swedish Winter	Paris, France; Berlin, Germany; Amsterdam, The Netherlands; Brussels, Belgium; Vienna, Austria; Madrid, Spain; Rome, Italy; Bern, Switzerland; Stockholm, Sweden; Oslo, Norway
+2:00	+120	EET - Eastern European, Russia Zone 1	Athens, Greece; Helsinki, Finland; Istanbul, Turkey; Jerusalem, Israel; Harare, Zimbabwe
+3:00	+180	BT - Baghdad, Russia Zone 2	Kuwait; Nairobi, Kenya; Riyadh, Saudi Arabia; Moscow, Russia; Tehran, Iran
+4:00	+240	ZP4 - Russia Zone 3	Abu Dhabi, UAE; Muscat; Tblisi; Volgograd; Kabul
+5:00	+300	ZP5 - Russia Zone 4	
+5:30	+330	IST - India Standard Time	New Delhi, Pune, Allahabad, India
+6:00	+360	ZP6 - Russia Zone 5	
+7:00	+420	WAST - West Australian Standard	
+8:00	+480	CCT - China Coast, Russia Zone 7	



GMT Offset in Hours	GMT Offset in Minutes	Common Time Zone References	Cities
+9:00	+540	JST - Japan Standard, Russia Zone 8	
+10:00	+600	EAST - East Australian Standard GST - Guam Standard Russia Zone 9	
+11:00	+660		
+12:00	+720	IDLE - International Date Line East NZST - New Zealand Standard NZT - New Zealand	Wellington, New Zealand; Fiji, Marshall Islands

For name creation guidelines and a list of reserved names, see [Object Names](#) in the *ExtremeXOS Concepts Guide*.

Example

The following command configures GMT offset for Mexico City, Mexico and disables automatic DST:

```
configure timezone -360 noautodst
```

The following four commands are equivalent, and configure the GMT offset and automatic DST adjustment for the US Eastern timezone, with an optional timezone ID of EST:

```
configure timezone name EST -300 autodst name EDT 60 begins every second
sunday march at 2:00 ends every first sunday november at 2:00
configure timezone name EST -300 autodst name EDT 60 begins every 1 1 4 at
2:00 ends every 5 1 10 at 2:00
configure timezone name EST -300 autodst name EDT
configure timezone -300 autodst
```

The following command configures the GMT offset and automatic DST adjustment for the Middle European timezone, with the optional timezone ID of MET:

```
configure timezone name MET 60 autodst name MDT begins every last sunday
march at 1 ends every last sunday october at 1
```

The following command configures the GMT offset and automatic DST adjustment for New Zealand. The ending date must be configured each year because it occurs on the first Sunday on or after March 5:

```
configure timezone name NZST 720 autodst name NZDT 60 begins every first
sunday october at 2 ends on 3/16/2002 at 2
```

History

This command was first available in ExtremeXOS 10.1.



Platform Availability

This command is available on all platforms.

create account

```
create account [admin | user | lawful-intercept] account-name {encrypted
password}
```

Description

Creates a new user account.

Syntax Description

admin	Specifies an access level for admin account type. This user has read and write privileges.
user	Specifies an access level for user account type. This user has read-only privileges.
lawful-intercept	Specifies an access level for lawful intercept account type.
<i>account-name</i>	Specifies a new user account name.
encrypted	Specifies the encrypted option.
<i>password</i>	Specifies a user password.


Default

N/A.

User Account Levels

By default, the switch is configured with two accounts with the access levels shown in the table below.



Account Name	Access Level
admin	This user can access and change all manageable parameters. The admin account cannot be deleted.
user	This user can view (but not change) all manageable parameters, with the following exceptions: <ul style="list-style-type: none"> You cannot view the user account database. You cannot view the SNMP community strings. You cannot view SSL settings. This user has access to the ping command.
lawful-intercept	This user has special lawful intercept and read-only privileges. <div style="border: 1px solid black; padding: 5px; margin-top: 10px;">  <p>Note Only a single lawful-intercept account can exist at any one time on the system.</p> </div>

You can use the default names (admin and user), or you can create new names and passwords for the accounts. Default accounts do not have passwords assigned to them. For name creation guidelines and a list of reserved names, see [Object Names](#) in the *ExtremeXOS Concepts Guide*.

Usage Guidelines

The switch can have a total of 16 user accounts.

The system must have one administrator account.

When you use the encrypted keyword, the following password that you specify in plain text is entered and displayed by the switch in an encrypted format. Administrators should not use the encrypted option and should enter the password in plain text. The encrypted option is used by the switch after encrypting the plain text password. The encrypted option should be used by the switch only to show, store, and load a system-generated encrypted password in configuration; this applies with the following commands: `save configuration`, `show configuration`, and `use configuration`.

The system prompts you to specify a password after you enter this command and to reenter the password. If you do not want a password associated with the specified account, press [Enter] twice.

You must have administrator privileges to change passwords for accounts other than your own. User names are not case-sensitive. Passwords are case-sensitive. User account names must have a minimum of 1 character and can have a maximum of 32 characters. Passwords must have a minimum of 0 characters and can have a maximum of 32 characters.

Note



If the account is configured to require a specific password format, the minimum is eight characters. See `configure account password-policy char-validation` for more information.



Example

The following command creates a new account named John2 with administrator privileges:

```
create account admin John2
```

History

This command was first available in ExtremeXOS 10.1.

The encrypted option was added in ExtremeXOS 11.5.

The lawful intercept option was added in ExtremeXOS 15.3.2.

Platform Availability

This command is available on all platforms.

delete account

```
delete account name
```

Description

Deletes a specified user account.

Syntax Description

<i>name</i>	Specifies a user account name.
-------------	--------------------------------

Default

N/A.

Usage Guidelines

Use the show accounts command to determine which account you want to delete from the system.

The show accounts output displays the following information in a tabular format:

- The user name.
- Access information associated with each user.
- User login information.
- Session information.

Depending on the software version running on your switch and the type of switch you have, additional account information may be displayed.



You must have administrator privileges to delete a user account. The system must have one administrator account; the command will fail if an attempt is made to delete the last administrator account on the system.

To ensure security, change the password on the default account, but do not delete it. The changed password will remain intact through configuration uploads and downloads.

If you must delete the default account, first create another administrator-level account.

Example

The following command deletes account John2:

```
delete account John2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable cli prompting

disable cli prompting

Description

Disables CLI prompting for the session.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Use this command to have all CLI user prompts automatically continue with the default answer.

This applies to the current session only.

To re-enable CLI prompting for the session, use the `enable cli prompting` command.



To view the status of CLI prompting on the switch, use the `show management` command.

Example

The following command disables prompting:

```
disable cli prompting
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

disable cli space-completion

disable cli space-completion

Description

Disables the ExtremeXOS feature that completes a command automatically with the spacebar. If you disable this feature, the [Tab] key can still be used for auto-completion.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables using the spacebar to automatically complete a command:

```
disable cli space-completion
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable clipaging

disable clipaging

Description

Disables pausing at the end of each show screen.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

The command line interface (CLI) is designed for use in a VT100 environment.

Most show command output will pause when the display reaches the end of a page. This command disables the pause mechanism and allows the display to print continuously to the screen.

CLI paging is only active on a per-shell session basis. In other words, when you enable or disable CLI paging from within the current configuration, it only affects that session. For new or existing sessions, paging is enabled by default. This setting cannot be saved.

To view the status of CLI paging on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for CLI paging.

Example

The following command disables clipaging and allows you to print continuously to the screen:

```
disable clipaging
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable idletimeout

disable idletimeout

Description

Disables the timer that disconnects idle sessions from the switch.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Timeout 20 minutes.

Usage Guidelines

When idle time-outs are disabled, console sessions remain open until the switch is rebooted or until you logoff.

Telnet sessions remain open until you close the Telnet client.

If you have an SSH2 session and disable the idle timer, the SSH2 connection times out after 61 minutes of inactivity.

To view the status of idle time-outs on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for idle time-outs.

Example

The following command disables the timer that disconnects all sessions to the switch:

```
disable idletimeout
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable cli prompting

enable cli prompting

Description

Enables CLI prompting for the session.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Use this command to enable CLI prompting from a disabled state.

To view the status of CLI prompting on the switch, use the `show management` command.

Example

The following command enables prompting:

```
enable cli prompting
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.



enable cli space-completion

enable cli space-completion

Description

Enables the ExtremeXOS feature that completes a command automatically with the spacebar. The [Tab] key can also be used for auto-completion.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

CLI space auto completion cannot be enabled while CLI scripting is enabled with the `enable cli scripting` command.

Example

The following command enables using the spacebar to automatically complete a command:

```
enable cli space-completion
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable clipaging

enable clipaging

Description

Enables the pause mechanism and does not allow the display to print continuously to the screen.



Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

The command line interface (CLI) is designed for use in a VT100 environment.

Most show command output will pause when the display reaches the end of a page.

To view the status of CLI paging on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for CLI paging.

If CLI paging is enabled and you use the `show tech` command to diagnose system technical problems, the CLI paging feature is disabled.

CLI paging is only active on a per-shell session basis. In other words, when you enable or disable CLI paging from within the current configuration, it only affects that session. For new or existing sessions, paging is enabled by default. This setting cannot be saved.

Example

The following command enables clipaging and does not allow the display to print continuously to the screen:

```
enable clipaging
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable idletimeout

```
enable idletimeout
```

Description

Enables a timer that disconnects Telnet, SSH2, and console sessions after a period of inactivity (20 minutes is default).



Syntax Description

This command has no arguments or variables.

Default

Enabled.

Timeout 20 minutes.

Usage Guidelines

You can use this command to ensure that a Telnet, Secure Shell (SSH2), or console session is disconnected if it has been idle for the required length of time.

This ensures that there are no hanging connections.

To change the period of inactivity that triggers the timeout for a Telnet, SSH2, or console session, use the `configure timezone` command.

To view the status of idle timeouts on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for idle timeouts. You can configure the length of the timeout interval.

Example

The following command enables a timer that disconnects any Telnet, SSH2, and console sessions after 20 minutes of inactivity:

```
enable idletimeout
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

history

history

Description

Displays a list of all the commands entered on the switch.



Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

ExtremeXOS software “remembers” all the commands you entered on the switch.

Use the history command to display a list of these commands.

Example

The following command displays all the commands entered on the switch:

```
history
```

If you use a command more than once, consecutively, the history will only list the first instance.

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

ping

```
ping {count count {start-size start-size} | continuous {start-size start-size} |
{start-size start-size {end-size end-size}}} {udp} {dont-fragment} {ttl ttl} {tos
tos} {interval interval} {vr vrid} {ipv4 host | ipv6 host} {from} {with record-
route}
```

Description

Enables you to send User Datagram Protocol (UDP) or Internet Control Message Protocol (ICMP) echo messages or to a remote IP device.



Syntax Description

<i>count</i>	Specifies the number of ping requests to send.
<i>start-size</i>	Specifies the size, in bytes, of the packet to be sent, or the starting size if incremental packets are to be sent.
continuous	Specifies that UDP or ICMP echo messages to be sent continuously. This option can be interrupted by pressing [Ctrl] + C.
<i>end-size</i>	Specifies an end size for packets to be sent.
udp	Specifies that the ping request should use UDP instead of ICMP.
dont-fragment	Sets the IP to not fragment the bit.
ttl	Sets the TTL value.
<i>tos</i>	Sets the TOS value.
<i>interval</i>	Sets the time interval between sending out ping requests.
vr	Specifies the virtual route to use for sending out the echo message. If not specified, VR-Default is used. NOTE: User-created VRs are supported only on the platforms listed for this feature in Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i> .
ipv4	Specifies IPv4 transport.
ipv6	Specifies IPv6 transport. NOTE: If you are contacting an IPv6 link local address, you must specify the VLAN you are sending the message from: ping ipv6 link-local address %vlan_name host .
<i>host</i>	Specifies a host name or IP address (either v4 or v6).
from	Uses the specified source address. If not specified, the address of the transmitting interface is used.
with record-route	Sets the traceroute information.

Default

N/A.

Usage Guidelines

The `ping` command is used to test for connectivity to a specific host.

You use the `ipv6` variable to ping an IPv6 host by generating an ICMPv6 echo request message and sending the message to the specified address. If you are contacting an IPv6 link local address, you must specify the VLAN you are sending the message from, as shown in the following example (you must include the % sign):

```
ping ipv6 link-local address %vlan_name host.
```

The `ping` command is available for both the user and administrator privilege level.

When the IPv6 host ping fails, the following error message displays:



Error: cannot determine outgoing interface. Link local address must be of form LLA%*vlan_name*.

Example

The following command enables continuous ICMP echo messages to be sent to a remote host:

```
ping continuous 123.45.67.8
```

History

This command was first available in ExtremeXOS 10.1.

The IPv6 variable was added in ExtremeXOS 11.2.

IPv6 error message modified in ExtremeXOS 15.2.

Platform Availability

This command is available on all platforms.

reboot

For modular switches:

```
reboot {time month day year hour min sec} {cancel} {msm slot_id} {slot slot-number }
```

For Summit series switches and SummitStack:

```
reboot {[time mon day year hour min sec] | cancel} {slot slot-number }
```

Description

Reboots the switch, SummitStack, or the module in the specified slot at a specified date and time.

Syntax Description

time	Specifies a reboot date in mm dd yyyy format and reboot time in hh mm ss format.
cancel	Cancels a previously scheduled reboot.
msm	Specifies rebooting the MSM module. NOTE: This variable is available only on modular switches.
<i>slot_id</i>	Specifies the slot--A or B--for an MSM module. NOTE: This variable is available only on modular switches.
<i>slot-number</i>	Specifies the slot number currently being used by the active stack node that is to be rebooted. NOTE: This variable is available only on SummitStack.



Default

N/A.

Usage Guidelines

If you do not specify a reboot time, the switch will reboot immediately following the command, and any previously scheduled reboots are cancelled.

Prior to rebooting, the switch returns the following message:

```
Do you want to save configuration changes to primary and reboot? (y -  
save and reboot, n - reboot without save, <cr> - cancel command)
```

To cancel a previously scheduled reboot, use the cancel option.

Modular switches only

The modules that can be rebooted are management switch fabric modules (MSM)/management modules (MM).

BlackDiamond 8800 series switches only

On the BlackDiamond 8800 series switches, if your default BootROM image becomes corrupted, you can force the MSM to boot from an alternate BootROM image by inserting a sharp object into the “A” and “R” holes on the MSM and applying slight pressure.

Refer to the hardware documentation for information on the MSM.

The reboot MSM option on the BlackDiamond 8800 series switches affects the entire module.

SummitStack only

The `reboot` command used without any parameters on the master node reboots all members of the same active topology to which the master node belongs.

This version can only be used on the master node.

The `reboot slot slot-number` command can be used on any active node. The command will reboot the active node that is currently using the specified slot number in the same active topology as the issuing node. This variation cannot be used on a node that is not in stacking mode.

The `reboot node-address node-address` command can be used on any node whether or not the node is in stacking mode. It will reboot the node whose MAC address is supplied.

The `reboot stack-topology {as-standby}` command reboots every node in the stack topology. The command can be issued from any node whether or not the node is in stacking mode. If the as-standby option is used, every node in the stack topology restarts with master-capability disabled. This option is useful when manually resolving a dual master situation.



Example

The following command reboots the switch at 8:00 AM on April 15, 2005:

```
reboot time 04 15 2005 08 00 00
```

History

This command was first available in ExtremeXOS 10.1.

The alternate BootROM image was added in ExtremeXOS 11.1.

The slot, node-address, stack-topology, and as-standby options were added in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

show accounts

show accounts

Description

Displays user account information for all users on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

You need to create a user account using the create account command before you can display user account information.

To view the accounts that have been created, you must have administrator privileges.

This command displays the following information in a tabular format:

- **User Name**—The name of the user. This list displays all of the users who have access to the switch.
- **Access**—This may be listed as R/W for read/write or RO for read only.
- **Login OK**—The number of logins that are okay.



- **Failed**—The number of failed logins.
- **Accounts locked out**—Account configured to be locked out after three consecutive failed login attempts (using the `configure account password-policy lockout-on-login-failures` command).

**Note**

This command does not show the failsafe account.

Example

The following command displays user account information on the switch:

```
show accounts pppuser
```

Output from this command looks similar to the following:

User Name	Access	LoginOK	Failed
admin	R/W	3	1
user	RO	0	0
dbackman	R/W	0	0
ron*	RO	0	0
nocteam	RO	0	0

(*) - Account locked

The following command displays the lawful intercept account distinguished by the "R/L" displayed in the Access column:

```
* (Private) X440e-24t.9 # show accounts
```

User Name	Access	LoginOK	Failed
admin	R/W	6	0
user	RO	0	0
myLIuser	R/L	N/A	N/A

History

This command was first available in ExtremeXOS 11.0.

Lawful intercept output was added in ExtremeXOS 15.3.2.

Platform Availability

This command is available on all platforms.



show accounts password-policy

show accounts password-policy

Description

Displays password policy information for all users on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

To view the password management information, you must have administrator privileges.

The `show accounts password-policy` command displays the following information in a tabular format:

- Global password management parameters applied to new accounts upon creation:
 - Maximum age—The maximum number of days for the passwords to remain valid.
 - History limit—The number of previous password that the switch scans prior to validating a new password.
 - Minimum length—The minimum number of characters in passwords.
 - Character validation—The passwords must be in the specific format required by the `configure account password-policy char-validation` command.
 - Lockout on login failures—If enabled, the system locks out users after 3 failed login attempts.
 - Accounts locked out—Number of accounts locked out.
- **User Name**—The name of the user. This list displays all of the users who have access to the switch.
- **Password Expiry Date**—Date the password for this account expires; may be blank.
- **Password Max. age**—The number of days originally allowed to passwords on this account; may show None.
- **Password Min. length**—The minimum number of characters required for passwords on this account; may show None.
- **Password History Limit**—The number of previous passwords the system scans to disallow duplication on this account; may show None.



Example

The following command displays the password management parameters configured for each account on the switch:

```
show accounts password-policy
```

Output from this command looks similar to the following:

```
-----
Accounts global configuration(applied to new accounts on creation)
-----
Password Max. age           : None
Password History limit     : None
Password Min. length       : None
Password Character Validation : Disabled
Accts. lockout on login failures: Disabled
Accounts locked out        : No
-----

User Name      Password  Password Password Password Password Flags
Expiry      Max. age Min. len History
Date                               Limit
-----
admin                None      None      None      None      ---
user                 None      None      None      None      ---
test Apr-17-2005    12       32       9        C--
-----

Flags: (C) Password character validation enabled, (L) Account locked out
(l) Account lockout on login failures enabled
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show banner

```
show banner { after-login | before-login }
```

Description

Displays the user-configured banners.



Syntax Description

after-login	Specifies the banner that is displayed after login.
before-login	Specifies the banner that is displayed before login.

Default

N/A.

Usage Guidelines

Use this command to display specific configured CLI banners.

If no keywords are specified, all configured banners are displayed. To display a specific banner, use the before-login or after-login keyword.

Example

The following command displays the configured CLI switch banners:

```
show banner
```

Output from this command varies depending on your configuration; the following is one example:

```
Before-login banner:
Extreme Networks Summit X670 Switch
#####
Unauthorized Access is strictly prohibited.
Violators will be prosecuted
#####
Acknowledge: Enabled
After-login banner:
Press any key to continue
```

History

This command was first available in ExtremeXOS 10.1.

The after-login option was added in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

show banner

```
show banner { after-login | before-login }
```



Description

Displays the user-configured banners.

Syntax Description

after-login	Specifies the banner that is displayed after login.
before-login	Specifies the banner that is displayed before login.

Default

N/A.

Usage Guidelines

Use this command to display specific configured CLI banners.

If no keywords are specified, all configured banners are displayed. To display a specific banner, use the before-login or after-login keyword.

Example

The following command displays the configured CLI switch banners:

```
show banner
```

Output from this command varies depending on your configuration; the following is one example:

```
Before-login banner:
Extreme Networks Summit X670 Switch
#####
Unauthorized Access is strictly prohibited.
Violators will be prosecuted
#####
Acknowledge: Enabled
After-login banner:
Press any key to continue
```

History

This command was first available in ExtremeXOS 10.1.

The after-login option was added in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.



show failsafe-account

show failsafe-account

Description

Displays whether the user configured a username and password for the failsafe account or shows the configured connection type access restrictions.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to view the failsafe account configuration.

The command shows the access permissions and whether or not the user configured a username and password. It does not show the configured username or password.

Example

The following command displays the failsafe account configuration:

```
show failsafe-account
```

Output from this command looks similar to the following when a failsafe account username and password have been configured with all connections types permitted for failsafe account access:

```
BD-8810.7 # show failsafe-account
User-Specified Failsafe Account Username and Password are in effect for these
connection types:
  - Serial Console
  - Control Fabric (inter-node)
  - Mgmt VR Telnet
  - Mgmt VR SSH
  - User VR Telnet
  - User VR SSH
BD-8810.8 #
```

History

This command was first available in ExtremeXOS 12.0.



Platform Availability

This command is available on all platforms.

show switch

```
show switch {detail}
```

Description

Displays the current switch information.

On a SummitStack, this command displays the Master and Backup node information if executed on the Master, and displays the current node and the Master node information if executed on any other node.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The show switch command displays:

- sysName, sysLocation, sysContact.
- MAC address.
- System type.
- System health check.
- Recovery mode.
- Watchdog state.
- Current date, time, system boot time, and time zone configuration.
- Any scheduled reboot information.
- System up time.
- Master and Backup information (available only on modular switches and SummitStack).
- Current state (available only on stand-alone switches).
 - OPERATIONAL.
 - OPERATIONAL (OverHeat).
 - FAILED.
- Software image information (primary/secondary image and version).
- Configuration information (primary/secondary configuration and version).

This information may be useful for your technical support representative if you have a problem.



On a SummitStack, the System UpTime may be useful when manually resolving the dual master situation. For more information, see the “Eliminating a Dual Master Situation Manually” section in the *ExtremeXOS Concepts Guide*.

Depending on the software version running on your switch, additional or different switch information may be displayed.

On a stack the following additional information will be available:

- System Type.
- System UpTime.
- Details of Master and Backup, or current node and Master.

Example

The following command displays current switch information:

```
show switch
```

Output from this command on the modular switches looks similar to the following:

```
SysName:          BD-8810Rack3
SysLocation:
SysContact:       support@extremenetworks.com, +1 888 257 3000
System MAC:       00:04:96:1D:00:C0
System Type:      BD-8810
SysHealth check:  Enabled (Normal)
Recovery Mode:    All
System Watchdog:  Enabled
Current Time:     Fri Feb 13 02:25:24 1925
Timezone:         [Auto DST Disabled] GMT Offset: 0 minutes, name is UTC.
Boot Time:        Wed Feb 11 21:39:56 1925
Boot Count:       159
Next Reboot:      None scheduled
System UpTime:    1 day 4 hours 45 minutes 28 seconds
Slot:             MSM-A *                               MSM-B
-----
Current State:    MASTER                                BACKUP (In Sync)
Image Selected:   secondary                             secondary
Image Booted:     primary                               primary
Primary ver:      12.0.0.4                              12.0.0.4
Secondary ver:    12.0.0.4                              12.0.0.4
Config Selected:  primary.cfg                             primary.cfg
Config Booted:    primary.cfg                             primary.cfg
primary.cfg       Created by ExtremeXOS version 11.6.0.30
574246 bytes saved on Wed Jul 30 19:39:55 1924
```

Output from this command on the stand-alone Summit family switch looks similar to the following:

```
SysName:          SummitX
SysLocation:
SysContact:       support@extremenetworks.com, +1 888 257 3000
```



```

System MAC:          00:04:96:26:6B:EC
System Type:        SummitX
SysHealth check:    Enabled (Normal)
Recovery Mode:      All
System Watchdog:    Enabled
Current Time:        Wed Apr 25 21:17:18 2012
Timezone:           [Auto DST Disabled] GMT Offset: 0 minutes, name is UTC.
Boot Time:          Wed Apr 25 21:13:54 2012
Boot Count:         951
Next Reboot:        None scheduled
System UpTime:      3 minutes 24 seconds
Current State:      OPERATIONAL
Image Selected:     secondary
Image Booted:       secondary
Primary ver:        12.0.0.4
Secondary ver:      12.0.0.4
Config Selected:    primary.cfg
Config Booted:      Factory Default
primary.cfg         Created by ExtremeXOS version 12.0.0.4
156281 bytes saved on Mon Apr 23 17:10:11 2012

```

The show switch detail command displays the same information shown above.

Output from this command on a stack looks similar to the following:

```

SysName:            Stack
SysLocation:
SysContact:         support@extremenetworks.com, +1 888 257 3000
System MAC:         02:04:96:27:B7:41
System Type:        X (Stack)
SysHealth check:    Enabled (Normal)
Recovery Mode:      All
System Watchdog:    Enabled
Current Time:        Tue Jan 30 14:22:41 2007
Timezone:           [Auto DST Disabled] GMT Offset: 0 minutes, name is UTC.
Boot Time:          Mon Jan 29 21:51:38 2007
Boot Count:         317
Next Reboot:        None scheduled
System UpTime:      16 hours 31 minutes 3 seconds
Slot:               Slot-4 *                               Slot-5
-----
Current State:      MASTER                                 BACKUP (In Sync)
Image Selected:     secondary                             secondary
Image Booted:       secondary                             secondary
Primary ver:        12.0.0.10                             12.0.0.10
Secondary ver:      12.0.0.13                             12.0.0.13
Config Selected:    primary.cfg
Config Booted:      primary.cfg
primary.cfg         Created by ExtremeXOS version 12.0.0.10
139108 bytes saved on Fri Jan 26 22:56:40 2007

```

History

This command was first available in ExtremeXOS 10.1.



This command was updated to support stacking in ExtremeXOS 12.0 and the System Type was added to the output from this version.

Platform Availability

This command is available on all platforms.

traceroute

```
traceroute {vr vrid} {ipv4 host} {ipv6 host} {ttl number} {from from} {port port} | icmp}
```

Description

Enables you to trace the routed path between the switch and a destination endstation.

Syntax Description

vr	Specifies a virtual router. NOTE: User-created VRs are supported only on the platforms listed for this feature in Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i> .
<i>vrid</i>	Specifies which virtual router. NOTE: User-created VRs are supported only on the platforms listed for this feature in Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i> .
ipv4	Specifies IPv4 transport.
ipv6	Specifies IPv6 transport.
<i>host</i>	Specifies the host of the destination endstation.
ttl <i>number</i>	Configures the switch to trace up to the time-to-live number of the switch.
from <i>from</i>	Uses the specified source address in the ICMP packet. If not specified, the address of the transmitting interface is used.
port <i>port</i>	Specifies the UDP port number.
icmp	Configures the switch to send ICMP echo messages to trace the routed path between the switch and a destination endstation.

Default

N/A.

Usage Guidelines

Use this command to trace the routed path between the switch and a destination endstation.

Each router along the path is displayed.



Example

The following command enables the traceroute function to a destination of 123.45.67.8:

```
traceroute 123.45.67.8
```

The following is sample output that displays when the traceroute fails:

```
traceroute to 10.209.10.37, 30 hops max
 1  0.0.0.0                                * !u          * !u          * !u
--- Packet Response/Error Flags ---
(*) No response, (!N) ICMP network unreachable, (!H) ICMP host unreachable,
(!P) ICMP protocol unreachable, (!F) ICMP fragmentation needed,
(!S) ICMP source route failed, (!u) Transmit error, network unreachable,
(!f) Transmit error, fragmentation needed, (!t) General transmit error
```

History

This command was first available in ExtremeXOS 10.1.

The IPv6 variable was added in ExtremeXOS 11.2.

The display when the command fails was added in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

unconfigure banner

```
unconfigure banner { after-login | before-login }
```

Description

Unconfigures a specified banner from CLI screens.

Syntax Description

after-login	Specifies the banner that is displayed after login.
before-login	Specifies the banner that is displayed before login.

Default

N/A.



Usage Guidelines

Use this command to unconfigure one of two different types of banners:

- CLI session before login.
- CLI session after login.

If no optional parameters are specified, all configured banners are erased. To delete a specific banner, the before-login or after-login keyword must be used.

Banners can also be cleared by configuring a banner with only a <ret> or \n character.

Example

The following command clears the after-login banner, Welcome to the switch:

```
unconfigure banner after-login [Return]
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.



4 Commands for Managing the Switch

NTP

PTP

SNMP

Telnet

TFTP

System Redundancy with Dual Management Modules Installed, Modular Switches Only

Power Supply Management

SNTP

clear cdp counters

clear cdp neighbor

clear network-clock ptp counters

configure cdp device-id

configure cdp frequency

configure cdp hold-time

configure node priority

configure ntp local-clock none

configure ntp local-clock stratum

configure ntp restrict-list

configure ntp server/peer add

configure ntp server/peer delete

configure power led motion-detector

configure power monitor

configure power supply

configure snmp access-profile

configure snmp add community

configure snmp add trapreceiver

configure snmp delete community

configure snmp delete trapreceiver

configure snmp sysContact

configure snmp sysLocation

configure snmp sysName

configure snmpv3 add access

configure snmpv3 add community

configure snmpv3 add filter

configure snmpv3 add filter-profile

```
configure snmpv3 add group user
configure snmpv3 add mib-view
configure snmpv3 add notify
configure snmpv3 add target-addr
configure snmpv3 add target-params
configure snmpv3 add user
configure snmpv3 add user clone-from
configure snmpv3 delete access
configure snmpv3 delete community
configure snmpv3 delete filter
configure snmpv3 delete filter-profile
configure snmpv3 delete group user
configure snmpv3 delete mib-view
configure snmpv3 delete notify
configure snmpv3 delete target-addr
configure snmpv3 delete target-params
configure snmpv3 delete user
configure snmpv3 engine-boots
configure snmpv3 engine-id
configure snmpv3 target-addr retry
configure snmpv3 target-addr timeout
configure snmp-client
configure snmp-client update-interval
configure ssh2 access-profile
configure telnet access-profile
configure telnet port
configure telnet vr
configure web http access-profile
create network-clock ptp
create ntp key
create snmp trap
delete network-clock ptp
delete ntp key
disable auto-provision
disable cdp ports
disable dhcp vlan
disable network-clock ptp
disable network-clock ptp boundary unicast-negotiation
disable ntp broadcast-server
disable snmp access
disable snmp access vr
disable snmp community
```



```
disable snmp traps
disable snmpv3
disable snmp-client
disable telnet
disable watchdog
enable auto-provision
disable ntp
disable ntp authentication
disable ntp broadcast-client
disable ntp broadcast-server
disable ntp vlan
enable cdp ports
enable dhcp vlan
enable network-clock ptp
enable network-clock ptp unicast-negotiation
enable network-clock ptp end-to-end transparent
enable snmp access
enable snmp access vr
enable snmp community
enable snmp traps
enable snmpv3
enable snmp-client
enable telnet
enable watchdog
enable ntp
enable ntp authentication
enable ntp broadcast-client
enable ntp broadcast-server
enable ntp vlan
exit
logout
quit
show access-list counters process
show auto-provision
show checkpoint-data
show dhcp-client state
show management
show network-clock ptp
show network-clock ptp (datasets)
show network-clock ptp (interface)
show network-clock ptp end-to-end-transparent ports
show network-clock ptp boundary unicast-master
```



```
show network-clock ptp boundary unicast-slave
show network-clock ptp counters
show node
show ntp
show ntp association
show ntp association statistics
show ntp key
show ntp restrict-list
show ntp server
show ntp sys-info
show ntp vlan
show odometers
show power
show power budget
show power controller
show power led motion-detector
show session
show snmp
show snmp vr_name
show snmpv3 access
show snmpv3 community
show snmpv3 context
show snmpv3 counters
show snmpv3 engine-info
show snmpv3 extreme-target-addr-ext
show snmpv3 filter
show snmpv3 filter-profile
show snmpv3 group
show snmpv3 mib-view
show snmpv3 notify
show snmpv3 target-addr
show snmpv3 target-params
show snmpv3 user
show snntp-client
telnet
telnet msm
telnet slot
tftp
tftp get
tftp put
```

This chapter describes commands for:



- Configuring Network Time Protocol (NTP) on the switch.
- Configuring Precision Time Protocol (PTP) on the switch.
- Configuring Simple Network Management Protocol (SNMP) parameters on the switch.
- Managing the switch using Telnet.
- Transferring files using the Trivial File Transfer Protocol (TFTP).
- Configuring system redundancy.
- Displaying power management statistics on the switch.
- Configuring Simple Network Time Protocol (SNTP) on the switch.
- Configuring power visualization.

NTP

NTP is a protocol for synchronizing clocks of servers or network entities using a TCP/IP-based network that has coherent variable latency.

It is designed particularly to resist the effects of variable latency by using a jitter buffer. NTP provides a coordinated Universal Time Clock (UTC). However, no information about time zones or daylight saving time is transmitted. NTP uses a hierarchical, semi-layered system of levels of clock sources. Each level of this hierarchy is termed a stratum and is assigned a layer number starting with 0 (zero) at the top. The stratum level defines its distance from the reference clock and exists to prevent cyclical dependencies in the hierarchy.

PTP

IEEE1588v2 (also known as Precision Time Protocol, PTP) is an industry-standard protocol that enables the precise transfer of frequency and time to synchronize clocks over packet-based Ethernet networks.

The locally available clock on each network device is synchronized with a grandmaster clock in the network. The devices are synchronized by exchange of timestamps with sub-nanoseconds granularity, to deliver very high accuracies of synchronization needed to ensure the stability of base station frequency and handovers. The timestamps between master and slave devices are exchanged through PTP event packets. The IPv4/UDP transport mechanism is used for PTP packets in ExtremeXOS 1588v2 implementation.

SNMP

Any network manager running SNMP can manage the switch, if the Management Information Base (MIB) is installed correctly on the management station.

Each network manager provides its own user interface to the management facilities.

The following SNMP parameters can be configured on the switch:

- Authorized trap receivers—An authorized trap receiver can be one or more network management stations on your network. The switch sends SNMP traps to all trap receivers. Entries in this list can



be created, modified, and deleted using the RMON2 trapDestTable MIB table, as described in RFC 2021, and the SNMPv3 tables.

- **SNMP INFORM**—SNMP INFORM allows for confirmation of a message delivery. When an SNMP manager receives an INFORM message from an SNMP agent, it sends a confirmation response back to the agent. If the message has not been received and therefore no response is returned, the INFORM message is resent. You can configure the number of retries to be made and the interval between retries.
- **SNMP access control**—**This feature allows the administrator to restrict SNMP access by using the access control list (ACL) and implementing an ACL policy. The administrator can configure an ACL policy to either permit or deny a specific list of IP address and subnet masks.**
- **Authorized managers**—An authorized manager can be either a single network management station, or a range of addresses (for example, a complete subnet) specified by a prefix and a mask.
- **Community strings**—The community strings allow a simple method of authentication between the switch and the remote network manager. The default read-only community string is public. The default read-write community string is private. The community strings for all authorized trap receivers must be configured on the switch for the trap receiver to receive switch-generated traps.
- **System contact (optional)**—The system contact is a text field that enables you to enter the name of the person(s) responsible for managing the switch.
- **System name (optional)**—The system name enables you to enter a name that you have assigned to this switch. The default name is the model name of the switch (for example, BD-1.2).
- **System location (optional)**—Using the system location field, you can find the location of the switch.



Note

If you specify volatile storage when configuring SNMP parameters, that configuration is not saved across a switch reboot.

Telnet

Telnet allows you to access the switch remotely using TCP/IP through one of the switch ports or a workstation with a Telnet facility.

If you access the switch via Telnet, you will use the command line interface (CLI) to manage the switch and modify switch configurations.

TFTP

ExtremeXOS supports the Trivial File Transfer Protocol (TFTP) based on RFC1350.

TFTP is a method used to transfer files from one network device to another. The ExtremeXOS TFTP client is a command line application used to contact an external TFTP server on the network. For example, ExtremeXOS uses TFTP to download software image files, switch configuration files, and access control lists (ACLs) from a server on the network to the switch.



System Redundancy with Dual Management Modules Installed, Modular Switches Only

If you install two MSMs/MMs in a modular switch, one assumes the role of primary and the other assumes the role of backup.

The primary MSM/MM provides all of the switch management functions including bringing up and programming the I/O modules, running the bridging and routing protocols, and configuring the switch. The primary also keeps synchronized with the backup MSM/MM in case the backup MSM/MM needs to take over the management functions if the primary MSM/MM fails.

Power Supply Management

On modular switches, ExtremeXOS monitors and manages power consumption on the switch by periodically checking the power supply units (PSUs) and testing them for failures.

To determine the health of the PSU, ExtremeXOS checks the voltage, current, and temperature of the PSU.

The power management capability of ExtremeXOS:

- Protects the system from overload conditions.
- Monitors all installed PSUs, even installed PSUs that are disabled.
- Enables and disables PSUs as required.
- Powers up or down I/O modules based on available power and required power resources.
- Logs power resource changes, including power budget, total available power, redundancy, and so on.
- Detects and isolates faulty PSUs.

On Summit family switches, ExtremeXOS reports when the PSU has power or has failed. Summit family switches support an internal power supply with a range of 90V to 240V AC power as well as an external redundant power supply. The Extreme Networks External Power System (EPS) allows you to add a redundant power supply to the Summit family switch to protect against a power supply failure. The EPS consists of a tray (EPS-T) that holds one or two EPS-160 power supplies. The EPS-160 provides 100V to 240V AC power. Each EPS-160 power supply provides one-to-one redundancy to an attached Summit family switch.

SNTP

ExtremeXOS supports the client portion of the Simple Network Time Protocol (SNTP) Version 3 based on RFC1769.

SNTP can be used by the switch to update and synchronize its internal clock from a Network Time Protocol (NTP) server. When enabled, the switch sends out a periodic query to the indicated NTP server, or the switch listens to broadcast NTP updates. In addition, the switch supports the configured setting for Greenwich Mean time (GMT) offset and the use of Daylight Saving Time.



clear cdp counters

```
configure cdp counters {ports ports_list }
```

Description

Clears the CDP counter statistics.

Syntax Description

ports	Specifies the ports to clear.
<i>ports_list</i>	Specifies the port list.

Default

N/A.

Usage Guidelines

Use this command to clear the CDP counter statistics.

Example

The following command clears the CDP ports counters:

```
clear cdp counters
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

clear cdp neighbor

```
clear cdp neighbor [device id device_id | all]
```

Description

Clears the CDP neighbor information.



Syntax Description

device id	Specifies the Device Identifier to be used in CDP.
<i>device_id</i>	Specifies the Device Identifier of neighbor.
all	Specifies all CDP neighbors.

Default

N/A.

Usage Guidelines

Use this command to clear the CDP neighbor information.

Example

The following command clears all CDP neighbor associations:

```
clear cdp neighbor all
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

clear network-clock ptp counters

```
clear network-clock ptp [boundary | ordinary] vlan [vlan_name {ipv4_address
[unicast-master | unicast-slave]} | all] counters
```

Description

This command clears the accumulated PTP packet counters. The clear can be performed on the following groups:

Per unicast-master or unicast-slave peer on a clock port.

Peers on a clock port.

Peers on all clock ports.



Note

This command is available only for Boundary and Ordinary clocks.



Syntax Description

network-clock	External Clock for Ethernet synchronization.
ptp	Precise Time Protocol.
boundary	Boundary clock.
ordinary	Ordinary clock.
vlan	VLAN.
<i>ipv4_address</i>	Peer IP address.
unicast-master	IP addresses that are masters to the local clock.
unicast-slave	IP addresses that are slaves to the local clock.
all	All VLAN.
counters	PTP message counts.

Default

N/A.

Usage Guidelines

Use this command to clear the accumulated PTP packet counters.

Example

N/A.

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is only available boundary and ordinary clocks on cell site routers (E4G-200 and E4G-400) and Summit X770.



Note

PTP commands can be used only with the Network Timing feature pack.

configure cdp device-id

```
configure cdp device-id [device_id | system-mac]
```

Description

Configures the device ID only in CDP.



Syntax Description

<i>device-id</i>	Unique device Identifier to be used in CDP. The default is MAC Address.
system-mac	Device Identifier will be system MAC address.

Default

MAC address.

Usage Guidelines

Use this command to configure the Device ID. If you do not configure it, the MAC address is used as the Device ID. This configuration of device ID is only used in the CDP .

Example

The following command configures the device ID as the MAC address:

```
configure cdp device-id system-mac
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

configure cdp frequency

```
configure cdp frequency seconds
```

Description

Enables CDP on a port.

Syntax Description

<i>seconds</i>	Specifies the transmit frequency in seconds. The range is 5,254 seconds. The default value is 60 seconds.
----------------	---

Default

60 seconds.



Usage Guidelines

Example

The following command configures the CDP frequency as two minutes:

```
configure cdp frequency 120
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

configure cdp hold-time

```
configure cdp hold-time seconds
```

Description

Configures the hold time of the neighbor information .

Syntax Description

<i>seconds</i>	Duration in seconds that receiver must keep this packet. The range is 10-255 and the default is 180 seconds.
----------------	--

Default

60 seconds.

Usage Guidelines

Use this command to configure the hold time of the neighbor information for which a receiving device should hold information before discarding it.

Example

The following command configures the CDP hold time as two minutes:

```
configure cdp hold-time 120
```



History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

configure node priority

```
configure node slot slot_id priority node_pri
```

Description

Configures the priority of the node.

Syntax Description

<i>slot_id</i>	Specifies the slot of the node. A is for the MSM/MM installed in slot A. B is for the MSM/MM installed in slot B.
<i>node_pri</i>	Specifies the priority of the node. The default 0 gives MSM-A a higher priority over MSM-B. The range is 1 to 100; 0 means you have not configured a node priority.

Default

Default node priority is 0.

Usage Guidelines

Use this command to configure the priority of the node. The lower the number, the higher the priority.

The node priority is part of the selection criteria for the primary node. The following list describes the parameters used to determine the primary node:

- Node state—The node state must be STANDBY to participate in leader election and to be selected primary. If the node is in the INIT, DOWN, or FAIL states, the node will not participate in leader election.
- Configuration priority—This is a user assigned priority. The configured priority is compared only after the node meets the minimum thresholds in each category for it to be healthy. Required processes and devices must not fail.
- Software health—This represents the percent of processes available.
- Health of secondary hardware components—This represents the health of switch components, such as the power supplies, fans, and so forth.
- Slot ID—The MSM/MM slot where the node is installed (MSM-A or MSM-B).

If you do not configure any priorities, MSM-A has a higher priority than MSM-B.



Example

The following command configures a priority of 2 for MSM-B:

```
configure node slot B priority 2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available only on modular switches.

configure ntp local-clock none

```
configure ntp local-clock none
```

Description

Removes the internal local clock from the clock source list.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

N/A.

Example

The following command removes the internal local clock from the clock source list:

```
configure ntp local-clock none
```

History

This command was first available in ExtremeXOS 12.7.



Platform Availability

This command is available on all platforms.

configure ntp local-clock stratum

```
configure ntp local-clock stratum stratum_number
```

Description

Configures the internal local clock with a stratum number. The stratum number defines the distance from the reference clock. The lower the number, the closer the switch is to the reference clock.

Syntax Description

<i>stratum_number</i>	Specifies the distance from the reference clock from 2 through 16, with 2 being closest and 16 being the farthest away.
-----------------------	---

Default

The local clock is disabled by default.

Usage Guidelines

The internal local clock is configured as a clock source with a given stratum number. Because the local clock is not as reliable as an external clock source with GPS or CDMA, the stratum number should be higher than the stratum number of the external clock source to allow the system to acquire the most reliable clock information from the clock source lists.

Example

The following command configures the local clock with a stratum number of 3:

```
configure ntp local-clock stratum 3
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on all platforms.

configure ntp restrict-list

```
configure ntp restrict-list [add | delete] network {mask} [permit | deny]
```



Description

Restricts a host or block of client IP addresses from getting NTP service. When NTP is enabled over a VLAN, an NTP server is configured, or a broadcast NTP server is in a VLAN, the VLAN's IP block or NTP server's IP address is automatically added into the system with a permit action.

Syntax Description

add	Restricts a client from getting NTP service.
delete	Removes a client from the restrict list.
<i>network</i>	Specifies a host or block of IP addresses.
<i>mask</i>	Specifies the subnet mask of the network.
permit	Specifies that a particular block of client IP addresses is permitted to get NTP service.
deny	Specifies that a particular block of client IP addresses is denied NTP service.

Default

All addresses are denied by default.

Usage Guidelines

N/A.

Example

The following command restricts a block of client IP addresses from getting NTP service:

```
configure ntp restrict-list add 132.25.82.3 deny
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on all platforms.

configure ntp server/peer add

```
configure ntp [server | peer] add [ip_address | host_name] {key keyid} {option [burst | initial-burst]}
```



Description

Configures an NTP server or peer.

Syntax Description

<i>ip_address</i>	Specifies the IP address of the NTP server or peer.
<i>host_name</i>	Specifies the host name of the NTP server or peer.
<i>keyid</i>	Specifies the key ID as a value from 1 to 65534.
burst	Follows the same burst mechanism when an NTP server is reachable.
initial-burst	Allows the system to send six burst packets when an NTP server becomes unreachable (discovered but unreachable).

Default

N/A.

Usage Guidelines

The initial-burst option is useful when a fast time synchronization is required at the initial stage.

Example

The following command adds an NTP server named “Missouri” with key 5 and an initial burst:

```
configure ntp server add 134.20.16.35 Missouri key 5 option initial-burst
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on all platforms.

configure ntp server/peer delete

```
configure ntp [server | peer] delete [ip_address | host_name]
```

Description

Removes an NTP server or peer from external clock source lists.



Syntax Description

<i>ip-address</i>	Specifies the IP address of the NTP server or peer.
<i>host_name</i>	Specifies the host name of the NTP server or peer.

Default

N/A.

Usage Guidelines

N/A.

Example

The following command removes an NTP peer from external clock source lists:

```
configure ntp peer delete
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on all platforms.

configure power led motion-detector

```
configure power led motion-detector [disable | enable {timeout seconds}]
```

Description

Configures the motion detector to control LEDs depending on motion near the front of the switch.

Syntax Description

disable	Disables the motion detector.
enable	Enables motion detector.
<i>seconds</i>	Specifies the number of seconds before the LEDs are turned off after motion is detected. The range is 1 to 600.

Default

Disable. The default number of seconds is 180.



Usage Guidelines

Use this command to enable or disable the motion detector to control the port LEDs.

- When the motion detector is enabled, the LEDs are turned on only when motion is detected. The length of time in seconds that they remain turned on is configurable.
- When the motion detector is disabled, the LEDs are always turned on.

To view the status and timeout settings, use the `show power led motion-detector` command.

Example

The following command configures the LEDs to turn on when motion is detected and turn off 60 seconds after motion is detected:

```
configure power led motion-detector enable timeout 60
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on the Summit X670 switch.

configure power monitor

```
configure power monitor poll-interval [off | seconds] change-action [none | [log  
| log-and-trap | trap] change-threshold watts]
```

Description

Configures the power visualization, which periodically polls for input power usage.

Syntax Description

seconds	Input power usage poll interval in seconds. If zero is configured, then the input power measurement is disabled.
change-action	The action to be taken whenever the power is increased or decreased by the configured threshold power value (none, log, log-and-trap, or trap).
watts	The power value in watts for the change threshold. The default value is 2 watts.

Default

The default poll interval is 60 seconds.



The default change action is none.

The default change threshold is 2 watts.

Usage Guidelines

Use this command to configure change actions to be taken when input power usage is increased or decreased by the configured threshold power value. The polling interval is also configurable, with a default value of 60 seconds.



Note

Input power usage values are only estimates.

Example

The following command configures a polling interval of 10 seconds, a change action of log-and-trap, and a change threshold of 3 watts:

```
configure power monitor poll-interval 10 change-action log-and-trap change-threshold 3
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the Summit X460 switch, the E4G-400 switch, the BlackDiamond X8 series switches, and the BlackDiamond 8800 series switches.

configure power supply

```
configure power supply ps_num {auto | on}
```

Description

Configures a power supply for either automatic power management, or forced on, regardless of the impact to the total available system power.



Syntax Description

<i>ps_num</i>	Specifies the slot number of the installed power supply unit (PSU) to which this command applies.
auto	Specifies that ExtremeXOS determines the enabled or disabled state of the PSU to maximize total system power. This is the default.
on	Specifies that the PSU be enabled even if ExtremeXOS determines it should be disabled. This action may reduce the total available system power and may result in one or more I/O modules powering down.

Default

The default setting is auto; ExtremeXOS either enables or disables the PSU in order to maximize total system power.

Usage Guidelines

If a switch has PSUs with a mix of both 220V AC and 110V AC inputs, ExtremeXOS maximizes system power by automatically taking one of two possible actions:

- If all PSUs are enabled then all PSUs must be budgeted at 110V AC to prevent overload of PSUs with 110V AC inputs.

OR

- If the PSUs with 110V AC inputs are disabled, then the PSUs with 220V AC inputs can be budgeted with a higher output per PSU.

ExtremeXOS computes the total available power using both methods and automatically uses the PSU configuration that provides the greatest amount of power to the switch. The following table lists combinations where ExtremeXOS maximizes system power by disabling the PSUs with 110V AC inputs.

Table 9: PSU Combinations Where 110V PSUs Are Disabled

Number of PSUs with 220VAC Inputs	Number of PSUs with 110VAC Inputs
2	1
3	1
3	2
4	1
4	2
5	1

Table 10: BlackDiamond X8 Series PSU Combinations Where 110V PSUs Are Disabled

Number of PSUs with 220VAC Inputs	Number of PSUs with 110VAC Inputs
1	1
2	1
3	1



Table 10: BlackDiamond X8 Series PSU Combinations Where 110V PSUs Are Disabled (continued)

Number of PSUs with 220VAC Inputs	Number of PSUs with 110VAC Inputs
3	2
4	1
4	2
4	3
5	1
5	2
5	3
6	1
6	2

For all other combinations of 220V AC and 110V AC PSUs, ExtremeXOS maximizes system power by enabling all PSUs and budgeting each PSU at 110V AC.

In addition to the PSU, you can specify the following options:

- `auto`—Specifies that ExtremeXOS determines the enabled or disabled state of the PSU to maximize total system power. This is the default.
- `on`—Specifies that the PSU be enabled even if ExtremeXOS determines it should be disabled. This action may reduce the total available system power and may result in one or more I/O modules powering down.

You can override automatic power supply management to enable a PSU with 110V AC inputs that ExtremeXOS disables if the need arises, such as for a planned maintenance of 220V AC circuits. If the combination of AC inputs represents one of those listed in the table above, you can turn on a disabled PSU using the `configure power supply ps_num on` command.

**Note**

If you override automatic power supply management, you may reduce the available power and cause one or more I/O modules to power down.

To resume using automatic power supply management on a PSU, use the `configure power supply ps_num auto` command. The setting for each PSU is stored as part of the switch configuration.

To display power supply status and power budget information use the `show power` and `show power budget` commands.

Example

The following command configures the PSU in slot 1 to be forced on when either 110V AC or 220V AC power input is present, overriding automatic power management:

```
configure power supply 1 on
```



The switch displays the following message:

```
In a mixed environment of 110V and 220V AC inputs, power management may
automatically disable 110V supplies to maximize the system power budget.
By specifying 'on', you wish to override power management and enable the
specified power supply. This may cause the system power budget to decrease
and one or more I/O cards may be powered off as a result.
Are you sure you want to continue? (y/n)
```

Enter y to continue.

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available only on modular switches.

configure snmp access-profile

```
configure snmp access-profile [ access_profile {readonly | readwrite} | [[add
rule ] [first | [[before | after] previous_rule]]] | delete rule | none ]
```

Description

Configures SNMP to use an ACL policy or ACL rule for access control.

Syntax Description

<i>access_profile</i>	Specifies an ACL policy.
readonly	Specifies that access granted by the specified policy is read only.
readwrite	Specifies that access granted by the specified policy is read/write.
add	Specifies that an ACL rule is to be added to the SNMP application.
<i>rule</i>	Specifies an ACL rule.
first	Specifies that the new rule is to be added before all other rules.
before	Specifies that the new rule is to be added before a previous rule.
after	Specifies that the new rule is to be added after a previous rule.
<i>previous_rule</i>	Specifies an existing rule in the application.
delete	Specifies that the named rule is to be deleted.
none	Specifies that all the rules or a policy file is to be deleted.



Default

SNMP access is enabled by default, with no ACL policies.

Usage Guidelines

You must be logged in as administrator to configure SNMP parameters. You can restrict SNMP access in the following ways:

- Implement an ACL policy. You create an ACL policy file that permits or denies a specific list of IP addresses and subnet masks for SNMP. You must create the ACL policy file before you can use this command. If the ACL policy file does not exist on the switch, the switch returns an error message indicating that the file does not exist.

In the ACL policy file for SNMP, the source-address field is the only supported match condition. Any other match conditions are ignored.

Use the none option to remove a previously configured ACL policy.

- Add an ACL rule to the SNMP application through this command. Once an ACL is associated with SNMP, all the packets that reach an SNMP module are evaluated with this ACL and appropriate action (permit or deny) is taken, as is done using policy files.

The permit or deny counters are also updated accordingly, regardless of whether the ACL is configured to add counters. To display counter statistics, use the `show access-list counters process snmp` command.

Only the following match conditions and actions are copied to the client memory. Others that may be in the rule are not copied.

Match conditions:

- Source-address—IPv4 and IPv6

- Actions:

Permit

Deny

When adding a new rule, use the first, before, and after previous_rule parameters to position it within the existing rules.

If the SNMP traffic does not match any of the rules, the default behavior is permit. To deny SNMP traffic that does not match any of the rules, add a deny all rule at the end of the rule list.

Creating an ACL Policy File

To create an ACL policy file, use the `edit policy` command. For more information about creating and implementing ACL policy files, see [Policy Manager](#) and [ACLs](#) in the *ExtremeXOS Concepts Guide*.

If you attempt to implement a policy that does not exist, an error message similar to the following appears:

```
Error: Policy /config/MyAccessProfile.pol does not exist on file system
```



If this occurs, make sure the policy you want to implement exists. To confirm the existence of the policies, use the `configure snmp add community` command. If the policy does not exist, create the ACL policy file.

Viewing SNMP Information

To display the current management configuration, including SNMP access related information, whether SNMP access is enabled or disabled, and whether any ACL or rules are configured for SNMP, use the following command: `show management`

Example

The following command applies the ACL policy file MyAccessProfile_2 to SNMP:

```
configure snmp access-profile MyAccessProfile_2
```

The following command applies the ACL rule DenyAccess to SNMP as the first rule in the list:

```
configure snmp access-profile add DenyAccess first
```

The following command deletes the ACL rule DenyAccess from the SNMP application:

```
configure snmp access-profile delete DenyAccess
```

To delete the use of all the ACL rules or a policy file by SNMP, use the following command:

```
configure snmp access-profile none
```

History

This command was first available in ExtremeXOS 11.6.

Support for individual ACL rules was added in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

configure snmp add community

```
configure snmp add community [readonly | readwrite] alphanumeric_string
```



Description

Adds an SNMP read or read/write community string.

Syntax Description

readonly	Specifies read-only access to the system.
readwrite	Specifies read and write access to the system.
<i>alphanumeric_string</i>	Specifies an SNMP community string name. See “Usage Guidelines” for more information.

Default

The default read-only community string is public. The default read/write community string is private.

Usage Guidelines

Community strings provide a simple method of authentication between a switch and a remote network manager. Read community strings provide read-only access to the switch. The default read-only community string is public. Read-write community strings provide read and write access to the switch. The default read/write community string is private. Sixteen read-only and sixteen read/write community strings can be configured on the switch, including the defaults.

An authorized trap receiver must be configured to use the correct community strings on the switch for the trap receiver to receive switch-generated traps. In some cases, it may be useful to allow multiple community strings so that all switches and trap receivers are not forced to use identical community strings. The `configure snmp add community` command allows you to add multiple community strings in addition to the default community string.

An SNMP community string can contain up to 32 characters.

We recommend that you change the defaults of the community strings. To change the value of the default read/write and read-only community strings, use the `configure snmp delete community` command.

Example

The following command adds a read/write community string with the value extreme:

```
configure snmp add community readwrite extreme
```

History

This command was first available in ExtremeXOS 10.1.



Platform Availability

This command is available on all platforms.

configure snmp add trapreceiver

```
configure snmp add trapreceiver [ip_address | ipv6_address] community [[hex
hex_community_name] | community_name] {port port_number} {from [src_ip_address |
src_ipv6_address]} {vr vr_name} {mode trap_mode}
```

Description

Adds the IP address of a trap receiver to the trap receiver list and specifies which SNMPv1/v2c traps are to be sent.

Syntax Description

<i>ip_address</i>	Specifies an SNMP trap receiver IPv4 address.
<i>ipv6_address</i>	Specifies an SNMP trap receiver IPv6 address.
<i>hex_community_name</i>	Specifies that the trap receiver is to be supplied as a colon separated string of hex octets.
<i>community_name</i>	Specifies the community string of the trap receiver to be supplied in ASCII format.
<i>port_number</i>	Specifies a UDP port to which the trap should be sent. Default is 162.
<i>src_ip_address</i>	Specifies the IPv4 address of a VLAN to be used as the source address for the trap.
<i>src_ipv6_address</i>	Specifies the IPv6 address of a VLAN to be used as the source address for the trap.
<i>vr_name</i>	Specifies the name of the virtual router.
<i>trap_mode</i>	Specifies the mode of the traps:enhanced—Contains extra varbinds at the end.standard—Does not contain extra varbinds.

Default

Trap receivers are in enhanced mode by default, and the version is SNMPv2c by default.

Usage Guidelines

The IP address can be unicast, multicast, or broadcast.

An authorized trap receiver can be one or more network management stations on your network. Authorized trap receivers must be configured on the switch for the trap receiver to receive switch-generated traps. The switch sends SNMP traps to all trap receivers configured to receive the specific trap group.



To view the SNMP trap receivers configured on the switch, use the `show management` command. The `show management` command displays information about the switch including the destination and community of the SNMP trap receivers configured on the switch.

Example

The following command adds the IP address 10.101.0.100 as a trap receiver with community string purple:

```
configure snmp add trapreceiver 10.101.0.100 community purple
```

The following command adds the IP address 10.101.0.105 as a trap receiver with community string green, using port 3003:

```
configure snmp add trapreceiver 10.101.0.105 community green port 3003
```

The following command adds the IP address 10.101.0.105 as a trap receiver with community string blue, and IP address 10.101.0.25 as the source:

```
configure snmp add trapreceiver 10.101.0.105 community blue from 10.101.0.25
```

History

This command was first available in ExtremeXOS 10.1.

The virtual router parameter was added in ExtremeXOS 12.3.

IPv6 support was added in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

configure snmp delete community

```
configure snmp delete community [readonly | readwrite] [all | alphanumeric_string]
```

Description

Deletes an SNMP read or read/write community string.



Syntax Description

readonly	Specifies read-only access to the system.
readwrite	Specifies read and write access to the system.
all	Specifies all of the SNMP community strings.
<i>alphanumeric_string</i>	Specifies an SNMP community string name. See “Usage Guidelines” for more information.

Default

The default read-only community string is public. The default read/write community string is private.

Usage Guidelines

You must have at least one community string for SNMP access. If you delete all of the community strings on your system, you will no longer have SNMP access, even if you have SNMP enabled.

The community strings allow a simple method of authentication between the switch and the remote network manager. There are two types of community strings on the switch. Read community strings provide read-only access to the switch. The default read-only community string is public. read/write community strings provide read and write access to the switch. The default read/write community string is private. Sixteen read-only and sixteen read-write community strings can be configured on the switch, including the defaults. The community string for all authorized trap receivers must be configured on the switch for the trap receiver to receive switch-generated traps. SNMP community strings can contain up to 32 characters.

For increased security, we recommend that you change the defaults of the read/write and read-only community strings.

Use the `configure snmp add` commands to configure an authorized SNMP management station.

Example

The following command deletes a read/write community string named extreme:

```
configure snmp delete community readwrite extreme
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.



configure snmp delete trapreceiver

```
configure snmp delete trapreceiver [[ip_address | ipv6_address] {port_number} | all]
```

Description

Deletes a specified trap receiver or all authorized trap receivers.

Syntax Description

<i>ip_address</i>	Specifies an SNMP trap receiver IPv4 address.
<i>ipv6_address</i>	Specifies an SNMP trap receiver IPv6 address.
<i>port_number</i>	Specifies the port associated with the receiver.
all	Specifies all SNMP trap receiver IP addresses.

Default

The default port number is 162.

Usage Guidelines

Use this command to delete a trap receiver of the specified IPv4 or IPv6 address, or all authorized trap receivers.

This command deletes only the first SNMPv1/v2c trap receiver whose IP address and port number match the specified value.

Example

The following command deletes the trap receiver 10.101.0.100 from the trap receiver list:

```
configure snmp delete trapreceiver 10.101.0.100
```

The following command deletes entries in the trap receiver list for 10.101.0.100, port 9990:

```
configure snmp delete trapreceiver 10.101.0.100 9990
```

Any entries for this IP address with a different community string will not be affected.

History

This command was first available in ExtremeXOS 10.1.

IPv6 support was added in ExtremeXOS 12.4.



Platform Availability

This command is available on all platforms.

configure snmp sysContact

```
configure snmp syscontact sysContact
```

Description

Configures the name of the system contact.

Syntax Description

<i>sysContact</i>	An alphanumeric string that specifies a system contact name.
-------------------	--

Default

N/A.

Usage Guidelines

The system contact is a text field that enables you to enter the name of the person(s) responsible for managing the switch. A maximum of 255 characters is allowed.

To view the name of the system contact listed on the switch, use the `show switch` command. The `show switch` command displays switch statistics including the name of the system contact.

Example

The following command defines FredJ as the system contact:

```
configure snmp syscontact FredJ
```

The following output from the `show switch` command displays FredJ as the system contact:

```
SysName:          engineeringlab
SysLocation:      englab
SysContact:       FredJ
```

History

This command was first available in ExtremeXOS 10.1.



Platform Availability

This command is available on all platforms.

configure snmp sysLocation

```
configure snmp syslocation sysLocation
```

Description

Configures the location of the switch.

Syntax Description

<i>sysLocation</i>	An alphanumeric string that specifies the switch location.
--------------------	--

Default

N/A.

Usage Guidelines

Use this command to indicate the location of the switch. A maximum of 255 characters is allowed.

To view the location of the switch on the switch, use the `show switch` command. The `show switch` command displays switch statistics including the location of the switch.

Example

The following command configures a switch location name on the system:

```
configure snmp syslocation englab
```

The following output from the `show switch` command displays englab as the location of the switch:

```
SysName:          engineeringlab
SysLocation:     englab
SysContact:      FredJ
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.



configure snmp sysName

configure snmp sysname *sysName*

Description

Configures the name of the switch.

Syntax Description

<i>sysName</i>	An alphanumeric string that specifies a device name.
----------------	--

Default

The default *sysName* is the model name of the device (for example, BlackDiamond8800).

Usage Guidelines

You can use this command to change the name of the switch. A maximum of 32 characters is allowed. The *sysName* appears in the switch prompt. On a SummitStack, the *sysName* appears in the prompt of all active nodes in the stack when there is a master node present in the stack.

To view the name of the system listed on the switch, use the `show switch` command. The `show switch` command displays switch statistics including the name of the system.

Example

The following command names the switch:

```
configure snmp sysname engineeringlab
```

The following output from the `show switch` command displays `engineeringlab` as the name of the switch:

```
SysName:          engineeringlab
SysLocation:     englab
SysContact:      FredJ
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.



configure snmpv3 add access

```
configure snmpv3 add access [[hex hex_group_name] | group_name] {sec-model
[snmpv1 | snmpv2c | usm]} {sec-level [noauth | authnopriv | priv]} {read-view
[[hex hex_read_view_name] | read_view_name]} {write-view [[hex
hex_write_view_name]] | write_view_name]} {notify-view [[hex
hex_notify_view_name]] | notify_view_name]} {volatile}
```

Description

Creates (and modifies) a group and its access rights.

Syntax Description

<i>hex_group_name</i>	Specifies the group name to add or modify. The value is to be supplied as a colon separated string of hex octets.
<i>group_name</i>	Specifies the group name to add or modify. The value is to be supplied in ASCII format.
sec-model	Specifies the security model to use.
snmpv1	Specifies the SNMPv1 security model.
snmpv2c	Specifies the SNMPv2c security model.
usm	Specifies the SNMPv3 User-based Security Model (USM).
sec-level	Specifies the security level for the group.
noauth	Specifies no authentication (and implies no privacy) for the security level.
authnopriv	Specifies authentication and no privacy for the security level.
priv	Specifies authentication and privacy for the security level.
read-view	Specifies the read view name: hex_read_view_name—Specifies a hex value supplied as a colon separated string of hex octets read_view_name—Specifies an ASCII value.
write-view	Specifies the write view name: hex_write_view_name—Specifies a hex value supplied as a colon separated string of hex octets write_view_name—Specifies an ASCII value.
notify-view	Specifies the notify view name: hex_notify_view_name—Specifies a hex value supplied as a colon separated string of hex octets notify_view_name—Specifies an ASCII value.
volatile	Specifies volatile storage.

Default

The default values are:

- sec-model—USM
- sec-level—noauth
- read view name—defaultUserView
- write view name— ""



- notify view name—defaultNotifyView
- non-volatile storage

Usage Guidelines

Use this command to configure access rights for a group. All access groups are created with a unique default context, "", as that is the only supported context.

Use more than one character when creating unique community strings and access group names.

A number of default groups are already defined. These groups are: admin, initial, v1v2c_ro, v1v2c_rw.

- The default groups defined are v1v2c_ro for security name v1v2c_ro, v1v2c_rw for security name v1v2c_rw, admin for security name admin, and initial for security names initial, initialmd5, initialsha, initialmd5Priv and initialshaPriv.
- The default access defined are admin, initial, v1v2c_ro, v1v2c_rw, and v1v2cNotifyGroup.

Example

In the following command, access for the group defaultROGroup is created with all the default values: security model usm, security level noauth, read view defaultUserView, no write view, notify view defaultNotifyView, and storage nonvolatile.

```
configure snmpv3 add access defaultROGroup
```

In the following command, access for the group defaultROGroup is created with the values: security model USM, security level authnopriv, read view defaultAdminView, write view defaultAdminView, notify view defaultAdminView, and storage nonvolatile.

```
configure snmpv3 add access defaultROGroup sec-model usm sec-level authnopriv
read-view defaultAdminView write-view defaultAdminView notify-view
defaultAdminView
```

History

This command was first available in ExtremeXOS 10.1.

The hex_read_view_name, hex_write_view_name, and hex_notify_view_name parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 add community



```
configure snmpv3 add community [[hex hex_community_index] | community_index]
[encrypted name community_name | name [[hex hex_community_name] | community_name]
{store-encrypted} ] user [[hex hex_user_name] | user_name] {tag [[hex
transport_tag] | transport_tag]} {volatile}
```

Description

Adds an SNMPv3 community entry.

Syntax Description

<i>hex_community_index</i>	Specifies the row index in the snmpCommunity table as a hex value supplied as a colon separated string of hex octets.
<i>community_index</i>	Specifies the row index in the snmpCommunity Table as an ASCII value.
<i>hex_community_name</i>	Specifies the community name as a hex value supplied as a colon separated string of hex octets.
<i>community_name</i>	Specifies the community name as an ASCII value.
<i>hex_user_name</i>	Specifies the USM user name as a hex value supplied as a colon separated string of hex octets.
<i>user_name</i>	Specifies the USM user name as an ASCII value.
tag	Specifies the tag used to locate transport endpoints in SnmpTargetAddrTable. When this community entry is used to authenticate v1/v2c messages, this tag is used to verify the authenticity of the remote entity. hex_transport_tag—Specifies a hex value supplied as a colon separated string of hex octet transport_tag—Specifies an ASCII value
volatile	Specifies volatile storage.

Default

N/A.

Usage Guidelines

Use this command to create or modify an SMMPv3 community in the community MIB.

Example

```
switch # configure snmp add community readonly extreme store-encrypted
switch # show snmpv3 community
Community Index      : extreme
Community Name      : hys{fnj (encrypted)
Security Name       : v1v2c_ro
Context EngineID    : 80:00:07:7c:03:00:04:96:27:b6:63
Context Name        :
Transport Tag       :
Storage Type        : NonVolatile
Row Status          : Active
switch # configure snmp add community readwrite extreme123
```



```

switch # show snmpv3 community
Community Index : extreme
Community Name  : hys{fnj (encrypted)
Security Name   : vlv2c_ro
Context EngineID : 80:00:07:7c:03:00:04:96:27:b6:63
Context Name    :
Transport Tag   :
Storage Type    : NonVolatile
Row Status      : Active
Community Index : extremel23
Community Name  : extremel23
Security Name   : vlv2c_rw
Context EngineID : 80:00:07:7c:03:00:04:96:27:b6:63
Context Name    :
Transport Tag   :
Storage Type    : NonVolatile
Row Status      : Active
switch # show configuration "snmp"
#
# Module snmpMaster configuration.
#
configure snmpv3 add community extreme encrypted name hys{fnj user vlv2c_ro
configure snmpv3 add community extremel23 name extremel23 user vlv2c_rw
The following command creates an entry with the community index comm_index,
community name comm_public, and user (security) name vlv2c_user:
configure snmpv3 add community comm_index name comm_public user vlv2c_user

```

History

This command was first available in ExtremeXOS. 10.1.

The hex_community_index, hex_community_name, hex_user_name, and hex_transport_tag parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 add filter

```

configure snmpv3 add filter [[hex hex_profile_name] | profile_name] subtree
object_identifier {/subtree_mask} type [included | excluded] {volatile}

```

Description

Adds a filter to a filter profile.



Syntax Description

<i>hex_profile_name</i>	Specifies the filter profile that the current filter is added to. The value is to be supplied as a colon separated string of hex octets.
<i>profile_name</i>	Specifies the filter profile that the current filter is added to in ASCII format.
<i>object_identifier</i>	Specifies a MIB subtree.
subtree_mask	Specifies a hex octet string used to mask the subtree. For example, f7a indicates 1.1.1.0.1.1.1.0.1.0.
included	Specifies that the MIB subtree defined by object identifier/mask is to be included.
excluded	Specifies that the MIB subtree defined by object identifier/mask is to be excluded.
volatile	Specifies volatile storage.

Default

The default values are:

- mask value—empty string (all 1s).
- type—included.
- storage—non-volatile.

Usage Guidelines

Use this command to create a filter entry in the snmpNotifyFilterTable. Each filter includes or excludes a portion of the MIB. Multiple filter entries comprise a filter profile that can eventually be associated with a target address. Other commands are used to associate a filter profile with a parameter name, and the parameter name with a target address.

This command can be used multiple times to configure the exact filter profile desired.

Example

The following command adds a filter to the filter profile prof1 that includes the MIB subtree 1.3.6.1.4.1/f0:

```
configure snmpv3 add filter prof1 subtree 1.3.6.1.4.1/f0 type included
```

History

This command was first available in ExtremeXOS 10.1.

The hex_profile_name parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.



configure snmpv3 add filter-profile

```
configure snmpv3 add filter-profile [[hex hex_profile_name] | profile_name] param
[[hex hex_param_name]] | param_name] {volatile}
```

Description

Associates a filter profile with a parameter name.

Syntax Description

<i>hex_profile_name</i>	Specifies the filter profile name. The value is to be supplied as a colon separated string of hex octets.
<i>profile_name</i>	Specifies the filter profile name in ASCII format.
<i>hex_param_name</i>	Specifies a parameter name to associate with the filter profile. The value to follow is to be supplies as a colon separated string of hex octets.
<i>param_name</i>	Specifies a parameter name to associate with the filter profile in ASCII format.
volatile	Specifies volatile storage.

Default

The default storage type is non-volatile.

Usage Guidelines

Use this command to add an entry to the snmpNotifyFilterProfileTable. This table associates a filter profile with a parameter name. The parameter name is associated with target addresses, and the filter profile is associated with a series of filters, so, in effect, you are associating a series of filters with a target address.

Example

The following command associates the filter profile prof1 with the parameter name P1:

```
configure snmpv3 add filter-profile prof1 param P1
```

History

This command was first available in ExtremeXOS 10.1.

The *hex_profile_name* and *hex_param_name* parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.



configure snmpv3 add group user

```
configure snmpv3 add group [[hex hex_group_name] | group_name] user [[hex
hex_user_name] | user_name] {sec-model [snmpv1 | snmpv2c | usm]} {volatile}
```

Description

Adds a user name (security name) to a group.

Syntax Description

<i>hex_group_name</i>	Specifies the group name to add or modify. The value is to be supplied as a colon separated string of hex octets.
<i>group_name</i>	Specifies the group name to add or modify in ASCII format.
<i>hex_user_name</i>	Specifies the user name to add or modify. The value to follow is to be supplies as a colon separated string of hex octets.
<i>user_name</i>	Specifies the user name to add or modify in ASCII format.
sec-model	Specifies the security model to use.
snmpv1	Specifies the SNMPv1 security model.
snmpv2c	Specifies the SNMPv2c security model.
usm	Specifies the SNMPv3 User-based Security Model (USM).
volatile	Specifies volatile storage.

Default

The default values are:

- sec-model—USM.
- non-volatile storage.

Usage Guidelines

Use this command to associate a user name with a group.

As per the SNMPv3 RFC, a security name is model independent while a username is model dependent. For simplicity, both are assumed to be same here. User names and security names are handled the same. In other words, if a user is created with the user name username, the security name value is the same, username.

Every group is uniquely identified by a security name and security model. So the same security name can be associated to a group name but with different security models.



Example

The following command associates the user userV1 to the group defaultRoGroup with SNMPv1 security:

```
configure snmpv3 add group defaultRoGroup user userV1 sec-model snmpv1
```

The following command associates the user userV3 with security model USM and storage type volatile to the access group defaultRoGroup:

```
configure snmpv3 add group defaultRoGroup user userV3 volatile
```

History

This command was first available in ExtremeXOS 10.1.

The hex_group_name and hex_user_name parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 add mib-view

```
configure snmpv3 add mib-view [[hex hex_view_name] | view_name] subtree  
object_identifier {subtree_mask} {type [included | excluded]} {volatile}
```

Description

Adds (and modifies) a MIB view.

Syntax Description

<i>hex_view_name</i>	Specifies the MIB view name to add or modify. The value is to be supplies as a colon separated string of hex octets.
<i>view_name</i>	Specifies the MIB view name to add or modify in ASCII format.
<i>object_identifier</i>	Specifies a MIB subtree.
<i>subtree_mask</i>	Specifies a hex octet string used to mask the subtree. For example, f7a indicates 1.1.1.0.1.1.1.0.1.0.
included	Specifies that the MIB subtree defined by subtree/mask is to be included.
excluded	Specifies that the MIB subtree defined by subtree/mask is to be excluded.
volatile	Specifies volatile storage.



Default

The default mask value is an empty string (all 1s). The other default values are included and non-volatile.

Usage Guidelines

Use this command to create a MIB view into a subtree of the MIB. If the view already exists, this command modifies the view to additionally include or exclude the specified subtree.

In addition to the created MIB views, there are three default views. They are: defaultUserView, defaultAdminView, and defaultNotifyView.

Example

The following command creates the MIB view allMIB with the subtree 1.3 included as non-volatile:

```
configure snmpv3 add mib-view allMIB subtree 1.3
```

The following command creates the view extremeMib with the subtree 1.3.6.1.4.1.1916 included as non-volatile:

```
configure snmpv3 add mib-view extremeMib subtree 1.3.6.1.4.1.1916
```

The following command creates a view vrrpTrapNewMaster which excludes VRRP notification .1 and the entry is volatile:

```
configure snmpv3 add mib-view vrrpTrapNewMaster 1.3.6.1.2.1.68.0.1/ff8 type  
excluded volatile
```

History

This command was first available in ExtremeXOS 10.1.

The hex_view_name parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 add notify

```
configure snmpv3 add notify [[hex hex_notify_name] | notify_name] tag [[hex  
hex_tag] | tag] {type [trap | inform]}{volatile}
```



Description

Adds an entry to the snmpNotifyTable.

Syntax Description

<i>hex_notify_name</i>	Specifies the notify name to add. The value is to be supplied as a colon separated string of hex octets.
<i>notify_name</i>	Specifies the notify name to add in ASCII format.
<i>hex_tag</i>	Specifies a string identifier for the notifications to be sent to the target. The value is supplied as a colon separated string of octets.
<i>tag</i>	Specifies a string identifier for the notifications to be sent to the target in ASCII format.
trap	Specifies an unconfirmed notification.
inform	Specifies a confirmed notification.
volatile	Specifies volatile storage. By specifying volatile storage, the configuration is not saved across a switch reboot.

Default

The default storage type is non-volatile.

The default type is trap.

Usage Guidelines

Use this command to add an entry to the snmpNotifyTable. When a notification is to be sent, this table is examined. For the target addresses that have been associated with the tags present in the table, notifications are sent based on the filters also associated with the target addresses.

Example

The following command sends notifications to addresses associated with the tag type1:

```
configure snmpv3 add notify N1 tag type1
```

History

This command was first available in ExtremeXOS 10.1.

The hex_notify_name and hex_tag parameters were added in ExtremeXOS 11.0.

The INFORM option was added in ExtremeXOS 12.5.3.



Platform Availability

This command is available on all platforms.

configure snmpv3 add target-addr

```
configure snmpv3 add target-addr [[hex hex_addr_name] | addr_name] param [[hex
hex_param_name] | param_name ] ipaddress [ ip_address | ipv4-with-mask
ip_and_tmask ] | [ ipv6_address | ipv6-with-mask ipv6_and_tmask ]] {transport-
port port_number} {from [src_ip_address | src_ipv6_address]} {vr vr_name} {tag-
list tag_list} {volatile}
```

Description

Adds and configures an SNMPv3 target address and associates filtering, security, and notifications with that address.

Syntax Description

<i>hex_addr_name</i>	Specifies a string identifier for the target address. The value is to be supplied as a colon separated string of hex octets.
<i>addr_name</i>	Specifies a string identifier for the target address in ASCII format.
<i>hex_param_name</i>	Specifies the parameter name associated with the target. The value is to be supplied as a colon separated string of hex octets.
<i>param_name</i>	Specifies the parameter name associated with the target in ASCII format.
<i>ip_address</i>	Specifies an SNMPv3 target IPv4 address.
ipv4-with-mask	Specify IPv4 address with hexadecimal mask.
<i>ip_and_tmask</i>	Specifies the IPv4 address and hexadecimal mask in form A.B.C.D/NN...
<i>ipv6_address</i>	Specifies an SNMPv3 target IPv6 address.
ipv6-with-mask	Specify IPv6 address with hexadecimal mask.
<i>ipv6_and_tmask</i>	Specifies an IPv6 address and hexadecimal mask in form A:B:C:D:E:F:G:H/NN...
<i>port_number</i>	Specifies a UDP port. Default is 162.
<i>src_ip_address</i>	Specifies the IPv4 address of a VLAN to be used as the source address for the trap.
<i>src_ipv6_address</i>	Specifies the IPv6 address of a VLAN to be used as the source address for the trap.
<i>vr_name</i>	Specifies the name of the virtual router.
tag-list	Specifies a list of comma separated string identifiers for the notifications to be sent to the target.
volatile	Specifies volatile storage. By specifying volatile storage, the configuration is not saved across a switch reboot.



Default

The default values are:

- transport-port—port 162.
- non-volatile storage.

If you do not specify tag-list the single tag defaultNotify, a pre-defined value in the snmpNotifyTable is used.

Usage Guidelines

Use this command to create an entry in the SNMPv3 snmpTargetAddressTable. The **param** parameter associates the target address with an entry in the snmpTargetParamsTable, which specifies security and storage parameters for messages to the target address, and an entry in the snmpNotifyFilterProfileTable, which specifies filter profiles to use for notifications to the target address. The filter profiles are associated with the filters in the snmpNotifyFilterTable.

The list of tag-lists must match one or more of the tags in the snmpNotifyTable for the trap to be sent out.

Example

The following command specifies a target address of 10.203.0.22 with the name A1, and associates it with the security parameters and target address parameter P1:

```
configure snmpv3 add target-addr A1 param P1 ipaddress 10.203.0.22
```

The following command specifies a target address of 10.203.0.22 with the name A1, and associates it with the security parameters and target address parameter P1, and the notification tags type1 and type2:

```
configure snmpv3 add target-addr A1 param P1 ipaddress 10.203.0.22 from  
10.203.0.23 tag-list type1,type2
```

History

This command was first available in ExtremeXOS 10.1.

The virtual router, IP address and hexadecimal mask parameters were added in ExtremeXOS 12.3.

IPv6 support was added in ExtremeXOS 12.4.

The **IPv4-with-mask** and **IPv6-with-mask** keywords were added in ExtremeXOS 15.3.2.

Platform Availability

This command is available on all platforms.



configure snmpv3 add target-params

```
configure snmpv3 add target-params [[hex hex_param_name] | param_name ] user
[[hex hex_user_name] | user_name ] mp-model [snmpv1 | snmpv2c | snmpv3] sec-model
[snmpv1 | snmpv2c | usm] {sec-level [noauth | authnopriv | priv]} {volatile}
```

Description

Adds and configures SNMPv3 target parameters.

Syntax Description

<i>hex_param_name</i>	Specifies the parameter name associated with the target. The value is to be supplied as a colon separated string of hex octets.
<i>param_name</i>	Specifies the parameter name associated with the target in ASCII format.
<i>hex_user_name</i>	Specifies a user name. The value is to be supplied as a colon separated string of hex octets.
<i>user_name</i>	Specifies a user name in ASCII format.
mp-model	Specifies a message processing model; choose from SNMPv1, SNMPv2, or SNMPv3.
sec-model	Specifies the security model to use.
snmpv1	Specifies the SNMPv1 security model.
snmpv2c	Specifies the SNMPv2c security model.
usm	Specifies the SNMPv3 User-based Security Model (USM).
sec-level	Specifies the security level for the group.
noauth	Specifies no authentication (and implies no privacy) for the security level.
authnopriv	Specifies authentication and no privacy for the security level.
priv	Specifies authentication and privacy for the security level.
volatile	Specifies volatile storage. By specifying volatile storage, the configuration is not saved across a switch reboot.

Default

The default values are:

- **sec-level—noauth**.
- non-volatile storage.

Usage Guidelines

Use this command to create an entry in the SNMPv3 snmpTargetParamsTable. This table specifies the message processing model, security level, security model, and the storage parameters for messages to any target addresses associated with a particular parameter name.



To associate a target address with a parameter name, see the command `configure snmpv3 add target-addr`.

Example

The following command specifies a target parameters entry named P1, a user name of guest, message processing and security model of SNMPv2c, and a security level of no authentication:

```
configure snmpv3 add target-params P1 user guest mp-model snmpv2c sec-model
snmpv2c sec-level noauth
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_param_name` and `hex_user_name` parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 add user

```
configure snmpv3 add user [ hex hex_user_name | user_name ] {engine-id engine_id}
{ authentication [md5 | sha] [ {auth-encrypted} localized-key auth_localized_key
| {auth-encrypted} hex hex_auth_password | auth_password] {privacy {des | 3des |
aes {128 | 192 | 256}} [ {privacy-encrypted} localized-key priv_localized_key |
{privacy-encrypted} hex hex_priv_password | <priv_password>]} } {volatile}
```

Description

Adds (and modifies) an SNMPv3 user.

Syntax Description

<i>hex_user_name</i>	Specifies the user name to add or modify. The value is to be supplied as a colon separated string of hex octets.
<i>user_name</i>	Specifies the user name to add or modify in ASCII format.
engine-id	SNMP engine id. If not specified, the user is created with the local engine id.
<i>engine_id</i>	Engine id (in hexadecimal)"; type="ostring_t
authentication	Specifies the authentication password or hex string to use for generating the authentication key for this user.
md5	Specifies RSA Data Security, Inc. MD5 Message-Digest Algorithm authentication.
sha	Specifies SHA authentication.



localized-key	Following value is a MD5 or SHA digest of the engine-id and user's password.
<i>auth_localized_key</i>	Authentication localized key (in hexadecimal) ;type="ostring_t"
privacy	Specifies the privacy password or hex string to use for generating the privacy key for this user.
des	Specifies the use of the 56-bit DES algorithm for encryption. This is the default.
3des	Specifies the use of the 168-bit 3DES algorithm for encryption.
aes	Specifies the use of the AES algorithm for encryption.
128	Specifies the use of the 128-bit AES algorithm for encryption.
192	Specifies the use of the 192-bit AES algorithm for encryption.
256	Specifies the use of the 256-bit AES algorithm for encryption.
<i>priv_localized_key</i>	Privacy localized key (in hexadecimal)"; type="ostring_t"
volatile	Specifies volatile storage. By specifying volatile storage, the configuration is not saved across a switch reboot.

Default

The default values are:

- authentication—no authentication.
- privacy—no privacy.
- non-volatile storage.

Usage Guidelines

Use this command to create or modify an SNMPv3 user configuration.

The default user names are: admin, initial, initialmd5, initialsha, initialmd5Priv, initialshaPriv. The initial password for admin is password. For the other default users, the initial password is the user name.

If hex is specified, supply a 16 octet hex string for RSA Data Security, Inc. MD5 Message-Digest Algorithm, or a 20 octet hex string for SHA.

You must specify authentication if you want to specify privacy. There is no support for privacy without authentication.



Note

3DES, AES 192, and AES 256 bit encryptions are proprietary implementations and may not work with some SNMP managers.



Example

The following command configures the user `guest` on the local SNMP Engine with security level `noauth` (no authentication and no privacy):

```
configure snmpv3 add user guest
```

The following command configures the user `authMD5` to use RSA Data Security, Inc. MD5 Message-Digest Algorithm authentication with the password `palertyu`:

```
configure snmpv3 add user authMD5 authentication md5 palertyu
```

The following command configures the user `authShapriv` to use SHA authentication with the hex key shown below, the privacy password `palertyu`, and volatile storage:

```
configure snmpv3 add user authShapriv authentication sha hex
01:03:04:05:01:05:02:ff:ef:cd:12:99:34:23:ed:ad:ff:ea:cb:11 privacy palertyu
volatile
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_user_name` parameter was added in ExtremeXOS 11.0.

Support for 3DES and AES was added in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.

configure snmpv3 add user clone-from

```
configure snmpv3 add user [[hex hex_user_name] | user_name] {engine-id
engine_id}clone-from [[hex hex_user_name] | user_name] {engine-id
clone_from_engine_id}
```

Description

Creates a new user by cloning from an existing SNMPv3 user.

Syntax Description

<i>hex_user_name</i>	Specifies the user name to add or to clone from. The value is to be supplies as a colon separated string of hex octets.
<i>user_name</i>	Specifies the user name to add or to clone from in ASCII format.



engine-id	SNMP engine ID
<i>engine_id</i>	Engine ID of the user to be added in hexadecimal format. Default: local engine ID"; type="ostring_t
<i>clone_from_engine_id</i>	Engine ID of the user to be cloned in hexadecimal (Default: local engine ID)"; type="ostring_t

Default

N/A.

Usage Guidelines

Use this command to create a new user by cloning an existing one. After you have successfully cloned the new user, you can modify its parameters using the following command:

```
configure snmpv3 add user [[hex hex_user_name] | user_name] {authentication {md5 | sha} [hexhex_auth_password | auth_password]} {privacy {des | 3des | aes {128 | 192 | 256}} [[hexhex_priv_password] | priv_password]} {volatile}
```

Users cloned from the default users will have the storage type of non-volatile. The default names are: admin, initial, initialmd5, initialsha, initialmd5Priv, initialshaPriv.

Example

The following command creates a user cloneMD5 with same properties as the default user initialmd5. All authorization and privacy keys will initially be the same as with the default user initialmd5.

```
configure snmpv3 add user cloneMD5 clone-from initialmd5
```

The following command adds a remote user named nmsuser2 belonging to the SNMP engine with engine-id 11:22:33 by cloning another remote user named nmsuser1 belonging to the SNMP engine with engine id AA:BB:CC:

```
conf snmpv3 add user nmsuser2 engine-id 11:22:33 clone-from nmsuser1 engine-id AA:BB:CC
```

History

This command was first available in ExtremeXOS 10.1.

The hex_user_name parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.



configure snmpv3 delete access

```
configure snmpv3 delete access [all-non-defaults | {[[hex hex_group_name] |
group_name] {sec-model [snmpv1 | snmpv2c | usm] sec-level [noauth | authnopriv |
priv]}}]
```

Description

Deletes access rights for a group.

Syntax Description

all-non-defaults	Specifies that all non-default (non-permanent) security groups are to be deleted.
<i>hex_group_name</i>	Specifies the group name to be deleted. The value is to be supplies as a colon separated string of hex octets.
<i>group_name</i>	Specifies the group name to be deleted in ASCII format.
sec-model	Specifies the security model to use.
snmpv1	Specifies the SNMPv1 security model.
snmpv2c	Specifies the SNMPv2c security model.
usm	Specifies the SNMPv3 User-based Security Model (USM).
sec-level	Specifies the security level for the group.
noauth	Specifies no authentication (and implies no privacy) for the security level.
authnopriv	Specifies authentication and no privacy for the security level.
priv	Specifies authentication and privacy for the security level.

Default

The default values are:

- sec-model—USM.
- sec-level—noauth.

Usage Guidelines

Use this command to remove access rights for a group. Use the all-non-defaults keyword to delete all the security groups, except for the default groups. The default groups are: admin, initial, v1v2c_ro, v1v2c_rw.

Deleting an access will not implicitly remove the related group to user association from the VACMSecurityToGroupTable. To remove the association, use the following command:

```
configure snmpv3 delete group {[[hex hex_group_name] |group_name] } user [all-non-
defaults | {[[hexhex_user_name] |user_name] {sec-model [snmpv1|snmpv2c|usm]}}]
```



Example

The following command deletes all entries with the group name userGroup:

```
configure snmpv3 delete access userGroup
```

The following command deletes the group userGroup with the security model snmpv1 and security level of authentication and no privacy (authnopriv):

```
configure snmpv3 delete access userGroup sec-model snmpv1 sec-level authnopriv
```

History

This command was first available in ExtremeXOS 10.1.

The hex_group_name parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 delete community

```
configure snmpv3 delete community [all-non-defaults | {[[hex hex_community_index]  
| community_index} | {name [[hex hex_community_name] | community_name}]
```

Description

Deletes an SNMPv3 community entry.

Syntax Description

all-non-defaults	Specifies that all non-default community entries are to be removed.
<i>hex_community_index</i>	Specifies the row index in the snmpCommunityTable. The value is to be supplied as a colon separated string of hex octets.
<i>community_index</i>	Specifies the row index in the snmpCommunityTable in ASCII format.
<i>hex_community_name</i>	Specifies the community name. The value is to be supplied as a colon separated string of hex octets.
<i>community_name</i>	Specifies the community name in ASCII format.

Default

The default entries are public and private.



Usage Guidelines

Use this command to delete an SMMPv3 community in the community MIB.

Example

The following command deletes an entry with the community index `comm_index`:

```
configure snmpv3 delete community comm_index
```

The following command creates an entry with the community name (hex) of `EA:12:CD:CF:AB:11:3C`:

```
configure snmpv3 delete community name hex EA:12:CD:CF:AB:11:3C
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_community_index` and `hex_community_name` parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 delete filter

```
configure snmpv3 delete filter [all | [[hex hex_profile_name] | profile_name]  
{subtree object_identifier}]
```

Description

Deletes a filter from a filter profile.

Syntax Description

all	Specifies all filters.
<i>hex_profile_name</i>	Specifies the filter profile of the filter to delete. The value is to be supplied as a colon separated string of hex octets.
<i>profile_name</i>	Specifies the filter profile of the filter to delete in ASCII format.
<i>object_identifier</i>	Specifies the MIB subtree of the filter to delete.

Default

N/A.



Usage Guidelines

Use this command to delete a filter entry from the snmpNotifyFilterTable. Specify all to remove all entries. Specify a profile name to delete all entries for that profile name. Specify a profile name and a subtree to delete just those entries for that filter profile and subtree.

Example

The following command deletes the filters from the filter profile prof1 that reference the MIB subtree 1.3.6.1.4.1:

```
configure snmpv3 delete filter prof1 subtree 1.3.6.1.4.1
```

History

This command was first available in ExtremeXOS 10.1.

The hex_profile_name parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 delete filter-profile

```
configure snmpv3 delete filter-profile [all | [hex hex_profile_name] | profile_name] {param[ [hex hex_param_name] | param_name ] }
```

Description

Removes the association of a filter profile with a parameter name.

Syntax Description

all	Specifies all filter profiles.
<i>hex_profile_name</i>	Specifies the filter profile name to delete. The value is to be supplied as a colon separated string of hex octets.
<i>profile_name</i>	Specifies the filter profile name to delete in ASCII format.
<i>hex_param_name</i>	Specifies to delete the filter profile with the specified profile name and parameter name. The value is to be supplied as a colon separated string of hex octets.
<i>param_name</i>	Specifies to delete the filter profile with the specified profile name and parameter name in ASCII format.



Default

The default storage type is non-volatile.

Usage Guidelines

Use this command to delete entries from the snmpNotifyFilterProfileTable. This table associates a filter profile with a parameter name. Specify all to remove all entries. Specify a profile name to delete all entries for that profile name. Specify a profile name and a parameter name to delete just those entries for that filter profile and parameter name.

Example

The following command deletes the filter profile prof1 with the parameter name P1:

```
configure snmpv3 delete filter-profile prof1 param P1
```

History

This command was first available in ExtremeXOS 10.1.

The hex_profile_name and hex_param_name parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 delete group user

```
configure snmpv3 delete group {[[hex hex_group_name] | group_name]} user [all-  
non-defaults | {[[hex hex_user_name] | user_name] {sec-model [snmpv1|snmpv2c|  
usm]}]}]
```

Description

Deletes a user name (security name) from a group.

Syntax Description

<i>hex_group_name</i>	Specifies the group name to delete or modify. The value is to be supplied as a colon separated string of hex octets.
<i>group_name</i>	Specifies the group name to delete or modify in ASCII format.
all-non-defaults	Specifies that all non-default (non-permanent) users are to be deleted from the group.
<i>hex_user_name</i>	Specifies the user name to delete or modify. The value is to be supplied as a colon separated string of hex octets.



<i>user_name</i>	Specifies the user name to delete or modify in ASCII format.
sec-model	Specifies the security model to use.
snmpv1	Specifies the SNMPv1 security model.
snmpv2c	Specifies the SNMPv2c security model.
usm	Specifies the SNMPv3 User-based Security Model (USM).

Default

The default value for `sec-model` is USM.

Usage Guidelines

Use this command to remove the associate of a user name with a group.

As per the SNMPv3 RFC, a security name is model independent while a username is model dependent. For simplicity, both are assumed to be same here. User names and security names are handled the same. In other words, if a user is created with the user name `username`, the security name value is the same, `username`.

Every group is uniquely identified by a security name and security model. So the same security name can be associated to a group name but with different security models.

The default groups are: `admin`, `initial`, `v1v2c_ro`, `v1v2c_rw`.

The default users are: `admin`, `initial`, `initialmd5`, `initialsha`, `initialmd5Priv`, `initialshaPriv`.

Example

The following command deletes the user `guest` from the group `UserGroup` for the security model `snmpv2c`:

```
configure snmpv3 delete group UserGroup user guest sec-model snmpv2c
```

The following command deletes the user `guest` from the group `userGroup` with the security model USM:

```
configure snmpv3 delete group userGroup user guest
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_group_name` and the `hex_user_name` parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.



configure snmpv3 delete mib-view

```
configure snmpv3 delete mib-view [all-non-defaults | {[[hex hex_view_name] | view_name] {subtree object_identifier}}]
```

Description

Deletes a MIB view.

Syntax Description

all-non-defaults	Specifies that all non-default (non-permanent) MIB views are to be deleted.
<i>hex_view_name</i>	Specifies the MIB view to delete. The value is to be supplied as a colon separated string of hex octets.
<i>view_name</i>	Specifies the MIB view name to delete in ASCII format.
<i>object_identifier</i>	Specifies a MIB subtree.

Default

N/A.

Usage Guidelines

Use this command to delete a MIB view. Views which are being used by security groups cannot be deleted. Use the all-non-defaults keyword to delete all the MIB views (not being used by security groups) except for the default views. The default views are: defaultUserView, defaultAdminView, and defaultNotifyView.

Use the `configure snmpv3 add mib-view` command to remove a MIB view from its security group, by specifying a different view.

Example

The following command deletes all views (only the permanent views will not be deleted):

```
configure snmpv3 delete mib-view all-non-defaults
```

The following command deletes all subtrees with the view name AdminView:

```
configure snmpv3 delete mib-view AdminView
```



The following command deletes the view AdminView with subtree 1.3.6.1.2.1.2

```
configure snmpv3 delete mib-view AdminView subtree 1.3.6.1.2.1.2
```

History

This command was first available in ExtremeXOS 10.1.

The hex_view_name parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 delete notify

```
configure snmpv3 delete notify [{{[[hex hex_notify_name] | notify_name]} | all-non-defaults]
```

Description

Deletes an entry from the snmpNotifyTable.

Syntax Description

<i>hex_notify_name</i>	Specifies the notify name to add. The value is to be supplied as a colon separated string of hex octets.
<i>notify_name</i>	Specifies the notify name to add in ASCII format.
all-non-defaults	Specifies that all non-default (non-permanent) notifications are to be deleted.

Default

N/A.

Usage Guidelines

Use this command to delete an entry from the snmpNotifyTable. When a notification is to be sent, this table is examined. For the target addresses that have been associated with the tags present in the table, notifications will be sent, based on the filters also associated with the target addresses.



Example

The following command removes the N1 entry from the table:

```
configure snmpv3 delete notify N1
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_notify_name` parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 delete target-addr

```
configure snmpv3 delete target-addr [ {[hex hex_addr_name] | addr_name} | all]
```

Description

Deletes SNMPv3 target addresses.

Syntax Description

<i>hex_addr_name</i>	Specifies an identifier for the target address. The value is to be supplied as a colon separated string of hex octets.
<i>addr_name</i>	Specifies a string identifier for the target address.
all	Specifies all target addresses.

Default

N/A.

Usage Guidelines

Use this command to delete an entry in the SNMPv3 `snmpTargetAddressTable`.

Example

The following command deletes target address named A1:

```
configure snmpv3 delete target-addr A1
```



History

This command was first available in ExtremeXOS 10.1.

The `hex_addr_name` parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure snmpv3 delete target-params

```
configure snmpv3 delete target-params [{[[hex hex_param_name] | param_name}] | all]
```

Description

Deletes SNMPv3 target parameters.

Syntax Description

<i>hex_param_name</i>	Specifies the parameter name associated with the target. The value is to be supplied as a colon separated string of hex octets.
<i>param_name</i>	Specifies the parameter name associated with the target in ASCII format.

Default

N/A.

Usage Guidelines

Use this command to delete an entry in the SNMPv3 `snmpTargetParamsTable`. This table specifies the message processing model, security level, security model, and the storage parameters for messages to any target addresses associated with a particular parameter name.

Example

The following command deletes a target parameters entry named P1:

```
configure snmpv3 delete target-params P1
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_param_name` parameter was added in ExtremeXOS 11.0.



Platform Availability

This command is available on all platforms.

configure snmpv3 delete user

```
configure snmpv3 delete user [all-non-defaults | [[hex hex_user_name] |  
user_name] {engine-id engine_id}
```

Description

Deletes an existing SNMPv3 user.

Syntax Description

all-non-defaults	Specifies that all non-default (non-permanent) users are to be deleted.
<i>hex_user_name</i>	Specifies the user name to delete. The value is to be supplied as a colon separated string of hex octets.
<i>user_name</i>	Specifies the user name to delete.
engine-id	SNMP engine ID
<i>engine-id</i>	Engine ID in hexadecimal (Default: local engine ID)"; type="ostring_t

Default

N/A.

Usage Guidelines

Use this command to delete an existing user.

Use the all-non-defaults keyword to delete all users, except for the default users. The default user names are: admin, initial, initialmd5, initialsha, initialmd5Priv, initialshaPriv.

Deleting a user does not implicitly remove the related group to user association from the VACMSecurityToGroupTable. To remove the association, use the following command:

```
configure snmpv3 delete group {[[hex hex_group_name] | group_name]} user [all-  
non-defaults | [[hex hex_user_name] | user_name] {sec-model [snmpv1|snmpv2c|  
usm]}}
```

Example

The following command deletes all non-default users:

```
configure snmpv3 delete user all-non-defaults
```



The following command deletes the user guest:

```
configure snmpv3 delete user guest
```

The following command deletes a remote user named ambiguoususer with engine id 11:22:33:

```
configure snmpv3 delete user ambiguoususer engine-id 11:22:33
```

History

This command was first available in ExtremeXOS 10.1.

The hex_user_name parameter was added in ExtremeXOS 11.0.

The engine_id keyword was added in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

configure snmpv3 engine-boots

```
configure snmpv3 engine-boots (1-2147483647)
```

Description

Configures the SNMPv3 Engine Boots value.

Syntax Description

(1-2147483647)	Specifies the value of engine boots.
----------------	--------------------------------------

Default

N/A.

Usage Guidelines

Use this command if the Engine Boots value needs to be explicitly configured. Engine Boots and Engine Time will be reset to one (1) if the Engine ID is changed. Engine Boots can be set to any desired value, but will latch on its maximum, 2147483647.



Example

The following command configures Engine Boots to 4096:

```
configure snmpv3 engine-boots 4096
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure snmpv3 engine-id

```
configure snmpv3 engine-id hex_engine_id
```

Description

Configures the SNMPv3 snmpEngineID.

Syntax Description

<i>hex_engine_id</i>	Specifies the colon delimited hex octet that serves as part of the snmpEngineID (5-32 octets).
----------------------	--

Default

The default snmpEngineID is the device MAC address.

Usage Guidelines

Use this command if the snmpEngineID needs to be explicitly configured. The first four octets of the ID are fixed to 80:00:07:7C, which represents Extreme Networks Vendor ID. Once the snmpEngineID is changed, default users are reverted back to their original passwords/keys, while non-default users are removed from the device.

In a chassis, the snmpEngineID is generated using the MAC address of the MSM/MM with which the switch boots first. For MSM/MM hitless failover, the same snmpEngineID is propagated to both of the MSMs/MMs.



Example

The following command configures the snmpEngineID to be 80:00:07:7C:00:0a:1c:3e:11:

```
configure snmpv3 engine-id 00:0a:1c:3e:11
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure snmpv3 target-addr retry

```
configure snmpv3 target-addr [[hex hex_addr_name] | addr_name] retry retry_count
```

Description

Configures SNMPv3 INFORM notification retries.

Syntax Description

<i>hex_addr_name</i>	Specifies a address name in hexadecimal format.
<i>addr_name</i>	Specifies the address name in ASCII format.
<i>retry_count</i>	Specifies the maximum number of times to resend an SNMPv3 inform.

Default

The retry default is 3.

Usage Guidelines

Use this command to configure the number of times an SNMPv3 INFORM message is to be resent to the (notification responder) manager when a response has not been received.

Example

The following command configures a retry count of 5 for the target address A1:

```
configure snmpv3 target-addr A1 retry 5
```



History

This command was first available in ExtremeXOS 12.5.3.

Platform Availability

This command is available on all platforms.

configure snmpv3 target-addr timeout

```
configure snmpv3 target-addr [[hex hex_addr_name] | addr_name] timeout
timeout_val
```

Description

Configures the SNMPv3 INFORM notification timeout.

Syntax Description

<i>hex_addr_name</i>	Specifies the address name in hexadecimal format.
<i>addr_name</i>	Specifies the address name in ASCII format.
<i>timeout_val</i>	Specifies the number of seconds.

Default

The timeout value default is 15 seconds.

Usage Guidelines

Use this command to configure how many seconds to wait for a response before resending an SNMPv3 INFORM.

Example

The following command configures a timeout value of 20 seconds for the target address A1:

```
configure snmpv3 target-addr A1 timeout 20
```

History

This command was first available in ExtremeXOS 12.5.3.

Platform Availability

This command is available on all platforms.



configure sntp-client

```
configure sntp-client [primary | secondary] host-name-or-ip {vr vr_name}
```

Description

Configures an NTP server for the switch to obtain time information.

Syntax Description

primary	Specifies a primary server name.
secondary	Specifies a secondary server name.
<i>host-name-or-ip</i>	Specifies a host name or IPv4 address or IPv6 address.
vr	Specifies use of a virtual router. NOTE: User-created VRs are supported only on the platforms listed for this feature in the Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i> .
<i>vr_name</i>	Specifies the name of a virtual router.

Default

N/A.

Usage Guidelines

Queries are first sent to the primary server. If the primary server does not respond within 1 second, or if it is not synchronized, the switch queries the second server. If the switch cannot obtain the time, it restarts the query process. Otherwise, the switch waits for the sntp-client update interval before querying again.

Example

The following command configures a primary NTP server:

```
configure sntp-client primary 10.1.2.2
```

The following command configures the primary NTP server to use the management virtual router VR-Mgmt:

```
configure sntp-client primary 10.1.2.2 vr VR-Mgmt
```

History

This command was first available in ExtremeXOS 10.1.

The **vr***vr_name* option was added in ExtremeXOS 11.0.



Platform Availability

This command is available on all platforms.

configure sntp-client update-interval

```
configure sntp-client update-interval update-interval
```

Description

Configures the interval between polls for time information from SNTP servers.

Syntax Description

<i>update-interval</i>	Specifies an interval in seconds.
------------------------	-----------------------------------

Default

64 seconds.

Usage Guidelines

None.

Example

The following command configures the interval timer:

```
configure sntp-client update-interval 30
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ssh2 access-profile

```
configure ssh2 access-profile [ access_profile | [[add rule ] [first | [[before | after] previous_rule]]] | delete rule | none ]
```



Description

Configures SSH2 to use an ACL policy or ACL rule for access control.

Syntax Description

<i>access_profile</i>	Specifies an ACL policy.
add	Specifies that an ACL rule is to be added to the Telnet application.
<i>rule</i>	Specifies an ACL rule.
first	Specifies that the new rule is to be added before all other rules.
before	Specifies that the new rule is to be added before a previous rule.
after	Specifies that the new rule is to be added after a previous rule.
<i>previous_rule</i>	Specifies an existing rule in the application.
delete	Specifies that one particular rule is to be deleted.
none	Specifies that all the rules or a policy file is to be deleted.

Default

N/A.

Usage Guidelines

You must be logged in as administrator to configure SSH2 parameters.

- Implement an ACL policy file that permits or denies a specific list of IP addresses and subnet masks for the SSH2 port. You must create the ACL policy file before you can use this command. If the ACL policy file does not exist on the switch, the switch returns an error message indicating that the file does not exist.

In the ACL policy file for SSH2, the “source-address” field is the only supported match condition. Any other match conditions are ignored.

Use the none option to remove a previously configured ACL.

Policy files can also be configured using the following command:

```
enable ssh2 {access-profile [access_profile | none]} {porttcp_port_number} {vr [vr_name | all | default]}
```

- Add an ACL rule to the SSH2 application through this command. Once an ACL is associated with SSH2, all the packets that reach an SSH2 module are evaluated with this ACL and appropriate action (permit or deny) is taken, as is done using policy files.

The permit or deny counters are also updated accordingly regardless of whether the ACL is configured to add counters. To display counter statistics, use the `show access-list counters process ssh2` command.

Only the following match conditions and actions are copied to the client memory. Others that may be in the rule are not copied.



Match conditions:

- Source-address—IPv4 and IPv6
- Actions—Permit or Deny

When adding a new rule, use the `first`, `before`, and `after previous_rule` parameters to position it within the existing rules.

If the SSH2 traffic does not match any of the rules, the default behavior is permit. To deny SSH2 traffic that does not match any of the rules, add a deny all rule at the end of the rule list.

Creating an ACL Policy File

To create an ACL policy file, use the `edit policy` command. For more information about creating and implementing ACL policy files, see [Policy Manager](#) and [ACLs](#) in the *ExtremeXOS Concepts Guide*.

If you attempt to implement a policy that does not exist on the switch, an error message similar to the following appears:

```
Error: Policy /config/MyAccessProfile.pol does not exist on file system
```

If this occurs, make sure the policy you want to implement exists on the switch. To confirm the policies on the switch, use the `configure snmp add community` command. If the policy does not exist, create the ACL policy file.

Example

The following command applies the ACL `MyAccessProfile_2` to SSH2:

```
configure ssh2 access-profile MyAccessProfile_2
```

The following command copies the ACL rule, `DenyAccess` to the SSH2 application in first place:

```
configure ssh2 access-profile add DenyAccess first
```

The following command removes the association of a single rule from the SSH2 application:

```
configure ssh2 access-profile delete DenyAccess
```

The following command removes the association of all ACL policies and rules from the SSH2 application:

```
configure ssh2 access-profile none
```

History

This command was first available in ExtremeXOS 12.5.



Platform Availability

This command is available on all platforms.

configure telnet access-profile

```
configure telnet access-profile [ access_profile | [[add rule ] [first | [[before
| after] previous_rule]]] | delete rule | none ]
```

Description

Configures Telnet to use an ACL policy or ACL rule for access control.

Syntax Description

<i>access_profile</i>	Specifies an ACL policy.
add	Specifies that an ACL rule is to be added to the Telnet application.
<i>rule</i>	Specifies an ACL rule.
first	Specifies that the new rule is to be added before all other rules.
before	Specifies that the new rule is to be added before a previous rule.
after	Specifies that the new rule is to be added after a previous rule.
<i>previous_rule</i>	Specifies an existing rule in the application.
delete	Specifies that one particular rule is to be deleted.
none	Specifies that all the rules or a policy file is to be deleted.

Default

Telnet is enabled with no ACL policies and uses TCP port 23.

Usage Guidelines

You must be logged in as administrator to configure Telnet parameters.

You can restrict Telnet access in the following ways:

- Implement an ACL policy file that permits or denies a specific list of IP addresses and subnet masks for the Telnet port. You must create the ACL policy file before you can use this command. If the ACL policy file does not exist on the switch, the switch returns an error message indicating that the file does not exist.

In the ACL policy file for Telnet, the “source-address” field is the only supported match condition. Any other match conditions are ignored.

Use the none option to remove a previously configured ACL.



- Add an ACL rule to the Telnet application through this command. Once an ACL is associated with Telnet, all the packets that reach a Telnet module are evaluated with this ACL and appropriate action (permit or deny) is taken, as is done using policy files.

The permit or deny counters are also updated accordingly regardless of whether the ACL is configured to add counters. To display counter statistics, use the `show access-list counters process telnet` command.

Only the following match conditions and actions are copied to the client memory. Others that may be in the rule are not copied.

Match conditions:

- Source-address—IPv4 and IPv6
- Actions—Permit or Deny

When adding a new rule, use the `first`, `before`, and `after previous_rule` parameters to position it within the existing rules.

If the Telnet traffic does not match any of the rules, the default behavior is permit. To deny Telnet traffic that does not match any of the rules, add a deny all rule at the end of the rule list.

Creating an ACL Policy File

To create an ACL policy file, use the `edit policy` command. For more information about creating and implementing ACL policy files, see [Policy Manager](#) and [ACLs](#) in the *ExtremeXOS Concepts Guide*.

If you attempt to implement a policy that does not exist on the switch, an error message similar to the following appears:

```
Error: Policy /config/MyAccessProfile.pol does not exist on file system
```

If this occurs, make sure the policy you want to implement exists on the switch. To confirm the policies on the switch, use the `configure snmp add community` command. If the policy does not exist, create the ACL policy file.

Viewing Telnet Information

To display the status of Telnet, including the current TCP port, the virtual router used to establish a Telnet session, and whether ACLs are controlling Telnet access, use the following command: `show management`.

Example

The following command applies the ACL policy `MyAccessProfile_2` to Telnet:

```
configure telnet access-profile MyAccessProfile_2
```



The following command applies the ACL rule DenyAccess to the Telnet application in the first position in the list:

```
configure telnet access-profile add DenyAccess first
```

The following command removes the association of a single ACL rule from the Telnet application:

```
configure telnet access-profile delete DenyAccess
```

The following command removes the association of an ACL policy or all ACL rules from the Telnet application:

```
configure telnet access-profile none
```

History

This command was first available in ExtremeXOS 11.2.

Support for ACL rules for Telnet was added in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

configure telnet port

```
configure telnet port [portno | default]
```

Description

Configures the TCP port used by Telnet for communication.

Syntax Description

<i>portno</i>	Specifies a TCP port number. The default is 23. The range is 1 through 65535. The following TCP port numbers are reserved and cannot be used for Telnet connections: 22, 80, and 1023.
default	Specifies the default Telnet TCP port number. The default is 23.

Default

The switch listens for Telnet connections on Port 23.



Usage Guidelines

You must be logged in as administrator to configure the Telnet port.

The portno range is 1 through 65535. The following TCP port numbers are reserved and cannot be used for Telnet connections: 22, 80, and 1023. If you attempt to configure a reserved port, the switch displays an error message similar to the following:

```
configure telnet port 22
Error: port number is a reserved port
```

If this occurs, select a port number that is not a reserved port.

The switch accepts IPv6 connections.

Example

The following command changes the port used for Telnet to port 85:

```
configure telnet port 85
```

The following command returns the port used for Telnet to the default port of 23:

```
configure telnet port default
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 connections was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure telnet vr

```
configure telnet vr [all | default | vr_name]
```

Description

Configures the virtual router used on the switch for listening for Telnet connections.



Syntax Description

all	Specifies to use all virtual routers for Telnet connections.
default	Specifies to use the default virtual router for Telnet connections. The default router is VR-Mgmt.
<i>vr_name</i>	Specifies the name of the virtual router to use for Telnet connections. NOTE: User-created VRs are supported only on the platforms listed for this feature in the Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i> ,

The default is all.

Usage Guidelines

You must be logged in as administrator to configure the virtual router.

The switch accepts IPv6 connections.

If you specify all, the switch listens on all of the available virtual routers for Telnet connections.

The *vr_name* specifies the name of the virtual router to use for Telnet connections.

If you specify a virtual router name that does not exist, the switch displays an error message similar to the following:

```
configure telnet vr vr-ttt ^ %% Invalid input detected at '^' marker.
```

Example

The following command configures the switch to listen for and receive Telnet requests on all virtual routers:

```
configure telnet vr all
```

History

This command was first available in ExtremeXOS 11.0.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure web http access-profile

```
configure web http access-profile [[[add rule ] [first | [[before | after]
previous_rule]]] | delete rule | none ]
```



Description

Configures HTTP to use an ACL rule for access control.

Syntax Description

add	Specifies that an ACL rule is to be added to the Telnet application.
<i>rule</i>	Specifies an ACL rule.
first	Specifies that the new rule is to be added before all other rules.
before	Specifies that the new rule is to be added before a previous rule.
after	Specifies that the new rule is to be added after a previous rule.
<i>previous_rule</i>	Specifies an existing rule in the application.
delete	Specifies that one particular rule is to be deleted.
none	Specifies that all the rules or a policy file is to be deleted.

Default

N/A.

Usage Guidelines

You must be logged in as administrator to configure HTTP parameters.

Use this command to restrict HTTP access by adding an ACL rule to the HTTP application. Once an ACL is associated with HTTP, all the packets that reach a HTTP module are evaluated with this ACL and appropriate action (permit or deny) is taken, as is done using policy files.

The permit or deny counters are also updated accordingly regardless of whether the ACL is configured to add counters. To display counter statistics, use the `This command is available on all platforms. http` command.

Only the following match conditions and actions are copied to the client memory. Others that may be in the rule are not copied.

Match conditions

- Source-address—IPv4 and IPv6

Actions

- Permit
- Deny

When adding a new rule, use the `first`, `before`, and `after previous_rule` parameters to position it within the existing rules.

If the SNMP traffic does not match any of the rules, the default behavior is permit. To deny SNMP traffic that does not match any of the rules, add a deny all rule at the end of the rule list.



Example

The following command copies the ACL rule, DenyAccess to the HTTP application in first place:

```
configure web http access-profile add DenyAccess first
```

The following command removes the association of the ACL rule DenyAccess from the HTTP application:

```
configure web http access-profile delete DenyAccess
```

The following command removes the association of all ACL rules from the HTTP application:

```
configure web http access-profile none
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

create network-clock ptp

```
create network-clock ptp [boundary | ordinary] {domain domain_number} | end-to-end-transparent ]
```

Description

Creates PTP clock instance and defines the mode of operation.

Syntax Description

boundary	Create the clock instance as a boundary clock.
ordinary	Create the clock instance as an ordinary clock.
<i>domain_number</i>	PTP domain number (default 0, range 0 to 255).
end-to-end-transparent	Create the clock instance as an end-to-end transparent clock.

Default

The PTP domain number defaults to 0 for boundary and ordinary clock instances.



Usage Guidelines

Use this command to create a PTP clock instance, and administratively configure the mode of operation of PTP on this instance. You can provision a boundary or ordinary clock instance to synchronize the node with another node with the most precise clock. In boundary clock configuration, the device synchronizes with the grand-master, or another boundary clock, and operates as a master clock for downstream nodes. In ordinary clock configuration, the device synchronizes with the grand-master, or another boundary clock, and acts as a slave. The ordinary clock is by default in the slave-only mode of operation, and does not propagate the clock downstream. The ordinary clock cannot have more than one clock port.

The end-to-end-transparent clock can be provisioned to correct for the residence delay incurred by PTP event packets passing through the switch (referred as residence time).

Note



You can create a maximum of two clock instances in the switch—one boundary clock and one end-to-end transparent clock, or one ordinary clock and one end-to-end transparent clock. The boundary and ordinary clock instances cannot be simultaneously provisioned in the switch.

After you enable a boundary clock, you cannot create an ordinary clock. However, you can delete the boundary clock instance and create a new one in order to change the domain number. To create an ordinary clock instance in the switch that has the boundary clock instance enabled, delete the boundary clock instance, save the configuration and reboot the switch. After the reboot, you can create and enable the ordinary clock instance.

Similarly, to create and enable a boundary clock in a switch that has an ordinary clock enabled, delete the ordinary clock instance, save the configuration and reboot the switch. After the reboot you can create and enable a boundary clock.

The following message is displayed when you create the boundary clock instance in a device with no prior clock instances:

Warning: The ordinary clock cannot be created after enabling the boundary clock. A delete followed by save and reboot are required to create the ordinary clock.

After you enable a boundary clock instance, if you delete the instance and try to create an ordinary clock instance, the above message is displayed as an error, and the ordinary clock instance is not created.

Example

The following command creates an ordinary clock on domain 5:

```
create network-clock ptp ordinary domain 5
```



The following command creates a boundary clock on default domain (domain 0):

```
create network-clock ptp boundary domain 0
```

The following command creates an end-to-end transparent clock:

```
create network-clock ptp end-to-end-transparent
```

History

This command was first available in ExtremeXOS 15.1.

The ordinary clock parameter was added in ExtremeXOS 15.1 Revision 2.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400) and Summit X770.



Note

PTP commands can be used only with the Network Timing feature pack.

create ntp key

```
create ntp key keyid md5 key_string
```

Description

Enables an NTP key for an NTP session.

Syntax Description

<i>keyid</i>	Specifies the key ID as a value from 1 to 65534.
<i>key_string</i>	Specifies an alphanumeric key string, from 5 to 32 numbers or characters, or a combination of both.

Default

N/A.

Usage Guidelines

N/A.



Example

The following command enables an NTP key using RSA Data Security, Inc. MD5 Message-Digest Algorithm encryption on the switch:

```
enable ntp key 1 md5 oklahoma
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on all platforms.

create snmp trap

```
create snmp trap severity severity event EventName msg
```

Description

Creates and sends an SNMP trap containing the information defined in the command.

Syntax Description

<i>severity</i>	Specifies one of the eight severity levels defined in the ExtremeXOS software. Enter one of the following values: critical, error, warning, notice, info, debug-summary, debug-verbose, debug-data.
<i>EventName</i>	Specifies the event name. Enter a name using alphanumeric characters.
<i>msg</i>	Specifies a message. Enter the message using alphanumeric characters.

Default

N/A.

Usage Guidelines

None.

Example

The following example sends a trap of severity info for event AAA with the message user XYZ logged in:

```
create snmp trap severity info event AAA "user XYZ logged in"
```



History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

delete network-clock ptp

```
delete network-clock ptp [boundary | ordinary | end-to-end-transparent]
```

Description

Delete a PTP clock instance.

Syntax Description

boundary	Delete the boundary clock instance.
ordinary	Delete the ordinary clock instance.
end-to-end-transparent	Delete the End-to-End transparent clock instance.

Usage Guidelines

Use this command to delete a PTP boundary, ordinary, or end-to-end transparent clock instance. We recommend that you delete the ordinary or boundary clock instance before you delete the end-to-end transparent clock instance.

Example

The following commands delete an ordinary clock, boundary clock and end-to-end transparent clock:

```
delete network-clock ptp ordinary
delete network-clock ptp boundary
delete network-clock ptp end-to-end-transparent
```

History

This command was first available in ExtremeXOS 15.1.

The `ordinary clock` parameter was added in ExtremeXOS 15.1 Revision 2.



Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400) and Summit X770.



Note

PTP commands can be used only with the Network Timing feature pack.

delete ntp key

```
delete ntp key [keyid | all]
```

Description

Deletes an NTP key; it cannot be used for outgoing or incoming NTP sessions.

Syntax Description

<i>keyid</i>	Specifies the key ID as a value from 1 to 65534.
all	Deletes all keys.

Default

N/A.

Usage Guidelines

N/A.

Example

The following command deletes NTP key 5 on the switch:

```
delete ntp key 5
```

The following command deletes all NTP keys on the switch:

```
delete ntp key all
```

History

This command was first available in ExtremeXOS 12.7.



Platform Availability

This command is available on all platforms.

disable auto-provision

disable auto-provision

Description

Disables the auto provision capability.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to disable the auto provision capability.

To display the status of auto provision on the switch, use the `show auto-provision` command.

Example

The following command disables the auto provision capability:

```
disable auto-provision
```

The following message is displayed:

```
# disable auto-provision  
This setting will take effect at the next reboot of this switch.
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.



disable cdp ports

```
disable cdp ports [port_list | all]
```

Description

Disables CDP on a port.

Syntax Description

<i>port_list</i>	Specifies the list of ports to disable CDP on.
all	Specifies that you disable CDP on all ports.

Default

Enabled.

Usage Guidelines

Example

The following command disables CDP on all ports on the switch:

```
disable cdp ports all
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

disable dhcp vlan

```
disable dhcp vlan [vlan_name | all]
```

Description

Disables the generation and processing of DHCP packets on a VLAN to obtain an IP address for the VLAN from a DHCP server.



Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
all	Specifies all VLANs.

Default

Disabled for all VLANs.

Usage Guidelines

None.

Example

The following command disables the generation and processing of DHCP packets on a VLAN named accounting:

```
disable dhcp vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable network-clock ptp

```
disable network-clock ptp [boundary | ordinary] [{vlan} vlan_name]
```

Description

Disable PTP on a particular clock instance, or on a specified vlan port (clock port) of the clock instance.

Syntax Description

boundary	Boundary clock instance.
ordinary	Ordinary clock instance.
<i>vlan_name</i>	Vlan name.



Default

PTP is disabled by default on a clock instance.

Usage Guidelines

Use this command to disable PTP on a clock instance.

Example

The following example disables the ordinary clock:

```
disable network-clock ptp ordinary
```

The following example disables the clock port lpbk-transit on the boundary clock:

```
disable network-clock ptp boundary vlan lpbk-transit
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400) and Summit X770.



Note

PTP commands can be used only with the Network Timing feature pack.

disable network-clock ptp boundary unicast-negotiation

```
disable network-clock ptp [boundary | ordinary] unicast-negotiation {vlan}
vlan_name
```

Description

Disable unicast negotiation property in the specified clock port. The unicast negotiation disabled clock port rejects the unicast signaling requests from other clock slaves.

Syntax Description

boundary	Boundary clock.
ordinary	Ordinary clock.



vlan_name	VLAN name.
all	All VLANs.

Default

N/A.

Usage Guidelines

The unicast negotiation feature is currently unsupported, and this command is retained to provide configuration compatibility to previous releases.

Example

Example output not yet available and will be provided in a future release.

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400) and Summit X770.



Note

PTP commands can be used only with the Network Timing feature pack.

disable ntp broadcast-server

```
disable ntp {vlan} vlan-name broadcast-server
```

Description

Prevents NTP from sending broadcast messages to a VLAN.

Syntax Description

<i>vlan-name</i>	Specifies the name of a particular VLAN.
------------------	--

Default

NTP does not send broadcast messages to a VLAN by default.



Usage Guidelines

N/A.

Example

The following command prevents NTP from sending broadcast messages to a VLAN called “Northwest”:

```
disable ntp vlan Northwest broadcast-server
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on all platforms.

disable snmp access

```
disable snmp access {snmp-v1v2c | snmpv3}
```

Description

Selectively disables SNMP on the switch.

Syntax Description

snmp-v1v2c	Specifies SNMPv1/v2c access only.
snmpv3	Specifies SNMPv3 access only.

Default

Enabled.

Usage Guidelines

Disabling SNMP access does not affect the SNMP configuration (for example, community strings). However, if you disable SNMP access, you will be unable to access the switch using SNMP.

This command allows you to disable either all SNMP access, v1/v2c access only, or v3 access only.

To allow access, use the following command:

```
enable snmp access {snmp-v1v2c | snmpv3}
```



Example

The following command disables all SNMP access on the switch:

```
disable snmp access
```

History

This command was first available in ExtremeXOS 10.1.

SNMPv3 was added to ExtremeXOS 12.2. It was also included in ExtremeXOS 11.6.4 and 12.1.2.

Platform Availability

This command is available on all platforms.

disable snmp access vr

```
disable snmp access vr [vr_name | all]
```

Description

Selectively disables SNMP access on virtual routers.

Syntax Description

<i>vr_name</i>	Specifies the virtual router name.
all	Specifies all virtual routers.

Default

Enabled on all virtual routers.

Usage Guidelines

Use this command to disable SNMP access on any or all virtual routers.

When SNMP access is disabled on a virtual router, the incoming SNMP request is dropped and an EMS message is logged.

To enable SNMP access on virtual routers use the `enable snmp access vr` command.

To display the SNMP configuration and statistics on a specified virtual router, use the `show snmp vr_name` command.



Example

The following command disables SNMP access on the virtual router vr-finance:

```
disable snmp access vr vr-finance
```

History

This command was first available in ExtremeXOS 12.4.2.

Platform Availability

This command is available on all platforms.

disable snmp community

```
disable snmp community alphanumeric-community-string
```

Description

Disables SNMP community strings on the switch.

Syntax Description

<i>alphanumeric-community-string</i>	Specifies the SNMP community string name.
--------------------------------------	---

Default

N/A.

Usage Guidelines

This command allows the administrator to disable an snmp community. It sets the rowStatus of the community to NotInService. When disabled, SNMP access to the switch using the designated community is not allowed.

Example

The following command disables the community string named extreme:

```
disable snmp community extreme
```



History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

disable snmp traps

disable snmp traps

Description

Prevents SNMP traps from being sent from the switch.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command does not clear the SNMP trap receivers that have been configured. The command prevents SNMP traps from being sent from the switch even if trap receivers are configured.

To view if SNMP traps are being sent from the switch, use the `show management` command. The `show management` command displays information about the switch including the enabled/disabled state of SNMP traps being sent.

Example

The following command prevents SNMP traps from being sent from the switch to the trap receivers:

```
disable snmp traps
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.



disable snmpv3

```
disable snmpv3 [default-group | default-user]
```

Description

Selectively disables SNMPv3 default-group or default-user access on the switch.

Syntax Description

default-group	Specifies SNMPv3 default-group.
default-user	Specifies SNMPv3 default-user.

Default

Enabled.

Usage Guidelines

This command is used to disable SNMPv3 default-group or default-user access.

Disabling SNMPv3 default-group access removes access to default-users and user-created users who are part of the default-group. The user-created authenticated SNMPv3 users (who are part of a user-created group) are able to access the switch. By disabling default-users access, the end-user is not able to access the switch/MIBs using SNMPv3 default-user.

The default groups are: admin, initial, v1v2c_ro, v1v2c_rw.

The default users are: admin, initial, initialmd5, initialsha, initialmd5Priv, initialshaPriv.

Example

The following command disables the default group on the switch:

```
disable snmp default-group
```

History

This command was available in ExtremeXOS 12.2.

It was also included in ExtremeXOS 11.6.4 and ExtremeXOS 12.1.2.

Platform Availability

This command is available on all platforms.



disable sntp-client

`disable sntp-client`

Description

Disables the SNTP client.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

SNTP can be used by the switch to update and synchronize its internal clock from a Network Time Protocol (NTP) server. After the SNTP client has been enabled, the switch sends out a periodic query to the indicated NTP server, or the switch listens to broadcast NTP updates. In addition, the switch supports the configured setting for Greenwich Mean Time (GMT) offset and the use of Daylight Savings Time (DST).

Example

The following command disables the SNTP client:

```
disable sntp-client
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable telnet

`disable telnet`

Description

Disables external Telnet services on the system.



Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

You must be logged in as an administrator to enable or disable Telnet.



Note

Telnet sessions between MSMs/MMs or the nodes of a stack are not affected by this command.

Example

With administrator privilege, the following command disables external Telnet services on the switch:

```
disable telnet
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable watchdog

```
disable watchdog
```

Description

Disables the system watchdog timer.

Syntax Description

This command has no arguments or variables.

Default

Enabled.



Usage Guidelines

The watchdog timer monitors the health of the switch hardware and software events. For example, the watchdog timer reboots the switch if the system cannot reset the watchdog timer. This can be caused by a long CPU processing loop, any unhandled exception, or a hardware problem with the communication channel to the watchdog. In most cases, if the watchdog timer expires, the switch captures the current CPU status and posts it to the console and the system log. In some cases, if the problem is so severe that the switch is unable to perform any action, the switch reboots without logging any system status information prior to reboot.

This command takes affect immediately.

The watchdog settings are saved in the configuration file.

To display the watchdog state of your system, use the `show switch` command.

Example

The following command disables the watchdog timer:

```
disable watchdog
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

enable auto-provision

enable auto-provision

Description

Enables the auto provision capability.

Syntax Description

This command has no arguments or variables.

Default

N/A.



Usage Guidelines

Use this command to enable the switch to obtain the IP address, gateway and config file via DHCP.

Following is the mandatory DHCP option configuration used for auto provision to work:

Standard Option:

```
IP address
Subnet mask
Gateway
Option 60:
Vendor identifier option
Option 43:
TFTP server IP address
Configuration file name
Optional DHCP option
SNMP trap receiver IP address
```

To display the status of auto provision on the switch, use the `show auto-provision` command.

Example

The following command enables the auto provision capability:

```
enable auto-provision
```

The following message is displayed:

```
X450a-24t.12 # enable auto-provision
This setting will take effect at the next reboot of this switch.
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

disable ntp

```
disable ntp
```

Description

Disables NTP globally on the switch.



Syntax Description

N/A.

Default

NTP is disabled by default.

Usage Guidelines

N/A.

Example

The following command disables NTP globally on the switch:

```
disable ntp
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on all platforms.

disable ntp authentication

disable ntp authentication

Description

Disables NTP authentication globally on the switch.

Syntax Description

N/A.

Default

NTP authentication is disabled by default.



Usage Guidelines

If authentication is disabled, NTP will not use any authentication mechanism to a server or from clients. To use authentication for a specific server, enable NTP authentication globally and then configure an RSA Data Security, Inc. MD5 Message-Digest Algorithm key index for the specific server.

Example

The following command disables NTP authentication globally on the switch:

```
disable ntp authentication
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on all platforms.

disable ntp broadcast-client

```
disable ntp broadcast-client
```

Description

Disables an NTP broadcast client on the switch.

Syntax Description

N/A.

Default

An NTP broadcast client is enabled by default.

Usage Guidelines

If the broadcast client function is enabled, the system can receive broadcast-based NTP messages and process them only if a VLAN is enabled for NTP and the VLAN is active.



Example

The following command disables an NTP broadcast client on the switch:

```
disable ntp broadcast client
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on all platforms.

disable ntp broadcast-server

```
disable ntp {vlan} vlan-name broadcast-server
```

Description

Prevents NTP from sending broadcast messages to a VLAN.

Syntax Description

<i>vlan-name</i>	Specifies the name of a particular VLAN.
------------------	--

Default

NTP does not send broadcast messages to a VLAN by default.

Usage Guidelines

N/A.

Example

The following command prevents NTP from sending broadcast messages to a VLAN called “Northwest”:

```
disable ntp vlan Northwest broadcast-server
```

History

This command was first available in ExtremeXOS 12.7.



Platform Availability

This command is available on all platforms.

disable ntp vlan

```
disable ntp [{vlan} vlan-name | all]
```

Description

Disables NTP on a VLAN.

Syntax Description

disable	Disables NTP on a VLAN.
<i>vlan-name</i>	Specifies the name of a particular VLAN on which to enable or disable NTP.
all	Enables or disables NTP on all VLANs.

Default

NTP is disabled on all VLANs by default.

Usage Guidelines

N/A.

Example

The following command disables NTP on all VLANs:

```
disable ntp vlan all
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on all platforms.

enable cdp ports

```
enable cdp ports [port_list | all]
```



Description

Enables CDP on a port.

Syntax Description

<i>port_list</i>	Specifies the list of ports to enable CDP on.
all	Specifies that you enable CDP on all ports.

Default

Enabled.

Usage Guidelines

Example

The following command enables CDP on all ports on the switch:

```
enable cdp ports all
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

enable dhcp vlan

```
enable dhcp vlan [vlan_name | all]
```

Description

Enables the generation and processing of DHCP packets on a VLAN to obtain an IP address for the VLAN from a DHCP server.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
all	Specifies all VLANs.



Default

Disabled for all VLANs.

Usage Guidelines

None.

Example

The following command enables the generation and processing of DHCP packets on a VLAN named accounting:

```
enable dhcp vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable network-clock ptp

```
enable network-clock ptp [boundary | ordinary] [{vlan} vlan_name]
```

Description

Enable PTP on a particular clock instance, or on a specified vlan port (clock port) of the clock instance.

Syntax Description

boundary	Boundary clock instance.
ordinary	Ordinary clock instance.
vlan_name	Vlan name.

Default

PTP is disabled by default on a clock instance.

Usage Guidelines

Use this command to enable PTP on the clock instance or on the specified VLAN (clock port).



Example

```
enable network-clock ptp boundary
enable network-clock ptp ordinary
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).



Note

PTP commands can be used only with the Network Timing feature pack.

enable network-clock ptp unicast-negotiation

```
enable network-clock ptp [boundary | ordinary] unicast-negotiation {vlan}
vlan_name]
```

Description

Enable unicast negotiation property in the specified clock port. The unicast negotiation enabled clock port responds to the unicast signaling requests from other clock slaves.

Syntax Description

boundary	Boundary clock.
ordinary	Ordinary clock.
vlan_name	VLAN name.

Default

N/A.

Usage Guidelines

The unicast negotiation feature is currently not supported, and this command is retained to provide configuration compatibility to previous releases.

Example

Example output not yet available and will be provided in a future release.



History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400) and Summit X770.



Note

PTP commands can be used only with the Network Timing feature pack.

enable network-clock ptp end-to-end transparent

```
[enable] network-clock ptp end-to-end-transparent ports port_list
```

Description

Enable PTP end-to-end-transparent clock functionality (1-step PHY timestamp) on the ports.

Syntax Description

<i>port_list</i>	List of physical ports.
------------------	-------------------------

Default

N/A.

Usage Guidelines

See Description.

Example

The following example enables end-to-end transparent clock on the front panel ports 1-3:

```
enable network-clock ptp end-to-end-transparent ports 1-3
```

History

This command was first available in ExtremeXOS 15.1.



Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400) and Summit X770.



Note

PTP commands can be used only with the Network Timing feature pack.

enable snmp access

```
enable snmp access {snmp-v1v2c | snmpv3}
```

Description

Selectively enables SNMP access on the switch.

Syntax Description

snmp-v1v2c	Specifies SNMPv1/v2c access only.
snmpv3	Specifies SNMPv3 access only.

Default

Enabled.

Usage Guidelines

To have access to the SNMP agent residing in the switch, at least one VLAN must have an IP address assigned to it.

Any network manager running SNMP can manage the switch for v1/v2c/v3, provided the MIB is installed correctly on the management station. Each network manager provides its own user interface to the management facilities.

For SNMPv3, additional security keys are used to control access, so an SNMPv3 manager is required for this type of access.

This command allows you to enable either all SNMP access, no SNMP access, v1/v2c access only, or v3 access only.

To prevent any SNMP access, use the following command :

```
disable snmp access {snmp-v1v2c | snmpv3}
```

ExtremeXOS 11.2 introduced the concept of safe defaults mode. Safe defaults mode runs an interactive script that allows you to enable or disable SNMP, Telnet, and switch ports. When you set up your switch for the first time, you must connect to the console port to access the switch. After logging in to the switch, you enter safe defaults mode. Although SNMP, Telnet, and switch ports are enabled by default, the script prompts you to confirm those settings.



If you choose to keep the default setting for SNMP—the default setting is enabled—the switch returns the following interactive script:

```
Since you have chosen less secure management methods, please remember to
increase the security of your network by taking the following actions: *
change your admin password * change your SNMP public and private strings
* consider using SNMPv3 to secure network management traffic
```

In addition, you can return to safe defaults mode by issuing the following command:

```
configure safe-default-script
```

If you return to safe defaults mode, you must answer the questions presented during the interactive script.

For more detailed information about safe defaults mode, see [This switch currently has all management methods enabled for convenience reasons.](#) in the *ExtremeXOS Concepts Guide*.

Example

The following command enables all SNMP access for the switch:

```
enable snmp access
```

History

This command was first available in ExtremeXOS 10.1.

SNMPv3 was added to ExtremeXOS 12.2. It was also included in ExtremeXOS 11.6.4 and 12.1.2.

Platform Availability

This command is available on all platforms.

enable snmp access vr

```
enable snmp access vr [vr_name | all]
```

Description

Selectively enables SNMP access on virtual routers.

Syntax Description

<code>vr_name</code>	Specifies the virtual router name.
<code>all</code>	Specifies all virtual routers.



Default

Enabled on all virtual routers.

Usage Guidelines

Use this command to enable SNMP access on any or all virtual routers.

To disable SNMP access on virtual routers, use the `disable snmp access vr` command.

To display the SNMP configuration and statistics on a specified virtual router, use the `show snmp vr_name` command.

Example

The following command enables SNMP access on the virtual router vr-finance:

```
enable snmp access vr vr-finance
```

History

This command was first available in ExtremeXOS 12.4.2.

Platform Availability

This command is available on all platforms.

enable snmp community

```
enable snmp community alphanumeric-community-string
```

Description

Enables SNMP community strings.

Syntax Description

<code><i>alphanumeric-community-string</i></code>	Specifies the SNMP community string name.
---	---

Default

N/A.



Usage Guidelines

This command allows the administrator to enable an snmp community that has been disabled. It sets the rowStatus of the community to Active.

Example

The following command enables the community string named extreme:

```
enable snmp community extreme
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

enable snmp traps

```
enable snmp traps
```

Description

Turns on SNMP trap support.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

An authorized trap receiver can be one or more network management stations on your network. The switch sends SNMP traps to all trap receivers.

To view if SNMP traps are being sent from the switch, use the `show management` command. The `show management` command displays information about the switch including the enabled/disabled state of SNMP traps being sent.



Example

The following command enables SNMP trap support on the switch:

```
enable snmp traps
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable snmpv3

```
enable snmpv3 [default-group | default-user]
```

Description

Selectively enables SNMPv3 default-group or default-user access on the switch.

Syntax Description

default-group	Specifies SNMPv3 default-group.
default-user	Specifies SNMPv3 default-user.

Default

Enabled.

Usage Guidelines

This command is used to enable SNMPv3 default-group or default-user access.

Enabling SNMPv3 default-group access activates the access to an SNMPv3 default-group and the user-created SNMPv3-user part of default-group. Enabling the SNMPv3 default-user access allows an end user to access the MIBs using SNMPv3 default-user. This command throws an error if the SNMPv3 access is disabled on the switch.

The default groups are: admin, initial, v1v2c_ro, v1v2c_rw.

The default users are: admin, initial, initialmd5, initialsha, initialmd5Priv, initialshaPriv.



Example

The following command enables the default users on the switch:

```
enable snmp default-user
```

History

This command was available in ExtremeXOS 12.2.

It was also included in ExtremeXOS 11.6.4 and ExtremeXOS 12.1.2.

Platform Availability

This command is available on all platforms.

enable sntp-client

enable sntp-client

Description

Enables the SNTP client.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

SNTP can be used by the switch to update and synchronize its internal clock from a Network Time Protocol (NTP) server. After the SNTP client has been enabled, the switch sends out a periodic query to the indicated NTP server, or the switch listens to broadcast NTP updates. In addition, the switch supports the configured setting for Greenwich Mean Time (GMT) offset and the use of Daylight Savings Time (DST).

Example

The following command enables the SNTP client:

```
enable sntp-client
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable telnet

enable telnet

Description

Enables external Telnet services on the system.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

You must be logged in as an administrator to enable or disable Telnet.

ExtremeXOS 11.2 introduces the concept of safe defaults mode. Safe defaults mode runs an interactive script that allows you to enable or disable SNMP, Telnet, and switch ports. When you set up your switch for the first time, you must connect to the console port to access the switch. After logging in to the switch, you enter safe defaults mode. Although SNMP, Telnet, and switch ports are enabled by default, the script prompts you to confirm those settings.

If you choose to keep the default setting for Telnet—the default setting is enabled—the switch returns the following interactive script:

```
Since you have chosen less secure management methods, please remember to
increase the security of your network by taking the following actions: *
change your admin password * change your SNMP public and private strings
* consider using SNMPv3 to secure network management traffic
```

In addition, you can return to safe defaults mode by issuing the following command:

```
configure safe-default-script
```

If you return to safe defaults mode, you must answer the questions presented during the interactive script.

For more detailed information about safe defaults mode, see [Use Safe Defaults Mode](#) in the *ExtremeXOS Concepts Guide*.



Example

With administrator privilege, the following command enables Telnet services on the switch:

```
enable telnet
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

enable watchdog

enable watchdog

Description

Enables the system watchdog timer.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

The watchdog timer monitors the health of the switch hardware and software events. For example, the watchdog timer reboots the switch if the system cannot reset the watchdog timer. This is caused by a long CPU processing loop, any unhandled exception, or a hardware problem with the communication channel to the watchdog. In most cases, if the watchdog timer expires, the switch captures the current CPU status and posts it to the console and the system log. In some cases, if the problem is so severe that the switch is unable to perform any action, the switch reboots without logging any system status information prior to reboot.

This command takes affect immediately.

The watchdog settings are saved in the configuration file.

To display the watchdog state of your system, use the `show switch` command.



Example

The following command enables the watchdog timer:

```
enable watchdog
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

enable ntp

enable ntp

Description

Enables NTP globally on the switch.

Syntax Description

N/A.

Default

NTP is disabled by default.

Usage Guidelines

N/A.

Example

The following command enables NTP globally on the switch:

```
enable ntp
```

History

This command was first available in ExtremeXOS 12.7.



Platform Availability

This command is available on all platforms.

enable ntp authentication

enable ntp authentication

Description

Enables NTP authentication globally on the switch.

Syntax Description

N/A.

Default

NTP authentication is disabled by default.

Usage Guidelines

If authentication is disabled, NTP will not use any authentication mechanism to a server or from clients. To use authentication for a specific server, enable NTP authentication globally and then configure an RSA Data Security, Inc. MD5 Message-Digest Algorithm key index for the specific server.

Example

The following command enables NTP authentication globally on the switch:

```
enable ntp authentication
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on all platforms.

enable ntp broadcast-client

enable ntp broadcast-client



Description

Enables an NTP broadcast client on the switch.

Syntax Description

N/A.

Default

An NTP broadcast client is enabled by default.

Usage Guidelines

If the broadcast client function is enabled, the system can receive broadcast-based NTP messages and process them only if a VLAN is enabled for NTP and the VLAN is active.

Example

The following command enables an NTP broadcast client on the switch:

```
enable ntp broadcast client
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on all platforms.

enable ntp broadcast-server

```
enable ntp {vlan} vlan-name broadcast-server {key keyid}
```

Description

Enables NTP to send broadcast messages with or without a key to a VLAN.

Syntax Description

<i>vlan-name</i>	Specifies the name of a particular VLAN on which to enable or disable NTP.
<i>keyid</i>	Specifies the key ID as a value from 1 to 65534.



Default

An NTP broadcast server is enabled by default.

Usage Guidelines

For the broadcast server function to work correctly, configure a VLAN to forward broadcast packets by using the `enable ipforwarding broadcast vlan-name` command. All broadcast clients will receive clock information from the broadcasted clock messages.

Example

The following command enables an NTP broadcast server on the switch:

```
enable ntp vlan toSW3 broadcast-server key 100
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on all platforms.

enable ntp vlan

```
enable ntp [{vlan} vlan-name | all]
```

Description

Enables NTP on a VLAN.

Syntax Description

enable	Enables NTP on a VLAN.
<i>vlan-name</i>	Specifies the name of a particular VLAN on which to enable or disable NTP.
all	Enables or disables NTP on all VLANs.

Default

NTP is disabled on all VLANs by default.

Usage Guidelines

N/A.



Example

The following command enables NTP on a VLAN named “Southwest”:

```
enable ntp vlan Southwest
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on all platforms.

exit

exit

Description

Logs out the session of a current user for CLI or Telnet.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to log out of a CLI or Telnet session.

When you issue this command, you are asked to save your configuration changes to the current, active configuration. Enter y if you want to save your changes. Enter n if you do not want to save your changes.

Example

The following command logs out the session of a current user for CLI or Telnet:

```
exit
```



A message similar to the following is displayed:

```
Do you wish to save your configuration changes to primary.cfg? (y or n)
```

Enter y if you want to save your changes. Enter n if you do not want to save your changes.

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

logout

logout

Description

Logs out the session of a current user for CLI or Telnet.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to log out of a CLI or Telnet session.

When you issue this command, you are asked to save your configuration changes to the current, active configuration. Enter y if you want to save your changes. Enter n if you do not want to save your changes.

Example

The following command logs out the session of a current user for CLI or Telnet:

```
logout
```



A message similar to the following is displayed:

```
Do you wish to save your configuration changes to primary.cfg? (y or n)
```

Enter y if you want to save your changes. Enter n if you do not want to save your changes.

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

quit

quit

Description

Logs out the session of a current user for CLI or Telnet.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to log out of a CLI or Telnet session.

When you issue this command, you are asked to save your configuration changes to the current, active configuration. Enter y if you want to save your changes. Enter n if you do not want to save your changes.

Example

The following command logs out the session of a current user for CLI or Telnet:

```
quit
```



A message similar to the following is displayed:

```
Do you wish to save your configuration changes to primary.cfg? (y or n)
```

Enter y if you want to save your changes. Enter n if you do not want to save your changes.

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show access-list counters process

```
show access-list counters process [snmp | telnet | ssh2 | http]
```

Description

Displays the access-list permit and deny statistics.

Syntax Description

snmp	Specifies statistics for SNMP.
telnet	Specifies statistics for Telnet.
ssh2	Specifies statistics for SSH2.
http	Specifies statistics for HTTP.

Default

N/A.

Usage Guidelines

Use this command to display the access-list permit and deny statistics. The permit and deny counters are updated automatically regardless of whether the ACL is configured to add counters.

Example

The following command displays permit and deny statistics for the SNMP application:

```
# sh access-list counter process snmp
```



Following is sample output for this command:

```

show access-list counter process snmp
=====
Access-list      Permit Packets      Deny Packets
=====
a1                10                  0
a3                0                   25
a2                0                   6
=====
Total Rules : 3

```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

show auto-provision

```
show auto-provision {{vr} vr_name}
```

Description

Displays the current state of auto provision on the switch.

Syntax Description

<i>vr_name</i>	Specifies the virtual router. This may be VR-Default or VR-Mgmt only
----------------	--

Default

N/A.

Usage Guidelines

Use this command to display the current state and the statistics of the auto provision feature on the switch.

Example

The following command displays all information on the current state of auto provision:

```
show auto-provision
```



Following is sample output for the command when the auto provision is enabled. When “Enabled” the feature can be “In progress”, “Done”, or “Failed.”

```
(Auto-Provision) switch # show auto-provision
-----
VR-Name      Auto-Provision Status  Number of attempts
-----
VR-Default   Enabled (In progress)  2
VR-Mgmt      Enabled (In progress)  1
switch # show auto-provision
-----
VR-Name      Auto-Provision Status  Number of attempts
-----
VR-Default   Enabled (Done)         0
VR-Mgmt      Enabled (Done)         0
```

The following command displays information on the current state of auto provision on VR-Mgmt.

```
show auto-provision vr "VR-Mgmt "
```

Following is sample output for the command when auto provision is disabled:

```
switch # show auto-provision vr "VR-Mgmt "
DHCP Auto-Provision      : Disabled
Number of attempts      : 0
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

show checkpoint-data

```
show checkpoint-data {process}
```

Description

Displays the status of one or more processes being copied from the primary MSM/MM to the backup MSM/MM.

Syntax Description

<i>process</i>	Specifies the name of the processes being copied.
----------------	---



Default

N/A.

Usage Guidelines

This command displays, in percentages, the amount of internal state copying completed by each process and the traffic statistics between the process on both the primary and the backup MSMs/MMs.

This command is also helpful in debugging synchronization problems that occur at run-time. To check the status of synchronizing the MSMs/MMs, use the `show switch` command.

Depending on the software version running on your switch and the type of switch you have, additional or different checkpoint status information may be displayed.

Example

The following command displays the checkpointing status and the traffic statics of all of the processes between the primary and the backup MSM:

```
show checkpoint-data
```

The following is sample output from this command:

Process	Tx	Rx	Errors	Sent	Total	%	Chkpt	Debug-info
devmgr	3812	1731	0	3	3	100%	ON OK	1 (00008853)
dirser	0	0	0	0	0	0%	ON OK	1 (000008D3)
ems	5	0	0	0	0	100%	ON OK	1 (000008D3)
nodemgr	0	0	0	0	0	0%	ON OK	1 (000008D3)
snmpSubagent	0	0	0	0	0	0%	ON OK	1 (000018D3)
snmpMaster	0	0	0	0	0	0%	ON OK	1 (000008D3)
cli	0	0	0	0	0	0%	ON OK	1 (000018D3)
edp	0	0	0	0	0	0%	ON OK	1 (000008D3)
cfgmgr	82	82	0	1	1	100%	ON OK	1 (000018D3)
elrp	0	0	0	0	0	0%	ON OK	1 (000008D3)
vlan	1047	1	0	0	0	100%	ON OK	1 (000008D3)
aaa	0	0	0	0	0	0%	ON OK	1 (000008D3)
fdb	957	2	0	0	0	100%	ON OK	1 (000008D3)
msgsrv	0	0	0	0	0	100%	ON OK	1 (000008D3)
eaps	0	0	0	0	0	0%	ON OK	1 (000008D3)
stp	1	0	0	0	0	0%	ON OK	1 (000008D3)
esrp	1	0	0	0	0	100%	ON OK	1 (000008D3)
polMgr	0	0	0	0	0	0%	ON OK	1 (000008D3)
mcmgr	2	2	0	0	0	100%	ON OK	1 (000008D3)
acl	0	0	0	0	0	100%	ON OK	1 (000008D3)
netLogin	0	0	0	0	0	0%	ON OK	1 (000008D3)
ospf	0	0	0	0	0	0%	ON OK	1 (000008D3)
netTools	1	0	0	0	0	100%	ON OK	1 (000008D3)
telnetd	0	0	0	0	0	0%	ON OK	1 (000008D3)
rtmgr	4	4	0	0	0	100%	ON OK	1 (000008D3)
vrrp	378	0	0	0	0	0%	ON OK	1 (000008D3)
tftpd	0	0	0	0	0	0%	ON OK	1 (000008D3)



tthttpd	0	0	0	0	0	0%	ON	OK	1 (000008D3)
rip	0	0	0	0	0	0%	ON	OK	1 (000008D3)
dosprotect	0	0	0	0	0	0%	ON	OK	1 (000008D3)
epm	0	0	0	0	0	0%	ON	OK	1 (000008D3)
hal	0	0	0	0	0	0%	ON	OK	1 (000008D3)
bgp	0	0	0	0	0	0%	ON	OK	1 (000008D3)
pim	0	0	0	0	0	0%	ON	OK	1 (000008D3)
etmon	185	185	0	0	0	100%	ON	OK	1 (000008D3)

To view the output for a specific process, use the process option. The following command displays detailed information for the STP process:

```
show checkpoint-data stp
```

The following is sample output from this command:

Process	Tx	Rx	Errors	Sent	Total	%	Chkpt	Debug-info
-----	-----	-----	-----	-----	-----	-----	-----	-----
stp	1	0	0	0	0	0%	ON OK	1 (000008D3)

History

This command was first available in ExtremeXOS 10.1.

An error count was added to the output in ExtremeXOS 11.1.

Platform Availability

This command is available only on modular switches and SummitStack.

show dhcp-client state

```
show dhcp-client state
```

Description

Displays the current DHCP/BOOTP client state for each vlan.

Syntax Description

This command has no arguments or variables.

Default

Displays the client state for all existing VLANs.



Usage Guidelines

None.

Example

The following command displays the DHCP/BOOTP status for all VLANs:

```
show dhcp-client state
```

Depending on your configurations, output from this command is similar to the following:

```
Client VLAN      Protocol Server      Current State
-----
Default          BOOTP    10.1.2.3Received IP address configured on vlan
accountingDHCP10.2.3.4DHCP state; Requesting
Mgmt              None      0.0.0.0
A total of 3 vlan(s) where displayed
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show management

show management

Description

Displays the SNMP and CLI settings configured on the switch and the SNMP statistics.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The following management output is displayed:



- Enable/disable state for Telnet, and SNMP access.
- Login statistics.
 - Enable/disable state for idle timeouts.
 - Maximum number of CLI sessions.
- SNMP community strings.
- SNMP trap receiver list.

For ExtremeXOS 11.0 and later, the following management output is also displayed:

- SNMP trap receiver source IP address.
- SNMP statistics counter.
- SSH access states of enabled, disabled, and module not loaded.
- CLI configuration logging.
- SNMP access states of v1, v2c disabled and v3 enabled.

If all three types of SNMP access are enabled or disabled, SNMP access is displayed as either Enabled or Disabled.

For ExtremeXOS 11.1 and later, the following management output is also displayed:

- Enable/disable state for RMON.

For ExtremeXOS 11.2 and later, the following management output is also displayed:

- Access-profile usage configured via Access Control Lists (ACLs) for additional Telnet and SSH2 security.

For ExtremeXOS 11.6 and later, the following management output is also displayed:

- CLI scripting settings
 - Enable/disable state.
 - Error message setting.
 - Persistence mode.

For ExtremeXOS 12.4 and later, the following management output is also displayed:

- Dropped SNMP packet counter.

For ExtremeXOS 12.5 and later, the following management output is also displayed:

- CLI prompting.
- SNMP INFORM.

Example

The following command displays configured SNMP settings on a BlackDiamond 12804 switch:

```
show management
```

The following is sample output from this command:

```
BD-12804.2 # show management
```




```

CLI idle timeout           : Enabled (20 minutes)
CLI max number of login attempts : 3
CLI max number of sessions  : 8
CLI paging                 : Enabled (this session only)
CLI space-completion       : Disabled (this session only)
CLI configuration logging   : Disabled
CLI scripting               : Disabled (this session only)
CLI scripting error mode    : Ignore-Error (this session only)
CLI persistent mode        : Persistent (this session only)
Telnet access              : Enabled (tcp port 23 vr all)
: Access Profile : not set
SSH Access                  : ssh module not loaded.
Web access                  : Disabled (tcp port 80)
: Access Profile : not set
Total Read Only Communities : 1
Total Read Write Communities : 1
RMON                       : Disabled
SNMP access                 : Enabled
: Access Profile : not set
SNMP Traps                  : Enabled
SNMP v1/v2c TrapReceivers  :
Destination                 Source IP Address  Flags  Timeout  Retries
10.120.91.89 /10550 10.127.9.147      2ET   -        -
10.120.91.89 /162
2EI
 60      5
Flags: Version: 1=v1 2=v2c
Mode: S=Standard E=Enhanced
Notification Type: T=Trap I=Inform
SNMP stats:      InPkts 0      OutPkts 134      Errors 0      AuthErrors 0
Gets 0      GetNexts 0      Sets 0      Drops 0
SNMP traps:      Sent 134      AuthTraps Enabled
SNMP inform:     Sent 20      Timeout 2

```

History

This command was first available in ExtremeXOS 10.1.

The trap receiver source IP address, SNMP counter statistics, SSH access, CLI logging, and SNMP access states were added to the output in ExtremeXOS 11.0.

The enabled/disabled state for RMON was added to the output in ExtremeXOS 11.1.

Additional Telnet and SSH2 information about ACL usage was added to the output in ExtremeXOS 11.2.

Information about CLI scripting including, the enabled/disabled state, error mode, and persistent mode was added to the output in ExtremeXOS 11.6.

The dropped SNMP packet counter (Drops) was added to the output in ExtremeXOS 12.4.

CLI prompting was added to the output in ExtremeXOS 12.5.

SNMP INFORM was added in ExtremeXOS 12.5.3.



Platform Availability

This command is available on all platforms.

show network-clock ptp

show network-clock ptp

Description

Show the PTP clock recovery state information. The clock recovery using PTP event messages undergoes the following servo state changes:

- Warmup - The local reference clock is in warmup state.
- Fast Loop - The local reference clock is being corrected and the correction is converging.
- Bridge - The local reference clock correction has been interrupted due to changes in the clocking information in the received PTP event messages or loss of PTP event messages.
- Holdover - Prolonged loss of PTP event messages puts the local reference clock correction to holdover state.
- Normal - The local reference clock correction has converged and the corrected clock is synchronous to the master clock information received in the PTP event messages.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

See Description.

Example

```
E4G-200.1 # show network-clock ptp
=====
Servo State Information
=====
Servo State           : 3 (Normal Loop)
Servo State Duration  : 10285 (sec)
=====
```

History

This command was first available in ExtremeXOS 15.1.



Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).



Note

PTP commands can be used only with the Network Timing feature pack.

show network-clock ptp (datasets)

```
show network-clock ptp [boundary | ordinary] {parent | port | time-property}
```

Description

Show PTP clock and port datasets such as parent, port, default, current and time-property.

Syntax Description

boundary	Specifies boundary clock.
ordinary	Specifies ordinary clock.
parent	Parent clock data set.
port	PTP port data set.
time-property	Time properties data set.

Default

N/A.

Usage Guidelines

Show PTP clock and port datasets such as parent, port, default, current and time-property. If no options are specified, the default, current, parent and time-property datasets are shown for the configured clock instance.

Example

The following command displays network clock information on a modular switch:

```
E4G-200.1 # show network-clock ptp boundary parent
Mode           : Boundary Clock
State          : Enabled
Instance       : 0
=====
Parent Data Set
=====
Parent Port Clock Identity      : 00:04:96:ff:fe:52:d2:42
Parent Port Number              : 2
```



```

Parent Active Peer Addr           : 22.55.0.2
Parent Stats                       : FALSE
Observed Parent Offset Scaled Log Variance : 0
Observed Parent Clock Phase Change Rate : 0
Grandmaster Identity              : 00:11:11:11:11:11:11:11
Grandmaster Clock Class           : 6
Grandmaster Clock Accuracy        : Within 100 ns
Grandmaster Clock Offset Scaled Log Variance : 0
Grandmaster Priority1             : 128
Grandmaster Priority2             : 128
=====
E4G-200.1 # show network-clock ptp boundary port
Mode          : Boundary Clock
State         : Enabled
Instance      : 0
=====
Port Data Set
=====
VLAN Name           : ptp-s1
Port Clock Identity : 00:04:96:ff:fe:52:d2:45
Port Number         : 1
Port State          : Slave
Log Min Delay Request Interval : -6 (1/64 sec)
Log Announce Interval : 1 (2 sec)
Log Announce Receipt Timeout  : 3 (8 sec)
Log Sync Interval   : -6 (1/64 sec)
Delay Mechanism     : End-to-End
Version Number      : 2
=====
E4G-200.1 # show network-clock ptp boundary time-property
Mode          : Boundary Clock
State         : Enabled
Instance      : 0
=====
Time Properties Data Set
=====
Current UTC Offset      : 34
Current UTC Offset Valid : TRUE
Leap 59                 : FALSE
Leap 61                 : FALSE
Time Traceable          : TRUE
Frequency Traceable     : TRUE
PTP Timescale           : TRUE
Time Source             : GPS
Calendar Time           : Mon Mar 12 12:58:47 2012 +0000
=====
E4G-200.1 # show network-clock ptp boundary
Mode          : Boundary Clock
State         : Enabled
Instance      : 0
=====
Default Data Set
=====
Two Step Flag          : FALSE
Clock Identity         : 00:04:96:ff:fe:52:d2:45
Number Of Ports        : 32
Clock Class            : 255
Clock Accuracy         : Within 25 us

```



```

Offset Scaled Log Variance      : 1
Priority1                       : 128
Priority2                       : 128
Domain Number                   : 0
Slave Only Flag                 : FALSE
=====
Current Data Set
=====
Steps Removed                   : 2
Offset From Master              : 0.000 (nsec)
Mean Path Delay                 : 296.185 (nsec)
=====
Parent Data Set
=====
Parent Port Clock Identity      : 00:04:96:ff:fe:52:d2:42
Parent Port Number              : 2
Parent Active Peer Addr        : 22.55.0.2
Parent Stats                    : FALSE
Observed Parent Offset Scaled Log Variance : 0
Observed Parent Clock Phase Change Rate : 0
Grandmaster Identity           : 00:11:11:11:11:11:11:11
Grandmaster Clock Class        : 6
Grandmaster Clock Accuracy     : Within 100 ns
Grandmaster Clock Offset Scaled Log Variance : 0
Grandmaster Priority1          : 128
Grandmaster Priority2          : 128
=====
Time Properties Data Set
=====
Current UTC Offset              : 34
Current UTC Offset Valid       : TRUE
Leap 59                        : FALSE
Leap 61                        : FALSE
Time Traceable                 : TRUE
Frequency Traceable            : TRUE
PTP Timescale                  : TRUE
Time Source                     : GPS
                                Calendar Time           : Mon Mar 12 12:58:47 2012 +0000
=====

```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400) and Summit X770.



Note

PTP commands can be used only with the Network Timing feature pack.



show network-clock ptp (interface)

```
show network-clock ptp [boundary | ordinary] [{vlan} vlan_name | vlan all]
```

Description

Show interface information for the given clock port that is added to the specified clock instance. This command is available only for boundary and ordinary clocks.

Syntax Description

boundary	Specifies boundary clock.
ordinary	Specifies ordinary clock.
vlan	VLAN.
vlan_name	VLAN name.
vlan all	All VLANs.

Default

N/A.

Usage Guidelines

Use this command to display all the clock ports that have PTP enabled and added to the specified clock instance.

Example

The following command displays PTP clock port interface information for all VLANs:

```
show network-clock ptp boundary vlan all
VLAN Name           : gm1-lpbk
Acceptable Master Option : Disabled
Unicast Negotiation  : Disabled
Port Type           : Slave (forced)
Log Announce Interval : 1
Log Announce Timeout : 3
Log Delay Request Interval : -6
Log Sync Interval    : -6
VLAN Name           : ptp-s1
Acceptable Master Option : Disabled
Unicast Negotiation  : Disabled
Port Type           : Master or Slave
Log Announce Interval : 1
Log Announce Timeout : 3
Log Delay Request Interval : -6
Log Sync Interval    : -6
VLAN Name           : ptp-m1
```



```

Acceptable Master Option      : Disabled
Unicast Negotiation          : Disabled
Port Type                     : Master (forced)
Log Announce Interval        : 1
Log Announce Timeout         : 3
Log Delay Request Interval    : -6
Log Sync Interval            : -6

```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400) and Summit X770.



Note

PTP commands can be used only with the Network Timing feature pack.

show network-clock ptp end-to-end-transparent ports

```
show network-clock ptp end-to-end-transparent ports port_list {detail}
```

Description

Show the details of end-to-end-transparent PTP instance configured on the specified ports. Additionally, the port information detail output includes the PTP configuration, and the timestamping mode on the port.

Syntax Description

<i>port_list</i>	List of physical ports.
detail	Display detailed information.

Default

N/A.

Usage Guidelines

See Description.

Example

```
E4G-200.1 # show network-clock ptp end-to-end-transparent ports 1,3,5,12
```



```

=====
==
Port      Flags      BC / OC VLAN
(or # VLANs)
=====
==
1          SD-----
3          SEeO-----
5          SEeO----- 2
12         N-----
-----
--
> indicates Port or VLAN Display Name truncated past 8 characters
Legend: BC - Boundary Clock, OC - Ordinary Clock
Flags : (D) PTP Disabled, (e) PTP End-to-End Transparent Clock,
(E) PTP Enabled, (N) PTP Not Supported, (O) One-step time-stamping,
(P) PTP Peer-to-Peer Transparent Clock, (S) PTP Supported,
(T) Two-step time-stamping
E4G-200.1 # show port 3 information detail
Port: 3
Virtual-router: VR-Default
Type: SF+_SR Unsupported Optic Module
Random Early drop: Unsupported
Admin state: Enabled with 10G full-duplex
Link State: Active, 10Gbps, full-duplex
Link Ups: 1 Last: Wed Feb 06 12:28:48 2013
Link Downs: 0 Last: --

VLAN cfg:
STP cfg:

Protocol:
Trunking: Load sharing is not enabled.

EDP: Enabled
ELSM: Disabled
Ethernet OAM: Disabled
Learning: Enabled
Unicast Flooding: Enabled
Multicast Flooding: Enabled
Broadcast Flooding: Enabled
Jumbo: Disabled

Flow Control: Rx-Pause: Enabled Tx-Pause: Disabled

Priority Flow Control: Disabled
Reflective Relay: Disabled
Link up/down SNMP trap filter setting: Enabled
Egress Port Rate: No-limit
Broadcast Rate: No-limit
Multicast Rate: No-limit
Unknown Dest Mac Rate: No-limit
QoS Profile: None configured
Ingress Rate Shaping : Unsupported
Ingress IPTOS Examination: Disabled
Ingress 802.1p Examination: Enabled
Ingress 802.1p Inner Exam: Disabled
Egress IPTOS Replacement: Disabled

```




```

Egress 802.1p Replacement: Disabled
NetLogin: Disabled
NetLogin port mode: Port based VLANs
Smart redundancy: Enabled
Software redundant port: Disabled
IPFIX: Disabled Metering: Ingress, All Packets, All
Traffic
IPv4 Flow Key Mask: SIP: 255.255.255.255 DIP:
255.255.255.255
IPv6 Flow Key Mask: SIP:
ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff
DIP:
ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff
Far-End-Fault-Indication: Disabled
Shared packet buffer: default
VMAN CEP egress filtering: Disabled
Isolation: Off
PTP Configured: Disabled
Time-Stamping Mode: None
Synchronous Ethernet: Unsupported
Dynamic VLAN Uplink: Disabled
VM Tracking Dynamic VLANs: Disabled
    
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400) and Summit X770.



Note

PTP commands can be used only with the Network Timing feature pack.

show network-clock ptp boundary unicast-master

```

show network-clock ptp [boundary | ordinary] unicast-master [{vlan} vlan_name |
vlan all]
    
```

Description

Show the unicast master table of the specified clock port or all clock ports. This command is available only for boundary and ordinary clocks.

Syntax Description

boundary	Specifies the boundary clock.
ordinary	Specifies the ordinary clock.



unicast-master	Specifies table of unicast masters.
vlan	VLAN.
<i>vlan_name</i>	VLAN name.
vlan all	Specifies all vlans.

Default

N/A.

Usage Guidelines

Use this command to display the unicast master table of the specified clock port or all clock ports. This command is available only for boundary and ordinary clocks.

Example

The following command displays the unicast master table for all clock ports:

```
E4G-200.1 # show network-clock ptp boundary unicast-master vlan all
VLAN Name                IP Address          Log Query
gm1-lpbk                  1.1.1.101           3
ptp-s1                     22.55.0.2           3
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400) and Summit X770.



Note

PTP commands can be used only with the Network Timing feature pack.

show network-clock ptp boundary unicast-slave

```
show network-clock ptp boundary unicast-slave [{vlan} vlan_name | vlan all]
```

Description

Show the unicast slave table of the specified clock port, or all clock ports. This command is only available for boundary clocks.



Syntax Description

boundary	Specifies the boundary clock.
unicast-slave	Specifies table of unicast masters.
vlan	VLAN.
<i>vlan_name</i>	VLAN name.
vlan all	Specifies all vlans.

Default

N/A.

Usage Guidelines

Use this command to display the unicast master table of the specified clock port or all clock ports. This command is only available for boundary clocks.

Example

The following command displays the unicast master table for all clock ports:

```
E4G-200.1 # show network-clock ptp boundary unicast-slave vlan all
VLAN Name                IP Address
ptp-m1                    7.1.1.1
ptp-m1                    7.1.2.1
ptp-m1                    88.3.5.1
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400) and Summit X770.



Note

PTP commands can be used only with the Network Timing feature pack.

show network-clock ptp counters

```
show network-clock ptp [boundary | ordinary] vlan [vlan_name {{ipv4_address}}
[unicast-master | unicast-slave]} | all {unicast-master | unicast-slave}]
counters
```



Description

This command displays the count of PTP packets that include event messages, signaling and management messages transmitted and received on each unicast-slave/unicast-master peers. Use the following queries for the counters:

- Per unicast-master or unicast-slave peer on a clock port.
- All unicast-master or unicast-slave peers on a clock port.
- Peers on a clock port.
- unicast-master or unicast-slave peers on all clock ports.
- Peers on all clock ports.

The refresh option is not supported for the counters. This command is available only for Boundary and Ordinary clocks.

Syntax Description

network-clock	External clock for ethernet synchronization.
ptp	Precise time protocol.
boundary	Boundary clock.
ordinary	Ordinary clock.
vlan	VLAN.
<i>ipv4_address</i>	Peer IP address.
unicast-master	IP addresses that are masters to the local clock.
unicast-slave	IP addresses that are slaves to the local clock.
all	All VLANS.
counters	PTP message counts.

Default

N/A.

Usage Guidelines

Use this command to display the count of PTP packets that include event messages, signaling and management messages transmitted and received on each unicast-slave/unicast-master peers.

Example

```
E4G-400.10 # show network-clock ptp boundary vlan lpbk-gm 1.1.1.101 unicast-
master counters
VLAN name:          lpbk-gm
Peer IP Address:    1.1.1.101
IN (packets)        OUT (packets)  Message Type
=====
211                  0  Announce
```



```

27041          0 Sync
0              0 FollowUp
0             26532 DelayReq
26489          0 DelayResp
0              0 Management
141            150 Signaling
0              Rejected
E4G-400.10 # show network-clock ptp boundary vlan all unicast-slave counters
VLAN name:      lpbk-slave
Peer IP Address: 15.1.1.1
IN (packets)    OUT (packets)  Message Type
=====
0               258  Announce
0               32958 Sync
0               32958 FollowUp
0               0 DelayReq
0               0 DelayResp
0               0 Management
0               0 Signaling
0               Rejected
VLAN name:      lpbk-slave
Peer IP Address: 15.1.1.2
IN (packets)    OUT (packets)  Message Type
=====
0               258  Announce
0               32960 Sync
0               32960 FollowUp
31569           0 DelayReq
0               31569 DelayResp
0               0 Management
166             0 Signaling
0               Rejected

```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400) and Summit X770.



Note

PTP commands can be used only with the Network Timing feature pack.

show node

```
show node {detail}
```

Description

Displays the status of the nodes in the system as well as the general health of the system.



Syntax Description

detail	Displays the information on a per-node basis rather than in a tabular format.
---------------	---

Default

N/A.

Usage Guidelines

Use this command to display the current status of the nodes and the health of the system. The information displayed shows the node configurations (such as node priority) and the system and hardware health computations. You can use this information to determine which node will be elected primary in case of a failover.

The following table lists the node statistic information collected by the switch.

Table 11: Node States

Node State	Description
BACKUP	In the backup state, this node becomes the primary node if the primary fails or enters the DOWN state. The backup node also receives the checkpoint state data from the primary.
DOWN	In the down state, the node is not available to participate in leader election. The node enters this state during any user action, other than a failure, that makes the node unavailable for management. Examples of user actions are: Upgrading the software Rebooting the system using the <code>reboot</code> command. Initiating an MSM/MM failover using the <code>run msm-failover</code> command. Synchronizing the MSM's/MM's software and configuration in non-volatile storage using the <code>synchronize</code> command
FAIL	In the fail state, the node has failed and needs to be restarted or repaired. The node reaches this state if the system has a hardware or software failure.
INIT	In the initial state, the node is being initialized. A node stays in this state when it is coming up and remains in this state until it has been fully initialized. Being fully initialized means that all of the hardware has been initialized correctly and there are no diagnostic faults.
MASTER	In the primary state, the node is responsible for all switch management functions.
STANDBY	In the standby state, leader election occurs—the primary and backup nodes are elected. The priority of the node is only significant in the standby state.

Example

The following command displays the status of the node, the priority of the node, and the general health of the system:

```
show node
```

The following is sample output from this command:

```
Node      State      Priority    SwHealth    HwHealth
```



```

-----
MSM-A  MASTER          0      49      7
MSM-B  BACKUP         0 49      7

```

If you specify the detail option, the same information is displayed on a per node basis rather than in a tabular format, as shown in the following example:

```

Node MSM-A information:
Node State:    MASTER
Node Priority: 0
Sw Health:    49
Hw Health:    7
Node MSM-B information:
Node State:    BACKUP
Node Priority: 0
Sw Health:    49
Hw Health:    7

```

History

This command was first available in an ExtremeXOS 10.1.

Platform Availability

This command is available only on modular switches and SummitStack.

show ntp

show ntp

Description

Displays the global NTP status of the switch.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

N/A.



Example

The following command shows the NTP status of the switch:

```
Switch# show ntp
NTP                : Enabled
Authentication     : Disabled
Broadcast-Client   : Disabled
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on all platforms.

show ntp association

```
show ntp association [{ip_address} | {host_name}]
```

Description

Shows all of the NTP clock source information, from a statically configured server, peer, or broadcast server. The NTP service updates the local clock from only one NTP server, with the best stability and stratum value which is considered as a system peer.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

N/A.

Example

```
The following command shows detailed information about the NTP server:
switch # show ntp association 1.us.pool.ntp.org
Remote IP          : 72.18.205.156   Local IP           : 10.45.203.74
Host Mode          : Client           Peer Mode          : Server
Version            : 3                Key ID             : 0
Stratum            : 3                Precision          : -18
```




```

Root Distance      : 0.05460          Root Dispersion   : 0.06429
Reachability      : 377              UnReachability    : 0
Peer Poll         : 6                Host Pool         : 6
Broadcast Offset  : 0.01985          TTL/Mode         : 0
Offset           : 0.012219         Delay            : 0.01985
Error Bound      : 0.07240          Filter Error     : 0.06711
Peer Flags       : Config, Broadcast Client
Reference Time   : d140506c.4ba4702e Fri, Apr 1 2011 6:23:56.295
Originate Time  : 00000000.00000000 Wed, Feb 6 2036 22:28:16.000
Receive Timestamp : d1405468.44694b54 Fri, Apr 1 2011 6:40:56.267
Transmit Timestamp : d1405468.44694b54 Fri, Apr 1 2011 6:40:56.267
Filter Order    :      0      1      2      3      4
5      6      7
Filter Delay    : 0.02116 0.02234 0.02104 0.02141 0.04269 0.01985
0.02100 0.02029
Filter Offset   : 0.009781 0.011904 0.011966 0.011380 0.000619 0.012219
0.012885 0.012667

```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on all platforms.

show ntp association statistics

```
show ntp association [{ip_address} | {host_name}] statistics
```

Description

Shows NTP-related statistics about a specific NTP server.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

N/A.



Example

The following command shows NTP-related statistics about the NTP server called “1.us.pool.ntp.org”:

```
Switch# show ntp association 1.us.pool.ntp.org statistics
Remote Host      : 1.us.pool.ntp.org
Local Interface  : 10.45.203.74
Time Last Received : 40 second
Time Until Next Send: 27 second
Reachability Change : 849 second
Packets Sent     : 18
Packets Received  : 18
Bad Authentication : 0
Bogus Origin     : 0
Duplicate        : 0
Bad Dispersion   : 0
Bad Reference Time : 0
Candidate Order  : 4
Peer Flags       : Config, Broadcast Client, Initial Burst
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on all platforms.

show ntp key

show ntp key

Description

Shows the NTP key index number, trusted or non-trusted, authentication type, and encrypted key string.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

N/A.



Example

The following command shows the NTP key index number, trusted or non-trusted, authentication type, and encrypted key string:

```
* X450a-24t.1 # show ntp key
Key Index      Trusted      Auth      Key String (encrypted)
=====
100            No           MD5       67:74:7d:78:6f:6c:67:5b:33
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on all platforms.

show ntp restrict-list

```
show ntp restrict-list {user | system | all}
```

Description

Show the NTP access list of the current system based on the source IP blocks.

Syntax Description

user	Displays the NTP access list of the current user.
system	Displays the for the current system.
all	Displays both user and system data.

Default

Displays all by default.

Usage Guidelines

N/A.

Example

The following command displays all NTP access list information:

```
switch # show ntp restrict-list all
IP Address      Mask      Count  Type  Action
```



```

=====
0.0.0.0          0.0.0.0          0 System Deny
10.1.1.0         255.255.255.0   0 System Permit
10.1.1.1         255.255.255.255 0 System Permit
10.45.200.0      255.255.252.0   0 System Permit
10.45.203.74     255.255.255.255 0 System Permit
69.65.40.29      255.255.255.255 37 System Permit
110.1.1.0        255.255.255.0   0 User Permit
127.0.0.1        255.255.255.255 0 System Permit
127.127.1.1      255.255.255.255 0 System Permit
173.9.142.98     255.255.255.255 37 System Permit
173.203.122.111  255.255.255.255 36 System Permit
216.93.242.12    255.255.255.255 35 System Permit

```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on all platforms.

show ntp server

show ntp server

Description

Shows the NTP servers configured on the switch, including the name, IP address, key ID, and index.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

N/A.

Example

The following command shows the NTP servers configured on the switch:

```

* switch # show ntp server
Name                IP Address          Type      Flags  Key Index

```



```

=====
0.us.pool.ntp.org      108.69.104.139      Server    I    -
Flags                  : (I) Initial Burst, (B) Burst

```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on all platforms.

show ntp sys-info

show ntp sys-info

Description

Shows the current system status based on the most reliable clock server or NTP server.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

N/A.

Example

The following command shows the current system status based on the most reliable clock server or NTP server:

```

switch # show ntp sys-info
System Peer      : 0.us.pool.ntp.org
System Peer Mode : Client
Leap Indicator   : 00
Stratum          : 3
Precision        : -20
Root Distance    : 0.09084 second
Root Dispersion  : 0.23717 second
Reference ID     : [216.93.242.12]
Reference time   : d140571d.e8389ff7  Fri, Apr  1 2011  6:52:29.907
System Flags     : Monitor, Ntp, Kernel, Stats

```



```

Jitter           : 0.004700 second
Stability        : 0.000 ppm
Broadcast Delay  : 0.007996 second
Auth Delay       : 0.000000 second

```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on all platforms.

show ntp vlan

show ntp vlan

Description

Shows the NTP status of each VLAN configured on the switch.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

N/A.

Example

The following command shows the NTP status of each VLAN configured on the switch:

```

Switch# show ntp vlan
Vlan           NTP Status  Broadcast Server  Key Index
=====
internet       Enabled     Disabled          -
test           Enabled     Disabled          -

```

History

This command was first available in ExtremeXOS 12.7.



Platform Availability

This command is available on all platforms.

show odometers

show odometers

Description

Displays a counter for each component of a switch that shows how long it has been functioning since it was manufactured.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The output from this command displays how long individual components in the switch have been functioning since it was manufactured. This odometer counter is kept in the EEPROM of each monitored component. On a modular switch, this means that even if you plug in the component into a different chassis, the odometer counter is available in the new switch chassis.

Monitored Components

On a modular switch, the odometer monitors the following components:

- Chassis.
- MSMs/MMs.
- I/O modules.
- Power controllers.

On the Summit family switches (whether or not included in a SummitStack), the odometer monitors the following components:

- Switch.
- XGM-2xn card.

Recorded Statistics

The following odometer statistics are collected by the switch:

- Service Days—The amount of days that the component has been running.
- First Recorded Start Date—The date that the component was powered-up and began running.



Depending on the software version running on your switch, the modules installed in your switch, and the type of switch you have, additional or different odometer information may be displayed.

Example

The following command displays how long each component of a switch has been functioning since its manufacture date:

```
show odometers
```

The following is sample output from the BlackDiamond X8 series switch:

```
BD-X8.4 # show odometers
Service First Recorded
Field Replaceable Units          Days      Start Date
-----
Chassis      : BD-X8                41      Oct-28-2011
Slot-1      : BDXA-10G48X            31      Oct-12-2011
Slot-2      : BDXA-10G48X            13      Nov-16-2011
Slot-3      :
Slot-4      :
Slot-5      :
Slot-6      : BDXA-40G24X           27      Nov-22-2011
Slot-7      : BDXA-40G24X           27      Nov-22-2011
Slot-8      : BDXA-40G24X           27      Nov-22-2011
FM-1       : BDXA-FM20T              33      Nov-15-2011
FM-2       : BDXA-FM20T              33      Nov-15-2011
FM-3       : BDXA-FM20T              29      Nov-15-2011
FM-4       : BDXA-FM20T              16      Nov-15-2011
MM-A       : BDX-MM1                30      Sep-08-2011
MM-B       : BDX-MM1                52      Sep-09-2011
```

The following is sample output from the BlackDiamond 8800 series switch:

```
Field Replaceable Units          Days      Start Date
-----
Chassis      : BD-8810                1848     Sep-21-2004
Slot-1      : G48P                    1819     Sep-27-2004
Slot-2      : G24X                      343      Sep-21-2004
Slot-3      : G48T                    1818     Sep-27-2004
Slot-4      :
Slot-5      : G8X                      1632     Sep-21-2004
Slot-6      : G8X                      1231     Sep-28-2004
Slot-7      :
Slot-8      : G48Te                    1435     Feb-06-2006
Slot-9      : G48Ta                    1484     Apr-13-2006
Slot-10     :
MSM-A       : MSM-48C                  1624     Oct-21-2004
MSM-B       : MSM-48C                  1218     Nov-30-2004
PSUCTRL-1  :                          1809     Nov-30-2004
PSUCTRL-2  :                          1821     Nov-29-2004
```



The following is sample output from a stand-alone Summit series switch:

```

Service  First Recorded
Field Replaceable Units          Days      Start Date
-----
Switch      : SummitX              7  Dec-08-2004
XGM-2xn-1   :

Service  First
Recorded
Field Replaceable Units          Days      Start Date
-----
Switch      : X(SS)              381  Oct-29-2009
VIM1-SS-1   :                    376  Jul-30-2009

```

History

This command was first available in ExtremeXOS 10.1.

Information about the power controller(s) for modular switches was added to the `show odometers` output in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show power

```
show power {ps_num} {detail}
```

Description

Displays the current status of the installed power supplies.

Syntax Description

<i>ps_num</i>	Specifies the slot number of the installed power supply.
detail	The detail option is reserved for future use.

Default

N/A.

Usage Guidelines

Use this command to view detailed information about the health of the power supplies.

This status information may be useful for your technical support representative if you have a network problem.



The switch collects the following power supply information:

- **State**—Indicates the current state of the power supply. Options are:
 - **Empty**—There is no power supply installed.
 - **Power Failed**—The power supply has failed.
 - **Powered Off**—The power supply is off.
 - **Powered On**—The power supply is on and working normally.

Modular switches only:

Located next to the “State” of the power supply, the following information provides more detailed status information. Options are:

- **Disabled for net power gain**—Indicates that the power supply is disabled in order to maximize the total available system power.
- **Configured ON**—Indicates that the user requested to enable a disabled power supply regardless of the affect on the total available system power.
- **Configured ON when present**—Indicates that the power supply slot is currently empty, but the user requested to enable the power supply regardless of the affect on the total available system power.
- **Unsupported**—Indicates that a 600/900 W AC PSU is inserted in a chassis other than the BlackDiamond 8806.
- **PartInfo**—Provides information about the power supply. Depending on your switch, options include:

Modular switches only:

- **Serial number**—A collection of numbers and letters, that make up the serial number of the power supply.
- **Part number**—A collection of numbers and letters that make up the part number of the power supply.

Summit family switches only:

- **Internal Power Supply (PowerSupply 1 information)**—The Summit family switches come with one power supply pre-installed at the factory. The Summit power supply is not user-replaceable; therefore, the part information display indicates internal power supply.
- **External Power Supply (PowerSupply 2-4 information)**—Displays information about the optional External Power System (EPS) that allows you to add a redundant power supply to the Summit family switches to protect against a power supply failure.

On modular switches, the output also includes the following information:

- **Revision**—Displays the revision number of the power supply.
- **Odometer**—Specifies how long the power supply has been operating.
- **Temperature**—Specifies, in Celsius, the current temperature of the power supply.
- **Input**—Specifies the input voltage and the current requirements of the power supply and whether the input is AC or DC.
- **Output 1 and Output 2**—Specifies the output voltage and the current supplied by the power supply. The values are only displayed if known for the platform.

The Summit X460, E4G-400, BlackDiamond X8 series, and the BlackDiamond 8800 series switches include Power Usage, which is only an estimate for the input power consumed. SummitStack displays the supplies associated to each active node that is present in a slot. The supplies are represented with



flags that describe whether the supply is providing power, has failed, or is providing no power, or if the supply has had its 48v power output automatically turned off because two or three external power supplies are available. For more information, see the [show power \(Stack Nodes Only\)](#) section in [SummitStack Feature Commands](#)

In ExtremeXOS 10.1 and earlier, use the `show powersupplies {detail}` command to view detailed health information about the power supplies.

Example

Modular switch example:

The following command displays the status of the power supply installed in slot 1 in a modular switch:

```
show power 1
```

The following is sample output from this command:

```
PowerSupply 1 information:
State:           Powered On
PartInfo:        PS 2336 5003J-00479 4300-00137
Revision:        2.0
Odometer:        90 days 5 hours
Temperature:     29.0 deg C
Fan 1:           6473 RPM
Fan 2:           6233 RPM
Input:           230.00 V AC
Output 1:        48.50 V, 7.25 A (48V/1104W Max)
Output 2:        12.44 V, 0.62 A (12V/48W Max)
```

If power management needs to disable a power supply to maximize the total available power, you see Disabled for net power gain next to the state of the power supply, as shown in the sample truncated output:

```
PowerSupply 1 information:
State:           Powered Off (Disabled for net power gain)
PartInfo:        PS 2336 0413J-00732 4300-00137
...
```

If you choose to always enable a power supply, regardless of the affect on the total available power, you see Configured ON next to the state of the power supply, as shown in the sample truncated output:

```
PowerSupply 1 information:
State:           Powered On (Configured ON)
PartInfo:        PS 2336 0413J-00732 4300-00137
```



If you install the 600/900 W AC PSU in a chassis other than a BlackDiamond 8806, you see unsupported next to the state of the power supply, as shown in this sample truncated output:

```
PowerSupply 3 information:
State:           Unsupported
PartInfo:        PS 2431 0622J-00013 4300-00161
```

Summit family switches example:

The following command displays the status of the power supplies installed in the Summit X440 series switch:

```
show power
```

The following sample output assumes that you have not installed an EPS:

```
PowerSupply 1 information:
State:           Powered On
PartInfo:        Internal Power Supply
PowerSupply 2 information:
State:           Empty
```

The following sample output assumes that you have installed an EPS:

```
PowerSupply 1 information:
State:           Powered On
PartInfo:        Internal Power Supply
PowerSupply 2 information:
State:           Powered On
PartInfo:        External Power Supply
```

The following sample output for the Summit X440 assumes a connection to an external EPS-C with three EPS-600LS modules installed:

```
show power
PowerSupply 1 information:
State:           Powered Off
PartInfo:        Internal Power Supply
Input:           0.00 V AC
PowerSupply 2 information:
State:           Powered On
PartInfo:        External Power Supply
Input:           0.00 V AC
PowerSupply 3 information:
State:           Powered On
PartInfo:        External Power Supply
Input:           0.00 V AC
PowerSupply 4 information:
State:           Powered On
```



```
PartInfo:      External Power Supply
Input:         0.00 V AC
```

For the BlackDiamond X8 switch, the command output is enhanced to show the power supply groupings as follows:

```
BD-X8.2 # show power
PowerSupply 1 information:
State           : Powered On
PartInfo        : H2500A2-EX 1109X-88834 800429-00
Revision        : 1.0
Odometer        : 200 days 6 hours
Temperature     : 29.0 deg C
Fan            : 8616 RPM
Input           : 220.00 V AC
Output 1        : 48.00 V, 9.10 A (48V/2160W Max)
Power Usage     : 526.26 W
PowerSupply 2 information:
State           : Powered On
PartInfo        : H2500A2-EX 1109X-88787 800429-00
Revision        : 1.0
Odometer        : 154 days 18 hours
Temperature     : 30.0 deg C
Fan            : 8616 RPM
Input           : 220.00 V AC
Output 1        : 47.99 V, 9.47 A (48V/2160W Max)
Power Usage     : 547.54 W
PowerSupply 3 information:
State           : Empty
PowerSupply 4 information:
State           : Empty
PowerSupply 5 information:
State           : Powered On
PartInfo        : H2500A2-EX 1121X-88454 800429-00
Revision        : 1.0
Odometer        : 85 days 8 hours
Temperature     : 46.0 deg C
Fan            : 8616 RPM
Input           : 220.00 V AC
Output 1        : 47.93 V, 9.22 A (48V/2160W Max)
Power Usage     : 532.42 W
PowerSupply 6 information:
State           : Powered On
PartInfo        : H2500A2-EX 1121X-88598 800429-00
Revision        : 1.0
Odometer        : 72 days 20 hours
Temperature     : 47.0 deg C
Fan            : 8616 RPM
Input           : 220.00 V AC
Output 1        : 47.99 V, 8.48 A (48V/2160W Max)
Power Usage     : 490.30 W
PowerSupply 7 information:
State           : Powered On
PartInfo        : H2500A2-EX 1109X-88765 800429-00
Revision        : 1.0
Odometer        : 135 days 18 hours
Temperature     : 46.0 deg C
```



```

Fan                : 8616 RPM
Input              : 220.00 V AC
Output 1           : 47.97 V, 9.06 A (48V/2160W Max)
Power Usage       : 523.62 W
PowerSupply 8 information:
State              : Empty
System Power Usage : 2620.14 W
Poll Interval     : 60 s
Change Threshold  : N/A
Change Action     : N/A
Note: Input Power Sources should be split among
PSU Groups A (1-4) and B (5-8) for best redundancy.

```

For stacking systems, the power detail display is enhanced as follows:

```

* Slot-1 Stack.2 # show power
PSU-1 or PSU-2 or
Internal External External External Power
Slots Type PSU PSU PSU PSU Usage
-----
Slot-1 X460-48p P - - - 113.88 W
Slot-2
Slot-3
Slot-4
Slot-5
Slot-6
Slot-7
Slot-8
Flags : (P) Power available, (F) Failed or no power,
(O) 48V powered off when 2 or 3 external PSUs are powered on,
(-) Empty
System Power Usage : 120 W
* Slot-1 Stack.2 # show power detail
Slot-1 PowerSupply 1 information:
State : Powered On
PartInfo : PSSF751301A- 1022A-40459 800382-00-01
Power Usage : 113.88W
Output 1 : 18.63 V, 4.37 A
Output 2 : 3 V, 3 A
Slot-1 PowerSupply 2 information:
State : Empty
System Power Usage : 120 W

```

History

This command was first available in an ExtremeXOS 10.1.

The syntax for this command was modified in ExtremeXOS 11.0 from `show powersupplies` to `show power {ps_num} {detail}`.

The output was modified to include power management details for modular switches in ExtremeXOS 11.3.

DC power output for the BlackDiamond 8800 series switches was added in ExtremeXOS 11.4.



For the BlackDiamond X8 switch, the command output was enhanced to show the power supply groupings in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms.

show power budget

show power budget

Description

Displays the power status and the amount of available and required power on a modular switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to view detailed information about the amount of power available on the switch.

This status information may be useful if the `show slot` command displays a state of Powered OFF for any I/O module, for monitoring power, or for power planning purposes.

The first table of the `show power budget` command displays:

- Slot number of the power supply.
- Current state of the power supply. Options are:
 - Empty—There is no power supply installed.
 - Power Failed—The power supply has failed.
 - Power Off—The power supply is off.
 - Power On—The power supply is on.
- Watts and voltage amounts of the power supply.
- Redundant power information. Redundant power is the amount of power available if power to one PSU is lost. If a switch has PSUs with a mix of both 220V AC and 110V AC inputs, the amount of redundant power shown is based on the worst-case assumption that power to a PSU with 220V AC input is lost.

The second table of the `show power budget` command displays:

- Slot number and name of the component installed in the slot. Options include:
 - I/O modules.



- MSMs/MMs.
- Fan trays.
- Current state of the module. Options include, among others:
 - Empty: There is no component installed.
 - Operational: The component is installed and operational.
 - Present: The component is installed but not operational.
 - Down: The module is installed, but the administrator has taken the module offline.
 - Power ON: There is sufficient system power to power up the module.
 - Powered OFF: There is insufficient system power to keep the module up and running, or there is a mismatch between the module configured for the slot and the actual module installed in the slot.
 - Booting: The module has completed downloading the software image and is now booting.
 - Initializing: The module is initializing.
- Watts and voltage amounts of the modules.
- Power Surplus or Power Shortfall.
 - If the amount of available power meets or exceeds the required port, the excess is displayed as the Power Surplus.
 - If the available power is insufficient to meet the required power, the deficit is displayed as Power Shortfall.
- Redundant power information. If the amount of redundant power meets or exceeds the required power, the system has (N+1) power.
 - Yes—The system has redundant (N+1) power.
 - No—The system does not have redundant (N+1) power.

The information contained in this display is for planning purposes since the system operates without redundant power as long as a power surplus is shown. However, if power is lost to a single PSU when the system is not redundant, I/O modules are powered down. Please refer to the section [Understanding Power Supply Management in Managing the Switch](#) in the *ExtremeXOS Concepts Guide*.

Depending on the software version running on your switch, the modules installed in your switch, and the type of switch you have, additional or different power information may be displayed.

Example

The following command displays the distribution of power and the available power on the switch:

```
show power budget
```

This display will appear on the BlackDiamond X8 switch as follows:

```
BD-X8.36 # show power budget
Watts
PS State                               at 48V
-----
1   Powered On                          2500.00
```




```

2    Powered On                2500.00
3    Powered On                2500.00
4    Powered On                2500.00
5    Powered On                2500.00
6    Powered On                2500.00
7    Powered On                2500.00
8    Powered On                2500.00
-----
Power Available:                20000.00
N+1 Redundant Power Available:  17500.00
N+N Redundant Power Available:  10000.00
Note: Input Power Sources should be split among
PS groups A (1-4) and B (5-8) for best redundancy.
Watts
Slots      Type                State      at 48V
-----
Slot-1    BDXA-10G48X             Operational  439.00
Slot-2    BDXA-10G48X             Operational  439.00
Slot-3    BDXA-10G48X             Operational  439.00
Slot-4    BDXA-10G48X             Operational  439.00
Slot-5    BDXA-10G48X             Operational  439.00
Slot-6    BDXA-10G48X             Operational  439.00
Slot-7    BDXA-10G48X             Operational  439.00
Slot-8    BDXA-40G24X             Operational  759.00
FM-1     BDXA-FM20T              Operational  479.00
FM-2     BDXA-FM20T              Operational  479.00
FM-3     BDXA-FM20T              Operational  479.00
FM-4     BDXA-FM20T              Operational  479.00
MM-A     BDXA-MM1                 Operational  199.00
MM-B     BDXA-MM1                 Operational  199.00
FanTray-1                Operational  249.00
FanTray-2                Operational  249.00
FanTray-3                Operational  249.00
FanTray-4                Operational  249.00
FanTray-5                Operational  249.00
-----
Power Required:                7391.00
Power Allocated:                7391.00
Power Surplus:                  12609.00
N+1 Redundant Power Supply(s) Present?:    Yes
N+N Redundant Power Supply(s) Present?:    Yes
BD-X8.37 #

```

History

This command was first available in ExtremeXOS 11.0.

Power over Ethernet (PoE) data (inline power) was added to the `show power budget` output in ExtremeXOS 11.1. PoE data is displayed when you install a G48P module in a BlackDiamond 8800 series switch.

Redundant (N+1) power information was added to the `show power budget` output in ExtremeXOS11.1.



Platform Availability

This command is available only on modular switches.

show power controller

```
show power controller {num}
```

Description

Displays the current status of the installed power supply controllers.

Syntax Description

<i>num</i>	Specifies the slot number of the installed power supply controller.
------------	---

Default

N/A.

Usage Guidelines

Use this command to view detailed information about the health of the power supply controllers. Power controllers collect data about the installed power supplies and report the results to the MSM/MM.

This status information may be useful for your technical support representative if you have a network problem.

The switch collects the following power supply controller information:

- State—Indicates the current state of the power supply controller. Options are:
 - Empty: There is no power supply controller installed.
 - Operational: The power supply controller is installed and operational.
 - Present: The power supply controller is installed.
- PartInfo—Provides information about the power supply controller including the:
 - Slot number where the power supply controller is installed.
 - Serial number, a collection of numbers and letters, that make up the serial number of the power supply controller.
 - Part number, a collection of numbers and letters that make up the part number of the power supply controller.
- Revision—Displays the revision number of the power supply controller.
- FailureCode—Specifies the failure code of the power supply controller.
- Odometer—Specifies the date and how long the power supply controller has been operating.
- Temperature—Specifies, in Celsius, the current temperature of the power supply controller.
- Status—Specifies the status of the power supply controller.



Example

The following command displays the status of the installed power supply controllers:

```
show power controller
```

The following is sample output from this command:

```
PSUCTRL-1 information:
State:          Operational
PartInfo:      PSUCTRL-1 04334-00021 450117-00-01
Revision:      1.0
FailureCode:   0
Odometer:      337 days 7 hours since Nov-30-2004
Temperature:   32.14 deg C
Status:        PSU CTRL Mode:   Master
PSUCTRL-2 information:
State:          Empty
```

If you have two power supply controllers installed, the switch displays output about both of the power supply controllers:

```
PSUCTRL-1 information:
State:          Operational
PartInfo:      PSUCTRL-1 04334-00021 450117-00-01
Revision:      1.0
FailureCode:   0
Odometer:      17 days 5 hours 30 minutes since Oct-19-2004
Temperature:   35.1 deg C
Status:        PSU CTRL Mode:   Master
PSUCTRL-2 information:
State:          Operational
PartInfo:      PSUCTRL-2 04334-00068 450117-00-01
Revision:      1.0
FailureCode:   0
Odometer:      4 days 13 hours since Sep-21-2004
Temperature:   33.56 deg C
Status:        PSU CTRL Mode:   Backup
```

For the BlackDiamond X8 switch, this command shows the Power Entry Circuit (PEC) boards as follows:

```
BD-X8.41 # show power controller
PSUCTRL-1 information:
State:          Operational
PartInfo:      EV-AC-PEC 1102G-00137 450357-00-01
Revision:      1.0
PSUCTRL-2 information:
State:          Operational
PartInfo:      EV-AC-PEC 1102G-00138 450357-00-01
Revision:      1.0
Note: Power Supplies 1-4 (group A) use PSUCTRL-1 and 5-8 (group B) use
```



```
PSUCTRL-2.  
BD-X8.42 #
```

Unlike the power controller on the BlackDiamond 8000 switch, the PEC is a passive board. Thus, the state shown can only be either Operational (meaning the card is present and in use) or Empty.

History

This command was first available in an ExtremeXOS 11.0.

Platform Availability

This command is available only on modular switches.

show power led motion-detector

```
show power led motion-detector
```

Description

Displays the status of the motion detector on the Summit X670 switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to display the status and timeout setting of the motion detector.

Example

The following command displays the status for motion detection on the switch:

```
show power led motion-detector
```

The following is sample output from this command:

```
X670-48x.1 # show power led motion-detector  
Motion detector      : Enabled  
Timeout              : 180 seconds
```



History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on the Summit X670 switch only.

show session

```
show session [{detail} {sessID}] {history}
```

Description

Displays the currently active Telnet and console sessions communicating with the switch.

Syntax Description

detail	Specifies more detailed session information.
<i>sessID</i>	Specifies a session ID number.
history	Displays a list of all sessions.

Default

N/A.

Usage Guidelines

The `show session` command displays the username and IP address of the incoming Telnet session, whether a console session is currently active, and the login time. Each session is numbered.

Beginning with ExtremeXOS 11.2, the switch accepts IPv6 connections. If the incoming session is from an IPv6 address, the show session output indicates IPv6.

You can specify the following options to alter the session output:

- `detail`—The output for all current sessions is displayed in a list format.
- `sessID`—The output for the specified session is displayed in a list format.
- `history`—Displays a list of current and previous sessions, including the user, type of session, location, and start and end time of the session.

The `show session` command fields are defined in the following table.

Table 12: Show Command Field Definitions

Field	Definition
#	Indicates session number.
Login Time	Indicates login time of session.



Table 12: Show Command Field Definitions (continued)

Field	Definition
User	Indicates the user logged in for each session.
Type	Indicates the type of session, for example: console, telnet, http, https.
Auth	Indicates how the user is logged in.
CLI Auth	Indicates the type of authentication (RADIUS and TACACS) if enabled.
Location	Indicates the location (IP address) from which the user logged in. The output also indicates if the location is an IPv6 address.

Example

The following command displays the active sessions on the switch:

```
show session
```

The following is sample output from this command:

```
CLI
#      Login Time                User      Type     Auth    Auth Location
=====
==
1      Thu Apr 28 20:16:56 2005 admin    console local   dis    serial
*2     Thu Apr 28 23:36:20 2005 admin    ssh2    local   dis    3001::20d:
88ff:fec5:ad40
3      Fri Apr 29 11:14:27 2005 admin    telnet  local   dis    10.255.44.55
```

The following command displays a list of current and previous sessions on the switch:

```
show session history
```

The following is sample output from this command:

```
Session History:
admin                console    serial
Mon Jun 21 09:19:00 2004
Mon Jun 21 10:00:16 2004
admin                console    serial
Tue Jun 22 07:28:11 2004
Tue Jun 22 11:46:48 2004
admin                console    serial
Wed Jun 23 10:05:44 2004
Wed Jun 23 14:11:47 2004
admin                console    serial
Thu Jun 24 07:07:25 2004
Thu Jun 24 07:08:55 2004
admin                console    serial
Thu Jun 24 13:30:07 2004 Active
```



History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show snmp

```
show snmp [get | get-next] object_identifier
```

Description

Displays the contents of an SNMP MIB object.

Syntax Description

<i>object_identifier</i>	Specifies the object identifier for an SNMP MIB object.
--------------------------	---

Default

N/A.

Usage Guidelines

Use the get option to establish an index into the SNMP MIB. After the get option is executed, you can use the get next option to step through the MIB objects.

Example

The following gets the contents of SNMP object 1.3.6.1.2.1.1.5.0:

```
show snmp get 1.3.6.1.2.1.1.5.0
system.5.0 = BD-12804
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.



show snmp vr_name

```
show snmp {vr} vr_name
```

Description

Displays the SNMP configuration and statistics on a virtual router.

Syntax Description

<code>vr_name</code>	Specifies the virtual router.
----------------------	-------------------------------

Default

N/A.

Usage Guidelines

Use this command to display the SNMP configuration and statistics on a virtual router.

Example

The following command displays configuration and statistics for the virtual router VR-Default:

```
show snmp vr VR-Default
```

Following is sample output for the command:

```
SNMP access                : Disabled
SNMP Traps                  : Enabled
SNMP v1/v2c TrapReceivers  :
Destination                Source IP Address      Flags
10.120.91.89 /162          2E
Flags: Version: 1=v1 2=v2c
Mode: S=Standard E=Enhanced
SNMP stats:   InPkts 300   OutPkts  300   Errors 0   AuthErrors 0
Gets  0   GetNexts 300   Sets  0   Drops  0
SNMP traps:   Sent  0   AuthTraps Enabled
show snmp vr <vr_name>
SNMP access : Enabled
SNMP ifMib ifAlias size: Extended
SNMP Traps: Enabled
SNMP TrapReceivers: None
SNMP stats:   InPkts 0   OutPkts  0   Errors 0   AuthErrors
0
0
Gets  0   GetNexts 0   Sets  0   Drops
0
SNMP traps:   Sent  0   AuthTraps Enable
```



History

This command was first available in ExtremeXOS 12.4.2.

The SNMP ifMib ifAlias size status was added in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

show snmpv3 access

```
show snmpv3 access {[[hex hex_group_name] | group_name]}
```

Description

Displays SNMPv3 access rights.

Syntax Description

hex	Specifies that the value to follow is to be supplied as a colon separated string of hex octets.
<i>group_name</i>	Specifies the name of the group to display.

Default

N/A.

Usage Guidelines

The `show snmpv3 access` command displays the access rights of a group. If you do not specify a group name, the command will display details for all the groups.

This command displays the SNMPv3 vacmAccessTable entries.

Example

The following command displays all the access details:

```
show snmpv3 access
```

The following is sample output from this command:

```
switch # show snmpv3 access
Group Name      : admin
Context Prefix  :
Security Model  : USM
```



```

Security Level : Authentication Privacy
Context Match : Exact
Read View     : defaultAdminView
Write View    : defaultAdminView
Notify View   : defaultNotifyView
Storage Type  : NonVolatile
Row Status    : Active
Group Name    : initial
Context Prefix :
Security Model : USM
Security Level : No-Authentication No-Privacy
Context Match : Exact
Read View     : defaultUserView
Write View    :
Notify View   : defaultNotifyView
Storage Type  : NonVolatile
Row Status    : Active
Group Name    : initial
Context Prefix :
Security Model : USM
Security Level : Authentication No-Privacy
Context Match : Exact
Read View     : defaultUserView
Write View    : defaultUserView
Notify View   : defaultNotifyView
Storage Type  : NonVolatile
Row Status    : Active
Group Name    : vlv2c_ro
Context Prefix :
Security Model : snmpv1
Security Level : No-Authentication No-Privacy
Context Match : Exact
Read View     : defaultUserView
Write View    :
Notify View   : defaultNotifyView
Storage Type  : NonVolatile
Row Status    : Active
Group Name    : vlv2c_ro
Context Prefix :
Security Model : snmpv2c
Security Level : No-Authentication No-Privacy
Context Match : Exact
Read View     : defaultUserView
Write View    :
Notify View   : defaultNotifyView
Storage Type  : NonVolatile
Row Status    : Active
Group Name    : vlv2c_rw
Context Prefix :
Security Model : snmpv1
Security Level : No-Authentication No-Privacy
Context Match : Exact
Read View     : defaultUserView
Write View    : defaultUserView
Notify View   : defaultNotifyView
Storage Type  : NonVolatile
Row Status    : Active
Group Name    : vlv2c_rw

```



```

Context Prefix      :
Security Model     : snmpv2c
Security Level     : No-Authentication No-Privacy
Context Match     : Exact
Read View         : defaultUserView
Write View        : defaultUserView
Notify View       : defaultNotifyView
Storage Type      : NonVolatile
Row Status        : Active
Group Name        : vlv2cNotifyGroup
Context Prefix    :
Security Model    : snmpv1
Security Level    : No-Authentication No-Privacy
Context Match    : Exact
Read View        :
Write View       :
Notify View      : defaultNotifyView
Storage Type     : NonVolatile
Row Status      : Active
Group Name      : vlv2cNotifyGroup
Context Prefix   :
Security Model   : snmpv2c
Security Level   : No-Authentication No-Privacy
Context Match   : Exact
Read View       :
Write View      :
Notify View     : defaultNotifyView
Storage Type    : NonVolatile
Row Status     : Active
Total num. of entries in vacmAccessTable : 9

```

The following command displays the access rights for the group group1:

```
show snmpv3 access group1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show snmpv3 community

```
show snmpv3 community
```

Description

Displays information about SNMP community strings.



Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays information about and status of the SNMP community on the switch. This information is available to Administrator Accounts.

Example

The following command displays the community:

```
show snmpv3 community
```

The following is sample output from this command:

```
switch # show snmpv3 community
Community Index   : private
Community Name    : private
Security Name     : v1v2c_rw
Context EngineID  : 80:00:07:7c:03:00:04:96:27:b6:7b
Context Name      :
Transport Tag     :
Storage Type      : NonVolatile
Row Status        : Active
Community Index   : public
Community Name    : public
Security Name     : v1v2c_ro
Context EngineID  : 80:00:07:7c:03:00:04:96:27:b6:7b
Context Name      :
Transport Tag     :
Storage Type      : NonVolatile
Row Status        : Active
Total num. of entries in snmpCommunityTable : 2
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.



show snmpv3 context

show snmpv3 context

Description

Displays information about the SNMPv3 contexts on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the entries in the View-based Access Control Model (VACM) context table (VACMContextTable).

Example

The following command displays information about the SNMPv3 contexts on the switch:

```
show snmpv3 context
```

The following is sample output from this command:

```
VACM Context Name :  
Note : This Version Supports one global context ("")
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show snmpv3 counters

show snmpv3 counters



Description

Displays SNMPv3 counters.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The `show snmpv3 counters` command displays the following SNMPv3 counters:

- `snmpUnknownSecurityModels`.
- `snmpInvalidMessages`.
- `snmpUnknownPDUHandlers`.
- `usmStatsUnsupportedSecLevels`.
- `usmStatsNotInTimeWindows`.
- `usmStatsUnknownUserNames`.
- `usmStatsUnknownEngineIDs`.
- `usmStatsWrongDigests`.
- `usmStatsDecryptionErrors`.

Issuing the command `clear counters` resets all counters to zero.

Example

The following command displays all the SNMPv3 counters:

```
show snmpv3 counters
```

The following is sample output from this command:

```
snmpUnknownSecurityModels      : 0
snmpInvalidMessages            : 0
snmpUnknownPDUHandlers         : 0
usmStatsUnsupportedSecLevels    : 0
usmStatsNotInTimeWindows       : 0
usmStatsUnknownUserNames       : 0
usmStatsUnknownEngineIDs       : 0
usmStatsWrongDigests           : 0
usmStatsDecryptionErrors       : 0
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show snmpv3 engine-info

show snmpv3 engine-info

Description

Displays information about the SNMPv3 engine on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The following show engine-info output is displayed:

- Engine-ID—Either the ID auto generated from MAC address of switch, or the ID manually configured.
- Engine Boots—Number of times the agent has been rebooted.
- Engine Time—Time since agent last rebooted, in centiseconds.
- Max. Message Size—Maximum SNMP Message size supported by the Engine (8192).

Example

The following command displays information about the SNMPv3 engine on the switch:

```
show snmpv3 engine-info
```

The following is sample output from this command:

```
SNMP Engine-ID          : 80:0:7:7c:3:0:30:48:41:ed:97 'H'  
SNMP Engine Boots      : 1  
SNMP Engine Time       : 866896  
SNMP Max. Message Size : 8192
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show snmpv3 extreme-target-addr-ext

```
show snmpv3 extreme-target-addr-ext [[hex hex_addr_name] | addr_name]
```

Description

Displays information about SNMPv3 target addresses enhanced or standard mode.

Syntax Description

<i>hex_addr_name</i>	Specifies an identifier for the target address. The value is to be supplied as a colon separated string of hex octets.
<i>addr_name</i>	Specifies a string identifier for the target address.

Default

N/A.

Usage Guidelines

Use this command to display entries in the SNMPv3 extremeTargetAddressExtTable.

Example

The following command displays the entry for the target address named A1:

```
show snmpv3 extreme-target-addr-ext A1
```

The following is sample output from this command:

```
Target Addr Name      : A1
Mode                  : Enhanced
IgnoreMPModel         : No
UseEventComm          : Yes
```

History

This command was first available in ExtremeXOS 10.1.



The `hex_addr_name` parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show snmpv3 filter

```
show snmpv3 filter {[hex hex_profile_name] | profile_name] {{subtree}}
object_identifier}
```

Description

Displays the filters that belong a filter profile.

Syntax Description

<i>hex_profile_name</i>	Specifies the filter profile to display. The value is to be supplied as a colon separated string of hex octets.
<i>profile_name</i>	Specifies the filter profile to display in ASCII format.
<i>object_identifier</i>	Specifies a MIB subtree.

Default

N/A.

Usage Guidelines

Use this command to display entries from the `snmpNotifyFilterTable`. If you specify a profile name and subtree, you will display only the entries with that profile name and subtree. If you specify only the profile name, you will display all entries for that profile name. If you do not specify a profile name, then all the entries are displayed.

Example

The following command displays the part of filter profile `prof1` that includes the MIB subtree `1.3.6.1.4.1`:

```
show snmpv3 filter prof1 subtree 1.3.6.1.4.1
```

The following is sample output from this command:

```
Profile Name      : prof1
Subtree           : 1.3.6.1.4.1
Mask              :
Type              : Included
```



```
Storage Type      : NonVolatile
Row Status       : Active
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_profile_name` parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show snmpv3 filter-profile

```
show snmpv3 filter-profile {[[hex hex_profile_name] | profile_name]} {param [[hex hex_param_name] | param_name]}
```

Description

Displays the association between parameter names and filter profiles.

Syntax Description

<i>hex_profile_name</i>	Specifies the filter profile name. The value is to be supplied as a colon separated string of hex octets.
<i>profile_name</i>	Specifies the filter profile name in ASCII format.
<i>hex_param_name</i>	Specifies the parameter name. The values is to be supplied as a colon separated string of hex octets.
<i>param_name</i>	Specifies the parameter name in ASCII format.

Default

N/A.

Usage Guidelines

Use this command to display the `snmpNotifyFilterProfileTable`. This table associates a filter profile with a parameter name. The parameter name is associated with target addresses, and the filter profile is associated with a series of filters, so, in effect, you are associating a series of filters with a target address.



Example

The following command displays the entry with filter profile prof1 with the parameter name P1:

```
show snmpv3 filter-profile prof1 param P1
```

The following is sample output of this command:

```
Filter Profile Params Name : p1
Name                      : prof1
Storage Type              : NonVolatile
Row Status                 : Active
```

History

This command was first available in ExtremeXOS 10.1.

The hex_profile_name and hex_param_name parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show snmpv3 group

```
show snmpv3 group {[[hex hex_group_name] | group_name] {user [[hex hex_user_name]
| user_name]}}
```

Description

Displays the user name (security name) and security model association with a group name.

Syntax Description

<i>hex_group_name</i>	Specifies the group name to display. The value is to be supplied as a colon separated string of hex octets.
<i>group_name</i>	Specifies the group name to display. The value is to be supplied in ASCII format.
<i>hex_user_name</i>	Specifies the user name to display. The value is to be supplied as a colon separated string of hex octets.
<i>user_name</i>	Specifies the user name to display. The value is to be supplied in ASCII format.

Default

N/A.



Usage Guidelines

The `show snmpv3 group` command displays the details of a group with the given group name. If you do not specify a group name, the command will display details for all the groups.

This command displays the SNMPv3 vacmSecurityToGroupTable.

Example

The following command displays information about all groups for every security model and user name:

```
show snmpv3 group
```

The following is sample output from this command:

```
switch # sh snmpv3 group
Group Name      : vlv2c_ro
Security Name   : vlv2c_ro
Security Model  : snmpv1
Storage Type    : NonVolatile
Row Status      : Active
Group Name      : vlv2c_rw
Security Name   : vlv2c_rw
Security Model  : snmpv1
Storage Type    : NonVolatile
Row Status      : Active
Group Name      : vlv2c_ro
Security Name   : vlv2c_ro
Security Model  : snmpv2c
Storage Type    : NonVolatile
Row Status      : Active
Group Name      : vlv2c_rw
Security Name   : vlv2c_rw
Security Model  : snmpv2c
Storage Type    : NonVolatile
Row Status      : Active
Group Name      : admin
Security Name   : admin
Security Model  : USM
Storage Type    : NonVolatile
Row Status      : Active
Group Name      : initial
Security Name   : initial
Security Model  : USM
Storage Type    : NonVolatile
Row Status      : Active
Group Name      : initial
Security Name   : initialmd5
Security Model  : USM
Storage Type    : NonVolatile
Row Status      : Active
Group Name      : initial
Security Name   : initialsha
Security Model  : USM
Storage Type    : NonVolatile
```



```

Row Status      : Active
Group Name      : initial
Security Name   : initialmd5Priv
Security Model  : USM
Storage Type    : NonVolatile
Row Status      : Active
Group Name      : initial
Security Name   : initialshaPriv
Security Model  : USM
Storage Type    : NonVolatile
Row Status      : Active
Total num. of entries in vacmSecurityToGroupTable : 10

```

The following command shows information about the group testgroup and user name testuser:

```
show snmpv3 group testgroup user testuser
```

The following is sample output from this command:

```

Group Name      : testgroup
Security Name    : testuser
Security Model   : USM
Storage Type     : NonVolatile
Row Status       : Active

```

History

This command was first available in ExtremeXOS 10.1.

The hex_group_name and hex_user_name parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show snmpv3 mib-view

```
show snmpv3 mib-view {[hex hex_view_name] | view_name} {subtree
object_identifier}}
```

Description

Displays a MIB view.



Syntax Description

<i>hex_view_name</i>	Specifies the name of the MIB view to display. The value is to be supplied as a colon separated string of hex octets.
<i>view_name</i>	Specifies the name of the MIB view to display. The value is to be supplied in ASCII format.
<i>object_identifier</i>	Specifies the object identifier of the view to display.

Default

N/A.

Usage Guidelines

The `show snmpv3 mib-view` command displays a MIB view. If you do not specify a view name, the command will display details for all the MIB views. If a subtree is not specified, then all subtrees belonging to the view name will be displayed.

This command displays the SNMPv3 vacmViewTreeFamilyTable.

Example

The following command displays all the view details:

```
show snmpv3 mib-view
```

The following is sample output from this command:

```
switch # sh snmpv3 mib-view
View Name      : defaultUserView
MIB Subtree    : 1
Mask           :
View Type      : Included
Storage Type   : NonVolatile
Row Status     : Active
View Name      : defaultUserView
MIB Subtree    : 1.3.6.1.6.3.16
Mask           :
View Type      : Excluded
Storage Type   : NonVolatile
Row Status     : Active
View Name      : defaultUserView
MIB Subtree    : 1.3.6.1.6.3.18
Mask           :
View Type      : Excluded
Storage Type   : NonVolatile
Row Status     : Active
View Name      : defaultUserView
MIB Subtree    : 1.3.6.1.6.3.15.1.2.2.1.4
Mask           :
View Type      : Excluded
```



```

Storage Type      : NonVolatile
Row Status        : Active
View Name         : defaultUserView
MIB Subtree       : 1.3.6.1.6.3.15.1.2.2.1.6
Mask              :
View Type         : Excluded
Storage Type      : NonVolatile
Row Status        : Active
View Name         : defaultUserView
MIB Subtree       : 1.3.6.1.6.3.15.1.2.2.1.9
Mask              :
View Type         : Excluded
Storage Type      : NonVolatile
Row Status        : Active
View Name         : defaultAdminView
MIB Subtree       : 1
Mask              :
View Type         : Included
Storage Type      : NonVolatile
Row Status        : Active
View Name         : defaultNotifyView
MIB Subtree       : 1
Mask              :
View Type         : Included
Storage Type      : NonVolatile
Row Status        : Active
Total num. of entries in vacmViewTreeFamilyTable : 8

```

The following command displays a view with the view name Review and subtree 1.3.6.1.2.1.1:

```
show snmpv3 mib-view Review subtree 1.3.6.1.2.1.1
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_view_name` parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show snmpv3 notify

This command displays the `snmpNotifyTable`.

```
show snmpv3 notify {[hex hex_notify_name] | notify_name}
```

Description

Displays the notifications that are set.



Syntax Description

<i>hex_notify_name</i>	Specifies the parameter name associated with the target. The value is to be supplied as a colon separated string of hex octets.
<i>notify_name</i>	Specifies the parameter name associated with the target. The value is to be supplied in ASCII format.

Default

N/A.

Usage Guidelines

Use this command to display entries from the SNMPv3 snmpNotifyTable. This table lists the notify tags that the agent will use to send notifications (traps).

If no notify name is specified, all the entries are displayed.

Example

The following command displays the notify table entries:

```
show snmpv3 notify
```

The following is sample output from this command:

```
BD-12804.2 # show snmpv3 notify
Notify Name      : defaultNotify
Tag              : defaultNotify
Type             : Trap
Storage Type     : NonVolatile
Row Status       : Active
Notify Name      : xyz
Tag              : xyz1
Type             : Inform
Storage Type     : NonVolatile
Row Status       : Active
Total entries in snmpNotifyTable : 2
```

History

This command was first available in ExtremeXOS 10.1.

The hex_notify_name parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.



show snmpv3 target-addr

```
show snmpv3 target-addr {[[hex hex_addr_name] | addr_name]}
```

Description

Displays information about SNMPv3 target addresses.

Syntax Description

<i>hex_addr_name</i>	Specifies an identifier for the target address. The value is to be supplied as a colon separated string of hex octets.
<i>addr_name</i>	Specifies a string identifier for the target address.

Default

N/A.

Usage Guidelines

Use this command to display entries in the SNMPv3 snmpTargetAddressTable. If no target address is specified, the entries for all the target addresses will be displayed.

To view the source IP address, use the `show management` command.

Example

The following command displays the entry for the target address named A1:

```
show snmpv3 target-addr A1
```

The following is sample output from this command:

```
Target Addr Name      : A1
TDomain               : 1.3.6.1.6.1.1
TAddress              : 10.201.31.234, 162
TMask                 :
Timeout               : 1500
Retry Count           : 0
Tag List              : defaultNotify
Params                : v1v2cNotifyParam1
Storage Type          : NonVolatile
Row Status            : Active
Storage Type          : NonVolatile
Row Status            : Active
```



History

This command was first available in ExtremeXOS 10.1.

This command was modified in ExtremeXOS 11.0 to display a list of tags if more than one was configured and to display the timeout value for the entry in the snmpTargetAddrTable. This command was also modified to support the hex_addr_name parameter.

Platform Availability

This command is available on all platforms.

show snmpv3 target-params

```
show snmpv3 target-params {[[hex hex_target_params] | target_params]}
```

Description

Displays the information about the options associated with the parameter name.

Syntax Description

<i>hex_target_params</i>	Specifies the parameter to display. The value is to be supplied as a colon separated string of hex octets.
<i>target_params</i>	Specifies the parameter name to display. The value is to be supplied in ASCII format.

Default

N/A.

Usage Guidelines

Use this command to display entries from the SNMPv3 snmpTargetParamsTable. This table specifies the message processing model, security level, security model, and the storage parameters for messages to any target addresses associated with a particular parameter name.

If no parameter name is specified, all the entries are displayed.

Example

The following command displays the target parameter entry named P1:

```
show snmpv3 target-params P1
```



The following is sample output from this command:

```
Target Params Name      : p1
MP Model                : snmpv2c
Security Model          : snmpv2c
User Name               : testuser
Security Level          : No-Authentication No-Privacy
Storage Type            : NonVolatile
Row Status              : Active
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_target_params` parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show snmpv3 user

```
show snmpv3 user {[hex hex_user_name] | user_name}
```

Description

Displays detailed information about the user.

Syntax Description

<i>hex_user_name</i>	Specifies the user name to display. The value is to be supplied as a colon separated string of hex octets.
<i>user_name</i>	Specifies the user name to display. The value is to be supplied in ASCII format.

Default

N/A.

Usage Guidelines

The `show snmpv3 user` command displays the details of a user. If you do not specify a user name, the command will display details for all the users. The authentication and privacy passwords and keys will not be displayed.

The user entries in SNMPv3 are stored in the USMUserTable, so the entries are indexed by EngineID and user name.



Example

The following command lists all user entries:

```
show snmpv3 user
```

The following is sample output from this command:

```
switch # sh snmpv3 user
Engine-ID       : 80:00:07:7c:03:00:04:96:27:b6:7b 'H'
User Name       : admin
Security Name   : admin
Authentication  : HMAC-MD5
Privacy         : DES
Storage Type    : NonVolatile
Row Status      : Active
Engine-ID       : 80:00:07:7c:03:00:04:96:27:b6:7b 'H'
User Name       : initial
Security Name   : initial
Authentication  : No-Authentication
Privacy         : No-Privacy
Storage Type    : NonVolatile
Row Status      : Active
Engine-ID       : 80:00:07:7c:03:00:04:96:27:b6:7b 'H'
User Name       : initialmd5
Security Name   : initialmd5
Authentication  : HMAC-MD5
Privacy         : No-Privacy
Storage Type    : NonVolatile
Row Status      : Active
Engine-ID       : 80:00:07:7c:03:00:04:96:27:b6:7b 'H'
User Name       : initialsha
Security Name   : initialsha
Authentication  : HMAC-SHA
Privacy         : No-Privacy
Storage Type    : NonVolatile
Row Status      : Active
Engine-ID       : 80:00:07:7c:03:00:04:96:27:b6:7b 'H'
User Name       : initialmd5Priv
Security Name   : initialmd5Priv
Authentication  : HMAC-MD5
Privacy         : DES
Storage Type    : NonVolatile
Row Status      : Active
Engine-ID       : 80:00:07:7c:03:00:04:96:27:b6:7b 'H'
User Name       : initialshaPriv
Security Name   : initialshaPriv
Authentication  : HMAC-SHA
Privacy         : DES
Storage Type    : NonVolatile
Row Status      : Active
Total num. of entries in usmUserTable : 6
```



The following command lists details for the specified user, testuser:

```
show snmpv3 user testuser
```

History

This command was first available in ExtremeXOS 10.1.

The hex_user_name parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show sntp-client

show sntp-client

Description

Displays the DNS configuration.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Displays configuration and statistics information of SNTP client.

Example

The following command displays the SNTP configuration:

```
show sntp-client
```

The following is sample output from this command:

```
SNTP client is enabled
SNTP time is valid
Primary server: 172.17.1.104
```



```

Secondary server: 172.17.1.104
Query interval: 64
Last valid SNTP update: From server 172.17.1.104, on Wed Oct 30 22:46:03 2003
SNTPC Statistics:
Packets transmitted:
to primary server:      1
to secondary server:    0
Packets received with valid time:
from Primary server:    1
from Secondary server:  0
from Broadcast server:  0
Packets received without valid time:
from Primary server:    0
from Secondary server:  0
from Broadcast server:  0
Replies not received to requests:
from Primary server:    0
from Secondary server:  0

```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

telnet

```
telnet {vr vr_name} [host_name | remote_ip] {port}
```

Description

Allows you to Telnet from the current command-line interface session to another host.

Syntax Description

vr	Specifies use of a virtual router.NOTE: User-created VRs are supported only on the platforms listed for this feature in the Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i> .
<i>vr_name</i>	Specifies the name of the virtual router.
<i>host_name</i>	Specifies the name of the host.
<i>remote_ip</i>	Specifies the IP address of the host.
<i>port</i>	Specifies a TCP port number. The default is port 23.

Default

- Telnet—enabled.



- Virtual router—Uses all virtual routers on the switch for outgoing Telnet requests.
- Port—23.

Usage Guidelines

Only VT100 emulation is supported.

Before you can start an outgoing Telnet session, you need to configure the switch IP parameters. To open a Telnet connection, you must specify the host IP address or the host name of the device you want to connect to. Check the user manual supplied with the Telnet facility if you are unsure of how to do this. Although the switch accepts IPv6 connections, you can only Telnet from the switch to another device with an IPv4 address.

You must configure DNS in order to use the `host_name` option.

Host Name and Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z).
- Numerals (0-9).
- Period (.).
- Dash (-) Permitted only for host names.
- Underscore (_) Permitted only for host names.
- Colon (:).

When naming or configuring an IP address for your network server, remember the requirements listed above.

Virtual Router Requirements

The `vr_name` option specifies the name of the virtual router. The valid virtual router names at system boot-up are VR-Mgmt, VR-Control, and VR-Default; however, you can only Telnet on VR-Mgmt and VR-Default. In ExtremeXOS 10.1, the valid virtual routers are VR-0, VR-1, and VR-2 respectively, and Telnet used VR-0 by default. For more information about virtual routers, see [Virtual Routers](#) in the *ExtremeXOS Concepts Guide*.

Example

The following command starts a Telnet client communication to the host at IP address 123.45.67.8:

```
telnet 123.45.67.8
```



The following command starts a Telnet client communication with a host named sales:

```
telnet sales
```

History

This command was first available in ExtremeXOS 10.1.

Support for the following virtual routers was added in ExtremeXOS 11.0: VR-Mgmt and VR-Default.

Platform Availability

This command is available on all platforms.

telnet msm

```
telnet msm [a | b]
```

Description

Allows you to Telnet to either the primary or the backup MSM regardless of which console port you are connected to.

Syntax Description

a	Specifies the MSM installed in slot A.
b	Specifies the MSM installed in slot B.

Default

N/A.

Usage Guidelines

Use this command to access either the primary or the backup MSM regardless of which console port you are connected to. For example, if MSM A is the primary MSM and you are connected to MSM A via its console port, you can access the backup MSM installed in slot B by issuing the telnet msm b command.

Example

The following example makes the following assumptions:

- The MSM installed in slot A is the primary.
- The MSM installed in slot B is the backup.
- You have a console connection to MSM B.



The following command accesses the primary MSM installed in slot A from the backup MSM installed in slot B:

```
My8800.6 # telnet msm b
Entering character mode
Escape character is '^]'.
telnet session telnet0 on /dev/ptyb0
login: admin
password:
ExtremeXOS
Copyright (C) 2000-2007 Extreme Networks. All rights reserved.
Protected by US Patent Nos: 6,678,248; 6,104,700; 6,766,482; 6,618,388;
6,034,957; 6,859,438; 6,912,592; 6,954,436; 6,977,891; 6,980,550; 6,981,174;
7,003,705; 7,012,082; 7,046,665; 7,126,923; 7,142,509; 7,149,217; 7,152,124;
7,154,861.
=====
You are connected to a Backup node. Only a limited command set is supported.
You may use "telnet msm A" to connect to the Master node to access
the full set of commands.
Press the <tab> or '?' key at any time for completions.
Remember to save your configuration changes.
My8800.1 >
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available only on modular switches.

telnet slot

telnet slot *slot-number*

Description

Allows you to Telnet from any source node to any other target node when both nodes are in the same active topology.

Syntax Description

<i>slot-number</i>	Specifies the number of a slot that is currently occupied by the intended target node.
--------------------	--

Default

N/A.



Usage Guidelines

When the target node accepts the related TCP connection, it prompts the user for a user ID and password. If the failsafe account is used, user ID and password authentication takes place on the specified node. Otherwise, authentication takes place on the master node, regardless of the source and target nodes used.

Telnet must be enabled on the SummitStack for this command to function.

Example

The following command accesses the node in slot 2:

```
Slot-1 Stack.6 # telnet slot 2
Entering character mode
Escape character is '^]'.
telnet session telnet0 on /dev/ptyb0
login: admin
password:
ExtremeXOS
Copyright (C) 2000-2007 Extreme Networks. All rights reserved.
Protected by US Patent Nos: 6,678,248; 6,104,700; 6,766,482; 6,618,388;
6,034,957; 6,859,438; 6,912,592; 6,954,436; 6,977,891; 6,980,550; 6,981,174;
7,003,705; 7,012,082; 7,046,665; 7,126,923; 7,142,509; 7,149,217; 7,152,124;
7,154,861.
=====
You are connected to a Backup node. Only a limited command set is supported.
You may use "telnet slot <slot_number>" to connect to the Master node to
access
the full set of commands.
Press the <tab> or '?' key at any time for completions.
Remember to save your configuration changes.
Slot-2 Stack.1 >
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on SummitStack.

tftp

```
tftp [host-name | ip-address] {-v vr_name [-g | -p] [{"-l [internal-memory local-
file-internal | memorycard local-file-memcard | local-file] {-r remote-file |-
r remote-file] {-l [internal-memory local-file-internal | memorycard local-file-
memcard | local-file]}]
```



Description

Allows you to TFTP from the current command line interface session to a TFTP server.

Syntax Description

<i>host-name</i>	Specifies the name of the remote host.
<i>ip-address</i>	Specifies the IP address of the TFTP server.
<i>vr_name</i>	Specifies the name of the virtual router. NOTE: User-created VRs are supported only on the platforms listed for this feature in the Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i> .
-g	Gets the specified file from the TFTP server and copies it to the local host.
-p	Puts the specified file from the local host and copies it to the TFTP server.
internal-memory	Specifies the internal memory card.
<i>local-file-internal</i>	Specifies the name of the core dump file located on the internal memory card.
memorycard	Specifies a removable storage device, which can be a compact flash card or a USB 2.0 storage device. NOTE: This parameter is available for compact flash cards on BlackDiamond 8800 and for USB 2.0 storage devices on Summit X460, X480, X670, X670V, and X770 switches.
<i>local-file-memcard</i>	Specifies the name of the file on a removable storage device, which can be a compact flash card or a USB 2.0 storage device. NOTE: This parameter is available for compact flash cards on BlackDiamond 8800 and for USB 2.0 storage devices on Summit X460, X480, X670, X670V, and X770 switches.
<i>local-file</i>	Specifies the name of the file (configuration file, policy file) on the local host.
<i>remote-file</i>	Specifies the name of the file on the remote host.

Default

If you do not specify a virtual router, VR-Mgmt is used.

If you do not specify the internal memory card or a removable storage device, the switch downloads or uploads the file from the switch local file system.

Usage Guidelines

NetASCII and mail file type formats are not supported.

TFTP Server Requirements

Extreme Networks recommends using a TFTP server that supports blocksize negotiation (as described in RFC 2348, TFTP Blocksize Option), to enable faster file downloads and larger file downloads. If the TFTP server does not support blocksize negotiation, the file size is limited to 32MB. Older TFTP servers that do not support blocksize negotiation have additional implementation limits that may decrease the maximum file size to only 16MB, which may be too small to install ExtremeXOS images.



If your TFTP server does not support blocksize negotiation, the switch displays a message similar to the following when you attempt a get (-g) or put (-p) operation:

```
Note: The blocksize option is not supported by the remote TFTP server.  
Without this option, the maximum file transfer size is limited to 32MB.  
Some older TFTP servers may be limited to 16MB file.
```

Using TFTP

Use TFTP to download a previously saved configuration file or policy file from the TFTP server to the switch. When you download a file, this command does not automatically apply it to the switch. You must specify that the downloaded file be applied to the switch. For example, if you download a configuration file, issue the `use configuration` command to apply the saved configuration on the next reboot. You must use the `reboot` command to activate the new configuration. If you download a policy file, use the `refresh policy` command to reprocess the text file and update the policy database.

You also use TFTP to upload a saved configuration file or policy file from the switch to the TFTP server.

If your download from the TFTP server to the switch is successful, the switch displays a message similar to the following:

```
Downloading megtest2.cfg to switch... done!
```

If your upload from the switch to the TFTP server is successful, the switch displays a message similar to the following:

```
Uploading megtest1.cfg to TFTPHost ... done!
```

Up to eight active TFTP sessions can run on the switch concurrently.

You must configure DNS in order to use the `host_name` option.

Host Name and Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z).
- Numerals (0-9).
- Period (.).
- Dash (-) Permitted only for host names.
- Underscore (_) Permitted only for host names.
- Colon (:).



When naming or configuring an IP address for your network server, remember the requirements listed above.

Local and Remote Filename Character Restrictions

This section provides information about the characters supported by the switch for local and remote filenames.

When specifying a local or remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z).
- Numerals (0-9).
- Period (.).
- Dash (-).
- Underscore (_).
- Slash (/) Permitted only for remote files.

When naming a local or remote file, remember the requirements listed above.

Virtual Router Requirements

The `vr_name` option specifies the name of the virtual router. The valid virtual router names at system boot-up are VR-Mgmt, VR-Control, and VR-Default; however, you can only TFTP on VR-Mgmt and VR-Default. In ExtremeXOS 10.1, the valid virtual routers are VR-0, VR-1, and VR-2 respectively. For more information about virtual routers, see “Virtual Routers” in the *ExtremeXOS Concepts Guide*.

Internal Memory and Core Dump Files

Core dump files have a `.gz` file extension. The filename format is: `core.process-name.pid.gz` where `process-name` indicates the name of the process that failed and `pid` is the numerical identifier of that process. If you have a modular switch and save core dump files to the external memory card, the filename also includes the affected MSM/MM: MSM-A or MSM-B.

If you configure and enable the switch to send core dump (debug) information to the internal memory card, specify the `internal-memory` option to transfer those files from the internal memory card to a TFTP server.

If you specify the `memorycard` option, you can use TFTP to copy and transfer files to and from removable storage devices (compact flash cards or USB 2.0 storage devices) on BlackDiamond 8800 series switches and Summit X460, X480, X670, X670V, and X770 switches.

If the switch has not saved any debug files, you cannot transfer other files to or from the internal memory. For example if you attempt to transfer a configuration file from the switch to the internal memory, the switch displays a message similar to the following:

```
Error: tftp transfer to internal-memory not allowed.
```



For information about configuring and sending core dump information to the internal memory card, see the `configure debug core-dumps` and `save debug tracefiles memorycard` commands.

For more detailed information about core dump files, see [Troubleshooting](#) in the *ExtremeXOS Concepts Guide*.

Other Useful Commands

On the Summit family switches and SummitStack, use the `download bootrom` command to upgrade the BootROM. This command utilizes TFTP to transfer the BootROM image file from your TFTP server to the switch. Only upgrade the BootROM when asked to do so by an Extreme Networks technical representative. For more information about this command, see `download bootrom`.

To upgrade the image, use the `download image` command. This command utilizes TFTP to transfer the software image file from your TFTP server to the switch. For more information about this command, see `download image`.

Example

The following command downloads the configuration file named XOS1.cfg from the TFTP server with an IP address of 10.123.45.67:

```
tftp 10.123.45.67 -v "VR-Default" -g -r XOS1.cfg
```

The following command uploads the configuration file named XOS2.cfg to the TFTP server with an IP address of 10.123.45.67:

```
tftp 10.123.45.67 -v "VR-Default" -p -r XOS2.cfg
```

The following command downloads a policy file to a removable storage device:

```
tftp 10.1.2.3 -g -l memorycard test.pol -r august23.pol
```

History

This command was first available in ExtremeXOS 10.1.

The memorycard option was added in ExtremeXOS 11.1.

The internal-memory option was added in ExtremeXOS 11.4.

Support for USB 2.0 storage devices was added in ExtremeXOS 12.5.3.

Platform Availability

This command is available on all platforms.



tftp get

```
tftp get [host-name | ip-address] {vr vr_name} [{internal-memory local-file-internal | memorycard local-file-memcard | local_file} {remote_file} | {remote_file} [{internal-memory local-file-internal | memorycard local-file-memcard | local_file}]] {force-overwrite}
```

Description

Allows you to use TFTP from the current command line interface session to copy the file from a TFTP server and copy it to a local host, including the switch, internal memory card, compact flash card, or USB 2.0 storage device.

Syntax Description

<i>host-name</i>	Specifies the name of the remote host.
<i>ip-address</i>	Specifies the IP address of the TFTP server.
<i>vr_name</i>	Specifies the name of the virtual router. NOTE: User-created VRs are supported only on the platforms listed for this feature in <i>Feature License Requirements</i> found in the <i>ExtremeXOS Concepts Guide</i> .
internal-memory	Specifies the internal memory card.
<i>local-file-internal</i>	Specifies the name of the core dump file located on the internal memory card.
memorycard	Specifies a removable storage device, which can be a compact flash card or a USB 2.0 storage device. NOTE: This parameter is available for compact flash cards on BlackDiamond 8800 series switches and for USB 2.0 storage devices on Summit X460, X480, X670, X670V, X770 switches.
<i>local-file-memcard</i>	Specifies the name of the file on a removable storage device, which can be a compact flash card or a USB 2.0 storage device. NOTE: This parameter is available for compact flash cards on BlackDiamond 8800 series switches and for USB 2.0 storage devices on Summit X460, X480, X670, X670V, and X770 switches.
<i>local_file</i>	Specifies the name of the file (configuration file, policy file) on the local host.
<i>remote_file</i>	Specifies the name of the file on the remote host.
force-overwrite	Specifies the switch to automatically overwrite an existing file.

Default

If you do not specify a virtual router, VR-Mgmt is used; if you transfer a file with a name that already exists on the system, the switch prompts you to overwrite the existing file.

If you do not specify the internal memory card or a removable storage device, the switch downloads or uploads the file from the switch local file system.

Usage Guidelines

NetASCII and mail file type formats are not supported.



By default, the switch prompts you to overwrite an existing file. For example, if you have a file named `test.cfg` on the switch and download a file named `test.cfg` from a TFTP server, the switch displays a message similar to the following:

```
test.cfg already exists, do you want to overwrite it? (y/n)
```

Enter `y` to download the file and overwrite the existing file. Enter `n` to cancel this action.

If you successfully download the file, the switch displays a message similar to the following:

```
Downloading test.cfg to switch... done!
```

If you cancel this action, the switch displays a message similar to the following:

```
Tftp download aborted.
```

If you specify the force-overwrite parameter, the switch automatically overwrites an existing file. For example, if you have a file named `test.cfg` on the switch and download a file named `test.cfg` from a TFTP server, the switch automatically overrides the existing file. If you successfully download the file, the switch displays a message similar to the following:

```
Downloading test.cfg to switch... done!
```

This command was introduced to simplify using TFTP to transfer configuration, policy, and if configured, core dump files from the switch to the TFTP server. You can continue to use the original `tftp` command introduced in ExtremeXOS 10.1.

For more information about TFTP, including:

- TFTP server requirements.
- How to use TFTP.
- Host name and remote IP address character restrictions.
- Local and remote filename character restrictions.
- Virtual router requirements.
- Internal memory and core dump files.
- Topics related to only modular switches.
- Other useful commands.

See the `tftp` command.

Example

The following command retrieves and transfers the file `test.pol` from a TFTP server with an IP address of `10.1.2.3` and renames the file `august23.pol` when transferred to a removable storage device:

```
tftp get 10.1.2.3 vr "VR-Mgmt" test.pol memory-card august23.pol
```

The following command retrieves the configuration file named `meg-upload.cfg` from a TFTP server with an IP address of `10.10.10.10`:

```
tftp get 10.10.10.10 vr "VR-Mgmt" meg_upload.cfg
```



History

This command was first available in ExtremeXOS 11.4.

Support for USB 2.0 storage devices was added in ExtremeXOS 12.5.3.

Platform Availability

This command is available on all platforms.

tftp put

```
tftp put [host-name | ip-address] {-vr vr_name} [{[internal-memory local-file-internal | memorycard local-file-memcard | local_file] {remote_file} | {remote_file} {[internal-memory local-file-internal | memorycard local-file-memcard | local_file]}]
```

Description

Allows you to use TFTP from the current command line interface session to copy the file from the local host, including the switch, internal memory card, compact flash card, or USB 2.0 storage device and put it on a TFTP server.

Syntax Description

<i>host-name</i>	Specifies the name of the remote host.
<i>ip-address</i>	Specifies the IP address of the TFTP server.
<i>vr_name</i>	Specifies the name of the virtual router. NOTE: User-created VRs are supported only on the platforms listed for this feature in the Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i> .
internal-memory	Specifies the internal memory card.
<i>local-file-internal</i>	Specifies the name of the core dump file located on the internal memory card.
memorycard	Specifies a removable storage device, which can be a compact flash card or a USB 2.0 storage device. NOTE: This parameter is available for compact flash cards on BlackDiamond 8800 series switches and for USB 2.0 storage devices on Summit X460, X480, X670, X670V, and X770 switches.
<i>local-file-memcard</i>	Specifies the name of the file on a removable storage device, which can be a compact flash card or a USB 2.0 storage device. NOTE: This parameter is available for compact flash cards on BlackDiamond 8800 and for USB 2.0 storage devices on Summit X460, X480, X670, X670V, and X770 switches.
<i>local_file</i>	Specifies the name of the file (configuration file, policy file) on the local host.
<i>remote_file</i>	Specifies the name of the file on the remote host.

Default

If you do not specify a virtual router, VR-Mgmt is used.



If you do not specify the internal memory card or a removable storage device, the switch downloads or uploads the file from the switch local file system.

Usage Guidelines

NetASCII and mail file type formats are not supported.

This command was introduced to simplify using TFTP to transfer configuration, policy, and if configured, core dump files from the switch to the TFTP server. You can continue to use the original `tftp` command introduced in ExtremeXOS 10.1.

For more information about TFTP, including:

- TFTP server requirements.
- How to use TFTP.
- Host name and remote IP address character restrictions.
- Local and remote filename character restrictions.
- Virtual router requirements.
- Internal memory and core dump files.
- Topics related to only modular switches.
- Other useful commands.

See the `tftp` command `tftp`.

Example

The following command transfers a saved, not currently used configuration file named XOS1.cfg from the switch to the TFTP server:

```
tftp put 10.123.45.67 vr "VR-Mgmt" XOS1.cfg
```

History

This command was first available in ExtremeXOS 11.4.

Support for USB 2.0 storage devices was added in ExtremeXOS 12.5.3.

Platform Availability

This command is available on all platforms.



5 Commands for Managing the ExtremeXOS Software

```
cp
disable xml-mode
enable xml-mode
ls
mv
restart process
rm
show heartbeat process
show memory
show memory process
show process
start process
terminate process
```

This chapter describes commands for:

- Working with the configuration and policy files used by the switch.
- Starting, stopping, and displaying information about processes on the switch.
- Viewing system memory resources.

Note



For information about downloading and upgrading a new software image, saving configuration changes, and upgrading the BootROM, see [Configuration and Image Commands](#).

Like any advanced operating system, ExtremeXOS gives you the tools to manage your switch and create your network configurations. With the introduction of ExtremeXOS, the following enhancements and functionality have been added to the switch operating system:

- File system administration—You can move, copy, and delete files from the switch. The file system structure allows you to keep, save, rename, and maintain multiple copies of configuration files on the switch. In addition, you can manage other entities of the switch such as policies and access control lists (ACLs).
- Configuration file management—You can oversee and manage multiple configuration files on your switch. In addition, you can upload, download, modify, and name configuration files used by the switch.
- Process control—You can stop and start processes, restart failed processes, and update the software for a specific process or set of processes.

- Memory protection—With memory protection, ExtremeXOS protects each process from every other process in the system. If one process experiences a memory fault, that process cannot affect the memory space of another process.
- CPU monitoring—You can monitor CPU utilization for Management Switch Fabric Modules (MSMs)/ Management Modules (MMs) and the individual processes running on the switch. Monitoring the workload of the CPU allows you to troubleshoot and identify suspect processes.



Note

Filenames are case-sensitive. For information on filename restrictions, please refer to the specific command in the *ExtremeXOS Command Reference Guide*.

cp

```
cp [internal-memory old_name_internal internal-memory new_name_internal |
internal-memory old_name_internal memorycard new_name_memorycard | memorycard
old_name_memorycard memorycard new_name_memorycard | memorycard
old_name_memorycard new_name | old_name memorycard new_name_memorycard | old_name
new_name]
```

Description

Copies an existing configuration, policy, or if configured, core dump file stored in the system.

Syntax Description

internal-memory	Specifies the internal memory card.
<i>old_name_internal</i>	Specifies the name of the core dump file located on the internal memory card that you want to copy.
<i>new_name_internal</i>	Specifies the name of the newly copied core dump file located on the internal memory card.
memorycard	Specifies a removable storage device, which can be a compact flash card or a USB 2.0 storage device. NOTE: This parameter is available for compact flash cards on BlackDiamond 8800 series switches and for USB 2.0 storage devices on Summit X460, X480, X670, X670V, and X770 switches.
<i>old_name_memorycard</i>	Specifies the name of the file you want to copy from a removable storage device. NOTE: This parameter is available for compact flash cards on BlackDiamond 8800 series switches and for USB 2.0 storage devices on Summit X460, X480, X670, X670V, and X770 switches.
<i>new_name_memorycard</i>	Specifies the name of the newly copied file located on a removable storage device. NOTE: This parameter is available for compact flash cards on BlackDiamond 8800 series switches and for USB 2.0 storage devices on Summit X460, X480, X670, X670V, and X770 switches.
<i>old_name</i>	Specifies the name of the configuration or policy file that you want to copy.
<i>new_name</i>	Specifies the name of the newly copied configuration or policy file.



Default

N/A.

Usage Guidelines

Use this command to make a copy of an existing file before you alter or edit the file. By making a copy, you can easily go back to the original file if needed.

When you copy a configuration or policy file, remember the following:

- XML-formatted configuration files have a .cfg file extension. The switch only runs .cfg files.
- ASCII-formatted configuration files have a .xsf file extension. See [Software Upgrade and Boot Options](#) in the *ExtremeXOS Concepts Guide* for more information.
- Policy files have a .pol file extension.
- Core dump files have a .gz file extension. See [Internal Memory and Core Dump Files](#) below.

When you copy a configuration or policy file from the system, make sure you specify the appropriate file extension. For example, when you want to copy a policy file, specify the filename and .pol.

When you copy a file on the switch, the switch displays a message similar to the following:

```
Copy config test.cfg to config test1.cfg on switch? (y/n)
```

Enter y to copy the file. Enter n to cancel this process and not copy the file.

When you enter y, the switch copies the file with the new name and keeps a backup of the original file with the original name. After the switch copies the file, use the `ls` command to display a complete list of files. In this example, the switch displays the original file named test.cfg and the copied file named test_rev2.cfg.

The following is sample output from the `ls` command:

```
...
-rw-r--r--  1 root    root      100980 Sep 23 09:16 test.cfg
-rw-r--r--  1 root    root      100980 Oct 13 08:47 test_rev2.cfg
...
```

When you enter n, the switch displays a message similar to the following:

```
Copy cancelled.
```

For the memorycard option, the source and/or destination is a compact flash card or USB 2.0 storage device. You must mount the compact flash card or USB storage device for this operation to succeed. The `cp` command copies a file from the switch to the compact flash card, the USB storage device, or a file already on one of those devices. If you copy a file from the switch to the removable storage device, and the new filename is identical to the source file, you do not need to re-enter the filename.

Case-sensitive Filenames

Filenames are case-sensitive. In this example, you have a configuration file named Test.cfg. If you attempt to copy the file with the incorrect case, for example test.cfg, the switch displays a message similar to the following:



```
Error: cp: /config/test.cfg: No such file or directory
```

Since the switch is unable to locate test.cfg, the file is not copied.

Local Filename Character Restrictions

This section provides information about the characters supported by the switch for local filenames.

When specifying a local filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z).
- Numerals (0-9).
- Period (.).
- Dash (-).
- Underscore (_).

When naming a local file, remember the requirements listed above.

Internal Memory and Core Dump Files

Core dump files have a .gz file extension. The filename format is: `core.process-name.pid.gz` where *process-name* indicates the name of the process that failed and *pid* is the numerical identifier of that process. If you have a BlackDiamond 8800 series switch and save core dump files to the compact flash card, the filename also includes the affected MSM/MM: MSM-A or MSM-B.

By making a copy of a core dump file, you can easily compare new debug information with the old file if needed.

When you configure and enable the switch to send core dump (debug) information to the internal memory card, specify the internal-memory option and associated internal-memory name options to copy an existing core dump file.

If you have an external compact flash card or a USB 2.0 storage device installed, you can copy the core dump file to that location. When you send core dump information to a removable storage device, specify the memorycard option and associated memorycard name options to copy an existing core dump file.

For information about configuring and sending core dump information to the internal memory card, see the `configure debug core-dumps` and `save debug tracefiles memorycard` commands.

For more detailed information about core dump files, see [Troubleshooting](#) in the *ExtremeXOS Concepts Guide*.

Modular Switches Only

This command also replicates the action from the primary MSM/MM to the backup MSM/MM. For example, when you copy a file on the primary MSM, the same file is copied to the backup MSM/MM.



Example

The following command makes a copy of a configuration file named test.cfg and gives the copied file a new name of test_rev2.cfg:

```
cp test.cfg test_rev2.cfg
```

The following command makes a copy of a configuration file named primary.cfg on the switch and stores the copy on the removable storage device with the same name, primary.cfg:

```
cp primary.cfg memorycard
```

The above command performs the same action as entering the following command:

```
cp primary.cfg memorycard primary.cfg
```

History

This command was first available in ExtremeXOS 11.0.

The memorycard option was added in ExtremeXOS 11.1.

The internal-memory option was added in ExtremeXOS 11.4.

Support for USB 2.0 storage devices was added in ExtremeXOS 12.5.3.

Platform Availability

This command is available on all platforms.

disable xml-mode

disable xml-mode

Description

Disables XML configuration mode on the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.



Usage Guidelines

Use this command to disable the XML configuration mode on the switch. XML configuration mode is not supported for end users.

See the command:

```
enable xml-mode
```

Example

The following command disables XML configuration mode on the switch:

```
disable xml-mode
```

History

This command was first available in an ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable xml-mode

```
enable xml-mode
```

Description

Enables XML configuration mode on the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables the XML configuration mode on the switch, however XML configuration mode is not supported for end users, and Extreme Networks strongly cautions you not to enable this mode. Use this command only under the direction of Extreme Networks.

If you inadvertently issue this command, the switch prompt will be changed by adding the text (xml) to the front of the prompt. If you see this mode indicator, please disable XML configuration mode by using the following command:




```
disable xml-mode
```

Example

The following command enables XML configuration mode on the switch:

```
enable xml-mode
```

History

This command was first available in an ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

ls

```
ls {[internal-memory | memorycard]} {file_name}
```

Description

Lists all configuration, policy, and if configured, core dump files in the system.

Syntax Description

internal-memory	Lists the core dump (debug) files that are present and saved in the internal memory card.
memorycard	Lists all of the files on a removable storage device, which can be a compact flash card or a USB 2.0 storage device. NOTE: This parameter is available for compact flash cards on BlackDiamond 8800 series switches and for USB 2.0 storage devices on Summit X460, X480, X670, X670V, and X770 switches.
<i>file_name</i>	Lists all the files that match the wildcard.

Default

N/A.

Usage Guidelines

When you use issue this command without any options, the output displays all of the configuration and policy files stored on the switch.



When you configure and enable the switch to send core dump (debug) information to the internal memory card, specify the `internal-memory` option to display the core dump files stored on the internal memory card. For more information about core dump files, see [Core Dump Files](#).

When you specify the `memorycard` option on a switch, the output displays all of the files stored on the removable storage device, including core dump files if so configured. For more information about core dump files, see [Core Dump Files](#).

When you specify the `file-name` option, the output displays all of the files that fit the wildcard criteria.

Understanding the Output

Output from this command includes the following:

- The first column displays the file permission using the following ten place holders:
 - The first place holder displays - for a file.
 - The next three place holders display r for read access and w for write access permission for the file owner.
 - The following three place holders display r for read access permission for members of the file owner's group.
 - The last three place holders display r for read access for every user that is not a member of the file owner's group.
- The second column displays how many links the file has to other files or directories.
- The third column displays the file owner.
- The remaining columns display the file size, date and time the file was last modified, and the file name.

Core Dump Files

Core dump files have a `.gz` file extension. The filename format is: `core.process-name.pid.gz` where `process-name` indicates the name of the process that failed and `pid` is the numerical identifier of that process. If you have a BlackDiamond 8800 series switch and save core dump files to the compact flash card, the filename also includes the affected MSM/MM: MSM-A or MSM-B.

When the switch has not saved any debug files, no files are displayed. For information about configuring and sending core dump information to the internal memory card, compact flash card, or USB 2.0 storage device, see the `configure debug core-dumps` and `save debug tracefiles memorycard` commands.

For more detailed information about core dump files, see [Troubleshooting](#) in the *ExtremeXOS Concepts Guide*.

Example

The following command displays a list of all current configuration and policy files in the system:

```
ls
```



The following is sample output from this command:

```
total 424
-rw-r--r-- 1 root root 50 Jul 30 14:19 hugh.pol
-rw-r--r-- 1 root root 94256 Jul 23 14:26 hughtest.cfg
-rw-r--r-- 1 root root 100980 Sep 23 09:16 megtest.cfg
-rw-r--r-- 1 root root 35 Jun 29 06:42 newpolicy.pol
-rw-r--r-- 1 root root 100980 Sep 23 09:17 primary.cfg
-rw-r--r-- 1 root root 94256 Jun 30 17:10 roytest.cfg
```

The following command displays a list of all current configuration and policy files on a removable storage device:

```
ls memorycard
```

The following is sample output from this command:

```
-rwxr-xr-x 1 root 0 15401865 Mar 30 00:03 bd10K-11.2.0.13.xos
-rwxr-xr-x 1 root 0 10 Mar 31 09:41 test-1.pol
-rwxr-xr-x 1 root 0 10 Apr 4 09:15 test.pol
-rwxr-xr-x 1 root 0 10 Mar 31 09:41 test_1.pol
-rwxr-xr-x 1 root 0 223599 Mar 31 10:02 v11_1_3.cfg
```

The following command displays a list of all configuration and policy files with a filename beginning with the letter “a:”

```
(debug) BD-12804.1 # ls a*
```

Following is sample output from this command:

```
-rw-r--r-- 1 root 0 2062 Jan 6 09:11 abc
-rw-rw-rw- 1 root 0 1922 Jan 7 02:19 abc.xsf
1k-blocks Used Available Use%
16384 496 15888 3%
```

The following command displays a list of all .tgz files:

```
(debug) BD-12804.24 # ls internal-memory *.tgz
```

Following is sample output from this command:

```
-rwxr-xr-x 1 root 0 79076 Jan 6 09:47 old_traces.tgz
1k-blocks Used Available Use%
49038 110 48928 0%
```



History

This command was first available in ExtremeXOS 10.1.

The memorycard option was added in ExtremeXOS 11.0.

The internal-memory option was added in ExtremeXOS 11.4.

The file-name option was added in ExtremeXOS 12.2.

Support for USB 2.0 storage devices was added in ExtremeXOS 12.5.3.

Platform Availability

This command is available on all platforms.

mv

```
mv [internal-memory old_name_internal internal-memory new_name_internal |
internal-memory old_name_internal memorycard new_name_memorycard | memorycard
old_name_memorycard memorycard new_name_memorycard | memorycard
new_name_memorycard new-name | old_name memorycard new_name_memorycard | old_name
new_name]
```

Description

Moves or renames an existing configuration, policy, or if configured, core dump file in the system.

Syntax Description

internal-memory	Specifies the internal memory card.
<i>old_name_internal</i>	Specifies the current name of the core dump file located on the internal memory card.
<i>new_name_internal</i>	Specifies the new name of the core dump file located on the internal memory card.
memorycard	Specifies a removable storage device, which can be a compact flash card or a USB 2.0 storage device. NOTE: This parameter is available for compact flash cards on BlackDiamond 8800 series switches and for USB 2.0 storage devices on Summit X460, X480, X670, X670V, and X770 switches.
<i>old_name_memorycard</i>	Specifies the current name of the file located on a removable storage device. Depending on your switch configuration, you can have configuration, policy, or cord dump files stored in this card. NOTE: This parameter is available for compact flash cards on BlackDiamond 8800 series switches and for USB 2.0 storage devices on Summit X460, X480, X670, X670V, and X770 switches.
<i>new_name_memorycard</i>	Specifies the new name of the file located a removable storage device. NOTE: This parameter is available for compact flash cards on BlackDiamond 8800 series switches and for USB 2.0 storage devices on Summit X460, X480, X670, X670V, and X770 switches.



<code>old_name</code>	Specifies the current name of the configuration or policy file on the system.
<code>new_name</code>	Specifies the new name of the configuration or policy file on the system.

Default

N/A.

Usage Guidelines

When you rename a file with a given extension, remember the following:

- XML-formatted configuration files have the .cfg file extension. The switch only runs .cfg files.
- ASCII-formatted configuration files have the .xsf file extensions. See [Software Upgrade and Boot Options](#) in the *ExtremeXOS Concepts Guide* for more information.
- Policy files have the .pol file extension.
- Core dump files have the .gz file extension. See [Internal Memory and Core Dump Files](#) for more information.

Make sure the renamed file uses the same file extension as the original file. If you change the file extensions, the file may be unrecognized by the system. For example, if you have an existing configuration file named test.cfg, the new filename must include the .cfg file extension.

You cannot rename an active configuration file (the configuration currently selected to boot the switch). To verify the configuration that you are currently using, issue the `show switch {detail}` command. If you attempt to rename the active configuration file, the switch displays a message similar to the following:

```
Error: Cannot rename current selected active configuration file.
```

When you rename a file, the switch displays a message similar to the following:

```
Rename config test.cfg to config megtest.cfg on switch? (y/n)
```

Enter y to rename the file on your system. Enter n to cancel this process and keep the existing filename.

The memorycard option moves files between a removable storage device and the switch. If you use the memorycard option for both the old-name and the new-name, this command just renames a file on the removable storage device.

Case-sensitive Filenames

Filenames are case-sensitive. In this example, you have a configuration file named Test.cfg. If you attempt to rename the file with the incorrect case, for example test.cfg, the switch displays a message similar to the following:

```
Error: mv: unable to rename `/config/test.cfg': No such file or directory
```

Since the switch is unable to locate test.cfg, the file is not renamed.

Local Filename Character Restrictions

This section provides information about the characters supported by the switch for local filenames.



When specifying a local filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z).
- Numerals (0-9).
- Period (.).
- Dash (-).
- Underscore (_).

When naming a local or remote file, remember the requirements listed above.

Internal Memory and Core Dump Files

Core dump files have a .gz file extension. The filename format is: `core.process-name.pid.gz` where process-name indicates the name of the process that failed and pid is the numerical identifier of that process. If you have a BlackDiamond 8800 series switch and save core dump files to the compact flash card, the filename also includes the affected MSM/MM: MSM-A or MSM-B.

When you configure the switch to send core dump (debug) information to the internal memory card, specify the internal-memory option to rename an existing core dump file. If you have a switch with a removable storage device installed, you can move and rename the core dump file to that location.

For information about configuring and sending core dump information to the internal memory card, see the `configure debug core-dumps` and `save debug tracefiles memorycard` commands.

Modular Switches Only

This command also replicates the action from the primary MSM/MM to the backup MSM/MM. For example, when you rename a file on the primary MSM/MM, the same file on the backup MSM/MM is renamed.

Example

The following command renames the configuration file named Testb91.cfg to Activeb91.cfg:

```
mv Testb91.cfg Activeb91.cfg
```

On a switch with a removable storage device installed, the following command moves the configuration file named test1.cfg from the switch to the removable storage device:

```
mv test1.cfg memorycard test1.cfg
```

If you do not change the name of the configuration file, you can also use the following command to move the configuration file test1.cfg from the switch to a removable storage device:

```
mv test1.cfg memorycard
```



On a switch with a removable storage device installed, the following command moves the policy file named `bgp.pol` from the removable storage device to the switch:

```
mv memorycard bgp.pol bgp.pol
```

History

This command was first available in ExtremeXOS 10.1.

Support for replicating information from the primary MSM to the backup MSM was introduced in ExtremeXOS 11.0.

The `memorycard` option was added in ExtremeXOS 11.1.

The `internal-memory` option was added in ExtremeXOS 11.4.

Support for USB 2.0 storage devices was added in ExtremeXOS 12.5.3.

Platform Availability

This command is available on all platforms.

restart process

```
restart process [class cname | name {msm slot}]
```

Description

Terminates and restarts the specified process during a software upgrade on the switch.

Syntax Description

<i>cname</i>	Specifies the name of the process to restart. With this parameter, you can terminate and restart all instances of the process associated with a specific routing protocol on all VRs. You can restart the OSPF routing protocol and associated processes.
<i>name</i>	Specifies the name of the process to terminate and restart. You can use this command with the following processes: <code>bgp</code> (not available on Summit X440-L2 series switches), <code>eapsexsshd</code> (available only when you have installed the SSH module), <code>isis</code> (not available on Summit X440-L2 series switches), <code>lldpnetLoginnetToolsospf</code> (not available on Summit X440-L2 series switches), <code>snmpSubagentsnmpMastertelnetdthttpdftpdvrrp</code> (not available on Summit X440-L2 series switches), and <code>xmld</code> .
<i>slot</i>	Specifies the MSM/MM where the process should be terminated and restarted. A specifies the MSM/MM installed in slot A, and B specifies the MSM/MM installed in slot B. NOTE: This parameter is available only on modular switches.

Default

N/A.



Usage Guidelines

Use this command to terminate and restart a process during a software upgrade on the switch. You have the following options:

- `cname`—Specifies that the software terminates and restarts all instances of the process associated with a specific routing protocol on all VRs.
- `name`—Specifies the name of the process.

Depending on the software version running on your switch and the type of switch you have, you can terminate and restart different or additional processes. To see which processes you can restart during a software upgrade, enter `restart process` followed by [Tab]. The switch displays a list of available processes.

SummitStack Only.

You can issue this command only from the master node. If you issue this command from any other node, the following message appears:

```
Error: This command can only be executed on Master.
```

To display the status of ExtremeXOS processes on the switch, including how many times a process has been restarted, use the `show process {name} {detail} {description} {slotslotid}` command. The following is a truncated sample of the `show process` command on a Summit switch:

Process Name	Version	Restart	State	Start Time
aaa	3.0.0.2	0	Ready	Thu Sep 1 17:00:52 2005
acl	3.0.0.2	0	Ready	Thu Sep 1 17:00:54 2005
bgp	Not Started	0	No license	Not Started
cfgmgr	3.0.0.21	0	Ready	Thu Sep 1 17:00:52 2005
cli	3.0.0.22	0	Ready	Thu Sep 1 17:00:52 2005
devmgr	3.0.0.2	0	Ready	Thu Sep 1 17:00:52 2005
dirser	3.0.0.2	0	Ready	Thu Sep 1 17:00:51 2005
dosprotect	3.0.0.1	0	Ready	Thu Sep 1 17:00:56 2005
eaps	3.0.0.8	0	Ready	Thu Sep 1 17:00:53 2005
...				

You can also use the `restart process` command when upgrading a software modular package. For more information, see the section [Upgrade a Modular Software Package](#) in *Software Upgrade and Boot Options* of the *ExtremeXOS Concepts Guide*.

Example

The following command stops and restarts the process `tftpd` during a software upgrade:

```
restart process tftpd
```



The following command stops and restarts all instances of the OSPF routing protocol for all VRs during a software upgrade:

```
restart process class ospf
```

History

This command was first available in ExtremeXOS 11.3.

Support for restarting the Link Layer Discovery Protocol (lldp), Open Shortest Path First (ospf), and network login (netLogin) processes was added in ExtremeXOS 11.3.

Support for Border Gateway Protocol (bgp) and Ethernet Automatic Protection Switching (eaps) was added in ExtremeXOS 11.4.

Support for MultiProtocol Label Switching (mpls) and Virtual Router Redundancy Protocol (vrrp) was added in ExtremeXOS 11.6.

Support for netTools was added in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

rm

```
rm {internal-memory | memorycard} file_name
```

Description

Removes/deletes an existing configuration, policy, or if configured, core dump file from the system.

Syntax Description

internal-memory	Specifies the internal memory card.
memorycard	Specifies a removable storage device, which can be a compact flash card or a USB 2.0 storage device. NOTE: This parameter is available for compact flash cards on BlackDiamond 8800 series switches and for USB 2.0 storage devices on Summit X460, X480, X670, X670V, and X770 switches.
<i>file_name</i>	Specifies the name of the configuration, policy file, or if configured, the core dump file.

Default

N/A.



Usage Guidelines

After you remove a configuration or policy file from the system, that file is unavailable to the system. For information about core dump files, see [Internal Memory Card and Core Dump Files](#).

You cannot remove an active configuration file (the configuration currently selected to boot the switch). To verify the configuration that you are currently using, issue the `show switch {detail}` command. If you attempt to remove the active configuration file, the switch displays a message similar to the following:

```
Error: Cannot remove current selected active configuration file.
```

When you delete a file from the switch, a message similar to the following appears:

```
Remove testpolicy.pol from switch? (y/n)
```

Enter y to remove the file from your system. Enter n to cancel the process and keep the file on your system.

Case-sensitive Filenames

Filenames are case-sensitive. In this example, you have a configuration file named Test.cfg. If you attempt to remove a file with the incorrect case, for example test.cfg, the system is unable to remove the file. The switch does not display an error message; however, the `ls` command continues to display the file Test.cfg. To remove the file, make sure you use the appropriate case.

Local Filename Character Restrictions

This section provides information about the characters supported by the switch for local filenames.

When specifying a local filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z).
- Numerals (0-9).
- Period (.).
- Dash (-).
- Underscore (_).

When naming a local or remote file, remember the requirements listed above.

The `memorycard` option removes/deletes an existing file on the removable storage device, including core dump files if configured. See the section [Internal Memory Card and Core Dump Files](#) for information about core dump files.

Internal Memory Card and Core Dump Files

When you delete a core dump file from the system, that file is unavailable.

When you configure the switch to send core dump (debug) information to the internal memory card, specify the `internal-memory` option to remove/delete the specified core dump file.

For information about configuring and sending core dump information to the internal memory card, see the `configure debug core-dumps` and `save debug tracefiles memorycard` commands.



You can use the * wildcard to delete core dump files from the internal memory card. You can also use the * wildcard to delete all of a particular file type from a removable storage device. Currently running and in-use files are not deleted.

If you configure the switch to write core dump files to the internal memory card and attempt to download a new software image, you might have insufficient space to complete the image download. When this occurs, you must decide whether to continue the software download or move or delete the core dump files from the internal memory. For example, if you have a switch with a removable storage device installed with space available, transfer the files to the storage device. Another option is to transfer the files from the internal memory card to a TFTP server. This frees up space on the internal memory card while keeping the core dump files.

Modular Switches Only

This command also replicates the action from the primary MSM/MM to the backup MSM/MM. For example, when you delete a file on the primary MSM/MM, the same file on the backup MSM/MM is deleted.

Example

The following command removes the configuration file named Activeb91.cfg from the system:

```
rm Activeb91.cfg
```

The following command removes all of the core dump files stored on the internal memory card:

```
rm internal-memory *
```

On a switch with a removable storage device installed, the following command removes the policy file named test.pol from the removable storage device:

```
rm memorycard test.pol
```

On a switch with a removable storage device installed, the following command removes all of the configuration files from the removable storage device:

```
rm memorycard *.cfg
```

History

This command was first available in ExtremeXOS 10.1.

Support for replicating information from the primary MSM to the backup MSM was introduced in ExtremeXOS 11.0.

The memorycard option was added in ExtremeXOS 11.1.



The internal-memory option was added in ExtremeXOS 11.4.

Support for USB 2.0 storage devices was added in ExtremeXOS 12.5.3.

Platform Availability

This command is available on all platforms.

show heartbeat process

```
show heartbeat process {name}
```

Description

Displays the health of the ExtremeXOS processes.

Syntax Description

<i>name</i>	Specifies the name of the process.
-------------	------------------------------------

Default

N/A.

Usage Guidelines

The software monitors all of the XOS processes running on the switch. This process monitor creates and terminates XOS processes on demand (for example, when you log in or log out of the switch) and restarts processes if an abnormal termination occurs (for example, if your system crashes). The process monitor also ensures that only version-compatible processes and processes with proper licenses are started.

The `show heartbeat process` command is a resource for providing background system health information because you can view the health of ExtremeXOS processes on the switch.

Use this command to monitor the health of the XOS processes. The switch uses two algorithms to collect process health information: polling and reporting. Both polling and reporting measure the heartbeat of the process. Polling occurs when a HELLO message is sent and a HELLO_ACK message is received. The two counts are the same. Reporting occurs when a HELLO_ACK message is sent only. Therefore, no HELLO messages are sent and the HELLO count remains at zero.

The `show heartbeat process` command displays the following information in a tabular format:

- Card—The name of the module where the process is running (modular switches only).
- Process Name—The name of the process.
- Hello—The number of hello messages sent to the process.



- HelloAck—The number of hello acknowledgement messages received by the process manager.
- Last Heartbeat Time—The timestamp of the last health check received by the process manager. (Unknown specifies kernel modules and they do not participate in heartbeat monitoring.)

This status information may be useful for your technical support representative if you have a network problem.

You may find it useful to capture the process information under normal operating conditions to establish a baseline. By having a baseline, if you experience a problem, you and your technical support representative can more easily identify the problem.

Example

To display the health of all processes on your system, use the following command:

```
show heartbeat process
```

The following is sample output from a modular switch:

Card	Process Name	Hello	HelloAck	Last Heartbeat Time
MSM-A	aaa	0	180324	Wed Dec 10 15:06:04 2003
MSM-A	acl	36069	36069	Wed Dec 10 15:05:57 2003
MSM-A	bgp	0	180348	Wed Dec 10 15:06:05 2003
MSM-A	cfgmgr	72139	72139	Wed Dec 10 15:06:02 2003
MSM-A	cli	60116	60116	Wed Dec 10 15:06:03 2003
MSM-A	devmgr	0	180339	Wed Dec 10 15:06:03 2003
MSM-A	dirser	0	180324	Wed Dec 10 15:06:03 2003
MSM-A	edp	36069	36069	Wed Dec 10 15:05:57 2003
MSM-A	ems	45087	45087	Wed Dec 10 15:06:03 2003
MSM-A	epm	0	0	Unknown
MSM-A	exacl	0	0	Unknown
			

The following is sample output from a Summit switch:

Process Name	Hello	HelloAck	Last Heartbeat Time
aaa	0	254328	Tue Feb 10 05:21:46 2004
acl	50867	50867	Tue Feb 10 05:21:43 2004
bgp	0	0	Wed Feb 4 08:03:18 2004
cfgmgr	25433	25433	Tue Feb 10 05:21:33 2004
cli	84779	84779	Tue Feb 10 05:21:47 2004
cna	20234	20234	Mon Feb 9 00:28:35 2004
devmgr	0	250507	Tue Feb 10 05:21:47 2004
dirser	0	254336	Wed Feb 4 08:03:18 2004
dosprotect	0	254335	Tue Feb 10 05:21:47 2004
eaps	0	254336	Tue Feb 10 05:21:48 2004
edp	50867	50867	Tue Feb 10 05:21:44 2004
elrp	50867	50867	Tue Feb 10 05:21:43 2004
ems	63584	63584	Tue Feb 10 05:21:44 2004
epm	0	0	Wed Feb 4 08:03:18 2004



```
esrp          50867    50867    Tue Feb 10 05:21:46 2004
...
```

To display the health of the STP process on your system, use the following command:

```
show heartbeat process stp
```

The following is sample output from a modular switch:

```
Card Process Name      Hello HelloAck      Last Heartbeat Time
-----
MSM-A stp              34921    34921    Wed Dec 10 11:54:37 2003
```

The following is sample output from the Summit switch:

```
Process Name      Hello HelloAck      Last Heartbeat Time
-----
stp              50870    50870    Tue Feb 10 05:22:13 2004
```

History

This command was first available in an ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show memory

```
show memory {slot [slotid | a | b]}
```

Description

Displays the current system memory information.

Syntax Description

slot a	Specifies the MSM module installed in slot A. NOTE: This parameter is available only on modular switches.
slot b	Specifies the MSM module installed in slot B. NOTE: This parameter is available only on modular switches.
slotid	Specifies slot number for the node in a stack. The value can be from 1 to 8.



Default

N/A.

Usage Guidelines

Viewing statistics on a regular basis allows you to see how well your network is performing. When you keep simple daily records, you see trends emerging and notice problems arising before they cause major network faults. This way, statistics can help you get the best out of your network.

This information may be useful for your technical support representative if you experience a problem.

Depending on the software version running on your switch or your switch model, additional or different memory information might be displayed.

You can also use the `show memory process name {slotslotid}` command to view the system memory and the memory used by the individual processes.

SummitStack Only

When you issue the command with out any parameters:

- From the stack manager or backup node, the stack displays the current system memory information for the master node and the back-up node in the Active Topology.
- From a standby node, the stack displays the current system memory information for the master node and the standby node in the Active Topology.

Modular Switches Only

When you issue the command without any parameters, the switch displays information about all of the MSMs/MMs installed in your system.

Reading the Output

The `show memory` command displays the following information in a tabular format:

- System memory information (both total and free).
- Current memory used by the individual processes.

The current memory statistics for the individual process also includes the following:

- The module (MSM A or MSM B) and the slot number of the MSM (modular switches only).
- The name of the process.

In general, the free memory count for an MSM/MM or Summit family switch decreases when one or more running processes experiences an increase in memory usage.

If you observe a continuous decrease in the free memory over an extended period of time, and you have not altered your switch configuration, please contact Extreme Networks Technical Support.



Example

The following command displays current system memory information for the MSM installed in slot A of a modular switch:

```
show memory slot a
```

The following is sample output from this command:

```
System Memory Information
-----
MSM-A    Total DRAM (KB): 524288
MSM-A    System      (KB): 45912
MSM-A    User        (KB): 102264
MSM-A    Free        (KB): 376112
Memory Utilization Statistics
-----
Card Slot Process Name      Memory (KB)
-----
MSM-A  9    aaa                7772
MSM-A  9    acl                6716
MSM-A  9    bgp               16708
MSM-A  9    cfmgr            3484
MSM-A  9    cli              33964
MSM-A  9    devmgr           3656
MSM-A  9    dirser           3072
MSM-A  9    eaps             9136
MSM-A  9    edp              4644
MSM-A  9    elrp             4608
MSM-A  9    ems              5832
MSM-A  9    epm              8084
MSM-A  9    esrp            11004
MSM-A  9    etmon           11356
MSM-A  9    exacl            13
MSM-A  9    exosmc           22
MSM-A  9    exosq            29
MSM-A  9    exsflow          8
MSM-A  9    exsnoop          15
MSM-A  9    exvlan           252
MSM-A  9    fdb              8760
MSM-A  9    hal              22624
MSM-A  9    mcmgr            13128
MSM-A  9    msgsrv           2972
MSM-A  9    netLogin         4564
MSM-A  9    netTools         4696
MSM-A  9    nettx            56
MSM-A  9    nodemgr          5388
MSM-A  9    ospf            12476
MSM-A  9    pim              10012
MSM-A  9    polMgr           3272
MSM-A  9    rip              10392
MSM-A  9    rtmgr            9748
MSM-A  9    snmpMaster       6400
MSM-A  9    snmpSubagent    8104
MSM-A  9    stp              6896
MSM-A  9    telnetd         3236
```




```

MSM-A 9  tftpd          3080
MSM-A 9  vlan           5816
MSM-A 9  vrrp           6584

```

The following command displays current system memory information for a Summit family switch:

```
show memory
```

The following is sample output from this command:

```

System Memory Information
-----
Total DRAM (KB): 262144
System      (KB): 25852
User       (KB): 96608
Free      (KB): 139684
Memory Utilization Statistics
-----
Process Name      Memory (KB)
-----
aaa                13468
acl                11420
bgp                 0
cfgmgr            8336
cli               41040
cna                 0
devmgr            8452
dirser            7068
dosprotect        8264
eaps              18784
edp                9780
elrp              10040
ems               10672
epm               15520
esrp              16728
etmon             18924
exacl              30
exdos              8
exfib              3
exosmc             29
exosnvram          4
exosq              36
exsflow           10
exsnoop           20
exsshd            9272
exvlan            290
fdb               12908
hal               64768
lldp              8816
mcmgr             17836
msgsrv            6960
netLogin          8924
netTools          11524
nettx              70
nodemgr           9636

```



```

ospf          18124
ospfv3        0
pim           15996
poe           8936
polMgr        7576
rip           17736
ripng         0
rtmgr         16016
snmpMaster    15416
snmpSubagent  26428
stp           10768
telnetd       8464
tftpd         7584
thttpd        11344
vlan          9660
vrrp          11184
xmld          9148

```

The following command displays current system memory information for a stack, where slot 1 is the master and slot 6 is the backup:

```

Slot-1 stacK.3 # show memory
System Memory Information
-----
Slot-1   Total DRAM (KB): 262144
Slot-1   System      (KB): 25476
Slot-1   User        (KB): 132256
Slot-1   Free         (KB): 104412
Slot-6   Total DRAM (KB): 262144
Slot-6   System      (KB): 25476
Slot-6   User        (KB): 122820
Slot-6   Free         (KB): 113848
Memory Utilization Statistics
-----
Card Slot Process Name      Memory (KB)
-----
Slot-1 1   aaa                2548
Slot-1 1   acl                2960
Slot-1 1   bgp                 0
Slot-1 1   brm                2428
Slot-1 1   cfgmgr             3256
Slot-1 1   cli               16932
Slot-1 1   devmgr             2708
Slot-1 1   dirser            1916
Slot-1 1   dosprotect         1972
Slot-1 1   eaps               6976
Slot-1 1   edp               2656
Slot-1 1   elrp              2640
Slot-1 1   elsm              2592
Slot-1 1   ems               2764
Slot-1 1   epm               3092
Slot-1 1   esrp              2844
Slot-1 1   etmon            16264
...
Slot-6 6   aaa                2440
Slot-6 6   acl                2872
Slot-6 6   bgp                 0

```



```

Slot-6 6   brm           2396
Slot-6 6   cfgmgr        2776
Slot-6 6   cli           16292
Slot-6 6   devmgr        2672
Slot-6 6   dirser        1836
Slot-6 6   dosprotect    1944
Slot-6 6   eaps          6924
Slot-6 6   edp           2624
Slot-6 6   elrp          2628
Slot-6 6   elsm          2564
Slot-6 6   ems           2744
Slot-6 6   epm           2976
Slot-6 6   esrp          2792
Slot-6 6   etmon         10068
...

```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show memory process

```
show memory process name {slot slotid}
```

Description

Displays the current system memory and that of the specified process.

Syntax Description

<i>name</i>	Specifies the name of the process.
<i>slotid</i>	In a modular switch: Specifies the slot number of the MSM/MM module:A specifies the MSM installed in slot A.B specifies the MSM installed in slot B. In a SummitStack, slotid specifies the slot number of the node in the stack topology. The value can be from 1 to 8.

Default

N/A.



Usage Guidelines

Viewing statistics on a regular basis allows you to see how well your network is performing. When you keep simple daily records, you see trends emerging and notice problems arising before they cause major network faults. This way, statistics can help you get the best out of your network.

This information may be useful for your technical support representative if you experience a problem.

Depending on the software version running on your switch or your switch model, additional or different memory information might be displayed.

You can also use the `show memory {slot [slotid | a | b]}` command to view the system memory and the memory used by the individual processes, even for all processes on all MSMs/MMs installed in modular switches.

SummitStack Only

When you issue the command with out any parameters:

- From the stack manager or backup node, the stack displays current system memory and that of the specified process running on the master node and the back-up node in the Active Topology.
- From a standby node, the stack displays current system memory and that of the specified process running on the master node and the standby node in the Active Topology.

Reading the Output

The `show memory process` command displays the following information in a tabular format:

- System memory information (both total and free).
- Current memory used by the individual processes.

The current memory statistics for the individual process also includes the following:

- The module (MSM A or MSM B) and the slot number of the MSM/MM (modular switches only).
- The name of the process.

Example

The following command displays system memory and VRRP memory usage:

```
show memory process vrrp
```

The following is sample output from a modular switch:

```
System Memory Information
-----
MSM-A    Total (KB): 512508 KB
MSM-A    Free  (KB): 395796 KB
Memory Utilization Statistics
-----
Card Slot Process Name      Memory (KB)
```



```
-----
MSM-A 9      vrrp          6596
```

The following is sample output from a Summit switch:

```
System Memory Information
-----
Total DRAM (KB): 262144
System      (KB): 25852
User        (KB): 96608
Free        (KB): 139684
Memory Utilization Statistics
-----
Process Name      Memory (KB)
-----
vrrp              11184
```

The following is sample output from a SummitStack:

```
Slot-1 stack.4 # show memory process "aaa"
System Memory Information
-----
Slot-1  Total DRAM (KB): 262144
Slot-1  System      (KB): 25476
Slot-1  User        (KB): 132276
Slot-1  Free         (KB): 104392
Slot-6  Total DRAM (KB): 262144
Slot-6  System      (KB): 25476
Slot-6  User        (KB): 122820
Slot-6  Free         (KB): 113848
```

Memory Utilization Statistics

```
-----
Card Slot Process Name      Memory (KB)
-----
Slot-1 1    aaa              2548
Slot-6 6    aaa              2440
```

History

This command was first available in an ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

show process

```
show process {name} {detail} {description} {slot slotid}
```



Description

Displays the status of the ExtremeXOS processes.

Syntax Description

<i>name</i>	Specifies the name of the process.
detail	Specifies more detailed process information.
description	Describes the name of all of the processes or the specified process running on the switch.
<i>slotid</i>	Specifies the slot number of the MSM/MM module:A specifies the MSM installed in slot A.B specifies the MSM installed in slot B. NOTE: This parameter is available only on modular switches. In a Summit stack, slotid specifies the slot number of a node in a stack topology.

Default

N/A.

Usage Guidelines

The ExtremeXOS process manager monitors all of the XOS processes. The process manager also ensures that only version-compatible processes are started.

Using this command without the optional keywords displays summary process information. When you specify the slot keyword, summary information is displayed for that particular slot only.

The `show process` and `show process slot slotid` commands display the following information in a tabular format:

- Card—The name of the module where the process is running (modular switches only).
- Process Name—The name of the process.
- Version—The version number of the process. Options are:
 - Version number—A series of numbers that identify the version number of the process. This is helpful to ensure that you have version-compatible processes and if you experience a problem.
 - Not Started—The process has not been started. This can be caused by not having the appropriate license or for not starting the process.
- Restart—The number of times the process has been restarted. This number increments by one each time a process stops and restarts.
- State—The current state of the process. Options are:
 - No License—The process requires a license level that you do not have. For example, you have not upgraded to that license, or the license is not available for your platform.
 - Ready—The process is running.
 - Stopped—The process has been stopped.
- Start Time—The current start time of the process. Options are:
 - Day/Month/Date/Time/Year—The date and time the process began. When a process terminates and restarts, the start time is also updated.



- Not Started—The process has not been started. This can be caused by not having the appropriate license or for not starting the process.

When you specify the detail keyword, more specific and detailed process information is displayed.

The show process detail and show process slot *slotid* detail commands display the following information in a multi-tabular format:

- Detailed process information.
- Memory usage configurations.
- Recovery policies.
- Process statistics.
- Resource usage.

This status information may be useful for your technical support representative if you have a network problem.

Depending on the software version running on your switch or your switch model, additional or different process information might be displayed.

You may find it useful to capture the process information under normal operating conditions to establish a baseline. By having a baseline, if you experience a problem, you and your technical support representative can more easily identify the problem.

SummitStack Only

When you run the command with out any parameters:

- From the stack manager or backup node, the stack displays the status of the ExtremeXOS processes running on the master node and the back-up node in the Active Topology.
- From a standby node, the stack displays the status of the ExtremeXOS processes running on the standby node and the master node in the Active Topology.

Example

To display the processes on your system, use the following command:

```
show process
```

The following is sample output from a modular switch:

Card	Process Name	Version	Restart	State	Start Time
MSM-A	aaa	3.0.0.2	0	Ready	Sat Dec 6 10:54:24 2003
MSM-A	acl	3.0.0.2	0	Ready	Sat Dec 6 10:54:25 2003
MSM-A	bgp	3.0.0.2	0	Ready	Sat Dec 6 10:54:24 2003
MSM-A	cfgmgr	3.0.0.20	0	Ready	Sat Dec 6 10:54:23 2003
MSM-A	cli	3.0.0.21	0	Ready	Sat Dec 6 10:54:23 2003
MSM-A	devmgr	3.0.0.2	0	Ready	Sat Dec 6 10:54:23 2003
MSM-A	dirser	3.0.0.2	0	Ready	Sat Dec 6 10:54:21 2003
MSM-A	edp	3.0.0.2	0	Ready	Sat Dec 6 10:54:24 2003



```

MSM-A ems          3.0.0.2      0    Ready    Sat Dec  6 10:54:23 2003
MSM-A epm          3.0.0.2      0    Ready    Sat Dec  6 10:54:21 2003
MSM-A exacl        3.0.0.2      0    Ready    Sat Dec  6 10:54:23 2003
MSM-A exosmc       3.0.0.2      0    Ready    Sat Dec  6 10:54:23 2003
MSM-A exosq        3.0.0.2      0    Ready    Sat Dec  6 10:54:22 2003
MSM-A exsnoop      3.0.0.2      0    Ready    Sat Dec  6 10:54:23 2003
MSM-A exvlan       3.0.0.2      0    Ready    Sat Dec  6 10:54:22 2003
MSM-A fdb          3.0.0.2      0    Ready    Sat Dec  6 10:54:24 2003
....

```

The following is sample output from a Summit switch:

Process Name	Version	Restart	State	Start Time
aaa	3.0.0.2	0	Ready	Thu Sep 1 17:00:52 2005
acl	3.0.0.2	0	Ready	Thu Sep 1 17:00:54 2005
bgp	Not Started	0	No license	Not Started
cfgmgr	3.0.0.21	0	Ready	Thu Sep 1 17:00:52 2005
cli	3.0.0.22	0	Ready	Thu Sep 1 17:00:52 2005
devmgr	3.0.0.2	0	Ready	Thu Sep 1 17:00:52 2005
dirser	3.0.0.2	0	Ready	Thu Sep 1 17:00:51 2005
dosprotect	3.0.0.1	0	Ready	Thu Sep 1 17:00:56 2005
eaps	3.0.0.8	0	Ready	Thu Sep 1 17:00:53 2005
edp	3.0.0.2	0	Ready	Thu Sep 1 17:00:53 2005
elrp	3.0.0.1	0	Ready	Thu Sep 1 17:00:53 2005
ems	3.0.0.2	0	Ready	Thu Sep 1 17:00:52 2005
epm	3.0.0.4	0	Ready	Thu Sep 1 17:00:49 2005
esrp	3.0.0.4	0	Ready	Thu Sep 1 17:00:53 2005
etmon	1.0.0.1	0	Ready	Thu Sep 1 17:00:55 2005
exacl	3.0.0.2	0	Ready	Thu Sep 1 17:00:50 2005
exdos	3.0.0.2	0	Ready	Thu Sep 1 17:00:50 2005
exfib	1.0.0.2	0	Ready	Thu Sep 1 17:00:51 2005
exosmc	3.0.0.2	0	Ready	Thu Sep 1 17:00:50 2005
exosnvram	3.0.0.3	0	Ready	Thu Sep 1 17:00:50 2005
exosq	3.0.0.2	0	Ready	Thu Sep 1 17:00:50 2005
exsflow	1.0.0.2	0	Ready	Thu Sep 1 17:00:51 2005
exsnoop	3.0.0.2	0	Ready	Thu Sep 1 17:00:50 2005
exvlan	3.0.0.2	0	Ready	Thu Sep 1 17:00:50 2005
fdb	4.0.0.2	0	Ready	Thu Sep 1 17:00:52 2005
hal	3.0.0.2	0	Ready	Thu Sep 1 17:00:52 2005
lACP	3.0.0.1	0	Ready	Thu Sep 1 17:00:53 2005
lldp	1.2.0.0	0	Ready	Thu Sep 1 17:00:53 2005
mcmgr	4.0.0.2	0	Ready	Thu Sep 1 17:00:54 2005
msgsrv	3.0.0.2	0	Ready	Thu Sep 1 17:00:52 2005
netLogin	1.0.0.0	0	Ready	Thu Sep 1 17:00:55 2005
netTools	3.0.0.2	0	Ready	Thu Sep 1 17:00:55 2005
nettx	3.0.0.2	0	Ready	Thu Sep 1 17:00:50 2005
nodemgr	3.0.0.2	0	Ready	Thu Sep 1 17:00:52 2005
ospf	3.0.0.2	0	Ready	Thu Sep 1 17:00:54 2005
ospfv3	Not Started	0	No license	Not Started
pim	3.0.0.2	0	Ready	Thu Sep 1 17:00:54 2005
poe	3.0.0.2	0	Ready	Thu Sep 1 17:00:56 2005
polMgr	3.0.0.2	0	Ready	Thu Sep 1 17:00:54 2005
rip	3.0.0.2	0	Ready	Thu Sep 1 17:00:54 2005
ripng	3.0.0.1	0	Ready	Thu Sep 1 17:00:54 2005
rtmgr	4.0.0.2	0	Ready	Thu Sep 1 17:00:53 2005
snmpMaster	3.0.0.2	0	Ready	Thu Sep 1 17:00:52 2005




```

snmpSubagent    3.0.0.2    0    Ready    Thu Sep 1 17:00:52 2005
stp              3.0.3.1    0    Ready    Thu Sep 1 17:00:53 2005
telnetd         3.0.0.2    0    Ready    Thu Sep 1 17:00:55 2005
tftpd           3.0.0.2    0    Ready    Thu Sep 1 17:00:55 2005
thttpd          3.0.0.1    0    Ready    Thu Sep 1 17:00:55 2005
vlan             3.1.0.2    0    Ready    Thu Sep 1 17:00:52 2005
vrrp            3.0.0.5    0    Ready    Thu Sep 1 17:00:55 2005
xmld            1.0.0.0    0    Ready    Thu Sep 1 17:00:56 2005
    
```

The following example specifies the process aaa along with the detail keyword:

```
show process aaa detail
```

The following is sample output from this command:

```

Name          PID      Path    Type Link Date          Build By
Peer
-----
--
aaa          284     ./aaa   App  Thu Dec 4 13:23:07 PST 2003  release-
manager 2
3
Virtual Router(s):
-----
--
Configuration:
Start Priority  SchedPolicy  Stack  TTY  CoreSize  Heartbeat  StartSeq
-----
--
1           0           0       0    0       0           1       1
Memory Usage Configuration:
Memory(KB) Zones: Green Yellow Orange Red
-----
--
0           0           0       0    0           0
Recovery policies
-----
--
failover-reboot
-----
--
Statistics:
ConnctionLost  Timeout  Start  Restart  Kill  Register  Signal  Hello  Hello
Ack
-----
--
0           0           0       0       0       1       0       0
173199
Memory Zone  Green  Yellow  Orange  Red
-----
--
Green       0       0       0       0
-----
--
Commands:
    
```



```

Start          Stop          Resume          Shutdown          Kill
-----
--
0              0              0              0              0
-----
--
Resource Usage:
UserTime SysTime PageReclaim PageFault Up Since          Up Date  Up
Time
-----
--
2.160000 0.560000    546           966    Sat Dec  6 10:54:24 2003 00/00/04
00:14:02
-----
--
Thread Name          Pid          Tid          Delay  Timeout Count
-----
--
tacThread            0            2051         10      0
radiusThread        0            1026         10      1
main                 0            1024         2        1
-----
--

```

The following example describes the name of all of the processes running on the switch:

```
show process description
```

The following is sample output from this command:

```

Process Name      Description
-----
aaa               Authentication, Authorization, and Accounting Server
acl               Access Control List Manager
bfd               IETF Bidirectional Forwarding Detection
bgp               Border Gateway Protocol
brm               Bandwidth Resource Manager
cfgmgr            Configuration Manager
cli               Cli Manager
devmgr            Device Manager
dirser            Directory Services
dosprotect        protects against Denial of Service attacks
dotlag            IEEE 802.lag; Connectivity Fault Management
eaps              Ethernet Automatic Protection Switching
edp               Extreme Discovery Protocol
elrp              Extreme Loop Recovery Protocol
elms              Extreme Link State Monitor
ems               Event Management System server application
epm               Extreme Process Manager
esrp              Extreme Standby Routing Protocol
ethoam            Ethernet OAM
etmon             Traffic monitoring and sampling utility
exacl             Access Control List Module
exdhcpsnoop       DHCP snooping module
exdos             Detection of potential Denial of Service attacks module

```



exfib	Routing interface to manage missing routes in ASIC
exosipv6	IPv6 Custom Interface Module
exosmc	Multicast Forwarding Module
exosnvram	Interface to non-volatile RAM
exosq	EXOS Queue Module
exsflow	Sflow interface to gather sflow samples
exsnoop	IGMP/MLD Snooping Module
exvlan	Layer 2 configuration module
fdb	Forwarding Data Base Manager
hal	Hardware Abstraction Layer
hclag	Health Check LAG
idMgr	Identity Manager
ipSecurity	IP Security
ipfix	IPFIX Traffic monitoring utility
isis	Intermediate System to Intermediate System Route Protocol
lacp	Link Aggregation Control Protocol
lldp	802.1AB; Station and Media Access Control Connectivity
Discover	
mcmgr	Multicast Cache Manager
mpls	Multi-Protocol Label Switching
msdp	Multicast Source Discovery Protocol
msgsrv	Message Server
netLogin	Network Login includes MAC, Web-Based and 802.1x
authentication	
netTools	Network Tools set includes ping/tracert/bootprelay/
dhcp/dns/sn	
nettx	Layer 2 forwarding engine module
nodemgr	Fault Tolerance Manager
ospf	Open Shortest Path First Routing Protocol
ospfv3	Open Shortest Path First Routing Protocol for IPv6
pim	Protocol Independent Multicast
poe	Power Over Ethernet Manager
polMgr	Policy Manager
rip	Routing Information Protocol
ripng	Routing Information Protocol for IPv6
rtmgr	Route Table Manager
snmpMaster	Simple Network Management Protocol - Master agent
snmpSubagent	Simple Network Management Protocol - Subagent
stp	Spanning Tree Protocol
synce	Synchronous Ethernet
telnetd	Telnet server
tftpd	Tftp server
thttpd	Web Server
upm	Universal Port Manager
vlan	VLAN Manager - L2 Switching application
vmt	Virtual Machine Tracking
vrrp	Virtual Router Redundancy Protocol (RFC 3768)
vsm	Virtual Switch Manager
xmlc	XML Client Manager
xmld	XML server

The following example shows the truncated output for the command on a stack:

```
Slot-1 stacK.7 # show process
Card  Process Name      Version  Restart  State          Start Time
-----
Slot-1 aaa                3.0.0.3    0    Ready          Thu Mar  1 11:29:34 2007
```



```

Slot-1 acl          3.0.0.2    0    Ready    Thu Mar  1 11:29:39 2007
Slot-1 bgp          Not Started 0    No license Not Started
Slot-1 brm          1.0.0.0    0    Ready    Thu Mar  1 11:29:42 2007
Slot-1 cfgmgr       3.0.0.21   0    Ready    Thu Mar  1 11:29:34 2007
Slot-1 cli          3.0.0.22   0    Ready    Thu Mar  1 11:29:34 2007
Slot-1 devmgr       3.0.0.2    0    Ready    Thu Mar  1 11:29:34 2007
Slot-1 dirser       3.0.0.2    0    Ready    Thu Mar  1 11:29:33 2007
Slot-1 dosprotect   3.0.0.1    0    Ready    Thu Mar  1 11:29:41 2007
Slot-1 eaps         3.0.0.8    0    Ready    Thu Mar  1 11:29:36 2007
Slot-1 edp          3.0.0.2    0    Ready    Thu Mar  1 11:29:36 2007
Slot-1 elrp         3.0.0.1    0    Ready    Thu Mar  1 11:29:35 2007
Slot-1 elsm         3.0.0.1    0    Ready    Thu Mar  1 11:29:35 2007
Slot-1 ems          3.0.0.2    0    Ready    Thu Mar  1 11:29:34 2007
Slot-1 epm          3.0.0.4    0    Ready    Thu Mar  1 11:29:30 2007
Slot-1 esrp         3.0.0.4    0    Ready    Thu Mar  1 11:29:37 2007
Slot-1 etmon        1.0.0.1    0    Ready    Thu Mar  1 11:29:40 2007
...
...
Slot-6 aaa          3.0.0.3    0    Ready    Thu Mar  1 11:29:22 2007
Slot-6 acl          3.0.0.2    0    Ready    Thu Mar  1 11:29:24 2007
Slot-6 bgp          Not Started 0    No license Not Started
Slot-6 brm          1.0.0.0    0    Ready    Thu Mar  1 11:29:26 2007
Slot-6 cfgmgr       3.0.0.21   0    Ready    Thu Mar  1 11:29:21 2007
Slot-6 cli          3.0.0.22   0    Ready    Thu Mar  1 11:29:21 2007
Slot-6 devmgr       3.0.0.2    0    Ready    Thu Mar  1 11:29:21 2007
Slot-6 dirser       3.0.0.2    0    Ready    Thu Mar  1 11:29:20 2007
Slot-6 dosprotect   3.0.0.1    0    Ready    Thu Mar  1 11:29:26 2007
Slot-6 eaps         3.0.0.8    0    Ready    Thu Mar  1 11:29:23 2007
Slot-6 edp          3.0.0.2    0    Ready    Thu Mar  1 11:29:23 2007
Slot-6 elrp         3.0.0.1    0    Ready    Thu Mar  1 11:29:22 2007
Slot-6 elsm         3.0.0.1    0    Ready    Thu Mar  1 11:29:22 2007
Slot-6 ems          3.0.0.2    0    Ready    Thu Mar  1 11:29:21 2007
Slot-6 epm          3.0.0.4    0    Ready    Thu Mar  1 11:29:19 2007
Slot-6 esrp         3.0.0.4    0    Ready    Thu Mar  1 11:29:23 2007
Slot-6 etmon        1.0.0.1    0    Ready    Thu Mar  1 11:29:25 2007
...

```

History

This command was first available in an ExtremeXOS 10.1.

The description keyword was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

start process

```
start process name {msm slot}
```



Description

Starts the specified process on the switch. Used to restart a process after it has been terminated.

Syntax Description

<i>name</i>	Specifies the name of the process to start. You can start the following processes: bgp (not available on Summit X440-L2 series switches) eapsexsshd (available only when you have installed the SSH module) isis (not available on Summit X440-L2 series switches) lldpnetLoginnetToolsospf (not available on Summit X440-L2 series switches) snmpMastersnmpSubagenttelnetdthttpdftpdvrrp (not available on Summit Summit X440-L2 series switches) xmlid
<i>slot</i>	Specifies the MSM/MM where the process should be started. A specifies the MSM installed in slot A, and B specifies the MSM installed in slot B. NOTE: This parameter is available only on modular switches.

Default

N/A.

Usage Guidelines

Use this command after you have stopped a process and you want to restart it. To stop a process, use the `terminate process` command.

You are unable to start a process that is already running. If you try to start a currently running process, an error message similar to the following appears:

Error: Process telnetd already exists!

Depending on the software version running on your switch and the type of switch you have, you can restart different or additional processes. To see which processes you can restart, enter `start process` followed by [Tab]. The switch displays a list of available processes.

To display the status of ExtremeXOS processes on the switch, including how many times a process has been restarted, use the `show process {name} {detail} {description} {slotslotid}` command.

You can also use the `start process` command when upgrading a software modular package. For more information, see the section [Upgrade a Modular Software Package](#) in *Software Upgrade and Boot Options* of the *ExtremeXOS Concepts Guide*.

SummitStack Only

You can issue this command only from the master node. If you issue this command from any other node, the following message appears:



Error: This command can only be executed on Master.

Note



After you stop a process, do not change the configuration on the switch until you start the process again. A new process loads the configuration that was saved prior to stopping the process. Changes made between a process termination and a process start are lost. Else, error messages can result when you start the new process.

Example

The following restarts the process tftpd:

```
start process tftpd
```

History

This command was first available in ExtremeXOS 11.0.

Support for restarting the Link Layer Discovery Protocol (lldp), Open Shortest Path First (ospf), and network login (netLogin) processes was added in ExtremeXOS 11.3.

Support for restarting the Border Gateway Protocol (bgp) was added in ExtremeXOS 11.4.

Support for restarting netTools was added in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

terminate process

```
terminate process name [forceful | graceful] {msm slot}
```

Description

Terminates the specified process on the switch.

Syntax Description

<i>name</i>	Specifies the name of the process to terminate. You can terminate the following processes: bgp (not available on Summit X440-L2 series switches), eapsexsshd (available only when you have installed the SSH module), isis (not available on Summit X440-L2 series switches), lldpnetLoginnetToolsospf (not available on Summit X440-L2 series switches), snmpMastersnmpSubagenttelnetdthttpdftpdvrrp (not available on Summit X440-L2 series switches), and xml.
forceful	Specifies a forceful termination.



graceful	Specifies a graceful termination.
<i>slot</i>	For a modular chassis, specifies the MSM/MM where the process should be terminated. A specifies the MSM installed in slot A, and B specifies the MSM installed in slot B. On a SummitStack, specifies the target node's slot number. The number is a value from 1 to 8. NOTE: This parameter is available only on modular switches and SummitStack.

Default

N/A.

Usage Guidelines

If recommended by Extreme Networks Technical Support personnel, you can stop a running process.

The forceful option quickly terminates a process on demand. Unlike the graceful option, the process is immediately shutdown without any of the normal process cleanup. The status of the operation is displayed on the console. After a successful forceful termination of a process, a message similar to the following appears:

```
Forceful termination success for snmpMaster
```

The graceful option terminates the process by allowing it to close all opened connections, notify peers on the network, and other types of process cleanup. After this phase, the process is finally terminated. After a successful graceful termination of a process, a message similar to the following appears:

```
Successful graceful termination for snmpSubagent
```

SummitStack Only

You can issue this command only from the master node. If you issue this command from any other node, the following message appears:

```
Error: This command can only be executed on Master.
```

To display the status of ExtremeXOS processes on the switch, including how many times a process has been restarted, use the `show process {name} {detail} {description} {slotslotid}` command.

Depending on the software version running on your switch and the type of switch you have, you can terminate different or additional processes. To see which processes you can terminate, enter `terminate process` followed by [Tab]. The switch displays a list of available processes.



To restart a process that has been terminated, use the `start process` command.

Note

Do not terminate a process that was installed since the last reboot unless you have saved your configuration. If you have installed a software module and you terminate the newly installed process without saving your configuration, your module may not be loaded when you attempt to restart the process with the `start process` command. To preserve a process's configuration during a terminate and (re)start cycle, save your switch configuration before terminating the process. Do not save the configuration or change the configuration during the process terminate and re(start) cycle. If you save the configuration after terminating a process, and before the process (re)starts, the configuration for that process is lost.

You can also use the `terminate process` command when upgrading a software modular package. For more information, see the section [Upgrade a Modular Software Package in Software Upgrade and Boot Options](#) of the *ExtremeXOS Concepts Guide*.

Example

The following initiates a graceful termination of the process `tftpd`:

```
terminate process tftpd graceful
```

History

This command was first available in ExtremeXOS 11.0.

Support for terminating the Link Layer Discovery Protocol (lldp), network login (netLogin), and Open Shortest Path First (ospf) processes was added in ExtremeXOS 11.3.

Support for terminating the Border Gateway Protocol (bgp) and Ethernet Automatic Protection Switching (eaps) processes was added in ExtremeXOS 11.4.

Support for terminating the MultiProtocol Label Switch (mpls) and Virtual Router Redundancy Protocol (vrrp) processes was added in ExtremeXOS 11.6.

Support for terminating netTools was added in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.



6 SummitStack Feature Commands

```
configure stacking alternate-ip-address
configure stacking easy-setup
configure stacking license-level
configure stacking mac-address
configure stacking master-capability
configure stacking priority
configure stacking protocol
configure stacking redundancy
configure stacking slot-number
configure stacking-support stack-ports
disable stacking
disable stacking-support
enable stacking
enable stacking-support
show power (Stack Nodes Only)
show stacking
show stacking configuration
show stacking detail
show stacking stack-ports
show stacking-support
synchronize stacking
unconfigure stacking
unconfigure stacking alternate-ip-address
unconfigure stacking license-level
unconfigure stacking-support
```

This chapter describes commands for:

- Configuring the SummitStack™ feature.
- Enabling and disabling a stack.
- Deploying a stack.
- Displaying stack details.

For an introduction to the SummitStack feature, see the *ExtremeXOS Concepts Guide*.

configure stacking alternate-ip-address

```
configure stacking alternate-ip-address [ipaddress netmask | ipNetmask] gateway
automatic configure stacking [node-address node-address | slot slot_number]
alternate-ip-address [ipaddress netmask | ipNetmask] gateway
```

Description

Configures an alternate management IP address, subnetwork, and gateway.

Syntax Description

<i>node-address</i>	Specifies the MAC address of a node in the stack. To view the MAC addresses for all nodes in a stack, enter the <code>show stacking</code> command. A node address or slot number is required unless the automatic keyword is specified.
<i>slot_number</i>	Specifies the slot number of the target node. To view the slot numbers, enter the <code>show stacking</code> command.
<i>ipaddress netmask</i>	Specifies the unique address that exists on the Management VLAN subnet as configured on the initial master node together with the subnetwork mask specified for the Management subnetwork. Example: 66.77.88.1 255.255.255.0.
<i>ipNetmask</i>	Specifies the unique address that exists on the Management VLAN subnet as configured on the initial master node, followed by a slash (/) character, followed by a decimal number that represents the number of leading one bits in the subnetwork address. An example is 66.77.88.1/24.
<i>gateway</i>	The address of an IP router. A default route is set up to reach this gateway.

Default

No alternate IP address is configured.

Usage Guidelines

If a Management subnetwork is configured and the alternate IP subnetwork does not exactly match the configured Management subnetwork, the information configured by one of the commands specified above is not used. The previously configured alternate IP address is removed if it was installed and subsequently a Management subnetwork is configured that does not exactly match the alternate IP subnetwork. In either case, an error message is logged. The alternate IP address is used if there is no configured Management subnetwork.

To use the command with the node address, the node must be in the stack topology; and to use the command with the slot number, the node must be in the active topology. This form of the command operates only on one node at a time. There are no checks to verify that the address is the one configured in the management VLAN subnet.

The command that does not require a node address or slot number specifies the automatic keyword. Usage of this form of the command causes an alternate IP address to be assigned to every node in the stack topology. The first address is the address specified in the [*ipaddress netmask* | *ipNetmask*]



parameter. The next address is the IP address plus one, and so on. Since there is a specified subnet mask, the address is checked to insure that the block of IP addresses fits within the specified subnet given the number of nodes in the stack topology. The range of addresses is tested to insure that each one is a valid IP unicast address. If the test fails, no node is configured and an error message is printed. Assignment is in the order in which nodes would currently appear in the `show stacking` display.

The configuration takes effect after the command is successfully executed.

The alternate IP address, subnetwork, and gateway are only used when the node is operating in stacking mode.

Example

To configure an alternate IP address for every node in the stack with a single command:

```
configure stacking alternate-ip-address 10.120.1.10/24 10.120.1.1 automatic
```

To configure an alternate IP address on a single node in the stack topology:

```
configure stacking node-address 00:04:96:26:6b:ed alternate-ip-address  
10.120.1.1/24 10.120.1.1
```

You may configure an alternate IP address using a slot number for a node that is currently occupying the related slot:

```
configure stacking slot 4 alternate-ip-address 10.120.1.13/24 10.120.1.1
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure stacking easy-setup

```
configure stacking easy-setup
```

Description

This command provides an easy way to initially configure the stacking parameters of all nodes in a new stack.



Syntax Description

This command does not have additional syntax.

Default

N/A.

Usage Guidelines

This command performs the following functions:

- Informs you of the stacking parameters that will be set.
- Informs you of the number of nodes that will be configured.
- Informs you whether minimal or no redundancy will be configured, and which slot will contain the master node.
- Informs you of the slot number that will be assigned to the node on which your management session is being run.
- If applicable, warns you that the current configuration file changes will be lost and you need to save the files.
- If the stack topology is a daisy chain, warns you that you should wire the stack as a ring before running this command.
- Requires you to confirm before the operation takes place. If you proceed, the command does the following:
 - Enables stacking on all nodes.
 - Configures the stacking MAC address using the factory address of the current node.
 - Configures a slot number for each node.
 - Configures redundancy to minimal in a ring topology or none in a daisy chain topology.
 - Configures the stacking protocol.
 - Reboots the stack topology.
- Selects the enhanced stacking protocol by default on Summit X460 and X480 series switches.

Stacking is enabled as if the `enable stacking {node-address node-address}` command was issued.

The stack mac-address is configured as if the `configure stacking mac-address` was issued on the current node.

Stack slot numbers are assigned as if the `configure stacking slot-number automatic` command was issued on the current node.

On a daisy chain topology, the master-capability is configured as if the `configure stacking redundancy none` command was issued. On a ring topology, the master-capability is configured as if the `configure stacking redundancy minimal` command was issued.

If you choose not to proceed with the setup, the following message is displayed:



Cancelled easy stack setup configuration.



Note

Summit Stack X440, X460, X480, X670, and X770 configured via `configure stacking easy-setup` use the enhanced stacking protocol by default.

Example

If you have an 8-node stack in a ring topology and have powered on all the nodes, the `show stacking` command shows the stack topology as a ring with all intended nodes present. If you have not changed any ExtremeXOS configuration, the command displays as follows:

```
* Switch.30 # configure stacking easy-setup
For every node in the 8-node stack, this command will:
- enable stacking
- configure a stack MAC address
- choose and configure a slot number (this node will be assigned to slot 1)
- configure redundancy to minimal (slot 1 will be the Master node)
Upon completion, the stack will automatically be rebooted into the new
configuration.
Warning: If stacking is already configured, this command will alter that
configuration.
Warning: There are unsaved configuration changes. You may wish to save them
before proceeding.
Do you wish to proceed? (y/N) y
Stacking configuration is complete. Rebooting...
```

If the 8-node stack topology is a daisy chain, and the user is logged into a node in the middle of the chain, the command output might appear as follows:

```
* Switch.30 # configure stacking easy-setup
For every node in the 8-node stack, this command will:
- enable stacking
- configure a stack MAC address
- choose and configure a slot number (this node will be assigned to slot 5)
- configure redundancy to none (slot 1 will be the master node)
Upon completion, the stack will automatically be rebooted into the new
configuration.
Warning: If stacking is already configured, this command will alter that
configuration.
Warning: This stack is a daisy chain. It is highly recommended that the stack
be connected as a ring before running this command.
Do you wish to proceed? (y/N) Yes
Stacking configuration is complete. Rebooting...
```

History

This command was first available in ExtremeXOS 12.0.



Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure stacking license-level

```
configure stacking {node-address node-address | slot slot-number} license-level
[core | advanced-edge | edge]
```

Description

Allows you to restrict the license level at which the node operates.

Syntax Description

<i>node-address</i>	Specifies the MAC address of a node in the stack. To view the MAC addresses for all nodes in a stack, enter the <code>show stacking</code> command.
<i>slot-number</i>	Specifies the slot number of the target node. To view the slot numbers, enter the <code>show stacking</code> command.

Default

No license level restriction is configured.

Usage Guidelines

This command causes a node to operate at a lower license level than the level that was purchased for the node.

Running this command does not change the installed license level. For example, if a stackable is configured with the Advanced Edge license and you configure a license level restriction of Edge, the unit is restricted to features available in the Edge license. However, you can remove the restriction and operate at the Advanced Edge level.

If the installed license level of the target node is lower than the level you are attempting to configure, the following message appears:

```
Warning: Switch will not operate at a license level beyond that which was purchased.
```

If the `node-address` or `slot` parameter is not specified, the command takes effect on every node in the stack topology.

This command takes effect after you restart the node. The following message appears after the command is executed:

```
This command will take effect at the next reboot of the specified node(s).
```



If you restart the node without configuring a license level restriction, the node operates at the purchased license level. To see the purchased license level of a node, run `show licenses` after logging in to the node.

The `show licenses` command displays the current license level in use as the Effective License Level:

```
Slot-2 Stack.1 # show licenses
Enabled License Level:
Advanced Edge
Enabled Feature Packs:
None
Effective License Level:
Edge
```

The `show stacking configuration` and `show stacking {node-address node-address | slotslot-number} detail` commands allow you to see the configured license level restriction and the restriction currently in use.

The Effective License Level appears only when stacking is enabled. The command is node-specific. The effective license level is the level at which the node is restricted to operate, and is not necessarily the level at which the entire stack is operating. This is because it is possible to have the restriction differ on each node, in which case one or more nodes may have failed because of the differing levels.

Example

To configure the stacking level Edge on all nodes in a stack:

```
configure stacking license-level edge
```

To configure stacking level Edge for a node:

```
configure stacking node-address 00:04:96:26:6b:ed license-level edge
```

To configure the stacking level Advanced Edge for an active node that currently occupies slot 4:

```
configure stacking slot 4 license-level advanced-edge
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure stacking mac-address

```
configure stacking {node-address node-address | slot slot-number} mac-address
```

Description

Selects a node in the stack whose factory assigned MAC address is to be used to form the stack MAC address.

The formed address is then configured on every node in the stack topology.

Syntax Description

<i>node-address</i>	Specifies the MAC address of a node in the stack. To view the MAC addresses for all nodes in a stack, enter the <code>show stacking</code> command.
<i>slot-number</i>	Specifies the slot number of an active node whose factory MAC address is to be used to form the stack MAC address. To view the slot numbers, enter the <code>show stacking</code> command.

Default

No stack MAC selection is configured.

Usage Guidelines

You must select a node whose factory assigned MAC address can be used to form a MAC address that represents the stack as a whole. The system forms the stack MAC address by setting the Universal / Local bit in the specified MAC address. This means that the stack MAC address is a locally administered address, and not the universal MAC address assigned to the selected node.

If you do not specify any node, the stack MAC address is formed from the factory assigned MAC address of the node from which you are running the command.

This command takes effect only after you restart the node. The following message appears after you run the command:

```
This command will take effect at the next reboot of the specified node(s).
```

If a stack node that has just joined the active topology detects that its stack MAC address is not configured or is different than the stack MAC address in use, it logs the following message at the Error log level:

```
The stack MAC address is not correctly configured on this node. The stack can not operate properly in this condition. Please correct and reboot.
```

If you have not configured (or inconsistently configured) the stack MAC address you might encounter difficulty in diagnosing the resulting problems. Whenever the master node (including itself) detects that one or more nodes in its active topology do not have the correct or any stack MAC address configured, it displays the following message to the console every five minutes until you configure a MAC address and restart the node(s):



The stack MAC address is either not configured or its configuration is not consistent within the stack. The stack can not operate properly in this condition. Please correct and reboot.

Example

To select the node to which you have logged in to supply the MAC address for stack MAC address formation:

```
configure stacking mac-address
```

To select a node other than the one to which you are logged in to supply the MAC address for stack MAC address formation:

```
configure stacking node-address 00:04:96:26:6b:ed mac-address
```

To select an active node to supply the MAC address for stack MAC address formation:

```
configure stacking slot 4 mac-address
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure stacking master-capability

```
configure stacking [node-address node_address | slot slot_number] master-capability [on | off]
```

Description

The command configures a node to be allowed to operate as either a backup or master, or prevents a node from operating as either.

The command controls the setting on the specified node only. To set the master capability for all nodes on a stack, you can use the command `configure stacking redundancy [none | minimal | maximal]`.



Syntax Description

<code>node_address</code>	Specifies the MAC address of a node in the stack. To view the MAC addresses for all nodes in a stack, enter the <code>show stacking</code> command.
<code>slot_number</code>	Specifies the slot number of the target active node. To view the slot numbers, enter the <code>show stacking</code> command.

Default

Master-capability is On.

Usage Guidelines

At least one node in the stack topology must be master-capable.

If you attempt to disable the master-capability of the only master capable node in a stack topology, the attempt is denied and following message appears:

```
Error: At least one node must have Master-capability configured "on".
```

This command is used to set up master-capability manually. It can also be used to adjust the result achieved when the `configure stacking redundancy [none | minimal | maximal]` command is used.

The setting takes effect the next time the node reboots. When this command is executed successfully, the following message appears:

```
This command will take effect at the next reboot of the specified node(s).
```

Example

To turn on the master capability for a node:

```
configure stacking node-address 00:04:96:26:6b:ed master-capability on
```

To turn on the master capability of an active node currently occupying slot 4:

```
configure stacking slot 4 master-capability on
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure stacking priority

```
configure stacking {node-address node-address | slot slot_number} priority
[node_pri | automatic]
```

Description

Configures a priority value to be used to influence master and backup election.

Syntax Description

<i>node-address</i>	Specifies the MAC address of a node in the stack. To view the MAC addresses for all nodes in a stack, enter the <code>show stacking</code> command.
<i>slot_number</i>	Specifies the slot number of the target node. To view the slot numbers, enter the <code>show stacking</code> command.
<i>node_pri</i>	Specifies the priority as a value between 1 and 100.

Default

Automatic priority.

Usage Guidelines

The node role election priority is a value that is internally calculated by ExtremeXOS for each node. This calculated value helps determine which nodes are elected as master and backup. For more information, see [Configuring the Master, Backup, and Standby Roles](#) in the *ExtremeXOS Concepts Guide*.

This command allows you to configure a priority value that affects the outcome of this calculation. You can configure the priority on any node in a stack topology. You can specify an integer node-pri value between 1 and 100. The larger the value, the greater the node role election priority.

If no node address or slot is specified, the command takes effect on all nodes at the next node role election cycle. Priority configuration has no operational effect on switches that are not in stacking mode.

If configured on every node, automatic priority commands ExtremeXOS to determine the node role election priority of each active node. Currently, the automatic priority algorithm chooses the master-capable node with the lowest slot number as master and the node with the second lowest slot number as backup. Extreme networks may alter this behavior in later releases.

If you have configured a node with automatic priority and if you have configured another node to use a node-pri value, the node with automatic priority uses zero as the node-priority value during the node role election.



Example

To allow ExtremeXOS to determine node role election priority:

```
configure stacking priority automatic
```

To configure the node priority for the stackable in slot 4:

```
configure stacking slot 4 priority 50
```

To configure the automatic priority algorithm for the stackable with node address 00:04:96:26:6b:ed:

```
configure stacking node-address 00:04:96:26:6b:ed priority automatic
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure stacking protocol

```
configure stacking protocol [standard | enhanced]
```

Description

Configures the stacking port protocol.

Syntax Description

standard	Specifies the standard protocol, which is supported on all SummitStack capable switches.
enhanced	Specifies the enhanced protocol, which is supported only on Summit X460, X480, and X670 switches. The enhanced protocol is required to support MPLS.

Default

Standard.



Usage Guidelines

Use this command to change the configured stacking protocol on a stack made up of Summit X460, X480, or X670 switches.



Note

You must reboot the switch to activate the protocol change.

If MPLS is enabled on the switch, you must disable MPLS before you can change the stacking protocol to standard.

To display the stacking protocol configuration, enter the show stacking configuration command.

Example

To configure a switch to use the enhanced protocol, enter the following command:

```
configure stacking protocol enhanced
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available only on Summit X440, X460, X480, and X670 switches.

configure stacking redundancy

```
configure stacking redundancy [none | minimal | maximal]
```

Description

This command sets a master-capability value for every node in the stack topology.

Syntax Description

none	Only one node has master-capability turned on and all other nodes have master-capability turned off.
minimal	Two nodes have master-capability turned on and all other nodes have master-capability turned off.
maximal	All nodes have master-capability turned on.

Default

Default value in an unconfigured stack is maximal.



Usage Guidelines

If there are more than eight nodes in the stack topology, the following message appears and the command is not executed:

```
ERROR: This command can only be used when the stack has eight nodes or less.
```

Since only eight nodes can be operational in an active topology at a time, you must disconnect the remaining nodes before configuring master-capability with this command.

If you are using the none or minimal redundancy configuration:

- The configured values of slot-number and priority decide the nodes on which the master-capability should be turned on.
- If the priority values are configured on the nodes, the highest priority node(s) are chosen.
- If the priority values of all nodes are set to automatic or to the same priority value, the node(s) with the lowest slot number(s) are chosen. Extreme Networks may change automatic priority behavior in a future release.

If there is a slot number tie or if the slot numbers were never configured, the following message appears and the command is not executed:

```
ERROR: Unique slot numbers must be configured before using this command.
```

The setting takes effect at the next restart of the node. The following message appears after the command is successfully executed:

```
This command will take effect at the next reboot of the specified node(s).
```

Redundancy configuration has no operational effect on a node that is not in stacking mode.

Example

To turn on master-capability on all nodes:

```
configure stacking redundancy maximal
```

To turn on master-capability on only one node:

```
configure stacking redundancy none
```

To turn on master-capability on two nodes:

```
configure stacking redundancy minimal
```

History

This command was first available in ExtremeXOS 12.0.



Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure stacking slot-number

```
configure stacking slot-number automatic configure stacking node-address  
node_address slot-number slot_number
```

Description

Configures a slot number on one or all nodes in the stack topology.

Syntax Description

<i>node_address</i>	Specifies the MAC address of a node in the stack. To view the MAC addresses for all nodes in a stack, enter the <code>show stacking</code> command.
<i>slot_number</i>	Specifies a number between 1 and 8 that is to be assigned as the slot number of the target node.

Default

The default slot-number for a node in stacking mode is 1.

Usage Guidelines

The configuration is stored on the affected node(s) immediately but does not take effect until the next reboot of the node(s). The configuration applies only when the node is running in stacking mode. To see the configured and active slot numbers of all nodes, use the `show stacking configuration` command.

If a node-address and a slot number are specified, then the node is configured with the specified slot number. There is no check for a duplicate slot number at this time; the number is simply assigned as requested.

If the automatic keyword is specified, then automatic slot number assignment is selected, and the assignment is performed on all nodes in the stack topology. If there are more than eight nodes in the stack topology, the assignment is only performed on the first eight nodes.

Automatic slot number assignment causes assignment of slot numbers starting from 1 and increasing up to 8. The nodes in the stack topology are assigned the numbers in the order in which they would appear currently in the `show stacking` command output. In a ring, slot number 1 is assigned to the current node, slot number 2 is assigned to the node connected to the current node's stack port 2, and so forth. In a daisy chain, slot 1 is assigned to the node at the end of the chain that begins with the node connected to the current node's stack port 1.



To see the resulting slot number assignment, run the `show stacking configuration` command.



Note

Failure to configure a node does not prevent configuration of the slot numbers on the other nodes, and does not affect the slot number assigned to each node.

If you enter the command with the automatic option, the following confirmation message appears:

```
Reassignment of slot numbers may make the stack incompatible with the
current configuration file. Do you wish to continue? (y/n)
```

When this command is executed successfully, the following message appears:

```
This command will take effect at the next reboot of the specified
node(s).
```

Example

To configure all slot-numbers for a stack:

```
configure stacking slot-number automatic
```

To configure slot number 4 for the node with MAC address 00:04:96:26:6b:ed:

```
configure stacking node-address 00:04:96:26:6b:ed slot-number 4
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure stacking-support stack-ports

```
configure stacking-support stack-port [stack-ports | all] selection [native {V80
| V160} | V320} | alternate]
```

Description

Selects the switch ports and speed for stack communications.



Syntax Description

<i>stack-ports</i>	Specifies the stacking port range to be configured. Valid stacking port entries are 1, 2, 1-2, and all.
native	Selects the specified stacking port, which is the native, dedicated port that only supports stacking.
v80	Specifies that the native stacking ports on a VIM3-40G4X or VIM4-40G4X option card operate at 80 Gbps.
v160	Specifies that the native stacking ports on a VIM3-40G4X or VIM4-40G4X option card operate at 160 Gbps.
v320	Specifies that the native stacking ports on a VIM3-40G4X or VIM4-40G4X operate at 320 Gbps.
alternate	Selects the alternate (Ethernet) stacking port associated with the specified stacking port. The alternate port numbers are listed in the following table.

Default

Switches with native stack ports: Native. This command does not apply to switches without native stack ports.

Native stacking ports on Summit switches with a VIM3-40G4X or VIM4-40G4X option card operate as one 40 Gbps port.

Usage Guidelines

The configuration entered with this command applies to only the local node and does not become active until after the following events:

- The stacking-support option is enabled (if applicable).
- The switch restarts.

The V80, V160, and V320 keywords apply only to Summit switches with an installed VIM3-40G4X or VIM4-40G4X option card. Each speed configuration requires a specific cabling configuration. For more information, see the Summit Family Switches Hardware Installation Guide.

The stacking-support option configures the switch to use stacking protocols. This option is automatically enabled on most platforms, but some platforms require you to manually enable the stacking-support option. The following table lists the Summit family switches and option card configurations that support Stacking Port Selection Control, and it lists which platforms require manual Stacking-Support Option Control.

¹ To operate in a stack, the `enable stacking-support` command must be entered for switch configurations for which this column displays Yes.

² The `configure stacking-support stack-ports` command is supported only on Summit switch configurations for which this column displays Yes.



Table 13: Summit Family Switch Support for Alternate Stack Ports

Summit Switch Model Number	Summit Switch Option Card	Alternate Port for Stack Port1	Alternate Port for Stack Port2	Stacking-Support Option Control ¹	Stacking Port Selection Control ²
X440-24t-10G, X440-24p-10G X440-48t-10G, X440-48p-10G	None None XGM2-2xf XGM2-2xn XGM2-2sf XGM2-2bt ³ None	25 49 25 ⁴	26 50 26c	Yes Yes No	No No Yes
X460-48t X460-48p	XGM3-2sf with either the XGM3 SummitStack V80 or XGM3 SummitStack module or neither	S1c	S2c	No	Yes
X460-24t X460-24x X460-24p		S1c	S2c	No	Yes
X460-48x		S1c	S2c	No	Yes
X480-48t X480-48x	VIM2-10G4X	S3	S4	Yes	No
	VIM2-SummitStack	N/A	N/A	N/A	N/A
	VIM2-SummitStack-V80	N/A	N/A	N/A	N/A
	VIM2-SummitStack128	N/A	N/A	N/A	N/A
	None	N/A	N/A	N/A	N/A
X480-24x	VIM2-10G4X	S3	S4	Yes	No
	VIM2-SummitStack	25	26	No	Yes
	VIM2-SummitStack-V80	25	26	No	Yes
	VIM2-SummitStack128	25	26	No	Yes
	None	25	26	Yes	No
X670-48x	None	47	48	Yes	Yes
X670V-48x	VIM4-40G4X	47	48	Yes	Yes
X770-32q	None	103	104	Yes	Yes

When the alternate stack port is selected for a native stack port and the switch is restarted, the native stack port remains visible in the CLI and can be configured. However, any configuration applied to the replaced stack port is ignored and does not affect switch operation.

¹ To operate in a stack, the `enable stacking-support` command must be entered for switch configurations for which this column displays Yes.

² The `configure stacking-support stack-ports` command is supported only on Summit switch configurations for which this column displays Yes.

³ SummitStack-V requires XGM2-2bt version 4 or later option cards.

⁴ This alternate port number requires an installed option card. You can configure the port without the option card, but the configuration does not apply until the switch restarts with the required option card.

¹ To operate in a stack, the `enable stacking-support` command must be entered for switch configurations for which this column displays Yes.

² The `configure stacking-support stack-ports` command is supported only on Summit switch configurations for which this column displays Yes.



An alternate stack port runs the stacking protocol and cannot operate on a link connected to a data port that is not configured as a stack port. Both ends of a stack link must be configured to use the stacking protocol. The stacking link must be directly connected to two the alternate stacking ports of two stacking switches. The direct connection is necessary because stacking protocols cannot pass through an intermediate switch.

After a data port is activated as an alternate stack port, all data port configuration commands still work, but they do not change the operation of the alternate stack port. The LEDs on an Ethernet port used as an alternate stacking port operate according to the behavior of the Ethernet port. The LEDs on the related (disabled) native stacking port remain dark.

Note

Commands that contain the stacking-support keyword operate only on the local switch; they do not apply to all switches in the stack. If an active stack topology has been formed, you can telnet to a slot elsewhere in the stack, log on to that switch, and use commands with the stacking-support keyword on that switch.

Example

The following command configures the switch to use the alternate stack port for Stack Port 1 after the next switch restart:

```
configure stacking-support stack-ports 1 selection alternate
```

The following command configures the switch to use both native stacking ports after the next switch restart:

```
configure stacking-support stack-ports 1-2 selection native
```

The following command configures stack ports 1 and 2 to operate as four 10 Gbps ports:

```
configure stacking-support stack-ports 1-2 selection native V80
```

History

This command was first available in ExtremeXOS 12.5.

The V80 and V160 keywords were added in ExtremeXOS 12.6.

The V320 keyword was added in ExtremeXOS 15.1 Revision 2.

Platform Availability

This command is available on the platforms listed in the preceding table. The V80 and V160 keywords are supported only on Summit switches with the VIM3-40G4X or VIM4-40G4X option card installed.



disable stacking

```
disable stacking {node-address node-address}
```

Description

This command disables the stacking on one or all nodes in the stack topology.

Syntax Description

node-address	Specifies the MAC address of a node in the stack. To view the MAC addresses for all nodes in a stack, enter the <code>show stacking</code> command.
---------------------	---

Default

Default value is stacking disabled.

Usage Guidelines

If you do not specify the node-address, stacking is disabled on all nodes in the stack topology.

If the node-address parameter is present, stacking is disabled on the node with the specified node-address. This is the MAC address assigned to the stackable by the factory.

A node in the stack topology that is disabled for stacking does not forward the customer's data through its stacking links and does not become a member of the active topology.

A disabled node becomes its own master and processes and executes its own configuration independently.

When this command is executed successfully, the following message appears:

```
This command will take effect at the next reboot of the specified node(s).
```

Use `show stacking configuration` command to see the current configuration of the stack. Verify the flags in `show stacking configuration` output to confirm that stacking is disabled on the specified node(s).

Example

The following example disables stacking on an 8 node stack:

```
* Switch.3 # disable stacking
This command will take effect at the next reboot of the specified node(s).
```



The following example disables stacking on the node with the factory assigned MAC address 00:04:96:26:6b:ed:

```
* Switch.3 # disable stacking node-address 00:04:96:26:6b:ed
This command will take effect at the next reboot of the specified node(s).
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable stacking-support

disable stacking-support

Description

This command disables the stacking-support option on a Summit X480, X670, or X770 switch with dual-purpose hardware.

Syntax Description

This command does not have additional syntax.

Default

Disabled.

Usage Guidelines

The Stacking-Support Option Control column in [Table 13: Summit Family Switch Support for Alternate Stack Ports](#) on page 380 displays Yes in the rows for switch configurations for which you can disable the stacking-support option.

After you disable the stacking-support option, you must reboot the switch to activate the configuration change.

If you disable the stacking-support option on a switch and reboot, stacking communication stops and the data ports listed in [Table 13: Summit Family Switch Support for Alternate Stack Ports](#) on page 380 will use Ethernet protocols instead of stacking protocols.



Example

To disable the stacking ports, enter the following command:

```
* switch # disable stacking-support
This setting will take effect at the next reboot of this switch.
```

History

This command was first available in ExtremeXOS 12.2.

Support for the Summit X480 switch configurations listed in the “Usage Guidelines” section was added in ExtremeXOS 12.5.

Platform Availability

This command is available only on Summit X480, X670, and X770 switches.

enable stacking

```
enable stacking {node-address node-address}
```

Description

This command enables stacking on one or all nodes.

Syntax Description

<i>node-address</i>	Specifies the MAC address of a node in the stack. To view the MAC addresses for all nodes in a stack, enter the <code>show stacking</code> command.
---------------------	---

Default

Default value is stacking disabled.

Usage Guidelines

This command enables stacking on one or all nodes. When a node is operating in stacking mode, QoS profile QP7 cannot be created.

If a `node-address` is not specified, this command first performs an analysis of the current stacking configuration on the entire stack. If the stack has not yet been configured for stacking operation, or if the configuration is self-inconsistent, the user is offered the option of invoking the easy setup function. The following message appears:

```
You have not yet configured all required stacking parameters. Would you
like to perform an easy setup for stacking operation? (y/N)
```



If you enter Yes, the easy setup procedure is invoked and you first see the following message:

```
Executing "configure stacking easy-setup" command...
```

If you enter No, the following message appears:

```
Stacking has been enabled as requested.
```

The following describes the operation performed if easy setup is neither offered nor selected.

If you do not enter any node-address, stacking is enabled on all nodes in the stack topology.

If the node-address parameter is present, stacking is enabled on the node with the specified node-address. This is the MAC address assigned to the stackable by the factory.

The `show stacking configuration` command shows the current configuration of this parameter as well as the value currently in use.

A node that is enabled for stacking attempts to join the active topology. If successful, it then negotiates a node role with the other nodes in the stack and becomes an operational node in the stack according to its role. The master node's configuration is applied to the node.

When this command is executed successfully, the following message appears:

```
This command will take effect at the next reboot of the specified node(s).
```

Example

To enable stacking on a stack:

```
* Switch.3 # enable stacking
This command will take effect at the next reboot of the specified node(s).
```

To enable stacking on node 5, with a MAC address 00:04:96:26:6b:ed:

```
* Switch.3 # enable stacking node-address 00:04:96:26:6b:ed
This command will take effect at the next reboot of the specified node(s).
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



enable stacking-support

enable stacking-support

Description

This command enables a Summit X480, X670, or X770 switch with dual-purpose hardware to participate in a stack.

Syntax Description

This command does not have additional syntax.

Default

Disabled.

Usage Guidelines

The Stacking-Support Option Control column in [Table 13: Summit Family Switch Support for Alternate Stack Ports](#) on page 380 displays Yes in the rows for switch configurations for which you can enable the stacking-support option.

After you enable the stacking-support option, you must reboot the switch to activate the configuration change.

If you enable the stacking-support option on a switch and reboot, data communications on the data ports listed in [Table 13: Summit Family Switch Support for Alternate Stack Ports](#) on page 380 stops, and the ports use stacking protocols instead of Ethernet protocols.

Example

To enable the stack ports, enter the following command:

```
* switch # enable stacking-support
This setting will take effect at the next reboot of this switch.
```

History

This command was first available in ExtremeXOS 12.2.

Support for the Summit X480 switch configurations listed in the “Usage Guidelines” section was added in ExtremeXOS 12.5.

Platform Availability

This command is available only on Summit X480, X670, and X770 switches.



show power (Stack Nodes Only)

show power

Description

Displays the number of power modules present and providing power in each slot.

Syntax Description

This command has no arguments or variables.

Default

Default value

Usage Guidelines

This command is available on all platforms. However, it produces completely different output on a stack. The following table describes the flags that appear when this command is executed on an active node.

Table 14: Flag Descriptions for the show power Command

Power Supply	Flag	Meaning
Internal	F	Failed or no Power.
Internal	P	Power available.
Internal	O	48V powered off (48p Summits only).
External (non 48P)	-	Empty.
External (non 48P)	F	Failed or no power.
External (non 48P)	P	Power available.
External (non 48P)		Power supply can never occupy this position.
External (48P only)	-	Empty or no power to all PSUs present.
External (48P only)	F	Failed or no power (at least 1 PSU has power).
External (48P only)	P	Power available.

All Summit family switches accept an external power chassis that holds only one power supply. For Summit family switches, the External PSU columns are left blank.

For slots without active nodes, the slot number appears and the remainder of the row is blank.



Example

The following are sample displays for this command:

```
Slot-2 Stack.40 # show power
Internal External External External
Slots Type      PSU      PSU      PSU      PSU
-----
Slot-1 X          F          P
Slot-2 X          P          -
Slot-3 X          -          P
Slot-4 X          P          P
Slot-5 X          P          -
Slot-6 X          -          P
Slot-7 X          P          P          F          -
Slot-8 X          P          -
Flags : (P) Power available, (F) Failed or no power,
(O) 48V powered off when 2 or 3 external PSU are powered on,
(-) Empty
Slot-2 Stack.41 #
Slot-2 Stack250.4 # show power
Internal External External External
Slots Type      PSU      PSU      PSU      PSU
-----
Slot-1 X          F          P
Slot-2 X          P          -
Slot-3 X          P          -
Slot-4 X          P          P
Slot-5 X          P          -
Slot-6 X          O          P          P          F
Slot-7
Slot-8 X          P          -
Flags : (P) Power available, (F) Failed or no power,
(O) 48V powered off when 2 or 3 external PSUs are powered on,
(-) Empty
Slot-2 StackX #
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms, but the output described in this section is available only on nodes configured for the SummitStack feature.

show stacking

show stacking



Description

The show stacking command shows a summary of the nodes in the stack topology.

The show stacking command shows all nodes that are in the stack topology.

Syntax Description

This command has no arguments or variables.

Default

There is no default value for this command.

Usage Guidelines

The asterisk (*) that precedes the node MAC address indicates the node on which this command is being executed, that is, the node to which the user is logged in.

The node MAC address is the address that is factory assigned to the stackable.

The slot number shown is the number currently in use by the related node. Since slot number configuration only takes effect during node initialization, a change in configured value alone does not cause a change to the slot number that is in use. Slot numbers show as hyphen (-) characters on nodes that have stacking disabled.

The Stack State shows the state values.

The Role is one of the following: Master, Backup, Standby, or *none*.

In a ring topology, the node on which this command is executed is always the first node displayed. The order of the nodes shown in the display is the order of their physical connection in the ring.

Even though the stack topology can be a ring, the active topology can simultaneously be a daisy chain because it is only a proper subset of the stack topology. If the node on which this command is executed is not active, the line

```
Active Topology is a ____
```

is replaced by the line

```
This node is not in an Active Topology.
```

The daisy chain topology is displayed in the order of physical connection. The master node detects the two nodes in the stack topology that have only one operating link, and these nodes become the ends of the stack. Such nodes always display at the top and bottom of the output.

It is possible for a node to be in Stabilizing or Waiting state and still be in the active topology. This is because it is possible for an active node to move to these states when a topology change is detected.



Once a node becomes active, the node remains an active node until it reboots or an overflow condition occurs.

The Flags have the following definitions:

- The C flag indicates that the related node is a candidate for membership of the same active topology to which the node on which the command is executed would belong.
- The A flag indicates that the related node is an active node in the active topology of which the node on which the command is run is also a candidate node. Being an active node is necessary but not sufficient for presence of the node in a slot. Once the node has fully initialized, the active node appears as Present in the show slot display.
- The O flag indicates that the related node is probably an active node in an active topology for which the node on which this command is being run is not a candidate.

The O flag is useful for the case where there is an inhibited link or a disabled or failed node that separates two active topologies. One active topology may contain the local node, and all other nodes in this active topology do not have the O flag set. All nodes that are members of an active topology that is separated by an inhibited link from the active topology that contains the local node have only the O flag set. All possibly active nodes have the O flag set if the local node is not a member of any active topology. For any node for which the O flag is set, the C and A flags are not set and vice-versa.

The following information is displayed:

- Stack Topology is a ring or daisy-chain.
- Active Topology is a ring or daisy-chain (or This node is not in an Active Topology.).
- For each node:
 - Node MAC address (factory assigned).
 - Slot number in use.
 - Stack State:
 - Disabled - Node is not configured for stacking.
 - Failed - Node can't come up in the stack because it has a duplicate slot number.
 - Overflow - The node has detected that there are more nodes in the stack topology than are allowed.
 - Listening - Initial state when attempting to join the stack. The node is checking to see if its configured slot number duplicates that of another node. The node cannot be an active node in this state.
 - Stabilizing - Node is waiting until it sees no new topology changes. The node may or may not be an active node in this state.
 - Waiting - Topology has stabilized, if the active topology is to be a ring, and stacking link blocking is being performed. The node may or may not be an active node in this state.
 - Active - The node is an active node and is fully programmed to operate in the active topology.
- Node role (master, backup, standby, or other transient node state).
- Flags describing the node's membership in the active topology.
- Whether or not the node is this node, that is, the node on which the command is run.



Example

The following example shows the output of show stacking command:

```
Slot-1 Stack.30 # show stacking
Stack Topology is a Ring
Active Topology is a Daisy-Chain
Node MAC Address      Slot  Stack State  Role    Flags
-----
*00:04:96:26:60:DD  1      Stabilizing  Master  CA-
00:04:96:26:60:EE  4      Stabilizing  Standby C--
00:04:96:26:60:FF  -      Disabled     Master  ---
00:04:96:26:60:AA  -      Disabled     Master  ---
00:04:96:26:60:88  -      Disabled     Master  ---
00:04:96:26:60:99  -      Disabled     Master  ---
00:04:96:26:60:BB  2      Stabilizing  Standby C--
00:04:96:26:60:CC  3      Active       Backup  CA-
(*) Indicates This Node
Flags: (C) Candidate for this active topology, (A) Active node,
(O) node may be in Other active topology
Slot-1 Stack.31 #
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show stacking configuration

show stacking configuration

Description

Shows how the nodes are configured in a stack topology. The configured values shown are the ones actually stored in the remote nodes at the time you issue this command.

Syntax Description

This command has no arguments or variables.

Default

N/A.



Usage Guidelines

Some stacking parameters do not take effect until the next restart, the configured values and the values currently being used are both shown. Specifically, this applies to the slot number, whether or not stacking is enabled, the master-capable configuration, the license level restriction, and the stack MAC configuration.

The only parameters that take effect without a reboot are the node priority and the alternate management IP subnetwork and gateway.

The Stack MAC in use line can display the following values:

- If the command is executed on the master node:
 - *none* if there is no stack MAC configured.
 - The stack MAC configured on the master node.
- If the command is executed on a non-master node:
 - *unknown*. The stack MAC address is only known by the executing master node. In this case, the M and m flags are not set. The i flag is set if there is a stack MAC configured locally.

Identified with the asterisk, the current node is the one on which the `show stacking configuration` command is executed.

A node identified with the ? character indicates that timely attempts to fetch the configuration information from the node have failed. There are two possible reasons for this display:

- Communications with the node have been lost, in which case the node will probably be removed from the stack topology shortly.
- The node is too busy to respond in time.

A row that displays the ? indicator shows the last values that were received from the node. If no values were ever received, all configured values show as not configured (-) or *none*. The node MAC address and the slot number that is currently in use are still displayed.

Example

The following example:

```
Slot-1 Stack.2 # show stacking configuration
Stack MAC in use: 02:04:96:26:6b:ed
Node          Slot          Alternate          Alternate
MAC Address   Cfg Cur Prio Mgmt IP / Mask         Gateway          Flags
Lic
-----
---
*00:04:96:26:6b:ed 1  1  Auto <none>          <none>          CcEeMm---
--
00:04:96:34:d0:b8 2  2  Auto <none>          <none>          CcEeMm--- --
* - Indicates this node
Flags: (C) master-Capable in use, (c) master-capable is configured,
(E) Stacking is currently Enabled, (e) Stacking is configured Enabled,
(M) Stack MAC in use, (m) Stack MACs configured and in use are the same,
(N) Stack link protocol Enhanced in use, (n) Stack link protocol Enhanced
configured,
(i) Stack MACs configured and in use are not the same or unknown,
```



(-) Not in use or not configured
 License level restrictions: (C) Core, (A) Advanced edge, or (E) Edge in use,
 (c) Core, (a) Advanced edge, or (e) Edge configured,
 (-) Not in use or not configured

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show stacking detail

```
show stacking {node-address node_address | slot slot_number} detail
```

Description

This command displays information about a specified node.

Syntax Description

<i>node_address</i>	Specifies the MAC address of a node in the stack. To view the MAC addresses for all nodes in a stack, enter the <code>show stacking</code> command.
<i>slot_number</i>	Specifies the slot number of the target node. To view the slot numbers, enter the <code>show stacking</code> command.

Default

N/A.

Usage Guidelines

If no node is specified, the output is generated for all nodes in the stack topology. If the specified node does not exist, an error message appears. The slot parameter is available only for active nodes in the same active topology as the node on which the command is run. The node-address parameter is always available.

Current information represents stacking states and configured values that are currently in effect. Configured information is that which takes effect at node reboot only. Thus, differences between values in use and values configured can be seen here. The advantages of this command over the `show stacking configuration` command is that the values in use and the configured values are fully expanded without the need for flags. You can also see the port state information of the node(s).



The roles values are: Master, Backup, Standby, and *none*.

License level restrictions can be Edge, Advanced Edge, or Core.

If one of the fields in the example below is missing on your switch, your switch does not support the feature that the field represents.

Example

The following is a sample output of this command:

```
Slot-1 Stack.33 # show stacking slot 1 detail
Stacking Node 00:04:96:26:6b:ec information:
Current:
Stacking                : Enabled
Role                    : Master
Priority                 : Automatic
Slot number             : 1
Stack state              : Active
Master capable?         : Yes
Stacking protocol       : Enhanced
License level restriction : <none>
In active topology?     : Yes
Factory MAC address     : 00:04:96:26:6b:ec
Stack MAC address       : 02:04:96:26:6b:ec
Alternate IP address     : <none>
Alternate gateway       : <none>
Stack Port 1:
State                   : Operational
Blocked?                : No
Control path active?    : Yes
Selection                : Alternate (23)
Stack Port 2:
State                   : Operational
Blocked?                : Yes
Control path active?    : Yes
Selection                : Native
Configured:
Stacking                : Enabled
Master capable?         : Yes
Slot number             : 1
Stack MAC address       : 02:04:96:26:6b:ec
Stacking protocol       : Enhanced
License level restriction : <none>
Stack Port 1:
Selection                : Alternate (23)
Stack Port 2:
Selection                : Native
```

History

This command was first available in ExtremeXOS 12.0.

The Stacking protocol and Stack Port Selection fields were added in ExtremeXOS 12.5.



Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show stacking stack-ports

show stacking stack-ports

Description

This command displays the port states of each node in the stack topology and the connections between the nodes.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The slot number shown is the slot number in use on stacking enabled nodes. If the node does not have stacking enabled, a hyphen character (-) is shown instead of a number.

The Port and Node MAC Address field values in the command display identify a particular stacking port. Each node MAC address appears twice in two consecutive rows in the output because each node has two stacking ports. On all platforms, the ports are labeled with the values 1 or 2. The order in which stacking ports appear in the display is the order in which they are physically connected.

The Select field indicates whether the stacking port is using a native stacking port or an alternate 10Gbps Ethernet port. If a number appears in this column, it represents the port number printed on the switch for a 10 Gbps Ethernet port. For more information, see the description for the `configure stacking-support stack-ports` command.

The Port State field for each port shows one of the following states:

- Link Down – port is not receiving a signal.
- No Neighbor – the port is receiving a signal but it is not identifying a stack neighbor.
- Overflow – 17 nodes (or more) are physically connected to this port.
- Inhibited – When you connected the link, active topologies were detected on both sides, and at least one slot number was duplicated. The stack merge is blocked.
- Operational – the port is operational in the stack. This is a necessary but insufficient condition for the port to be used for control path or user data. For example, a node with stacking Failed state may still show its port states as Operational.



The Flags field contains the following flag definitions:

- C - The control path is active on this port. Note that the user data path over the stack links follows the control path.
- B - The port is blocked from transmitting traffic that is to be flooded to multiple non-stacking ports. This flag is only set in an active ring topology on two adjacent ports. In the example below, the active topology is a daisy chain, so no ports are blocked.

Identified with the asterisk, the current node is the one on which the show stacking command was executed. The stack topology is shown in a particular order. In a ring topology, the current node is always the first node, the next node is the node connected to the port 2 of the first node, and the last node is the node connected to the port 1 of the current node. In a daisy chain, the order shown depends on the connection of the node on which the command executes:

- The first node is the one at the far end of the daisy-chain connected to the current node port 1.
- The last node is the one at the far end of the daisy-chain connected to the current node port 2.
- The previous node is the one at the near end of the daisy-chain connected to the current node port 1.
- The next node is the one at the near end of the daisy-chain connected to the current node port 2.
- If there is no node connected to the current node port 1, the current node is the first node.
- If there is no node connected to the current node port 2, the current node is the last node.

The port speed is the unidirectional speed of the port.

Note



Some VIM names include speed ratings which are 4 times the unidirectional stacking port speed. For example, the actual stacking port speed for VIM1-SummitStack512 is 128 Gbps. The 512 Gbps rating for the VIM is the unidirectional rate X 2 (bidirectional) X 2 (ports).

Example

The following example shows the command output for a stack that is operating in a ring and uses both native and alternate stack ports:

```
Slot-1 Stack.9 # show stacking stack-ports
Stack Topology is a Ring
Slot Port Select Node MAC Address Port State Flags Speed
-----
*1 1 23 00:04:96:26:6b:ec Operational C- 10G
*1 2 Native 00:04:96:26:6b:ec Operational CB 64G
2 1 Native 00:04:96:18:7d:e8 Operational CB 64G
2 2 24 00:04:96:18:7d:e8 Operational C- 10G
3 1 23 00:04:96:27:c5:12 Operational C- 10G
3 2 Native 00:04:96:27:c5:12 Operational C- 64G
4 1 Native 00:04:96:26:6b:34 Operational C- 64G
4 2 24 00:04:96:26:6b:34 Operational C- 10G
* - Indicates this node
Flags: (C) Control path is active, (B) Port is Blocked
Slot-1 Stack.10 #
```



The following example shows the command output for stacks that use the 512 Gbps stacking ports:

```
Slot-1 Stack.2 # show stacking stack-ports
Stack Topology is a Ring
Slot Port Select Node MAC Address Port State Flags Speed
-----
*1 1 Native 00:04:96:35:8b:a5 Operational C- 128G
*1 2 Native 00:04:96:35:8b:a5 Operational C- 128G
2 2 Native 00:04:96:35:a8:b0 Operational C- 128G
2 1 Native 00:04:96:35:a8:b0 Operational C- 128G
```

Note



Although the VIM1-SummitStack512 option card has four physical ports, the physical ports are grouped into two pairs, forming two logical ports. The show stacking stack-ports command displays the status of the logical ports.

History

This command was first available in ExtremeXOS 12.0.

The Select column was added in ExtremeXOS 12.5.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show stacking-support

show stacking-support

Description

This command displays the configured and current states of configuration options configured on the local node with the stacking-support keyword.

Syntax Description

This command has no arguments or variables.

Default

N/A.



Usage Guidelines

The display parameters are described in the following table:

Display Item	Description
Stack Port column	Displays rows for Stack Port 1 and Stack Port 2.
Native column	Indicates whether the switch has native stack ports. A Yes entry indicates that the switch has native stacking ports. A No entry indicates that no native stacking ports are present. An asterisk indicates that the native stack port is selected.
Alternate column	Displays the port numbers for data ports that can operate as alternate stack ports. A No entry indicates that there are no data ports that can operate as alternate stacking ports. An asterisk indicates that the data port number is selected as an alternate stack port.
Configured column	Indicates the configured state for stack ports 1 and 2, which can be Native or Alternate. This column also indicates the configured state of stacking-support option. For platform configurations with dual-purpose hardware (that supports stack ports or data ports), this column displays either Enabled (stack ports enabled) or Disabled (stack ports disabled, data ports active). For platform configurations without dual-purpose hardware, this column displays N/A, which indicates that stacking-support option cannot be disabled on this switch. This is the configuration that becomes active the next time the switch boots if the stacking-support option is enabled. For more information, see the command descriptions for the enable stacking-support and disable stacking-support commands.
Current column	Indicates the selection that is currently in effect for stack ports, which can be Native, Alternate, or N/A. N/A indicates that the port selection is not applicable to this switch hardware configuration. The column also indicates the current operating state of the stacking-support option, which can be Enabled, Disabled, or N/A. An N/A entry indicates that no option card is present. This is the configuration that is active now.

Example

The following example shows the stack port selection and stacking-support option configuration after the unconfigure stacking-support command has been executed and before a subsequent reboot has been initiated:

```
show stacking-support
Stacking Support Settings
Stack   Available Ports
Port   Native  Alternate  Configured  Current
-----
1      Yes *   23         Native      Native
2      Yes     24 *       Native      Alternate
stacking-support:           Disabled    Enabled
Flags: * - Current stack port selection
```



NOTE: This node must be rebooted before the configured settings will take effect.

The following example shows that the stacking-support option is disabled and will remain disabled when the switch reboots:

```
show stacking-support
Stacking Support Settings
Stack   Available Ports
Port   Native  Alternate  Configured  Current
-----  -
1      No      S3         Native      N/A
2      No      S4         Native      N/A
stacking-support:      Disabled  Disabled
Flags: * - Current stack port selection
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms listed in [Summit Family Switch Support for Alternate Stack Ports](#) that support alternate stack port selection or permit disabling of the stacking-support option.

synchronize stacking

```
synchronize stacking {node-address node_address | slot slot_number}
```

Description

This command copies certain NVRAM based configuration parameters to the target node.

Syntax Description

<i>node_address</i>	Specifies the MAC address of a node in the stack. To view the MAC addresses for all nodes in a stack, enter the <code>show stacking</code> command.
<i>slot_number</i>	Specifies the slot number of the target node. To view the slot numbers, enter the <code>show stacking</code> command.

Default

N/A.

Usage Guidelines

This command synchronizes the following NVRAM-based parameters:



- Stacking mode.
- Stack MAC address.
- Failsafe account and password.
- Failsafe account access point permissions (whether the failsafe account is allowed over the stacking links, console port, or management port).
- The selected partition.

These parameters are copied from the executing node's NVRAM to the target node's NVRAM.

Example

Example for the synchronize stacking command output:

```
Slot-2 Stack.3 > synchronize stacking slot 3
Are you sure you want to synchronize the specified slot with this slot's
stacking configuration? (y/n) Yes
Synchronized configuration will take effect at the next reboot of the
specified node(s).
Slot-2 Stack.4 >
Slot-2 Stack.4 > synchronize stacking node 00:04:96:27:87:10
Are you sure you want to synchronize the specified node with this node's
stacking configuration? (y/n) Yes
Synchronized configuration will take effect at the next reboot of the
specified node(s).
Slot-2 Stack.5 >
Slot-2 Stack.5 > synchronize stacking
Are you sure you want to synchronize all remote nodes with this node's
stacking configuration? (y/n) Yes
Synchronized configuration will take effect at the next reboot of the
specified node(s).
Slot-2 Stack.6 >
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure stacking

```
unconfigure stacking {node-address node_address | slot slot_number}
```

Description

This command resets most stacking parameters to the default or unconfigured values.



Syntax Description

<i>node_address</i>	Specifies the MAC address of a node in the stack. To view the MAC addresses for all nodes in a stack, enter the <code>show stacking</code> command.
<i>slot_number</i>	Specifies the slot number of the target node. To view the slot numbers, enter the <code>show stacking</code> command.

Default

N/A.

Usage Guidelines

Run this command from any node. If you do not specify a target node, the stacking parameters on all nodes are reset.

This command resets the stacking parameters shown in the following table.

Table 15: Stacking Configuration Items, Time of Effect and Default Value

Configuration Item	Takes Effect	Default Value
Stacking Mode	At boot time	Disabled
Slot Number	At boot time	1
Master-Capable	At boot time	Yes
License Restriction	At boot time	Not Configured
Priority	At next master election	Automatic
Alternate IP Address	Immediately	Not Configured
Stack MAC	At boot time	Not Configured

This command does not reset the stacking parameters configured with the following commands that use the `stacking-support` keyword:

- `configure stacking-support`
- `disable stacking-support`
- `enable stacking-support`

Example

To unconfigure the stacking parameters of all nodes in the stack topology:

```
unconfigure stacking
```

History

This command was first available in ExtremeXOS 12.0.



Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure stacking alternate-ip-address

```
unconfigure stacking {node-address node_address | slot slot_number} alternate-ip-address
```

Description

Removes the configured alternate management IP address from the specified node.

If no node is specified, the alternate management IP address is removed from every node. The change takes effect immediately for all nodes operating in stacking mode.

Syntax Description

<i>node_address</i>	Specifies the MAC address of a node in the stack. To view the MAC addresses for all nodes in a stack, enter the <code>show stacking</code> command.
<i>slot_number</i>	Specifies the slot number of the target active node. To view the slot numbers, enter the <code>show stacking</code> command.

Default

N/A.

Usage Guidelines

Run this command from any node.

Example

To unconfigure stacking alternate-ip-address on a node:

```
unconfigure stacking node-address 00:04:96:26:6b:ed alternate-ip-address
```

To unconfigure the stacking alternate IP address configured on the active node in slot 4:

```
unconfigure stacking slot 4 alternate-ip-address
```

History

This command was first available in ExtremeXOS 12.0.



Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure stacking license-level

```
unconfigure stacking {node-address node_address | slot slot_number} license-level
```

Description

This command removes a previously configured license level restriction.

Syntax Description

<i>node_address</i>	Specifies the MAC address of a node in the stack. To view the MAC addresses for all nodes in a stack, enter the <code>show stacking</code> command.
<i>slot_number</i>	Specifies the slot number of the target node. To view the slot numbers, enter the <code>show stacking</code> command.

Default

N/A.

Usage Guidelines

If no node is specified, the licensing restriction is removed from all nodes in the stack topology.

After the command is executed, the following message appears:

```
This command will take effect at the next reboot of the specified node(s).
```

Example

To unconfigure the stacking license level on a node:

```
unconfigure stacking node-address 00:04:96:26:6b:ed license-level
```

To unconfigure the stacking license level configured on slot 4:

```
unconfigure stacking slot 4 license-level
```



History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure stacking-support

unconfigure stacking-support

Description

This command resets the stacking parameters configured with commands that use the stacking-support keyword.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Run this command from the local node on which you want to reset stacking-support parameters.

This command resets the stacking parameters configured with the following commands that use the stacking-support keyword:

- `configure stacking-support`
- `disable stacking-support`
- `enable stacking-support`

Example

To unconfigure the stacking-support parameters on the local node, use the following command:

```
unconfigure stacking-support
The stacking-support configuration has been reset.
The defaults will take effect at the next reboot of this switch.
```



History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms listed in [Summit Family Switch Support for Alternate Stack Ports](#) that support alternate stack port selection or permit disabling of the stacking-support option.



7 Commands for Configuring Slots and Ports on a Switch

```
clear counters ports
clear lacp counters
clear counters edp
clear slot
configure ces add peer ipaddress
configure ces add peer ipaddress fec-id-type pseudo-wire
configure ces add peer mac-address
configure ces add service
configure ces delete peer
configure ces delete service
configure ces delete peer
configure ces filler-pattern
configure ces jitter-buffer
configure ces lops-threshold
configure ces payload-size
configure ces peer ipaddress static-pw
configure ces qosprofile
configure ces ttl
configure edp advertisement-interval
configure forwarding external-tables
configure forwarding switching-mode
configure ip-fix domain
configure ip-fix flow-key ipv4
configure ip-fix flow-key ipv6
configure ip-fix flow-key nonip
configure ip-fix ip-address
configure ip-fix ports
configure ip-fix ports flow-key ipv4 mask ipaddress
configure ip-fix ports flow-key ipv6 mask ipaddress
configure ip-fix ports record
configure ip-fix source ip-address
configure ip-mtu vlan
configure jumbo-frame-size
configure lacp member-port priority
configure mirror add ports anomaly
```

```
configure mirror add
configure mirror delete
configure mirror description
configure mirror name
configure mirror to port
configure mlag peer interval
configure mlag peer ipaddress
configure MLAG peer lacp-mac
configure mlag ports convergence-control
configure network-clock clock-source input
configure network-clock clock-source output
configure network-clock ptp (priority)
configure network-clock ptp announce interval
configure network-clock ptp announce timeout
configure network-clock ptp boundary add vlan
configure network-clock ptp boundary add unicast-slave
configure network-clock ptp boundary delete unicast-slave
configure network-clock ptp delete
configure network-clock ptp delay-request-interval
configure network-clock ptp end-to-end transparent
configure network-clock ptp ordinary add
configure network-clock ptp sync-interval
configure network-clock ptp add unicast-master
configure network-clock ptp delete unicast-master
configure network-clock sync-e
configure network-clock sync-e clock-source
configure port description-string
configure ports auto off
configure ports auto on
configure ports auto-polarity
configure ports display-string
configure ports dwdm channel
configure ports dwdm channel none
configure ports eee enable
configure ports far-end-fault-indication
configure ports isolation
configure ports link-scan interval
configure ports mode
configure ports partition
configure ports preferred-medium
configure ports redundant
configure ports tdm cable-length
```



```
configure ports tdm clock-source
configure ports tdm display-string
configure ports tdm framing
configure ports tdm idle-code
configure ports tdm line-coding
configure ports tdm recovered-clock
configure ports tdm signaling
configure ports tdm trunk-conditioning
configure ports wan-phy clocking
configure ports wan-phy framing
configure ports wan-phy loopback
configure ports wan-phy trace-path
configure ports wan-phy trace-section
configure sharing add ports
configure sharing address-based custom
configure sharing delete ports
configure sharing health-check member-port add tcp-tracking
configure sharing health-check member-port delete tcp-tracking
configure sharing health-check member-port tcp-tracking
configure sharing lacp activity-mode
configure sharing lacp defaulted-state-action
configure sharing lacp system-priority
configure sharing lacp timeout
configure sharing port-based key
configure slot module
configure slot restart-limit
configure tdm hierarchy
configure tdm service circuit add port
configure tdm service circuit delete port
configure tdm service circuit seized-code
create ces psn
create mirror to port
create mlag peer
create tdm service circuit
delete ces
delete mirror name
delete mlag peer
delete tdm service circuit
disable ces
disable edp ports
disable flow-control ports
disable ip-fix ports
```



```
disable jumbo-frame ports
disable learning port
disable mirror
disable mlag port
disable network-clock ptp end-to-end-transparent ports
disable network-clock sync-e
disable port
disable ports tdm
configure slot module
disable slot
disable smartredundancy
disable snmp traps port-up-down ports
enable ces
enable | disable ces peer ipaddress
enable edp ports
enable flow-control ports
enable | disable ip-fix
enable ip-fix ports
enable jumbo-frame ports
enable learning port
enable mirror
enable mlag port peer id
enable network-clock ptp
enable network-clock ptp end-to-end-transparent ports
enable network-clock sync-e
enable port
enable ports tdm
enable ports tdm loopback
enable sharing grouping
enable slot
enable smartredundancy
enable snmp traps port-up-down ports
restart ports
run failover
run msm-failover
show cdp
show cdp counters
show cdp neighbor
show cdp ports
show ces
show ces clock-recovery
show ces errors
```



```
show ces peer
show dwdm channel-map
show edp
show ip-fix
show lacp
show lacp counters
show lacp lag
show lacp member-port
show mirror
show mirroring
show mlag peer
show mlag ports
show network-clock clock-source
show network-clock sync-e ports
show port eee
show ports
show ports anomaly
show ports buffer
show ports collisions
show ports configuration
show port information
show ports ip-fix
show ports link-scan
show ports packet
show ports redundant
show ports sharing
show ports tdm alarms
show ports tdm configuration
show ports tdm errors
show ports tdm information
show ports tdm no-refresh
show ports transceiver information
show ports transceiver information detail
show ports utilization
show ports wan-phy configuration
show ports wan-phy errors
show ports wan-phy events
show ports wan-phy overhead
show sharing distribution port-based
show sharing health-check
show sharing port-based keys
show slot
```




```
show tdm hierarchy
show tdm service
unconfigure ip-fix
unconfigure ip-fix flow-key
unconfigure ip-fix ip-address
unconfigure ip-fix ports
unconfigure ip-fix ports flow-key mask
unconfigure ip-fix source ip-address
unconfigure mlag peer interval
unconfigure mlag peer ipaddress
unconfigure network-clock sync-e
unconfigure network-clock sync-e clock-source
unconfigure port description-string
unconfigure ports display string
unconfigure ports redundant
unconfigure ports tdm display string
unconfigure ports tdm recovered-clock
unconfigure ports wan-phy
unconfigure slot
```

This chapter describes commands related to:

- Enabling, disabling, and configuring individual ports.
- Configuring port speed (Fast Ethernet ports only) and half- or full-duplex mode.
- Creating link aggregation groups on multiple ports.
- Displaying port statistics.
- Configuring mirroring.
- Configuring software-controlled redundant ports and Smart Redundancy.
- Configuring Extreme Discovery Protocol.
- Configuring time division multiplexing (TDM).

By default, all ports on the switch are enabled. After you configure the ports to your specific needs, you can select which ports are enabled or disabled.

Fast Ethernet ports can connect to either 10BASE-T or 100BASE-T networks. By default, the ports autonegotiate (automatically determine) the port speed. You can also configure each port for a particular speed (either 10 Mbps or 100 Mbps).

In general Gigabit Ethernet ports with fiber interfaces are statically set, and their speed cannot be modified. However, there are two SFPs supported by Extreme Networks that can have a configured speed:

- 100 FX SFPs, which must have their speed configured to 100 Mbps.
- 100FX/1000LX SFPs, which can be configured at either speed.



The switch comes configured to use autonegotiation to determine the port speed and duplex setting for each port. You can choose to manually configure the duplex setting and the speed of 10/100 Mbps ports, and you can manually configure the duplex setting on gigabit Ethernet ports.

All ports on the switch (except gigabit Ethernet ports) can be configured for half-duplex or full-duplex operation. The ports are configured to autonegotiate the duplex setting, but you can manually configure the duplex setting for your specific needs.

Flow control is supported only on gigabit Ethernet ports. It is enabled or disabled as part of autonegotiation. If autonegotiation is set to off, flow control is disabled. When autonegotiation is turned on, flow control is enabled. See the *ExtremeXOS Concepts Guide* for more detailed information on flow control on Extreme Networks devices.

Link aggregation, or load sharing, with Extreme Network switches allows you to increase bandwidth and resilience between switches by using a group of ports to carry traffic in parallel between switches. The sharing algorithm allows the switch to use multiple ports as a single logical port. For example, VLANs see the link aggregation group (LAG) as a single logical port. The algorithm also guarantees packet sequencing between clients.

ExtremeXOS software supports two broad categories of load sharing, or link aggregation: static load sharing and dynamic load sharing.

If a port in a link aggregation group fails, traffic is redistributed to the remaining ports in the LAG. If the failed port becomes active again, traffic is redistributed to include that port.

You can view port status on the switch using the `show ports` commands. These commands, when used with specific keywords and parameters, allow you to view various issues such as collision statistics, link speed, flow control, and packet size. Beginning with ExtremeXOS software version 11.3, these port information displays show real-time statistics, or you can configure the display to show a snapshot of real-time statistics (as in earlier versions of the software).

Beginning with ExtremeXOS version 11.6 software, you can configure WAN PHY OAM on those interfaces that connect 10G Ethernet ports to the SONET/SDH network.

Commands that require you to enter one or more port numbers use the parameter `port_list` in the syntax. On a modular switch or SummitStack, a `port_list` can be a list of slots and ports. On a stand-alone switch, a `port_list` can be one or more port numbers. For a detailed explanation of port specification, see [Port Numbering in Command Reference Overview](#).

For synchronous Ethernet (SyncE), the following ports are supported on each platform:

- X460-24X: Input Ports 1-28, Output Ports 1 - 28.
- X460-48X: Input Ports 1-48, Output Ports 1 - 48.
- E4G-200: All Ethernet ports.
- E4G-400: All Ethernet ports including XGM3S ports if present.

The E4G-200 and E4G-400 have clock sources beyond SyncE. The clock which drives all of the ports on a switch may be selected from:

- SyncE.
- PTP – an optional 1588v2 module.



- TDM – an optional module that has multiple T1/E1 interfaces for TDM/Ethernet interworking.
- BITS – Building Integrated Timing Supply. A connector capable of receiving a timing signal provided by other building equipment.

clear counters ports

clear counters ports

Description

Clears the counters associated with the ports.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines



Note

If you use the clear counters command with no keyword, the system clears the counters for all applications.

This command clears the counters for the ports, including the following:

- Statistics.
- Transmit errors.
- Receive errors.
- Collisions.
- Packets.

Example

The following command clears the counters on all ports:

```
clear counters ports
```

History

This command was first available in ExtremeXOS 11.3.



Platform Availability

This command is available on all platforms.

clear lacp counters

clear lacp counters

Description

Clears the counters associated with Link Aggregations Control Protocol (LACP).

Syntax Description

This command has no parameters or variables.

Default

N/A.

Usage Guidelines

This command clears the following counters for LACP; it sets these counters back to 0 for every LACP port on the device:

- LACP PDUs dropped on non_LACP ports.
- Stats:
 - Rx - Accepted.
 - Rx - Dropped due to error in verifying PDU.
 - Rx - Dropped due to LACP not being up on this port.
 - Rx - Dropped due to matching own MAC.
 - Tx - Sent Successfully.
 - Tx - Transmit error.

Example

The following command clears the LACP counters on all ports:

```
clear lacp counters
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.



clear counters edp

```
clear counters edp {ports ports}
```

Description

Clears the counters associated with Extreme Discovery Protocol (EDP).

Syntax Description

<i>ports</i>	Specifies one or more ports or slots and ports.
--------------	---

Default

If you do not specify a port, the EDP counters will be cleared for all ports.

Usage Guidelines

This command clears the following counters for EDP protocol data units (PDUs) sent and received per EDP port:

- Switch PDUs transmitted.
- VLAN PDUs transmitted.
- Transmit PDUs with errors.
- Switch PDUs received.
- VLAN PDUs received.
- Received PDUs with errors.

Example

The following command clears the EDP counters on all ports:

```
clear counters edp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

clear slot

```
clear slot slot
```



Description

Clears a slot of a previously assigned module type.

Syntax Description

<i>slot</i>	Specifies the slot number.
-------------	----------------------------

Default

N/A.

Usage Guidelines

All configuration information related to the slot and the ports on the module is erased. If a module is present when you issue this command, the module is reset to default settings.

If a slot is configured for one type of module, and a different type of module is inserted, the inserted module is put into a mismatch state (where the inserted module does not match the configured slot), and is not brought online. To use the new module type in a slot, the slot configuration must be cleared or configured for the new module type. Use the `enable mirroring to port tagged` command to configure the slot.

Example

The following command clears slot 2 of a previously assigned module type:

```
clear slot 2
```

The following command clears slot 4 of a previously assigned module type in a stack:

```
clear slot 4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on modular switches and SummitStack.

configure ces add peer ipaddress

```
configure ces ces_name add peer ipaddress ipaddress fec-id-type pseudo-wire pw_id
{static-pw transmit-label outgoing_pw_label receive-label incoming_pw_label} {lsp
lsp_name}
```



Description

Statically configures a new MPLS TDM PW for the specified CES.

Syntax Description

<i>ces_name</i>	Specifies the circuit emulation service (CES).
<i>ipaddress</i>	Specifies an IP address.
<i>pw_id</i>	Specifies the pseudowire ID.
fec-id-type pseudo-wire	Specifies the ???.
static-pw transmit-label	Specifies the static pseudowire transmit label.
<i>outgoing_pw_label</i>	Specifies the outgoing static label. The <i>outgoing_pw_label</i> must match the peer's configured incoming PW label.
receive-label	Specifies the label that you apply to the static PW on egress.
<i>incoming_pw_label</i>	Specifies the egress label.
lsp	Specifies the Label Switch Path through the network.
<i>lsp_name</i>	Specifies the name of the LSP.

Default

N/A.

Usage Guidelines

This command statically configures a new MPLS TDM PW for the specified CES. Both the outgoing (MPLS ingress) and incoming (MPLS egress) PW labels must be specified. The peer must be similarly configured with a static PW that has the reverse PW label mappings. Locally, the *incoming_pw_label* must be unique and is allocated out of the static label space. The *outgoing_pw_label* must match the peer's configured incoming PW label.

Optionally, you can configure the PW to use any type of tunnel LSP: LDP, RSVP-TE, or Static. In the case of RSVP-TE and LDP, those protocols must be configured and enabled, and an LSP must be established, before traffic can be transmitted over the static PW. For Static LSPs, only the MPLS ingress LSP (or outgoing LSP) is specified. Unlike signaled PWs, there is no end-to-end PW communication that is used to verify that the PW endpoint is operational, and in the case of static LSPs, that the data path to the PW endpoint is viable.

In the event of a network fault, if a secondary RSVP-TE LSP is configured or the routing topology changes such that there is an alternate LDP LSP, the static PW will automatically switch LSPs in order to maintain connectivity with the PW endpoint. Static LSPs can be protected proactively by configuring BFD to verify the static LSPs IP next hop connectivity.

Optionally, the underlying LSP for the PW can be explicitly specified using a named LSP. When a named LSP is explicitly specified, only the specified named LSP is used to carry the PW. In the event that a specified named LSP is withdrawn, the CES remains operationally down until the named LSP is restored.



Example

The following command adds a static pseudowire to:

```
configure iproute add 10.1.1.0/24 lsp lsp598
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on the E4G-200 and E4G-400 platforms.

configure ces add peer ipaddress fec-id-type pseudo-wire

```
configure ces ces_name add peer ipaddress ipaddress [fec-id-type pseudo-wire
pw_id {lsp lsp_name} | udp-port local src_udp_port remote dst_udp_port vlan
vlan_name]
```

Description

This command is used to configure a new MPLS TDM PW for the specified CES. The signaled PW parameters are passed to the peer using Targeted LDP over the IP network. The peer specified identifies the endpoint of the PW. The *pw_id* parameter uniquely identifies PW service and cannot conflict with any other configured service in the network. The value is signaled and used to negotiate the PW labels between the two PW endpoint peers.

The PW is immediately signaled once the associated service information is known (e.g., TDM T1 Port, TDM E1 Port, Ethernet VLAN, etc.) and provided the administrative status of the CES is enabled. The enable setting for the CES peer does not affect the signaling of the PW. This setting only affects the preferred forwarding status. Configuration information associated with the service is signaled to the peer and both PW endpoints must have a compatible service attachment. If the service attachments are not compatible, the PW is not established and an error message is logged.

Optionally, the underlying LSP for the PW can be explicitly specified using a named LSP. When a named LSP is explicitly specified, only the specified named LSP is used to carry the PW. In the event that a specified named LSP is withdrawn, the switch signals a PW Status Code of (PSN Facing-TX Fault, Forwarding Preference-Standby). The CES remains operationally down until the named LSP is restored.

Syntax Description

ces	Circuit Emulation Service.
<i>ces_name</i>	Alphanumeric string identifying the CES pseudowire.
IPaddress <i>ipaddress</i>	Specifies the peer IPv4 address.



fec-id-type	FEC ID type.
pseudo-wire	Pseudowire FEC.
<i>pw_id</i>	Specifies the pseudowire VPN ID. The range is 1 - 4294967295.
lsp	Specifies the LSP.
<i>lsp_name</i>	Alpha numeric string identifying the LSP.

Default

N/A.

Usage Guidelines

Use this command to configure a new MPLS TDM PW for the specified CES.

Example

```
create ces ces-test psn mpls
configure ces "ces-test" add peer ipaddress 1.1.1.1 fec-id-type pseudo-wire
100
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

configure ces add peer mac-address

```
configure ces ces_name add peer mac-address mac_address ecid local tx_ecid remote
rx_ecid vlan vlan_name
```

Description

Manually adds an Ethernet (MEF-8) peer (far-end) for the specified CES pseudowire. The `<peer mac-address, rx_ecid>` is used to de-multiplex CES pseudo-wires in the CE-bound direction. The `<switch mac-address, tx_ecid>` parameters are used to encapsulate the packets in the PE-bound direction.



Syntax Description

<i>ces_name</i>	Alphanumeric string identifying the CES pseudowire.
<i>mac_address</i>	Peer MAC address.
<i>tx_ecid</i>	The local MEF-8 emulated circuit identifier of the CES pseudowire.
<i>rx_ecid</i>	The remote MEF-8 emulated circuit identifier of the CES pseudowire.
<i>vlan_name</i>	VLAN name of the layer-2 forwarding domain of the CES pseudowire.

Default

N/A.

Usage Guidelines

See Description.

Example

```
create ces ces-test psn mef
configure ces "ces-test" add peer mac-address 00:49:00:11:22:54 ecid local
2000 remote 3000 vlan v1
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

configure ces add service

```
configure ces ces_name add service service_name
```

Description

Adds the TDM service to the specified CES pseudowire.

Syntax Description

<i>ces_name</i>	Alphanumeric string identifying the CES pseudowire.
<i>service_name</i>	Provider-provisioned CES service name.



Default

N/A.

Usage Guidelines

Structure-agnostic (SAToP) pseudo-wires are configured by adding a structure-agnostic TDM service. Structure-aware (CESoP) pseudo-wires are configured by adding a structure-aware TDM service.

Example

```
#create tdm service circuit service-test
#configure tdm service circuit service-test add port 35 unframed
#create ces ces-test psn udp
#configure ces ces-test add service service-test
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

configure ces delete peer

```
configure ces ces_name delete peer [ipaddress ipaddress | mac-address mac_address]
```

Description

Deletes the peer of the specified CES pseudowire.

Syntax Description

<i>ces_name</i>	Alphanumeric string identifying the CES pseudowire.
-----------------	---

Default

N/A.

Usage Guidelines

Deletes the peer of the specified CES pseudowire.



Example

```
#configure ces ces-test delete peer ip-address 1.1.1.1
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

configure ces delete service

```
configure ces ces_name delete service
```

Description

Deletes the TDM service from the specified CES pseudowire.

Syntax Description

<i>ces_name</i>	Alphanumeric string identifying the CES pseudowire.
-----------------	---

Default

N/A.

Usage Guidelines

N/A.

Example

```
configure ces ces-test delete service
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).



configure ces delete peer

```
configure ces ces_name delete peer [ipaddress ipaddress | mac-address mac_address]
```

Description

Deletes the peer of the specified CES pseudowire.

Syntax Description

<i>ces_name</i>	Alphanumeric string identifying the CES pseudowire.
-----------------	---

Default

N/A.

Usage Guidelines

Deletes the peer of the specified CES pseudowire.

Example

```
#configure ces ces-test delete peer ip-address 1.1.1.1
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

configure ces filler-pattern

```
configure ces ces_name filler-pattern byte_value
```

Description

Configures the filler pattern of the specified CES pseudowire.



Syntax Description

<i>ces_name</i>	Alphanumeric string identifying the CES pseudowire.
<i>byte_value</i>	Filler pattern to be played out on a timeslot bound to the CES pseudowire. Allowed values are between 0 and 255, with a default of 255.

Default

The default value is 255.

Usage Guidelines

Note that this pattern is played out only on timeslots bound to the CES pseudowire. For unused timeslots, i.e., for timeslots not bound to a CES pseudowire, the fixed pattern of 0xff will be played out.

Example

```
E4G400#configure ces ces-test filler-pattern 100
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

configure ces jitter-buffer

```
configure ces ces_name jitter-buffer min_jbf {max max_jbf}
```

Description

Configures the jitter-buffer value to be used in the CE-bound direction for the specified CES pseudowire.

Syntax Description

<i>ces_name</i>	Alphanumeric string identifying the CES pseudowire.
ms	Jitter-buffer value in milliseconds between 1 and 200.



Default

E1 and T1 Unframed Min. 3 ms and max 6 ms E1 Basic Framing Min. 3 ms and max 6 ms E1 MF Min. 4 ms and max 8 ms T1 SF and ESF (without signaling) Min. 3 ms and max 6 ms T1 SF and ESF (with signaling) Min. 6 ms and max 12 ms

Usage Guidelines

The allowed values for the jitter-buffer are between 1 and 200 milliseconds.

Example

```
E4G200#configure ces ces-test jitter-buffer 2000 max 5000
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

configure ces lops-threshold

```
configure ces ces_name lops-threshold [entry num_packets_for_entry {exit num_packets_for_exit} | exit num_packets_for_exit]
```

Description

Configures the LOPS (loss of packet state) threshold for the specified CES pseudowire. The threshold can be specified for entry, for exit, or for both.

Syntax Description

<i>ces_name</i>	Alphanumeric string identifying the CES pseudowire.
<i>num_packets_for_entry</i>	The number of consecutive packet misses required to enter the LOPS. The value can be between 1 and 15, with a default of 8.
<i>num_packets_for_exit</i>	The number of consecutive packets required to exit the LOPS. The value can be between 1 and 10, with a default of 8.

Default

The default is 8 for both entry and exit thresholds.



Usage Guidelines

The number of consecutive packet misses required to enter the LOPS can be between 1 and 15, and the number of consecutive packets required to exit the LOPS can be between 1 and 10.

Example

```
E4G-400#configure ces "ces-test" lops-threshold entry 2 exit 5
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

configure ces payload-size

```
configure ces ces_name payload-size bytes
```

Description

Configures the payload-size value in the PE-bound direction for the specified CES pseudowire.

Syntax Description

<i>ces_name</i>	Alphanumeric string identifying the CES pseudowire.
<i>bytes</i>	Payload size in bytes between 1 and 576 (default is 256 for E1 and 192 for T1).

Default

The default size for a CES pseudowire transporting structure-agnostic E1 TDM service is 256 bytes, and the default size is 192 bytes for structure-agnostic T1 TDM service.

Usage Guidelines

The allowed value for the payload size is between 1 and 512 bytes.

Example

```
E4G200#configure ces ces-test payload 512
```



History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

configure ces peer ipaddress static-pw

```
configure ces ces_name peer ipaddress ipaddress static-pw [{transmit-label
outgoing_pw_label} {receive-label incoming_pw_label}]
```

Description

Change labels for TDM Circuit Emulation Service over MPLS Static PW.

Syntax Description

<i>ces_name</i>	Specifies the circuit emulation service (CES).
<i>ipaddress</i>	Specifies an IP address.
static-pw transmit-label	Specifies the static pseudowire transmit label.
<i>outgoing_pw_label</i>	Specifies the outgoing static label. The <i>outgoing_pw_label</i> can be any value, and must match the peer's configured incoming PW label.
receive-label	Specifies the label that you apply to the static PW on egress .
<i>incoming_pw_label</i>	Specifies the egress label. The <i>incoming_pw_label</i> must be locally unique, and is allocated out of the static label space

Default

N/A.

Usage Guidelines

This command allows a network administrator to modify the current configuration of a PW for the specified *ces_mpls_name*. The *incoming_pw_label* must be locally unique, and is allocated out of the static label space. The *outgoing_pw_label* can be any value, and must match the peer's configured incoming PW label.

Example

The following command changes the PW name:

```
configure ces new_ces peer ipaddress 10.1.1.0/24 static-pw transmit-label
new_pw lsp lsp598
```



History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on the E4G-200 and E4G-400 platforms.

configure ces qosprofile

```
configure ces ces_name qosprofile qosprofile
```

Description

Configures the QoS (quality of service) profile to be associated to a Circuit Emulation Service pseudo-wire.

Syntax Description

<i>ces_name</i>	Alphanumeric string identifying the CES pseudowire.
<i>qosprofile</i>	Name of switch fabric QoS profile (default is QP1).

Default

The default QoS profile is QP1.

Usage Guidelines

Use this command to configure the QoS profile for a CES pseudo-wire.

Example

```
E4G200#configure ces ces-test qosprofile qp4
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).



configure ces ttl

```
configure ces ces_name ttl ttl_value
```

Description

Configures the TTL (time-to-live) for the specified CES pseudo-wire.

Syntax Description

<i>ces_name</i>	Alphanumeric string identifying the CES pseudowire.
<i>ttl_value</i>	Time-To-Live for the specified CES pseudowire. The TTL can be between 1 and 254, with a default of 254.

Default

The default TTL is 254.

Usage Guidelines

The following error message will be displayed if the TTL is configured for incompatible CES pseudo-wires (when the PSN type of the CES pseudowire is MEF-8):

```
Error: TTL option is incompatible with the configured CES pseudo-wire  
Packet-switched Network (PSN) type.
```

Example

```
E4G-400#configure ces ces-test ttl 10
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

configure edp advertisement-interval

```
configure edp advertisement-interval timer holddown-interval timeout
```

Description

Sets the advertisement interval and hold down interval for EDP.



Syntax Description

<i>timer</i>	Specifies the advertisement interval in seconds.
<i>timeout</i>	Specifies the hold down interval in seconds.

Default

The default setting for timer is 60 seconds, and for timeout is 180 seconds.

Usage Guidelines

Extreme Discover Protocol (EDP) is used to gather information about neighbor Extreme Networks switches. EDP-enabled ports advertise information about the Extreme switch to other switches on the interface and receive advertisements from other Extreme switches. Information about other Extreme switches is discarded after the hold down interval timeout value is reached without receiving another advertisement.

Example

The following command configures the EDP advertisement-interval to 2 minutes and the hold down interval to 6 minutes:

```
configure edp advertisement-interval 120 holddown-interval 360
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure forwarding external-tables

```
configure forwarding external-tables [13-only {ipv4 | ipv4-and-ipv6 | ipv6} | 12-only | acl-only | 12-and-13 | 12-and-13-and-acl | 12-and-13-and-ipmc | none]
```

Description

Customizes the use of the external memory.



Syntax Description

l3-only	Programs the external lookup table to store Layer 3 routes only. (Default is IPv4 only).
l3-only ipv4-and-ipv6	Program the external lookup table for IPv4 routes and hosts, and IPv6 routes.
l3-only ipv6	Program the external lookup table for IPv6 routes and hosts.
l2-only	Programs the external lookup table to store Layer 2 MAC FDB only.
acl-only	Programs the external lookup table to store access-lists only.
l2-and-l3	Programs the external lookup table to store Layer 2 MAC FDB and Layer 3 routes (this is the default).
l2-and-l3-and-acl	Programs the external lookup table to store Layer 2 MAC FDB, Layer 3 routes, and access-lists.
l2-and-l3-and-ipmc	Programs the external lookup table to store Layer 2 MAC FDB, Layer 3 routes, and IP multicast groups.
none	Specifies that the external lookup table is not used.

Default

l2-and-l3.

Usage Guidelines

Use this command to set the use of the external forwarding table memory. This external memory can be configured in various ways to support extending either one internal table (such as, Layer 2) or many internal tables (such as Layer 2, Layer 3, and ACL).

Following are the table limits for each of the options on the BlackDiamond 8900 xl-series modules and Summit X480 series switches:

Layer 3 only (IPv4)	512K IPv4 routes.
Layer 3 only (IPv4and IPv6)	464K IPv4 routes and 48K IPv6 routes prefix length 0 to 64 bits.
Layer 3 only (IPv6)	240K IPv6 routes prefix length 0 to 128 bits.
Layer 2 only	512K MAC FDB entries.
ACL only	60K ACLs.
Layer 2 and Layer 3	256K MAC FDB entries + 256K IPv4 routes.
Layer 2, Layer 3, and ACL	128K MAC FDB entries + 128K IPv4 routes + 56K ACLs.
Layer 2, Layer 3, and IP multicast	128K MAC FDB entries + 128K IPv4 routes + 12000 IP multicast FDBs (SourceIP, groupIP, vlanId).
None	External tables will not be used.

The configuration applies to the entire system.

Changing the use of the external memory cannot be done during runtime. After a change in configuration, a reboot is required for the change to take effect.



To display the current configuration, use the `show forwarding configuration` command.

Example

The following command configures the external tables to store ACLs:

```
BD-8810.1 # config forwarding external-tables acl-only
WARNING: This command will take effect after a save and reboot.
```

History

This command was first available in ExtremeXOS 12.4.

The l2-and-l3-and-ipmc option was added in ExtremeXOS 12.5.

The l3-only ipv4-and-ipv6 and l3-only ipv6 options were added in ExtremeXOS 15.2.

Platform Availability

This command is available on Summit X480 series switches and BlackDiamond 8800 series switches for the configuration of the external tables on the 8900 xl-series modules.

configure forwarding switching-mode

```
configure forwarding switching-mode [cut-through | store-and-forward]
```

Description

Configures the switching mode as either cut-through or store-and-forward.

Syntax Description

cut-through	Specifies that a packet can begin being transmitted prior to its being received in full.
store-and-forward	Specifies that a packet is transmitted only after the entire packet has been received and stored in the packet memory.

Default

Store-and-forward.

Usage Guidelines

Use this command to configure the switch to begin transmitting a packet before its entire contents have been received. This reduces the forwarding latency of the switch.



Cut-through mode cannot be achieved for packet sizes that are less than or equal to 384 bytes.

To display the switch mode settings, use the `show forwarding configuration` command.

When issued on a BlackDiamond 8800 series switch, this command does not affect module types other than those listed in the Platform Availability section below.

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on BlackDiamond X8 switches and BlackDiamond 8800 series switches with 8900-10G24X-c and 8900-MSM-128 series modules.

The BlackDiamond 8900-G96T-c I/O module supports cut-through switching only in the switching fabric and impacts cross-chip and cross-slot switching only.

configure ip-fix domain

```
configure ip-fix domain domain_id
```

Description

Configures an observation domain ID.

Syntax Description

<i>domain_id</i>	Specifies a decimal integer.
------------------	------------------------------

Default

Domain 0.

Usage Guidelines

Use this command to set an observation domain ID that is used in the flow records sent to the collector. The collector can then use this ID to correlate records to their origin.

The entire switch operates as one domain.



Example

The following command configures a domain ID of 4 for the switch:

```
configure ip-fix domain 4
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on BlackDiamond 8900 G96Tc, G48T-xl, G48X-xl, and 10G8X-xl modules and Summit X460 and X480 switches.

configure ip-fix flow-key ipv4

```
configure ip-fix flow-key ipv4 {src-ip} {src-port} {dest-ip} {dest-port}
{protocol} {tos}
```

Description

Configures the settings for the flow key(s) for IPv4.

Syntax Description

src-ip	Specifies the source IP address field as part of the flow key.
src-port	Specifies the L4 source port field as part of the low key.
dest-ip	Specifies the destination IP address field as part of the flow key.
dest-port	Specifies the L4 destination port field as part of the flow key.
protocol	Specifies the L4 protocol field as part of the flow key.
tos	Specifies the type of service field as part of the flow key.

Default

All flow keys.

Usage Guidelines

Use this command to specify which of the designated flow-keys to use. This overrides the default which is all keys. The template sent to the Collector (per the IPFIX standard) contains only the keys used. Then, on a per port basis, you can define masks for the IPv4 source and destination address fields, for instance, to aggregate flows based on subnets. (see `configure ip-fix ports flow-key ipv4 mask ipaddress`)



The size of the field (in bits) for each key is as follows:

- Source IP Address (32).
- Destination IP Address (32).
- L4 Source Port (16).
- L4 Destination Port (16).
- L4 Protocol (8).
- TOS (DSCP + ECN) (8).

To unconfigure, use the `unconfigure ip-fix flow-key` command.

To display the flow keys use the `show ip-fix` command.

Example

The following command configures IPv4 traffic to use the source IP address and L4 protocol:

```
configure ip-fix flow-key ipv4 src-ip protocol
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on BlackDiamond 8900 G96Tc, G48T-xl, G48X-xl, and 10G8X-xl modules and Summit X460 and X480 switches.

configure ip-fix flow-key ipv6

```
configure ip-fix flow-key ipv6 {src-ip} {src-port} {dest-ip} {dest-port} {next-hdr} {tos} {flow-label}
```

Description

Configures the settings for the flow key(s) for IPv6.

Syntax Description

src-ip	Specifies the source IP address field as part of the flow key.
src-port	Specifies the L4 source port field as part of the flow key.
dest-ip	Specifies the destination IP address field as part of the flow key.
dest-port	Specifies the L4 destination port field as part of the flow key.
next-hdr	Specifies the next header field as part of the flow key.



itos	Specifies type of service field as part of the flow key.
flow-label	Specifies IPv6 flow-label field as part of the flow key.

Default

All flow keys.

Usage Guidelines

Use this command to specify which of the designated flow-keys to use. This overrides the default which is all keys. The template sent to the Collector (per the IPFIX standard) contains only the keys used. Then, on a per port basis, you can define masks for the IPv6 source and destination address fields, for instance, to aggregate flows based on subnets.

The size of the field (in bits) for each key is as follows:

- Source IP Address (128).
- Destination IP Address (128).
- L4 Source Port (16).
- L4 Destination Port (16).
- Next Header (8).
- IPv6 Flow Label (20).
- TOS (DSCP + ECN) (8).

To unconfigure, use the `unconfigure ip-fix flow-key` command.

To display the configured flow keys, use the `show ip-fix` command.

Example

The following command configures IPv6 traffic to use the destination IP address and next header:

```
configure ip-fix flow-key ipv6 dest-ip next-hdr
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on BlackDiamond 8900 G96Tc, G48T-xl, G48X-xl, and 10G8X-xl modules and Summit X460 and X480 switches.



configure ip-fix flow-key nonip

```
configure ip-fix flow-key nonip {src-mac} {dest-mac} {ethertype} {vlan-id}
{priority} {tagged}
```

Description

Configures the settings for the flow key(s) for non-IP type data.

Syntax Description

src-mac	Specifies the source MAC address field as part of the flow key.
dest-mac	Specifies the destination MAC address field as part of the flow key.
ethertype	Specifies the Ethertype field as part of the flow key.
vlan-id	Specifies the VLAN ID field as part of the flow key.
priority	Specifies the VLAN priority field as part of the flow key.
tagged	Specifies the VLAN tagged field as part of the flow key.

Default

All flow keys.

Usage Guidelines

Use this command to specify which of the designated flow-keys to use. This overrides the default which is all keys. The template sent to the Collector (per the IPFIX standard) contains only the keys used.

The size of the field (in bits) for each key is as follows:

- Source MAC Address (48).
- Destination MAC Address (48).
- Ethertype (16).
- VLAN ID (12).
- VLAN Priority (3).
- VLAN Tagged (1).

To unconfigure, use the `unconfigure ip-fix flow-key` command.

To display the configured flow keys, use the `show ip-fix` command.

Example

The following command configures non-IP traffic to use the source MAC address and VLAN ID:

```
configure ip-fix flow-key src-mac vlan-id
```



History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on BlackDiamond 8900 G96Tc, G48T-xl, G48X-xl, and 10G8X-xl modules and Summit X460 and X480 switches.

configure ip-fix ip-address

```
configure ip-fix ip-address ipaddress {protocol [sctp | tcp | udp]} {L4-port
portno} {vr vrname}
```

Description

Identifies the collector and how communication with it is handled.

Syntax Description

<i>ipaddress</i>	Specifies the IP address.
sctp	Specifies SCTP.
tcp	Specifies TCP.
udp	Specifies UDP. This is the default.
<i>portno</i>	Specifies the number of an L4 port. The default is 4739.
<i>vrname</i>	Specifies a VR.

Default

The protocol field will default to UDP. The L4-port field will default to 4739. The VR field will default to VR-Mgmt.

Usage Guidelines

Use this command to specify the IP address, port number, transport protocol and VR for a collector.

To unconfigure the settings, use the `unconfigure ip-fix ip-address` command.

To display the collector settings, use the `show ip-fix` command.

Example

The following command specifies a collector with an IP address of 1.1.1.1, and transport protocol of TCP:

```
configure ip-fix ip-address 1.1.1.1 protocol tcp
```



History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on BlackDiamond 8900 G96Tc, G48T-xl, G48X-xl, and 10G8X-xl modules and Summit X460 and X480 switches.

configure ip-fix ports

```
configure ip-fix ports port_list [ingress | egress | ingress-and-egress]
```

Description

Configures metering on ingress and/or egress ports.

Syntax Description

<i>port_list</i>	Specifies the ports.
ingress	Specifies ingress ports only.
egress	Specifies egress ports only.
ingress-and-egress	Specifies both ingress and egress ports.

Default

Ingress.

Usage Guidelines

Use this command to configure metering on ingress and/or egress ports.

Example

The following command configures metering on port 2 egress:

```
configure ip-fix ports 2 egress
```

History

This command was first available in ExtremeXOS 12.5.



Platform Availability

This command is available on BlackDiamond 8900 G96Tc, G48T-xl, G48X-xl, and 10G8X-xl modules and Summit X460 and X480 switches.

configure ip-fix ports flow-key ipv4 mask ipaddress

```
configure ip-fix ports port_list flow-key ipv4 mask [source | destination]
ipaddress value
```

Description

Defines masks for the IPv4 source and destination address flow keys.

Syntax Description

<i>port_list</i>	Specifies the ports.
source	Specifies a source IP address.
destination	Specifies a destination IP address.
<i>value</i>	Specifies the IP address mask (in standard format).

Default

N/A.

Usage Guidelines

Use this command to define masks for the IPv4 source and destination address flow keys on a per port basis. For example, this can be used to minimize the information sent to the collector and aggregate flows.

Example

The following command defines a mask for source IP address flow key 255.255.0.0 on port 2.1:

```
configure ip-fix ports 2.1 flow-key ipv4 mask source ipaddress 255.25.0.0
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on BlackDiamond 8900 G96Tc, G48T-xl, G48X-xl, and 10G8X-xl modules and Summit X460 and X480 switches.



configure ip-fix ports flow-key ipv6 mask ipaddress

```
configure ip-fix ports port_list flow-key ipv6 mask [source | destination]
ipaddress value
```

Description

Defines masks for the IPv6 source and destination address flow keys.

Syntax Description

<i>port_list</i>	Specifies the ports.
source	Specifies a source IP address.
destination	Specifies a destination IP address.
<i>value</i>	Specifies the IP address mask (in standard format).

Default

N/A.

Usage Guidelines

Use this command to define masks for the IPv6 source and destination address flow keys on a per port basis. For example, this can be used to minimize the information sent to the collector and aggregate flows.

Example

The following command defines a mask for the source IP address flow key ff::0 on port 2.1:

```
configure ip-fix ports 2.1 flow-key ipv6 mask source ipaddress ff::0
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on BlackDiamond 8900 G96Tc, G48T-xl, G48X-xl, and 10G8X-xl modules and Summit X460 and X480 switches.

configure ip-fix ports record

```
configure ip-fix ports port_list record [all | dropped-only | non-dropped]
```



Description

Configures metering on all, dropped only, or non dropped traffic.

Syntax Description

<i>port_list</i>	Specifies the ports.
all	Specifies all packets.
dropped-only	Specifies only dropped packets.
non-dropped	Specifies only non-dropped packets.

Default

All.

Usage Guidelines

Use this command to configure metering on all packets, only dropped packets, or only non-dropped packets.

Example

The following command configures metering dropped packets only:

```
configure ip-fix ports 2:1 record dropped-only
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on BlackDiamond 8900 G96Tc, G48T-xl, G48X-xl, and 10G8X-xl modules and Summit X460 and X480 switches.

configure ip-fix source ip-address

```
configure ip-fix source ip-address ipaddress {vr vrname}
```

Description

Configures the source IP address used to communicate to the collector.



Syntax Description

<i>ipaddress</i>	Specifies the source IP address to be used in IPFIX packets.
<i>vrname</i>	Specifies a virtual router name.

Default

Switch IP address of the interface the traffic egresses.

Usage Guidelines

Use this command to specify the source IP address and VR to use when sending from the switch to a given collector. Otherwise, the default is used.

There is one collector.

To reset to the default of the switch IP address, use the `unconfigure ip-fix ip-address` command.

Example

The following command configures an IP address of 1.1.1.1 and VR of finance:

```
configure ip-fix source ip-address 1.1.1.1 finance
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on BlackDiamond 8900 G96Tc, G48T-xl, G48X-xl, and 10G8X-xl modules and Summit X460 and X480 switches.

configure ip-mtu vlan

```
configure ip-mtu mtu vlan vlan_name
```

Description

Sets the maximum transmission unit (MTU) for the VLAN.



Syntax Description

<i>mtu</i>	Specifies the IP maximum transmission unit (MTU) value. Range is from 1500 to 9194.
<i>vlan_name</i>	Specifies a VLAN name.

Default

The default IP MTU size is 1500.

Usage Guidelines

BlackDiamond 8000 c-, e-, and xl-series modules, BlackDiamond X8 series switches, and Summit X440, X460, X480, X670, and X770 series switches support IP fragmentation and path MTU discovery.

Use this command to enable jumbo frame support or for IP fragmentation with jumbo frames. Jumbo frames are Ethernet frames that are larger than 1522 bytes, including 4 bytes used for CRC. Both endstations involved in the transfer must be capable of supporting jumbo frames. The switch does not perform IP fragmentation or participate in MTU negotiation on behalf of devices that do not support jumbo frames.

When enabling jumbo frames and setting the MTU size for the VLAN, keep in mind that some network interface cards (NICs) have a configured maximum MTU size that does not include the additional 4bytes of CRC included in a jumbo frame configuration. Ensure that the NIC maximum MTU is at or below the maximum MTU size configured on the switch. Frames that are larger than the MTU size configured on the switch are dropped at the ingress port.

If you use IP fragmentation with jumbo frames and you want to set the MTU size greater than 1500, all ports in the VLAN must have jumbo frames enabled.

Example

The following command sets the MTU size to 2000 for VLAN sales:

```
configure ip-mtu 2000 vlan sales
```

History

This command was available in ExtremeXOS 11.0.

Platform Availability

This command is available on BlackDiamond X8 switches, the BlackDiamond 8000 c-, e-, xl-, and xm-series modules, and the Summit X440, X460, X480, X670, and X770 series switches (whether or not included in a SummitStack).



configure jumbo-frame-size

```
configure jumbo-frame-size framesize
```

Description

Sets the maximum jumbo frame size for the switch.

Syntax Description

<i>framesize</i>	Specifies a maximum transmission unit (MTU) size for a jumbo frame. The range is 1523 to 9216; the default is 9216.
------------------	---

Default

Jumbo frames are disabled by default. The default size setting is 9216.

Usage Guidelines

Jumbo frames are used between endstations that support larger frame sizes for more efficient transfers of bulk data. Both endstations involved in the transfer must be capable of supporting jumbo frames.

The framesize keyword describes the maximum jumbo frame size “on the wire,” and includes 4 bytes of cyclic redundancy check (CRC) plus another 4 bytes if 802.1Q tagging is being used.

To enable jumbo frame support, you must configure the maximum transmission unit (MTU) size of a jumbo frame that will be allowed by the switch.



Note

Extreme Networks recommends that you set the MTU size so that fragmentation does not occur.

Some network interface cards (NICs) have a configured maximum MTU size that does not include the additional 4 bytes of CRC. Ensure that the NIC maximum MTU size is at or below the maximum MTU size configured on the switch. Frames that are larger than the MTU size configured on the switch are dropped at the ingress port.

Example

The following command configures the jumbo frame size to 5500:

```
configure jumbo-frame-size 5500
```

History

This command was first available in ExtremeXOS 10.1.



Platform Availability

This command is available on all platforms.

configure lacp member-port priority

```
configure lacp member-port port priority port_priority
```

Description

Configures the member port of an LACP to ensure the order that ports are added to the aggregator. The lower value you configure for the port's priority, the higher priority that port has to be added to the aggregator.

Syntax Description

<i>port</i>	Specifies the LACP member port that you are specifying the priority for.
<i>port_priority</i>	Specifies the priority you are applying to this member port to be assigned to the LACP aggregator. The range is from 0 to 65535; the default is 0. The lower configured value has higher priority to be added to the aggregator.

Default

The default priority is 0.

Usage Guidelines

The port must be added to the LAG prior to configuring it for LACP. The default value is 0, or highest priority.

You can configure the port priority to ensure the order in which LAG ports join the aggregator. If you do not configure this parameter, the lowest numbered ports in the LAG are the first to be added to the aggregator; if there are additional ports configured for that LAG, they are put in standby mode.

Use this command to override the default behavior and ensure the order in which LAG ports are selected. Also, if more than one port is configured with the same priority, the lowest numbered port joins the aggregator.

Example

The following command sets the port priority for the LAG port 5:1 to be 55 (which will probably put that port in standby initially):

```
configure lacp member-port 5:1 priority 55
```

History

This command was first available in ExtremeXOS 11.4.



Platform Availability

This command is available on all platforms.

configure mirror add ports anomaly

```
configure mirror add ports port_list anomaly
```

Description

Mirrors detected anomaly traffic to the mirror port.

Syntax Description

<i>port_list</i>	Specifies the list of ports.
------------------	------------------------------

Default

N/A.

Usage Guidelines

The command mirrors detected anomaly traffic to the mirror port. You must enable a mirror port and enable protocol anomaly protection on the slot that has the port to be monitored before using this command. After configuration, only detected anomaly traffic from these ports are dropped or mirrored to the mirror port, and legitimate traffic is not affected.

This command takes effect after enabling anomaly-protection.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

configure mirror add




```
configure mirror mirror_name add [vlan name | port port {ingress | egress | ingress-and-egress | anomaly}]
```

Description

Specifies mirror source filters for an instance.



Syntax Description

<i>mirror_name</i>	Specifies a VLAN.
vlan	Specifies a VLAN.
<i>name</i>	Specifies a VLAN name.
port	Specifies a port or slot and port.
<i>port</i>	Specifies particular ports or slots and ports.
ingress	<p>Specifies packets be mirrored as they are received on a port.</p> <hr/> <p>Note  This parameter is available only on BlackDiamond X8 series switches, BlackDiamond 8800 series switches, SummitStack, and Summit family switches. For BlackDiamond X8 series switches, BlackDiamond 8800 series switches, SummitStack, and Summit family switches, it is available only with port-based mirroring.</p>
egress	<p>Specifies packets be mirrored as they are sent from a port.</p> <hr/> <p>Note  This parameter is available only on BlackDiamond X8 series switches, BlackDiamond 8800 series switches, SummitStack, and the Summit family switches. For BlackDiamond X8 series switches, BlackDiamond 8800 series switches, SummitStack, and Summit family switches, it is available only with port-based mirroring.</p>
ingress-and-egress	<p>Specifies all forwarded packets be mirrored. This is the default setting on BlackDiamond 8800 series switches, SummitStack, and Summit family switches for port-based mirroring.</p> <hr/> <p>Note  This parameter is available only on BlackDiamond X8 series switches, BlackDiamond 8800 series switches, SummitStack, and Summit family switches. For BlackDiamond X8 series switches, BlackDiamond 8800 series switches, SummitStack, and Summit family switches, it is available only with port-based mirroring.</p>

Default

N/A.

Usage Guidelines

You must enable port-mirroring using the `enable mirroring to port` command before you can configure the mirroring filter definitions.

Port mirroring configures the switch to copy all traffic associated with one or more ports to a monitor port on the switch. The switch uses a traffic filter that copies a group of traffic to the monitor port.

Up to 128 mirroring filters can be configured with the restriction that a maximum of 16 of these can be configured as VLAN and/or virtual port (port + VLAN) filters.



One monitor port or 1 monitor port list can be configured. A monitor port list may contain up to 16 ports.

Frames that contain errors are not mirrored.

For general guideline information and information for various platforms, see [Guidelines for Mirroring](#) in the *ExtremeXOS Concepts Guide* or Usage Guidelines in the `enable mirroring to port` command discussion.

Example

The following example sends all traffic coming into a BlackDiamond 8800 series switch on slot 3, port 2 to the mirror port:

```
configure mirror add port 3:2 ingress
```

The following example sends all traffic coming into a switch on port 11 and the VLAN default to the mirror port:

```
configure mirror add port 11 vlan default
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

configure mirror delete

```
configure mirror mirror_name delete [all | vlan name | port port | vlan name port port | port port vlan name]
```

Description

Deletes mirror source filters for an instance.

Syntax Description

<i>mirror_name</i>	Specifies all mirroring filter definitions.
all	Specifies all mirroring filter definitions.
<i>port</i>	Specifies a port or a slot and port.



<i>port</i>	Specifies particular ports or slots and ports.
vlan	Specifies a VLAN.
<i>name</i>	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

On a modular switch, *port_list* must be a slot and port in the form *slot:port*. For a detailed explanation of port specification, see [Port Numbering in Command Reference Overview](#)

Example

The following example deletes the mirroring filter on an BlackDiamond 8800 series switch defined for slot 7, port 1:

```
configure mirroring delete ports 7:1
```

History

This command was first available in ExtremeXOS 10.1.

The VLAN option was added in ExtremeXOS 11.0.

The VLAN mirroring capability was added to the BlackDiamond 8800 series switch in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure mirror description

```
configure mirror mirror_name description [ mirror-desc | none ]
```

Description

Creates, edits or deletes a mirroring instance description string.



Syntax Description

<i>mirror_name</i>	Specifies the mirror name.
description	Specifies the mirror description to create or edit.
none	Specifies t.

Default

Disabled.

Usage Guidelines

Use this command to create, edit or delete a mirroring instance description string.

Example

The following example ... :

```
configure mirror description
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

configure mirror name

```
configure mirror mirror_name name new_name
```

Description

Updates or specifies the "to port" definitions for a named mirroring instance .

Syntax Description

<i>mirror_name</i>	Specifies the mirror name.
name	Specifies the mirror name.

Default

Disabled.



Usage Guidelines

Use this command to update or specify the "to port" definitions for a named mirroring instance.

Example

```
configure mirror name
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

configure mirror to port

```
configure mirror mirror_name {to [port port | port_list port_list | loopback port  
port ] {remote-tag rtag | port none}
```

Description

Updates or specifies the "to port" definitions for a named mirroring instance.

Syntax Description

<i>mirror_name</i>	Specifies the mirror name.
port	Specifies the mirror output port.
port_list	Specifies the list of ports where traffic is to be mirrored.
loopback port	Specifies an otherwise unused port required when mirroring to a <i>port_list</i> . The loopback-port is not available for switching user data traffic.
<i>port</i>	Specifies a single loopback port that is used internally to provide this feature.
remote-tag	Specifies the value of the VLAN ID used by the mirrored packets when egressing the monitor port.
port	Specifies ? .
none	Specifies ? .

Default

Disabled.



Usage Guidelines

Use this command to update, or specify the "to port" definitions for a named mirroring instance.

Example

The following example configures a mirror instance to port 3, slot 4 :

```
configure mirror to port 3:4
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

configure mlag peer interval

```
configure mlag peer peer_name interval msec
```

Description

Configures the length of time between health check hello packets.

Syntax Description

<i>peer_name</i>	Specifies an alpha numeric string identifying the MLAG peer.
<i>msec</i>	Specifies an MLAG peer health-check hello interval in milliseconds. The range is 50-10000ms. The default is 1000ms.

Default

The interval default is 1000 milliseconds.

Usage Guidelines

Use this command to configure the length of time between health check hello packets exchanged between MLAG peer switches. After three health check hellos are lost, the MLAG peer switch is declared to be failed, triggering an MLAG topology change.



Example

The following command sets an interval of 700 milliseconds on the switch101 peer. switch:

```
configure mlag peer switch101 interval 700
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

configure mlag peer ipaddress

```
configure mlag peer peer_name ipaddress peer_ip_address {vr VR}
```

Description

Associates an MLAG peer switch with an MLAG peer structure.

Syntax Description

<i>peer_name</i>	Specifies an alpha numeric string identifying the MLAG peer.
<i>peer_ip_address</i>	Specifies an IPv4 or IPv6 address.
<i>VR</i>	Specifies a virtual router.

Default

N/A.

Usage Guidelines

Use this command to associate an MLAG peer structure with an MLAG peer switch IP address.

The specified IP address must be contained within an existing direct route. If not, the following error message is displayed:

```
ERROR: Specified IP address is not on directly attached subnet in VR.
```

The link connecting MLAG peer switches should use load sharing. If it does not, a output similar to the following is displayed:

```
Note: VLAN v1 will be used as the Inter-Switch Connection to the MLAG peer mp1. Warning: The VLAN v1 does not have a load share port configured yet. It is recommended that the Inter-Switch Connection use load sharing.
```



Example

The following command associates the MLAG peer structure switch101 with the MLAG peer switch IP address 1.1.1.1 on VR-USER:

```
configure mlag peer switch101 ipaddress 1.1.1.1 vr "VR-USER"
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

configure MLAG peer lacp-mac

```
configure mlag peer peer_name lacp-mac [auto | lacp_mac_address]
```

Description

Configures MLAG LACP MAC on each of the MLAG peer switches. This MAC address will be used as the system identifier in the LACPDUs sent over the MLAG ports.

Syntax Description

mlag	Multi-switch link aggregation used to combine remote ports and local ports to a common logical connection.
<i>peer_name</i>	Alphanumeric string identifying the MLAG peer.
lacp-mac	MAC address to be used as the system identifier in LACPDU for MLAG ports.
auto	System identifier in LACPDU automatically uses switch MAC of MLAG peer with higher IP address for ISC control VLAN (default).
<i>lacp_mac_address</i>	MAC address.

Default

Auto.

Usage Guidelines

This command is used to configure the System Identifier used in LACPDU for MLAG ports. The same value has to be configured on both the MLAG peers.



Example

```
configure mlag peer "peer1" lACP-mac auto
configure mlag peer "peer1" lACP-mac 00:01:02:03:04:05
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

configure mlag ports convergence-control

```
configure mlag ports convergence-control [conserve-access-lists | fast]
```

Description

Sets a preference for having a fast convergence time or conserving access lists.

Syntax Description

conserve-access-lists	Specifies that conserving access lists is preferred over low traffic convergence time.
fast	Specifies that low traffic convergence time is preferred at the expense of the number of user access lists.

Default

Conserve-access-lists.

Usage Guidelines

Achieving fast convergence times on local MLAG port state changes (down and up), independent of the number of FDB entries learned on the MLAG port, requires the use of ACLs. This limits the number of ACLs you have available. This command allows you to set your preference for having either fast convergence time or conserving available access lists for your users.

Note



Configuring fast convergence-control limits the number of ACLs that can be supported by the switch. You must ensure that the system has sufficient user ACLs free when fast mode is selected. Configuring conserve-access-lists convergence-control may increase convergence times on MLAG port failures.

Fast convergence configuration has global significance in that it applies to all MLAG groups that are currently configured and those that may be configured in the future.



Example

The following command specifies a priority of conserving access lists over low traffic convergence time:

```
configure mlag ports convergence-control conserve-access-lists
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

configure network-clock clock-source input

```
configure network-clock clock-source input {[sync-e | ptp | tdm | [bits-rj45 | bits-bnc]} {quality-level value} | region [E1 | T1]
```

Description

Configure the input network clock source as Sync-E or PTP or TDM or BITS A region can also be configured using this command.

Syntax Description

sync-e	Synchronous Ethernet (ITU-T standard) (Default).
ptp	Precision Time Protocol.
tdm	Time Division Multiplexing.
bits-rj45	External bits clock from RJ45.
bits-bnc	External bitsw clock from BNC connector.
<i>value</i>	Value of the quality level of the clock (T1 default QL_ST3, E1 default QL_SEC).
E1	Specifies the European and Asian clock region selection (Default).
T1	Specifies the North American clock region selection.

Default

The default input clock source is Synchronous Ethernet. The default region is E1.

Usage Guidelines

The E4G-200 and E4G-400 have clock sources beyond SyncE. The clock which drives all of the ports on a switch may be selected from:



- SyncE (Synchronous Ethernet).
- PTP – Precision Time Protocol, an optional 1588v2 module.
- TDM – an optional module which has multiple T1/E1 interfaces for TDM/Ethernet interworking.
- BITS – Building Integrated Timing Supply. A connector capable of receiving a timing signal provided by other building equipment.

Example

The following command configures the region as T1:

```
configure network-clock clock-source input region t1
```

The following command configures the network-clock source as TDM:

```
configure network-clock clock-source input tdm
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on X460-24x, X460-48x E4G-200 and E4G-400 switches. On X460 platforms only Sync-E option is available as clock-source.

configure network-clock clock-source output

```
configure network-clock clock-source output {bits-bnc-1 [1pps | 8KHz] bits-bnc-2 [E1 | T1 | 10MHz]}
```

Description

Configure the output network clock source as bits bnc 1 or 2.

Syntax Description

bits-bnc-1	Bits Clock 1 BNC connector.
1pps	1pps output clock.
8KHz	8 KHz output clock (default).
bits-bnc-2	Bits Clock 2 BNC connector.
E1	E1 (2.048 MHz) output clock (default).
T1	T1 (1.544 MHz) output clock.
10MHz	10 MHz output clock.



Default

The default output clock source for Bits Clock 1 BNC connector is 8 KHz, and the default for Bits Clock 2 BNC connector is E1 (2.048 MHz).

Usage Guidelines

The E4G-200 and E4G-400 have clock sources beyond SyncE. The clock which drives all of the ports on a switch may be selected from:

- SyncE (Synchronous Ethernet).
- PTP – Precision Time Protocol, an optional 1588v2 module.
- TDM – an optional module which has multiple T1/E1 interfaces for TDM/Ethernet interworking.
- BITS – Building Integrated Timing Supply. A connector capable of receiving a timing signal provided by other building equipment.

Example

The following command configures the region as T1:

```
configure network-clock clock-source input region t1
```

The following command configures the network-clock source as TDM:

```
configure network-clock clock-source input tdm
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on X460-24x, X460-48x, E4G-200 and E4G-400 switches. On X460 platforms only sync-e option is available as clock-source.

configure network-clock ptp (priority)

```
configure network-clock ptp [boundary | ordinary] [priority1 | priority2]  
priority
```

Description

Assign priority1 and priority2 values for a PTP clock instance.



Syntax Description

boundary	Boundary clock instance.
ordinary	Ordinary clock instance.
priority1	Priority1 of the clock.
priority2	Priority2 of the clock.
<i>priority</i>	Value of priority. 0 is the highest priority. The default value is 128, the range is 0-255.

Default

The default value of the priority is 128. The range is 0-255.

Usage Guidelines

Use this command to assign priority1 and priority2 values for a PTP clock instance. This command is available only for boundary and ordinary clocks.

Example

The following example assigns priority1 and priority2 values for the boundary clock:

```
configure network-clock ptp boundary priority1 50
configure network-clock ptp boundary priority2 128
```

The following example assigns priority1 and priority2 values for ordinary clock:

```
configure network-clock ptp ordinary priority1 10
configure network-clock ptp ordinary priority2 200
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).



Note

PTP commands can be used only with the Network Timing feature pack.



configure network-clock ptp announce interval

```
configure network-clock ptp [boundary | ordinary] announce interval
seconds_log_base_2 {vlan} vlan_name
```

Description

Configure the value of the PTP announce interval time on the port(s) for sending announce messages. This command is available only for boundary and ordinary clocks. The interval time is the time between successive announce messages.

Syntax Description

boundary	Boundary clock instance.
ordinary	Ordinary clock instance.
<i>seconds_log_base_2</i>	The log base 2 value of the seconds between successive announce messages. For example, to specify 8 seconds, use 3. The default value is 1, and the range is -2 to 4.
<i>vlan_name</i>	Name of the VLAN to apply the command to.

Default

The default value of the announce interval is log base 2 (1), or 2 seconds.

Usage Guidelines

Use this command to configure the value of PTP announce interval time on the port(s) for sending announce messages. The announce interval is set using log base 2 values. The range is -2 to 4. This command is available only for boundary and ordinary clocks.

Example

The following example configures announce interval to be 1/second on the clock port lpbk-gm of the ordinary clock:

```
configure network-clock ptp ordinary announce interval 0 vlan lpbk-gm
```

The following example configures announce interval to be 2/second on the clock-port lpbk-gm of the boundary clock:

```
configure network-clock ptp boundary announce interval -1 vlan lpbk-gm
```

History

This command was first available in ExtremeXOS 15.1.



Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).



Note

PTP commands can be used only with the Network Timing feature pack.

configure network-clock ptp announce timeout

```
configure network-clock ptp [boundary | ordinary] announce timeout
seconds_log_base_2 {vlan} vlan_name
```

Description

Configure the value of the timeout for receiving PTP announce messages. This command is available only for boundary clocks and ordinary clocks.

Syntax Description

boundary	Boundary clock instance.
ordinary	Ordinary clock instance.
<i>seconds_log_base_2</i>	The log base 2 value of the seconds of the timeout for PTP announce messages. For example, to specify 8 seconds, use 3. The default value is 3, and the range is 2 to 8.
<i>vlan_name</i>	VLAN name to which the command is to be applied.

Default

The default value of the announce timeout is log base 2 (3), or 8 seconds.

Usage Guidelines

Use this command to configure the value of the timeout for PTP announce messages. The clock port should not have unicast static slaves or masters added for this configuration to be applied.

Example

The following example configures announce timeout interval to be 16 seconds on the clock port lpbk-gm of the ordinary clock:

```
configure network-clock ptp ordinary announce timeout 4 vlan lpbk-gm
```



The following example configures announce timeout interval to be 4 seconds on the clock port lpbk-gm of the boundary clock:

```
configure network-clock ptp boundary announce timeout 2 vlan lpbk-gm
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).



Note

PTP commands can be used only with the Network Timing feature pack.

configure network-clock ptp boundary add vlan

```
configure network-clock ptp boundary add {vlan} vlan_name {one-step | two-step}
{slave-only | master-only}
```

Description

Add a VLAN to a PTP boundary clock instance as a clock port. You can configure the clock port as slave-only port or master-only port.

Syntax Description

boundary	Boundary clock instance.
<i>vlan_name</i>	Name of the specific VLAN to be added to or deleted from the PTP clock instance.
one-step	1-step protocol mode (default).
two-step	2-step protocol mode.
slave-only	Force clock port to be slave.
master-only	Force clock port to be master.

Default

The default configuration of clock port is master or slave mode.

Usage Guidelines

Use this command to add a VLAN to a PTP boundary clock instance as a clock port. You can configure the clock port as slave-only port, or master-only port. The slave-only clock port has the PTP port state



forced to slave. The slave port does not respond to signaling messages from other slaves, and Sync/DelayResponse event messages are not generated by the slave-only ports.

The master-only clock port has the PTP port state forced to master. The master port generates Sync/DelayResponse event messages to downstream slave clocks.

The default configuration of clock port is master or slave mode. In this mode, the clock port state is based on the Best Master Clock (BMC) algorithm running on the port. The BMC algorithm decides the clock port state transition to master or slave depending on the event messages received on the clock port from the associated unicast master(s)/slave(s).

The following restrictions apply on the VLAN to be added as clock port:

- The VLAN should be tagged.
- Loopback-mode should be enabled on the VLAN.
- IPv4 address should be configured and IP forwarding must be enabled on the VLAN.
- The VLAN should not have front panel ports added.

Example

The following example adds a vlan 'lpbk-gm' as a slave clock port to boundary clock:

```
configure network-clock ptp boundary add vlan lpbk-gm one-step slave-only
```

The following example adds a vlan 'lpbk-ord' as a master clock port to boundary clock in two-step protocol mode:

```
configure network-clock ptp boundary add vlan lpbk-ord two-step master-only
```

The following example adds a vlan 'lpbk-transit' as a clock port to boundary clock whose master/slave state is decided by BMCA:

```
configure network-clock ptp boundary add vlan lpbk-transit one-step
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).



Note

PTP commands can be used only with the Network Timing feature pack.



configure network-clock ptp boundary add unicast-slave

```
configure network-clock ptp boundary add unicast-slave ipv4_address {vlan}
vlan_name
```

Description

Add an entry to the unicast slave table for a PTP clock instance. This command is available only for boundary clocks.

Syntax Description

boundary	Boundary clock instance.
unicast-slave	IP addresses that are potential slaves to the local clock.
<i>ipv4_address</i>	IPv4 address.
vlan	VLAN.
<i>vlan_name</i>	Name of the specific VLAN to add for the PTP clock instance.

Default

The default configuration of clock port is slave mode.

Usage Guidelines

Use this command to add an entry to the unicast slave table for a PTP clock instance.

Example

The following command adds a static unicast slave entry to the PTP clock port lpbk-slave in the boundary clock:

```
configure network-clock ptp boundary add unicast-slave 192.168.15.20 vlan
lpbk-slave
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).



Note

PTP commands can be used only with the Network Timing feature pack.



configure network-clock ptp boundary delete unicast-slave

```
configure network-clock ptp boundary delete unicast-slave ipv4_address {vlan}
vlan_name
```

Description

Deletes an entry from the unicast slave table for a PTP clock instance. This command is available only for boundary clocks.

Syntax Description

boundary	Boundary clock instance.
unicast-slave	IP addresses that are potential slave to the local clock.
<i>ipv4_address</i>	IPv4 address.
vlan	VLAN.
<i>vlan_name</i>	Name of the specific VLAN to add for the PTP clock instance.

Default

The default configuration of clock port is slave mode.

Usage Guidelines

Use this command to delete an entry from the unicast slave table for a PTP clock instance.

Example

The following command removes a static unicast master entry from boundary clock:

```
configure network-clock ptp boundary delete unicast-slave 192.168.1.1 vlan
lpbk-gm
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).



Note

PTP commands can be used only with the Network Timing feature pack.



configure network-clock ptp delete

```
configure network-clock ptp [boundary | ordinary] delete [{vlan} vlan_name | vlan all]
```

Description

Deletes a given clock port (VLAN), or all clock ports added to the specified PTP clock instance.

Syntax Description

boundary	Boundary clock instance.
ordinary	Ordinary clock instance.
vlan	VLAN.
<i>vlan_name</i>	Name of the specific VLAN to be added to or deleted from the PTP clock instance.
all	Add or delete all VLANs.

Default

N/A.

Usage Guidelines

Use this command to delete a given clock port (VLAN), or all clock ports added to the specified PTP clock instance.

Example

The following example deletes all clock ports from the boundary clock:

```
configure network-clock ptp boundary delete vlan all
```

The following example deletes vlan lpbk-gm from the ordinary clock:

```
configure network-clock ptp ordinary delete vlan lpbk-gm
```

History

This command was first available in ExtremeXOS 15.1.



Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).



Note

PTP commands can be used only with the Network Timing feature pack.

configure network-clock ptp delay-request-interval

```
configure network-clock ptp [boundary | ordinary] delay-request-interval
seconds_log_base_2 {vlan} vlan_name
```

Description

Configures the value of PTP delay request interval time to send successive delay request messages when the port is in the master state.

Syntax Description

boundary	Boundary clock instance.
ordinary	Ordinary clock instance.
<i>seconds_log_base_2</i>	The log base 2 value of the seconds of the timeout for PTP announce messages. For example, to specify 8 seconds, use 3. The default value is -6, and the range is -7 to 5.
<i>vlan_name</i>	VLAN name to which the command is to be applied.

Default

The default value of the announce delay request interval is log base 2 (-6), or 64 delay request messages per second.

Usage Guidelines

Use this command to configure the value of PTP delay request interval time to send delay request messages when the port is in the master state. The clock port should not have unicast static slaves or masters added to apply the configuration.

Set the delay request interval using log base 2 values. The range is -7 to 5. This command is available only for boundary and ordinary clocks.

Example

The following command configures the delay request message rate of 32/second on the clock port lpbk-gm of the ordinary clock:

```
configure network-clock ptp ordinary delay-request-interval -5 vlan lpbk-gm
```



The following command configures the delay request message rate of 128/second on the clock port lpbk-gm of the boundary clock:

```
configure network-clock ptp boundary delay-request-interval -7 vlan lpbk-gm
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).



Note

PTP commands can be used only with the Network Timing feature pack.

configure network-clock ptp end-to-end transparent

```
configure network-clock ptp end-to-end-transparent [add | delete] ports port_list
{one-step}
```

Description

Adds or deletes the physical port(s) to or from the end-to-end-transparent clock.

Syntax Description

add	Add ports.
delete	Delete ports.
ports	Physical ports.
<i>port_list</i>	List of ports to be added or deleted.
one-step	One step operation.

Default

N/A.

Usage Guidelines

Use this command to add or delete the physical port(s) to, or from, the end-to-end-transparent clock. The fiber only 1G ports, 10G ports, and stack ports cannot be added to the End-to-End transparent clock.



Example

The following example configures end-to-end transparent clock on the front panel ports:

```
configure network-clock ptp end-to-end-transparent add ports 1-4 one-step
```

The following example deletes the front panel ports from the end-to-end transparent clock:

```
configure network-clock ptp end-to-end-transparent delete ports 2-4
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).



Note

PTP commands can be used only with the Network Timing feature pack.

configure network-clock ptp ordinary add

```
configure network-clock ptp ordinary add {vlan} vlan_name {one-step | two-step}
slave-only
```

Description

Add a VLAN to a PTP ordinary clock instance as a clock port.

Syntax Description

ordinary	Ordinary clock instance.
<i>vlan_name</i>	Name of the specific VLAN to be added to or deleted from the PTP clock instance.
one-step	1-step protocol mode (default).
two-step	2-step protocol mode.
slave-only	Force clock port to be a slave.

Default

The default protocol mode on the clock port is one-step.



Usage Guidelines

Use this command to add a VLAN to a PTP ordinary clock instance as a clock port. The ordinary clock master (grand-master) mode of operation is not supported.

The following restrictions apply on the VLAN to be added as clock port:

- The VLAN should be tagged.
- Loopback-mode should be enabled on the VLAN.
- IPv4 address should be configured and IP forwarding must be enabled on the VLAN.
- The VLAN should not have front panel ports added.

Example

The following example adds a vlan 'lpbk-gm' as a slave clock port to ordinary clock:

```
configure network-clock ptp ordinary add vlan lpbk-gm one-step slave-only
```

The following example adds a vlan 'lpbk-gm2' as a slave clock port to ordinary clock in two-step mode:

```
configure network-clock ptp ordinary add vlan lpbk-gm2 two-step slave-only
```

History

This command was first available in ExtremeXOS 15.1 Revision 2.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).



Note

PTP commands can be used only with the Network Timing feature pack.

configure network-clock ptp sync-interval

```
configure network-clock ptp [boundary | ordinary] sync-interval  
seconds_log_base_2 {vlan} vlan_name
```

Description

Configure the value of PTP Sync message interval time to send Sync event message to the slaves when the port is in the master state.



Syntax Description

boundary	Boundary clock instance.
ordinary	Ordinary clock instance.
sync-interval	Time between successive sync messages.
<i>seconds_log_base_2</i>	The log base 2 value of the seconds of the timeout for PTP announce messages. For example, to specify 2 seconds, use 1. The default value is -6, and the range is -7 to 1.
<i>vlan_name</i>	VLAN name to which the command is to be applied.

Default

The default value of the *seconds_log_base_2* parameter is -6.

Usage Guidelines

Use this command to configure the announce message rate on the clock port. The clock port should not have unicast static slaves or masters added for this configuration to be applied.

This command is available only for boundary and ordinary clocks.

Example

The following command configures the sync message rate of 2/second on the clock port lpbk-gm of the boundary clock:

```
configure network-clock ptp boundary sync-interval -1 vlan lpbk-gm
```

The following command configures the sync message rate of 8/second on the clock port lpbk-gm of the ordinary clock:

```
configure network-clock ptp ordinary sync-interval -3 vlan lpbk-gm
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).



Note

PTP commands can be used only with the Network Timing feature pack.



configure network-clock ptp add unicast-master

```
configure network-clock ptp [boundary | ordinary] add unicast-master ipv4_address
{query-interval seconds_log_base_2} [{vlan} vlan_name]
```

Description

Adds an entry to the unicast master table for a PTP clock instance. This command is available for boundary and ordinary clocks.

Syntax Description

boundary	Boundary clock.
ordinary	Ordinary clock.
add unicast-master	IP addresses that are potential master to the local clock.
<i>ipv4_address</i>	IPv4 address.
query-interval	Mean interval between requests from a node for a unicast Announce message
<i>seconds_log_base_2</i>	The log base 2 value in seconds of the mean interval between requests from a node for a unicast Announce message. For example, to specify 8 seconds between requests, use 3. The default value is 1, and the range is -2 to 4.
<i>vlan_name</i>	VLAN name to which the command is to be applied.

Default

The default value of the query interval is log base 2 (1), or 2 seconds mean interval between requests from the node for a unicast Announce Message.

Usage Guidelines

Use this command to add an entry to the unicast master table for a PTP clock instance. This command is available only for boundary clocks. The mean interval between requests from the node for a unicast Announce message can be configured from 1/4 second to 16 seconds, with a default of 2 seconds.

Example

The following command adds a static unicast master entry to the PTP clock port lpbk-gm in the boundary clock:

```
configure network-clock ptp boundary add unicast-master 192.168.1.1 query-
interval 0 vlan lpbk-gm
```



The following command adds a static unicast master entry to the PTP clock port lpbk-master in the ordinary clock clock:

```
configure network-clock ptp ordinary add unicast-master 192.168.15.10 vlan
lpbk-master
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).



Note

PTP commands can be used only with the Network Timing feature pack.

configure network-clock ptp delete unicast-master

```
configure network-clock ptp [boundary | ordinary] delete unicast-master
ipv4_address [{vlan} vlan_name]
```

Description

Delete an entry from the unicast master table for a PTP clock instance. This command is available only for boundary clocks.

Syntax Description

boundary	Boundary clock.
ordinary	Ordinary clock.
<i>ipv4_address</i>	IPv4 address of a master to the local clock.
<i>vlan_name</i>	VLAN name to which the command is to be applied.

Default

N/A.

Usage Guidelines

Use this command to delete an entry from the unicast master table for a PTP clock instance. This command is available only for boundary clocks.



Example

The following command removes a static unicast master entry from boundary clock:

```
configure network-clock ptp boundary delete unicast-master 192.168.1.1 vlan
lpbk-gm
```

The following command removes a static unicast master entry from ordinary clock:

```
configure network-clock ptp ordinary delete unicast-master 192.168.15.10 vlan
lpbk-master
```

History

This command was first available in ExtremeXOS 15.1 Revision 2.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).



Note

PTP commands can be used only with the Network Timing feature pack.

configure network-clock sync-e

```
configure network-clock sync-e [source-1 | source-2] port port
```

Description

Configures synchronous Ethernet on a particular port to be a source 1 or source 2 for synchronizing clock.

Syntax Description

source-1	Source 1 external input clock
source-2	Source 2 external input clock
<i>port</i>	port
<i>port</i>	100Mbps/1G port Copper/Fiber Ports

Default

None of the ports are source 1 or 2.



Usage Guidelines

Use this command to configure SyncE on a particular port to be a primary master or secondary master for synchronizing the clock.

The following lists the ports supported on each platform:

- X460-24X: Ports 1 - 28.
- X460-48X: Ports 1 - 48.
- E4G-200: All Ethernet ports.
- E4G-400: All Ethernet ports including XGM3S ports if present.

If you attempt to configure SyncE on a management port, the following message is displayed:

```
ERROR: Synchronous Ethernet is not supported on the Mgmt port.
```

If you attempt to configure more than one source-1 or source-2 port, the following message is displayed:

```
ERROR: Only one port can be configured as source-1/source-2.
```

If you attempt to configure SyncE on a port that is not supported, the following message is displayed:

```
ERROR: Cannot Configure Synchronous Ethernet on ports
```

To unconfigure SyncE, use the `unconfigure network-clock sync-e` command.

To display SyncE settings, use the `show network-clock sync-e ports` command.

Example

The following command configures port 2 as SyncE source.

```
configure network-clock sync-e source-1 port 2
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on Summit X460-24X and X460-48X switches and on E4G-200 and E4G-400 switches.

configure network-clock sync-e clock-source

```
configure network-clock sync-e clock-source [source-1|source-2] {quality-level value}
```



Description

Configures synchronous ethernet clock-source to be a source 1 or source 2 for synchronizing clock.

Syntax Description

source-1	Source 1 external input clock.
source-2	Source 2 external input clock.
<i>value</i>	Value of the Quality level of the clock (T1 default QL_ST3, E1 default QL_SEC).

Default

None of the ports are source 1 or 2.

Usage Guidelines

Use the following command to configure "source-1" as the Synchronous Ethernet clock-source. It generates ESMC messages with quality level QL_PRC.

```
configure network-clock sync-e clock-source source-1 quality-level QL_PRC
```

If no quality-level is specified, the default value is used (T1 default QL_ST3, E1 default QL_SEC).

Example

```
E4G-400.33 # configure network-clock sync-e clock-source source-1 quality-  
level QL_PRC
```

configure port description-string

```
configure ports port_list description-string string
```

Description

Configures a description string setting up to 255 characters.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>string</i>	Specifies a port description of up to 255 characters per port. You cannot use the following characters: ' " ' "<" ">" ":" "<space>" "&"

Default

None.



Usage Guidelines

Use this command to configure a port description of up to 255 characters per port.

In case that user configures a string longer than 64 chars, the following warning will be displayed:

```
Port description strings longer than 64 chars are only accessible through
SNMP if the following command is issued: configure snmp ifmib ifalias
size extended
```

Some characters are not permitted as they have special meanings. These are: ‘ ‘, “<”, “>”, “:”, “<space>”, “&”. The first character should be alphanumeric. This new field is CLI accessible only via “show port info detail” but is also accessible via the SNMP ifAlias object of IfXTable from IF-MIB (RFC 2233) and the XML API. In order to access the value via SNMP the following command should be issued: `configure snmp ifmib ifalias size extended`.

Example

The following command configures the port:

```
configure ports 1:3 description-string CorporatePort_123
```

History

This command was available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

configure ports auto off

```
configure ports port_list {medium [copper | fiber]} auto off speed speed duplex
[half | full]
```

Description

Manually configures port speed and duplex setting configuration on one or more ports on a switch.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
medium	Specifies the medium as either copper or fiber. Note: This parameter applies to combo ports only on the Summit family switches.
<i>speed</i>	Specifies the port speed as either 10, 100, 1000 (1 Gigabit), or 10000 (10 Gigabit) Mbps ports.



duplex [half]	Specifies half duplex; transmitting and receiving data one direction at a time.
duplex [full]	Specifies full duplex; transmitting and receiving data at the same time.

Default

Auto on for 1G ports.

Auto off for 10G ports.

Usage Guidelines

You can manually configure the duplex setting and the speed on 10/100 and 10/100/1000 Mbps and fiber SFP gigabit Ethernet ports.

In general, SFP gigabit Ethernet ports are statically set to 1 Gbps, and their speed cannot be modified. However, there are SFPs supported by Extreme Networks that can have a configured speed:

- 100 FX SFPs, which must have their speed configured to 100 Mbps.
- 100FX/1000LX SFPs, which can be configured at either speed (available only on the BlackDiamond 8800 series switches, SummitStack, and Summit family switches).
- SFP+ optics, must have their speed configured to 10G auto off.

The 8900-10G24X-c, BDXA-10G48X, Summit X670-48x, and X670V-48x ports default to 10G auto off. The user is recommended to change the configuration to 1G auto on (or auto off) when a 1G SFP optic is plugged into the port. In the event, the user does not change user configuration, the port links up at 1G, but a warning message is logged for the user to take corrective action and change the user configuration. Similarly, when the port has been configured to 1G auto on (or auto off) and a SFP+ optic is inserted, the port links up with the correct speed and auto configuration, but a log message recommends that the user configuration be changed to match the optic inserted.

In certain interoperability situations, it is necessary to turn autonegotiation off on a fiber gigabit Ethernet port. Even though a gigabit Ethernet port runs only at full duplex and gigabit speeds, the command that turns off autonegotiation must still include the duplex setting.

Gigabit Ethernet ports support flow control only when autonegotiation is turned on. When autonegotiation is turned off, flow control is not supported. (See the ExtremeXOS Concepts Guide for more detailed information on flow control on Extreme Networks devices.)

Summit Family Switches Only

When configuring combination ports you can specify the medium as copper or fiber. If the medium is not specified for combination ports then the configuration is applied to the current primary medium.



The current primary medium is displayed in the Media Primary column of the `show ports configuration` command output.

Note

The keyword `medium` is used to select the configuration medium for combination ports. If the `port_list` contains any non-combination ports, the command is rejected.



When upgrading a switch running ExtremeXOS 12.3 or earlier software to ExtremeXOS 12.4 or later, saved configurations from combo ports (copper or fiber) are applied only to combo ports fiber medium. When downgrading from ExtremeXOS 12.4 or later to ExtremeXOS 12.3 or earlier, saved configurations from combo ports (copper or fiber) are silently ignored. Therefore, you need to reconfigure combo ports during such an upgrade or downgrade.

Example

The following example turns autonegotiation off for slot 2, port 1 at full duplex on a modular switch:

```
configure ports 2:1 auto off speed 100 duplex full
```

The following example turns autonegotiation off for port 2 with copper medium and a port speed of 100 Mbps at full duplex:

```
configure ports 2 medium copper auto off speed 100 duplex full
```

History

This command was first available in ExtremeXOS 10.1.

The `medium` parameter was added in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

configure ports auto on

```
configure ports port_list {medium [copper|fiber]} auto on [{speed speed} {duplex [half | full]}] | [{duplex [half | full]} {speed speed}]
```

Description

Enables autonegotiation for the particular port type.



Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
medium	Specifies the medium as either copper or fiber. Note: This parameter applies to combo ports only on the Summit family switches.
<i>speed</i>	Specifies the port speed as either 10, 100, 1000 (1 Gigabit), or 10000 (10 Gigabit) Mbps ports.
duplex [half]	Specifies half duplex; transmitting and receiving data one direction at a time.
duplex [full]	Specifies full duplex; transmitting and receiving data at the same time.

Default

Auto on for 1 Gbps ports.

Auto off for 10 Gbps ports.

Usage Guidelines

The type of ports enabled for autonegotiation are 802.3u for 10/100 Mbps ports or 802.3z for gigabit Ethernet ports.

Flow control on gigabit Ethernet ports is enabled or disabled as part of autonegotiation. If autonegotiation is set to off, flow control is disabled. When autonegotiation is turned on, flow control is enabled. See the *ExtremeXOS Concepts Guide* for more detailed information on flow control on Extreme Networks devices.

Summit Family Switches Only

When configuring combo ports you can specify the medium as copper or fiber. If the medium is not specified for combination ports then the configuration is applied to the current primary medium. The current primary medium is displayed in the Media Primary column of the `show ports configuration` command output.

Note

The keyword `medium` is used to select the configuration medium for combination ports. If the `port_list` contains any non-combination ports, the command is rejected.



When upgrading a switch running ExtremeXOS 12.3 or earlier software to ExtremeXOS 12.4 or later, saved configurations from combo ports (copper or fiber) are applied only to combo ports fiber medium. When downgrading from ExtremeXOS 12.4 or later to ExtremeXOS 12.3 or earlier, saved configurations from combo ports (copper or fiber) are silently ignored. Therefore, you need to reconfigure combo ports during such an upgrade or downgrade.

Example

The following command configures the switch to autonegotiate for slot 1, ports 2 and 4 on a modular switch:

```
configure ports 1:2, 1:4 auto on
```

The following command configures the switch to autonegotiate for port 2, with copper medium at a port speed of 100 Mbps at full duplex:

```
configure ports 2 medium copper auto on speed 100 duplex full
```

History

This command was first available in ExtremeXOS 10.1.

The speed and duplex parameters were added in ExtremeXOS 11.6.

The medium parameter was added in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

configure ports auto-polarity

```
configure ports port_list auto-polarity [off | on]
```

Description

Configures the autopolarity detection feature on the specified Ethernet ports.

Syntax Description

<i>port_list</i>	Specifies one or more ports on the switch.
off	Disables the autopolarity detection feature on the specified ports.
on	Enables the autopolarity detection feature on the specified ports.

Default

Enabled.



Usage Guidelines

This feature applies to only the 10/100/1000 BASE-T ports on the switch and copper medium on Summit combination ports.

When autopolarity is disabled on one or more Ethernet ports, you can verify that status by using the following command:

```
show ports information detail
```

Example

The following command disables the autopolarity detection feature on ports 5 to 7 on a Summit series switch:

```
configure ports 5-7 auto-polarity off
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure ports display-string

```
configure ports port_list display-string string
```

Description

Configures a user-defined string for a port or group of ports.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>string</i>	Specifies a user-defined display string.

Default

The null string is the default.



Usage Guidelines

The display string can be up to 15 characters. Display strings do not need to be unique for each port—you can assign the same string to multiple ports. For example, you could give all the ports that connected to a particular department a common display string.

The string is displayed in certain commands such as the `show ports information` command.



Note

Do not use a port number as a display string. For example, do not assign the display string “2” to port2.

Example

The following command configures the user-defined string corporate for port 1 on a stand-alone switch:

```
configure ports 1 display-string corporate
```

The following command configures the user-defined string corporate for ports 3, 4, and 5 on slot 1 on a modular switch:

```
configure ports 1:3-5 display-string corporate
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ports dwdm channel

```
configure port all | port_list dwdm channel channel_number
```

Description

Selects the DWDM channel frequency for the selected ports.

Syntax Description

all	Specifies all ports.
<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>channel_number</i>	Specifies the channel number, which corresponds to one of 102 available channel frequencies.



Default

Channel number – 21.

Usage Guidelines

The following table lists the available frequencies and the channel number you must specify to select each frequency.

Table 16: TX Wavelengths and Channel Assignments for the Tunable DWDM XFP

TX Wavelength	Channel	TX Wavelength	Channel	TX Wavelength	Channel	TX Wavelength	Channel
1568.77 nm	11	1558.17 nm	24	1547.72 nm	37	1537.40 nm	50
1568.36 nm	1150	1557.77 nm	2450	1547.32 nm	3750	1537.00 nm	5050
1567.95 nm	12	1557.36 nm	25	1546.92 nm	38	1536.61 nm	51
1567.54 nm	1250	1556.96 nm	2550	1546.52 nm	3850	1536.22 nm	5150
1567.13 nm	13	1556.55 nm	26	1546.12 nm	39	1535.82 nm	52
1566.72 nm	1350	1556.15 nm	2650	1545.72 nm	3950	1535.43 nm	5250
1566.31 nm	14	1555.75 nm	27	1545.32 nm	40	1535.04 nm	53
1565.90 nm	1450	1555.34 nm	2750	1544.92 nm	4050	1534.64 nm	5350
1565.50 nm	15	1554.94 nm	28	1544.53 nm	41	1534.25 nm	54
1565.09 nm	1550	1554.54 nm	2850	1544.13 nm	4150	1533.86 nm	5450
1564.68 nm	16	1554.13 nm	29	1543.73 nm	42	1533.47 nm	55
1564.27 nm	1650	1553.73 nm	2950	1543.33 nm	4250	1533.07 nm	5550
1563.86 nm	17	1553.33 nm	30	1542.94 nm	43	1532.68 nm	56
1563.45 nm	1750	1552.93 nm	3050	1542.54 nm	4350	1532.29 nm	5650
1563.05 nm	18	1552.52 nm	31	1542.14 nm	44	1531.90 nm	57
1562.64 nm	1850	1552.12 nm	3150	1541.75 nm	4450	1531.51 nm	5750
1562.23 nm	19	1551.72 nm	32	1541.35 nm	45	1531.12 nm	58
1561.83 nm	1950	1551.32 nm	3250	1540.95 nm	4550	1530.72 nm	5850
1561.42 nm	20	1550.92 nm	33	1540.56 nm	46	1530.33 nm	59
1561.01 nm	2050	1550.52 nm	3350	1540.16 nm	4650	1529.94 nm	5950
1560.61 nm	21	1550.12 nm	34	1539.77 nm	47	1529.55 nm	60
1560.20 nm	2150	1549.72 nm	3450	1539.37 nm	4750	1529.16 nm	6050
1559.79 nm	22	1549.32 nm	35	1538.98 nm	48	1528.77 nm	61
1559.39 nm	2250	1548.91 nm	3550	1538.58 nm	4850	1528.38 nm	6150
1558.98 nm	23	1548.51 nm	36	1538.19 nm	49		
1558.58 nm	2350	1548.11 nm	3650	1537.79 nm	4950		



The supported channel numbers are not contiguous. If you specify a channel number that is not listed in the preceding table, the following error message appears:

```
Error: DWDM Channel configuration failed. Channel number 100 is out of
configurable range. The channel range for the Optical module in port
<port number> is 11 .. 6150.
```

If the optical module in one of the ports in the specified list does not support DWDM, the following error message is displayed:

```
Error: No TDWDM Optics on port <port number>.
```

If the optical module in one of the ports in the specified port list is not an Extreme supported optical module, the following error message is displayed:

```
Error: DWDM Channel configuration failed. Optical module is not Extreme
Networks certified. For DWDM channel configuration, Extreme Network
Certified DWDM module is required.
```

To display the configuration, use the `show ports configuration` or the `show ports information detail` command.

Example

The following command configures DWDM channel 21 on a modular port 1:1:

```
configure port 1:1 dwdm channel 21
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on: BlackDiamond 8800 switches with 10G8Xc, 10G4Xc, or 8900 - 10G8X-xl modules and S-10G1Xc option cards, and Summit X480 switches with VIM2-10G4X modules.

configure ports dwdm channel none

```
configure port all | port_list dwdm channel none
```

Description

Configures the default DWDM channel number.



Syntax Description

all	Specifies all ports.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Channel number - 21.

Usage Guidelines

Use this command to configure the default DWDM channel number to the DWDM optical module inserted in the given port. This default channel number of 21 and will be mapped to the appropriate corresponding channel number of the vendor specific channel. If a non-tunable DWDM optic is present, then the DWDM configuration is silently removed from the software.

Example

The following command configures the default DWDM channel 21 on a supported modular port 1:1:

```
configure port 1:1 dwdm channel none
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on BlackDiamond 8800 switches with 10G8Xc, 10G4Xc, or 8900 - 10G8X-xl modules and S-10G1Xc option cards, and Summit X480 switches with VIM2-10G4X modules.

configure ports eee enable

```
configure ports port_list enable eee [on | off]
```

Description

Enables or disables EEE on the physical layer.



Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
on	Specifies that the port advertises to its link partner that it is EEE capable at certain speeds
off	Specifies that the port speed as either 10, 100, 1000 (1 Gigabit), or 10000 (10 Gigabit) Mbps ports.

Default

Off.

Usage Guidelines

Use this command to enable EEE on the switch. **Enable on** specifies that the port advertises to its link partner that it is EEE capable at certain speeds. If both sides, during auto-negotiation, determine that they both have EEE on and are compatible speed wise, they will determine other parameters (how long it takes to come out of sleep time, how long it takes to wake up) and the link comes up. During periods of non-activity, the link will shut down parts of the port to save energy. This is called LPI for low power idle. When one side sees it must send something, it wakes up the remote and then transmits.

Example

The following example turns the EEE feature on:

```
config port portlist eee enable on
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

EEE is supported on the following Extreme Networks platforms:

- BlackDiamond 8K
- BD X - 10G48T.



Note

EEE is only supported at 10G on this card.

- Summit Series 670V-48T.



Note

EEE is only supported at 10G on this switch.

- Summit Series 440 - all copper ports will support EEE.
- Summit series E4G400 - EEE is implemented through autogrEEEn.
- E4G200 - EEE is implemented through autogrEEEn.



configure ports far-end-fault-indication

```
configure ports port_list far-end-fault-indication [on | off]
```

Description

Enables/disables far-end-fault-indication (FEFI).

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
------------------	---

Default

FEFI is disabled by default.

Usage Guidelines

Use this command to enable FEFI. When enabled, FEFI signals to the remote end on detecting single link (RX) failure/recovery.

The command is supported on non-combination SFP ports, that is, ports 1-12 on the Summit X480-24x, and all ports on the Summit X480-48x. If an attempt is made to use this command on combination ports, it is rejected.

Since the SFPs can be plugged and removed, a warning message is displayed in the syslog when a media type changes from FX/LX and FX to another media type. The warning message indicates that FEFI mode is enabled but is not useful for the new media type.

The current status is displayed in the `show ports information` detail command.

Example

The following command enables FEFI on port 1:

```
configure ports 1 far-end-fault-indication on
```

History

This command was first available in ExtremeXOS 12.1.3.

Platform Availability

This command is available on the Summit X480-24x and X480-48x platforms.



configure ports isolation

```
configure ports port_list isolation[on|off]
```

Description

Enables isolation mode on a per-port basis.

Syntax Description

<i>port_list</i>	Specifies one or more ports, or slots and ports.
isolation	Specifies that Isolated ports are not allowed to inter-communicate.
on	Turns on isolation. Isolated ports are not allowed to inter-communicate.
off	Turns off isolation. This is the default setting.

Default

Isolation is off by default.

Usage Guidelines

Use this command to enable isolation mode on a per-port basis. You can issue the command on a single port or on a master port of a load share group. If you issue the command on a non-master port of a load share group the command will fail. When a port load share group is formed, all of the member ports assume the same isolation setting as the master port.

Example

The following command enables isolation mode on slot 1, ports 2 and 4 on a modular switch:

```
configure ports 1:2, 1:4 isolation on
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.



configure ports link-scan interval

```
configure ports link-scan interval [ milliseconds | default ] { slot [ slot | all ] }
```



Description

Configures the link-scan interval. The configure command allows the user to set the interval in a range between the default for the platform and 500 ms. A higher interval can free up CPU cycles when fast link detection is not a requirement.

Syntax Description

ports	Ports.
link-scan	Configure link scan attributes for polling port status.
interval	Configure amount of time between polling port status
milliseconds	Interval in milliseconds. Range is 50 to 500 for most platforms. The minimum interval depends on the default for the platform."; type="int"; range="[50,500]
default	Default interval (50 ms for most platforms).
slot	Slot number (default all slots)"; capability="slot_available"
all	All slots.

Default

50 ms.

Usage Guidelines

Use this command to configure the link-scan interval.

Example

```
# sh ports link-scan
Slot  Interval (ms)
-----
  1      50 (default)
  2      300
  3      50 (default)
  4      50 (default)
  5
  6
  7
  8      200
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.



configure ports mode

```
configure ports port_list mode {lan | wan-phy}
```

Description

Configures the 10G port in WAN PHY or LAN mode.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
lan	Specifies LAN mode.
wan-phy	Specifies WAN PHY mode.

Default

LAN mode.

Usage Guidelines

Use this command to configure either LAN or WAN PHY mode on a 10G port.

Example

The following command configures the WAN PHY option on a 10G port:

```
configure ports 1:3 mode wan-phy
```

History

This command was available in ExtremeXOS 12.3.

Platform Availability

This command is available on all 10G XFP ports only on Summit X480 series switches.

configure ports partition

```
configure ports [port_list | all] partition [4x10G | 1x40G]
```

Description

Partitions a 40G port into a 4x10G or 1x40G mode.



Syntax Description

<i>port_list</i>	Specifies one or more ports.
all	Specifies all ports.
4x10G	Specifies partitioning ports into four 10G ports.
1x40G	Specifies partitioning ports into one 40G port.

Default

1x40G.

Usage Guidelines

Use this command to partition a 40G port into either four 10G ports or one 40G port.

After you make a configuration change, you must do one of the following to apply the change:

- For BlackDiamond X8 series switches and BlackDiamond 8900-40G6X-xm modules, you can disable and then enable the affected slot, which applies the change without affecting other modules
- For BlackDiamond X8 series switches, BlackDiamond 8900-40G6X-xm modules and Summit X670 switches you can reboot the switch



Note

A configuration change is not applied until the affected slot is disabled and enabled or the switch is rebooted. Port will be removed from all vlans including "default" vlan.

If you attempt to configure a switch that does not support this command, the system returns the following error message:

```
ERROR: Port 1:1 does not support Port Partition mode
```

Example

The following example partitions port 6:1 into 4 10G ports:

```
configure ports 6:1 partition 4x10G
Warning: Configuration will be lost on ports 6:1
No configuration changes should be made on these ports until after
the next reboot.
This command will only take effect after save configuration and either reboot
switch or disable and enable slots 6.
Are you sure you want to continue? (y/N)
```

History

This command was available in ExtremeXOS 12.6.



Platform Availability

This command is available on BlackDiamond X8 switches, BlackDiamond 8900-40G6X-xm modules, and Summit X480 switches with a VIM3-40G3X option card, and X670V switches with a VIM4-40G4X option card.

configure ports preferred-medium

```
configure ports port_list preferred-medium [copper | fiber] {force}
```

Description

Configures the primary uplink port to use a preferred medium.

Syntax Description

<i>port_list</i>	Specifies the port number. On a stand-alone switch, this value is just the port number, and on a modular switch, this value is the slot and port number.
copper	Specifies that the port should always use the 10/100/1000 BASE-T connection whenever a link is established even when a fiber link is also present.
fiber	Specifies that the port should always use the 1 gigabit Ethernet fiber connection whenever a link is established even when a copper link is also present.
force	Disables automatic failover. If the specified preferred medium is not present, the link does not come up even if the secondary medium is present.

Default

The default is fiber.

Usage Guidelines

You specify either copper or fiber for the specified port. The switch evaluates the copper energy and the fiber signal at the time these ports come online. If both are present, the configured preferred medium is chosen; however, if only one is present, the switch brings up that medium and uses this medium the next time the switch is rebooted. When a failure occurs and the uplinks are swapped, the switch continues to keep that uplink assignment until another failure occurs or until the assignment is changed using the CLI.

If you use the force option, it disables automatic failover. If you force the preferred-medium to fiber and the fiber link goes away, the copper link is not used, even if available.

To display the preferred medium, use the show port information detail command (you must use the detail variable to display the preferred medium).



Example

The following establishes copper port 4 as the primary uplink on the Summit series switch and fiber port 4 as the redundant uplink port:

```
configure ports 4 preferred-medium copper
```

Copper port 4 becomes the primary uplink until a failure occurs on that link. At that time, fiber port 4 becomes the primary uplink and copper port 4 becomes the redundant port. This assignment stays in place until the next failure.

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available only on Summit family switches and SummitStack.


configure ports redundant

```
configure ports primaryPort redundant secondaryPort {link [on | off]}
```

Description

Configures a software-controlled redundant port.

Syntax Description

<i>primaryPort</i>	Specifies one primary port or slot and port.
redundantPort <i>secondaryPort</i>	Specifies one or redundant port or slot and port.
link	Specifies state of link: on—Specifies keeping the redundant port active, but block traffic off—Specifies forcing the link down on the redundant port
	Note The default value is off.

Default

N/A.

Usage Guidelines

The first port specifies the primary port. The second port specifies the redundant port.



A software-controlled redundant port is configured to back up a specified primary port; both ports are on the same device. The redundant port tracks the link state of the associated primary port, and if the link on the primary port fails, the redundant port establishes a link and becomes active. You can back up a specified Ethernet port with a redundant, dedicated Ethernet port.

You configure the redundant link to be always physically up but logically blocked or to be always physically down. The default is off, or the redundant link is down.

The following criteria must be considered when configuring a software-controlled redundant port:

- You can configure only one redundant port for each primary port.
- You cannot have any Layer 2 protocols configured on any of the VLANs that are present on the ports. (You will see an error message if you attempt to configure software redundant ports on ports with VLANs running Layer 2 protocols.)
- The primary and redundant port must have identical VLAN memberships.
- The master port is the only port of a load-sharing group that can be configured as either a primary or redundant port. (The entire trunk must go down before the software-controlled redundant port takes effect.)
- Only one side of the link should be configured as redundant.

Example

The following command configures a software-controlled redundant port:

```
configure ports 1:3 redundant 2:3
```

History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure ports tdm cable-length

```
configure ports port_list tdm cable-length [ short-haul [110 | 220 | 330 | 440 | 550 | 660] | long-haul line-build-out [0db | 75db | 150db | 225db]]
```

Description

Configures the cable length and receiver gain used on the specified time division multiplexing (TDM) ports. This option is applicable only for T1 ports.



Syntax Description

<i>port_list</i>	Specifies a port or group of ports.
short-haul	Short-haul cable length (less than 660 feet).
110	110 feet (Default).
220	220 feet.
330	330 feet.
440	440 feet.
550	550 feet.
660	660 feet.
long-haul	Long-haul cable length (greater than 660 feet).
line-build-out	Receiver gain in decibels.
0db	0 decibels.
75db	-7.5 decibels.
150db	-15 decibels.
225db	-22.5 decibels.

Default

The default selection for T1 is 110 feet.

Usage Guidelines

The cable length is given in decibels (long-haul) or feet (short-haul). A short-haul cable length is less than 660 feet, while a long-haul is greater than 660 feet. For a long-haul cable, the line built-out is the receiver gain in decibels.

This option is applicable only for T1 ports:

Error: TDM port cable length configuration is applicable onhly to T1 TDM hierarchy.

Example

```
E4G-400.34 # configure port 35 tdm cable-length long-haul line-build-out 150db
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).



configure ports tdm clock-source

```
configure ports port_list tdm clock-source [line | network | [adaptive | differential]] ces ces_name]
```

Description

Configures the transmit clock source of the specified TDM port.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
line	TDM line recovered clock (clock slave mode) (Default).
adaptive	Adaptive clock recovery
differential	Differential clock recovery.
network	Network clock.
<i>ces_name</i>	Alpha-numeric string identifying Circuit Emulation Service pseudowire recovered clock.

Default

The default clock-source is line.

Usage Guidelines

Currently the hardware limits the CES pseudowire recovered clock (Adaptive/Differential) to be supplied only to the port terminating the pseudo-wire. When the [adaptive | differential] switch is used, the backend EXOS process processing the command will return the following error message if the *port_list* parameter expands to multiple ports:

```
Error: The CES pseudo-wire <ces_name> cannot be used to supply clock to multiple ports.
```

The backend EXOS process processing the command will return the following error message if the port specified in the *port_list* parameter does not terminate the CES pseudo-wire when the [adaptive | differential] switch is used:

```
Error: CES pseudo-wire <ces_name> does not terminate on port specified for using Adaptive or Differential clock recovery.
```

Example

```
E4G-400#configure port 40 tdm clock-source network
```



History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers E4G-200 and E4G-400.

configure ports tdm display-string

```
configure ports port_list tdm display-string string
```

Description

Configures a user-defined string on the specified TDM ports.

Syntax Description

<i>port_list</i>	Specifies a port or group of ports.
<i>string</i>	Character string to be associated with the ports.

Default

The null string is the default.

Usage Guidelines

This command can be used to associate a user-defined string to one or more TDM ports.

The display string can be up to 15 characters. Display strings do not need to be unique for each port—you can assign the same string to multiple ports. For example, you could give all the ports that connected to a particular department a common display string.

The string is displayed in certain commands such as the `show ports information` command.



Note

Do not use a port number as a display string. For example, do not assign the display string “2” to port2.

Example

```
E4G-200#configure pot 15 tdm display-string tdm-port-15
```



History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

configure ports tdm framing

```
configure ports port_list tdm framing [d4 | esf | [basic | mf] {crc4} | unframed]
```

Description

Configures the framing used on the specified TDM ports.

Syntax Description

<i>port_list</i>	Specifies a port or group of ports.
d4	T1 Super-Frame.
esf	T1 Extended Super-Frame.
basic	E1 Basic Frame.
mf	E1 Multi-Frame.
crc4	E1 Framing with CRC-4.
unframed	Unframed (Default).

Default

The default selection for framing (E1/T1) is unframed.

Usage Guidelines

If framing options configured are not compatible with the chosen TDM hierarchy, the following error message will be printed, and the configuration will be rejected.

```
Error: TDM port framing option(s) specified on port 31 is incompatible with the configured TDM hierarchy (T1).
```

or

```
Error: TDM port framing option(s) specified on port 31 is incompatible with the configured TDM hierarchy (E1).
```



Example

```
E4G-200#configure port 15 tdm framing basic crc4
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

configure ports tdm idle-code

```
configure ports port_list tdm idle-code idle_code
```

Description

Configures the idle-code to be transmitted in the 4-bits telephony line signaling coding of the TDM channel (DS0 timeslots), when the channel is not connected.

Syntax Description

<i>port_list</i>	Specifies a port or group of ports.
tdm	Time Division Multiplexing.
idle-code	Idle code signaled when the timeslots of a port are not connected.
<i>idle-code</i>	Idle code value between 0 and 15 to be transmitted in the 4-bits telephony line signaling coding of the TDM channel. The default value is 15.

Default

The default idle-code is 15.

Usage Guidelines

If the line coding options configured are not compatible with the chosen TDM hierarchy, the following error message is printed, and the configuration is rejected:

```
Error: Idle Code cannot be configured on unframed TDM ports or on framed TDM ports with signaling disabled.
```



Example

```
E4G400#configure port 35 tdm idle-code 15
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

configure ports tdm line-coding

```
configure ports port_list tdm line-coding [b8zs | hdb3 | ami]
```

Description

Configures the line coding scheme used on the specified time division multiplexing (TDM) ports.

Syntax Description

<i>port_list</i>	Specifies a port or group of ports.
tdm	Time Division Multiplexing.
line-coding	Zero code suppression variety of TDM ports.
b8zs	T1 B8ZS line coding (default for T1).
hdb3	E1 HDB3 line coding (default for E1).
ami	E1 AMI line coding.

Default

The default selection for T1 is B8ZS, and the default selection for E1 is HDB3.

Usage Guidelines

If the line coding options configured are not compatible with the chosen TDM hierarchy, the following error message will be printed and the configuration will be rejected:

```
Error: TDM port line coding configuration is incompatible with the
configured TDM hierarchy (T1).
```

or

```
Error: TDM port line coding configuration is incompatible with the
configured TDM hierarchy (E1).
```



Example

```
E4G400#configure port 35 tdm line-coding HDB3
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

configure ports tdm recovered-clock

```
configure ports port_list tdm recovered-clock {quality-level value}
```

Description

Configures the clock recovery on the specified TDM ports.

Syntax Description

tdm	Time Division Multiplexing.
recovered-clock	Clock recovered from the port.
quality-level	Quality level of the clock (T1 default is QL_ST3, E1 default is QL_SEC).
<i>value</i>	Value of quality level.

Default

The T1 default is QL_ST3, E1 default is QL_SEC.

Usage Guidelines

Use this command to configure the clock recovery on the specified TDM ports.

Example

```
E4G200#configure port 20 rdm recovered-clock quality-level QL_PRC
```

History

This command was first available in ExtremeXOS 15.1.



Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

configure ports tdm signaling

```
configure ports port_list tdm signaling [bit-oriented | robbed-bit | none]
```

Description

Configures the signaling on the specified TDM ports.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
tdm	Time Division Multiplexing.
signaling	Signaling mode of the TDM ports.
bit-oriented	E1 channel associated signaling mode.
robbed-bit	T1 robbed bit signaling.
none	No signaling (default).

Default

The default is no signaling (none).

Usage Guidelines

Configures the signaling on the specified TDM ports. If an incompatible setting is attempted, an error message is displayed:

```
Error: TDM port signaling option is incompatible with the configured TDM hierarchy (T1).
```

or

```
Error: TDM port signaling option is incompatible with the configured TDM hierarchy (E1).
```

Example

Example output not yet available and will be provided in a future release.

History

This command was first available in ExtremeXOS 15.1.



Platform Availability

This command is available on the cell site routers E4G-200 and E4G-400.

configure ports tdm trunk-conditioning

```
configure tdm service circuit service_name trunk-conditioning trunk_conditioning
```

Description

Configures the trunk-conditioning value to be transmitted in the TDM channel (DS0 timeslots) during alarm conditions.

Syntax Description

<i>service_name</i>	TDM service name.
trunk-conditioning	Trunk conditioning signal for the timeslots of a port.
<i>trunk_conditioning</i>	Trunk conditioning value between 0 and 255 to be transmitted in the TDM channel during alarm conditions (default is 255, the range is 0-255).

Default

The default trunk-conditioning value is 255.

Usage Guidelines

Note the following error type:

```
Error: Trunk conditioning cannot be configured on unframed TDM ports.
```

Example

```
E4G-200.45 # configure tdm ser circuit test trunk-conditioning 255
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).



configure ports wan-phy clocking

```
configure ports port_list wan-phy clocking [line | internal]
```

Description

Configures the clocking source for the specified WAN PHY port.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
line	Specifies the port clocking is recovered from the received bitstream.
internal	Specifies the port clock is an internal, free-run clock.

Default

Line.

Usage Guidelines

Each WAN PHY port can be configured for the source it will use for clocking; line clocking uses received clocking and internal clocking uses an internal clock.

XFP ports must be in WAN PHY mode. To configure the mode, use the `configure ports port_list mode {lan | wan-phy}` command.

Example

The following command configures an LW XENPAK WAN PHY port to use the free-run internal clock:

```
configure ports 1:3 wan-phy clocking internal
```

History

This command was available in ExtremeXOS 11.6.

Platform Availability

This command is available on 10G XFP ports on Summit X480 series switches.

configure ports wan-phy framing

```
configure ports port_list wan-phy framing [sonet | sdh]
```



Description

Configures the framing type for the specified WAN PHY port.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
sonet	Specifies the port framing type as SONET.
sdh	Specifies the port framing type as SDH.

Default

SONET.

Usage Guidelines

Each WAN PHY port can be configured for framing that complies with either the SONET standard or the SDH standard. (SONET is primarily a U.S. standard, and SDH is the international version.)

XFP ports must be in WAN PHY mode. To configure the mode, use the `configure ports port_list mode {lan | wan-phy}` command.

Example

The following command configures an LW XENPAK WAN PHY port to use SONET framing:

```
configure ports 1:3 wan-phy framing sonet
```

History

This command was available in ExtremeXOS 11.6.

Platform Availability

This command is available on all 10G ports on 10G XFP ports only on Summit X480 series switches.

configure ports wan-phy loopback

```
configure ports port_list wan-phy loopback [line | off]
```

For Summit X480 series switches only:


```
configure ports port_list wan-phy loopback {off | internal | line}
```

Description

Configures the loopback options for the specified WAN PHY port.



Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
line	Specifies the signal received from the optical media returns back to the transmitter.
off	Specifies no loopback.
internal	Specifies that the signal received from the system is sent back to the system instead of the optical module.
 Note This parameter is available only on Summit X480 series switches.	

Default

Off.

Usage Guidelines

Configuring loopback on a WAN PHY port is used for diagnostics and network troubleshooting.

XFP ports must be in WAN PHY mode. To configure the mode, use the `configure ports port_list mode {lan | wan-phy}` command.

Example

The following command configures the loopback option on an WAN PHY capable port as off:

```
configure ports 1:3 wan-phy loopback off
```

History

This command was available in ExtremeXOS 11.6.

Platform Availability

This command is available on 10G XFP ports only on Summit X480 series switches.

configure ports wan-phy trace-path

```
configure ports port_list wan-phy trace-path id_string
```

Description

Configures the path trace identifier for the specified WAN PHY port. Path trace is a maintenance feature of WAN PHY.



Syntax Description

<i>port_list</i>	Specifies the port number. On a stand-alone switch, this value is just the port number, and on a modular switch, this value is the slot and port number.
<i>id_string</i>	Enter an alphanumeric 16-character string.

Default

The IEEE default value, which has no string representation.

Usage Guidelines

The path trace message is used for troubleshooting. One byte of the path overhead associated with each WAN PHY interface SONET/SDH frame carries information identifying the originating path terminating equipment (PTE).

XFP ports must be in WAN PHY mode. To configure the mode, use the `configure ports port_list mode {lan | wan-phy}` command.

Example

The following command configures a patch trace ID:

```
configure ports 1:3 wan-phy trace-path bear3
```

History

This command was available in ExtremeXOS 11.6.

Platform Availability

This command is available on 10G XFP ports only on Summit X480 series switches.

configure ports wan-phy trace-section

```
configure ports port_list wan-phy trace-section id_string
```

Description

Configures the section trace identifier for the specified WAN PHY port. Section trace is a maintenance feature of WAN PHY.



Syntax Description

<i>port_list</i>	Specifies the port number. On a stand-alone switch, this value is just the port number, and on a modular switch, this value is the slot and port number.
<i>id_string</i>	Enter an alphanumeric 16-character string.

Default

The IEEE default value, which has no string representation.

Usage Guidelines

The section trace message is used for troubleshooting. The J0 transmit octets allow a receiver to verify its continued connection to the transmitter; the J0 octet transports a 16-octet continuously repeating section trace message.

The first transmitted section trace octet is J0 transmit 15, which contains the delineation octet; the default value is 137 (hexadecimal 89). The last transmitted section trace octet is J0 transmit 0; the default value for J0 transmit 0 through 14 is 0.

XFP ports must be in WAN PHY mode. To configure the mode, use the `configure ports port_list mode {lan | wan-phy}` command.

Example

The following command configures a section trace ID for an LW XENPAK WAN PHY port:

```
configure ports 1:3 wan-phy trace-section beta4
```

History

This command was available in ExtremeXOS 11.6.

Platform Availability

This command is available on 10G XFP ports only on Summit X480 series switches.

configure sharing add ports

```
configure sharing port add ports port_list
```

Description

Adds ports to a load-sharing, or link aggregation, group. By using link aggregation, you use multiple ports as a single logical port. Link aggregation also provides redundancy because traffic is redistributed to the remaining ports in the link aggregation group (LAG) if one port in the group goes down.



Syntax Description

<i>port</i>	Specifies the logical port for a load-sharing group or link aggregation group (LAG). This number also functions as the LAG Group ID.
<i>port_list</i>	Specifies one or more ports or slots and ports to be grouped in the LAG.

Default

N/A.

Usage Guidelines

Use this command to dynamically add ports to a load-sharing group, or link aggregation group (LAG).



Note

You must create a LAG (or load-sharing group) before you can configure the LAG. To create a LAG, see .

VMAN ports can belong to LAGs. If any port in the LAG is enabled for VMAN, all ports in the group are automatically enabled to handle jumbo size frames. Also, VMAN is automatically enabled on all ports of the untagged LAG.

To verify your configuration, use the `show ports sharing` command.



Note

All ports that are designated for the LAG must be removed from all VLANs prior to configuring the LAG.

Summit family switches

The following guidelines apply to link aggregation on the Summit family switches:

- On the Summit X670 in EXOS 15.2 and later, 32 ports per LAG are supported with LACP support and custom load sharing algorithm only. Additionally, a static LAG can contain up to 16 ports and an LACP LAG can include up to 32 ports.
- One static LAG can contain up to 8 ports.
- An LACP LAG can include a maximum of 16 ports; out of these up to 8 can be selected links and the remaining 8 will be standby links.
- A Health Check LAG can contain up to 8 ports.

BlackDiamond 8800 series switch and SummitStack only

The following guidelines apply to link aggregation on the BlackDiamond 8800 series switch:

- A static LAG can include a maximum of 8 ports.
- An LACP LAG can include a maximum of 16 ports; out of these up to 8 can be selected links and the remaining 8 will be standby links.
- A Health Check LAG can include a maximum of 8 ports.



BlackDiamond X8 series switch only

The following guidelines apply to link aggregation on the BlackDiamond X8 series switch:

- "With distributed ARP mode on, the maximum number of aggregator ports on BlackDiamond X8 is 16.
- With LACP support, but with custom load sharing algorithm only, 64 ports per LAG.

Example

The following example adds port 3:13 to the LAG with the logical port 3:9 on a modular switch:

```
configure sharing 3:9 add port 3:13
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure sharing address-based custom

```
configure sharing address-based custom [ipv4 [L3-and-L4 | source-only |
destination-only | source-and-destination] | hash-algorithm [xor | crc-16]]
```

Description

This command configures the part of the packet examined by the switch when selecting the egress port for transmitting link aggregation, or load-sharing, data.

Syntax Description

ipv4	IPv4 hash configuration for custom load sharing and L2VPN sharing.
L3-and-L4	Indicates that the switch should examine the IP source and destination address and the TCP or UDP source and destination port number.
source-only	Indicates that the switch should examine the IP source address only.
destination-only	Indicates that the switch should examine the IP destination address only.
source-and-destination	Indicates that the switch should examine the IP source and destination address.
hash-algorithm	Hash algorithm for custom load sharing and L2VPN sharing.
xor	Use exclusive-OR for load sharing hash computation.
crc-16	Use CRC-16 for load sharing hash computation.



Default

Algorithm: L3-and-L4.

Hash algorithm: xor.

Usage Guidelines

This command specifies the part of the packet header that the switch examines to select the egress port for address-based load-sharing trunks. The address-based load-sharing setting is global and applies to all load-sharing trunks, or LAGs, that are address-based and configured with a custom algorithm. You change this setting by issuing the command again with a different option.

The addressing information examined is based on the packet protocol as follows:

- IPv4 packets—Uses the source and destination IPv4 addresses and Layer4 port numbers as specified with this command.
- IPv6 packets—Uses the source and destination IPv6 addresses and Layer4 port numbers.
- MPLS packets—Uses the top, second, and reserved labels and the source and destination IP addresses.
- Non-IP Layer 2—Uses the VLAN ID, the source and destination MAC addresses, and the ethertype.

The xor hash algorithm guarantees that the same egress port is selected for traffic distribution based on a pair of IP addresses, Layer4 ports, or both, regardless of which is the source and which is the destination.

For IP-in-IP and GRE tunneled packets, the switch examines the inner header to determine the egress port.

To verify your configuration, use the `show ports sharing` command.

Example

The following example configures the switch to examine the source IP address:

```
configure sharing address-based custom ipv4 source-only
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.

configure sharing delete ports

```
configure sharing port delete ports port_list
```



Description

Deletes ports from a link aggregation, or load-sharing, group.

Syntax Description

<i>port</i>	Specifies the logical port for a load-sharing group or a link aggregation group (LAG). This number also functions as the LAG Group ID.
<i>port_list</i>	Specifies one or more ports or slots and ports to be grouped in the LAG.

Default

N/A.

Usage Guidelines

Use this command to dynamically delete ports from a load-sharing group, or link aggregation group (LAG). This command applies to static and dynamic link aggregation.

Example

The following example deletes port 3:12 from the LAG with the logical port, or LAG Group ID, 3:9:

```
configure sharing 3:9 delete port 3:12
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure sharing health-check member-port add tcp-tracking

```
configure sharing health-check member-port port add tcp-tracking IP Address {tcp-port TC Port frequency sec misses count}
```

Description

Configures monitoring for each member port of a health check LAG.



Syntax Description

<i>port</i>	Specifies the member port.
<i>IP Address</i>	Specifies the IP address to monitor.
<i>TCP Port</i>	Specifies the TCP port to watch. The default is port 80.
<i>sec</i>	Specifies the frequency in seconds at which tracking takes place. The default is 10 seconds.
<i>count</i>	Specifies the number of misses before a connection loss is reported. The default is 3 misses.

Default

N/A.

Usage Guidelines

To configure a health check LAG, you first create a health check type of LAG using the `enable sharing grouping` command. Then use this command to configure the monitoring for each member port. You can configure each member port to track a particular IP address, but only one IP address per member port.

To display the monitoring configuration for a health check LAG, use the `show sharing health-check` command.

To display the link aggregation configured on a switch, use the `show ports sharing` command.

Example

The following commands configure four different member ports:

```
# configure sharing health-check member-port 10 add track-tcp 10.1.1.1 tcp-
port 23
# configure sharing health-check member-port 11 add track-tcp 10.1.1.2 tcp-
port 23
# configure sharing health-check member-port 12 add track-tcp 10.1.1.3
# configure sharing health-check member-port 13 add track-tcp 10.1.1.4
```

When the TCP port, seconds, or counts are not specified, they default to the values described in the Syntax Description.

History

This command was first available in ExtremeXOS 12.1.3.

Platform Availability

This command is available on all platforms.



configure sharing health-check member-port delete tcp-tracking

```
configure sharing health-check member-port port delete tcp-tracking IP Address
{tcp-port TC Port}
```

Description

Unconfigures monitoring for each member port of a health check LAG.

Syntax Description

<i>port</i>	Specifies the member port.
<i>IP Address</i>	Specifies the IP address.
<i>TCP Port</i>	Specifies the TCP port.

Default

N/A.

Usage Guidelines

Use this command to remove the monitoring configuration on the ports of a health check link aggregation group. Each port must be unconfigured separately, specifying the IP address and TCP port.

Example

The following command removes the configuration setting on port 12 that monitors IP address 10.1.1.3:

```
# configure sharing health-check member-port 12 delete track-tcp 10.1.1.3
```

History

This command was first available in ExtremeXOS 12.1.3.

Platform Availability

This command is available on all platforms.

configure sharing health-check member-port tcp-tracking

```
configure sharing health-check member-port port [disable | enable] tcp-tracking
```



Description

Enables or disables configured monitoring on a member port of a health check LAG.

Syntax Description

<i>port</i>	Specifies the member port.
-------------	----------------------------

Default

N/A.

Usage Guidelines

This disables/enables monitoring on a particular member port. When monitoring is disabled, the member port is added back to the LAG if it has not already been added. This allows a member port to be added back to LAG even though connectivity to the host is down.

Example

The following command disables port 12:

```
configure sharing health-check member-port 12 disable tcp-tracking
```

History

This command was first available in ExtremeXOS 12.1.3.

Platform Availability

This command is available on all platforms.

configure sharing lacp activity-mode

```
configure sharing port lacp activity-mode [active | passive]
```

Description

Configures whether the switch sends LACPDUs periodically (active) or only in response to LACPDUs sent from the partner on the link (passive).



Syntax Description

<i>port</i>	Specifies the master logical port for the LAG you are setting the activity mode for.
active	Enter this value to have the switch periodically sent LACPDUs for this LAG.
passive	Enter this value to have the switch only respond to LACPDUs for this LAG.

Default

Active.

Usage Guidelines

You must enable sharing and create the LAG prior to assigning this LACP activity mode.



Note

One side of the link must be in active mode in order to pass traffic. If you configure your side in the passive mode, ensure that the partner link is in LACP active mode.

To verify the LACP activity mode, use the `show lacp lag group-id detail` command.

If you attempt to enter a port number that is different than a LAG group ID, the system returns the following error message:

```
ERROR: LAG group Id does not exist
```



Note

In ExtremeXOS version 11.3, the activity mode cannot be changed from active.

Example

The following command changes the activity mode to passive for the specified LAG group ID:

```
configure sharing 5:1 lacp activity-mode passive
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure sharing lacp defaulted-state-action

```
configure sharing port lacp defaulted-state-action [add | delete]
```



Description

Configures a defaulted LAG port to be removed from the aggregator.

Syntax Description

<i>port</i>	Specifies the master logical port for the LAG you are setting the default action for.
add	Enter this value to have the switch add defaulted ports to the aggregator for this LAG.
delete	Enter this value to have the switch delete defaulted ports from the aggregator for this LAG.

Default

Delete.

Usage Guidelines

You must enable sharing and create the LAG prior to configuring this LACP parameter.

You can configure whether you want a defaulted LAG port removed from the aggregator or added back into the aggregator. If you configure the LAG to remove ports that move into the default state, those ports are removed from the aggregator and the port state is set to unselected.



Note

In ExtremeXOS version 11.3, defaulted ports in the LAG are always removed from the aggregator; this is not configurable.

If you configure the LAG to add the defaulted port into the aggregator, the system takes inventory of the number of ports currently in the aggregator:

- If there are fewer ports in the aggregator than the maximum number allowed, the system adds the defaulted port to the aggregator (port set to selected and collecting-distributing).
- If the aggregator has the maximum ports, the system adds the defaulted port to the standby list (port set to standby).



Note

If the defaulted port is assigned to standby, that port automatically has a lower priority than any other port in the LAG (including those already in standby).

To verify the LACP default action, use the `show lacp lag group-id detail` command.

If you attempt to enter a port number that is different than a LAG group ID, the system returns the following error message:

```
ERROR: LAG group Id does not exist
```



Note

To force the LACP trunk to behave like a static sharing trunk, use this command to add ports to the aggregator.



Example

The following command deletes defaulted ports from the aggregator for the specified LAG group ID:

```
configure sharing 5:1 lacp defaulted-state-action delete
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure sharing lacp system-priority

```
configure sharing port lacp system-priority priority
```

Description

Configures the system priority used by LACP for each LAG to establish the station on which end assumes priority in determining those LAG ports moved to the collecting/distributing state of the protocol. That end of the LAG with the lowest system priority is the one that assumes control of the determination. This is optional; if you do not configure this parameter, LACP uses system MAC values to determine priority. If you choose to configure this parameter, enter a value between 1 and 65535.

Syntax Description

<i>port</i>	Specifies the master logical port for the LAG you are setting the priority for.
<i>priority</i>	Enter the value you want for the priority of the system for the LACP. The range is 0 to 65535; there is no default.

Default

N/A.

Usage Guidelines

The LACP uses the system MAC values to assign priority to one of the systems, and that system then determines which LAG ports move into the collecting/distributing state and exchange traffic. That end of the LAG with the lowest system priority is the one that assumes control of the determination. If you wish to override the default LACP system priority for a specific LAG, use this command to assign that LAG a specific LACP priority. Enter a value between 0 and 65535.

You must enable sharing and create the LAG prior to assigning this LACP priority.

To verify the LACP system priority, use the `show lacp` command.



To change the system priority you previously assigned to a specific LAG, issue the `configure sharing lacp system-priority` command using the new priority you want. To remove the assigned system priority entirely and use the LACP priorities, issue the `configure sharing lacp system-priority` command using a value of 0.

Example

The following command assigns LAG 10 an LACP system priority of 3:

```
configure sharing 10 lacp system-priority 3
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

configure sharing lacp timeout

```
configure sharing port lacp timeout [long | short]
```

Description

Configures the timeout used by each LAG to stop transmitting once LACPDUs are no longer received from the partner link. You can configure this timeout value to be either 90 seconds, long, or 3 seconds, short.

Syntax Description

<i>port</i>	Specifies the master logical port for the LAG you are setting the timeout value for.
long	Enter this value to use 90 seconds as the timeout value.
short	Enter this value to use 3 seconds as the timeout value.

Default

Long.

Usage Guidelines

You must enable sharing and create the LAG prior to assigning this LACP timeout value.

To verify the LACP timeout value, use the `show lacp lag group-id detail` command.



If you attempt to enter a port number that is different than a LAG group ID, the system returns the following error message:

```
ERROR: LAG group Id does not exist
```



Note

In ExtremeXOS version 11.3, the timeout value is set to long and cannot be changed.

Example

The following command changes the timeout value for the specified LAG group ID to short:

```
configure sharing 5:1 lacp timeout short
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure sharing port-based key

```
configure sharing [ load_sharing_key | default ] ports port_list
```

Description

Sets the *load_sharing_key* for all ports in the *port_list*.

Syntax Description

<i>load_sharing_key</i>	Specifies the load sharing key. Valid load sharing keys are in the range [0-15].
default	Unconfigures and resets the load sharing keys for ports in the <i>port_list</i> to default values.
ports	Specifies the logical port for a load-sharing group.
<i>port_list</i>	Specifies one or more ports or slots and ports to be grouped in the LAG.

Default

N/A.

Usage Guidelines

This command sets the *load_sharing_key* for all ports in the *port_list*. **default** unconfigures and resets the load sharing keys for ports in *port_list* to default values.



Configured load sharing keys are displayed in the output of the `show configuration hal` command.

Example

The following example causes all packets received on ports in slot 1 to choose the lowest port number in all aggregators for distribution.:

```
configure sharing port-based key 0 ports 1
```



Note

If you attempts to configure a load sharing key on a BDx module that is not configured to use packet-based hashing, the following error message is displayed:

```
Error: Slot 1 is not configured for packet - based fabric hashing.
```

Use the `configure forwarding fabric hash packet slot slot` command before configuring load sharing keys on a slot.

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

configure slot module

```
configure slot slot module module_type
```

Description

Configures a slot for a particular I/O module card in a modular switch. On a stack, this command configures a slot for a particular type of node.

Syntax Description

<i>slot</i>	Specifies the slot number.
<i>module_type</i>	Specifies the type of module or node for which the slot should be configured. The list of modules you can enter will vary depending on the type of switch and version of ExtremeXOS you are running. Certain modules are supported only with specific ExtremeXOS Technology Releases. On a stack, module type will be the list of switches that support SummitStack.



Default

If a slot has not been configured for a particular type of I/O module, then any type of module is accepted in that slot, and a default port and VLAN configuration is automatically generated.

Usage Guidelines

The `configure sharing lacp timeout` command displays different module parameters depending on the type of modular switch you are configuring and the version of ExtremeXOS running on the switch.

You can also preconfigure the slot before inserting the module card. This allows you to begin configuring the module and ports before installing the card in the chassis.

If a slot has not been configured for a particular type of I/O module, then any type of module is accepted in that slot, and a default port and VLAN configuration is automatically generated. If a slot is configured for one type of module, and a different type of module is inserted, the inserted module is put into a mismatch state, and is not brought online. To use the new module type in a slot, the slot configuration must be cleared or configured for the new module type.

Upon powering up the chassis, or when an I/O module is hot-swapped, ExtremeXOS automatically determines the system power budget and protects the switch from any potential overpower configurations. If power is available, ExtremeXOS powers on and initializes the module. When ExtremeXOS detects that a module will cause an overpower condition, the module remains powered down, and is not initialized. An entry is made to the system log indicating the condition.

On a stack, the module type must be a switch that supports SummitStack.

Example

The following command configures slot 2 for a 10/100/1000, 60-port, copper module:

```
configure slot 2 module G60T
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available only on modular switches and SummitStack.

configure slot restart-limit

```
configure slot slot_number restart-limit num_restarts
```

Description

Configures the number of times a slot can be restarted on a failure before it is shut down.



Syntax Description

<i>slot_number</i>	Specifies the slot number.
<i>num_restarts</i>	Specifies the number of times the slot can be restarted. The range is from 0 to 10,000.

Default

The default is 5.

Usage Guidelines

This command allows you to configure the number of times a slot can be restarted on a failure before it is shut down. If the number of failures exceeds the restart-limit, the module goes into a “Failed” state. If that occurs, use the `disable slot` and `enable slot` commands to restart the module.

Example

The following command configures slot 2 on the switch to be restarted up to 3 times upon a failure:

```
configure slot 2 restart-limit 3
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on modular switches and SummitStack.

configure tdm hierarchy

```
configure tdm hierarchy [t1 | e1]
```

Description

Configures the time division multiplexing (TDM) hierarchy to be used by the physical layer carrier system of the switch. This can be either the North American hierarchy (e.g., T1, T3, OC-3, etc.) or the European hierarchy (E1, E3, STM-1, etc.). The configuration of the TDM hierarchy is a global configuration affecting all TDM ports present in the switch. When the TDM hierarchy is changed, the command warns the user with the following messages if any CES is already configured:

Message: A save and reboot are required before the changes take effect. Upon reboot, TDM port parameters will be reset to defaults. Error: Cannot change TDM hierarchy with CES pseudo-wire(s) configured. All CES pseudo-wire(s) must be deleted before changing TDM hierarchy. Error: Cannot change TDM hierarchy with TDM service(s) configured. All TDM service(s) must be deleted before changing TDM hierarchy.



Syntax Description

tdm	Time Division Multiplexing.
hierarchy	Physical layer carrier system.
t1	T1, North American digital carrier hierarchy.
e1	E1, European digital carrier hierarchy (default).

Default

The default hierarchy is E1.

Usage Guidelines

Recommended sequence of CLI commands for changing TDM hierarchy (e.g., changing E1 (default) to T1):

- Before reboot:

```
configure tdm hierarchy t1 save configuration reboot
```

Example

```
E4G-200#delete ces all
E4G-200#delete tdm service circuit all
E4G-200#configure tdm hierarchy t1
A Save and Reboot are required to make the changes effect at the next reboot
of this switch. Upon reboot, TDM port parameters will be reset to defaults
E4G-200#save configuration
E4G-200#reboot
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

configure tdm service circuit add port

```
configure tdm service circuit service_name add port port {time-slots
[time_slot_list | all]}
```

Description

Adds a <port, time-slots> or *port* to the specified TDM service.



Syntax Description

tdm	Time Division Multiplexing.
service	Provider provisioned TDM service.
circuit	TDM circuit service.
<i>service_name</i>	TDM service name.
add	Add specified TDM circuit as a service.
<i>port</i>	TDM port number.
<i>time_slot_list</i>	TDM time slot list.
all	All TDM time slots.

Default

N/A.

Usage Guidelines

A structure-aware TDM circuit can be created by adding time slots from a framed T1/E1 port to the service. A structure-agnostic TDM bit-stream is created by adding an unframed T1/E1 port to the service.

Adding an unframed TDM port to a bundle would not require the time-slots token. The command would be `configure tdm service circuit service_name add port port`. The time-slots token is required for framed ports only. The following error message will be displayed for T1. For E1 the error message is not required since the CLI performs the range check (range is 1-32).

```
Error: Time slot(s) specified is outside the valid range. The valid time slots range for T1 is [1-24].
```

If the user specifies time-slot token for an unframed port, the following error is displayed:

```
Error: Time slot(s) cannot be specified for an unframed T1 port.
```

or

```
Error: Time slot(s) cannot be specified for an unframed E1 port.
```

Example

```
E4G-400#configure port 35 tdm framing basic crc
E4G-400#createt dm service circuit service-test
E4G-400#configure tdm service circuit service-test add port 35 time-slots all
```

History

This command was first available in ExtremeXOS 15.1.



Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

configure tdm service circuit delete port

```
configure tdm service circuit service_name delete port port
```

Description

Deletes a *port*, *time-slots* or *port* to the specified TDM service.

Syntax Description

<i>service_name</i>	TDM service name.
delete	Delete specified TDM circuit as a service.
<i>port</i>	TDM port number.

Default

N/A.

Usage Guidelines

A structure-aware TDM circuit can be created by adding time slots from a framed T1/E1 port to the service. A structure-agnostic TDM bit-stream is created by adding an unframed T1/E1 port to the service.

Adding an unframed TDM port to a bundle would not require the time-slots token. The command would be `configure tdm service circuit service_name add port port`. The time-slots token is required for framed ports only. The following error message will be displayed for T1. For E1 the error message is not required since the CLI performs the range check (range is 1-32).

```
Error: Time slot(s) specified is outside the valid range. The valid time slots range for T1 is [1-24].
```

If the user specifies time-slot token for an unframed port, the following error is displayed:

```
Error: Time slot(s) cannot be specified for an unframed T1 port.
```

or

```
Error: Time slot(s) cannot be specified for an unframed E1 port.
```

Example

```
E4G200#configure tdm service circuit service-test delete port 35
```



History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

configure tdm service circuit seized-code

```
configure tdm service circuit service_name seized-code seized_code
```

Description

Configures the seized-code to be transmitted in the 4-bits telephony line signaling coding of the TDM channel (DS0 timeslots), when the channel is connected.

Syntax Description

<i>seized_code</i>	The seized code value between 0 and 15 to be transmitted in the 4-bits telephony line signaling coding of the TDM channel (default is 15).
--------------------	--

Default

The default seized-code is 15.

Usage Guidelines

Note the following error type:

```
Error: Seized Code cannot be configured on unframed TDM ports or on framed TDM ports with signaling disabled.
```

Example

```
config ports 2 tdm seized-code 12
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).



create ces psn

```
create ces ces_name psn [mef8 | udp | mpls]
```

Description

Creates a circuit emulation service (CES) pseudowire with the specified name and packet switched network (PSN) tunnel type.

Syntax Description

ces	Circuit emulation service.
<i>ces_name</i>	Alphanumeric string identifying the CES pseudowire.
mef8	MEF-8 PSN tunnel for the CES pseudowire.
udp	IPv4/UDP PSN tunnel for the CES pseudowire.
mpls	MPLS tunnel for the CES pseudowire.

Default

N/A.

Usage Guidelines

Use this command to create a CES pseudowire, specifying the name and PSN tunnel type.

Example

```
#create ces ces-test psn mpls
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

create mirror to port

```
create mirror mirror_name {to [port port | port_list port_list loopback-port  
port ] { remote-tag rtag }} {description mirror-desc}
```



Description

Creates a named mirror instance with an optional description, and optional "to port" definition.

Syntax Description

<i>mirror_name</i>	Specifies the mirror name.
port	Specifies the mirror output port.
<i>port_list</i>	Specifies the list of ports where traffic is to be mirrored.
loopback-port	Specifies an otherwise unused port required when mirroring to a <i>port_list</i> . The loopback-port is not available for switching user data traffic.
<i>port</i>	Specifies a single loopback port that is used internally to provide this feature.
remote-tag	Specifies the value of the VLAN ID used by the mirrored packets when egressing the monitor port.
description	Specifies a description of the named mirror instance.
<i>mirror-desc</i>	The specified mirror description.

Default

Disabled.

Usage Guidelines

Use this command to create a named mirror instance with an optional description and optional "to port" definitions. You can create 15 named instances (the instance "DefaultMirror" is created automatically).

Example

The following example creates a mirror instance on port 3, slot 4 :

```
create mirror to port 3:4
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

create mlag peer

```
create mlag peer peer_name
```



Description

Creates an MLAG peer switch association structure.

Syntax Description

<i>peer_name</i>	Specifies an alpha numeric string identifying the MLAG peer.
------------------	--

Default

N/A.

Usage Guidelines

This command creates an MLAG peer switch association structure.

You must use a unique name for the peer switch. If you attempt to create an MLAG peer with a name that already exists, the following error message is displayed:

```
ERROR: MLAG peer with specified name already exists
```

Example

The following command creates a peer switch structure switch101:

```
create mlag peer switch101
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

create tdm service circuit

```
create tdm service circuit service_name
```

Description

Creates a TDM service with the specified name.



Syntax Description

service	Provider provisioned TDM service.
circuit	TDM circuit service.
<i>service_name</i>	TDM service name.

Default

N/A.

Usage Guidelines

Currently, only TDM circuit services are supported. A structure-aware TDM circuit can be created by adding time slots of a framed T1/E1 port to the service. A structure-agnostic TDM bit-stream is created by adding an unframed T1/E1 port to the service.

Example

```
#create tdm service circuit service-test
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

delete ces

```
delete ces [ces_name | all]
```

Description

Deletes the specified circuit emulation service (CES) pseudowire.

Syntax Description

ces	Circuit emulation service.
<i>ces_name</i>	Alphanumeric string identifying the CES pseudowire.
all	All CES pseudo-wires in the switch.



Default

N/A.

Usage Guidelines

Use this command to delete a specific CES pseudowire, or to delete all CES pseudo-wires in the switch.

Example

This example deletes all CES pseudo-wires in the switch:

```
delete ces all
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

delete mirror name

```
delete mirror mirror_name | all]
```

Description

Deletes a user-defined mirroring instance, and unconfigures the "DefaultMirror" instance.

Syntax Description

<i>mirror_name</i>	Specifies a specific mirror name to delete.
all	Specifies that you delete all named mirror instances.

Default

Disabled.

Usage Guidelines

Use this command to delete a user-defined mirroring instance and unconfigure the "DefaultMirror" instance. Mirroring instances must be in the "disabled" state in order to be deleted. The **all** command will fail if any mirroring instance is in the "enabled" state.



Example

The following example deletes all mirroring instances:

```
delete mirror all
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

delete mlag peer

```
delete mlag peer peer_name
```

Description

Deletes a peer switch from the MLAG structure.

Syntax Description

<i>peer_name</i>	Specifies an alpha numeric string identifying the MLAG peer.
------------------	--

Default

N/A.

Usage Guidelines

This command deletes an MLAG peer switch from the association structure.

Before you delete an MLAG peer switch, you must disable it. If it is not disabled, the following error message is displayed:

```
ERROR: MLAG ports currently associated with peer. First disable MLAG ports using "disable mlag port <port>" before deleting MLAG peer
```

Example

The following command deletes a peer switch structure switch101:

```
delete mlag peer switch101
```



History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

delete tdm service circuit

```
delete tdm service circuit [service_name | all]
```

Description

Deletes the specified TDM service.

Syntax Description

service	Provider provisioned TDM service.
circuit	TDM circuit service.
<i>service_name</i>	TDM service name.
all	All TDM services.

Default

N/A.

Usage Guidelines

Use this command to delete a specific TDM service or to delete all TDM services.

Example

To delete all TDM services, enter the following command:

```
delete tdm service circuit all
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).



disable ces

```
disable ces [ces_name | all]
```

Description

Disables the administrative status of the specified circuit emulation service (CES) pseudowire.

Syntax Description

ces	Circuit emulation service.
<i>ces_name</i>	Alphanumeric string identifying the CES pseudowire.
all	All CES pseudo-wires in the switch.

Default

The default administrative state is enable.

Usage Guidelines

Use this command to disable a specific CES pseudowire, or to disable all CES pseudo-wires in the switch.

Example

This example disables all CES pseudo-wires in the switch:

```
disable ces all
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

disable edp ports

```
disable edp ports [ports | all]
```

Description

Disables the Extreme Discovery Protocol (EDP) on one or more ports.



Syntax Description

<i>ports</i>	Specifies one or more ports or slots and ports.
all	Specifies all ports on the switch.

Default

Enabled.

Usage Guidelines

You can use the `disable edp ports` command to disable EDP on one or more ports when you no longer need to locate neighbor Extreme Networks switches.

Example

The following command disables EDP on slot 1, ports 2 and 4 on a modular switch:

```
disable edp ports 1:2, 1:4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable flow-control ports

```
disable flow-control [tx-pause {priority priority} | rx-pause {qosprofile qosprofile}] ports [all | port_list]
```

Description

Disables specified flow control configurations.

Syntax Description

tx-pause	Specifies transmission pause processing.
<i>priority</i>	Specifies all priorities or single priorities--dot1p priority for tagged packets and internal priority for untagged packets. Used with priority flow control only.
rx-pause	Specifies reception pause processing.



<i>qosprofile</i>	Specifies a QoS profile (“qp1” “qp2” “qp3” “qp4” “qp5” “qp6” “qp7” “qp8”) to pause for priority flow control packet reception. Used with priority flow control only.
all	Specifies all ports or slots.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Disabled.

Usage Guidelines

IEEE 802.3x-Flow Control

Use this command to disable the processing of IEEE 802.3x pause flow control messages received from the remote partner. Disabling rx-pause processing avoids dropping packets in the switch and allows for better overall network performance in some scenarios where protocols such as TCP handle the retransmission of dropped packets by the remote partner.

To disable RX flow-control, TX flow-control must first be disabled. Refer to the `disable flow-control ports` command. If you attempt to disable RX flow-control with TX flow-control enabled, an error message is displayed.

IEEE 802.1Qbb-Priority Flow Control

Use this command to disable the processing of IEEE 802.1Qbb priority flow control messages received from the remote partner. Disabling TX stops the port from transmitting PFC packets for that priority, regardless of congestion. Disabling RX stops the processing of PFC packets received on that port for the specific QoS profile.

IEEE 802.3x

The following command disables the tx flow-control feature on ports 5 through 7 on a Summit switch:

```
disable flow-control tx-pause ports 5-7
```

IEEE 802.1Qbb

The following command disables TX for priority 3 on port 3 on a Summit X460 switch:

```
disable flow-control tx-pause priority 3 ports 3
```

The following command disables RX for QoS profile qp4 on port 6 of a Summit X460 switch:

```
disable flow-control rx-pause qosprofile qp4 port 6
```



History

This command was first available in ExtremeXOS 12.1.3.

The priority function (PFC) was added in ExtremeXOS 12.5.

Platform Availability

IEEE 802.3x

The basic TX-pause and RX-pause functions of this command are available on the BlackDiamond X8 switches, BlackDiamond 8000 series modules and the Summit family switches.

IEEE 802.1Qbb

The priority function (PFC) is available only on 10G ports and in some cases on specific models of the following newer platforms indicated by the part number:

- BlackDiamond 8900-10G24X-c modules (manufacturing number 800397-00)
- BlackDiamond 8900-40G6X-xm modules
- 40G ports and 10G ports when in 4x10 partition mode
- BlackDiamond X8 switches
- Summit X460 switches
- 10G ports
- Summit X670 switches
- Summit X770 switches
- 10G ports Summit X670V switches
- 10G and 40G ports

disable ip-fix ports

```
disable ip-fix ports [port_list | all]
```

Description

Disables IPFIX metering on the port.

Syntax Description

<i>port_list</i>	Specifies the ports.
all	Specifies all ports.

Default

The default is disabled.



Usage Guidelines

Use this command to turn off IPFIX metering on a port.

Example

The following command disables the IPFIX metering support on the port:

```
disable ip-fix ports 2:1
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on BlackDiamond 8900 G96Tc, G48T-xl, G48X-xl, and 10G8X-xl modules and Summit X460 and X480 switches.

disable jumbo-frame ports

```
disable jumbo-frame ports [all | port_list]
```

Description

Disables jumbo frame support on a port.

Syntax Description

all	Specifies all ports.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Disabled.

Usage Guidelines

Use this command to disable jumbo frames on individual ports.



Example

The following command disables jumbo frame support on a BlackDiamond 8810 switch:

```
disable jumbo-frame ports all
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable learning port

```
disable learning {drop-packets | forward-packets} port [port_list | all]
```

Description

Disables MAC address learning on one or more ports for security purposes.

Syntax Description

drop-packets	Specifies that packets with unknown source MAC addresses be dropped. When disable learning is configured, this is the default behavior.
forward-packets	Specifies that packets with unknown source MAC addresses be forwarded.
port	Specifies the port.
<i>port_list</i>	Specifies one or more ports or slots and ports.
all	Specifies all ports and slots.

Default

Enabled.

Usage Guidelines

Use this command in a secure environment where access is granted via permanent forwarding databases (FDBs) per port.



Example

The following command disables MAC address learning on port 4:3 on a modular switch:

```
disable learning ports 4:3
```

History

This command was first available in ExtremeXOS 10.1.

Support for drop-packets and forward-packets options on the BlackDiamond 8800 switches was included in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

disable mirror

```
disable mirror mirror_name | all
```

Description

Disables a mirror instance.

Syntax Description

<i>mirror_name</i>	Specifies the mirror name.
all	Specifies all mirror instance are deleted.

Default

Disabled.

Usage Guidelines

Use this command to disable mirrors. Disabling an instance only changes the state, its configuration remains as defined (a change from current operation, which loses some configuration parameters).

Example

The following example disable a mirror instance named "mirror1" :

```
disable mirror mirror1
```



History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

disable mlag port

```
disable mlag port port
```

Description

Removes a local port or LAG from an MLAG.

Syntax Description

<i>port</i>	Specifies a local member port of the MLAG group.
-------------	--

Default

N/A.

Usage Guidelines

Use this command to remove a local port or LAG from an MLAG.

Example

The following command binds the local member port 2 to the peer switch switch101 with an identifier of 101:

```
disable mlag port 2
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.



disable network-clock ptp end-to-end-transparent ports

```
disable network-clock ptp end-to-end-transparent ports port_list
```

Description

Disable PTP end-to-end-transparent clock functionality (1-step PHY timestamp) on the specified ports.

Syntax Description

<i>port_list</i>	Specifies a port or group of ports.
------------------	-------------------------------------

Default

N/A.

Usage Guidelines

Use this command to disable 1-step transparent clock on the specified ports.

Example

The following example disables end-to-end-transparent clock on front panel port 2:

```
disable network-clock ptp end-to-end-transparent ports 2
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers E4G-200 and E4G-400.

disable network-clock sync-e

```
disable network-clock sync-e port [port_list | all]
```

Description

Disables synchronous Ethernet (SyncE) on port(s).



Syntax Description

<i>port_list</i>	Specifies a port or group of ports.
all	Specifies all ports.

Default

Disabled.

Usage Guidelines

Use this command to disable SyncE on one or more ports.

Example

The following command disables SyncE on port 2:

```
disable network-clock sync-e port 2
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on Summit X460-24X and X460-48X switches and on E4G-200 and E4G-400 switches.

disable port

```
disable port [port_list | all]
```

Description

Disables one or more ports on the switch.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
all	Specifies all ports on the switch.

Default

Enabled.



Usage Guidelines

Use this command for security, administration, and troubleshooting purposes.

When a port is disabled, the link is brought down.

Example

The following command disables ports 3, 5, and 12 through 15 on a stand-alone switch:

```
disable ports 3,5,12-15
```

The following command disables slot 1, ports 3, 5, and 12 through 15 on a modular switch:

```
disable port 1:3,1:5,1:12-1:15
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable ports tdm

```
disable ports [port_list | all] tdm
```

Description

Disables the specified TDM ports.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
all	Disable all TDM ports on the switch.
tdm	Indicates that command applies to time division multiplexing (TDM) ports.

Default

The default is enable.



Usage Guidelines

This command disables TDM ports only and has no effect on Ethernet ports. If this command is used on unsupported ports, the following error message appears:

Error: There are no valid TDM ports in the specified port list.

Example

```
E4G200#disable port 15 tdm
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers E4G-200 and E4G-400.

configure slot module

configure slot *slot* **module** *module_type*

Description

Configures a slot for a particular I/O module card in a modular switch. On a stack, this command configures a slot for a particular type of node.

Syntax Description

<i>slot</i>	Specifies the slot number.
<i>module_type</i>	Specifies the type of module or node for which the slot should be configured. The list of modules you can enter will vary depending on the type of switch and version of ExtremeXOS you are running. Certain modules are supported only with specific ExtremeXOS Technology Releases. On a stack, module type will be the list of switches that support SummitStack.

Default

If a slot has not been configured for a particular type of I/O module, then any type of module is accepted in that slot, and a default port and VLAN configuration is automatically generated.

Usage Guidelines

The `configure sharing lacp timeout` command displays different module parameters depending on the type of modular switch you are configuring and the version of ExtremeXOS running on the switch.



You can also preconfigure the slot before inserting the module card. This allows you to begin configuring the module and ports before installing the card in the chassis.

If a slot has not been configured for a particular type of I/O module, then any type of module is accepted in that slot, and a default port and VLAN configuration is automatically generated. If a slot is configured for one type of module, and a different type of module is inserted, the inserted module is put into a mismatch state, and is not brought online. To use the new module type in a slot, the slot configuration must be cleared or configured for the new module type.

Upon powering up the chassis, or when an I/O module is hot-swapped, ExtremeXOS automatically determines the system power budget and protects the switch from any potential overpower configurations. If power is available, ExtremeXOS powers on and initializes the module. When ExtremeXOS detects that a module will cause an overpower condition, the module remains powered down, and is not initialized. An entry is made to the system log indicating the condition.

On a stack, the module type must be a switch that supports SummitStack.

Example

The following command configures slot 2 for a 10/100/1000, 60-port, copper module:

```
configure slot 2 module G60T
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available only on modular switches and SummitStack.

disable slot


```
disable slot slot | FM-1 | FM-2 | FM-3 | FM-4> {offline}
```

Description

Disables slot and leaves that module in a power down state. These can be either regular slots or fabric slots (FM-1 through FM-4).



Syntax Description

<code>slot</code>	Specifies the slot to be disabled.
<code>offline</code>	Specifies that the slot be disabled offline.
<div style="display: flex; align-items: flex-start;">  <div> <p>Note</p> <p>This variable is supported only on the BlackDiamond 8800 series switches; that is, those modular switches that support offline diagnostics.</p> </div> </div>	

Default

Enabled.

Usage Guidelines

This command allows the user to disable a slot. When the user types this command, the I/O card in that particular slot number is brought down, and the slot is powered down. The LEDs on the card go OFF.

When a fabric slot (e.g., FM-1) is disabled, it is powered off and the bandwidth it provides is unavailable. Disabling an active fabric slot reroutes the switch fabric traffic before powering off the inserted FM blade. Thus, if there are four active fabric modules when one is disabled, there should be no traffic loss.

A disabled slot can be re-enabled using the `enable slot` command. When the slot is re-enabled, the software on the I/O module is updated to match the software on the primary MSM/MM.

The `show slot` command, if invoked after the user disables the slot, shows this slot state as “Power Off/Disabled.”

If there is no I/O card present in a slot when the user disables the slot, the slot still goes to the “Disable” state. If a card is inserted in a slot that has been disabled, the card does not come up and stays in the “Power Off/Disabled” state until the slot is enabled by using the `enable slot` command. below.

If you do not save the configuration before you do a switch reboot, the slot will be re-enabled upon reboot. If you save the configuration after disabling a slot, the slot will remain disabled after a reboot.

On Power over Ethernet (PoE) modules, disabling a slot also disables any inline power that is flowing to that slot.

BlackDiamond 8800 series switch only

This command applies only to the data, or I/O ports on slots holding an MSM. The slots holding an MSM on the BlackDiamond 8810 switch are 5 and possibly 6; the slots holding an MSM on the BlackDiamond 8806 switch are 3 and possibly 4. Use the offline parameter to run the diagnostics offline.



Example

The following command disables slot 5 on the switch:

```
disable slot 5
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available only on modular switches.

disable smartredundancy

```
disable smartredundancy port_list
```

Description

Disables the Smart Redundancy feature.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
------------------	---

Default

Enabled.

Usage Guidelines

The Smart Redundancy feature works in concert with the software-controlled redundant feature. When Smart Redundancy is disabled, the switch attempts only to reset the primary port to active if the redundant port fails. That is, if you disable Smart Redundancy, the traffic does not automatically return to the primary port once it becomes active again; the traffic continues to flow through the redundant port even after the primary port comes up again.

Example

The following command disables the Smart Redundancy feature on ports 1:1 to 1:4 on a modular switch:

```
disable smartredundancy 1:1-4
```



History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

disable snmp traps port-up-down ports

```
disable snmp traps port-up-down ports [port_list | all]
```

Description

Disables port up/down trap reception for specified ports.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
all	Specifies all ports on the switch.

Default

Enabled.

Usage Guidelines

Use this command to stop receiving SNMP trap messages when a port transitions between being up and down.

Example

The following command stops ports 3, 5, and 12 through 15 on a stand-alone switch from receiving SNMP trap messages when the port goes up/down:

```
disable snmp traps port-up-down ports 3,5,12-15
```

History

This command was first available in ExtremeXOS 10.6.

Platform Availability

This command is available on all platforms.



enable ces

```
enable ces [ces_name | all]
```

Description

Enables the administrative status of the specified circuit emulation service (CES) pseudowire.

Syntax Description

<i>ces_name</i>	Alphanumeric string identifying the CES pseudowire.
all	All CES pseudowires in the switch.

Default

The default administrative state is enable.

Usage Guidelines

Use this command to enable a specific CES pseudowire, or to enable all CES pseudowires in the switch.

Example

This example enables all CES pseudowires in the switch:

```
enable ces all
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

enable | disable ces peer ipaddress

```
[enable | disable] ces ces_name peer ipaddress ipaddress
```

Description

Enables/disables an IPv4 peer (far-end) for the specified CES pseudowire. This command only supports to MPLS PSN transport CES pseudowires.



This command enables or disables a specific PW configured for the specified *ces_name*. By default, the PW is enabled when it is configured. When a PW is disabled, the switch signals preferred forwarding status of Standby to the PW peer and the switch ceases transmitting or receiving data packets over the PW. The PW label is not withdrawn. Packets may still be sent and received over the PW's G-ACh. When the PW is re-enabled, the switch signals preferred forwarding status Active to the PW peer.

Syntax Description

ces	Circuit Emulation Service.
<i>ces_name</i>	Alphanumeric string identifying the CES pseudowire.
peer	Enable/disable the peer of the CES pseudowire.
<i>ipaddress</i>	IPv4 Address; type=ipv4_t.

Default

The default administrative state is enable.

Usage Guidelines

Use this command to enable/disable an IPv4 peer (far-end) for the specified CES pseudowire. This command only supports to MPLS PSN transport CES pseudo-wires.

Example

Example output not yet available and will be provided in a future release.

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

enable edp ports

```
enable edp ports [ports | all]
```

Description

Enables the Extreme Discovery Protocol (EDP) on one or more ports.



Syntax Description

<i>ports</i>	Specifies one or more ports or slots and ports.
all	Specifies all ports on the switch.

Default

Enabled.

Usage Guidelines

On a modular switch or SummitStack, *ports* can be a list of slots and ports. On a stand-alone switch, *ports* can be one or more port numbers. For a detailed explanation of port specification, see [Port Numbering in Command Reference Overview](#)

EDP is useful when Extreme Networks switches are attached to a port.

The EDP is used to locate neighbor Extreme Networks switches and exchange information about switch configuration. When running on a normal switch port, EDP is used to by the switches to exchange topology information with each other. Information communicated using EDP includes the following:

- Switch MAC address (switch ID).
- Switch software version information.
- Switch IP address.
- Switch VLAN information.
- Switch port number.
- Switch port configuration data: duplex, and speed.

Example

The following command enables EDP on slot 1, port 3 on a modular switch:

```
enable edp ports 1:3
```

History

This command was first available in ExtremeXOS 10.1.

The port configuration data was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.



enable flow-control ports

```
enable flow-control [tx-pause {priority priority} | rx-pause {qosprofile qosprofile}] ports [all | port_list]
```

Description

Enables flow control or priority flow control (PFC) on the specified ports.

Syntax Description

tx-pause	Specifies transmit pause frames.
<i>priority</i>	Specifies all priorities or single priorities--dot1p priority for tagged packets and internal priority for untagged packets. Used with priority flow control only.
rx-pause	Specifies received pause frames.
<i>qosprofile</i>	Specifies a QoS profile ("qp1" "qp2" "qp3" "qp4" "qp5" "qp6" "qp7" "qp8") to pause for priority flow control packet reception. Used with priority flow control only.
all	Specifies all ports or slots.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Disabled.

Usage Guidelines

With autonegotiation enabled, the Summit family switches and the BlackDiamond 8800 series switches advertise the ability to support pause frames. This includes receiving, reacting to (stopping transmission), and transmitting pause frames. However, the switch does not actually transmit pause frames unless it is configured to do so.

IEEE 802.3x-Flow Control

IEEE 802.3x flow control provides the ability to configure different modes in the default behaviors.

Use this command to configure the switch to transmit link-layer pause frames when congestion is detected. This stops all traffic on the configured port when there is buffer congestion for any traffic type. Use it also to configure the switch to return to the default behavior of processing received pause frames.

To enable TX flow-control, RX flow-control must first be enabled. If you attempt to enable TX flow-control with RX flow-control disabled, an error message is displayed.



IEEE 802.1Qbb-Priority Flow Control

IEEE 802.1Qbb priority flow control provides the ability to configure the switch to transmit link-layer pause frames to stop only a portion of the traffic when congestion is detected.

When IEEE 802.1Qbb priority flow control is enabled on a port, IEEE 802.3x pause functionality is no longer available on that port.

Priority is established for reception of PFC packets with a QoS profile value on the ExtremeXOS switch and for transmission with a priority value added to the PFC packet.

- **QoS profile**—Ingress traffic is associated with a QoS profile for assignment to one of eight hardware queues in the system that define how the traffic flows with respect to bandwidth, priority, and other parameters. By default, there are two QoS profiles (QP1 and QP8) defined in these supported platforms and PFC works with this default. To segregate the ingress traffic with more granularity, you will want to define other QoS profiles.
- **Priority**—The traffic that is paused is based on the priority bits in the VLAN header for tagged packets. You can specify this transmit priority independently from the QoS profile to associate it with the reception of a PFC packets thus giving flexibility in the configuration of the network.

It is suggested that the priority in the VLAN header match the QoS profile priority when traffic ingresses at the edge of the network so that the traffic can be more easily controlled as it traverses through the network.



Note

On Summit X670 and X670V switches, the PFC feature does not support fabric flow control messages on alternate stack ports or SummitStack-V80 native stack ports.

IEEE 802.3x

The following command enables the TX flow-control feature on ports 5 through 7 on a Summit switch:

```
enable flow-control tx-pause ports 5-7
```

IEEE 802.1Qbb

The following command enables the priority flow control feature on a Summit switch:

```
enable flow-control tx-pause priority 3 ports 2
```

History

This command was first available in ExtremeXOS 12.1.3.

IEEE 802.1Qbb priority flow control (PFC) was added in ExtremeXOS 12.5.



Platform Availability

IEEE 802.3x

The basic TX-pause and RX-pause functions of this command are available on BlackDiamond X8 switches, BlackDiamond 8000 series modules and Summit family switches.

IEEE 802.1Qbb

The priority function (PFC) is available only on 10G ports and on specific models of the following newer platforms indicated by the part number:

- BlackDiamond X8 switches.
- BlackDiamond 8900-10G24X-c modules (manufacturing number 800397-00).
- BlackDiamond 8900-40G6X-xm modules, 40G ports and 10G ports when in 4x10 partition mode.
- Summit X460 switches 10G ports.
- Summit X670 switches, 10G ports.
- Summit X670V switches, 10G and 40G ports.
- Summit X770 40G and 10G ports.

enable | disable ip-fix

[enable | disable] ip-fix

Description

Enables or disables IPFIX globally.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Use this command to enable or disable IPFIX globally. When used, it overrides the individual port enable command. It is provided to simplify debugging.

Example

The following command enables IPFIX globally on the switch:

```
enable ip-fix
```



History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on BlackDiamond 8900 G96Tc, G48T-xl, G48X-xl, and 10G8X-xl modules and Summit X460 and X480 switches.

enable ip-fix ports

```
enable ip-fix ports [port_list | all] {ipv4 | ipv6 | non-ip | all_traffic}
```

Description

Enables IPFIX on the ports.

Syntax Description

<i>port_list</i>	Specifies the ports.
all	Specifies all ports.
ipv4	Meter IPv4 traffic.
ipv6	Meter IPv6 traffic.
non-ip	Meter non-IP layer 2 traffic.
all_traffic	Meter IPv4, IPv6, and non-IP traffic. This is the default.

Default

The default is disabled. When enabled, the default is **all_traffic**.

Usage Guidelines

In addition to enabling IPFIX support of the port, use this command to check that the port being enabled has hardware support for IPFIX.

If the port does not support IPFIX, an error message similar to the following is displayed:

```
Error: IPFIX is not supported by hardware on port <slot>:<port>
```

The port specified must be either a physical port that is not part of a trunk LAG, or it must be the LAG master port. When it is neither of these, the following message is displayed:

```
Error: Master port 3:1 is enabled for IPFIX but Member port 5:2 is not capable.
```

When the specified port is the LAG master port, all ports in that LAG must support the IPFIX feature in hardware. If one or more ports in the LAG do not support IPFIX, the following error messages is displayed:



Error: IPFIX is not supported on port 5:1.

Example

The following command enables IPFIX metering of IPv6 traffic on port 2:1:

```
enable ip-fix ipv6 ports 2:1
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on BlackDiamond 8900 G96Tc, G48T-xl, G48X-xl, and 10G8X-xl modules and Summit X460 and X480 switches.

enable jumbo-frame ports

```
enable jumbo-frame ports [all | port_list]
```

Description

Enables support on the physical ports that will carry jumbo frames.

Syntax Description

all	Specifies ports.
<i>port_list</i>	Specifies one or more slots and ports.

Default

Disabled.

Usage Guidelines

Increases performance to back-end servers or allows for VMAN 802.1Q encapsulations.

You can configure the maximum size of a jumbo frame if you want to use a different size than the default value of 9216. Use the `configure jumbo-frame-size` command to configure the size.

This setting is preserved across reboots.



Example

The following command enables jumbo frame support on a BlackDiamond 8810 switch:

```
enable jumbo-frame ports all
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable learning port

```
enable learning port [all | port_list]
```

Description

Enables MAC address learning on one or more ports.

Syntax Description

all	Specifies all ports.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Enabled.

Usage Guidelines

N/A.

Example

The following command enables MAC address learning on slot 1, ports 7 and 8 on a modular switch:

```
enable learning ports 1:7-8
```

History

This command was first available in ExtremeXOS 10.1.



Platform Availability

This command is available on all platforms.

enable mirror

```
enable mirror mirror_name
```

Description

Enables a mirror instance.

Syntax Description

<i>mirror_name</i>	Specifies the mirror name.
--------------------	----------------------------

Default

Disabled.

Usage Guidelines

Use this command to enable a mirror instance. An instance may be enabled without source filters defined (per current function), but no traffic will be mirrored until source filters are added.

Example

The following example enables a mirror instance named "mirror1":

```
enable mirror mirror1
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

enable mlag port peer id

```
enable mlag port port peer peer_name id identifier
```



Description

Binds a local port or LAG to an MLAG.

Syntax Description

<i>port</i>	Specifies a local member port of the MLAG group.
<i>peer_name</i>	Specifies an alpha numeric string identifying the MLAG peer.
<i>identifier</i>	Specifies a unique MLAG identifier value. The range is 1 to 65000.

Default

N/A.

Usage Guidelines

Use this command to bind a local port or LAG to an MLAG that is uniquely identified by the MLAG ID value. The MLAG ID can be any number from 1 to 65000.

The specified port number may be a single port or the master port of a load sharing group but may not be a load sharing member port. If it is, a message similar to the following is displayed:

```
ERROR: Port 2 is a member of a load share group. Use the load share master port (10) instead.
```

A port can be part of only one MLAG. If you try to add it to another MLAG, a message similar to the following is displayed:

```
ERROR: Port 2 is already part of an MLAG Id 101
```

Once the MLAG group binding is made, any change to load sharing on MLAG ports is disallowed.

The MLAG peer must exist or the command will fail.

Example

The following command binds the local member port 2 to the peer switch switch101 with an identifier of 101:

```
enable mlag port 2 peer switch101 id 101
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.



enable network-clock ptp

```
[enable | disable] network-clock ptp [boundary | ordinary] {{vlan} vlan_name}
```

Description

Use this command to enable PTP on the clock instance or on the specified VLAN (clock port).

Syntax Description

boundary	Boundary clock.
ordinary	Ordinary clock.
vlan	VLAN.
<i>vlan_name</i>	VLAN name.

Default

N/A.

Usage Guidelines

Use this command to enable PTP on the clock instance or on the specified VLAN (clock port).

Example

The following example enables the ordinary clock:

```
disable network-clock ptp ordinary
```

The following example enables the clock port lpbk-transit on the boundary clock:

```
disable network-clock ptp boundary vlan lpbk-transit
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers E4G-200 and E4G-400.



enable network-clock ptp end-to-end-transparent ports

```
enable network-clock ptp end-to-end-transparent ports port_list
```

Description

Enables PTP end-to-end-transparent clock functionality (1-step PHY timestamp) on the specified ports.

Syntax Description

<i>port_list</i>	Specifies a port or a group of ports.
------------------	---------------------------------------

Default

N/A.

Usage Guidelines

Use this command to enable 1-step transparent clock on the specified ports.

Example

The following example enables end-to-end transparent clock on the front panel ports 1-3:

```
enable network-clock ptp end-to-end-transparent ports 1-3
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

enable network-clock sync-e

```
enable network-clock sync-e port [port_list | all]
```

Description

Enables synchronous Ethernet (SyncE) on port(s).



Syntax Description

<code>port_list</code>	Specifies a port or group of ports.
<code>all</code>	Specifies all ports.

Default

Disabled.

Usage Guidelines

Use this command to enable SyncE on one or more ports.

The following lists the individual platforms and the ports supported on each platform:

- X460-24X: Ports 1 - 28.
- X460-48X: Ports 1 - 48.
- E4G-200: All Ethernet ports.
- E4G-400: All Ethernet ports including XGM3S 10G ports if present.

If you attempt to enable SyncE on a port or ports that are not supported, one of the following messages is displayed:

```
ERROR: Cannot enable Synchronous Ethernet on ports
ERROR: Cannot enable Synchronous Ethernet on some/all ports
```

To disable SyncE, use the `disable network-clock sync-e` command.

To display SyncE settings, use the `show network-clock sync-e ports` command.

Example

The following command enables SyncE on port 2:

```
enable network-clock sync-e port 2
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on Summit X460-24X and X460-48X switches and on E4G-200 and E4G-400 switches.



enable port

```
enable port [port_list | all]
```

Description

Enables a port.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
all	Specifies all ports on the switch.

Default

All ports are enabled.

Usage Guidelines

Use this command to enable the port(s) if you disabled the port(s) for security, administration, or troubleshooting purposes.

Example

The following command enables ports 3, 5, and 12 through 15 on the stand-alone switch:

```
enable ports 3,5,12-15
```

The following command enables slot 1, ports 3, 5, and 12 through 15 on the modular switch:

```
enable port 1:3, 1:5, 1:12-1:15
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable ports tdm

```
enable ports [port_list | all] tdm
```



Description

Enables the specified TDM ports.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
all	Enable all TDM ports on the switch.
tdm	Indicates that command applies to time division multiplexing (TDM) ports.

Default

The default is enable.

Usage Guidelines

This command enables TDM ports only and has no effect on Ethernet ports. If this command is used on an invalid port list, a message similar to the following appears:

```
Error: There are no valid TDM ports in the specified port list.
```

Example

Example output not yet available and will be provided in a future release.

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers E4G-200 and E4G-400.

enable ports tdm loopback

```
enable ports port_list tdm loopback [local | network [line | payload]]
```

Description

Enables the local or network loopback mode used on the specified time division multiplexing (TDM) ports.

Syntax Description

<i>port_list</i>	Specifies a port or group of ports.
tdm	Time Division Multiplexing.



loopback	Loopback mode.
local	Local loopback.
network	Network loopback (transmit received data).
line	Loopback framing and data received from the far end.
payload	Loopback data received from the far end.

Default

The loopback is disabled by default.

Usage Guidelines

Use this command to enable local or network loopback mode on the specified TDM ports.

Example

Example output not yet available and will be provided in a future release.

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

enable sharing grouping

```
enable sharing port grouping port_list {algorithm [address-based {L2 | L3 | L3_L4 | custom} | port-based ]} {lacp | health-check}
```

Description

Enables the switch to configure port link aggregation, or load sharing. By using link aggregation, you use multiple ports as a single logical port. Link aggregation also provides redundancy because traffic is redistributed to the remaining ports in the LAG if one port in the group goes down. LACP allows the system to dynamically configure the LAGs.

The **port-based** keyword was added to the command to support the creation of port-based load sharing groups.

Syntax Description

<i>port</i>	Specifies the master logical port for a load-sharing group or link aggregation group (LAG).
<i>port_list</i>	Specifies one or more ports or slots and ports to be grouped to the logical port.



address-based	Specifies link aggregation by address-based algorithm.
L2	Specifies address-based link aggregation by Layer 2. This is the default value. NOTE: This parameter is available only on BlackDiamond 8800 series switches, SummitStack, and Summit family switches.
L3	Specifies address-based link aggregation by Layer 3. NOTE: The L3 algorithm will be deprecated. Selection of L3 behaves the same as L3_L4. The inclusion of Layer4 ports for distribution is not available on a per group basis. The inclusion of Layer4 ports for distribution is controlled globally for all LAGs in a switch via the <code>configure forwarding sharing [L3 L3_L4]</code> command.
L3_L4	Specifies address-based link aggregation by Layer 3 IP plus Layer4 port. NOTE: The inclusion of Layer4 ports for distribution is not available on a per group basis. The inclusion of Layer4 ports for distribution is controlled globally for all LAGs in a switch via the "configure forwarding sharing [L3 L3_L4]" command.
custom	Selects the custom link aggregation algorithm configured with the following command: <code>configure sharing address-based custom [ipv4 [L3-and-L4 source-only destination-only source-and-destination] hash-algorithm [xor crc-16]]</code> . The configuration of the custom option applies to all LAGs on the switch. This option is supported only on BlackDiamond X8, 8900 c- and xl-series modules, Summit X460, X480, X670 and X770 switches, and a SummitStack that hosts at least one of these supported Summit switches.
port-based	Supports the creation of port-based load sharing groups.
lACP	Specifies dynamic link aggregation, or load sharing, using the LACP.
health-check	Specifies a health check type of link aggregation group.

Default

Disabled.

Usage Guidelines

Link aggregation, or load sharing, allows you to increase bandwidth and availability between switches by using a group of ports to carry traffic in parallel between switches. The aggregation algorithm allows the switch to use multiple ports as a single logical port. For example, VLANs see the link aggregation group (LAG) as a single logical port. Groups can span multiple modules.



Note

All ports that are designated for the LAG must be removed from all VLANs prior to configuring the LAG.

You can enable and configure dynamic link aggregation, using LACP or health-check link aggregation. Static link aggregation is the default link aggregation method.



Note

Always verify the LACP configuration by issuing the `show ports sharing` command. Look for the ports listed as being in the aggregator.



If a port in a LAG fails, traffic is redistributed to the remaining ports in the LAG. If the failed port becomes active again, traffic is redistributed to include that port.

Link aggregation must be enabled on both ends of the link, or a network loop will result.

Any attempt to enable sharing on ports that have an MLAG configuration is denied with following error message:

ERROR: Sharing configuration on MLAG ports cannot be modified. Use "disable mlag port <port>" to remove port from MLAG group first.



Note

See the *ExtremeXOS Concepts Guide* for information on the interaction of port-based ACLs and LAGs of ports.

LAGs are defined according to the following rules:

- Although you can reference only the logical port of a LAG to a Spanning Tree Domain (STPD), all the ports of a load-sharing group actually belong to the specified STPD.
- When using link aggregation, you should always reference the logical port of the LAG when configuring or viewing VLANs. VLANs configured to use other ports in the LAG will have those ports deleted from the VLAN when link aggregation becomes enabled.

Link aggregation, or load-sharing, algorithms allow you to select the distribution technique used by the LAG to determine the output port selection. Algorithm selection is not intended for use in predictive traffic engineering.

ExtremeXOS switches use address based algorithms to determine which physical port in the LAG to use for forwarding traffic out of the switch. Refer to `configure sharing address-based custom` for more information on using addressing information.

For Port-based load sharing:



Note

If you attempt to create a port-based load sharing group with more than 16 possible aggregator ports, the following message will be displayed: **Error: The system can have a maximum of 16 ports in a load sharing group with the configured algorithm.** This message indicates enforcement of the limit of 16 aggregator ports in a port-based LAG. Existing error messages are also used to enforce the 16 aggregator port limit for port-based load sharing groups modified by the `configure sharing port add portsport_list` command.

Summit X440, X460, X480, X670 and X770 series switches and SummitStack only

The following guidelines apply to link aggregation on the Summit X440, X460, X480, X670 and X770 series switches:

- For all Summit switches other than Summit X670 or a SummitStack consisting only of Summit X670s, a static LAG can contain up to 8 ports.
- For a standalone Summit X670 only, a static LAG can contain up to 32 ports.
- A SummitStack consisting only of Summit X670s can contain up to 64 ports in a static LAG.



- An LACP LAG can include twice the number of ports as a static LAG; out of these half can be selected links and any remaining ports will be standby links.
- A Health Check LAG may contain the same number of ports as a static LAG.
- The maximum number of LAGs is 128.
- The default load-sharing algorithm is L2 address-based aggregation. Any broadcast, multicast, or unknown unicast packet is distributed across all members in the LAG.
- The available address-based parameters on the Summit X460, X480, and X670 series switches are L2 for Layer 2 and L3_L4 for Layer 3 plus Layer4. If the packet is not IP, the switch applies the Layer 2 algorithm, which is the default setting.
- The custom keyword is supported only on Summit X460, X480, and X670 switches and a SummitStack that contains at least one of these supported Summit switches. If the custom keyword is specified on a SummitStack that includes a mix of Summit X460, X480, or X670 switches with other Summit family switches, a warning message appears and the individual switches use algorithms as follows:
 - Summit X460, X480, and X670 switches forward unicast traffic using the custom algorithm.
 - All other Summit family switches forward unicast traffic using the L3_L4 algorithm.
 - All Summit family switches forward non-unicast traffic (broadcast, multicast, and unknown unicast packets) using a separate internal hash algorithm.

If the enable sharing grouping command is entered with the custom keyword on a SummitStack that does not contain a Summit X460, X480, or X670 switch, the command fails and an error message appears.

BlackDiamond 8800 series switch only

The following guidelines apply to link aggregation on the BlackDiamond 8800 series switch:

- A static LAG can include a maximum of 8 ports.
- An LACP LAG can include a maximum of 16 ports; out of these up to 8 can be selected links and the remaining 8 will be standby links.
- A Health Check LAG can include a maximum of 8 ports.
- The available address-based parameters on the BlackDiamond 8800 series switch are L2 for Layer 2 and L3_L4 for Layer 3 plus Layer4.

The maximum number of LAGs is 128. If the packet is not IP, the switch applies the Layer 2 algorithm, which is the default setting. Beginning with ExtremeXOS software version 11.2, the switch can use IPv6 addresses.

- Broadcast, multicast, or unknown unicast packets are distributed across all members of a LAG. The distribution of these packets depends on the type of the traffic. Broadcast, L2 multicast and unknown unicast traffic distribution is based on the source and destination MAC addresses. IP multicast traffic distribution is based on the source and destination IP addresses. This behavior is not configurable.
- The custom keyword is supported on BlackDiamond 8800 switches and the custom algorithms operate only on the BlackDiamond 8900 c- and xl-series modules. If the custom keyword is specified on a BlackDiamond 8800 switch that includes a mix of BlackDiamond 8000 series modules, the individual modules use algorithms as follows:
 - BlackDiamond 8900 c- and xl-series modules forward unicast traffic using the custom algorithm.
 - All other modules forward unicast traffic using the L3_L4 algorithm.



- All modules forward non-unicast traffic (broadcast, multicast, and unknown unicast packets) using a separate internal hash algorithm.
- If the enable sharing grouping command is entered with the custom keyword on a BlackDiamond 8800 switch that does not contain a BlackDiamond 8900 c- or xl-series module, the command fails and an error message similar to the following appears:

```
BD-8810.1 # enable sharing 2:1 grouping 2:1-2 algorithm address-based
custom Error: No slots currently in the system support the custom
algorithm
```

- If the enable sharing grouping command is entered with the custom keyword on a BlackDiamond 8800 switch that contains BlackDiamond 8900 c- or xl- series modules and modules that do not support the custom keyword, a warning message similar to the following appears:

Warning: The custom algorithm is not supported on slot 1, 2, 3, 5 and 10. See XOS Command Reference for details.

BlackDiamond X8 series switch only

The following guidelines apply to link aggregation on the BlackDiamond X8 series switch:

- A static LAG can include a maximum of 64 ports with distributed ARP mode off. With distributed ARP mode on, a static LAG can include a maximum of 16 ports.
- An LACP LAG can contain twice the number of ports as a static LAG; out of these half can be selected links and any remaining ports will be standby links. The maximum number of LAGS is 384.

Example

The following example defines a static link aggregation group (LAG) on a modular switch that contains ports 9 through 12 on slot 3, ports 7 through 10 on slot 5, and uses port 3:9 as the logical port:

```
enable sharing 3:9 grouping 3:9-3:12, 5:7-5:10
```

In this example, logical port 3:9 represents physical ports 3:9 through 3:12 and 5:7 through 5:10.

The following example defines a dynamic LAG on a stand-alone switch containing ports 10 through 15, with port 10 being the logical port:

```
enable sharing 10 grouping 10-15 lacp
```

The following example selects the custom option on a BlackDiamond 8810 switch:

```
BD-8810.1 # enable sharing 2:1 grouping 2:1-2 algorithm address-based custom
```

The following example defines a health check LAG containing ports 10 through 13 with port 10 as the master logical port and specifies address-based link aggregation by Layer 3 IP plus Layer4 port:

```
enable sharing 10 grouping 10,11,12,13 algorithm address L3_L4 health-check
```



To configure a health-check LAG, refer to the `configure sharing health-check member-port add tcp-tracking` command.

History

This command was first available in ExtremeXOS 10.1.

The address-based algorithm was added in ExtremeXOS 11.0.

The L2 and L3 optional parameters were added in ExtremeXOS 11.1.

IPv6-compatibility was added in ExtremeXOS 11.2.

Dynamic link aggregation, using LACP, was added in ExtremeXOS 11.3.

The L3_L4 optional parameter was added in ExtremeXOS 11.5.

SummitStack functionality was added in ExtremeXOS 12.0.

Health-check link aggregation was added in ExtremeXOS 12.1.3.

The custom keyword was added in ExtremeXOS 12.3.

LACP support with custom load sharing algorithm only:

- 64 ports per LAG for BlackDiamond X8 was added in ExtremeXOS 15.2.
- 32 ports per LAG for X670 was added in ExtremeXOS 15.2

The port-based keyword was added in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

enable slot

```
enable slot slot | FM-1 | FM-2 | FM-3 | FM-4>
```

Description

Enables slots. These can be either regular slots or fabric slots (FM-1 through FM-4).

Syntax Description

<code>slot</code>	Specifies the slot to be enabled.
-------------------	-----------------------------------

Default

Enabled.



Usage Guidelines

This command allows the user to enable a slot that has been previously disabled using the `disable slot` command.



Note

On the BlackDiamond 8800 series switch, this command only applies to the data, or I/O, ports on slots holding an MSM (slot 5 and possibly 6 on the BlackDiamond 8810; slot 3 and possibly 4 on the BlackDiamond 8806 switch).

When the user enters the enable command, the disabled I/O card in the specified slot is brought up, and the slot is made operational, if possible, or goes to the appropriate state as determined by the card state machine. The LEDs on the card are brought ON as usual. When the slot is enabled, the software on the I/O module is updated to match the software on the primary MSM/MM.

After the user enables the slot, the `show slot` command shows the state as “Operational” or will display the appropriate state if the card could not be brought up successfully. Note that there is no card state named “Enable” and the card goes to the appropriate states as determined by the card state machine when the `enable slot` command is invoked.

Only slots that have their state as “disabled” can be enabled using this command. If this command is used on slots that are in states other than “disabled,” the card state machine takes no action on these slots.

To enable inline power to a slot, the slot must be enabled as well as inline power for that slot. Use the `enable inline-power` command to enable inline power.



Note

If your chassis has an inline power module and there is not enough power to supply a slot, that slot will not be enabled; the slot will not function in data-only mode without enough power for inline power.

Example

The following command enables slot 5 on the switch:

```
enable slot 5
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available only on modular switches.



enable smartredundancy

```
enable smartredundancy port_list
```

Description

Enables the Smart Redundancy feature on the primary port.

Syntax Description

<code>port_list</code>	Specifies one or more ports or slots and ports.
------------------------	---

Default

Enabled.

Usage Guidelines

You must configure the software-controlled redundant port using the `configure ports redundant` command prior to enabling Smart Redundancy.

The Smart Redundancy feature works in concert with the software-controlled redundant port feature. With Smart Redundancy enabled on the switch, when the primary port becomes active the switch redirects all traffic to the primary port and blocks the redundant port again. If you disable Smart Redundancy, the primary port is blocked because traffic is now flowing through the redundant port.

Example

The following command enables the Smart Redundancy feature on slot 1, port 4 on a modular switch:

```
enable smartredundancy 1:4
```

History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

enable snmp traps port-up-down ports

```
enable snmp traps port-up-down ports [port_list | all]
```



Description

Enables port up/down trap reception for specified ports.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
all	Specifies all ports on the switch.

Default

Enabled.

Usage Guidelines

Use this command to begin receiving SNMP trap messages when a port transitions between being up and down.

Example

The following command enables ports 3, 5, and 12 through 15 on a stand-alone switch to receive SNMP trap messages when the port goes up/down:

```
enable snmp traps port-up-down ports 3,5,12-15
```

History

This command was first available in ExtremeXOS 10.6.

Platform Availability

This command is available on all platforms.

restart ports

```
restart ports [all | port_list]
```

Description

Resets autonegotiation for one or more ports by resetting the physical link.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.



Default

N/A.

Usage Guidelines

N/A.

Example

The following command resets autonegotiation on slot 1, port 4 on a modular switch:

```
restart ports 1:4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

run failover

```
run failover {force}
```

Description

Causes a user-specified node failover.

Syntax Description

force	Force failover to occur.
--------------	--------------------------

Default

N/A.

Usage Guidelines

Use this command to cause the primary MSM/MM to failover to the backup MSM/MM on modular switches, or the Master node to failover to the Backup node in SummitStack.

Before you initiate failover, use the `show switch {detail}` command to confirm that the nodes are in sync and have identical software and switch configurations. If the output shows MASTER and BACKUP (InSync), the two MSMs/MMs or nodes are in sync.



If the MSM/MM's software and configuration are not in sync, and both MSMs/MMs are running ExtremeXOS 11.0 or later, or the master and backup SummitStack nodes software and configuration are not in sync and are running ExtremeXOS 12.0 or later, use the `synchronize` command to get the two MSMs/MMs or nodes in sync. This command ensures that the backup has the same software in flash as the master.

**Note**

Both the backup and the master MSMs or nodes must be running ExtremeXOS 11.0 or later to use the `synchronize` command.

If the MSMs/MMs are not in sync, and one MSM/MM is running ExtremeXOS 10.1 or earlier, specify the force option of the `run failover` command. By specifying force, failover occurs regardless of the version of software running on the MSMs/MMs.

Example

The following command causes a failover on a modular switch or SummitStack:

```
run failover
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on modular switches and SummitStack.

run msm-failover

```
run msm-failover {force}
```

Description

Causes a user-specified node failover. This command is not supported on a SummitStack. To do a failover on a stack, use the command `run failover`.

Syntax Description

force	Force failover to occur.
--------------	--------------------------

Default

N/A.



Usage Guidelines

This command is being replaced with the `run failover` command. For usage guidelines, see the description for the `run failover` command.

Example

The following command causes a user-specified MSM failover:

```
run msm-failover
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available only on modular switches.

show cdp

show cdp

Description

Displays the interval between advertisements, the hold time and the version of the advertisement.

Syntax Description

There are no arguments or keywords for this command.

Default

N/A.

Usage Guidelines

Use this command to display the interval between advertisements, the hold time and the version of the advertisement.

Example

The following command displays specific information on the CDP feature:

```
# show cdp
    CDP Transmit time   : 60 seconds
```



```

CDP Hold time      : 180 seconds
CDP Device ID     : 00:04:96:27:80:03
CDP V1 enabled ports: 1  2  3

```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

show cdp counters

```
show cdp counters {ports port_list}
```

Description

Displays CDP port counter statistics.

Syntax Description

ports	Displays CDP port statistics.
<i>port_list</i>	Port list of CDP ports.

Default

N/A.

Usage Guidelines

Use this command to display CDP port counter statistics.

Example

The following command displays counter statistic for CPD ports:

```

# show cdp counters ports 10
Port    Tx PDU      Rx PDU      Rx Drop      Rx Drop
Count   Count       Count       Checksum Err Invalid
-----
10      3729        3729        0            0

```

History

This command was first available in ExtremeXOS 15.4.



Platform Availability

This command is available on all platforms.

show cdp neighbor

```
show cdp neighbor {detail}
```

Description

Displays information about neighbors.

Syntax Description

detail	Displays detailed information.
---------------	--------------------------------

Default

N/A.

Usage Guidelines

Use this command to display CDP neighbor information.

Example

The following command displays neighbor information for the CDP feature:

```
# show cdp neighbors

Device ID Local-Interface  Hold time  Capability  Platform  Port ID
-----
R1        Port19              180        R           cisco     Ethernet0
-----
Capability Codes: R - Router, T - Trans Bridge, B - Source Route
Bridge,
                S - Switch, H - Host, I - IGMP, r - Repeater

# show cdp neighbors detail

Device ID           : R1
Port ID (outgoing port): Ethernet0
IP Addresses        : 192.168.10.1
Platform            : cisco
Interface           : Port 19
Hold time           : 180seconds

Version:
Cisco Internetwork Operating System Software
IOS (tm) 1600 Software (C1600-NY-L), Version 11.2(12)P, RELEASE SOFTWARE
```



```
(fc1)
  Copyright (c) 1986-1998 by cisco Systems, Inc.
  Compiled Tue 03-Mar-98 06:33 by dsmith
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

show cdp ports

```
show cdp ports {port_list}
```

Description

Displays information about neighbors in ports.

Syntax Description

<i>port_list</i>	Specifies the port list, separated by a comma.
------------------	--

Default

N/A.

Usage Guidelines

Use this command to display port information.

Example

The following command displays port information on the CDP feature:

```
# show cdp ports
Port  Device-ID  Hold time  Remote CDP  Port
ID                                     Version
-----  -
19      R1           180        Version-1   Ethernet0
```

History

This command was first available in ExtremeXOS 15.4.



Platform Availability

This command is available on all platforms.

show ces

```
show ces {ces_name} {detail}
```

Description

Displays the specified circuit emulation services (CES) pseudowire parameters.

Syntax Description

ces	Circuit Emulation Service.
<i>ces_name</i>	Alphanumeric string identifying CES pseudowire.
detail	Displays detailed information.

Default

The default is no detail.

Usage Guidelines

Use this command to display the CES pseudowire parameters.

Example

```
Switch.1 # show ces
CES Name   Pseudowire  CES      Service   Peer Address   State
ID         Flags       Name
=====
==
cesop10    0 EACUTL--  serBun10  11.100.100.219  Up
satop10    0 EASFT-R-  serBun11  00:04:96:51:5F:31  Down
mpls_ces   0 EACMs---  serBun_10  10.10.10.1      Sgnl
=====
==
CES Flags      : (A) Oper Active, (b) CESoPSN Basic, (c) CESoPSN with CAS,
(d) Differential Timestamping, (e) Struct-Agnostic EloPSN,
(E) Admin Enabled, (f) Fragmentation, (F) MEF8,
(L) PW Local Alarm, (M) MPLS, (r) RTP Header,
(R) PW Remote Alarm, (s) LDP Signaled PW,
(t) Struct-Agnostic TloPSN, (T) Static PW, (U) IPv4/UDP
-----
Total number of CES PWs configured :      3
Total number of CES PWs active     :      2
E4G-400.14 # sh ces detail
CES PW Name   : cesop10
```



```

Type                : Static
PW ID               : 0                      Admin State      : Enabled
PSN Transport       : IPv4/UDP              Oper State       : Enabled
Transport Type      : E1, Structured-agnostic
Signaling Mode      : None
Service Name        : serBun10
Payload Size        : 256 bytes
Packet Latency      : 1000 us
Jitter Buffer        : 3000 us (max: 6000 us)
Clock Recovery       : Adaptive
LOPS Entry Thresh   : 8
LOPS Exit Thresh    : 8
Filler Pattern      : 255 (0xFF)
QOS Profile         : QP1
DSCP Value          : 63
TTL Value           : 254
Source Address      : 1.1.2.62
Peer Address        : 1.1.1.101
Local UDP Port      : 3000
Remote UDP Port     : 3100
PW Admin State      : Enabled
PW State            : Up
PW Uptime           : 0d:0h:1m:15s
PW Installed        : TRUE
Local PW Status     : No Faults
Remote PW Status    : No Faults
PW Rx Pkts         : 28474                  PW Tx Pkts       : 388938
CES PW Name        : satop11
Type                : Static
PW ID               : 0                      Admin State      : Enabled
PSN Transport       : MEF8                  Oper State       : Enabled
Transport Type      : T1, Structured-agnostic
Signaling Mode      : None
Service Name        : serBun11
Payload Size        : 256 bytes
Packet Latency      : 1000 us
Jitter Buffer        : 3000 us (max: 6000 us)
Clock Recovery       : None
LOPS Entry Thresh   : 8
LOPS Exit Thresh    : 8
Filler Pattern      : 255 (0xFF)
QOS Profile         : QP1
Source Address      : 02:04:96:7f:0b:65
Peer Address        : 00:00:00:00:00:02
Local MEF8 ECID     : 2000
Remote MEF8 ECID    : 2100
PW Admin State      : Enabled
PW State            : Up
PW Uptime           : 0d:0h:0m:56s
PW Installed        : TRUE
Local PW Status     : No Faults
Remote PW Status    : No Faults
Transport VLAN      : v1
PW Rx Pkts         : 28474                  PW Tx Pkts       : 388938
E4G-200.2 # show ces "mpls-ces" detail
CES PW Name        : mpls-ces
Type                : Signaled
PW ID               : 101                   Admin State      : Enabled
    
```



```

PSN Transport           : MPLS                Oper State : Enabled
Transport Type         : E1, Structure-locked
Signaling Mode         : None
Service Name           : mpls-cesop-s5
Payload Size           : 48 bytes
Packet Latency         : 1000 us
Jitter Buffer           : 3000 us (max: 6000 us)
Clock Recovery         : None
LOPS Entry Thresh     : 8
LOPS Exit Thresh      : 8
Filler Pattern         : 255 (0xFF)
QOS Profile            : QP1
PW Label TTL           : 4
Peer Address           : 2.2.2.2
PW Admin State        : Enabled
PW State               : Up
PW Uptime              : 0d:2h:9m:4s
PW Installed           : TRUE
Local PW Status        : Fault, Att-Rx, PW-Rx
Remote PW Status       : Fault, PW-Tx
Transport LSP         : LDP LSP (Not Configured)
Next Hop I/F          : vtl
Next Hop Addr         : 40.0.0.2           Tx Label      : 0x77
PW Rx Label           : 0x7A             PW Tx Label   : 0x79
PW Signaled Params    :
Payload Size          : Local=48         Remote=48
Bit-Rate              : Local=6          Remote=6
Fragmentation         : Local=Disabled   Remote=Disabled
Signaling              : Local=Disabled   Remote=Disabled
RTP Header            : Local=Disabled   Remote=Disabled
Signaling Error       : None
PW Rx Pkts           : 28474            PW Tx Pkts    : 388938

```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

show ces clock-recovery

```
show ces {ces_name} clock-recovery
```

Description

Displays the CES pseudowire clock recovery information.



Syntax Description

ces	Circuit Emulation Service.
<i>ces_name</i>	Alphanumeric string identifying CE.
clock-recovery	Displays clock recovery information.

Default

N/A.

Usage Guidelines

Use this command to display the CES pseudowire clock recovery information.

Example

```
CES PW Clock Master   : ces_1
Timed TDM port(s)    : 39
Algorithm              : Adaptive
Recovery state        : Free Running
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

show ces errors

```
show ces {ces_name} errors {total | intervals | day-intervals | current {no-refresh}}
```

Description

Displays the specified circuit emulation services (CES) pseudowire error counters.

Syntax Description

ces	Circuit Emulation Service.
<i>ces_name</i>	Alphanumeric string identifying the CES pseudowire.
errors	Displays errors.



intervals	Displays errors for 15-minute intervals.
day-intervals	Displays errors for 1-day intervals.
current	Displays current 15-minute interval errors (default).
no-refresh	Page by page display without auto-refresh.

Default

The default displays current error statistics.

Usage Guidelines

The current switch displays the error statistics of the ongoing 15-minute interval. The total switch displays the cumulative error statistics.

Example

Only one space is given between the first and second field to accommodate the output within 80 characters.

```
E4G-400.3 # show ces {<ces_name>} errors current
CES Statistics
CES      Missing Reorder      JBU      MD      UAS      SES  Err      Fail
MFP
Name     Packets  Packets                               Secs    Cnt
=====
==
cesopl           0      0      0      0      0           0
0 0
satopl23>       0      0      0      0      0           0
0 0
=====
==
> indicates Circuit Emulation Service Name truncated past 8 characters
U->page up  D->page down ESC->exit
Legend: JBU - Jitter Buffer Underrun, MD - Misorder Drops,
MFP - Malformed Packets, SES - Severely Errored Seconds,
UAS - Unavailable Seconds
E4G-400.3 # show ces {<ces_name>} errors current no-refresh
CES Statistics
CES      Missing Reorder      JBU      MD      UAS      SES  Err      Fail
MFP
Name     Packets  Packets                               Secs    Cnt
=====
==
cesopl           0      0      0      0      0           0
0 0
satopl23>       0      0      0      0      0           0
0 0
=====
==
> indicates Circuit Emulation Service Name truncated past 8 characters
Legend: JBU - Jitter Buffer Underrun, MD - Misorder Drops,
```



MFP - Malformed Packets, SES - Severely Errored Seconds,
 UAS - Unavailable Seconds

E4G-400.3 # show ces {<ces_name>} errors total

```
CES Name : cesopl0
Missing Packets      : 3455
Reordered Packets   : 343243650
Jitter Buffer Underrun Packets : 43550
Misordered Dropped Packets : 0
Malformed Packets   : 0
Errored Seconds     : 23240
Severely Errored Seconds : 0
Unavailable Seconds : 0
Failure Counts      : 0
```

E4G-400.1 # show ces errors intervals

```
CES Statistics - Interval
CES Name\ Missing ReOrder JBU MD UAS SES ERR
Fail MFP
Interval Packets Packets                               Secs  Cnt
=====
====
mef8-c1 >\
1      0      0      0      0      900      0      0      0      0
2      0      0      0      0      900      0      0      0      0
3      0      0      0      0      900      0      0      0      0
4      0      0      0      0      882      17      1      1      0
```

```
-----
mef8-c2 >\
1      0      0      0      0      900      0      0      0      0
2      0      0      0      0      900      0      0      0      0
3      0      0      0      0      900      0      0      0      0
4      0      0      0      0      891      10      1      1      0
```

E4G-400.2 # show ces errors day-intervals

```
CES Statistics - Day Interval
CES Name\ Missing ReOrder JBU MD UAS SES ERR
Fail MFP
Interval Packets Packets                               Secs  Cnt
=====
===
mef8-ces3\
1      0      0      0      0 3582  17      1      1      0
2      0      0      0      0 3231  47      1      1      0
3      0      0      0      0 3582  17      1      1      0
4      0      0      0      0 3231  47      1      1      0
```

```
---
mef8-ces4\
1      0      0      0      0 3591  10      1      1      0
2      0      0      0      0 3231  40      1      1      0
3      0      0      0      0 3582  10      1      1      0
4      0      0      0      0 3231  40      1      1      0
```



History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

show ces peer

```
show ces peer [ipaddress ipaddress | mac-address mac_address]
```

Description

Displays the specified circuit emulation services (CES) peer information.

Syntax Description

ces	Circuit Emulation Service.
peer	Displays CES peer information.
ipaddress	Peer IPV4 address.
<i>ipaddress</i>	IPv4 address.
mac-address	Peer MAC address.
<i>mac_address</i>	MAC address; type=mac_t.

Default

N/A.

Usage Guidelines

Use this command to display the CES peer information.

Example

```
Switch.1 # show ces peer ipaddress 11.100.100.219
CES Name   Pseudowire  CES      Service   Peer Addr
Id         Flags       Name
=====
cesop10    0  EACUTL--  serBun10  11.100.100.219
=====
CES Flags   : (A) Oper Active, (C) Circuit Emulation Service over Packet,
              (E) Admin Enable, (F) MEF8, (L) PW Local Alarm,
              (M) MPLS, (R) PW Remote Alarm, (s) LDP Signaled PW,
```



```

(S) Structure-Agnostic TDM over Packet, (T) Static PW,
(U) IPv4/UDP
Switch.1 # show ces peer mac-address 00:04:96:51:5F:31
CES Name   Pseudowire  CES      Service   Peer Addr
Id         Flags       Name
=====
==
satop10           0  EASFT-R-  serBun_10  00:04:96:51:5F:31
01234567> 0123456789 01234567 01234567> 01234567890123456
=====
==
CES Flags      : (A) Oper Active, (C) Circuit Emulation Service over Packet,
                  (E) Admin Enable, (F) MEF8, (L) PW Local Alarm,
                  (M) MPLS, (R) PW Remote Alarm, (s) LDP Signaled PW,
                  (S) Structure-Agnostic TDM over Packet, (T) Static PW,
                  (U) IPv4/UDP
    
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

show dwdm channel-map

```
show dwdm channel-map { channel_first { - channel_last } } {port port_num}
```

Description

Displays the channel scheme adopted for mapping the DWDM wavelengths.

Syntax Description

<i>channel_first</i>	Specifies the starting channel number.
<i>channel_last</i>	Specifies the ending channel number.
<i>port_num</i>	Specifies the port for which the status is to be displayed.

Default

N/A.

Usage Guidelines

Use this command to display the wavelength and the supportability of the channel by the optical module in the port.



Example

The following command displays information for channels 50 through 60 on port 3:1:

```
show dwdm channel-map 50 - 60 port 3:1
```

The following is sample output for this command:

```
=====
Channel #      Wavelength (nm)      Port 3:1
=====
50             1537.40              Supported
51             1536.61              Supported
52             1535.82              Supported
...           .....
58             1531.12              Supported
59             1530.33              Supported
60             1529.55              Supported
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on BlackDiamond X8 switches, BlackDiamond 8800 switches with 10G8Xc, 10G4Xc, 8900-10G8X-xl modules or S-10G1Xc option cards, and Summit X480 switches with XGM2-2xf option cards or VIM2-10G4X modules.

show edp

```
show edp {ports [all | ports] {detail}}
```

Description

Displays connectivity and configuration information for neighboring Extreme Networks switches.

Syntax Description

all	Specifies all ports.
<i>ports</i>	Specifies one or more ports or slots and ports.
detail	Show detailed information.

Default

N/A.



Usage Guidelines

On a modular switch, *ports* can be a list of slots and ports. On a stand-alone switch, *port_list* can be one or more port numbers. For a detailed explanation of port specification, see [Port Numbering in Command Reference Overview](#). To clear the counters, use the `clear lacp counters` command.

The neighbor-ID value is eight bytes. The first two bytes are always set to 00:00; the last six bytes are set to the neighbor's system MAC address.

Use the `show edp` command to display neighboring switches and configurations. This is most effective with Extreme Networks switches.

Example

The following command displays the configuration of the switch:

```
show edp
```

Following is sample output from this command:

```
EDP advert-interval      :60 seconds
EDP holddown-interval    :180 seconds
EDP enabled on ports     :1:1  1:2  1:3  1:4  1:5  1:6  3:1  3:2  3:3  3:4
```

Following is sample output from the `show edp ports 1:1` command:

Port	Neighbor	Neighbor-ID	Remote	Age	Num
Port	Vlans				
1:1	Oban	00:00:00:30:48:41:ed:97	1:1	54	1

The following command displays the connectivity and configuration of neighboring Extreme Networks switches:

```
show edp ports 1:1 detail
```

Following is sample output from this command:

```
=====  
Port 1:1: EDP is Enabled  
Tx stats: sw-pdu-tx=2555      vlan-pdu-tx=1465      pdu-tx-err=0  
Rx stats: sw-pdu-rx=2511      vlan-pdu-rx=2511      pdu-rx-err=0  
Time of last transmit error: None  
Time of last receive error:  None  
Remote-System:      Oban      Age = 41  
Remote-ID:          00:00:00:30:48:41:ed:97  
Software version:   11.1.0.19  
Remote-Port:        1:1
```



```

Port Type:           Ethernet
Auto Negotiation:    OFF
Flow Control:        SYMMETRIC/ASYMMETRIC
Duplex Speed:        Configured = HALF      Actual = HALF
Port Speed (MB):     Configured = ERROR    Actual = 100 Mbps
Remote-Vlans:
test (4094) Age = 41
=====

```

Note



The output differs if the port is connected to a port running ExtremeWare® software; the output shown above is displayed when both connected ports are running ExtremeXOS software.

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show ip-fix

show ip-fix

Description

Displays IPFIX information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to display IPFIX information regarding the global state, collector information and which ports are enabled.



Example

The following command displays IPFIX information:

```
show ip-fix
```

Following is sample output on a BlackDiamond 8800 switch:

```
BD-8810.2 # show ip-fix
Global IPFIX State: Enabled
Domain:                0
Source (Exporter) IP:  Default(Switch IP)      Virtual Router:      "VR-
Mgmt "
Collector IP:          10.66.9.32
Transport Protocol/Port/State:  SCTP, Port 4739, NA
Flow Keys:
IPv4:  "Source-IP"  "Source L4 Port"  "Dest-IP"  "Dest L4 Port"  "L4
Protocol"  "TOS"
Unused:
IPv6:  "Source-IP"  "Source L4 Port"  "Dest-IP"  "Dest L4 Port"  "Next
Header"  "TOS"  "Flow Label"
Unused:
Non-IP: "Source MAC"  "Dest MAC"  "Ethertype"  "VLAN ID"  "Priority"  "Tagged"
Unused:
IPFIX Enabled on Port(s):  3:1, 4:1
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on BlackDiamond 8900 G96Tc, G48T-xl, G48X-xl, and 10G8X-xl modules and Summit X460 and X480 switches.

show lacp

```
show lacp
```

Description

Displays LACP, or dynamic link aggregation, settings on the switch.

Syntax Description

This command has no arguments or variables.



Default

N/A.

Usage Guidelines

This command displays the following information about the LACP LAGs configured on the switch:

- Up or Down.
- Enabled or disabled (not configurable).
- System MAC.
 - MAC address for the system, which is used for LACP priority in the absence of a specifically configured priority.
- LACP PDUs dropped on non-LACP ports.
- LAG.
 - Identifies the particular LAG. This number comes from logical port assigned to the LAG and is the LAG group ID.
- Actor Sys-Pri.
 - Shows the system priority for that LAG.
 - If this number is lower than the number displayed for the Partner Sys-Pri, the system you are working on is the controlling partner in the LAG.
- Actor Key.
 - Automatically generated LACP key.
- Partner MAC.
 - Identifies the MAC address for the system connecting to the LAG on the remote end.
- Partner Sys-Pri.
 - Shows the system priority for that LAG on the remote end.
 - If this number is lower than the number displayed for the Actor Sys-Pri, the system at the remote end is the controlling partner in the LAG.
- Partner Key.
 - LACP key automatically generated by the system to which this aggregator is connected.
 - If this number is lower than the number displayed for the Actor Key, the partner system is the controlling partner in the LAG.
- Agg Count.
 - Identifies the number of ports added to the aggregator for that LAG.

Example

The following command displays the LACP LAGs on the switch:

```
show lacp
```

The following is sample output from this command on a modular switch:

```
LACP Up                : Yes
LACP Enabled           : Yes
System MAC             : 00:04:96:10:33:60
LACP PDUs dropped on non-LACP ports : 0
```



Lag	Actor	Actor	Partner	Partner	Partner	Agg
Sys-Pri	Key	MAC	Sys-Pri	Key	Count	

--						
2:1	90	0x07d1	00:01:30:f9:9c:30	601	0x1391	2
4:5	100	0x0fa5	00:01:30:f9:9c:30	321	0x1f47	16
4:9	677	0x0fa9	00:01:30:f9:9c:30	87	0x0fa9	8

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

show lacp counters

show lacp counters

Description

Displays all LACP, or dynamic link aggregation, counters for all member ports in the system.

Syntax Description

This command has no parameters or variables.

Default

N/A.

Usage Guidelines

This command displays the following information for all link aggregation groups (LAGs):

- LACP PDUs dropped on non-LACP ports.
- LACP bulk checkpointed messages sent.
- LACP bulk checkpointed messages received.
- LACP PDUs checkpointed sent.
- LACP PDUs checkpointed received.
- LAG group ID.
- Member port.
- Packets received.
- Packets dropped from PDU error.
- Packets dropped because LACP is not enabled on this port.
- Packets dropped because sender's system MAC address matches that of receiver.



- Packets successfully transmitted.
- Packets with errors during transmission.

Example

The following command displays LACP counters:

```
show lacp counters
```

The following is sample output from this command on a modular switch:

```
LACP PDUs dropped on non-LACP ports : 519392
LACP Bulk checkpointed msgs sent    : 1
LACP Bulk checkpointed msgs recv   : 0
LACP PDUs checkpointed sent        : 575616
LACP PDUs checkpointed recv        : 0
Lag      Member      Rx      Rx Drop  Rx Drop  Rx Drop  Tx      Tx
Group    Port             Ok       PDU Err  Not Up   Same MAC Sent Ok  Xmit Err
-----
--
1:1      1:1              2169    0         0         0         2170    0
1:2      2169             0        0         0         2170     0
1:3      2169             0        0         0         2170     0
1:4      2169             0        0         0         2170     0
1:5      2169             0        0         0         2170     0
1:6      2169             0        0         0         2170     0
1:7      2169             0        0         0         2170     0
1:8      2168             0        0         0         2169     0
=====
==
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

show lacp lag

```
show lacp lag group-id {detail}
```

Description

Displays LACP, or dynamic link aggregation, settings for the specified LAG.



Syntax Description

<i>group-id</i>	Specifies the LAG group ID you want to display. This is the number of the port you configured as the logical port of the LAG.
detail	Show detailed information.

Default

N/A.

Usage Guidelines

This command displays the following information about the specified LACP LAG:

- LAG
 - Identifies the particular LAG. This number comes from logical port assigned to the LAG and is the LAG group ID.
- Actor Sys-Pri
 - Shows the system priority for that LAG.
 - If this number is lower than the number displayed for the Partner Sys-Pri, the system you are working on is the controlling partner in the LAG.
- Actor Key
 - Automatically generated LACP key.
- Partner MAC
 - Identifies the MAC address for the system connecting to the LAG on the remote end.
- Partner Sys-Pri
 - Shows the system priority for that LAG on the remote end.
 - If this number is lower than the number displayed for the Actor Sys-Pri, the system at the remote end is the controlling partner in the LAG.
- Partner Key
 - LACP key automatically generated by the system to which this aggregator is connected.
 - If this number is lower than the number displayed for the Actor Key, the partner system is the controlling partner in the LAG.
- Agg Count
 - Identifies the number of ports added to the aggregator for that LAG.
- Member port
- Port priority
- Rx State—Receiving state of the port
 - Idle.
 - Initialized.
 - Current—Receiving LACP PDUs.
 - Expired.
 - Defaulted.
- Sel Logic—Selection state of the port
 - Selected—Ports with a matching admin key on the remote end.
 - Unselected—Ports that failed to meet with a matching admin key on the remote end.



- Standby—Ports that exceed the number of ports that can be active in the LAG simultaneously. These ports can be moved into selected mode if one of the currently selected ports in the LAG goes down.
- Mux State—Ability to transmit and collect data of the port.
 - Waiting—Selected port that is waiting for LACP to determine if it can join the aggregator.
 - Attached—Ports ready to be added to the aggregator.
 - Collecting-Dist—Ports that are added to the aggregator and are transferring data.
 - Detached—Ports that cannot be added to the aggregator.
- Actor Flag—Mux state of the port.
 - A—Activity.
 - T—Timeout.
 - G—Aggregation.
 - S—Synchronization.
 - C—Collecting.
 - D—Distributing.
 - F—Defaulted.
 - E—Expired.
- Partner Port
 - The operational value of the port number assigned to this link by partner.
- Up—Yes or no.
- Enabled—Yes or no.
- Unack count.
- Wait-for-count.
- Current timeout.
- Activity mode.
- Defaulted action.
- Receive state.
- Transmit state.
- Selected count—Number of selected ports in the LAG.
- Standby count—Number of standby ports in the LAG.
- LAG Id flag
 - S—Displays information on controlling partner of LAG.
 - T—Displays information on controlled partner of LAG.

Example

The following command displays information on the specified LACP LAG:

```
show lacp lag 4:9
```

The following is sample output from this command on a modular switch:

```
Lag          Actor      Actor      Partner          Partner  Partner  Agg      Actor
```



```

-----
          Sys-Pri  Key      MAC                Sys-Pri  Key      Count  Mac
-----
4:9      2110    0x0fa9  00:04:96:10:33:60  2110    0x0fa9    16
00:22:33:44:55:66

Port list:

Member   Port      Rx      Sel      Mux      Actor
Partner
Port     Priority  State   Logic    State    Flags    Port
-----
--
4:9      300      Current Selected Collect-Dist A-GSCD-- 4009
4:10     301      Current Selected Collect-Dist A-GSCD-- 4010
4:11     302      Current Standby   Detached  A-G----- 4011
4:12     303      Current Standby   Detached  A-G----- 4012
4:29     200      Current Selected Collect-Dist A-GSCD-- 4029
4:30     0        Current Selected Collect-Dist A-GSCD-- 4030
4:31     202      Current Selected Collect-Dist A-GSCD-- 4031
4:32     203      Current Selected Collect-Dist A-GSCD-- 4032
8:7      101      Current Selected Collect-Dist A-GSCD-- 8013
8:8      10       Current Selected Collect-Dist A-GSCD-- 8014
8:9      9        Current Selected Collect-Dist A-GSCD-- 8015
8:10     8        Current Selected Collect-Dist A-GSCD-- 8016
8:11     7        Current Selected Collect-Dist A-GSCD-- 8017
8:12     6        Current Selected Collect-Dist A-GSCD-- 8018
8:13     5        Current Selected Collect-Dist A-GSCD-- 8019
8:14     3        Current Selected Collect-Dist A-GSCD-- 8020
8:15     0        Current Selected Collect-Dist A-GSCD-- 8043
8:16     3        Current Selected Collect-Dist A-GSCD-- 8044
8:17     2        Idle     Unselected Detached  ----- 0
8:18     37       Idle     Unselected Detached  ----- 0
8:19     36       Idle     Unselected Detached  ----- 0
8:20     35       Idle     Unselected Detached  ----- 0
=====
==
Actor Flags: A-Activity, T-Timeout, G-Aggregation, S-Synchronization
              C-Collecting, D-Distributing, F-Defaulted, E-Expired

```

The following command displays detailed information on the specified LACP LAG:

```
show lacp lag 4:9 detail
```

The following is sample output from this command on a modular switch:

```

Lag      Actor   Actor   Partner           Partner  Partner  Agg   Actor
          Sys-Pri Key     MAC                Sys-Pri  Key     Count
Mac
-----
4:9      2110    0x0fa9  00:04:96:10:33:60  2110    0x0fa9    16
00:22:33:44:55:66
Up       : Yes
Enabled  : Yes
Unack count : 0

```



```

Wait-for-count      : 0
Current timeout     : Long
Activity mode       : Active
Defaulted Action    : Delete
Receive state       : Enabled
Transmit state      : Enabled
Selected count      : 16
Standby count       : 2
LAG Id flag         : Yes
S.pri:2110, S.id:00:01:30:f9:9c:30, K:0x0fa9
T.pri:2110, T.id:00:04:96:10:33:60, L:0x0fa9
    
```

Port list:

Member	Port	Rx	Sel	Mux	Actor	
Partner	Priority	State	Logic	State	Flags	Port

--						
4:9	300	Current	Selected	Collect-Dist	A-GSCD--	4009
4:10	301	Current	Selected	Collect-Dist	A-GSCD--	4010
4:11	302	Current	Standby	Detached	A-G-----	4011
4:12	303	Current	Standby	Detached	A-G-----	4012
4:29	200	Current	Selected	Collect-Dist	A-GSCD--	4029
4:30	0	Current	Selected	Collect-Dist	A-GSCD--	4030
4:31	202	Current	Selected	Collect-Dist	A-GSCD--	4031
4:32	203	Current	Selected	Collect-Dist	A-GSCD--	4032
8:7	101	Current	Selected	Collect-Dist	A-GSCD--	8013
8:8	10	Current	Selected	Collect-Dist	A-GSCD--	8014
8:9	9	Current	Selected	Collect-Dist	A-GSCD--	8015
8:10	8	Current	Selected	Collect-Dist	A-GSCD--	8016
8:11	7	Current	Selected	Collect-Dist	A-GSCD--	8017
8:12	6	Current	Selected	Collect-Dist	A-GSCD--	8018
8:13	5	Current	Selected	Collect-Dist	A-GSCD--	8019
8:14	3	Current	Selected	Collect-Dist	A-GSCD--	8020
8:15	0	Current	Selected	Collect-Dist	A-GSCD--	8043
8:16	3	Current	Selected	Collect-Dist	A-GSCD--	8044
8:17	2	Idle	Unselected	Detached	-----	0
8:18	37	Idle	Unselected	Detached	-----	0
8:19	36	Idle	Unselected	Detached	-----	0
8:20	35	Idle	Unselected	Detached	-----	0

=====
 ==

Actor Flags: A-Activity, T-Timeout, G-Aggregation, S-Synchronization
 C-Collecting, D-Distributing, F-Defaulted, E-Expired

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.



show lacp member-port

```
show lacp member-port port {detail}
```

Description

Displays LACP, or dynamic link aggregation, settings for the specified port that is a member of any LAG.

Syntax Description

<i>port</i>	Specifies the port number.
detail	Show detailed information.

Default

N/A.

Usage Guidelines

This command displays the following information about the specified port:

- Member Port.
- Port Priority.
- Rx State—Receiving state of the port.
 - Idle.
 - Initialized.
 - Current—Receiving LACP PDUs.
 - Expired.
 - Defaulted.
- Sel Logic—Selection state of the port.
 - Selected—Ports with a matching admin key on the remote end.
 - Unselected—Ports that failed to meet with a matching admin key on the remote end.
 - Standby—Ports that exceed the number of ports that can be active in the LAG simultaneously. These ports can be moved into selected mode if one of the currently selected ports in the LAG goes down.
- Mux State—Ability to transmit and collect data of the port.
 - Waiting—Selected port that is waiting for LACP to determine if it can join the aggregator.
 - Attached—Ports ready to be added to the aggregator.
 - Collecting-Dist—Ports that are added to the aggregator and are transferring data.
 - Detached—Ports that cannot be added to the aggregator.
- Actor Flag.
 - A—Activity.
 - T—Timeout.
 - G—Aggregation.



- S—Synchronization.
- C—Collecting.
- D—Distributing.
- F—Defaulted.
- E—Expired.
- Partner Port.
 - The operational value of the port number assigned to this link by partner.
- Up or Down—LACP protocol running or not on specified port.
- Enabled or disabled (not configurable).
- Link State—Link state on this port up or down.
- Actor Churn—True or false.
- Partner Churn—True or false.
- Ready_N—Ready to be added to aggregator.
- Wait pending.
- Ack pending.
- LAG Id.
 - S—Displays information on controlling partner of LAG.
 - T—Displays information on controlled partner of LAG.
- Stats.
 - Rx - Accepted.
 - Rx - Dropped due to error in verifying PDU.
 - Rx - Dropped due to LACP not being up on this port.
 - Rx - Dropped due to matching own MAC.
 - Tx - Sent Successfully.
 - Tx - Transmit error.

Example

The following command displays LACP information on the specified port:

```
show lacp member-port 4:9
```

The following is sample output from this command on a modular switch:

```

Member      Port      Rx      Sel      Mux      Actor
Partner
Port      Priority  State   Logic    State    Flags    Port
-----
--
4:9       300      Current Selected Collect-Dist A-GSCD-- 4009
=====
==
Actor Flags: A-Activity, T-Timeout, G-Aggregation, S-Synchronization
C-Collecting, D-Distributing, F-Defaulted, E-Expired

```



The following command displays detailed LACP information on the specified port:

```
show lacp member-port 4:9 detail
```

The following is sample output from this command on a modular switch:

```

Member      Port      Rx      Sel      Mux      Actor
Partner
Port        Priority  State   Logic    State    Flags    Port
-----
--
4:9         300      Current Selected Collect-Dist A-GSCD-- 4009
Up          : Yes
Enabled     : Yes
Link State  : Up
Actor Churn : False
Partner Churn : False
Ready_N    : Yes
Wait pending : No
Ack pending : No
LAG Id:
S.pri:2110, S.id:00:01:30:f9:9c:30, K:0x0fa9, P.pri:300 , P.num:4009
T.pri:2110, T.id:00:04:96:10:33:60, L:0x0fa9, Q.pri:300 , Q.num:4009
Stats:
Rx - Accepted                : 2174
Rx - Dropped due to error in verifying PDU : 0
Rx - Dropped due to LACP not being up on this port : 0
Rx - Dropped due to matching own MAC      : 0
Tx - Sent successfully        : 2175
Tx - Transmit error          : 0
=====
==
Actor Flags: A-Activity, T-Timeout, G-Aggregation, S-Synchronization
C-Collecting, D-Distributing, F-Defaulted, E-Expired
    
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

show mirror

```
show mirror mirror_name | mirror_name_li | all | enabled
```

Description

Displays various show output for mirror instances.



Syntax Description

<i>mirror_name</i>	Displays the mirror name.
<i>mirror_name_li</i>	Displays mirror instance name for Lawful Intercept account only. You must have lawful intercept user privileges to specify this variable.
all	Displays all mirror instances.
enabled	Displays only enabled mirror instances.

Default

N/A.

Usage Guidelines

Use this command to display mirror statistics and determine if mirroring is enabled or disabled on the switch.

Example

The following command displays switch mirroring statistics:

```
show mirror all
Mirror MyDebugMirror (Disabled)
  Description: This mirror sends traffic to PC-based Wireshark
  Mirror to ports: 4-5 Loopback port: 6
  Number of Mirroring filters:7
  Mirroring filters configured:
    Port number 7 in all vlans (ingress)
    Port number 8 in vlan Default (ingress-and-egress)
    Port number 9 in all vlans (egress)
    Port number 10 in all vlans (ingress-and-egress)
    Port number 11 in vlan Default (ingress)
    Port number 12 in vlan Default (ingress)
    Port number 13 in vlan Default (ingress)
  Active ACL target action: Yes

Mirror DefaultMirror (Enabled)
  Description: This mirror is used for debugging VRRP
  Mirror to port: 15 is up
  Number of Mirroring filters:1
  Mirroring filters configured:
    Port number 17 in vlan to_oslo (ingress)
  Active ACL target action: Yes
```

The following example displays output for a lawful intercept user session:

```
* X460-24p.9 > show mirror
DefaultMirror (Disabled)
  Description: Default Mirror Instance, created automatically
  Mirror to port: -

law_mirror (Enabled)
  Description: user for lawful intercept
```



```
Mirror to port: 3
Source filter instances used : 2
  Port 7, all vlans, ingress only
  Port 8, all vlans, ingress only

main_mirror  (Enabled)
Description:
  Mirror instance for Admin
  Mirror to port: 2
  Source filter instances used : 1
  Port 10, all vlans, ingress only

Mirrors defined:          3
Mirrors enabled:         2 (0 with egress filters)
HW filter instances used: 3 (Maximum 128)
```

History

This command was first available in ExtremeXOS 15.3.

The *mirror_name_li* variable was added in ExtremeXOS 15.3.2.

Platform Availability

This command is available on all platforms.

show mirroring

show mirroring

Description

Displays the port-mirroring configuration on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

You must enable mirroring on the switch prior to configuring mirroring, and you must configure mirroring to display mirroring statistics. Use the `enable mirroring to port` command to enable mirroring and the `configure mirroring add` command to configure mirroring.



You can use this command to display mirroring statistics and determine if mirroring is enabled or disabled on the switch.

Example

The following command displays switch mirroring statistics:

```
show mirroring
```

BlackDiamond 8800 series switches and Summit family switches only

Following is sample output from this command for a BlackDiamond 8810 switch or SummitStack that is configured for port-based mirroring for single monitor ports:

```
Mirror port: 3:15 is up
Number of Mirroring filters: 3
Mirror Port configuration:
  Port number 3:12 in all vlans ingress only
  Port number 5:4 in all vlans egress only
  Port number 8:30 in all vlans
```

Following is sample output from this command on a Summit series switch that is configured for mirroring a virtual port:

```
Mirror port: 12 is down
Number of Mirroring filters: 1
Mirror Port configuration:
  Port number 3 in vlan peggy.
```

Following is a sample output from this command for a Summit switch as enabled for multiple monitor ports:

```
Mirror ports: 5-7 Loopback port: 3
Number of Mirroring filters: 2
Mirror Port configuration:
  Port number 1 in all vlans
  Port number 2 in all vlans
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.



show mlag peer

```
show mlag peer {peer_name}
```

Description

Displays information about an MLAG.

Syntax Description

<i>peer_name</i>	Specifies an alpha numeric string identifying the MLAG peer switch.
------------------	---

Default

N/A.

Usage Guidelines

Use this command to display configured items, MLAG peer switch state, MLAG group count, and health-check statistics.

Example

The following command displays information for an MLAG peer switch:

```
BD-8810.5 # show mlag peer
```

Following is sample output for the command:

```
Multi-switch Link Aggregation Peers:
```

```
MLAG Peer      : leftBD8k
VLAN           : isc
Local IP Address : 1.1.1.2
MLAG ports     : 2
Checkpoint Status : Up
Rx-Hellos      : 184
Rx-Checkpoint Msgs : 12
Rx-Hello Errors : 0
Hello Timeouts  : 1
Up Time        : 0d:0h:0m:10s
Local MAC      :00:04:96:11:22:44
Peer MAC       :00:04:96:11:22:33
Config'd LACP MAC :None
Current LACP MAC:00 :04:96:11:22:33
Virtual Router  : VR-Default
Peer IP Address : 1.1.1.1
Tx-Interval    : 1000 ms
Peer Tx-Interval : 1000 ms
Tx-Hellos      : 184
Tx-Checkpoint Msgs : 12
Tx-Hello Errors : 0
Checkpoint Errors : 0
Peer Conn.Failures : 1
Multi-switch Link Aggregation Peers:
```

```
MLAG Peer      : rightBD8k
VLAN           : isc
Local IP Address : 1.1.1.1
MLAG ports     : 2
Virtual Router  : VR-Default
Peer IP Address : 1.1.1.2
Tx-Interval    : 1000 ms
```



```

Checkpoint Status : Up                Peer Tx-Interval : 1000 ms
Rx-Hellos         : 167              Tx-Hellos        : 167
Rx-Checkpoint Msgs : 12              Tx-Checkpoint Msgs : 12
Rx-Hello Errors   : 0                Tx-Hello Errors   : 0
Hello Timeouts    : 1                Checkpoint Errors : 0
Up Time           : 0d:0h:0m:7s      Peer Conn.Failures : 1
Local MAC         : 00:04:96:11:22:44 Peer MAC          :
00:04:96:11:22:33
Config'd LACP MAC :None              Current LACP MAC  :
00:04:96:11:22:33
    
```

Following is sample output when an MLAG peer has been created but the IP address is yet to be configured:

```

* switch # show mlag peer switch101

Multi-switch Link Aggregation Peers:

MLAG peer           : switch101
VLAN                :                Virtual
Router              :
Local IP address    :                Peer IP
address             :
MLAG groups         : 0                Tx-
Interval            :                : N/A
Checkpoint Status   : Down              Peer Tx-
Interval            : N/A
Rx-Hellos           : 0                Tx-
Hellos              : 0
Rx-Checkpoint Msgs : 0                Tx-Checkpoint
Msgs                : 0
Rx-Hello Errors     : 0                Tx-Hello
Errors              : 0
Hello Timeouts      : 0                Checkpoint
Errors              : 0
Up Time: N/A        : 0                Peer
Conn.Failures       : 0
Local MAC           : 00:04:96:11:22:44 Peer MAC          :
00:04:96:11:22:33
Config'd LACP MAC   :None              Current LACP MAC  :
00:04:96:11:22:33
    
```

Following is sample output displaying LACP MAC for a MLAG peer:

```

* switch # show mlag peer switch101

Multi-switch Link Aggregation Peers:

MLAG Peer           :S2
VLAN                :isc                Virtual Router      :VR-Default
Local IP Address     :1.1.1.1            Peer IP Address     :1.1.1.2
MLAG ports           :0                Tx-Interval         :1000 ms
Checkpoint Status    :Up                Peer Tx-Interval    :1000 ms
Rx-Hellos            :153379            Tx-Hellos           :153895
Rx-Checkpoint Msgs   :6                Tx-Checkpoint Msgs  :14
    
```




```

Rx-Hello Errors      :0                Tx-Hello Errors      :0
Hello Timeouts       :0                Checkpoint Errors    :0
Up Time              :1d:17h:45m:8s    Peer Conn. Failures  :0
Local MAC            :00:04:96:11:22:44 Peer MAC              :
00:04:96:11:22:33
Config'd LACP MAC    :None              Current LACP MAC     :
00:04:96:11:22:33    Config'd LACP MAC

```

is the configured LACP MAC using the

```
configure {mLAG peer} <peer_name> lACP-mac <lACP_mac_address>
```

command.

If no MAC is configured,

```
Config'd LACP MAC
```

is shown as None.

If same LACP MAC is configured on both the switches, the current LACP MAC will be the same as

```
Config'd LACP MAC.
```

If LACP MAC is not configured on any of the MLAG peers or if a different MAC is configured on the peers,

```
Current LACP MAC
```

is different from

```
Config'd LACP MAC
```

and is selected from Local MAC/Peer MAC combination.

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

show mlag ports

```
show mlag ports {port_list}
```

Description

Displays information about each MLAG group.

Syntax Description

<i>port_list</i>	Specifies one or more ports.
------------------	------------------------------



Default

N/A.

Usage Guidelines

Use this command to display information about each MLAG group including local port number, local port status, remote MLAG port state, MLAG peer name, MLAG peer status, local port failure count, remote MLAG port failure count, and MLAG peer failure count.

Local and remote link state and fail counts reflect the status of the entire LAG when a LAG is used in conjunction with an MLAG. For example, if 1 and 2 ports in a local LAG on the switch associated with an MLAG is down, the local link state will still show as active and the associated local fail count will be incremented. The remote fail count shown at MLAG neighboring switch will also be incremented.

Example

The following command displays information for an MLAG group:

```
BD-8810.5 # show mlag ports
```

Following is sample output for the command:

```

Local                               Local  Remote
MLAG                               Peer   Status  Fail   Fail
Id    Port    State  Link  Peer   Count  Count
Count
=====
==
2      1:1    A      Up    leftBD8K      0
0
1      1:2    A      Up    leftBD8K      0
0
=====
==
Local Link State: A - Active, D - Disabled, R - Ready, NP - Port not present
Remote Link      : Up - One or more links are active on the remote switch,
Down - No links are active on the remote switch,
N/A - The peer has not communicated link state for this MLAG
port
Number of Multi-switch Link Aggregation Groups : 2
Convergence control                          : Fast

```

The following command displays information about an MLAG group on ports 1 and 2:

```
show mlag port 1,2
```

Following is sample output for the command:

```

Local                               Local  Remote

```



```

MLAG Local Link Remote Peer Fail Fail
Id Port State Link Peer Status Count Count
=====
100 1 A Up switch101 Up 0 2
101 2 A Down switch101 Up 0 1
=====
Local Link State: A - Active, D - Disabled, R - Ready, NP - Port not present
Remote Link: Up - One or more links are active on the remote switch,
Down - No links are active on the remote switch,
N/A - The peer has not communicated link state for this MLAG
group
Number of Multi-switch Link Aggregation Groups: 2
Convergence Control : Conserve Access Lists

```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

show network-clock clock-source

```
show network-clock clock-source
```

Description

Displays the configured network clock source information.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

The E4G-200 and E4G-400 have clock sources beyond SyncE. The clock which drives all of the ports on a switch may be selected from:

- SyncE (Synchronous Ethernet).
- PTP – Precision Time Protocol, an optional 1588v2 module.
- TDM – an optional module which has multiple T1/E1 interfaces for TDM/Ethernet interworking.
- BITS – Building Integrated Timing Supply. A connector capable of receiving a timing signal provided by other building equipment.



Example

The following is an example for an E4G-400 switch:

```
E4G-400.1 # show network-clock clock-source
Input Clock Source      : Sync-E
Input Quality Level    : QL-SEC (11)
Input Region           : E1
Input Clock Status     : LOCAL CLOCK
Output Bits BNC 1 Clock Source : 8KHz
Output Bits BNC 2 Clock Source : E1
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on X460-24X, X480-48X, E4G-200 and E4G-400 switches.

show network-clock sync-e ports

```
show network-clock sync-e port [port_list] {details}
```

Description

Displays information about synchronous Ethernet (SyncE).

Syntax Description

<i>port_list</i>	Specifies the port(s).
details	Show additional details.

Default

N/A.

Usage Guidelines

Use this command to display the SyncE configuration and port state on one or more ports. The following command will display the SyncE configuration and port state for all ports:

```
show network-clock sync-e ports
```



Example

The following command displays SyncE information for ports 1 - 12 on a Summit X460-48x switch:

```
show network-clock sync-e ports 1-12
```

Following is sample output:

```
=====
Port  Flags      Input          Configured      Output          Source
Clock Quality Clock Quality Clock Quality (Port Number)
=====
1      S2          QL-SSU-A
2      S1          QL-PRC
3      S
4      S
5      S
6      S              QL-PRC          SOURCE-1 (2)
7      S
8      S
9      S
10     S
11     S
12     S
=====
Flags   : (S) SyncE Supported, (N) SyncE Not Supported,
        (E) SyncE Enabled,   (f) Forced Source Clock,
        (1) Clock Source 1, (2) Clock Source 2,
```

The following examples show output from an E4G-200 switch:

```
E4G-200.1 # show network-clock sync-e ports 1-8
=====
Port  Flags      Input          Configured      Output          Source
Clock Quality Clock Quality Clock Quality (Port Number)
=====
1      SE1          QL-PRC
2      SE2f         QL-SEC
3      SE
4      SE
5      SE
6      SE
7      SE
8      SE
=====
Flags   : (S) SyncE Supported, (N) SyncE Not Supported,
        (E) SyncE Enabled,   (f) Forced Source Clock,
        (1) Clock Source 1, (2) Clock Source 2,
E4G-200.2 # show network-clock sync-e ports 1-8 detail
Port Number          : 1
SyncE Configured     : SOURCE-1
Type of Input Clock  : QL-PRC   (Primary Reference Clock)
Configured Input Clock : None
SyncE Region         : E1      (European and Asian clock)
```



```

region)
Port Number                : 2
SyncE Configured          : SOURCE-2
Type of Input Clock       :
Configured Input Clock    : QL-SEC      (Synchronous Digital
Hierarchy Equipment Clock)
SyncE Region              : E1         (European and Asian clock
region)
Port Number                : 3
SyncE Output Status       : Enabled
Type of Output Clock      : QL-SEC      (Synchronous Digital
Hierarchy Equipment Clock)
Current Input Source (port) : SOURCE-2  (2)
Port Number                : 4
SyncE Output Status       : Enabled
Type of Output Clock      : QL-SEC      (Synchronous Digital
Hierarchy Equipment Clock)
Current Input Source (port) : SOURCE-2  (2)
Port Number                : 5
SyncE Output Status       : Enabled
Type of Output Clock      : QL-SEC      (Synchronous Digital
Hierarchy Equipment Clock)
Current Input Source (port) : SOURCE-2  (2)
Port Number                : 6
SyncE Output Status       : Enabled
Type of Output Clock      : QL-SEC      (Synchronous Digital
Hierarchy Equipment Clock)
Current Input Source (port) : SOURCE-2  (2)
Port Number                : 7
SyncE Output Status       : Enabled
Type of Output Clock      : QL-SEC      (Synchronous Digital
Hierarchy Equipment Clock)
Current Input Source (port) : SOURCE-2  (2)
Port Number                : 8
SyncE Output Status       : Enabled
Type of Output Clock      : QL-SEC      (Synchronous Digital
Hierarchy Equipment Clock)
Current Input Source (port) : SOURCE-2  (2)

```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on Summit X460-24x and X460-48x switches and on E4G-200 and E4G-400 switches.

show port eee

```
show port port_list eee
```



Description

Displays the EEE statistics for the specified ports.

Syntax Description

<i>port_list</i>	Optionally specifies the list of ports, or slots and ports, for which EEE information is displayed.
------------------	---

Default

N/A.

Usage Guidelines

This command shows various EEE statistics for the specified ports.

Example

The following command displays EEE information:

```

show port <portlist> eee
Port EEE Statistics Monitor                               Sat Dec 29 04:10:53
2012
Port      Link  EEE      Rx      Rx      Tx
Tx
          State State    Events  Duration (uS)  Events  Duration
(uS)
=====
==
3         A    E         1      4515403         1
4515460
4         A    E         1      4515948         1
4515966
=====
==

```

> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback

EEE State: E-Enabled, D-Disabled, NA-Not Available
0->Clear Counters U->page up D->page down ESC->exit

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms that support EEE.



show ports

```
show ports {port_list | tag tag} {no-refresh}
```

Description

Displays port summary statistics.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
tag	Specifies an 802.1Q or 802.1ad tag value.
no-refresh	Specifies a static snapshot of the data.

Default

N/A.

Usage Guidelines

Use this command to display the port number, display string, and some of the port states in tabular form.

The VLAN name is displayed only if that port contains a single VLAN. If the port contains more than one VLAN, then the number of VLANs is displayed.

The tag value may be associated with either a VMAN or a VLAN.

Example

The following command displays on slot 2-3 on port 1 and slot 12 on port 10 on a modular switch:

```
show ports 1:2-3,10:12
```

Following is sample output from this command:

```
show ports 1:2-3,10:12
Port Summary Monitor                               Thu Feb 14 14:19:50 2008
Port  Display          VLAN Name          Port  Link  Speed  Duplex
#    String             (or # VLANs)      State State Actual Actual
=====
1:2   2nd-Floor-Lab       Lab-Backbone      E    A    1000   FULL
1:3   Building2           Building2         E    A D
10:12 AllBackboneLANs (34)  E    R                FULL
=====
Port State: D-Disabled, E-Enabled
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback,
```




```
D-ELSM enabled but not up
U->page up D->page down ESC->exit
```

Restricted optics will show similarly to unsupported optics (!) in the `show ports` command by use of a '\$'. Additionally, when a 3rd party license has been installed, previously restricted optics are shown using a '%', because those optics are un-restricted by means of the license. Here is an example output:

```
show port 2:*          conf
Port Configuration Monitor                               Wed Sep  5  20:35:54
2012
Port   Virtual   Port Link Auto  Speed   Duplex  Flow Load
Media                                     Cfg Actual Cfg Actual Cntrl Master Pri
      router   State State Neg  Cfg Actual Cfg Actual Cntrl Master Pri
Red
=====
==
2:1    VR-Default  E    R   OFF 40000   FULL          $Q
+SR4
2:2    VR-Default  E    NP  OFF 10000   FULL
NONE
2:3    VR-Default  E    NP  OFF 10000   FULL
NONE
2:4    VR-Default  E    NP  OFF 10000   FULL
NONE
2:5    VR-Default  E    R   OFF 40000   FULL
NONE
2:6    VR-Default  E    NP  OFF 10000   FULL
NONE
2:7    VR-Default  E    NP  OFF 10000   FULL
NONE
2:8    VR-Default  E    NP  OFF 10000   FULL
NONE
2:9    VR-Default  E    R   OFF 40000   FULL          %Q
+LR4
2:10   VR-Default  E    NP  OFF 10000   FULL
NONE
2:11   VR-Default  E    NP  OFF 10000   FULL
NONE
2:12   VR-Default  E    NP  OFF 10000   FULL
NONE
2:13   VR-Default  E    R   OFF 40000   FULL          Q
+SR4
2:14   VR-Default  E    NP  OFF 10000   FULL
NONE
=====
== > indicates Port
      Display Name truncated past 8 characters   Link State:
      A-Active, R-Ready, NP-Port Not Present, L-Loopback   Port State:
      D-Disabled, E-Enabled, Media: !-Unsupported, $-Restrict., %-
Unrestrict.   Media Red: * -
      use "show port info detail" for redundant media type   0->Clear
Counters U->page up D->page down ESC->exit
```

```
show port 2:* conf no
Port Configuration
```



```

Port   Virtual   Port Link Auto   Speed   Duplex   Flow   Load
Media  router     State State Neg   Cfg Actual Cfg Actual Cntrl Master Pri
Red
=====
==
2:1    VR-Default E     R    OFF 40000   FULL           $Q
+SR4
2:2    VR-Default E     NP   OFF 10000   FULL
NONE
2:3    VR-Default E     NP   OFF 10000   FULL
NONE
2:4    VR-Default E     NP   OFF 10000   FULL
NONE
2:5    VR-Default E     R    OFF 40000   FULL
NONE
2:6    VR-Default E     NP   OFF 10000   FULL
NONE
2:7    VR-Default E     NP   OFF 10000   FULL
NONE
2:8    VR-Default E     NP   OFF 10000   FULL
NONE
2:9    VR-Default E     R    OFF 40000   FULL           %Q
+LR4
2:10   VR-Default E     NP   OFF 10000   FULL
NONE
2:11   VR-Default E     NP   OFF 10000   FULL
NONE
2:12   VR-Default E     NP   OFF 10000   FULL
NONE
2:13   VR-Default E     R    OFF 40000   FULL           Q
+SR4
2:14   VR-Default E     NP   OFF 10000   FULL
NONE
2:15   VR-Default E     NP   OFF 10000   FULL
NONE
2:16   VR-Default E     NP   OFF 10000   FULL
NONE
2:17   VR-Default E     R    OFF 40000   FULL
NONE
2:18   VR-Default E     NP   OFF 10000   FULL
NONE
2:19   VR-Default E     NP   OFF 10000   FULL
NONE
=====
==
>
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port not present, L-Loopback
Port State: D-Disabled, E-Enabled
Media: !-Unsupported Transceiver $-Restricted/%-Unrestricted
Transceiver
Media Red: * - use "show port info detail" for redundant media type

```

History

This command was first available in ExtremeXOS 12.1.



The tag value was added in ExtremeXOS 12.4.4.

Show output for 3rd party restricted optics support was added in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

show ports anomaly

```
show ports {port_list | tag tag} anomaly {no-refresh}
```

Description

Displays statistics of anomaly violation events in real time.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>tag</i>	Specifies an 802.1Q or 802.1ad tag value.
no-refresh	Specifies a static snapshot of data.

Default

N/A.

Usage Guidelines

If you do not specify a port number, range of ports, or tag value, statistics are displayed for all ports. To clear the counters, use the `clear counters ports` command. The default display is a constantly refreshing real-time display. If you specify the no-refresh parameter, the system displays a snapshot of the data at the time you issue the command.

This command takes effect after enabling anomaly-protection.

The tag value may be associated with either a VMAN or a VLAN.

Example

The following command displays real-time anomaly statistics on slot 2, all ports on a modular switch:

```
show ports 2:* anomaly
```



Following is sample output from this command:

```

Port Statistics Thu Nov  9 22:44:31 2006
Port Link      Rx Pkt ===== Anomaly Violation =====
State Count L3 Count      L4 Count      ICMP Count      Frag Count
=====
==
2:1 A 191585 1 2 0 0
2:2 R 0 0 0 0 0
2:3 R 0 0 0 0 0
2:4 R 0 0 0 0 0
2:5 R 0 0 0 0 0
2:6 R 0 0 0 0 0
2:7 R 0 0 0 0 0
2:8 R 0 0 0 0 0
2:9 R 0 0 0 0 0
2:10 R 0 0 0 0 0
2:11 R 0 0 0 0 0
2:12 A 178024 0 0 0 0
2:13 A 196956 0 0 0 0
2:14 R 0 0 0 0 0
2:15 R 0 0 0 0 0
2:16 R 0 0 0 0 0
2:17 R 0 0 0 0 0
=====
==
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
0->Clear Counters  U->page up  D->page down  ESC->exit

```

History

This command was first available in ExtremeXOS 12.0.

The tag value was added in ExtremeXOS 12.4.4.

Platform Availability

This command is available on all platforms.

show ports buffer

```
show ports port_list buffer
```

Description

Displays the packet buffer organization for the specified ports.



Syntax Description

<i>port_list</i>	Optionally specifies the list of ports, or slots and ports, for which packet buffer information is displayed. If the <i>port_list</i> is omitted then packet buffer information is displayed for all ports in the system.
------------------	---

Default

N/A.

Usage Guidelines

This command shows the packet buffer organization for the specified ports.

The *port_list* can span multiple ranges. The packet buffer description for each such port range is displayed.

Since ports and packet buffer are grouped by the hardware, the command displays the range of ports that share the same packet buffer.

The Total Packet Buffer Size for the port range is displayed in bytes, along with an indication of whether or not the user has configured over-commitment of the packet buffer (not overcommitted by default).

The amount of Reserved Buffer allocated to each port and QoS Profile is shown for the ports in the user-specified *port_list*. To configure the reserved buffer, use the `configure qosprofile qosprofile maxbuffer percentage ports port_list` command.

The Total Shared Buffer Size displayed is the Total Packet Buffer Size minus the total Reserved Buffer allocated to all ports and QoS profiles in the port range. Note that some packet buffer is also reserved to internal ports.

For each port, the maximum of the Total Shared Buffer Size that the port is allowed to use (Max Shared Buffer Usage) is shown both as an absolute number of bytes and as a percentage of the Total Shared Buffer Size. A port's Max Shared Buffer Usage may be configured using the command `configure ports {port_list} shared-packet-bufferpercentage`

Note the configured percentage may be different than the displayed percentage. This is because more recent hardware can only allocate shared packet buffer in steps, while older hardware can precisely allocate the requested percentage.

The more recent hardware dynamically adjusts each port's shared buffer usage limit based on simultaneous usage by multiple ports and QoS profiles, automatically providing fair usage of the shared buffer among the ports and QoS profiles that are currently demanding buffer space. This allows larger packet buffer usage bursts on a port when other ports are not using shared buffer. This dynamic adjustment cannot be observed with this command since only the maximum possible limits are displayed.

The VLAN name is displayed only if that port contains a single VLAN. If the port contains more than one VLAN, then the number of VLANs is displayed.

The tag value may be associated with either a VMAN or a VLAN.



Example

The following command displays on slot 2-3 on port 1 and slot 12 on port 10 on a modular switch:

```
BD-X8.34 # show ports 1:1,2:1-2 buffer
Packet Buffer Allocation for ports in range 1:1-12,1:25-36
Total Packet Buffer Size: 9584640 bytes, Not Overcommitted
Total Shared Buffer Size: 9051328
  Port 1:1  Max Shared Buffer Usage: 8055632 bytes (89%)
    QP1: Reserved Buffer: 3328 bytes
    QP8: Reserved Buffer: 1664 bytes
    MCQ: Reserved Buffer: 1664 bytes
Packet Buffer Allocation for ports in range 2:1-12,2:25-36
Total Packet Buffer Size: 9584640 bytes, Not Overcommitted
Total Shared Buffer Size: 9051328
  Port 2:1  Max Shared Buffer Usage: 1810224 bytes (20%)
    QP1: Reserved Buffer: 3328 bytes
    QP8: Reserved Buffer: 1664 bytes
    MCQ: Reserved Buffer: 1664 bytes
  Port 2:2  Max Shared Buffer Usage: 1810224 bytes (20%)
    QP1: Reserved Buffer: 3328 bytes
    QP8: Reserved Buffer: 1664 bytes
    MCQ: Reserved Buffer: 1664 bytes
BD-X8.35 #
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

show ports collisions

```
show ports {mgmt | port_list | tag tag} collisions {no-refresh}
```

Description

Displays real-time collision statistics.

Syntax Description

mgmt	Specifies the management port.
<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>tag</i>	Specifies an 802.1Q or 802.1ad tag value.
no-refresh	Specifies a static snapshot of data.



Default

Real-time statistics.

Usage Guidelines

If you do not specify a port number, range of ports, or tag value, statistics are displayed for all ports. To clear the counters, use the `clear counters ports` command. The default display is a constantly refreshing real-time display. If you specify the no-refresh parameter, the system displays a snapshot of the data at the time you issue the command.

The tag value may be associated with either a VMAN or a VLAN.

This status information may be useful for your technical support representative if you have a network problem.

Example

The following command displays real-time collision statistics on slot 1, ports 1 and 2 on a modular switch:

```
show ports 1:1-2 collisions
```

Following is sample output from this command:

```
Port Collision Monitor
Port      Link      Collision Histogram
State 1   2   3   4   5   6   7   8   9  10  11  12  13  14  15  16
=====
==
1:1      A      0   0   0   0   0   0   0   0   0   0   0   0   0   0   0
1:2      R      0   0   0   0   0   0   0   0   0   0   0   0   0   0   0
=====
==
Link State: A-Active R-Ready, NP-Port not present, L-Loopback
```

The numbers 1 to 16 represent the number of collisions encountered prior to successfully transmitting the packet; this is applicable only for half-duplex links.

History

This command was first available in ExtremeXOS 10.1.

The no-refresh variable was added in ExtremeXOS 11.3.

The tag value was added in ExtremeXOS 12.4.4.

Platform Availability

This command is available on all platforms.



show ports configuration

```
show ports {mgmt | port_list | tag tag} configuration {no-refresh}
```

Description

Displays port configuration statistics, in real time or snapshot.

Syntax Description

mgmt	Specifies the management port.
<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>tag</i>	Specifies an 802.1Q or 802.1ad tag value.
no-refresh	Specifies a static snapshot of data.

Default

Real-time statistics.

Usage Guidelines

If you do not specify a port number, range of ports, or tag value, configuration statistics are displayed for all ports. If you specify the no-refresh parameter, the system displays a snapshot of the data at the time you issue the command.

The tag value may be associated with either a VMAN or a VLAN.

This status information may be useful for your technical support representative if you have a network problem.

This command displays port configuration, which includes:

- Virtual router.
- Port state.
- Link state.
- Autonegotiation information.
- Link speed.
- Duplex mode.
- Flow control.



- Load sharing information.
- Link media information.

Note



On 10 Gbps ports, the Media Primary column displays NONE when no module is installed, and SR, LR, or ER depending on the module installed when there is one present. Combination ports on Summit X460, X480, and X670 series switches display Autonegotiation, Link speed, and Duplex mode information only for the current primary medium.

Example

The following command displays the port configuration for all ports on a Summit family switch:

```
show ports configuration
Port Configuration Monitor                               Fri Apr 13 10:22:29
2007
Port      Virtual      Port  Link  Auto   Speed      Duplex   Flow  Load  Media
router    State State Neg   Cfg  Actual Cfg Actual Cntrl Master Pri Red
=====
==
1         VR-Default  E     R     ON    AUTO      AUTO                    NONE
UTP
2         VR-Default  E     R     ON    AUTO      AUTO                    NONE
UTP
3         VR-Default  E     R     ON    AUTO      AUTO                    NONE
UTP
4         VR-Default  E     R     ON    AUTO      AUTO                    NONE
UTP
5         VR-Default  E     R     ON    AUTO      AUTO                    NONE
6         VR-Default  E     R     ON    AUTO      AUTO                    NONE
7         VR-Default  E     R     OFF   100      FULL                    SX
8         VR-Default  E     R     ON    AUTO      AUTO                    NONE
9         VR-Default  E     R     ON    AUTO      AUTO                    NONE
10        VR-Default  E     R     ON    AUTO      AUTO                    NONE
11        VR-Default  E     R     ON    AUTO      AUTO                    NONE
12        VR-Default  E     R     ON    AUTO      AUTO                    NONE
13        VR-Default  E     R     ON    AUTO      AUTO                    NONE
14        VR-Default  E     R     ON    AUTO      AUTO                    NONE
15        VR-Default  E     R     ON    AUTO      AUTO                    NONE
=====
==
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
Port State: D-Disabled, E-Enabled, Media: !-Unsupported Optic Module
Media Red: * - use "show port info detail" for redundant media type
0->Clear Counters  U->page up  D->page down  ESC->exit
```

The following command displays the port configuration for slot 2, port 2 on a modular switch:

```
show ports 2:2 configuration
```



Following sample output includes flags that indicate the summary status of a non-Extreme optical device. "Licensed" means the optics license is installed (the 40G+ device is unrestricted, but still unsupported). "Unsupported" means the non-Extreme device is unsupported by Extreme, but is unrestricted because it does not require a license (1G/10G). "Restricted" means the non-Extreme 40G+ device will be restricted because no optics license is installed.

```
show ports 3-4 configuration no-refresh
Port Configuration
Port   Virtual   Port   Link   Auto   Speed   Duplex   Flow
Load   Media
      router   State  State  Neg   Cfg Actual  Cfg Actual  Cntrl
Master Pri Red
=====
=====
3     VR-Default  E      R     OFF   10000
FULL
4     VR-Default  E      R     OFF   10000
FULL
=====
=====
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
Port State: D-Disabled, E-Enabled, Media: %-Licensed,!-Unsupported, $-
Restricted
Media Red: * - use "show port info detail" for redundant media type
0->Clear Counters U->page up D->page down ESC->exit
```

History

This command was first available in ExtremeXOS 10.1.

The Port not present and Media variables were added in ExtremeXOS 11.2.

The no-refresh variable was added in ExtremeXOS 11.3.

The tag value was added in ExtremeXOS 12.4.4.

This command was updated to include flags that indicate the summary status of a non-Extreme optical device in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

show port information

```
show port {mgmt | port_list | tag tag} information {detail}
```

Description

Displays detailed system-related information.



Syntax Description

mgmt	Specifies the management port.
<i>port_list</i>	Specifies one or more ports of slots and ports.
<i>tag</i>	Specifies an 802.1Q or 802.1ad tag value.
detail	Specifies detailed port information.

Default

N/A.

Usage Guidelines

This command displays information, including the following:

- Port number.
- Port configuration.
 - Virtual router.
 - Type of port.
 - Admin state.
 - Link state and speed.
 - Link Up/Down Transition.
 - VLAN configuration.
 - STP configuration.
 - Trunking, or load sharing.
 - EDP.
 - ELSM (disabled; or if enabled, the ELSM link state is shown as well).
 - Load balancing.
 - Learning.
 - Egress flooding.
 - Jumbo frames.
 - Link port up/down traps.
 - Port isolation status.
 - QoS profiles.
 - VMAN status.
 - Smart Redundancy status.
 - SRP status.
 - Additional platform-specific information.

If you do not specify a port number, range of ports, or tag value, detailed system-related information is displayed for all ports. The data is displayed in a table format.

The tag value may be associated with either a VMAN or a VLAN.

This status information may be useful for your technical support representative if you have a network problem.



The detail parameter is used to provide more specific port information. The data is called out with written explanations versus displayed in a table format.



Note

The keyword detail displays slightly different information depending on the platform and configuration you are working with.

The link filter counter displayed with the detail keyword is calculated at the middle layer on receiving an event. The link filter up indicates the number of link transitions from down to up at the middle layer filter.

Example

The following command displays port system-related information on a BlackDiamond 8810 switch:

```
show port 1:5-7 information
Port      Flags          Link          Link Num Num  Num  Jumbo QOS
Load
State     ELSM UPS   STP VLAN Proto Size  profile Master
=====
=====
1:5       Em-----e--fMB---b active    -    1    1    1    1    9216 none
1:6       Em-----e--fMB---x ready      -    0    1    1    1    9216 none
1:7       Em-----e--fMB---p ready      -    0    1    1    1    9216 none>
indicates Port Display Name truncated past 8 characters
Flags : a - Load Sharing Algorithm address-based, D - Port Disabled,
e - Extreme Discovery Protocol Enabled, E - Port Enabled,
g - Egress TOS Enabled, i - Isolation
j - Jumbo Frame Enabled, l - Load Sharing Enabled,
m - MACLearning Enabled, n - Ingress TOS Enabled,
o - Dot1p Replacement Enabled, P - Software redundant port(Primary),
R - Software redundant port(Redundant),
q - Background QOS Monitoring Enabled,
s - diffserv Replacement Enabled,
v - Vman Enabled, f - Unicast Flooding Enabled,
M - Multicast Flooding Enabled, B - Broadcast Flooding Enabled
L - Extreme Link Status Monitoring Enabled
O - Ethernet OAM Enabled
w - MACLearning Disabled with Forwarding
b - Rx and Tx Flow Control Enabled, x - Rx Flow Control Enabled
p - Priority Flow Control Enabled
```

The following command displays detailed information for port 1 on a Summit X670 switch:

```
X670V-48x.2 # show port 1 information detail
Port: 1
Virtual-router: VR-Default
Type: SF+_SR Unsupported Optic Module
Random Early drop: Unsupported
Admin state: Enabled with 10G full-duplex
Link State: Active, 10Gbps, full-duplex
Link Ups: 1 Last: Wed Feb 06 12:28:48 2013
Link Downs: 0 Last: --

VLAN cfg:
```



```

STP cfg:

Protocol:
Trunking:          Load sharing is not enabled.

EDP:              Enabled
ELSM:             Disabled
Ethernet OAM:     Disabled
Learning:         Enabled
Unicast Flooding: Enabled
Multicast Flooding: Enabled
Broadcast Flooding: Enabled
Jumbo:           Disabled

Flow Control:    Rx-Pause: Enabled          Tx-Pause:      Disabled

Priority Flow Control: Disabled
Reflective Relay: Disabled
Link up/down SNMP trap filter setting: Enabled
Egress Port Rate: No-limit
Broadcast Rate:  No-limit
Multicast Rate:  No-limit
Unknown Dest Mac Rate: No-limit
QoS Profile:     None configured
Ingress Rate Shaping :      Unsupported
Ingress IPTOS Examination: Disabled
Ingress 802.1p Examination: Enabled
Ingress 802.1p Inner Exam: Disabled
Egress IPTOS Replacement: Disabled
Egress 802.1p Replacement: Disabled
NetLogin:        Disabled
NetLogin port mode: Port based VLANs
Smart redundancy: Enabled
Software redundant port: Disabled
IPFIX:           Disabled          Metering:      Ingress, All Packets, All

Traffic
          IPv4 Flow Key Mask:      SIP: 255.255.255.255          DIP:
255.255.255.255
          IPv6 Flow Key Mask:      SIP:
ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff
          DIP:
ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff
          Far-End-Fault-Indication: Disabled
          Shared packet buffer:      default
          VMAN CEP egress filtering: Disabled
          Isolation:                 Off
          PTP Configured:            Disabled
          Time-Stamping Mode:        None
          Synchronous Ethernet:      Unsupported
          Dynamic VLAN Uplink:        Disabled
          VM Tracking Dynamic VLANs:  Disabled
X670V-48x.3 #
    
```

The following example output displays the current isolation mode of the ports:

```

show port 21 info detail
Port: 21
    
```



```

Virtual-router: VR-Default
Type:          SF+_SR Unsupported Optic Module
Random Early drop:  Unsupported
Admin state:   Enabled with 10G full-duplex
Link State:   Active, 10Gbps, full-duplex
Link Ups:     1          Last: Wed Feb 06 12:28:48 2013
Link Downs:   0          Last: --

VLAN cfg:
STP cfg:

Protocol:
Trunking:     Load sharing is not enabled.

EDP:          Enabled
ELSM:         Disabled
Ethernet OAM: Disabled
Learning:     Enabled
Unicast Flooding: Enabled
Multicast Flooding: Enabled
Broadcast Flooding: Enabled
Jumbo:        Disabled

Flow Control:  Rx-Pause: Enabled          Tx-Pause: Disabled

Priority Flow Control: Disabled
Reflective Relay: Disabled
Link up/down SNMP trap filter setting: Enabled
Egress Port Rate: No-limit
Broadcast Rate:  No-limit
Multicast Rate:  No-limit
Unknown Dest Mac Rate: No-limit
QoS Profile:    None configured
Ingress Rate Shaping :  Unsupported
Ingress IPTOS Examination: Disabled
Ingress 802.1p Examination: Enabled
Ingress 802.1p Inner Exam: Disabled
Egress IPTOS Replacement: Disabled
Egress 802.1p Replacement: Disabled
NetLogin:      Disabled
NetLogin port mode: Port based VLANs
Smart redundancy: Enabled
Software redundant port: Disabled
IPFIX:         Disabled          Metering: Ingress, All Packets, All
Traffic
          IPv4 Flow Key Mask: SIP: 255.255.255.255          DIP:
255.255.255.255
          IPv6 Flow Key Mask: SIP:
ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff
          DIP:
ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff
Far-End-Fault-Indication: Disabled
Shared packet buffer: default
VMAN CEP egress filtering: Disabled
Isolation: Off
PTP Configured: Disabled
Time-Stamping Mode: None

```



```
Synchronous Ethernet:      Unsupported
Dynamic VLAN Uplink:       Disabled
VM Tracking Dynamic VLANs:  Disabled
```

```
show port 21-23 info
Port      Flags          Link          ELSM Link Num Num  Num  Jumbo
QOS      Load
State     /OAM UPS   STP VLAN  Proto Size
profile Master
=====
=====
21      Emj-----e--fMB---xi active   - / - 1    0
      1 0 1600 none 22
      Emj-----e--fMB---xi active   - / - 1    0
      1 0 1600 none 23
      Emj-----e--fMB---x- ready    - / - 0    0
      0 0 1600 none
=====
=====
```

```
> indicates Port Display Name truncated
past 8 characters
```

```
Flags : a - Load Sharing Algorithm address-based, D - Port Disabled,
        e - Extreme Discovery Protocol Enabled, E - Port Enabled,
        g - Egress TOS Enabled, i - Isolation, j - Jumbo Frame Enabled,
        l - Load Sharing Enabled, m - MACLearning Enabled,
        n - Ingress TOS Enabled,
        o - Dot1p Replacement Enabled,
        P - Software redundant port(Primary),
        R - Software redundant port(Redundant),
        q - Background QOS Monitoring Enabled,
        s - diffserv Replacement Enabled,
        v - Vman Enabled, f - Unicast Flooding Enabled,
        M - Multicast Flooding Enabled, B - Broadcast Flooding Enabled

        L - Extreme Link Status Monitoring Enabled
        O - Ethernet OAM Enabled
        w - MACLearning Disabled with Forwarding
        b - Rx and Tx Flow Control Enabled, x - Rx Flow Control Enabled

        p - Priority Flow Control Enabled
```

The following example displays the current status of non-Extreme optical devices:

```
# sh ports 3 info detail
Port: 3
      Virtual-router: VR-Default
      Type:           Q+SR4      (Licensed)
or
                                   (Unsupported)
or
                                   (Restricted, 10 Gbps
maximum)
      Random Early drop:      Unsupported
      Admin state:           Enabled
```



```

Link State:      Ready
Link Ups:        0          Last:  --
Link Downs:     0          Last:  --
VLAN cfg:
    Name: Default, Internal Tag = 1, MAC-limit = No-limit,
Virtual router: VR-Default
STP cfg:
    s0(disable), Tag=(none), Mode=802.1D,
State=FORWARDING
Protocol:
    Name: Default          Protocol: ANY          Match all
protocols.
Trunking:        Load sharing is not enabled.
EDP:             Enabled
ELSM:           Disabled
Ethernet OAM:    Disabled
Learning:        Enabled
Unicast Flooding: Enabled
Multicast Flooding: Enabled
Broadcast Flooding: Enabled
Jumbo:          Disabled

Flow Control:    Rx-Pause: Enabled          Tx-Pause:
Disabled
Priority Flow Control: Disabled
Reflective Relay: Disabled
Link up/down SNMP trap filter setting: Enabled
Egress Port Rate: No-limit (Restricted, 10 Gbps
maximum)
Broadcast Rate: No-limit
Multicast Rate: No-limit
Unknown Dest Mac Rate: No-limit
QoS Profile:     None configured
Ingress Rate Shaping : Unsupported
Ingress IPTOS Examination: Disabled
    
```

The following output displays show port info detail that has port-specific tag.

```

Port: 3
    Virtual-router: None
    Type:           BASET
    Random Early drop: Unsupported
    Admin state:    Enabled with 10G full-duplex
    Link State:     Active, 1Gbps, full-duplex
    Link Ups:       0          Last: Wed Jun 05 00:29:24 2013
    Link Downs:    0          Last: --

VLAN cfg:          Name: test, 802.1Q Tag = 200, MAC-limit = No-
limit,
Virtual router:    VR-Test
    Port-specific VLAN ID: 10, 11, 200
    Name: test2, 802.1Q Tag = 300, MAC-limit = No-limit,
Virtual router:    VR-Default
STP cfg:
Protocol:
Trunking:          Load sharing is not enabled.

EDP:               Disabled
ELSM:             Disabled
    
```




```

Ethernet OAM:           Disabled
Learning:               Enabled
Unicast Flooding:      Enabled
Multicast Flooding:    Enabled
Broadcast Flooding:    Enabled
Jumbo:                  Enabled, MTU= 9216

Flow Control:   Rx-Pause: Enabled           Tx-Pause: Disabled
Priority Flow Control: Disabled
Reflective Relay: Disabled
Link up/down SNMP trap filter setting: Enabled
Egress Port Rate: No-limit
Broadcast Rate: No-limit
Multicast Rate: No-limit
Unknown Dest Mac Rate: No-limit
QoS Profile: None configured
Ingress Rate Shaping : Unsupported
Ingress IPTOS Examination: Disabled
Ingress 802.1p Examination: Enabled
Ingress 802.1p Inner Exam: Disabled
Egress IPTOS Replacement: Disabled
Egress 802.1p Replacement: Disabled
NetLogin: Disabled
NetLogin port mode: Port based VLANs
Smart redundancy: Enabled
Software redundant port: Disabled
IPFIX: Disabled           Metering: Ingress, All Packets,
All Traffic
IPv4 Flow Key Mask: SIP: 255.255.255.255           DIP:
255.255.255.255
IPv6 Flow Key Mask: SIP:
ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff
DIP: ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff

Far-End-Fault-Indication: Disabled
Shared packet buffer: default
VMAN CEP egress filtering: Disabled
Isolation: Off
PTP Configured: Disabled
Time-Stamping Mode: None
Synchronous Ethernet: Unsupported
Dynamic VLAN Uplink: Disabled
VM Tracking Dynamic VLANs: Disabled
    
```

History

This command was first available in ExtremeXOS 10.1.

Information on ingress rate shaping was added in ExtremeXOS 11.0.

Network Login, Smart Redundancy, and rate limiting were added in ExtremeXOS 11.1.

Information on unicast, multicast, and broadcast flooding; the Port not present parameter; and autopolarity status were added in ExtremeXOS 11.2.



The netlogin parameters were added in ExtremeXOS 11.3.

The output command was modified in ExtremeXOS 12.1 so that when learning is disabled with the disabled learning port command, a new w flag appears in the output.

Link Ups and Link Downs information was added to the output and the tag value was added to the command syntax in ExtremeXOS 12.4.4.

Industrial Temperature detail added to Type field in ExtremeXOS 15.1.2.



Note

The Industrial Temperature detail is only displayed only when detail is used. Without it, the output is compressed and the optic type is not displayed.

The show output was enhanced to display the SNMP ifMib ifAlias accessible string size information.

The show output was updated to include the current status of non-Extreme optical devices in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

Information on WAN PHY ports is available only on Summit X480 series switches, whether or not included in a SummitStack.

show ports ip-fix

```
show ports {port_list | tag tag} ip-fix {detail | no-refresh}
```

Description

Displays IPFIX information on specified port(s).

Syntax Description

<i>port_list</i>	Specifies the port(s).
<i>tag</i>	Specifies an 802.1Q or 802.1ad tag value.
detail	Specifies display in detailed format.
no-refresh	Specifies a static page by page display of the data without auto-refresh.

Default

N/A.



Usage Guidelines

Use this command to display IPFIX information regarding the global state, collector information, and which ports are enabled.

The tag value may be associated with either a VMAN or a VLAN.

Example

The following command displays IPFIX information for port 3:1:

```
show port 3:1 ip-fix no-refresh
```

Following is sample output on a BlackDiamond 8800 switch:

```
BD-8810.1 # show port 3:1 ip-fix no
Port IPFIX Metering
Port      Link  IPFIX          Byte      Pkt Flow Record  Premature
State    State          Count      Count      Count      Exports
=====
==
3:1      A    E              0          0          0          0
=====
==
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present L-Loopback
IPFIX State: D-Disabled E-Enabled
IPFIX Enabled on Port(s):    3:1
```

Flow Record Count—how many flow records have been sent from this port.

Premature Exports—the result of a hash bucket being full and the hardware having to export a flow prematurely to make room for a new flow.

History

This command was first available in ExtremeXOS 12.5.

The tag value was added in ExtremeXOS 12.4.4.

Platform Availability

This command is available on BlackDiamond 8900 G96Tc, G48T-xl, G48X-xl, and 10G8X-xl modules and Summit X460 and X480 switches.



show ports link-scan

```
show ports link-scan { slot [ <slot> | all ] }
```



Description

Use this command to display the link scan attributes for polling port status.

Syntax Description

ports	Ports.
link-scan	Configure link scan attributes for polling port status.
slot	Slot number (default all slots).
<i>s</i> lot	Slot number"; type="dmIOSlotNoNames
all	All slots.

Default

All slots.

Usage Guidelines

This command displays the link scan attributes for polling port status.

Example

```
# sh ports link-scan
Slot   Interval (ms)
-----
 1      50 (default)
 2     300
 3      50 (default)
 4      50 (default)
 5
 6
 7
 8     200
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

show ports packet

```
show ports {mgmt | port_list | tag tag} packet {no-refresh}
```



Complete

Displays a snapshot or real-time histogram of packet statistics.

Syntax Description

mgmt	Specifies the management port.
<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>tag</i>	Specifies an 802.1Q or 802.1ad tag value.
no-refresh	Specifies a static snapshot of data.

Default

Real-time statistics.

Usage Guidelines

If you do not specify a port number, range of ports, or tag value, the system displays information for all ports; if you specify the `no-refresh` parameter, the system displays a snapshot of the data at the time you issue the command. To clear the counters, use the `clear counters ports` command.

The tag value may be associated with either a VMAN or a VLAN.

This status information may be useful for your technical support representative if you have a network problem.

The following packet statistics are displayed:

- Port number.
- Link state.
- Packet size.

Example

The following command displays packet statistics for slot 1, port 1, slot 2, port 1, and slot 5, ports 1 through 8 on a modular switch:

```
show ports 1:1, 2:1, 5:1-5:8 packet
```

Following is sample output from this command:

```

Port      Link
State    0-64   65-127  128-255  256-511  512-1023  1024-1518  Jumbo
=====
==
1:1  A      0       0         0         0         0         0         0
2:1  R      0       0         0         0         0         0         0
5:1  R      0       0         0         0         0         0         0
5:2  R      0       0         0         0         0         0         0

```



```

5:3 R      0      0      0      0      0      0      0      0
5:4 R      0      0      0      0      0      0      0      0
5:5 R      0      0      0      0      0      0      0      0
5:6 R      0      0      0      0      0      0      0      0
5:7 R      0      0      0      0      0      0      0      0
5:8 R      0      0      0      0      0      0      0      0
=====
==
Link State: A-Active, R-Ready, NP-Port not present, L-Loopback

```

History

This command was first available in ExtremeXOS 10.1.

The Port Not Present variable was added in ExtremeXOS 11.2.

The no-refresh variable was added in ExtremeXOS 11.3.

The tag value was added in ExtremeXOS 12.4.4.

Platform Availability

This command is available on all platforms.

show ports redundant

show ports redundant

Description

Displays detailed information about redundant ports.

Syntax

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.



Example

The following command displays information on software-controlled redundant ports on the switch:

```
show ports redundant
```

Following is sample output from this command:

```
Primary: *1:1          Redundant: 3:1, Link on/off option: OFF  
Flags: (*)Active, (!) Disabled, (g) Load Share Group
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

show ports sharing

```
show ports sharing
```

Description

Displays port load-sharing groups, or link aggregation groups (LAGs).

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Output from this command displays the following information:

- **Config Master**—The port that is configured as the master logical port of the link aggregation group (LAG). This number is also the LAG group ID.
- **Current Master**—In LACP, this is the port that is currently the LAG group ID, or master logical port for the LAG.
- **Agg Control**—This is the aggregation control for the specified LAG; it can be either static, LACP or health-check. In LACP, it is the aggregation control for the specified LAG.



- Ld Share Algorithm—The algorithm used for the link aggregation. The available link aggregation algorithms vary among platforms; see the *ExtremeXOS Concepts Guide* for more information.
- Ld Share Group—The specific ports that belong to each LAG, or the port numbers in the trunk. A port can belong to only one LAG, either static or dynamic.
- Agg Mbr—In LACP, this shows whether the port has been added to the aggregator or not; it will be either Y for yes or - for no.
- Link State—This is the current status of the link.
- Link Up transitions—Number of times the link has cycled through being up, then down, then up.

Example

Following is sample outlook displaying link aggregation on a Summit series switch:

```
show ports sharing
Load Sharing Monitor
Config      Current      Agg          Ld Share     Ld Share     Agg      Link      Link Up
Master      Master      Control      Algorithm    Group        Mbr      State     transitions
=====
      1 1hlth-chk L2          1            Y            A         1
L2          1            -            A            1
L2          1            Y            A            1
5           5            LACP        L2           5            Y         A         3
L2          6            Y            A            3
L2          7            Y            A            3
L2          8            Y            A            3
L2          9            -            A            3
L2         10            -            A            3
12          Static      L2          12           -            R         0
L2         13            -            R            0
=====
Link State: A-Active, D-Disabled, R-Ready, NP-Port not present, L-Loopback
Load Sharing Algorithm: (L2) Layer 2 address based, (L3) Layer 3 address
based,
(L3_L4) Layer 3 address and Layer4 port based
Number of load sharing trunks: 3
```

Following is sample outlook for a BlackDiamond 8800 switch that uses a custom load sharing algorithm:

```
BD-8810.8 # show port sharing
Load Sharing Monitor
Config      Current      Agg          Ld Share     Ld Share     Agg      Link      Link Up
Master      Master      Control      Algorithm    Group        Mbr      State     Transitions
=====
2:1        2:1          Static      L2           2:1          Y         A         1
L2         2:2          Y            A            1
3:1        3:1          Static      L3_L4        3:1          Y         A         1
L3_L4     3:2          Y            A            1
4:1        4:1          Static      custom       4:1          Y         A         1
custom     4:2          Y            A            1
=====
Link State: A-Active, D-Disabled, R-Ready, NP-Port not present, L-Loopback
Load Sharing Algorithm: (L2) Layer 2 address based, (L3) Layer 3 address based
(L3_L4) Layer 3 address and Layer4 port based
```



(custom) User-selected address-based configuration
 Custom Algorithm Configuration: ipv4 source-only, xor

The 'custom' algorithm is not used for traffic ingressing on current slot 1, 2, 3, 5 and 10. Refer to XOS Command Reference.
 Number of load sharing trunks: 3

History

This command was first available in ExtremeXOS 10.1.

The LACP feature was added in ExtremeXOS 11.3.

The Health Check LAG was added in ExtremeXOS 12.1.3.

The round-robin algorithm was added in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

show ports tdm alarms

```
show ports {port_list} tdm alarms {no-refresh}
```

Description

Displays the TDM port alarms.

Syntax Description

tdm	Time Division Multiplexing.
alarms	Display the alarms for the tdm ports.
no-refresh	Page by page display without auto-refresh.

Default

N/A.

Usage Guidelines

Use this command to display the TDM port alarms.



Example

```
E4G-400.2 # sh ports {<port_list>} tdm alarms
TDM Alarms                                     Wed Apr 4 15:45:52 2011
Port      Tx RAI      Rx RAI      Tx AIS  Rx AIS  LOF      LOS
(Yellow)  (Yellow)  (Blue)   (Blue)  (Red)   Near End
=====
==
1          None       None      Alarm   None    Alarm   Alarm
2          None       None      None    None    None    None
=====
==
> indicates Port Display Name truncated past 8 characters
U->page up  D->page down ESC->exit
Legend: AIS - Alarm Indication Signal, LOF - Loss of Frame,
LOS - Loss of Signal, RAI - Remote Alarm Indication
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

show ports tdm configuration

```
show ports {port_list} tdm configuration {no-refresh}
```

Description

Displays the specified TDM port configuration.

Syntax Description

tdm	Time Division Multiplexing.
configuration	Displays TDM ports configuration.
no-refresh	Page by page display without auto-refresh.

Default

N/A.

Usage Guidelines

The configured cable length column will be displayed only if the hierarchy selected is T1/ANSI.



Example

```
E4G-400.2 # sh ports {<port_list>} tdm configuration
TDM Port Configuration Monitor                               Wed Apr 4 15:56:34 2011
Port      Flags          Port  Link  Line  Cable
State State Coding Len(ft)
=====
==
1          e-cCLt-----  E    R    HDB3
2          TS--PtR-----  D    R    B8ZS  550-660
3          TU-rPt-----  E    R    B8ZS  >660
4          eB-CLt-----  E    R    AMI
=====
==
> indicates Port Display Name truncated past 8 characters
Link State: (A) Active, (L) Loopback, (NP) Port not present, (R) Ready
Port State: (D) Disabled, (E) Enabled
Flags      : (B) E1 Basic Frame, (c) CRC-4, (C) Channel Associated Signaling,
(e) E1 Hierarchy, (F) Extended Super Framed,
(L) Line Recovered Clock, (M) Multi-framed,
(P) CES Pseudo-wire Recovered Clock,
(r) Robbed Bit signaling, (R) Clock Recovery Enabled,
(S) Super-framed, (t) TDM Circuit, (T) T1 Hierarchy,
(U) Unframed
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

show ports tdm errors

```
show ports {port_list} tdm errors {near-end} {total | intervals | current {no-
refresh}}
```

Description

Displays the specified TDM port error counters. The current switch lists the error statistics of the on-going 15-minute interval. The intervals switch displays the error statistics for the last 96 instances of 15-minute intervals. The total switch displays the cumulative error statistics.

Syntax Description

tdm	Time Division Multiplexing.
errors	Displays errors.
near-end	Displays near-end errors (Default).



total	Displays cumulative errors.
intervals	Displays errors for 15-minute intervals.
current	Displays current 15-minute interval errors (Default).
no-refresh	Page by page display without auto-refresh.

Default

The default will display near-end current error statistics.

Usage Guidelines

Use the current switch lists the error statistics of the on-going 15-minute interval. The intervals switch displays the error statistics for the last 96 instances of 15-minute intervals. The total switch displays the cumulative error statistics.

Example

Only one space is given between the first and second field to accommodate the output within 80 characters.

```
E4G-400.3 # show port {port_list} tdm errors near-end current
Port Statistics                                     Tue Jul 31 10:20:56 2012
Port          LCV/          PCV/  Unavail      Error  Burst      Severe
Slip
BPV          CRC   Secs          Secs   Err Secs  Err Secs      Secs
=====
==
31           0           0           0           0           0
0           0
32           0           0           0           0           0
0           0
=====
==
> indicates Port Display Name truncated past 8 characters
U->page up  D->page down  ESC->exit
Legend: BPV - Bipolar Violations, LCV - Line Code Violations,
PCV - Path Code Violations
E4G-400.3 # show port {port_list} tdm errors near-end current no-refresh
Port          LCV/          PCV/  Unavail      Error  Burst      Severe
Slip
BPV          CRC   Secs          Secs   Err Secs  Err Secs      Secs
=====
==
31           0           0           0           0           0
0           0
32           0           0           0           0           0
0           0
=====
==
> indicates Port Display Name truncated past 8 characters
Legend: BPV - Bipolar Violations, LCV - Line Code Violations,
```



```

PCV - Path Code Violations
E4G-400.3 # show port {<port_list>} tdm errors near-end intervals
Port/          LCV/          PCV/  Unavail    Error  Burst    Severe
Slip
Interval      BPV           CRC   Secs       Secs   Err Secs  Err Secs
Secs
=====
==
31/
1              0             0     0          89     0         0         0
2              0             8     0          0      0         0         12
3             456           0     0          0      0         0         0
96             0             0     0          0      0         0         0
0
-----
--
32/
1              0             0     0          0      0         0         0
96             0             0     0          0      0         0         0
0
=====
==
> indicates Port Display Name truncated past 8 characters
Legend: BPV - Bipolar Violations, LCV - Line Code Violations,
PCV - Path Code Violations
E4G-400.3 # show port {port_list} tdm errors near-end total
Port: 31
Line Code Violations / Bipolar Violations : 837483
Path Code Violations / CRC                : 289037434
Unavailable Seconds                       : 38389
Errored Seconds                           : 284
Bursty Errored Seconds                   : 297
Severely Errored Seconds                  : 82
Controlled Slip Seconds                   : 28
    
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

show ports tdm information

```
show ports {port_list} tdm information {detail}
```

Description

Displays the specified TDM port information.



Syntax Description

tdm	Time Division Multiplexing.
information	Displays TDM ports information.
detail	Display detailed TDM ports information.

Default

N/A.

Usage Guidelines

Use this command to display general or detailed information on a specified list of TDM ports.

Example

```
E4G-400.3 # sh port {port_list} tdm information
Port      Flags          Link   Line   Clock   #
State    Coding   Source   PW
=====
==
31      EeMcCRt----- ready  HDB3   Line    1
32      ETS-C-t----- ready  B8ZS   CCR     2
33      ETF-C-ts----- active B8ZS   CCR     3
34      DTU-r-tl----- ready  B8ZS   CCR     1
35      EeU---t----- ready  AMI    CCR     1
36      EeB-CRt----- ready  HDB3   CCR     1
=====
==
> indicates Port Display Name truncated past 8 characters
Flags : (B) E1 Basic Frame, (c) CRC-4, (C) Channel Associated Signaling,
(D) Port Disabled, (e) E1 Hierarchy, (E) Port Enabled,
(F) Extended Super Framed, (l) T1 Long-Haul Cable Length,
(M) Multi-framed, (r) Robbed Bit signaling,
(R) Clock Recovery Enabled, (s) T1 Short-Haul Cable Length,
(S) Super-framed, (t) TDM Circuit, (T) T1 Hierarchy,
(U) Unframed

Clock Source:
(CCR) CES pseudo-wire Recovered Clock
E4G-400.5 # show port {port_list} tdm information detail
Port: 1
Admin state          : Enabled
Link State           : Ready
Link Counter         : Up      0 time(s)
Framing              : E1 Multi-frame, CRC-4 Enabled
Signaling            : Channel Associated Signaling
Line Coding          : HDB3
Clock Source         : Line
Clock Recovery       : Disabled
Recovered Clock Quality : QL_DNU
Transmit Clock Quality : QL_DNU
Protocol             : TDM Circuit
```



```

Loopback                : Disabled
Idle Code                : 10
TDM Service Configuration:
Name : cesopBundle1
Time slots : 1-10, 11-20
Name : cesopBundle2
Time slots : 21-30
Port: 2
Admin state              : Enabled
Link State               : Ready
Link Counter             : Up          0 time(s)
Framing                  : E1 Unframed
Line Coding               : HDB3
Clock Source              : CES pseudo-wire recovered (satop21)
Clock Recovery           : Enabled
Recovered Clock Quality  : QL_PRC
Transmit Clock Quality   : QL_DNU
Protocol                  : TDM Circuit
Loopback                 : Enabled (Network Payload)
Idle Code                : 10
TDM Service Configuration:
Name : satopBundle1
Port: 3
Admin state              : Enabled
Link State               : Ready
Link Counter             : Up          0 time(s)
Framing                  : T1 Unframed
Line Coding               : B8ZS
Cable Length             : Short-Haul, 550-660 feet
Clock Source              : CES pseudo-wire recovered (satop22)
Clock Recovery           : Disabled
Recovered Clock Quality  : QL_DNU
Transmit Clock Quality   : QL_DNU
Protocol                  : TDM Circuit
Loopback                 : Enabled (Network Payload)
Idle Code                : 10
TDM Service Configuration:
Name : satopBundle2
Port: 4
Admin state              : Enabled
Link State               : Ready
Link Counter             : Up          0 time(s)
Framing                  : T1 Unframed
Line Coding               : B8ZS
Cable Length             : Long-Haul, > 660 feet
Receiver Gain            : 22.5dB
Clock Source              : CES pseudo-wire recovered (satop23)
Clock Recovery           : Enabled
Recovered Clock Quality  : QL_PRS
Transmit Clock Quality   : QL_PRS
Protocol                  : TDM Circuit
Loopback                 : Enabled (Local)
Idle Code                : 10
TDM Service Configuration:
Name : satopBundle3
Port: 5
Admin state              : Enabled
Link State               : Ready

```



```

Link Counter           : Up           0 time(s)
Framing                : T1 Framed
Signaling              : Robbed Bit
Line Coding            : B8ZS
Cable Length           : Long-Haul, > 660 feet
Receiver Gain          : 22.5dB
Clock Source           : CES pseudo-wire recovered (cesop1)
Clock Recovery         : Disabled
Recovered Clock Quality : QL_DNU
Transmit Clock Quality  : QL_DNU
Protocol               : TDM Circuit
Loopback               : Enabled (Local)
Idle Code              : 10
TDM Service Configuration:
Name : cesopBundle1
Time slots : 1-10, 11-20

```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

show ports tdm no-refresh

```
show ports {port_list} tdm {no-refresh}
```

Description

Displays the specified TDM port information.

Syntax Description

tdm	Time Division Multiplexing.
no-refresh	Page by page display without auto-refresh.

Default

N/A.

Usage Guidelines

Use this command to display the specified TDM port information in a page-by-page display without auto-refresh.



Example

```
E4G-400.2 # show ports tdm no-refresh
Port Summary
Port  Display      Service Name      Port  Link
#    String        (or #Services)   State State
=====
==
1          serBundle1       E     R
2          26                E     R
=====
==
Port State: D-Disabled, E-Enabled
Link State: A-Active, R-Ready, NP-Port not present, L-Loopback
E4G-400.2 # show ports tdm
Port Summary Monitor                               Tue Dec 13 11:51:26 2011
Port  Display      Service Name      Port  Link
#    String        (or #Services)   State State
=====
==
1          serBundle1       E     R
2          26                E     R
=====
==
Port State: D-Disabled, E-Enabled
Link State: A-Active, R-Ready, NP-Port not present, L-Loopback
U->page up  D->page down ESC->exit
E4G-400.2 # show ports 1 tdm
Port Summary Monitor                               Tue Dec 13 11:51:26 2011
Port  Display      Service Name      Port  Link
#    String        (or #Services)   State State
=====
==
1          serBundle1       E     R
=====
==
Port State: D-Disabled, E-Enabled
Link State: A-Active, R-Ready, NP-Port not present, L-Loopback
U->page up  D->page down ESC->exit
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

show ports transceiver information

```
show ports {port_list | tag tag} transceiver information
```



Description

Displays basic information about the optical transceiver.

Syntax Description

<i>port_list</i>	Specifies the port number(s).
<i>tag</i>	Specifies an 802.1Q or 802.1ad tag value.

Default

N/A.

Usage Guidelines

Digital Diagnostic Monitoring Interface (DDMI) provides critical system information about the installed optical modules. Use this command to monitor the condition of XFP, SFP, and SFP+ optical transceiver modules.

The tag value may be associated with either a VMAN or a VLAN.

If you try to execute this command on one of the ports in the port list that is non-compliant with DDMI, the following error message is displayed and the command does not go through:

```
Port 3:1 This command is not supported on this port. All ports and
transceiver of the ports requested in the command need to support DDMI.
```

If you try to execute this command on one of the ports in the port list on which the transceiver is non-compliant with DDMI, the following error message is displayed:

```
Port 3:1 This media/transceiver does not support enhanced digital
diagnostic monitoring interface (DDMI). All ports and transceiver of the
ports requested in the command need to support DDMI.
```

For more detailed information, use the `show ports transceiver information detail` command.

Example

The following display shows output for port 1:1-2 on a BlackDiamond 8810 switch:

```
BD-8810.2 # sh port 1:1-2 transceiver information
Port      Temp      TxPower  RxPower  TxBiasCurrent  Voltage-Aux1  Voltage-Aux2
(Celcius) (dBm)    (dBm)    (mA)      (Volts)        (Volts)
=====
==
1:1       30.60     -25.20   -18.70    0.40           5.09          5.07
1:2       30.60     -25.20   -18.70    0.40           5.09          N/A
=====
==
```



N/A indicates that the parameter is not applicable to the optics connected to the port

The following display shows output for port 25 on a Summit X480 switch:

```
X480-24x(10G4X).1 # sh ports 25 transceiver information
Port      Temp          TxPower
RxPower   TxBiasCurrent Voltage-Aux1 Voltage-Aux2
          (Celsius)      (dBm)        (dBm)
(mA)
(Volts)
=====
===
25         32.00         -3.35         -2.68
7.67      3.35          N/A
=====
===
```

History

This command was first available in ExtremeXOS 12.4.

Support for the Summit switches was added in ExtremeXOS 12.5.

The tag value was added in ExtremeXOS 12.4.4.

Support for SFP and SFP+ optics was added in ExtremeXOS 12.5.3.

Platform Availability

This command is available on the following platforms:

- BlackDiamond 8800—10G8Xc, 10G4Xc modules and S-10G1Xc option cards with 10G XFP optics.
- BlackDiamond 8900—10G8X-xl modules and S-10G1Xc option cards with 10G XFP optics.
- Summit X460 switch—SFP ZX and LX100 optics.
- Summit X480 switches and X480 VIM2-10G4X—XFP and SFP ZX and LX100 optics.

—SFP+ optics ER/LR.

- Summit X670 switches—SFP ZX and LX100 1G optics.

—SFP+ 10G optics.

show ports transceiver information detail

```
show ports {port_list | tag tag} transceiver information detail
```

Description

Displays detailed information about the optical transceiver.



Syntax Description

<i>port_list</i>	Specifies the port number(s).
<i>tag</i>	Specifies an 802.1Q or 802.1ad tag value.

Default

N/A.

Usage Guidelines

Digital Diagnostic Monitoring Interface (DDMI) provides critical system information about the installed optical modules. Use this command to monitor the condition of XFP, SFP, and SFP+ optical transceiver modules.

The tag value may be associated with either a VMAN or a VLAN.

If you try to execute this command on one of the ports in the port list that is non-compliant with DDMI, the following error message is displayed and the command does not go through:

```
Port 3:1 This command is not supported on this port. All ports and
transceiver of the ports requested in the command need to support DDMI.
```

If you try to execute this command on one of the ports in the port list on which the transceiver is non-compliant with DDMI, the following error message is displayed:

```
Port 3:1 This media/transceiver does not support enhanced digital
diagnostic monitoring interface (DDMI). All ports and transceiver of the
ports requested in the command need to support DDMI.
```

Example

```
BD-8810.2 # show port 2:* conf
Port Configuration Monitor                               Wed Sep 19 23:00:19
2012
Port      Virtual      Port Link Auto   Speed      Duplex  Flow  Load  Media
          router      State State Neg   Cfg Actual  Cfg Actual Cntrl Master Pri
Red
=====
==
2:1      VR-Default  E      A   OFF  40000  40000  FULL  FULL  SYM
%Q_UNKWN
2:2      VR-Default  E      NP  OFF  10000         FULL
2:3      VR-Default  E      NP  OFF  10000         FULL
2:4      VR-Default  E      NP  OFF  10000         FULL
2:5      VR-Default  E      R   OFF  40000         FULL
2:6      VR-Default  E      NP  OFF  10000         FULL
2:7      VR-Default  E      NP  OFF  10000         FULL
2:8      VR-Default  E      NP  OFF  10000         FULL
2:9      VR-Default  E      R   OFF  40000         FULL  %Q+LR4
2:10     VR-Default  E      NP  OFF  10000         FULL
2:11     VR-Default  E      NP  OFF  10000         FULL
2:12     VR-Default  E      NP  OFF  10000         FULL
```



```

2:13   VR-Default E      R   OFF 40000      FULL      %Q+SR4
2:14   VR-Default E      NP  OFF 10000      FULL      NONE
=====
==
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
Port State: D-Disabled, E-Enabled
Media: !/$/%/* - use "show port info detail" for more information
0->Clear Counters U->page up D->page down ESC->exit

```



Note

When a license is not installed, the restricted transceivers use a '\$'.

```

(pacman debug) BD-8810.6 # show port 2:* conf no
Port Configuration
Port      Virtual      Port Link Auto   Speed      Duplex      Flow Load  Media
         router      State State Neg   Cfg Actual   Cfg Actual Cntrl Master Pri
Red
=====
==
2:1      VR-Default E      R   OFF 40000      FULL      $Q
+SR4
2:2      VR-Default E      NP  OFF 10000      FULL      NONE
2:3      VR-Default E      NP  OFF 10000      FULL
NONE
2:4      VR-Default E      NP  OFF 10000      FULL
NONE
2:5      VR-Default E      R   OFF 40000      FULL
NONE
2:6      VR-Default E      NP  OFF 10000      FULL
NONE
2:7      VR-Default E      NP  OFF 10000      FULL
NONE
2:8      VR-Default E      NP  OFF 10000      FULL
NONE
2:9      VR-Default E      R   OFF 40000      FULL      $Q+LR4
2:10     VR-Default E      NP  OFF 10000      FULL      NONE
2:11     VR-Default E      NP  OFF 10000      FULL
NONE
2:12     VR-Default E      NP  OFF 10000      FULL      NONE
2:13     VR-Default E      R   OFF 40000      FULL      $Q+SR4
2:14     VR-Default E      NP  OFF 10000      FULL
NONE
2:15     VR-Default E      NP  OFF 10000      FULL      NONE
2:16     VR-Default E      NP  OFF 10000      FULL      NONE
2:17     VR-Default E      R   OFF 40000      FULL      NONE
2:18     VR-Default E      NP  OFF 10000      FULL      NONE
2:19     VR-Default E      NP  OFF 10000      FULL
NONE
=====
==
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port not present, L-Loopback
Port State: D-Disabled, E-Enabled
Media: !/$/%/* - use "show port info detail" for more information

```



```
BD-8810.4 # show port 2:1 info detail
Port: 2:1
  Virtual-router: VR-Default
  Type:          Q+SR4 ($ - Restricted Transceiver)
  Random Early drop:      Unsupported
  Admin state:   Enabled with 40G full-duplex
  Link State:    Ready (local fault)
  Link Ups:     1          Last: Wed Sep 05 04:38:19 2012
  Link Downs:   1          Last: Wed Sep 05 20:35:04 2012
```

```
BD-8810.4 # show port 2:9 info detail
Port: 2:9
  Virtual-router: VR-Default
  Type:          Q+LR4 (% - Unrestricted Transceiver)
  Random Early drop:      Unsupported
  Admin state:   Enabled with 40G full-duplex
  Link State:    Ready (local fault)
  Link Ups:     1          Last: Wed Sep 05 04:38:19 2012
  Link Downs:   1          Last: Wed Sep 05 20:35:04 2012
  Virtual-router: VR-Default
  Type:          SF+LR (! - Unsupported Transceiver)
  Random Early drop:      Unsupported
  Admin state:   Enabled with 10G full-duplex
  Link State:    Ready
  Link Ups:     0          Last: --
  Link Downs:   0          Last: --
```

The following example displays the current status of non-Extreme optical devices:

```
show port 3 transceiver information detail
Port : 3
  Media Type          : Q+SR4 (Licensed)
  or
  (Unsupported)
  or
  (Restricted, 10 Gbps maximum)
  Vendor Name         : NOTEXT,INC
  Part Number         : NOTEXT000EN-S001
  Serial Number       : NE7M90299

  Temp (Celsius)      : 0.00      Status      :
Low Alarm
  Low Warn Threshold  : 0.00      High Warn Threshold :
0.00
  Low Alarm Threshold : 0.00      High Alarm Threshold :
0.00

  Voltage AUX-1/Vcc (Volts) : 0.00      Status      :
Low Alarm
  Low Warn Threshold  : 0.00      High Warn Threshold :
0.00
  Low Alarm Threshold : 0.00      High Alarm Threshold :
0.00

  Tx Power (dBm)      : -inf      Status      :
Low Alarm
  Low Warn Threshold  : -inf      High Warn Threshold : -
inf
```



```

                Low Alarm Threshold   : -inf           High Alarm Threshold : -
inf
Press <SPACE> to continue or <Q> to quit:

```

History

This command was first available in ExtremeXOS 12.4.

Support for the Summit switches was added in ExtremeXOS 12.5.

The tag value was added in ExtremeXOS 12.4.4.

Support for SFP and SFP+ optics was added in ExtremeXOS 12.5.3.

Show output was updated in 15.3.

Show output was updated to include the current status of non-Extreme optical devices in ExtremeXOS 15.4.

Platform Availability

This command is available on the following platforms:

- BlackDiamond 8800—10G8Xc, 10G4Xc modules and S-10G1Xc option cards with 10G XFP optics.
- BlackDiamond 8900—10G8X-xl modules and S-10G1Xc option cards with 10G XFP optics.
- Summit X460 switch—SFP ZX and LX100 optics.
- Summit X480 switches and X480 VIM2-10G4X—XFP and SFP ZX and LX100 optics.

—SFP+ optics ER/LR.

- Summit X670 switches—SFP ZX and LX100 1G optics.

—SFP+ 10G optics.

show ports utilization

```

show ports {mgmt | port_list | tag tag | stack-ports stacking-port_list}
utilization {bandwidth | bytes | packet}

```

Description

Displays real-time port utilization information. The total utilization displays as real-time information, constantly refreshing, and the parameter displays show a snapshot of the activity on the port when you issue the command.

Syntax Description

mgmt	Specifies the management port.
<i>port_list</i>	Specifies one or more ports or slots and ports.



<i>tag</i>	Specifies an 802.1Q or 802.1ad tag value.
<i>stacking-port_list</i>	Specifies one or more stacking slots and ports.
bandwidth	Specifies port utilization as percentage of bandwidth.
bytes	Specifies port utilization in bytes per second.
packet	Specifies port utilization in packets per second.

Default

N/A.

Usage Guidelines

The software continuously monitors port utilization and calculates bandwidth as a function of each port's maximum link capacity.

The total utilization display presents real-time statistics. Use the *spacebar* to toggle the real-time displayed information for packets, bytes, and bandwidth in that order. When you use a parameter (packets, bytes, or bandwidth) with the command, the display for the specified type shows a snapshot per port when you issued the command. When the show ports utilization command is run with the bandwidth, bytes, or packets options, the command may need to be repeated a few times in order for the ExtremeXOS software to gather enough statistics to calculate appropriate values.

If you do not specify a port number, range of ports, or tag value, port utilization information is displayed for all ports.

The tag value may be associated with either a VMAN or a VLAN.

This status information may be useful for your technical support representative if you have a network problem.

Example

The following command displays utilization statistics for port 1 on a stand-alone switch:

```
show ports 1 utilization
```

The following command displays utilization statistics for slot 3, port 1 on a modular switch:

```
show ports 3:1 utilization
```

The following example shows sample output from the show ports utilization packets command:

```
Link Utilization Averages                               Mon Oct  6 22:38:25 2008
Port   Link   Rx           Peak Rx           Tx           Peak Tx
State  pkts/sec   pkts/sec   pkts/sec   pkts/sec
=====
```




```

==
1:1      A          47          191          0          0
1:2      A          0           0           0           0
2:1      R          0           0           0           0
2:2      R          0           0           0           0
3:1      R          0           0           0           0
3:2      R          0           0           0           0
4:1      R          0           0           0           0
4:2      R          0           0           0           0
5:1      R          0           0           0           0
5:2      R          0           0           0           0
6:1      R          0           0           0           0
6:2      R          0           0           0           0
7:1      R          0           0           0           0
7:2      R          0           0           0           0
=====
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
Spacebar->toggle screen U->page up D->page down ESC->exit

```



Note

Use the *spacebar* to toggle this real-time display for all ports from packets to bytes to bandwidth, in that order.

The following example shows sample output from the show ports utilization bytes command:

```

Link Utilization Averages                               Mon Oct  6 22:39:22 2008
Port   Link   Rx           Peak Rx           Tx           Peak Tx
State  bytes/sec  bytes/sec  bytes/sec  bytes/sec
=====
==
1:1      A          0           0           0           63
1:2      A          0           63          63           63
2:1      R          0           0           0           0
2:2      R          0           0           0           0
3:1      R          0           0           0           0
3:2      R          0           0           0           0
4:1      R          0           0           0           0
4:2      R          0           0           0           0
5:1      R          0           0           0           0
5:2      R          0           0           0           0
6:1      R          0           0           0           0
6:2      R          0           0           0           0
7:1      R          0           0           0           0
7:2      R          0           0           0           0
=====
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
Spacebar->toggle screen U->page up D->page down ESC->exit

```



The following example shows sample output of the show ports utilization bandwidth command:

```

Link Utilization Averages
Port      Link      Link  Rx      Peak Rx      Mon Oct  6 22:39:46 2008
State     Speed    % bandwidth  % bandwidth  % bandwidth  % bandwidth
=====
==
1:1       A         100      0.00      0.03         0.00         0.00
1:2       A         100      0.00      0.00         0.00         0.00
2:1       R          0        0.00      0.00         0.00         0.00
2:2       R          0        0.00      0.00         0.00         0.00
3:1       R          0        0.00      0.00         0.00         0.00
3:2       R          0        0.00      0.00         0.00         0.00
4:1       R          0        0.00      0.00         0.00         0.00
4:2       R          0        0.00      0.00         0.00         0.00
5:1       R          0        0.00      0.00         0.00         0.00
5:2       R          0        0.00      0.00         0.00         0.00
6:1       R          0        0.00      0.00         0.00         0.00
6:2       R          0        0.00      0.00         0.00         0.00
7:1       R          0        0.00      0.00         0.00         0.00
7:2       R          0        0.00      0.00         0.00         0.00
=====
==
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
Spacebar->toggle screen U->page up D->page down ESC->exit
    
```

History

This command was first available in ExtremeXOS 11.3.

The tag value was added in ExtremeXOS 12.4.4.

Platform Availability

This command is available on all platforms. The stack-ports option is available only on SummitStack.

show ports wan-phy configuration

```
show ports {port_list | tag tag} wan-phy configuration
```

Description

Displays the configuration of the specified WAN PHY port.

Syntax Description

<i>port_list</i>	Specifies the port number. On a stand-alone switch, this value is just the port number, and on a modular switch, this value is the slot and port number.
<i>tag</i>	Specifies an 802.1Q or 802.1ad tag value.



Default

Real-time information.

Usage Guidelines

This command displays the configuration for the specified WAN PHY ports. This command displays:

- Port number.
- Link.
- Framing.
- Clocking.
- Loopback.
- Section trace.
- Path trace.

If you do not specify a port number, range of ports, or tag value, the configuration is displayed for all WAN PHY ports. The default display is a constantly refreshing real-time display. If you specify the no-refresh parameter, the system displays a snapshot of the data at the time you issue the command.

The tag value may be associated with either a VMAN or a VLAN.

Example

The following command displays the configuration for all LW XENPAK WAN PHY ports:

```
show ports wan-phy configuration
```

Following is sample output from this command:

```
WAN PHY Port Configuration
Port Link Framing Clocking LoopBack SectionTrace PathTrace
=====
1 A SONET Internal off abcdefghijklmn 01234567890
2 R SONET Internal Line abcdefghijklmn 01234567890
3R SONET Line off abcdefghijklmn 01234567890
4A SONET Internal Line abcdefghijklmn 01234567890
5R SONET Line off abcdefghijklmn 01234567890
6A SONET Internal Line abcdefghijklmn 01234567890
7R SDH Line Line abcdefghijklmn 01234567890
=====
```

History

This command was available in ExtremeXOS 11.6.

The tag value was added in ExtremeXOS 12.4.4.



Platform Availability

This command is available on all 10G XFP ports only on the Summit X480 series switches.

show ports wan-phy errors

```
show ports {port_list | tag tag} wan-phy errors {no-refresh}
```

Description

Displays the error information of the specified WAN PHY port.

Syntax Description

<i>port_list</i>	Specifies the port number. On a stand-alone switch, this value is just the port number, and on a modular switch, this value is the slot and port number.
<i>tag</i>	Specifies an 802.1Q or 802.1ad tag value.
no-refresh	Specifies a static snapshot of data.

Default

Real-time information.

Usage Guidelines

This command displays the counters associated with parity errors for the specified WAN PHY ports. The display shows the following:

- Port number.
- B1, B2, and B3.
- Far end path block error count.
- Far end line BIP error count.

If you do not specify a port number, range of ports, or tag value, the errors are displayed for all WAN PHY ports. The default display is a constantly refreshing real-time display. If you specify the no-refresh parameter, the system displays a snapshot of the data at the time you issue the command.

The tag value may be associated with either a VMAN or a VLAN.

Example

The following command displays the errors for the LW XENPAK WAN PHY ports 1 to 4:

```
show ports 1-4 wan-phy errors
```



Following is sample output from this command:

```

WAN PHY Port Errors
Port      B1          B2          B3          FELBE          FEPLE
=====
1         0123456789 0123456789 0123456789 0123456789 0123456789
2         0123456789 0123456789 0123456789 0123456789 0123456789
3         0123456789 0123456789 0123456789 0123456789 0123456789
4         0123456789 0123456789 0123456789 0123456789 0123456789
=====
===
0->Clear Counters  U->page up  D->page down  ESC->exit
FELBE = Far End Path Block Error Count
FEPBE = Far End Line BIP Error Count

```

History

This command was available in ExtremeXOS 11.6.

The tag value was added in ExtremeXOS 12.4.4.

Platform Availability

This command is available on all 10G ports on 10G XFP ports only on the Summit X480 series switches only.

show ports wan-phy events

```
show ports {port_list | tag tag} wan-phy events {no-refresh}
```

Description

Displays the events information of the specified WAN PHY port.

Syntax Description

<i>port_list</i>	Specifies the port number. On a stand-alone switch, this value is just the port number, and on a modular switch, this value is the slot and port number.
<i>tag</i>	Specifies an 802.1Q or 802.1ad tag value.
no-refresh	Specifies a static snapshot of data.

Default

Real-time information.



Usage Guidelines

This command displays the events for the specified WAN PHY ports. The display shows whether each parameter is in the alarm or normal state; the alarm state includes alarm, warning, and error states.

This command displays notices for the following events:

- LOS: loss of signal.
- SEF: Severely Error Frame.
- LOF: loss of frame.
- LOP: loss of pointer.
- AIS-L: alarm indication signal, line.
- AIS-P: alarm indication signal, path.
- RDI-L: remote deflection indication, line.
- FarEnd PLM-P/LCD-P: Far end path label mismatch / loss of code-group delineation.
- FarEnd AIS-P/LCD-P: Far end path alarm indication signal / path loss of pointer.
- PLM-P: Path label mismatch.
- LCD-P: path loss of code-group delineation.

If you do not specify a port number, range of ports, or tag value, the errors are displayed for all WAN PHY ports. The default display is a constantly refreshing real-time display. If you specify the no-refresh parameter, the system displays a snapshot of the data at the time you issue the command.

The tag value may be associated with either a VMAN or a VLAN.

Example

The following command displays the events for the LW XENPAK WAN PHY ports 5 to 6:

```
show ports 5-6 wan-phy events
```

Following is sample output from this command:

```
WAN PHY Port Events:
Port LOS SEF LOF LOP AIS-L AIS-P RDI-L FarEnd FarEnd PLM-P LCD-P
      PLM-P/LCD-P AIS-P/LOP-P
=====
5A  NA  N   N   N   N   N   N   N   N   A
6N  N   A   N   N   N   N   A   N   A   N
=====
Event Status: A-Alarm N-Normal
```

History

This command was available in ExtremeXOS 11.6.

The tag value was added in ExtremeXOS 12.4.4.



Platform Availability

This command is available on 10G XFP ports only on the Summit X480 series switches only.

show ports wan-phy overhead

```
show ports {port_list | tag tag} wan-phy overhead {no-refresh}
```

Description

Displays selected WAN PHY OAM overhead information for the specified WAN PHY port.

Syntax Description

<i>port_list</i>	Specifies the port number. On a stand-alone switch, this value is just the port number, and on a modular switch, this value is the slot and port number.
<i>tag</i>	Specifies an 802.1Q or 802.1ad tag value.
no-refresh	Specifies a static snapshot of data.

Default

Real-time information.

Usage Guidelines

This command displays selected WAN PHY OAM overhead for the specified WAN PHY ports. The overhead number comes directly from the WAN PHY framing; for each overhead, both transmit and receive numbers are displayed.



Note

Overhead values display in hexadecimal format.

If you do not specify a port number, range of ports, or tag value, the errors are displayed for all WAN PHY ports. The default display is a constantly refreshing real-time display. If you specify the no-refresh parameter, the system displays a snapshot of the data at the time you issue the command.

The tag value may be associated with either a VMAN or a VLAN.

Example

The following command displays the events for the LW XENPAK WAN PHY ports 5 to 6:

```
show ports 5-6 wan-phy overhead
```



Following is sample output from this command:

```

WAN PHY Port Overhead Info
Port      C2      G1      K1      K2      M1      J0(0-15)      J1(0-15)
Rx      Rx      Rx      Rx      Rx      Rx      Rx
=====
5          0      0      0      0      F8  abcdefghijklmn  ExtremeNetworks
6         1A      0      0      0      F8  ExtremeNetworks  abcdefghijklmn
=====
Overhead values for C2, G1, K1, K2, and M1 are displayed in hexadecimal
representation.
Overhead values for J0 and J1 are displayed in ascii alphanumeric
representation with a dot representing a non-alphanumeric character.

```

History

This command was available in ExtremeXOS 11.6.

The tag value was added in ExtremeXOS 12.4.4.

Platform Availability

This command is available on 10G XFP ports only on the Summit X480 series switches only.

show sharing distribution port-based

```
show sharing distribution port port-based {ports port_list}
```

Description

Displays the load sharing distribution to member ports in the group specified by *port* for traffic received on ports in *port_list*.

Syntax Description

ports	Specifies the group of member ports to display for the load sharing.
<i>port_list</i>	Specifies the group of ports from which traffic is received.

Default

N/A.

Usage Guidelines

Use this command to display the load sharing distribution to member ports in the group specified by *port* for traffic received on ports in *port_list*. The selected member ports displayed are the results



of the calculation using the keys for the ports in the *port_list*, and the list of aggregator ports for the load sharing group. This command serves as port-based load sharing calculator for convenience.

Example

The following output shows the egress member ports selected for distribution in a load sharing group with master port 5:1, and aggregator ports 5:1 and 6:1 for packets received on ports 1:1-1:8, with default keys as shown.

```
BD-X8.4 # show sharing port-based keys ports 1:1-8
 1:1:  0   1:2:  1   1:3:  2   1:4:  3   1:5:  4   1:6:  5   1:7:
6     1:8:  7
```

```
BD-X8.5 # show sharing distribution 5:1 port-based keys ports 1:1-8
 1:1 -> 5:1
 1:2 -> 6:1
 1:3 -> 5:1
 1:4 -> 6:1
 1:5 -> 5:1
 1:6 -> 6:1
 1:7 -> 5:1
 1:8 -> 6:1
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

show sharing health-check

show sharing health-check

Description

Displays the configured health check LAGs on a switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.



Usage Guidelines

Use this command to display the health-check LAGs that have been configured on the switch.

Example

The following is sample output from this command:

```
(debug) BD-8810.1 # show sharing health-check
Member  Agg Admin Track          Track
Group   Port  Mbr State IP Addr          TCP Port Miss Freq State  Dn  Up
=====
==
2:8     2:1*  Y   En   30.1.1.1         23           3    3   Up    0   1
2:2*   Y   En   30.1.1.2         23           3    3   Up    0   1
2:3*   Y   En   30.1.1.3         23           3    3   Up    0   1
2:8*   -   En   30.1.1.8         80           3   10  Down  0   0
2:11*  Y   -   -                -            -    -   -     -   -
2:12*  -   En   44.1.3.2         80           3    4  Down  0   0
2:16   -   En   30.1.1.16        80           3   10  Dis   0   0
2:20   2:20* Y   En   192.1.1.1        80           10   3   Up    0   1
2:21*  Y   En   192.1.1.2        80           10   3   Up    0   1
=====
==
Member Port Flags: (*)Active, (!) Disabled
```

History

This command was first available in ExtremeXOS 12.13.

Platform Availability

This command is available on all platforms.

show sharing port-based keys

```
show sharing port-based keys {ports port_list}
```

Description

Displays the load sharing key values for all ports in the *port_list*. These values may be either default values, or configured values.

Syntax Description

ports	Specifies the ports for the load sharing.
<i>port_list</i>	Specifies one or more ports or slots and ports.



Default

N/A.

Usage Guidelines

Use this command to display the health-check LAGs that have been configured on the switch.

Example

The following output shows the default load sharing key values on a BlackDiamondf-X8 switch with a BDXA-10G48X module in slot 1 and a BDXA-40G24X module in slot 4 where slots 1 and 4 have not been configured to use a packet-based fabric hash. Keys assigned for port-based fabric hashing are displayed so that you can consider whether or not these keys are acceptable for a given application.

```
BD-X8.1 # show sharing
      port-based keysSlot 1 keys fixed for fabric
      hashing.Use "configure forwarding
      fabric hash packet slot 1" to configure. 1:1: 0
      1:2: 1 1:3: 2 1:4: 3 1:5: 4 1:6: 5 1:7: 6 1:8: 7
      1:9: 8 1:10: 9 1:11: 10 1:12: 11 1:13: 0 1:14: 1 1:15: 2
1:16: 3
      1:17: 4 1:18: 5 1:19: 6 1:20: 7 1:21: 8 1:22: 9 1:23: 10
1:24: 11
      1:25: 0 1:26: 1 1:27: 2 1:28: 3 1:29: 4 1:30: 5 1:31:
6 1:32: 7
      1:33: 8 1:34: 9 1:35: 10 1:36: 11 1:37: 0 1:38: 1 1:39: 2
1:40: 3
      1:41: 4 1:42: 5 1:43: 6 1:44: 7 1:45: 8 1:46: 9 1:47: 10
1:48: 11
      Slot 4 keys fixed for fabric hashing.
      Use "configure forwarding fabric hash packet slot 4" to configure.
4:1: 0 4:5: 1 4:9: 2 4:13: 3 4:17: 4 4:21: 5 4:25: 0
4:29: 1
      4:33: 2 4:37: 3 4:41: 4 4:45: 5 4:49: 0 4:53: 1 4:57: 2
4:61: 3
      4:65: 4 4:69: 5 4:73: 0 4:77: 1 4:81: 2 4:85: 3 4:89: 4
4:93: 5
```

The following output shows the default load sharing key values on a BlackDiamond-X8 switch with a BDXA-10G48X module in slot 1, and a BDXA-40G24X module in slot 4, where slots 1 and 4 have been configured to use a packet-based fabric hash. Note the assignment of consecutive values in the allowable range [0-15] to consecutive available ports in a slot.

```
BD-X8.3 # show sharing port-based keys
      1:1: 0 1:2: 1 1:3: 2 1:4: 3 1:5: 4 1:6: 5 1:7: 6 1:8:
7
      1:9: 8 1:10: 9 1:11: 10 1:12: 11 1:13: 12 1:14: 13 1:15: 14 1:16:
15
      1:17: 0 1:18: 1 1:19: 2 1:20: 3 1:21: 4 1:22: 5 1:23: 6 1:24:
7
      1:25: 8 1:26: 9 1:27: 10 1:28: 11 1:29: 12 1:30: 13 1:31: 14 1:32:
15
      1:33: 0 1:34: 1 1:35: 2 1:36: 3 1:37: 4 1:38: 5 1:39: 6
```



```

1:40: 7
 1:41: 8  1:42: 9  1:43: 10  1:44: 11  1:45: 12  1:46: 13  1:47: 14  1:48:
15
 4:1: 0  4:5: 1  4:9: 2  4:13: 3  4:17: 4  4:21: 5  4:25: 6  4:29:
7
4:33: 8  4:37: 9  4:41: 10  4:45: 11  4:49: 12  4:53: 13  4:57: 14  4:61:
15
4:65: 0  4:69: 1  4:73: 2  4:77: 3  4:81: 4  4:85: 5  4:89: 6  4:93:
7

```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

show slot

```
show slot {slot {detail} | detail }
```

Description

Displays the slot-specific information.

Syntax Description

<i>slot</i>	Specifies a slot on a modular switch or SummitStack.
detail	Specifies detailed port information.

Default

N/A.

Usage Guidelines

The `show slot` command displays the following information:

- The slot number.
- The type of module installed in the slot.
- The type of module configured for the slot.
- The state of the module, whether the power is down, if the module is operational, if a diagnostic being run, if there is a mismatch between the slot configuration and the module in the slot.



- The number of ports on the module.
- The current number of times the module has been restarted after a failure and the configured restart-limit.

**Note**

You may see slightly different information displayed depending on the platform and configuration you are using.

If you do not specify a slot number, information for all slots is displayed.

The display also includes a notice of insufficient power, should that arise.

The `show slot` command displays the following states, among others:

- Empty. (BlackDiamond 8800 series switch only—This also displays if you have a module in the chassis that is unsupported by the current software you are running.)
- Down.
- Power ON.
- Powered OFF.
- Booting.
- Initializing.
- VLAN sync.
- FDB sync.
- ACL sync.
- RT sync.
- Operational.

Show slot command on a SummitStack

The output of the show slot command is different for a SummitStack. The output of this command displays eight rows. If a node in the Active Topology is not assigned a slot number, the state of the slot is shown as Empty. A node that shows as a slot has successfully joined the active topology. This means, ExtremeXOS software can communicate with this node and does not mean that the node has been successfully brought up. The card state in the display indicates whether the slot was successfully started. If the card state is Operational, then the node is being used in the stack to carry user data as configured.

The output of the command displays the slot number, type of the Summit in that slot, state of the slot, and the number of ports.

The number of ports does not include the stacking links. It includes the option card ports regardless of whether the option card is installed.

This command is not available on Summit family switches operating in non stacking mode.

Example

You see slightly different displays, depending on the platform.



BlackDiamond X8 Switch

The following example displays output for a BlackDiamond X8 switch:

```
BD-X8.2 # show slot
Slots      Type                Configured          State      Ports  Flags
-----
-
Slot-1     BDXA-10G48X          BDXA-10G48X       Operational  48    MB
Slot-2     BDXA-10G48X          BDXA-10G48X       Operational  48    MB
Slot-3
Slot-4
Slot-5
Slot-6     BDXA-40G24X          BDXA-40G24X       Operational  96    MB
Slot-7     BDXA-40G24X          BDXA-40G24X       Operational  96    MB
Slot-8     BDXA-40G24X          BDXA-40G24X       Operational  96    MB
FM-1       BDXA-FM20T           BDXA-FM20T        Operational  0     MB
FM-2       BDXA-FM20T           BDXA-FM20T        Operational  0     MB
FM-3       BDXA-FM20T           BDXA-FM20T        Operational  0     MB
FM-4       BDXA-FM20T           BDXA-FM20T        Operational  0     MB
MM-A       BDX-MM1              BDX-MM1            Operational  0
MM-B       BDX-MM1              BDX-MM1            Operational  0
Flags : M - Backplane link to Master is Active
        B - Backplane link to Backup is also Active
        D - Slot Disabled
        I - Insufficient Power (refer to "show power budget")
```

Like management modules, the fabric modules on the BlackDiamond X8 switch are not configured and thus always show blank in the Configured column. The fabric modules do not have backplane links to the Master or Backup management module and so never show the M or B flag. They can be disabled, and there might be insufficient power to bring them up. None of the flags ever show up for the management modules, since there is always sufficient power to bring them up at the start of the day, and if power should become insufficient later on, the switch will have rebooted.

Here is an example of the detail slot display:

```
BD-X8.6 # show slot fm-4
FM-4 information:
State:                Failed
Download %:           0
Flags:
Last Error:           Wrong FM Type for Current Mode
Restart count:        6 (limit 5)
Serial number:        800434-00-00 000000000000
Hw Module Type:       BDXA-FM480
Configured Type:
Ports available:      0
Recovery Mode:        Reset
Debug Data:           Peer=Power ON
Flags : M - Backplane link to Master is Active
        B - Backplane link to Backup is also Active
        D - Slot Disabled
        I - Insufficient Power (refer to "show power budget")
BD-X8.7 #
```



The above display shows a fabric module slot that has failed due to incompatibility with the current fabric module mode. In this case the Last Error field appears and shows the following message:

```
Wrong FM Type for Current Mode
```

In the case of an I/O blade that has failed because there is no available switch fabric, the Last Error would contain the following:

```
No Fabric Module Is Available
```

In the case of an I/O blade that is incompatible with the current fabric mode, the Last Error would contain:

```
I/O Card Incompatible with Fabric Mode
```

Switches on a SummitStack

The following example displays module information for all slots in a stack:

```
* Slot-7 Stack.1 # show slot
Slots      Type              Configured          State              Ports
-----
Slot-1     SummitX            SummitX             Operational        26
Slot-2     X                  X                   Operational        50
Slot-3     X                  X                   Operational        26
Slot-4     SummitX            SummitX450-24x     Operational        26
Slot-5     X                  X                   Operational        26
Slot-6     X                  X                   Operational        26
Slot-7     X                  X                   Operational        26
Slot-8                                Empty               0
```

The following example displays module information for a specific slot on the stack:

```
* Slot-7 Stack.91 # show slot 1 detail
Slot-1 information:
State:                Operational
Download %:           100
Restart count:        0 (limit 5)
Serial number:        800187-00-02 0635G-00074
Hw Module Type:       SummitX450-24x
SW Version:           12.0.0.17
SW Build:              v1170b17
Configured Type:      SummitX450-24x
Ports available:      26
Recovery Mode:        Reset
Node MAC:              02:04:96:27:87:17
Current State:        STANDBY
Image Selected:       secondary
Image Booted:         secondary
Primary ver:          12.0.0.16
```



```

Secondary ver:      12.0.0.17
Config Selected:   primary.cfg

```

The following example displays detailed module information for all slots on a stack:

```

* Slot-7 Stack.90 # show slot detail
Slot-1 information:
State:              Operational
Download %:        100
Restart count:     0 (limit 5)
Serial number:     800187-00-02 0635G-00074
Hw Module Type:    SummitX
SW Version:        12.0.0.17
SW Build:          v1170b17
Configured Type:   SummitX
Ports available:   26
Recovery Mode:     Reset
Node MAC:          02:04:96:27:87:17
Current State:     STANDBY
Image Selected:    secondary
Image Booted:      secondary
Primary ver:       12.0.0.16
Secondary ver:     12.0.0.17
Config Selected:   primary.cfg
Slot-2 information:
State:              Operational
Download %:        100
Restart count:     0 (limit 5)
Serial number:     800163-00-04 0635G-01187
Hw Module Type:    SummitX
SW Version:        12.0.0.17
SW Build:          v1170b17
Configured Type:   SummitX
Ports available:   50
Recovery Mode:     Reset
Node MAC:          02:04:96:27:87:17
Current State:     STANDBY
Image Selected:    secondary
Image Booted:      secondary
Primary ver:       12.0.0.16
Secondary ver:     12.0.0.17
Config Selected:   primary.cfg
Slot-3 information:
State:              Operational
Download %:        100
Restart count:     0 (limit 5)
Serial number:     800152-00-04 0630G-00736
Hw Module Type:    SummitX
SW Version:        12.0.0.17
SW Build:          v1170b17
Configured Type:   SummitX
Ports available:   26
Recovery Mode:     Reset
Node MAC:          02:04:96:27:87:17
Current State:     STANDBY
Image Selected:    secondary
Image Booted:      secondary

```




```
Primary ver:          12.0.0.16
Secondary ver:       12.0.0.17
Config Selected:    primary.cfg
Slot-4 information:
State:              Operational
Download %:         100
Restart count:      0 (limit 5)
Serial number:      0635G-00073 S450-24X
Hw Module Type:    SummitX
SW Version:         12.0.0.17
SW Build:           v1170b17
Configured Type:    SummitX
Ports available:    26
Recovery Mode:      Reset
Node MAC:           02:04:96:27:87:17
Current State:      STANDBY
Image Selected:     secondary
Image Booted:       secondary
Primary ver:        12.0.0.16
Secondary ver:      12.0.0.17
Config Selected:    primary.cfg
Slot-5 information:
State:              Operational
Download %:         100
Restart count:      0 (limit 5)
Serial number:      800153-00-04 0646G-00683
Hw Module Type:    SummitX
SW Version:         12.0.0.17
SW Build:           v1170b17
Configured Type:    SummitX
Ports available:    26
Recovery Mode:      Reset
Node MAC:           02:04:96:27:87:17
Current State:      STANDBY
Image Selected:     secondary
Image Booted:       secondary
Primary ver:        12.0.0.16
Secondary ver:      12.0.0.17
Config Selected:    primary.cfg
Slot-6 information:
State:              Operational
Download %:         100
Restart count:      0 (limit 5)
Serial number:      800153-00-04 0646G-00691
Hw Module Type:    SummitX
SW Version:         12.0.0.17
SW Build:           v1170b17
Configured Type:    SummitX
Ports available:    26
Recovery Mode:      Reset
Node MAC:           02:04:96:27:87:17
Current State:      BACKUP
Image Selected:     secondary
Image Booted:       secondary
Primary ver:        12.0.0.16
Secondary ver:      12.0.0.17
Config Selected:    primary.cfg
Slot-7 information:
```



```

State:                Operational
Download %:          100
Restart count:       0 (limit 5)
Serial number:       800153-00-01 0603G-00741
Hw Module Type:     SummitX
SW Version:          12.0.0.17
SW Build:            v1170b17
Configured Type:     SummitX
Ports available:     26
Recovery Mode:       Reset
Node MAC:            02:04:96:27:87:17
Current State:       MASTER
Image Selected:      secondary
Image Booted:        secondary
Primary ver:         12.0.0.16
Secondary ver:       12.0.0.17
Config Selected:     primary.cfg
Slot-8 information:
State:                Empty
Restart count:       0 (limit 5)
Serial number:
Hw Module Type:
Configured Type:
Ports available:     0
Recovery Mode:       Reset
Node MAC:            00:00:00:00:00:00
Current State:
Image Selected:
Image Booted:
Primary ver:
Secondary ver:
Config Selected:
    
```

BlackDiamond 8800 series switch only

The following example displays module information for all slots:

Slots	Type	Configured	State	Ports	Flags
Slot-1			Empty	0	
Slot-2	G24X	G24X	Operational	24	M
Slot-3			Empty	0	
Slot-4			Empty	0	
Slot-5	G8X	G8X	Operational	8	M
Slot-6			Empty	0	
Slot-7		G48P	Empty	48	
Slot-8	G48P		Operational	48	M
Slot-9	10G4X	10G4X	Powered OFF	4	I
Slot-10			Empty	0	
MSM-A	MSM-48C		Operational	0	SMSM-
B			Empty	0	

Flags : M - Backplane link to Master MSM is Active
 B - Backplane link to Backup MSM is also Active



D - Slot Disabled
 I - Insufficient Power (refer to "show power budget")

The following example displays module information for a specified slot on a BlackDiamond 8810 switch:

```
BD-8810.3 # show slot 2
Slot-2 information:
State:                Operational
Download %:          100
Flags:               MB
Restart count:       0 (limit 5)
Serial number:       800114-00-04 04364-00013
Hw Module Type:     G48P
SW Version:          12.1.0.56
SW Build:            v1210b56
Configured Type:    G48P
Ports available:     48
Recovery Mode:       Reset
Flags : M - Backplane link to Master is Active
B - Backplane link to Backup is also Active
D - Slot Disabled
I - Insufficient Power (refer to "show power budget")
```

History

This command was first available in ExtremeXOS 10.1.

This command was first available on SummitStack in ExtremeXOS 12.0.

Platform Availability

This command is available only on modular switches and SummitStack.

show tdm hierarchy

show tdm hierarchy

Description

Displays the TDM hierarchy (T1/E1).

Syntax Description

tdm	Time Division Multiplexing.
hierarchy	Physical Layer Carrier System.

Default

N/A.



Usage Guidelines

Use this command to display the TDM hierarchy.

Example

```
Booted      : E1
Selected   : E1
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

show tdm service

```
show tdm service {circuit} {service_name}
```

Description

Displays the specified TDM service interface information.

Syntax Description

tdm	Time Division Multiplexing.
service	Provider provisioned TDM service.
circuit	TDM circuit service.
<i>service_name</i>	TDM service name.

Default

N/A.

Usage Guidelines

Use this command to display the TDM service interface information.

Example

```
Switch.1 # show tdm service
```



```

Service      CES PW      Flags
Name         Name
=====
==
serBun10    cesop10     EAt-----
serBun11    satop11     DRt-----
=====
==
Flags: (A) Active, (D) Admin Disabled, (E) Admin Enabled, (R) Ready,
(t) TDM Circuit
Number of TDM Service(s): 2
Switch.1 # show tdm service circuit
TDM Circuit Service Name: serBun10
Admin State      : Enabled
Oper State       : Enabled
Seized Code      : 0xa
Trunk Conditioning : 0xc8
Framing          : Unframed
Idle Pattern     : 0xff
Ports            : 39*
CES PW Name      : c1
Flags: (*) Active
TDM Circuit Service Name: serBun11
Admin State      : Enabled
Oper State       : Disabled
Seized Code      : 0xa
Trunk Conditioning : 0xc8
Framing          : Unframed
Idle Pattern     : 0xff
Ports            : 31
CES PW Name      : satop11
Flags: (*) Active

```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

unconfigure ip-fix

unconfigure ip-fix

Description

Unconfigures IPFIX globally.

Syntax Description

This command has no arguments or variables.



Default

N/A.

Usage Guidelines

Use this command to unconfigure IPFIX globally by removing all port and collector configuration and disabling IPFIX on all ports.

Example

The following command removes all IPFIX configuration:

```
unconfigure ip-fix
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on BlackDiamond 8900 G96Tc, G48T-xl, G48X-xl, and 10G8X-xl modules and Summit X460 and X480 switches.

unconfigure ip-fix flow-key

```
unconfigure ip-fix flow-key
```

Description

Unconfigures IPFIX flow key configuration and resets to the default.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to remove the flow-key settings that were configured using the following commands and reset to all the available keys:

```
configure ip-fix flow-key ipv4 configure ip-fix flow-key ipv6 configure ip-fix  
flow-key nonip
```



Example

The following command removes IPFIX flow key settings:

```
unconfigure ip-fix flow-key
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on BlackDiamond 8900 G96Tc, G48T-xl, G48X-xl, and 10G8X-xl modules and Summit X460 and X480 switches.

unconfigure ip-fix ip-address

```
unconfigure ip-fix ip-address
```

Description

Unconfigures the collector settings.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to unconfigure the collector settings and reset to the default.

Example

The following command returns to the default:

```
unconfigure ip-fix ip-address
```

History

This command was first available in ExtremeXOS 12.5.



Platform Availability

This command is available on BlackDiamond 8900 G96Tc, G48T-xl, G48X-xl, and 10G8X-xl modules and Summit X460 and X480 switches.

unconfigure ip-fix ports

```
unconfigure ip-fix ports port_list
```

Description

Unconfigures IPFIX on a port or group of ports.

Syntax Description

<i>port_list</i>	Specifies the ports.
------------------	----------------------

Default

N/A.

Usage Guidelines

Use this command to unconfigure IPFIX on a port or group of ports. This restores the configuration to the defaults for those ports. The global enable/disable of IPFIX is not affected by this command.

Example

The following command unconfigures IPFIX on port 2:

```
unconfigure ip-fix ports 2
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on BlackDiamond 8900 G96Tc, G48T-xl, G48X-xl, and 10G8X-xl modules and Summit X460 and X480 switches.

unconfigure ip-fix ports flow-key mask

```
unconfigure ip-fix ports port_list flow-key mask
```



Description

Unconfigures the IPv4 and IPv6 masks.

Syntax Description

<code>port_list</code>	Specifies the ports.
------------------------	----------------------

Default

N/A.

Usage Guidelines

Use this command to remove masks for the IPv4 and IPv6 source and destination address fields on ports. These masks were defined using one or more of the following commands:

```
configure ip-fix ports flow-key ipv4 mask ipaddress
configure ip-fix ports flow-key ipv6 mask ipaddress
```

Example

The following command removes a mask on port 2:1:

```
unconfigure ip-fix ports 2:1 flow-key mask
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on BlackDiamond 8900 G96Tc, G48T-xl, G48X-xl, and 10G8X-xl modules and Summit X460 and X480 switches.

unconfigure ip-fix source ip-address

```
unconfigure ip-fix source ip-address
```

Description

Unconfigures the source IP address used to communicate to the collector.

Syntax Description

This command has no arguments or variables.



Default

N/A.

Usage Guidelines

Use this command to unconfigure the collector and reset to the default.

Example

The following command returns to the default:

```
unconfigure ip-fix source ip-address
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on BlackDiamond 8900 G96Tc, G48T-xl, G48X-xl, and 10G8X-xl modules and Summit X460 and X480 switches.

unconfigure mlag peer interval

```
unconfigure mlag peer peer_name interval
```

Description

Unconfigures the length of time between health check hello packets.

Syntax Description

<i>peer_name</i>	Specifies an alpha numeric string identifying the MLAG peer.
------------------	--

Default

The interval default is 1000 milliseconds.

Usage Guidelines

Use this command to unconfigure the length of time between health check hello packets exchanged between MLAG peer switches and reset to the default.



Example

The following command unconfigures the interval on the switch101 peer. switch:

```
unconfigure mlag peer switch101 interval
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

unconfigure mlag peer ipaddress

```
unconfigure mlag peer peer_name ipaddress
```

Description

Unconfigures an MLAG peer switch IP address from an MLAG structure.

Syntax Description

<i>peer_name</i>	Specifies an alpha numeric string identifying the MLAG peer.
------------------	--

Default

N/A.

Usage Guidelines

Use this command to disassociate an MLAG peer structure with an MLAG peer switch IP address.

Example

The following command disassociates the MLAG peer structure switch101 with the MLAG peer switch IP address:

```
unconfigure mlag peer switch101 ipaddress
```

History

This command was first available in ExtremeXOS 12.5.



Platform Availability

This command is available on all platforms.

unconfigure network-clock sync-e

```
unconfigure network-clock sync-e [port port]
```

Description

Unconfigures synchronous Ethernet on a particular port to be a source 1 or source 2 for synchronizing clock.

Syntax Description

<code>source-1</code>	Source 1 external input clock
<code>source-2</code>	Source 2 external input clock
port	port
<i>port</i>	100Mbps/1G port Copper/Fiber Ports

Default

N/A.

Usage Guidelines

Use this command to unconfigure SyncE on a particular port.

Example

The following command unconfigures SyncE as a primary master on port 2:

```
unconfigure network-clock sync-e port 2
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on Summit X460-24X and X460-48X switches and on E4G-200 and E4G-400 switches.



unconfigure network-clock sync-e clock-source

```
unconfigure network-clock sync-e clock-sourceunconfigure network-clock sync-e  
clock-source
```

Description

Unconfigures the ethernet clock-source that is configured.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

Use this command to unconfigure the clock-source which was previously configured as source-1 or source-2.

Example

The following command unconfigures the Sync-E clock source configured earlier:

```
unconfigure network-clock sync-e clock-source
```

History

This Command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on Summit X460-24X and X460-48X switches and on E4G-200 and E4G-400 switches.

unconfigure port description-string

```
unconfigure ports port_list description string
```

Description

Unconfigures a description string setting.



Syntax Description

<code>port_list</code>	Specifies one or more ports or slots and ports.
------------------------	---

Default

N/A.

Usage Guidelines

Use this command to unconfigure a port description.

Example

The following command unconfigures the port description string:

```
unconfigure ports 1:3
```

History

This command was available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

unconfigure ports display string

```
unconfigure ports port_list display-string
```

Description

Clears the user-defined display string from one or more ports.

Syntax Description

<code>port_list</code>	Specifies one or more ports or slots and ports.
------------------------	---

Default

N/A.

Usage Guidelines

This command removes the display string that you configured using the `configure ports display-string` command.



Example

The following command clears the user-defined display string from slot 2, port 4 on a modular switch:

```
unconfigure ports 2:4 display-string
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

unconfigure ports redundant

```
unconfigure ports port_list redundant
```

Description

Clears a previously configured software-controlled redundant port.

Syntax Description

<i>port_list</i>	This refers to the primary port of the redundant pair and specifies one or more ports or slots and ports.
------------------	---

Default

N/A.

Usage Guidelines

The list of port numbers or the port display string specifies the primary port(s).

Example

The following command unconfigures a software-controlled redundant port on a modular switch:

```
unconfigure ports 2:3 redundant
```

History

This command was available in ExtremeXOS 11.1.



Platform Availability

This command is available on all platforms.

unconfigure ports tdm display string

```
unconfigure ports port_list tdm display-string
```

Description

Clears the user-defined display string from one or more ports.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
tdm	Indicates a time division multiplexing (TDM) port.

Default

N/A.

Usage Guidelines

Use this command to clear a user defined display string from TDM ports.

Example

Example output not yet available and will be provided in a future release.

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers E4G-200 and E4G-400.

unconfigure ports tdm recovered-clock

```
unconfigure ports port_list tdm recovered-clock
```

Description

Unconfigures the clock recovery from the specified TDM ports.



Syntax Description

tdm	Time Division Multiplexing.
recovered-clock	Clock recovered from the port.

Default

N/A.

Usage Guidelines

Use this command to unconfigure the clock recovery from the specified TDM ports.

Example

Example output not yet available and will be provided in a future release.

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

unconfigure ports wan-phy

```
unconfigure ports [port_list | all] wan-phy
```

Description

Resets the configuration parameters of the specified WAN PHY port to default values.

Syntax Description

<i>port_list</i>	Specifies the port number. On a stand-alone switch, this value is just the port number, and on a modular switch, this value is the slot and port number.
all	Specified all WAN PHY ports.

Default

N/A.



Usage Guidelines

XFP ports must be in WAN PHY mode. To configure the mode, use the `configure ports port_list mode {lan | wan-phy}` command.

When the all option is used, the command constructs a port_list of all applicable ports—LW-XENPAK ports and/or XFP ports that have been WAN-PHY enabled.

The default configurable WAN PHY OAM parameters are as follows:

- Framing—SONET.
- Clock source—line.
- Section trace—the IEEE default value, which has no string representation.
- Path trace—the IEEE default value, which has no string representation.
- Loopback—off.

Example

The following command unconfigures the configurable WAN PHY OAM parameters on a single LW XENPAK port on a modular switch:

```
unconfigure ports 2:1 wan-phy
```

History

This command was available in ExtremeXOS 11.6.

Platform Availability

This command is available on all 10G ports on 10G XFP ports only on the Summit X480 series switches.

unconfigure slot

```
unconfigure slot slot
```

Description

Clears a slot of a previously assigned module type.

Syntax Description

<i>slot</i>	Specifies a slot on a modular switch.
-------------	---------------------------------------

Default

N/A.



Usage Guidelines

If you issue the `unconfigure ports wan-phy` command on a slot containing a module with any ports configured for software-controlled redundancy, this command wipes away all software-controlled redundancy on both ports; both ports return to normal. Refer to the *ExtremeXOS Concepts Guide* for more information on software-controlled redundant ports.

Example

The following command clears slot 4 of a previously assigned module type:

```
unconfigure slot 4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available only on modular switches and SummitStack.



8 Universal Port Management Commands

```
configure log target upm filter
configure log target upm match
configure upm event
configure upm profile maximum execution-time
configure upm timer after
configure upm timer at
configure upm timer profile
create log target upm
create upm profile
create upm timer
delete log target upm
delete upm profile
delete upm timer
disable log target upm
disable upm profile
edit upm profile
enable log target upm
enable upm profile
run upm profile
show log configuration target upm
show upm events
show upm history
show upm history exec-id
show upm profile
show upm timers
unconfigure upm event
unconfigure upm timer
```

This chapter describes commands for:

- Configuring universal port profiles and triggers.
- Managing profiles and triggers.

For an introduction to universal port features, see the *ExtremeXOS Concepts Guide*.

configure log target upm filter

```
configure log target upm {upm_profile_name} filter filter-name {severity
[[severity] {only}}}
```

Description

Configures a log target to receive events that conform to a specific EMS filter and severity level requirements.

Syntax Description

<i>upm_profile_name</i>	Specifies a UPM log target to configure.
<i>filter-name</i>	Assigns an EMS filter to the specified log target.
<i>severity</i>	Specifies the minimum severity level for events sent to the log target.
only	Specifies that only events at the specified severity are sent to the log target.

Default

N/A.

Usage Guidelines

Events that meet the criteria established in the EMS filter and the optional severity requirements are forwarded to the UPM log target profile. You can further restrict the forwarded events with the following command:

```
configure log target upm {upm_profile_name} match {any | regex}.
```

Example

The following example configures UPM log target testprofile1 to receive events that meet the criteria defined in EMS filter testfilter1:

```
configure log target upm testprofile1 filter testfilter1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support



the Universal Port feature, see [Universal Port Management Commands](#) in the *ExtremeXOS Concepts Guide*.

configure log target upm match

```
configure log target upm {upm_profile_name} match {any | regex}
```

Description

Configures a log target to receive only those events that meet the specified match criteria.

Syntax Description

<i>upm_profile_name</i>	Specifies the UPM log target to be configured.
any	Matches any event. Use this option to remove a limitation configured with the <i>regex</i> option.
<i>regex</i>	Specifies an expression that must be contained in all forwarded events.

Default

N/A.

Usage Guidelines

This command further restricts the events selected by the command: `configure log target upm {upm_profile_name} filter filter-name {severity [[severity] {only}}}`.

Example

The following example configures UPM log target testprofile1 to receive events that meet the criteria contain the text warning:

```
configure log target upm testprofile1 match warning
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see [Universal Port Management Commands](#) in the *ExtremeXOS Concepts Guide*.



configure upm event

```
configure upm event upm-event profile profile-name ports port_list
```

Description

Configures a pre-defined event that triggers the named profile.

Syntax Description

<i>upm-event</i>	Specifies a pre-defined event type: device-detect, device-undetected, user-authenticate, user-unauthenticated.
<i>profile-name</i>	Specifies the profile to be configured.
<i>port-list</i>	Attaches the UPM profile to the specified port(s).

Default

N/A.

Usage Guidelines

This command configures a profile to be executed when the specified event occurs on the specified port(s).

You can configure multiple user profiles on the same port(s).

Example

The following example shows how to configure a profile on port 1:1, called "profile 1" that is triggered by the event "device-detect":

```
# configure upm event device-detect profile "p1" ports 1:1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see [Universal Port Management Commands](#) in the *ExtremeXOS Concepts Guide*.



configure upm profile maximum execution-time

```
configure upm profile profile-name maximum execution-time seconds
```

Description

Defines a maximum execution period for a profile.

Syntax Description

<i>seconds</i>	Defines the execution period in seconds. The range is 2 to 4294967295 seconds.
----------------	--

Default

30 seconds.

Usage Guidelines

If you make a mistake while configuring a profile and the profile loops, it will loop until the end of the maximum execution period. While testing new profiles, consider configuring a relatively short execution time so that any accidental loops do not create long delays during testing.

Example

The following example sets the execution period to 10 seconds:

```
# configure upm profile test maximum execution-time 10
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see [Universal Port Management Commands](#) in the *ExtremeXOS Concepts Guide*.

configure upm timer after

```
configure upm timer timer-name after time-in-secs {every seconds}
```



Description

Creates and names a UPM timer that is activated after the specified time in seconds.

Syntax Description

<i>timer-name</i>	Specifies the name of the UPM timer to be created.
<i>time-in-secs</i>	Configures the interval after which the UPM timer is activated.
<i>seconds</i>	Configures the UPM timer to be activated after every instance of the specified interval.

Default

N/A.

Usage Guidelines

Use this command to configure a timer that activates after the specified time. This is useful for deployment in CLI scripts, because you do not know what the current time will be when the script executes.

When a switch configuration is saved or restored, the UPM timers are activated only at the predetermined timings that were originally configured with the start time.

The periodic timer configured with the every keyword and the one-time timer configured with only the after keyword have a maximum range of one year in seconds (31,622,400 seconds).

Example

The following example configures the UPM timer "A" to be activated every 10 seconds, after an interval of 20 seconds:

```
# configure upm timer "timerA" after 20 every 10
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see [Universal Port Management Commands](#) in the *ExtremeXOS Concepts Guide*.



configure upm timer at

```
configure upm timer timer-name at month day year hour min secs {every seconds}
```

Description

Use this command to configure the time setting on a UPM timer.

Syntax Description

<i>timer-name</i>	Specifies the name of the UPM timer to be created.
<i>month</i>	Configures the month when the UPM timer is activated.
<i>day</i>	Configures the day when the UPM timer is activated.
<i>year</i>	Configures the year when the UPM timer is activated.
<i>hour</i>	Configures the hour when the UPM timer is activated.
<i>min</i>	Configures the minute when the UPM timer is activated.
<i>secs</i>	Configures the second when the UPM timer is activated.
<i>seconds</i>	Configures the UPM timer to be activated at every instance of the specified interval.

Default

N/A.

Usage Guidelines

Use this command to when you know the exact time you want an event to execute. If you use this command without the every keyword, the timer is activated once at the specified time. The every keyword configures a periodic timer that is activated at every instance of the time specified in seconds.

When a switch configuration is saved or restored, the UPM timers are activated only at the predetermined timings that were originally configured with the start time.

Example

The following example shows how to configure a timer, T1, that is activated every 10 seconds beginning at 1400 hours on October 16, 2006:

```
# configure upm timer "t1" at 10 16 2006 14 00 00 every 10
```

History

This command was first available in ExtremeXOS 11.6.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see [Universal Port Management Commands](#) in the *ExtremeXOS Concepts Guide*.

configure upm timer profile

```
configure upm timer timer-name profile profileName
```

Description

Associates a profile with a UPM timer.

Syntax Description

<i>timer-name</i>	Specifies the name of the UPM timer to be associated with the named profile.
<i>profileName</i>	Specifies the name of the profile to be associated with the UPM timer.

Default

N/A.

Usage Guidelines

Each timer can be attached to only one profile. Once a timer is configured to a profile, it must be unconfigured from that profile before it can be configured to a different profile.

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see [Universal Port Management Commands](#) in the *ExtremeXOS Concepts Guide*.

create log target upm

```
create log target upm {upm_profile_name}
```



Description

Creates a new UPM target profile.

Syntax Description

<code>upm_profile_name</code>	Specifies the name of an existing UPM profile.
-------------------------------	--

Default

N/A.

Usage Guidelines

After configuration, a UPM log target links an EMS filter with a UPM profile. This command creates the UPM log target.

The default configuration for a new log target binds the target to the EMS filter `defaultFilter`, which is used for all system events. To configure the log target, use the command: `configure log target upm {upm_profile_name} filterfilter-name {severity [[severity] {only}]}`.

The default status of a new UPM log target is disabled. To enable the log target, use the command: `enable log target upm {upm_profile_name}`.

To view the log target, use the command: `show log configuration target upm {upm_profile_name}`.

Example

The following example creates a new UPM log target named `testprofile1`:

```
create log target upm testprofile1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see [Universal Port Management Commands](#) in the *ExtremeXOS Concepts Guide*.

create upm profile

```
create upm profile profile-name
```



Description

Creates a new profile of a specified type.

Syntax Description

<code>profile-name</code>	Specifies the UPM profile to be created.
---------------------------	--

Default

N/A.

Usage Guidelines

Use this command to create a profile and name it. The maximum profile size is 5000 characters.

A UPM profile name must begin with an alphabetical character and may contain alphanumeric characters and underscores (`_`), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see [Object Names](#) in the *ExtremeXOS Concepts Guide*.

There is a limited capability to edit the profile with this command. If you enter a period (`.`) as the first and the only character on a line, you terminate the editing of the file. Use the command: `edit upm profile profile-name` for block mode capability.

Example

The following example shows how to create a profile named “P2”:

```
# create upm profile p2
enab por 2:
dis por 3:1
.
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see [Universal Port Management Commands](#) in the *ExtremeXOS Concepts Guide*.



create upm timer

```
create upm timer timer-name
```

Description

Creates and names a UPM timer.

Syntax Description

<code>timer-name</code>	Specifies the name of the UPM timer to be created.
-------------------------	--

Default

N/A.

Usage Guidelines

You can create UPM timers with a name. A profile can be associated with eight timers, but a timer can be bound to only one profile at any point in time. You can create a maximum of 32 timers. A name space for the timers is available to help when you are typing the commands.

A UPM timer name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see [Object Names](#) in the *ExtremeXOS Concepts Guide*.

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see [Universal Port Management Commands](#) in the *ExtremeXOS Concepts Guide*.

delete log target upm

```
delete log target upm {upm_profile_name}
```

Description

Deletes the specified UPM log target.



Syntax Description

<code>upm_profile_name</code>	Specifies the name of the UPM log target to be deleted.
-------------------------------	---

Default

N/A.

Usage Guidelines

This command deletes the log target and any configurations applied to that target. To disable a target and retain the target configuration, use the following command:

```
disable log target upm {upm_profile_name}.
```

Example

The following command deletes the UPM log target testprofile1:

```
delete log target upm testprofile1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see [Universal Port Management Commands](#) in the *ExtremeXOS Concepts Guide*.

delete upm profile

```
delete upm profile profile-name
```

Description

Deletes the specified profile.

Syntax Description

<code>profile-name</code>	Specifies the UPM profile to be deleted.
---------------------------	--



Default

N/A.

Example

The following command deletes a UPM profile called `sample_1`:

```
delete upm profile sample_1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see [Universal Port Management Commands](#) in the *ExtremeXOS Concepts Guide*.

delete upm timer

```
delete upm timer timer-name
```

Description

Deletes the specified UPM timer.

Syntax Description

<code>timer-name</code>	Specifies the name of the UPM timer to be deleted.
-------------------------	--

Default

N/A.

Usage Guidelines

You can delete a UPM timer by specifying its name.

History

This command was first available in ExtremeXOS 11.6.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see [Universal Port Management Commands](#) in the *ExtremeXOS Concepts Guide*.

disable log target upm

```
disable log target upm {upm_profile_name}
```

Description

Disables the specified UPM log target.

Syntax Description

<code>upm_profile_name</code>	Specifies the name of the UPM log target to be disabled.
-------------------------------	--

Default

N/A.

Usage Guidelines

This command disables the log target and retains any configurations applied to that target. To delete a target and any configuration applied to the target, use the following command:

```
delete log target upm {upm_profile_name}.
```

Example

The following command disables the UPM log target testprofile1:

```
disable log target upm testprofile1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see [Universal Port Management Commands](#) in the *ExtremeXOS Concepts Guide*.



disable upm profile

disable upm profile *profile-name*

Description

Disables the use of the specified Universal Port profile on the switch.

Syntax Description

<i>profile-name</i>	Specifies the UPM profile to be disabled.
---------------------	---

Default

A UPM profile is enabled by default.

Example

The following command disables a UPM profile called sample_1 on the switch:

```
disable upm profile sample_1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see [Universal Port Management Commands](#) in the *ExtremeXOS Concepts Guide*.

edit upm profile

edit upm profile *profile-name*

Description

Allows you to edit the specified profile.

Syntax Description

<i>profile-name</i>	Specifies the UPM profile to be edited.
---------------------	---



Default

N/A.

Usage Guidelines

Use the command to have VI-like editor features for editing the profile. Changes appear when you close the file for editing, not when you save it.

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see [Universal Port Management Commands](#) in the *ExtremeXOS Concepts Guide*.

enable log target upm

```
enable log target upm {upm_profile_name}
```

Description

Enables the specified UPM log target.

Syntax Description

<code>upm_profile_name</code>	Specifies the name of the UPM log target to be enabled.
-------------------------------	---

Default

N/A.

Usage Guidelines

UPM log targets are disabled when they are created.

Example

The following command enables the UPM log target testprofile1:

```
enable log target upm testprofile1
```



History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see [Universal Port Management Commands](#) in the *ExtremeXOS Concepts Guide*.

enable upm profile

```
enable upm profile profile-name
```

Description

Enables the use of the specified Universal Port profile on the switch.

Syntax Description

<i>profile-name</i>	Specifies the UPM profile to be enabled.
---------------------	--

Default

A UPM profile is enabled by default.

Example

The following command enables a UPM profile called example on the switch:

```
enable upm profile example
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see [Universal Port Management Commands](#) in the *ExtremeXOS Concepts Guide*.



run upm profile

```
run upm profile profile-name {event event-name} {variables variable-string}
```

Description

Executes the specified Universal Port profile on the switch.

Syntax Description

<i>profile-name</i>	Specifies the UPM profile to be run.
<i>event-name</i>	Specifies an event type for the specified profile. Valid event types are device-detect, device-undetect, user-authenticate, and user-unauthenticated.
<i>variable-string</i>	Specifies a string of variable names and the assigned variable values to be used in the profile. The format is: <i>var_name1=value_1</i> ; <i>var_name2=value_2</i> ; <i>var_name3=value_3</i> . Each variable name is followed by the equal sign (=), the variable value, and a semicolon (;).

Default

N/A.

Example

The following command runs a UPM profile called example on the switch:

```
run upm profile example
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see [Universal Port Management Commands](#) in the *ExtremeXOS Concepts Guide*.

show log configuration target upm

```
show log configuration target upm {upm_profile_name}
```



Description

Displays a UPM target profile configuration.

Syntax Description

<i>upm_profile_name</i>	Specifies the name of the UPM target profile you want to view.
-------------------------	--

Default

N/A.

Usage Guidelines

None.

Example

The following example displays the configuration for the UPM log target named testprofile1:

```
show log configuration target upm testprofile1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see [Universal Port Management Commands](#) in the *ExtremeXOS Concepts Guide*.

show upm events

```
show upm events event-type
```

Description

Displays UPM events of the specified type.

Syntax Description

<i>event-type</i>	Displays events of the specified type for all profiles. Valid values for event-type are: device-detectdevice-removeuser-authenticateuser-unauthenticated.
-------------------	---



Default

N/A.

Usage Guidelines

Use this command to display the following types of events:

- device-detect.
- device-remove.
- user-authenticate.
- user-unauthenticated.

Example

The following command displays device-detect events:

```
show upm event device-detect
```

The output of the command is similar to the following:

```

-----
UPM Profile          PortList
-----
profile1             3
-----

```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see [Universal Port Management Commands](#) in the *ExtremeXOS Concepts Guide*.

show upm history

```
show upm history {profile profile-name | event upm-event | status [pass | fail] |
timer timer-name | detail}
```

Description

Displays (in a tabular column) a list of UPM profile events executed on the switch.



Syntax Description

<i>profile-name</i>	Displays UPM events for the specified profile.
<i>upm-event</i>	Displays UPM events that were triggered by the specified event.
status [pass fail]	Displays UPM events that meet the specified status, which is either pass or fail.
<i>timer-name</i>	Displays UPM events that were triggered by the specified timer.
detail	Displays additional detail on UPM events.

Default

N/A.

Usage Guidelines

This is useful for trouble shooting and testing.

Example

The following example shows what appears when no UPM events have been triggered:

```
* VLAB-R3-BD8808.2 # show upm history
-----
--
Exec   Event/           Profile           Port Status Time Launched
Id     Timer/ Log filter
-----
--
--
Number of UPM Events in Queue for execution: 0
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see [Universal Port Management Commands](#) in the *ExtremeXOS Concepts Guide*.

show upm history exec-id

```
show upm history exec-id number
```



Description

Displays information about an instance of a UPM profile executed on the switch.

Syntax Description

<i>number</i>	Specifies the execution identifier for the event you want to view.
---------------	--

Default

N/A.

Usage Guidelines

To view the execution identifiers for which you can display information, enter the show upm history command.

Example

The following example shows information for the event identified as 8006:

```
* BD-8808.4 # show upm history exec 8006
UPM Profile: p1
Event: User Request      , Time run: 2006-10-18 11:56:15
Execution Identifier: 8006      Execution Status: Pass
Execution Information:
1 # enable cli scripting
2 # set var EVENT.NAME USER-REQUEST
3 # set var EVENT.TIME 1161172575
4 # set var EVENT.PROFILE p1
5 # enable por 1:1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see [Universal Port Management Commands](#) in the *ExtremeXOS Concepts Guide*.

show upm profile

```
show upm profile name
```



Description

Displays a list of the UPM profiles on the system and some of their configuration information, or the contents of a specified profile.

Syntax Description

<i>name</i>	Displays the contents of the specified profile.
-------------	---

Default

N/A.

Usage Guidelines

To see a list of all UPM profiles on a switch, use the command without the name option. The resulting display shows the names of the profiles on the system and their status, active or disabled.

Use the name option to see the contents of a specific profile.

Example

The output of the command is similar to the following:

```
* BD-8808.36 # show upm profile
=====
==
UPM Profile Events Ports Flags
=====
==
p1 UPM Timer(t1) e
p1 device detect 1:1 e
p2 e
=====
==
Number of UPM Profiles: 2
Number of UPM Events in Queue for execution: 0
Flags: d - disabled, e - enabled
Event name: log-message(Log filter name) - Truncated to 20 chars
* BD-8808.37 # show upm profile "p1"
Created at : 2010-04-11 04:07:41
Last edited at : 2010-04-11 04:07:41
*****Profile Contents Begin*****
ena por 1:1
*****Profile Contents Ends*****
Profile State: Enabled
Profile Maximum Execution Time: 30
Events and ports configured on the profile:
=====
Event                               Port list/Log filter
=====
device-detect                        1:1
device-undetected                    :
```



```

user-authenticated      :
user-unauthenticated    :
=====

```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see [Universal Port Management Commands](#) in the *ExtremeXOS Concepts Guide*.

show upm timers

```
show upm timers
```

Description

Displays a list of the UPM timers on the system and some of their configuration information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to see a list of all UPM timers on a switch. The resulting display shows timer configurations, the associated profile, and flags to indicate the timer status. Flags are defined below:

- a - active.
- d - disabled.
- p - periodic.

Example

This command displays UPM timer configuration:

```
show upm timers
```



The output of this command is similar to the following:

```
* BD-8808.43 # show upm timers
Current Time: 2006-10-16 14:03:55
-----
--
UPM Profile Flags Next Execution
Timer Name time
-----
--
t1 pl ep 2006-10-16 14:04:00(Every 10 secs)
timerA
-----
--
Flags: e - Profile is enabled, d: Profile is disabled
o -Timer is non-periodic, p - Timer is periodic
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see [Universal Port Management Commands](#) in the *ExtremeXOS Concepts Guide*.

unconfigure upm event

```
unconfigure upm event upm-event profile profile-name ports port_list
```

Description

Unconfigures the event from the specified profile and port list.

Syntax Description

<i>upm-event</i>	Specifies the type of event to be unconfigured.
<i>profile-name</i>	Specifies the profile from which the event is unconfigured.
<i>port-list</i>	Unconfigures the UPM profile from the specified port list.

Default

N/A.



Usage Guidelines

This command removes an event from the specified profile and port list.

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see [Universal Port Management Commands](#) in the *ExtremeXOS Concepts Guide*.

unconfigure upm timer

```
unconfigure upm timer timer-name profile profile-name
```

Description

Removes a UPM profile from a UPM timer.

Syntax Description

<i>timer-name</i>	Unconfigures the specified UPM timer and deactivates any running timer.
<i>profile-name</i>	Removes the specified profile from the UPM timer.

Default

N/A.

Usage Guidelines

Use this command to unconfigure a timer setting. This command does not delete the timer.



Note

The specified timer is stopped by this command, even if it has been activated.

History

This command was first available in ExtremeXOS 11.6.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see [Universal Port Management Commands](#) in the *ExtremeXOS Concepts Guide*.



9 CLI Scripting Commands

```
configure cli mode
configure cli mode scripting
configure cli script timeout
delete var
delete var key
disable cli scripting
disable cli scripting output
ELSE
enable cli scripting
enable cli scripting output
ENDIF
ENDWHILE
IF ... THEN
load var key
return
save var key
set var
show var
WHILE ... DO
```

This chapter describes commands for:

- Enabling and disabling CLI scripting.
- Managing script variables.
- Creating conditional loops.

For an introduction to CLI scripting features, see the *ExtremeXOS Concepts Guide*.

configure cli mode

```
configure cli mode [persistent | non-persistent]
```

Description

Configures the persistent nature of command execution for non-persistent commands.

Syntax Description

persistent	Configures command execution to be persistent.
non-persistent	Configures command execution to be not persistent.

Default

The default mode is non-persistent.

Usage Guidelines

All ExtremeXOS commands can operate in persistent mode, and a subset of the ExtremeXOS command set can operate in non-persistent mode. Commands that are executed in persistent mode become part of the saved switch configuration that persists when the switch is rebooted. Commands that are executed in non-persistent mode configure temporary changes that are not saved in the switch configuration and do not persist when the switch is rebooted.

Most commands operate only in persistent mode. The subset of commands that operate in non-persistent mode are called non-persistent-capable commands. The Universal Port feature uses the non-persistent-capable commands to configure temporary changes that could create security issues if the switch were rebooted or reset. The use of non-persistent-capable commands in scripts and Universal Port profiles allows you to make temporary configuration changes without affecting the default configuration the next time the switch is started.

The `configure cli mode` command affects only the non-persistent-capable commands, which are listed in the Universal Port chapter in the *Extreme XOS Concepts Guide*. By default, all commands operate in persistent mode with the following exceptions:

- In Universal Port dynamic profiles, the non-persistent-capable commands operate in non-persistent mode unless preceded by the `configure cli mode persistent` command in the profile.
- In the CLI, CLI scripts, and static profiles, the non-persistent-capable commands operate in non-persistent mode only when preceded by the `configure cli mode non-persistent` command.

You can use the `configure cli mode persistent` command and the `configure cli mode non-persistent` command to change the mode of operation for non-persistent-capable commands multiple times within a script, profile, or configuration session.

Example

The following example sets command execution to be persistent:

```
configure cli mode persistent
```

History

This command was first available in ExtremeXOS 11.6.



Platform Availability

This command is available on all platforms.

configure cli mode scripting

```
configure cli mode scripting [abort-on-error | ignore-error]
```

Description

Configures the error handling process for CLI scripting on the switch.

Syntax Description

abort-on-error	Configures Cli scripts to be aborted if a CLI error occurs.
ignore-error	Configures the script to be executed when CLI errors occur.

Default

CLI: ignore-error Static profiles: abort-on-error Dynamic profiles: abort-on-error

Usage Guidelines

You can change the error-handling options within the scripts.

Example

The following command configures the switch to ignore syntax errors in CLI scripts:

```
configure cli mode scripting ignore-error
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure cli script timeout

```
configure cli script timeout timeout
```



Description

Configures the maximum time a script can run.

Syntax Description

<code>timeout</code>	Defines the timeout period in seconds.
----------------------	--

Default

Regular script: no time limit default.xsf: 500 seconds autoexec.xsf: 500 seconds

Usage Guidelines

This command configures the maximum run time for all scripts, including default.xsf and autoexec.xsf, which are described in *Software Upgrade and Boot Options* in the *ExtremeXOS Concepts Guide*. If no timeout period is configured, regular scripts do not timeout, and the default.xsf and autoexec.xsf scripts time out after 500 seconds.

If a script does not finish running in the configured time, command execution stops and an error message is logged. If the timer expires while a command is executing, the command execution continues and all following commands are not executed.

If the timer command is executed inside a script, the timer is reset. If the command is issued more than once inside a script the last timer command executed resets the timer. The timer is valid only for that session. The use of nested scripts does not extend the execution period. When the parent script reaches the timeout value, the parent script and all nested scripts terminate.

To configure a different timeout value for autoexec.xsf or default.xsf, the configure cli script timeout command should be the first command in the script.

When a script timeout value is configured, the following variables are created: \$CLI.SCRIPT_TIMEOUT and \$CLI.SCRIPT_TIME_REMAINING. If no timeout value is configured for a session, the variables are not created.

You can use the \$CLI.SCRIPT_TIMEOUT variable to adjust the timeout value. The \$CLI.SCRIPT_TIME_REMAINING variable returns the time remaining. When a timeout value is configured, the variable values are as follows:

- If no script is running, both \$CLI.SCRIPT_TIME_REMAINING and \$CLI.SCRIPT_TIMEOUT show the configured timeout value.
- If a script is aborted due to timeout, the \$CLI.SCRIPT_TIME_REMANING variable returns the value0.
- If a script finishes execution (before the timeout value is reached) the \$CLI.SCRIPT_TIME_REMANING variable returns the remaining time.



Example

The following command configures the switch to terminate a script after 120 seconds:

```
configure cli script timeout 120
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

delete var

delete var *varname*



Note

This is a script command and operates only in scripts or on the command line when scripting is enabled with the following command: `enable cli scripting {permanent}`.

Description

Deletes a variable.

Syntax Description

<i>varname</i>	Specifies the name of the scripting variable to be deleted.
-----------------------	---

Default

N/A.

Usage Guidelines

The format of a local variable (case insensitive) is: \$VARNAME.

Example

The following example deletes local variable x:

```
delete var x
```



History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

delete var key

delete var key *key*



Note

This is a script command and operates only in scripts or on the command line when scripting is enabled with the following command: `enable cli scripting {permanent}`.

Description

Deletes the variables that have been saved using a key.

Syntax Description

<i>key</i>	Specifies that variables associated with the specified key must be deleted.
------------	---

Default

N/A.

Usage Guidelines

CLI scripting must be enabled to use this command. The user is responsible for generating unique keys for each variable. The system has a limited amount of memory to store these variables.

Example

The following command deletes all variables associated with the key “red:”

```
delete var key red
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.



disable cli scripting

```
disable cli scripting {permanent}
```

Description

Disables the use of the CLI scripting commands. When used without the permanent option, it disables the CLI scripting commands for the current session and is a per session setting. The permanent option affects new sessions only and is saved across switch reboots.

Syntax Description

permanent	Disables the CLI scripting commands for new sessions only; this setting is saved across switch reboots.
------------------	---

Default

CLI scripting commands are disabled by default.

Usage Guidelines

You can disable the CLI scripting commands for the session only after this feature has been enabled.

Example

The following command disables the CLI scripting commands for the current session:

```
disable cli scripting
```

History

This command was first available in ExtremeXOS 11.6.

The permanent option was added in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

disable cli scripting output

```
disable cli scripting output
```



Description

Disables the display of CLI commands and responses during script operation.

Syntax Description

This command has no arguments or variables.

Default

During interactive script sessions: CLI scripting output enabled.

During load script command operation: CLI scripting output disabled.

Usage Guidelines

When the CLI scripting output is disabled, the only script output displayed is the `show var {varname}` command and its output. All other commands and responses are not displayed.

When the `load script filename {arg1} {arg2} ... {arg9}` command is entered, the software disables CLI scripting output until the script is complete, and then CLI scripting output is enabled. Use the `enable cli scripting output` and `disable cli scripting output` commands to control what a script displays when you are troubleshooting.

Example

The following command disables CLI scripting output for the current session or until the `enable cli scripting output` command is entered:

```
disable cli scripting output
```

History

This command was first available in ExtremeXOS 12.1.2.

Platform Availability

This command is available on all platforms.

ELSE

ELSE



Note

This is a script statement and operates only in scripts when scripting is enabled with the following command: `enable cli scripting {permanent}`.



Description

Command block to be executed if the condition specified in the associated IF statement is not met.

Syntax Description

statements	Actions to be executed when the conditions specified in the associated IF statement are not met.
------------	--

Default

N/A.

Usage Guidelines

CLI scripting must be enabled before using this command.

This command must be preceded by IF *_expression* THEN *statements* and followed by ENDIF.

You can insert comments by using a number sign (#).

Example

The following example executes the show switch command if the value of the variable x is greater than 2, and execute the show vlan command otherwise:

```
IF ($x > 2) THEN
    show switch
ELSE
    show vlan
ENDIF
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

enable cli scripting

```
enable cli scripting {permanent}
```



Description

Enables the use of CLI scripting commands. When used without the permanent option, it enables the CLI scripting commands for the current session and is a per session setting. The permanent option enables the CLI scripting commands for new sessions only and is saved across switch reboots.

Syntax Description

permanent	Enables the CLI scripting commands for new sessions only; this setting is saved across switch reboots.
------------------	--

Default

The CLI scripting commands are disabled by default.

Usage Guidelines

You must enable the CLI scripting commands on the switch to use the scripting keywords in the script, and before you can configure or execute a script.



Note

CLI scripting commands cannot be enabled when CLI space auto completion is enabled with the `enable cli space-completion` command.

Example

The following command enables the CLI scripting commands for the current session:

```
enable cli scripting
```

History

This command was first available in ExtremeXOS 11.6.

The permanent option was added in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

enable cli scripting output

```
enable cli scripting output
```



Description

Enables the display of CLI commands and responses during script operation.

Default

During interactive script sessions: CLI scripting output enabled.

During load script command operation: CLI scripting output disabled.

Usage Guidelines

When the CLI scripting output is enabled, all script commands and responses are displayed.

When the `load script filename {arg1} {arg2} ... {arg9}` command is entered, the software disables CLI scripting output until the script is complete, and then CLI scripting output is enabled. Use the `enable cli scripting output` and `disable cli scripting output` commands to control what a script displays when you are troubleshooting.

Example

The following command enables CLI scripting output for the current session or until the `disable cli scripting output` command is entered:

```
enable cli scripting output
```

History

This command was first available in ExtremeXOS 12.1.2.

Platform Availability

This command is available on all platforms.

ENDIF

ENDIF



Note

This is a script command and operates only in scripts when scripting is enabled with the following command: `enable cli scripting {permanent}`.

Description

Causes the IF construct to be terminated.



Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The ENDIF command should be used after the `IF _expression THEN statement(s)` command.

You can insert comments by using a number sign (#). CLI scripting must be enabled to use this command.

Example

The following example executes the show switch command if the value of the variable is greater than 2 and execute the show vlan command otherwise:

```
IF ($x > 2) THEN
    show switch
ELSE
    show vlan
ENDIF
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

ENDWHILE

ENDWHILE



Note

This is a script command and operates only in scripts when scripting is enabled with the following command: `enable cli scripting {permanent}`.



Description

Causes the WHILE construct to be terminated.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The **ENDWHILE** command must be used after a corresponding **WHILE** *_expression* **DO***statement(s)* command.

You can insert comments by using a number sign (#). CLI scripting must be enabled to use this command.

Example

This example creates 10 VLANs, named x1 to x10:

```
set var x 1

WHILE ($x <= 10) DO
    create vlan v$x
    set var x ($x + 1)

ENDWHILE
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.



IF ... THEN

IF (expression) THEN



Note

This is a script command and operates only in scripts when scripting is enabled with the following command: `enable cli scripting {permanent}`.

Description

Optionally executes a code block based on the condition supplied.

Syntax Description

expression	Specifies the condition for which the statements should be executed.
statements	Actions to be executed when the specified conditions are met.

Default

N/A.

Usage Guidelines

This command is usually followed by statements that are executed if the condition evaluates to true.

It can also be followed by an ELSE block, which is executed if the condition evaluates to false.

The IF construct should be terminated by an **ENDIF** command.

The expression must be enclosed in parentheses.

The IF construct can be nested inside other IF and WHILE constructs. Nesting is supported up to five levels. If there is incorrect nesting of IF conditions, an error message is displayed. If a user tries to execute more than five nested IF conditions, an error message is displayed.

The operators mentioned in [Using Operators](#) can be used in an expression in an IF condition.

You can insert comments by using a number sign (#).

Example

The following example executes the show switch command if the value of the variable is greater than 2 and executes the show vlan command otherwise:

```
IF ($x > 2) THEN
    show switch
```



```

ELSE
    show vlan
ENDIF

```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

load var key

```
load var key key [var1 var2 ...]
```



Note

This is a script command and operates only in scripts or on the command line when scripting is enabled with the following command: `enable cli scripting {permanent}`.

Description

Imports the specified set of variables associated with a key into the current session.

Syntax Description

<i>key</i>	Specifies the key associated with the variables to be imported.
<i>var1 var2</i>	Specifies the variables to be imported. The first variable is mandatory, up to four more optional variables can be specified.

Default

N/A.

Usage Guidelines

The specified key should have created by the user. Also, the variables specified should have been saved using that key.

Attempting to use this command with a non-existent key results in an error message being displayed.



Example

The following example imports the variables “username,” “ipaddr,” and “vlan” from the key “blue:”

```
load var key blue username ipaddr vlan
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

return

return *statusCode*



Note

This is a script command and operates only in scripts or on the command line when scripting is enabled with the following command: `enable cli scripting {permanent}`.

Description

Exits the current script and sets the \$STATUS variable.

Syntax Description

<i>statusCode</i>	Specifies a integer value to which the \$STATUS variable is set.
-------------------	--

Default

N/A.

Usage Guidelines

When used in nested scripts, this command allows you to terminate the current script, set the \$STATUS variable, return to the parent script, and evaluate the \$STATUS variable in the parent script. For more information on the \$STATUS variable, see [Using CLI Scripting](#) in the ExtremeXOS Concepts Guide.

Example

The following example exits the current script and sets the \$STATUS variable to -200:

```
return -200
```



History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

save var key

```
save var key key [var1 var2 ...]
```



Note

This is a script command and operates only in scripts or on the command line when scripting is enabled with the following command: `enable cli scripting {permanent}`.

Description

Saves the specified variables to the specified key.

Syntax Description

<i>key</i>	Specifies the key to which the specified variables are saved.
<i>var1 var2</i>	Specifies the variables to save, The first variable is mandatory, up to four more optional variables can be specified.

Default

N/A.

Usage Guidelines

The variables saved by the SAVE VAR command are represented by the specified key and can be retrieved and restored in the context in which this profile was applied. They are available to rollback events like user-unauthenticate and device-undetected. The key option allows the user to save data for a unique key and retrieve the saved data based on this key. The user is responsible for generating unique keys for each variable. The system has a limited amount of memory to store these variables.

Example

The following example saves the variables “username,” “ipaddr,” and “vlan” to the key “blue:”

```
save var key blue username ipaddr vlan
```



History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

set var

set var *varname* *_expression*



Note

This is a script command and operates only in scripts or on the command line when scripting is enabled with the following command: `enable cli scripting {permanent}`.

Description

Creates and sets the CLI scripting variable to the desired value.

Syntax Description

<i>varname</i>	Specifies the name of the CLI scripting variable. Valid format is \$VARNAME (case insensitive, character string up to 32 characters).
<i>_expression</i>	Specifies the <i>_expression</i> whose value should be evaluated and used to set the variable.

Default

N/A.

Usage Guidelines

The format of a local variable (case insensitive) is: \$VARNAME.

An error message is displayed if the user attempts to use a variable name with a length greater than 32 characters.

If a variable already exists, it is overwritten. No error message is displayed.

Example

The following examples show some ways you can manipulate variables:

```
Set var x 100
Set var x ($x + 2)
```




```
Set var y ($x - 100)
Set var y ($(x) - 100)
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

show var

```
show var {varname}
```



Note

This is a script command and operates only in scripts or on the command line when scripting is enabled with the following command: `enable cli scripting {permanent}`.

Description

Displays the current session variables or the named variable.

Syntax Description

<i>varname</i>	Displays the variable specified, if present.
----------------	--

Default

N/A.

Usage Guidelines

Use this command to see the list of current session variables. The display includes the variable name and value.

Example

The output of this command is similar to the following:

```
Switch.7 # show var
-----
Count : 4
-----
-----
```



variableName	variableValue
-----	-----
CLI.SESSION_TYPE	serial
CLI.USER	admin
STATUS	0
x	66
-----	-----

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

WHILE ... DO

WHILE (**_expression**) DO



Note

This is a script command and operates only in scripts when scripting is enabled with the following command: `enable cli scripting {permanent}`.

Description

Executes a code block while the specified condition is true.

Syntax Description

expression	Specifies the condition for which the statements should be executed while the condition is true.
statements	Set of statements to be executed while the condition is true.

Default

N/A.

Usage Guidelines

This command is usually followed by statements to be executed while the condition is true and the entire construct is terminated by an ENDWHILE command.

The `_expression` must be enclosed in parentheses.



Nesting is supported up to five levels. An error message is displayed if there is incorrect nesting of WHILE conditions. An error message is displayed if a user tries to execute more than five WHILE conditions.

Ctrl-C can be used to break out of a WHILE loop(s). Breaking out of any number of WHILE loops always clears all the WHILE loops .

The operators mentioned in [Using Operators](#) can be used in an `_expression` in a WHILE condition.

You can insert comments by using a number sign (#).

Example

This example creates 10 VLANs, named x1 to x10:

```
set var x 1

WHILE ($x <= 10) DO
    create vlan v$x
    set var x ($x + 1)
ENDWHILE
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.



10 Commands for Configuring LLDP

```
clear lldp neighbors
configure lldp med fast-start repeat-count
configure lldp ports management-address
configure lldp ports port-description
configure lldp ports system-capabilities
configure lldp ports system-description
configure lldp ports system-name
configure lldp ports vendor-specific avaya-extreme call-server
configure lldp ports vendor-specific avaya-extreme dot1q-framing
configure lldp ports vendor-specific avaya-extreme file-server
configure lldp ports vendor-specific avaya-extreme poe-conservation-request
configure lldp ports vendor-specific dot1 port-vlan-ID
configure lldp ports vendor-specific dot1 port-protocol-vlan-ID
configure lldp ports vendor-specific dot1 vlan-name
configure lldp ports vendor-specific dot3 link-aggregation
configure lldp ports vendor-specific dot3 mac-phy
configure lldp ports vendor-specific dot3 max-frame-size
configure lldp ports vendor-specific dot3 power-via-mdi
configure lldp ports vendor-specific med capabilities
configure lldp ports vendor-specific med location-identification
configure lldp ports vendor-specific med policy application
configure lldp ports vendor-specific med power-via-mdi
configure lldp reinitialize-delay
configure lldp snmp-notification-interval
configure lldp transmit-delay
configure lldp transmit-hold
configure lldp transmit-interval
disable lldp ports
disable snmp traps lldp
disable snmp traps lldp-med
enable lldp ports
enable snmp traps lldp
enable snmp traps lldp-med
show lldp
show lldp neighbors
show lldp statistics
unconfigure lldp
```

This chapter describes commands for doing the following:

- Configuring LLDP
- Managing LLDP
- Displaying LLDP information

For an introduction to LLDP, see the *ExtremeXOS Concepts Guide*.



Note

Data Center Bridging Exchange (DCBX) commands that contain the `lldp` keyword are described in [Data Center Solution Commands](#).

clear lldp neighbors

```
clear lldp neighbors [all | port port_list]
```

Description

Clears the LLDP neighbor information collected for one or all ports on the switch.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
------------------	---

Default

N/A.

Usage Guidelines

LLDP neighbor information for each port is automatically cleared after the period defined by the TTL TLV if no update LLDP protocol data unit (LLDPDU) is received. This command immediately clears the LLDP neighbor information for the specified ports.

Example

The following command clears the LLDP information collected for all ports on the switch:

```
clear lldp neighbors all
```

History

This command was first available in ExtremeXOS 12.4.4.



Platform Availability

This command is available on all platforms.

configure lldp med fast-start repeat-count

```
configure lldp med fast-start repeat-count count
```

Description

The fast-start feature is automatically enabled when you enable the LLDP MED capabilities TLV. This command configures how many times, from 1 to 10, the switch sends out an LLDP MED packet with an interval of 1 second.

Syntax Description

<i>count</i>	Specifies the number of times the switch transmits LLDP MED TLVs each second (once it detects a neighbor transmitting LLDP MED TLVs). The range is 1 to 10.
--------------	---

Default

3.

Usage Guidelines

When the switch detects a MED-capable device, this count determines how many times the switch sends a LLDP MED TLVs with an interval of 1 second. The fast-start feature enables the MED-capable device to quickly learn information; this command changes the value from the default 3. The fast-start feature is automatically enabled when you enable the LLDP MED capabilities TLV.

Note



After you configure the LLDP MED capability TLV, the fast-start feature automatically runs. To configure the LLDP MED capability TLV, use the `configure lldp ports [all | port_list] [advertise | no-advertise] vendor-specific med capabilities` command.

Example

The following command configures fast learning on the switch to a value of 2:

```
configure lldp med fast-start repeat-count 2
```

History

This command was first available in ExtremeXOS 11.5.



Platform Availability

This command is available on all platforms.

configure lldp ports management-address

```
configure lldp ports [all | port_list] [advertise | no-advertise] management-address
```

Description

Configures the LLDP port to advertise or not to advertise management address information to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

With ExtremeXOS, you can only add one management address TLV per LLDPDU and the information must be the IP address configured on the management VLAN. If no IP address is assigned to the management VLAN, the system sends the system MAC address. LLDP does not send out IPv6 addresses in this field.

Example

The following command advertises the management address information for port 1:5:

```
configure lldp ports 1:5 advertise management-address
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.



configure lldp ports port-description

```
configure lldp ports [all | port_list] [advertise | no-advertise] port-  
description
```

Description

Configures the LLDP port to advertise or not advertise port description information to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

N/A.

Example

The following command configures port 1:7 to not advertise the port description information to neighbors:

```
configure lldp ports 1:7 no-advertise port-description
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp ports system-capabilities

```
configure lldp ports [all | port_list] [advertise | no-advertise] system-  
capabilities
```



Description

Configures the LLDP port to advertise or not to advertise its system capabilities to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

When at least one VLAN exists with more than two ports, bridging is sent to enabled.

When at least one VLAN on the switch has IP forwarding enabled, the system automatically sets the router bit.

Example

The following command configures all ports to advertise system capability information to neighbors:

```
configure lldp ports all advertise system-capabilities
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp ports system-description

```
configure lldp ports [all | port_list] [advertise | no-advertise] system-  
description
```

Description

Configures the LLDP port to advertise or not to advertise its system description to its neighbors.



Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

Advertise.

Usage Guidelines

Although not mandatory according to the standard, this TLV is included in the LLDPDU by default when you enable LLDP.

When enabled, the system sends the following image (from the show version command) in the system description TLV:

```
ExtremeXOS version 11.2.0.12 v1120b12 by release-manager
on Fri Mar 18 16:01:08 PST 2005
```

Example

The following command configures port 1:4 through port 1:8 to not advertise the system description information to neighbors:

```
configure lldp ports 1:4 - 1:8 no-advertise system-description
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp ports system-name

```
configure lldp ports [all | port_list] [advertise | no-advertise] system-name
```

Default

Configures the LLDP port to advertise or not to advertise its system name to its neighbors.



Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

N/A.

Example

The following command configures port 1:6 to advertise the system name to neighbors:

```
configure lldp ports 1:4 - 1:8 advertise system-name
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific avaya-extreme call-server

The Avaya phone uses this proprietary LLDP TLV to learn the IP address(es) of the call server(s) to use.

```
configure lldp ports [all | port_list] [advertise | no-advertise] vendor-specific  
avaya-extreme call-server ip_address_1 {ip_address_2 {ip_address_3 {ip_address_4  
{ip_address_5 {ip_address_6 {ip_address_7 {ip_address_8}}}}}}}}}
```


Description

Configures the LLDP port to advertise or not advertise up to 8 call server IP addresses to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.



advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.
ip_address_1 8	Specifies IP address of up to 8 call servers.
<div style="display: flex; align-items: flex-start;"> <div style="margin-right: 10px;">  </div> <div> <p>Note NOTE: This parameter does not apply when you configure the no-advertise parameter.</p> </div> </div>	

Default

No advertise.

Usage Guidelines

The Avaya phone uses this proprietary LLDP TLV for addressing information. You can configure the IP address for up to 8 call servers in a single TLV.

Example

The following command configures ports 1-5 to advertise two call server IP addresses to neighbors:

```
configure lldp ports 1-5 advertise vendor-specific avaya-extreme call-server
10.10.10.10 10.11.10.10
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific avaya-extreme dot1q-framing

```
configure lldp ports [all | port_list] [advertise | no-advertise] vendor-specific
avaya-extreme dot1q-framing [tagged | untagged | auto]
```

Description

Configures the LLDP port to advertise or not advertise the 802.1q framing configuration to its neighbors. The Avaya phone uses this proprietary LLDP TLV information. In addition to this LLDP TLV, you must enable LLLDP as well as configure both the LLDP MED capabilities TLV and the LLDP network policy TLV.



Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.
tagged	Specifies to use tagging. NOTE: This parameter applies only when you use the advertise parameter.
untagged	Specifies not to use tagging. NOTE: This parameter applies only when you use the advertise parameter.
auto	Specifies following a predetermined sequence (see Usage Guidelines below). NOTE: This parameter applies only when you use the advertise parameter.

Default

No advertise.

Usage Guidelines

Before configuring this LLDP TLV, you must take the following steps:

- Enable LLDP using the `enable lldp ports` command.
- Enable the LLDP MED capabilities TLV using the `configure lldp ports vendor-specific med capabilities` command.
- Enable the LLDP MED network policy TLV using the `configure lldp ports vendor-specific med policy application` command.

This TLV is used to exchange information about Layer 2 priority tagging between the network connectivity device (switch) and the Avaya phone.

If you configure the TLV to advertise tagging, the phone uses tagging information, which it retrieves from the `configure lldp ports vendor-specific med policy application` command. If you configure the TLV to advertise untagged, the phone does not use any tagging, including 802.1q priority tagging.

If you configure the TLV to advertise auto, the phone cycles through the following sequence until an action is successful:

- Uses the configuration advertised by the LLDP MED network policy TLV, as configured by the `configure lldp ports vendor-specific med policy application` command.
- Uses the priority tagged frames configured by the phone's server.
- Sends the traffic untagged.



Example

The following command configures all ports to advertise the dot1q framing as untagged to neighbors:

```
configure lldp ports all advertise vendor-specific avaya-extreme dot1q-
framing untagged
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific avaya-extreme file-server

```
configure lldp ports [all | port_list] [advertise | no-advertise] vendor-specific
avaya-extreme file-server ip_address_1 {ip_address_2 {ip_address_3
{ip_address_4}}
```

Description

Configures the LLDP port to advertise or not advertise up to 4 file server IP addresses to its neighbors. The Avaya phone uses this proprietary LLDP TLV to learn the IP address(es) of the file server(s) to use.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.
<i>ip_address_1...4</i>	Specifies IP address of up to 4 file servers. NOTE: This parameter does not apply when you configure the no-advertise parameter.

Default

No advertise.

Usage Guidelines

The Avaya phone uses this proprietary LLDP TLV for addressing information. You can configure the IP address for up to 4 file servers in a single TLV.



Example

The following command configures all ports to advertise two file server IP addresses to neighbors:

```
configure lldp ports 1-5 advertise vendor-specific avaya-extreme call-server
10.20.10.10 10.12.10.10
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific avaya-extreme poe-conservation-request

```
configure lldp ports [all | port_list] [advertise | no-advertise] vendor-specific
avaya-extreme poe-conservation-request
```

Description

Configures the LLDP port to advertise or not advertise a requested conservation level. By default, the requested conservation value on this proprietary LLDP TLV is 0, which is no power conservation. This LLDP TLV is sent out only on PoE-capable Ethernet ports.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

The switch sends this proprietary LLDP TLV to request a PD to go into a certain power conservation level or request the PD to go to the maximum conservation level. This LLDP TLV is transmitted only on PoE-capable ports.



When configured to advertise, the switch sends this TLV with a requested conservation power level of 0, which requests no power conservation. To temporarily change this conservation level, use the SNMP `IldpXAvExLocPortXPoEPSEPortReqLevel` object to set a new value; the reconfigured value is not saved over a reboot. (This SNMP object can be set from 0 to 243 or 255.)

Example

The following command configures all ports to advertise the currently requested conservation level to neighbors:

```
configure lldp ports all advertise vendor-specific avaya-extreme poe-
conservation-request
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific dot1 port-vlan-ID

```
configure lldp ports [all | port_list] [advertise | no-advertise] vendor-specific
dot1 port-vlan-ID
```

Description

Configures the LLDP port to advertise or not advertise port vlan ID information to its neighbors. This allows a VLAN bridge port to advertise the port VLAN identifier that is associated with untagged or priority-tagged frames.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.



Usage Guidelines

The port VLAN ID TLV allows the port to transmit the VLAN ID associated with untagged VLANs. There can be only one port VLAN ID in each LLDPDU.

If no untagged VLANs are configured on the specified port, the TLV is not added to the LLDPDU, even if you configured this to advertise.

Example

The following command configures all ports to advertise port vlan ID information to neighbors:

```
configure lldp ports all advertise vendor-specific dot1 port-vlan-ID
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific dot1 port-protocol-vlan-ID

```
configure lldp ports [all | port_list] [advertise | no-advertise] vendor-specific
dot1 port-protocol-vlan-ID {vlan [all | vlan_name]}
```

Description

Configures the LLDP port to advertise or not advertise port VLAN information to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.
all	Specifies all VLANs on the port.
<i>vlan_name</i>	Specifies the VLAN on the port that you want to advertise.

Default

No advertise.



Usage Guidelines

When configured to advertise, the switch inserts a port and protocol VLAN ID TLV for each VLAN configured on the ports. The port and protocol VLAN ID TLV allows the port to advertise if it supports protocol and/or tagged VLANs, along with the associated tagged values. A separate TLV is sent for each VLAN that you want to advertise.

By default, once you configure this TLV, the system sends all protocol-based VLANs on the port. However, the LLDPDU cannot exceed 1500 bytes, so you should configure the port to advertise only the specified VLANs.



Note

The total LLDPDU size is 1500 bytes; any TLVs after that limit are dropped.

This TLV does not send information on the type of protocol that the VLAN has enabled; it just says whether the port is enabled or disabled for protocol-based VLANs. As Extreme Networks devices are always capable of supporting protocol-based VLANs, once you configure this TLV, the system always advertises support for these VLANs.

Example

The following command configures all ports to advertise port and protocol VLAN information to neighbors for all VLANs on all ports:

```
configure lldp ports all advertise vendor-specific dot1 port-protocol-vlan-id
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific dot1 vlan-name

```
configure lldp ports [all | port_list] [advertise | no-advertise] vendor-specific  
dot1 vlan-name {vlan [all | vlan_name]}
```

Description

Configures the LLDP port to advertise or not advertise VLAN name information to its neighbors. Use this TLV to advertise information for the tagged VLANs you want to specify on the port. This allows an IEEE 802.1Q-compatible 802 LAN station to advertise the assigned name of any VLAN with which it is configured.



Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.
vlan	Specifies all VLANs on the port.
<i>vlan_name</i>	Specifies the VLAN on the port that you want to advertise.

Default

No advertise.

Usage Guidelines

The VLAN name TLV sends the VLAN name and the tag used; it associates a name to a tag for the specified VLAN. This allows an IEEE 802.1Q-compatible 802 LAN station to advertise the assigned name of any VLAN with which it is configured.

You can enable this TLV for tagged and untagged VLANs. When you enable this TLV for tagged VLANs, the TLV advertises the IEEE 802.1Q tag for that VLAN. (For untagged VLANs, the internal tag is advertised.) You can specify exactly which VLANs to advertise.

When configured to advertise, the switch inserts a VLAN name TLV for every VLAN configured on the ports. By default, once you configure this TLV, the system sends all VLAN names on the port. However, each VLAN name can require up to 32 bytes and the LLDPDU cannot exceed 1500 bytes, so you should configure the port to advertise only the specified VLANs, using the keyword `vlan_name`.



Note

The total LLDPDU size is 1500 bytes; any TLVs after that limit are dropped.

Example

The following command configures all ports to not advertise VLAN name information to neighbors:

```
configure lldp ports all no-advertise vendor-specific dot1 vlan-name
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.



configure lldp ports vendor-specific dot3 link-aggregation

```
configure lldp ports [all | port_list] [advertise | no-advertise] vendor-specific dot3 link-aggregation
```

Description

Configures the LLDP port to advertise or not advertise link-aggregation capabilities to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

When configured, this TLV is added to each LLDP port LLDPDU indicating the link-aggregation capabilities, status, and value of the master port of the load-sharing group.

Example

The following command configures port 1:12 to not advertise link-aggregation capabilities to neighbors:

```
configure lldp ports 1:12 no-advertise vendor-specific dot3 link-aggregation
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific dot3 mac-phy

```
configure lldp ports [all | port_list] [advertise | no-advertise] vendor-specific dot3 mac-phy
```



Description

Configures the LLDP port to advertise or not advertise MAC and physical layer capabilities to its neighbors. The capabilities include duplex and bit rate.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

When configured, the system advertises information about the speed capabilities, as well as autonegotiation support and status, of the LLDP port.

Example

The following command configures all ports to advertise MAC/PHY capabilities to neighbors:

```
configure lldp ports all advertise vendor-specific dot3 mac-phy
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific dot3 max-frame-size

```
configure lldp ports [all | port_list] [advertise | no-advertise] vendor-specific  
dot3 max-frame-size
```

Description

Configures the LLDP port to advertise or not advertise its maximum frame size to its neighbors.



Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

When jumbo frames are not enabled on the specified port, the TLV reports a value of 1518 once you configure it to advertise. If jumbo frames are enabled, the TLV inserts the configured value for the jumbo frames.

Example

The following command configures ports 1:12 and 1:13 to advertise the maximum frame size to neighbors:

```
configure lldp ports 1:12 - 1:13 advertise vendor-specific dot3 max-frame-size
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific dot3 power-via-mdi

```
configure lldp ports [all | port_list] [advertise | no-advertise] vendor-specific dot3 power-via-mdi {with-classification}
```

Description

Configures the LLDP port to advertise or not advertise Power over Ethernet (PoE) capabilities to its neighbors.



Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.
with-classification	Specifies to use LLDP for Data Link Layer Classification. This option is available only on PoE+ ports.

Default

No advertise.

Usage Guidelines

When configured, the system includes this TLV. We recommend enabling this TLV only on PoE-capable ports.

The following information is transmitted for LLDP ports with this TLV:

- Support PoE or not
- Port class
 - Power sourcing equipment (PSE)
 - Powered device (PD)
- Power pairs used to supply power
 - Signal
 - Spare
- Power status
- Support pairs control or not
- Power class
 - Class0
 - Class1
 - Class2
 - Class2
 - Class3
 - Class4

Data link layer classification allows fine-grained dynamic re-allocation of power based on changing needs. This feature is enabled by enabling LLDP (transmit and receive) and configuring transmission of the power-via-MDI TLV. The ExtremeXOS software sends an LLDPDU containing a power-via-MDI TLV within 10 seconds of DLL classification being enabled. A PD may request a new power value using an LLDPDU. The allocated power might be changed if a request is received and approved after a power review. The software responds with an allocated power value within 10 seconds of receipt of an LLDPDU with a different requested power from a PD. Power allocation can be controlled to a



granularity of 0.1 watts. When DLL classification is enabled, it takes precedence over physical classification.



Note

For more information on advertising power support, see the `configure lldp ports vendor-specific med power-via-mdi` command.

Example

The following command configures all ports to advertise power capabilities to neighbors:

```
configure lldp ports all advertise vendor-specific dot3 power-via-mdi
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific med capabilities

```
configure lldp ports [all | port_list] [advertise | no-advertise] vendor-specific med capabilities
```

Description

Configures the LLDP port to advertise or not advertise MED capabilities. This TLV must be enabled before any of the other MED TLVs can be enabled. Also, this TLV must be set to no-advertise after all other MED TLVs are set to no-advertise.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.



Usage Guidelines

This command enables the LLDP media endpoint discovery (MED) capabilities TLV, which allows LLDP-MED network connectivity devices to definitively determine that particular endpoints support LLDP MED, and if so, to discover which LLDP MED TLVs the particular endpoint devices are capable of supporting and to which specific device class the device belongs to.

This TLV must be enabled before any of the other MED TLVs can be enabled; and this TLV must be set to no-advertise after all other MED TLVs are set to no-advertise.

As with all the LLDP MED TLVs, the switch sends this TLV only after it detects a MED-capable device on the port. The switch does not automatically send this TLV after it is enabled; the switch must first detect a MED-capable device on the port.



Note

Network connectivity devices wait to detect LLDP MED TLVs from endpoints before they send out LLDP MED TLVs; so L2 network connectivity devices do not exchange LLDP MED messages.

The following information is included in the LLDP MED capabilities TLV when it is transmitted:

- The supported LLDP MED TLVs—For Extreme Networks devices, these are capabilities, network policy, location, and extended power (extended power only advertised only on PoE-capable ports).
- The MED device type—For Extreme Networks devices, this is advertised as a network connectivity device (set to 4).

Example

The following command configures all ports to advertise MED capabilities to neighbors:

```
configure lldp ports all advertise vendor-specific med capabilities
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific med location-identification

```
configure lldp ports [all | port_list] [advertise | no-advertise] vendor-specific med location-identification [coordinate-based hex_value | civic-based hex_value | ecs-elin elin]
```



Description

Configures the LLDP port to advertise or not advertise MED location information. You configure up to 3 different location identifiers.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies to not send the information to neighbors.
coordinate-based	Specifies using the coordinate-based location identifier. This value is exactly 16 bytes long; see RFC 3825 for details.
<i>hex_value</i>	Enter a hexadecimal value with each byte separated by a colon. Or, you can obtain this value from a network management application. NOTE: This parameter is not used when the no-advertise parameter is configured.
civic-based	Specifies using the civic-based location identifier. This value must have a minimum length of 6 bytes; see RFC3825 for details.
ecs-elin	Specifies using the ecs location identifier. (Emergency Call Service, as defined in the TIA-TSB-146.)
<i>elin</i>	Enter a numerical string; the range is 10 to 25 characters. Or, you can obtain this value from a network management application. (See the TIA-TSB-146 standard for a definition of these numbers; also, the network management application must be able to handle the LLDP MED MIB.) NOTE: This parameter is not used when the no-advertise parameter is configured.

Default

No advertise.

Usage Guidelines

You might need to use a specific format for your specific VoIP implementation; see the VoIP manufacturer's manual for details.

You must configure the LLDP MED capabilities TLV before configuring this TLV. Configure the LLDP MED capabilities TLV using the `configure lldp ports [all | port_list] [advertise | no-advertise] vendor-specific med capabilities` command.

As with all the LLDP MED TLVs, the switch sends this TLV only after it detects a MED-capable device on the port. The switch does not automatically send this TLV after it is enabled; the switch must first detect a MED-capable device on the port.



Example

The following command configures all ports to advertise MED location information to neighbors using the ECS format:

```
configure lldp ports all advertise vendor-specific med location-
identification ecs-elin 423233455676
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific med policy application

```
configure lldp ports [all | port_list] [advertise | no-advertise] vendor-specific
med policy application [voice | voice-signaling | guest-voice | guest-voice-
signaling | softphone-voice | video-conferencing | streaming-video | video-
signaling] vlan vlan_name dscp dscp_value {priority-tagged}
```

Description

Configures the LLDP port to advertise or not advertise MED network policy TLVs. This TLV advertises VLAN configuration and associated Layer 2 and Layer 3 attributes that apply for a set of specific applications on that port. You can advertise up to 8 TLVs, each for a specific application, per port/VLAN. Each application type can exist only once per port. This TLV tells the endpoint the specific VLAN to use for the specific application, along with its unique priority.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.
advertise	Specifies to send the information to neighbors.
voice	Specifies voice application on specified port/VLAN(s).
voice-signaling	Specifies voice signaling application on specified port/VLAN(s).
guest-voice	Specifies guest voice application on specified port/VLAN(s).
guest-voice- signaling	Specifies guest voice signaling application on specified port/VLAN(s).
softphone-voice	Specifies soft phone voice application on specified port/VLAN(s).



video-conferencing	Specifies videoconferencing application on specified port/VLAN(s).
streaming-video	Specifies streaming video application on specified port/VLAN(s).
video-signaling	Specifies video signaling application on specified port/VLAN(s).
<i>vlan_name</i>	Specifies the VLAN the specified application is using. NOTE: This parameter does not apply when the no-advertise parameter is configured.
<i>dscp_value</i>	Specifies the DSCP value for the specified application. This is a 6-bit value from 0 to 63. NOTE: This parameter does not apply when the no-advertise parameter is configured.
priority-tagged	Use this if you want priority tagging, and the VLAN is configured as untagged on the port. (The endpoint sends out frames for the specified application with a tag of 0.) NOTE: This parameter does not apply when the no-advertise parameter is configured.

Default

No advertise.

Usage Guidelines

This command enables the LLDP MED network policy TLV, which allows network connectivity devices and endpoint devices to advertise VLAN configuration and associated Layer 2 and Layer 3 attributes that apply for a set of specific application on that port. This TLV can be enabled on a per port/VLAN basis. Each application type can exist only once on a port.

You can enable the transmission of a TLV policy for each application. A maximum of 8 TLVs can be enabled, and each can have a unique DSCP value and/or priority tagging.

You must configure the LLDP MED capabilities TLV before configuring this TLV. Configure the LLDP MED capabilities TLV using the `configure lldp ports [all | port_list] [advertise | no-advertise] vendor-specific med capabilities` command.

As with all the LLDP MED TLVs, the switch sends this TLV only after it detects a MED-capable device on the port. The switch does not automatically send this TLV after it is enabled; the switch must first detect a MED-capable device on the port.

The following information is transmitted for LLDP ports with this TLV:

- Application type
 - Used as configured.
- Unknown policy flag
 - Set to 0.
- Tagged flag
 - Set to tagged for tagged VLANs; set to untagged for untagged VLANs. By default, set to 0.
- VLAN ID
 - Copied from the VLAN. However, if you configure the priority-tagged parameter, this value is set to 0.
- Layer 2 priority



Copied from the VLAN priority.

- DSCP value

Uses the value configured in the dscp parameter.



Note

See the documentation provided by the manufacturer of connected devices regarding values.

Example

The following command configures all ports to advertise videoconferencing on the VLAN video with a DSCP of 7 to neighbors:

```
configure lldp ports all advertise vendor-specific med policy application
video-conferencing vlan video dscp 7
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

configure lldp ports vendor-specific med power-via-mdi

```
configure lldp ports [all | port_list] [advertise | no-advertise] vendor-specific
med power-via-mdi
```

Description

Configures the LLDP port to advertise or not advertise MED power requirement details. This TLV can only be enabled on a PoE-capable port and is used for advanced power management between the MED network connectivity and endpoint devices.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.



Default

No advertise.

Usage Guidelines

When enabled, this LLDP MED TLV advertises fine-grained power requirement details about PoE settings and support. This TLV can be enabled only on a PoE-capable port; the switch returns an error message if this TLV is configured for a non-PoE-capable port.

You must configure the LLDP MED capabilities TLV before configuring this TLV. Configure the LLDP MED capabilities TLV using the `configure lldp ports [all | port_list] [advertise | no-advertise] vendor-specific med capabilities` command.

As with all the LLDP MED TLVs, the switch sends this TLV only after it detects a MED-capable device on the port. The switch does not automatically send this TLV after it is enabled; the switch must first detect a MED-capable device on the port.



Note

For additional information on power support, see the `configure lldp ports vendor-specific dot3 power-via-mdi` command.

The following information is transmitted for LLDP MED PoE-capable ports with this TLV:

- Power type
Set to PSE.
- Power source
Set to primary power source.
- Power priority
Taken from PoE port configuration.
- Power value
Taken from PoE port configuration.

Example

The following command configures all ports to advertise MED power information to neighbors:

```
configure lldp ports all advertise vendor-specific med power-via-mdi
```

History

This command was first available in ExtremeXOS 11.5.



Platform Availability

This command is available on all platforms.

configure lldp reinitialize-delay

```
configure lldp reinitialize-delay seconds
```

Description

Configures the delay before the receive state machine is reinstalled once the LLDP transmit mode has been disabled.

Syntax Description

<i>seconds</i>	Specifies the delay that applies to the reinitialization attempt. The range is 1 to 10 seconds.
----------------	---

Default

2 seconds.

Usage Guidelines

N/A.

Example

The following command configures a reinitialization delay of 10 seconds:

```
configure lldp reinitialize-delay 10
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp snmp-notification-interval

```
configure lldp snmp-notification-interval seconds
```



Description

Configures the allowed interval at which Simple Network Management Protocol (SNMP) notifications are sent.

Syntax Description

<i>seconds</i>	Specifies the interval at which LLDP SNMP notifications are sent. The range is 5 to 3600 seconds.
----------------	---

Default

5 seconds.

Usage Guidelines

This is a global timer. If one port sends a notification, no notifications for other ports go out for the configured interval.

Example

The following command configures an interval of 60 seconds for LLDP SNMP notifications:

```
configure lldp snmp-notification-interval 60
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp transmit-delay

```
configure lldp transmit-delay [ auto | seconds]
```

Description

Configures the delay time between successive frame transmissions initiated by a value change or status change in any of the LLDP local systems Management Information Base (MIB).

The auto option uses a formula ($0.25 * \text{transmit-interval}$) to calculate the number of seconds.



Syntax Description

auto	Uses the formula $(0.25 * \text{transmit-interval})$ to calculate the seconds.
<i>seconds</i>	Specifies the interval at which LLDP notifications are sent. The range is 1 to 8291.

Default

2 seconds.

Usage Guidelines

This is the timer between triggered updates.

Example

The following command configures the delay between LLDP frame transmissions for triggered updates to be automatically calculated:

```
configure lldp transmit-delay auto
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp transmit-hold

```
configure lldp transmit-hold hold
```

Description

Calculates the actual time-to-live (TTL) value used in the LLDPDU messages.

The formula is $\text{transmit-interval} * \text{transmit-hold}$; by default the TTL value is $(30 * 4) 120$ seconds.

Syntax Description

<i>hold</i>	Used to calculate the TTL value; the range is 2 to 10.
-------------	--

Default

4.



Usage Guidelines

N/A.

Example

The following command configures the transmit-hold value (which is used to calculate the TTL of the LLDP packets) to 5:

```
configure lldp transmit-hold 5
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure lldp transmit-interval

```
configure lldp transmit-interval seconds
```

Description

Configures the periodic transmittal interval for LLDPDUs.

Syntax Description

<i>seconds</i>	Specifies the time between LLDPDU transmissions. The range is 5 to 32768.
----------------	---

Default

30 seconds.

Usage Guidelines

N/A.

Example

The following command configures a transmittal interval of 20 seconds for LLDPDUs.

```
configure lldp transmit-interval 20
```



History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

disable lldp ports

```
disable lldp ports [all | port_list] {receive-only | transmit-only}
```

Description

Disables LLDP transmit mode, receive mode, or transmit and receive mode on the specified port or ports.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
receive-only	Specifies that only the receive mode for LLDP is disabled.
transmit-only	Specifies that only the transmit mode for LLDP is disabled.

Default

Disabled.

Usage Guidelines

If you do not specify an option, both LLDP modes (transmit and receive) are disabled.

Example

The following example disables the LLDP receive mode on ports 1:2 to 1:6.

```
disable lldp ports 1:2-1:6 receive-only
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.



disable snmp traps lldp

```
disable snmp traps lldp {ports [all | port_list]}
```

Description

Disables the sending of LLDP-specific SNMP traps on the specified port or ports.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Disabled.

Usage Guidelines

If you do not specify any ports, the system stops sending LLDP traps from all ports on the switch.

Example

The following example disables sending LLDP SNMP traps on all switch ports:

```
disable snmp traps lldp ports all
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

disable snmp traps lldp-med

```
disable snmp traps lldp-med {ports [all | port_list]}
```

Description

Disables the sending of LLDP MED-specific SNMP traps on the specified port or ports.



Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Disabled.

Usage Guidelines

If you do not specify any ports, the system stops sending LLDP MED traps from all ports on the switch.

Example

The following example disables sending LLDP MED SNMP traps on all switch ports:

```
disable snmp traps lldp-med ports all
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

enable lldp ports

```
enable lldp ports [all | port_list] {receive-only | transmit-only}
```

Description

Enables LLDP transmit mode, receive mode, or transmit and receive mode. If the transmit-only or receive-only option is not specified, both transmit and receive modes are enabled.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
receive-only	Specifies that the port only receives LLDP messages.
transmit-only	Specifies that the port only transmits LLDP messages.



Default

Disabled.

Usage Guidelines

If you do not specify an option, the port is enabled to both transmit and receive LLDP messages.

Once the port is enabled for LLDP in one mode and you issue another `enable lldp ports` command for another mode, that second mode replaces the original mode. For example, you might originally enable several ports to only receive LLDP messages and then want those ports to both receive and transmit LLDP messages. In that case, you issue the `enable lldp ports` command with no variables (and the receive-and-transmit mode replaces the receive-only mode).

To verify the port setting for LLDP, use the `show lldp {port [all |port_list]} {detailed}` command.

Example

The following example enables LLDP transmit and receive mode on port 1:4.

```
enable lldp port 1:4
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable snmp traps lldp

```
enable snmp traps lldp {ports [all | port_list]}
```

Description

Enables the transmission of LLDP SNMP trap notifications.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.



Default

Disabled.

Usage Guidelines



Note

To enable SNMP traps for LLDP MED TLVs, you must issue a separate command; use the `enable snmp traps lldp-med {ports [all | port_list]}`.

If you do not specify any ports, the system sends LLDP traps for all ports.



Note

The Avaya-Extreme proprietary TLVs do not send traps.

Example

The following command enables LLDP SNMP traps for all ports:

```
enable snmp traps lldp ports all
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable snmp traps lldp-med

```
enable snmp traps lldp-med {ports [all | port_list]}
```

Description

Enables the transmission of LLDP SNMP trap notifications related to LLDP MED extension TLVs.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Disabled.



Usage Guidelines

If you do not specify any ports, the system sends LLDP-MED traps for all ports.

Example

The following command enables LLDP-MED SNMP traps for all ports:

```
enable snmp traps lldp-med ports all
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

show lldp

```
show lldp {port [all | port_list]} {detailed}
```

Description

Displays LLDP configuration information for the specified port or ports.

Use the detailed keyword to display the configured VLANs on the port and the enabled VLAN-specific TLVs.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
detailed	Shows information on the configured VLANs on the port.

Default

N/A.

Usage Guidelines

Use the detailed variable to display information regarding configured VLANs on the ports and any enabled VLAN-specific TLVs.



Example

The following example displays LLDP configuration information for the switch:

```
# show lldp
LLDP transmit interval          : 30 seconds
LLDP transmit hold multiplier   : 4 (used TTL = 120 seconds)
LLDP transmit delay             : 2 seconds
LLDP SNMP notification interval : 5 seconds
LLDP reinitialize delay         : 2 seconds
LLDP-MED fast start repeat count : 4
LLDP Port Configuration:
Port    Rx      Tx      SNMP      --- Optional enabled transmit TLVs
--
Mode    Mode    Notification LLDP  802.1  802.3  MED  AvEx  DCBX
=====
=
1       Enabled Enabled  --      --D--  ---   ----  ----  IB
2       Enabled Enabled  --      --D--  ---   ----  ----  IB
=====
=
Notification: (L) lldpRemTablesChange, (M) lldpXMedTopologyChangeDetected
LLDP Flags   : (P) Port Description, (N) System Name, (D) System Description
              (C) System Capabilities, (M) Mgmt Address
802.1 Flags  : (P) Port VLAN ID, (p) Port & Protocol VLAN ID, (N) VLAN Name
802.3 Flags  : (M) MAC/PHY Configuration/Status, (P) Power via MDI
              (+) Power via MDI with DLL Classification for PoE+,
              (L) Link Aggregation, (F) Frame Size
MED Flags    : (C) MED Capabilities, (P) Network Policy,
              (L) Location Identification, (p) Extended Power-via-MDI
AvEx Flags   : (P) PoE Conservation Request, (C) Call Server, (F) File Server
              (Q) 802.1Q Framing
DCBX Flags   : (I) IEEE 802.1Qaz DCBX, (B) Baseline v1.01 DCBX
```

The following example includes detailed information on the LLDP configuration for port 1:1:

```
# show lldp port 1:1 detailed
LLDP transmit interval          : 30 seconds
LLDP transmit hold multiplier   : 4 (used TTL = 120 seconds)
LLDP transmit delay             : 2 seconds
LLDP SNMP notification interval : 5 seconds
LLDP reinitialize delay         : 2 seconds
LLDP-MED fast start repeat count : 4
LLDP Port Configuration:
Port    Rx      Tx      SNMP      --- Optional enabled transmit TLVs
--
Mode    Mode    Notification LLDP  802.1  802.3  MED  AvEx  DCBX
=====
=
1:1     Enabled Enabled  --      --D--  ---   ----  CLP-  ----  IB
VLAN: Default
VLAN: voice
AvEx Call-Server: IP Address(es)=10.0.0.20, 10.0.0.21
AvEx File-Server: IP Address(es)=10.0.0.20, 10.0.0.21, 10.0.0.22
AvEx 802.1Q Framing: Mode=tagged
MED LCI: Location Format=ECS ELIN based
```



```

1234567890
MED Policy: Application=voice
VLAN=voice, DSCP=40
DCBX: Priority 4, iSCSI
DCBX: Priority 3, FCoE
DCBX: Priority 3, FIP
=====
=
Notification: (L) lldpRemTablesChange, (M) lldpXMedTopologyChangeDetected
LLDP Flags : (P) Port Description, (N) System Name, (D) System Description
(C) System Capabilities, (M) Mgmt Address
802.1 Flags : (P) Port VLAN ID, (p) Port & Protocol VLAN ID, (N) VLAN Name
802.3 Flags : (M) MAC/PHY Configuration/Status, (P) Power via MDI
(L) Link Aggregation, (F) Frame Size
MED Flags : (C) MED Capabilities, (P) Network Policy,
(L) Location Identification, (p) Extended Power-via-MDI
AvEx Flags : (P) PoE Conservation Request, (C) Call Server, (F) File Server
(Q) 802.1Q Framing
DCBX Flags : (I) IEEE 802.1Qaz DCBX, (B) Baseline v1.01 DCBX

```

History

This command was first available in ExtremeXOS 11.2.

The information on fast-start repeat count, MED, AvEx, and notification was added in ExtremeXOS 11.5.

An additional flag was added for PoE+ in ExtremeXOS 12.5.

The display was updated for DCBX in ExtremeXOS 12.6.

Platform Availability

This command is available on all platforms.

show lldp neighbors

```
show lldp {port [all | port_list]} neighbors {detailed}
```

Description

Displays the information related to the LLDP neighbors detected on the specified port or ports.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
detailed	Shows detailed information on the neighbors.



Default

N/A.

Usage Guidelines

You must use the detailed parameter to display detailed information about the received LLDP TLVs.

Example

The following example displays LLDP neighbor information for all switch ports:

```
# show lldp port all neighbors
Port      Neighbor Chassis ID      Neighbor Port ID      TTL      Age
=====
1:2       00:04:96:26:A4:70      1:1                    120      7
2:6       (5.1)10.201.41.146     00:04:0D:EC:EA:5C     120      3
2:7       (5.1)10.201.41.147     00:04:0D:ED:41:9B     120      3
2:10     00:01:30:F9:9E:80      8:10                   120     15
=====
NOTE: The Chassis ID and/or Port ID might be truncated to fit the screen.
```

The following command lists detailed LLDP neighbor information for all switch ports:

```
# show lldp all neighbors detailed
-----
LLDP Port 1:2 detected 1 neighbor
Neighbor: 00:04:96:26:A4:70/1:1, age 12 seconds
- Chassis ID type: MAC address (4)
Chassis ID      : 00:04:96:26:A4:70
- Port ID type: ifName (5)
Port ID        : "1:1"
- Time To Live: 120 seconds
- System Description: "ExtremeXOS version 12.0.0.6 v1200b6 by release-ma\
nager on Mon Mar 19 00:37:59 PDT 2007"
-----
LLDP Port 2:6 detected 1 neighbor
Neighbor: (5.1)10.201.41.146/00:04:0D:EC:EA:5C, age 8 seconds
- Chassis ID type: Network address (5); Address type: IPv4 (1)
Chassis ID      : 10.201.41.146
- Port ID type: MAC address (3)
Port ID        : 00:04:0D:EC:EA:5C
- Time To Live: 120 seconds
- System Name: "AVAECEA5C"
- System Capabilities : "Bridge, Telephone"
Enabled Capabilities: "Bridge, Telephone"
- Management Address Subtype: IPv4 (1)
Management Address      : 10.201.41.146
Interface Number Subtype : System Port Number (3)
Interface Number        : 1
Object ID String         : "1.3.6.1.4.1.6889.1.69.2.3"
- IEEE802.3 MAC/PHY Configuration/Status
Auto-negotiation        : Supported, Enabled (0x03)
Operational MAU Type    : 100BaseTXFD (16)
```



```

- MED Capabilities: "MED Capabilities, Network Policy, Inventory"
MED Device Type : Endpoint Class III (3)
- MED Network Policy
Application Type : Voice (1)
Policy Flags      : Known Policy, Tagged (0x1)
VLAN ID          : 0
L2 Priority       : 6
DSCP Value       : 46
- MED Hardware Revision: "9650D01A"
- MED Firmware Revision: "hb96xxual_20r30s.bin"
- MED Software Revision: "ha96xxual_20r30s.bin"
- MED Serial Number: "06N537900335"
- MED Manufacturer Name: "Avaya"
- MED Model Name: "9650"
- Avaya/Extreme Conservation Level Support
Current Conservation Level: 0
Typical Power Value      : 0.0 Watts
Maximum Power Value     : 0.0 Watts
Conservation Power Level : 1=0.0W
- Avaya/Extreme Call Server(s): 69.26.36.53
- Avaya/Extreme IP Phone Address: 10.201.41.146 255.255.255.0
Default Gateway Address   : 10.201.41.1
- Avaya/Extreme CNA Server: 0.0.0.0
- Avaya/Extreme File Server(s): 10.201.41.36
- Avaya/Extreme IEEE 802.1q Framing: Tagged
-----
LLDP Port 2:7 detected 1 neighbor
Neighbor: (5.1)10.201.41.147/00:04:0D:ED:41:9B, age 8 seconds
- Chassis ID type: Network address (5); Address type: IPv4 (1)
Chassis ID      : 10.201.41.147
- Port ID type: MAC address (3)
Port ID        : 00:04:0D:ED:41:9B
- Time To Live: 120 seconds
- System Name: "AVAED419B"
- System Capabilities : "Telephone"
Enabled Capabilities: "Telephone"
- Management Address Subtype: IPv4 (1)
Management Address   : 10.201.41.147
Interface Number Subtype : System Port Number (3)
Interface Number     : 1
Object ID String      : "1.3.6.1.4.1.6889.1.69.2.5"
- IEEE802.3 MAC/PHY Configuration/Status
Auto-negotiation     : Supported, Enabled (0x03)
Operational MAU Type  : 100BaseTXFD (16)
- MED Capabilities: "MED Capabilities, Network Policy, Inventory"
MED Device Type : Endpoint Class III (3)
- MED Network Policy
Application Type : Voice (1)
Policy Flags      : Known Policy, Tagged (0x1)
VLAN ID          : 0
L2 Priority       : 6
DSCP Value       : 46
- MED Hardware Revision: "9610D01A"
- MED Firmware Revision: "hb96xxual_20r30s.bin"
- MED Software Revision: "ha96xxual_20r30s.bin"
- MED Serial Number: "06N538825133"
- MED Manufacturer Name: "Avaya"
- MED Model Name: "9610"

```



```

- Avaya/Extreme Conservation Level Support
Current Conservation Level: 0
Typical Power Value      : 0.0 Watts
Maximum Power Value      : 0.0 Watts
Conservation Power Level : 1=0.0W
- Avaya/Extreme Call Server(s): 69.26.36.53
- Avaya/Extreme IP Phone Address: 10.201.41.147 255.255.255.0
Default Gateway Address  : 10.201.41.1
- Avaya/Extreme CNA Server: 0.0.0.0
- Avaya/Extreme File Server(s): 10.201.41.36
- Avaya/Extreme IEEE 802.1q Framing: Tagged
-----
LLDP Port 2:10 detected 1 neighbor
Neighbor: 00:01:30:F9:9E:80/8:10, age 20 seconds
- Chassis ID type: MAC address (4)
Chassis ID      : 00:01:30:F9:9E:80
- Port ID type: ifName (5)
Port ID         : "8:10"
- Time To Live: 120 seconds
- System Description: "ExtremeXOS version 12.0.0.6 v1200b6 by release-ma\
nager on Mon Mar 19 00:43:19 PDT 2007"

```

History

This command was first available in ExtremeXOS 11.2. Information on the LLDP MED extension and Avaya-Extreme proprietary TLVs was added in ExtremeXOS 11.5.

Additional PoE+ information can appear when present in a Power via MDI TLV received from a neighbor starting in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

show lldp statistics

```
show lldp {port [all | port_list]} statistics
```

Description

Displays statistical counters related to the specified port or ports.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

N/A.



Usage Guidelines

The following counters are presented with the standard command (taken from the IEEE 802.1ab MIB definition):

- Last table change time: Last time an entry in the LLDP database was added, changed or deleted.
- Number of table inserts: The number of times the complete set of information advertised by a particular neighbor has been inserted into tables.
- Number of table deletes: The number of times the complete set of information advertised by a particular neighbor has been deleted from tables.
- Number of table drops: The number of times the complete set of information advertised by a particular neighbor could not be stored in memory because of insufficient resources.
- Number of table age outs: The number of times the complete set of information advertised by a particular neighbor has been deleted from tables because the information timeliness interval has expired.
- Tx Total: The number of LLDP frames transmitted by this switch on the indicated port.
- Tx Total Length Exceeded: The number of LLDP frames sent out on this port that could not hold all the information configured because the total frame length would exceed the maximum LDDPDU size of 1500 bytes.
- Rx Total: The number of valid LLDP frames received by this switch on the indicated port, while this LLDP agent is enabled.
- Rx Discarded: The number of LLDP frames received by this switch on the indicated port, and then discarded for any reason.
- Rx Errors: The number of invalid LLDP frames received by this switch on the indicated port, while this LLDP agent is enabled.
- TLVs Discarded: The number of LLDP TLVs discarded for any reason by this switch on the indicated port.
- TLVs Unrecognized: The number of LLDP TLVs received on the given port that are not recognized by the switch.

Example

The following example lists statistical counters for all ports on the switch:

```
# show lldp port all statistics
Last table change time   : Fri Dec 17 10:42:33 2004
Number of Table Inserts  : 3
Number of Table Deletes  : 0
Number of Table Drops    : 0
Number of Table Age Outs : 0
Port      Tx          Tx LengthRx Rx          Rx          TLVs          TLVs
Total     Exceeded  TotalDiscarded  Errors      Discarded  Unrecogn.
=====
=====
1:1       189          05654          0           0           0           0
2:2       188          0565          0           0           0           0
```



History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

unconfigure lldp

```
unconfigure lldp {ports [all | port_list]}
```

Description

Leaves LLDP enabled and configured; restores the LLDP timer default values.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

When you issue the global `unconfigure lldp`, only the LLDP timers are reset to default values. All the configured TLVs remain on the ports remain, and LLDP remains enabled.

When you use the keyword `ports`, the TLVs for each port are returned to the five default TLVs. LLDP remains enabled.

Example

The following command restores LLDP factory default TLVs for ports 1:4 to 1:8:

```
unconfigure lldp ports 1:4 - 1:8
```

History

This command was first available in ExtremeXOS 11.2.

The keyword `port` was changed to `ports` in ExtremeXOS 11.5.



Platform Availability

This command is available on all platforms.



11 Commands for OAM

```
clear counters bfd
clear counters cfm segment <segment_name>
clear counters cfm segment all
clear counters cfm segment all frame-delay
clear counters cfm segment all frame-loss
clear counters cfm segment frame-delay
clear counters cfm segment frame-loss
clear counters cfm segment frame-loss mep
configure bfd vlan
configure bfd vlan authentication
configure cfm domain add association integer
configure cfm domain add association string
configure cfm domain add association vlan-id
configure cfm domain add association vpn-id oui index
configure cfm domain association add
configure cfm domain association add remote-mep
configure cfm domain association delete
configure cfm domain association delete remote-mep
configure cfm domain association destination-mac-type
configure cfm domain association end-point add group
configure cfm domain association ports end-point ccm
configure cfm domain association end-point delete group
configure cfm domain association ports end-point mepid
configure cfm domain association ports end-point sender-id-ipaddress
configure cfm domain association end-point transmit-interval
configure cfm domain association ports end-point
configure cfm domain association remote-mep mac-address
configure cfm domain delete association
configure cfm domain md-level
configure cfm group add rmep
configure cfm group delete rmep
configure cfm segment add domain association
configure cfm segment delete domain association
configure cfm segment dot1p
configure cfm segment frame-delay dot1p
configure cfm segment frame-delay/frame-loss transmit interval
configure cfm segment frame-delay window
```

```
configure cfm segment frame-loss dot1p
configure cfm segment frame-loss window
configure cfm segment frame-loss mep
configure cfm segment frame-loss consecutive
configure cfm segment frame-loss ses-threshold
configure cfm segment threshold
configure cfm segment timeout
configure cfm segment transmit-interval
configure cfm segment window
create cfm domain dns md-level
create cfm domain mac md-level
create cfm domain string md-level
create cfm segment destination
delete cfm domain
delete cfm segment
disable cfm segment frame-delay measurement
disable cfm segment frame-loss measurement mep
enable/disable bfd vlan
enable cfm segment frame-delay measurement
enable cfm segment frame-loss measurement mep
ping mac port
show bfd
show bfd counters
show bfd session client
show bfd session counters vr all
show bfd session detail vr all
show bfd session vr all
show bfd vlan
show bfd vlan counters
show cfm
show cfm detail
show cfm groups
show cfm segment
show cfm segment frame-delay
show cfm segment frame-delay/frame-loss mep id
show cfm segment frame-delay statistics
show cfm segment frame-loss
show cfm segment frame-loss statistics
show cfm segment mep
traceroute mac port
unconfigure bfd vlan
unconfigure cfm domain association end-point transmit-interval
```



Operation, Administration, and Maintenance (OAM) includes functions used to detect network faults, measure network performance and distribute fault-related information.

This chapter describes commands that are part of the following features.

Connectivity Fault Management (CFM)—This feature, discussed in the emerging IEEE 802.1ag specification, allows you to detect, verify, and isolate connectivity failures in virtual bridged LANs. Part of this specification is a toolset to manually check connectivity, which is sometimes referred to as Layer 2 ping.

Hierarchical networks, or domains, and test connectivity within that domain are created by sending Layer 2 messages, known as Connectivity Check Messages (CCMs). You use these domains to send loopback messages and link trace messages.

Y.1731 Compliant Frame Delay and Delay Variance Measurement—This feature is based on the ITU-T Y.1731 standard and deals with the Ethernet Delay Measurement (ETH-DM) function.

ExtremeXOS software supports:

- Two-way delay measurement—Delay Measurement Message (DMM) and Delay Measurement Reply (DMR).
- Continuous (proactive) measurement of frame delay and frame-delay variation.
- On-demand measurement of frame delay and frame-delay variation.

EFM OAM—Unidirectional Link Fault Management—IEEE 802.3ah, the Ethernet in the First Mile (EFM) standard, includes mechanisms for network OAM to facilitate metro Ethernet network operation and troubleshooting to match traditional carrier network technologies. This section covers that portion of EFM that deals with the unidirectional link fault indication on a 1G link that has the capability to transmit and receive independently.

Bidirectional Forwarding Detection (BFD)—This feature is a hello protocol that provides rapid detection of failures in the path and informs the clients (routing protocols) to initiate the route convergence. It is independent of media, routing protocols and data protocols.

clear counters bfd

```
clear counters bfd {session | interface}
```

Description

Clears the counters associated with BFD specific settings.

Syntax Description

This command has no arguments or variables.

Default

N/A.



Usage Guidelines

Use this command to clear the counters in the BFD session or interface (VLAN). If neither session or interface are specified, the command clears all counters in BFD.

Example

The following command clears all counters in BFD:

```
clear counters bfd
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

clear counters cfm segment <segment_name>

```
clear counters cfm segment segment_name
```

Description

This command clears both frame-delay and frame-loss information for segment with given segment name.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
---------------------	---

Default

N/A.

Usage Guidelines

Use this command to clear both frame-delay and frame-loss information for segment with given segment name.

Example

```
E4G-200.56 # clear co cfm seg cs2
E4G-200.57 #
```



```

E4G-200.57 # sho cfm seg cs2
CFM Segment Name      : cs2
Domain Name           : dom1
Association            : a2
MD Level              : 1
Destination MAC       : 00:04:96:52:a7:38
Frame Delay:
DMM Transmission      : In Progress
Transmission Mode     : On Demand
Total Frames to be sent : 45
Frames Transmitted    : 0
Pending Frames        : 40
Frames Received       : 0
DMM Tx Interval       : 10 secs
DMR Rx Timeout        : 50 msec
Alarm Threshold       : 10 %
Clear Threshold       : 95 %
Measurement Window Size : 60
Class of Service      : 6
Tx Start Time         : None
Min Delay             : None
Max Delay             : None
Last Alarm Time       : None
Alarm State           : None
Lost Frames           : 0
Frame Loss:
LMM Tx Interval       : 10 secs
SES Threshold         : 1.000000e-02
Consecutive Available Count : 4
Measurement Window Size : 1200
Class of Service      : 6
Total Configured MEPS : 1
Total Active MEPS     : 1
MEP ID                : 2
LMM Transmission      : In Progress
Transmission Mode     : On Demand
Total Frames to be sent : 45
Frames Transmitted    : 0
Pending Frames        : 40
Frames Received       : 0
Availability Status    : Idle
Unavailability Start Time : None
Unavailability End Time : None
Press <SPACE> to continue or <Q> to quit:
Tx Start Time         : None

```

```

-----
Total Configured Segments : 11
Total Active Segments    : 11
E4G-200.58 #
E4G-200.58 #
E4G-200.58 #
E4G-200.58 #
E4G-200.58 #
E4G-200.58 #
E4G-200.58 #
E4G-200.58 #

```



History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms.

clear counters cfm segment all

clear counters cfm segment all

Description

This command clears both frame-delay and frame-loss information for all existing segments.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

Use this command to clear both frame-delay and frame-loss information for all existing segments.

Example

```
E4G-200.53 # clear co cfm seg all
E4G-200.54 # sho cfm seg
CFM Segment Name           : cs10
Domain Name                 : dom1
Association                 : a10
MD Level                   : 1
Destination MAC             : 00:04:96:52:a7:38
Frame Delay:
DMM Transmission           : In Progress
Transmission Mode         : On Demand
Total Frames to be sent    : 45
Frames Transmitted         : 0
Pending Frames            : 42
Frames Received           : 0
DMM Tx Interval           : 10 secs
DMR Rx Timeout            : 50 msec
Alarm Threshold            : 10 %
Clear Threshold            : 95 %
Measurement Window Size    : 60
Class of Service           : 6
```



```

Tx Start Time           : None
Min Delay               : None
Max Delay               : None
Last Alarm Time        : None
Alarm State            : None
Lost Frames            : 0
Frame Loss:
LMM Tx Interval        : 10 secs
SES Threshold          : 1.000000e-02
Consecutive Available Count : 4
Measurement Window Size : 1200
Class of Service       : 6
Total Configured MEPs  : 1
Total Active MEPs      : 1
MEP ID                 : 10
LMM Transmission       : In Progress
Transmission Mode      : On Demand
Total Frames to be sent : 45
Frames Transmitted     : 0
Pending Frames        : 42
Frames Received        : 0
Availability Status    : Idle
Unavailability Start Time : None
Unavailability End Time : None
Tx Start Time         : None
CFM Segment Name      : cs11
Domain Name           : dom1
Association            : all
MD Level              : 1
Destination MAC       : 00:04:96:52:a7:38
Frame Delay:
DMM Transmission      : In Progress
Transmission Mode     : On Demand
Total Frames to be sent : 45
Frames Transmitted     : 0
Pending Frames        : 42
Frames Received        : 0
DMM Tx Interval       : 10 secs
DMR Rx Timeout        : 50 msec
Alarm Threshold        : 10 %
Clear Threshold       : 95 %
Measurement Window Size : 60
Class of Service       : 6
Tx Start Time         : Mon Mar 12 10:26:39 2012
Min Delay             : Mon Mar 12 10:26:49 2012
Max Delay             : Mon Mar 12 10:26:49 2012
Last Alarm Time      : None
Alarm State          : None
Lost Frames          : 0
Frame Loss:
LMM Tx Interval      : 10 secs
SES Threshold        : 1.000000e-02
Consecutive Available Count : 4
Measurement Window Size : 1200
Class of Service     : 6
Total Configured MEPs  : 1
Total Active MEPs     : 1
MEP ID               : 11

```



```

LMM Transmission      : In Progress
Transmission Mode     : On Demand
Total Frames to be sent : 45
Frames Transmitted    : 0
Pending Frames       : 42
Frames Received       : 0
Availability Status    : Idle
Unavailability Start Time : None
Unavailability End Time  : None
Tx Start Time         : None
CFM Segment Name      : cs12
Domain Name           : dom1
Association            : a12
MD Level              : 1
Destination MAC       : 00:04:96:52:a7:38
Frame Delay:
DMM Transmission      : In Progress
Transmission Mode     : On Demand
Total Frames to be sent : 45
Frames Transmitted    : 0
Pending Frames       : 42
Frames Received       : 0
DMM Tx Interval       : 10 secs
DMR Rx Timeout        : 50 msec
Alarm Threshold       : 10 %
Clear Threshold       : 95 %
Measurement Window Size : 60
Class of Service      : 6
Tx Start Time         : Mon Mar 12 10:26:39 2012
Min Delay             : Mon Mar 12 10:26:49 2012
Max Delay             : Mon Mar 12 10:26:39 2012
Last Alarm Time       : None
Alarm State           : None
Lost Frames           : 0
Frame Loss:
LMM Tx Interval       : 10 secs
SES Threshold         : 1.000000e-02
Consecutive Available Count : 4
Measurement Window Size : 1200
Class of Service      : 6
Total Configured MEPs : 1
Total Active MEPs     : 1
MEP ID                : 12
LMM Transmission      : In Progress
Transmission Mode     : On Demand
Total Frames to be sent : 45
Frames Transmitted    : 1
Pending Frames       : 41
Frames Received       : 1
Availability Status    : Available
-----
Total Configured Segments : 11
Total Active Segments     : 11
E4G-200.55 #
E4G-200.55 #
E4G-200.55 #
E4G-200.55 #

```



History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms.

clear counters cfm segment all frame-delay

clear counters cfm segment all frame-delay

Description

This command clears only frame-delay information for all existing segments.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

Use this command to clear only frame-delay information for all existing segments.

Example

```
E4G-200.70 # clear co cfm seg all frame-delay
E4G-200.71 #
E4G-200.71 #
E4G-200.71 #
E4G-200.71 # sho cfm segment
CFM Segment Name           : cs10
Domain Name                 : dom1
Association                 : a10
MD Level                   : 1
Destination MAC            : 00:04:96:52:a7:38
Frame Delay:
DMM Transmission           : In Progress
Transmission Mode         : On Demand
Total Frames to be sent    : 45
Frames Transmitted         : 1
Pending Frames             : 30
Frames Received           : 1
DMM Tx Interval           : 10 secs
DMR Rx Timeout            : 50 msec
Alarm Threshold            : 10 %
```



```

Clear Threshold           : 95 %
Measurement Window Size  : 60
Class of Service         : 6
Tx Start Time            : Mon Mar 12 10:28:59 2012
Min Delay                 : Mon Mar 12 10:28:59 2012
Max Delay                 : Mon Mar 12 10:28:59 2012
Last Alarm Time          : None
Alarm State               : Not Set
Lost Frames               : 0
Frame Loss:
LMM Tx Interval           : 10 secs
SES Threshold             : 1.000000e-02
Consecutive Available Count : 4
Measurement Window Size  : 1200
Class of Service         : 6
Total Configured MEPs    : 1
Total Active MEPs        : 1
MEP ID                    : 10
LMM Transmission          : In Progress
Transmission Mode        : On Demand
Total Frames to be sent  : 45
Frames Transmitted       : 4
Pending Frames           : 30
Frames Received          : 4
Availability Status       : Available
Unavailability Start Time : None
Unavailability End Time  : None
Tx Start Time            : Mon Mar 12 10:28:29 2012
CFM Segment Name         : cs11
Domain Name               : dom1
Association                : all
MD Level                  : 1
Destination MAC           : 00:04:96:52:a7:38
Frame Delay:
DMM Transmission         : In Progress
Transmission Mode        : On Demand
Total Frames to be sent  : 45
Frames Transmitted       : 1
Pending Frames           : 30
Frames Received          : 1
DMM Tx Interval           : 10 secs
DMR Rx Timeout           : 50 msec
Alarm Threshold           : 10 %
Clear Threshold           : 95 %
Measurement Window Size  : 60
Class of Service         : 6
Tx Start Time            : Mon Mar 12 10:28:59 2012
Min Delay                 : Mon Mar 12 10:28:59 2012
Max Delay                 : Mon Mar 12 10:28:59 2012
Last Alarm Time          : None
Alarm State               : Not Set
Lost Frames               : 0
Frame Loss:
LMM Tx Interval           : 10 secs
SES Threshold             : 1.000000e-02
Consecutive Available Count : 4
Measurement Window Size  : 1200
Class of Service         : 6

```



```

Total Configured MEPs      : 1
Total Active MEPs         : 1
MEP ID                    : 11
LMM Transmission          : In Progress
Transmission Mode         : On Demand
Total Frames to be sent   : 45
Frames Transmitted        : 12
Pending Frames            : 30
Frames Received           : 12
Availability Status       : Available
Unavailability Start Time : None
Unavailability End Time   : None
Tx Start Time             : Mon Mar 12 10:27:09 2012
CFM Segment Name          : cs12
Domain Name                : dom1
Association                : a12
MD Level                  : 1
Destination MAC           : 00:04:96:52:a7:38
Frame Delay:
DMM Transmission          : In Progress
Transmission Mode         : On Demand
Total Frames to be sent   : 45
Frames Transmitted        : 1
Pending Frames            : 30
Frames Received           : 1
DMM Tx Interval           : 10 secs
DMR Rx Timeout            : 50 msec
Alarm Threshold           : 10 %
Clear Threshold           : 95 %
Measurement Window Size   : 60
Class of Service          : 6
Tx Start Time             : Mon Mar 12 10:28:59 2012
-----
Total Configured Segments : 11
Total Active Segments     : 11
E4G-200.72 #
E4G-200.72 #
E4G-200.72 #

```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms.

clear counters cfm segment all frame-loss

```
clear counters cfm segment all frame-loss
```



Description

This command clears only frame-loss information for all existing segments.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

Use this command to clear only frame-loss information for all existing segments.

Example

```
E4G-200.72 # clear co cfm seg all frame-loss
E4G-200.73 #
E4G-200.73 #
E4G-200.73 #
E4G-200.73 # sho cfm segment
CFM Segment Name      : cs10
Domain Name           : dom1
Association            : a10
MD Level              : 1
Destination MAC       : 00:04:96:52:a7:38
Frame Delay:
DMM Transmission      : In Progress
Transmission Mode     : On Demand
Total Frames to be sent : 45
Frames Transmitted    : 2
Pending Frames        : 29
Frames Received       : 2
DMM Tx Interval       : 10 secs
DMR Rx Timeout        : 50 msec
Alarm Threshold        : 10 %
Clear Threshold       : 95 %
Measurement Window Size : 60
Class of Service      : 6
Tx Start Time         : Mon Mar 12 10:28:59 2012
Min Delay              : Mon Mar 12 10:29:09 2012
Max Delay              : Mon Mar 12 10:29:09 2012
Last Alarm Time       : None
Alarm State           : Not Set
Lost Frames           : 0
Frame Loss:
LMM Tx Interval       : 10 secs
SES Threshold         : 1.000000e-02
Consecutive Available Count : 4
Measurement Window Size : 1200
Class of Service      : 6
```



```

Total Configured MEPs      : 1
Total Active MEPs         : 1
MEP ID                    : 10
LMM Transmission          : In Progress
Transmission Mode         : On Demand
Total Frames to be sent   : 45
Frames Transmitted        : 0
Pending Frames            : 29
Frames Received           : 0
Availability Status       : Idle
Unavailability Start Time : None
Unavailability End Time   : None
Tx Start Time             : None
CFM Segment Name         : cs11
Domain Name               : dom1
Association                : all
MD Level                  : 1
Destination MAC           : 00:04:96:52:a7:38
Frame Delay:
DMM Transmission          : In Progress
Transmission Mode         : On Demand
Total Frames to be sent   : 45
Frames Transmitted        : 2
Pending Frames            : 29
Frames Received           : 2
DMM Tx Interval           : 10 secs
DMR Rx Timeout            : 50 msec
Alarm Threshold           : 10 %
Clear Threshold           : 95 %
Measurement Window Size   : 60
Class of Service          : 6
Tx Start Time             : Mon Mar 12 10:28:59 2012
Min Delay                 : Mon Mar 12 10:29:09 2012
Max Delay                 : Mon Mar 12 10:28:59 2012
Last Alarm Time           : None
Alarm State               : Not Set
Lost Frames               : 0
Frame Loss:
LMM Tx Interval           : 10 secs
SES Threshold             : 1.000000e-02
Consecutive Available Count : 4
Measurement Window Size   : 1200
Class of Service          : 6
Total Configured MEPs    : 1
Total Active MEPs        : 1
MEP ID                    : 11
LMM Transmission          : In Progress
Transmission Mode         : On Demand
Total Frames to be sent   : 45
Frames Transmitted        : 0
Pending Frames            : 28
Frames Received           : 0
Availability Status       : Idle
Unavailability Start Time : None
Unavailability End Time   : None
Tx Start Time             : Mon Mar 12 10:29:19 2012
CFM Segment Name         : cs12
Domain Name               : dom1

```



```

Association                : a12
MD Level                   : 1
Destination MAC            : 00:04:96:52:a7:38
Frame Delay:
DMM Transmission           : In Progress
Transmission Mode         : On Demand
Total Frames to be sent    : 45
Frames Transmitted         : 2
Pending Frames             : 29
Frames Received            : 2
DMM Tx Interval           : 10 secs
DMR Rx Timeout             : 50 msec
Alarm Threshold            : 10 %
Clear Threshold            : 95 %
Measurement Window Size   : 60
Class of Service           : 6
Tx Start Time              : Mon Mar 12 10:28:59 2012
-----

```

```

Total Configured Segments  : 11
Total Active Segments      : 11
E4G-200.74 #
E4G-200.74 #
E4G-200.74 #

```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms.

clear counters cfm segment frame-delay

```
clear counters cfm segment segment_name frame-delay
```

Description

This command clears only frame-delay information for segment with given segment name.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
---------------------	---

Default

N/A.



Usage Guidelines

Use this command to clear only frame-delay information for segment with given segment name.

Example

```

E4G-200.59 # clear co cfm seg cs10 frame-delay
E4G-200.60 #
E4G-200.60 #
E4G-200.60 #
E4G-200.60 # sho cfm seg cs10
CFM Segment Name           : cs10
Domain Name                 : dom1
Association                  : a10
MD Level                    : 1
Destination MAC             : 00:04:96:52:a7:38
Frame Delay:
DMM Transmission           : In Progress
Transmission Mode         : On Demand
Total Frames to be sent   : 45
Frames Transmitted        : 1
Pending Frames            : 34
Frames Received           : 1
DMM Tx Interval           : 10 secs
DMR Rx Timeout            : 50 msec
Alarm Threshold            : 10 %
Clear Threshold           : 95 %
Measurement Window Size   : 60
Class of Service          : 6
Tx Start Time             : Mon Mar 12 10:28:19 2012
Min Delay                  : Mon Mar 12 10:28:19 2012
Max Delay                  : Mon Mar 12 10:28:19 2012
Last Alarm Time           : None
Alarm State               : Not Set
Lost Frames               : 0
Frame Loss:
LMM Tx Interval           : 10 secs
SES Threshold              : 1.000000e-02
Consecutive Available Count : 4
Measurement Window Size   : 1200
Class of Service          : 6
Total Configured MEPs     : 1
Total Active MEPs         : 1
MEP ID                    : 10
LMM Transmission         : In Progress
Transmission Mode         : On Demand
Total Frames to be sent   : 45
Frames Transmitted        : 8
Pending Frames            : 34
Frames Received           : 8
Availability Status       : Available
Unavailability Start Time : None
Unavailability End Time   : None
Tx Start Time             : Mon Mar 12 10:27:09 2012
-----
Total Configured Segments : 11

```



```
Total Active Segments      : 11
E4G-200.61 #
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms.

clear counters cfm segment frame-loss

```
clear counters cfm segment segment_name frame-loss
```

Description

This command clears only frame-loss information for segment with given segment name for all associated MEPs.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
---------------------	---

Default

N/A.

Usage Guidelines

Use this command to clear only frame-loss information for segment with given segment name for all associated MEPs.

Example

```
E4G-200.61 # clear co cfm seg cs10 frame-loss
E4G-200.62 #
E4G-200.62 #
E4G-200.62 #
E4G-200.62 # sho cfm seg cs10
CFM Segment Name      : cs10
Domain Name           : dom1
Association            : a10
MD Level              : 1
Destination MAC       : 00:04:96:52:a7:38
Frame Delay:
DMM Transmission      : In Progress
```




```

Transmission Mode           : On Demand
Total Frames to be sent     : 45
Frames Transmitted          : 1
Pending Frames              : 34
Frames Received             : 1
DMM Tx Interval             : 10 secs
DMR Rx Timeout              : 50 msec
Alarm Threshold             : 10 %
Clear Threshold             : 95 %
Measurement Window Size    : 60
Class of Service            : 6
Tx Start Time               : Mon Mar 12 10:28:19 2012
Min Delay                   : Mon Mar 12 10:28:19 2012
Max Delay                   : Mon Mar 12 10:28:19 2012
Last Alarm Time             : None
Alarm State                 : Not Set
Lost Frames                 : 0
Frame Loss:
LMM Tx Interval             : 10 secs
SES Threshold               : 1.000000e-02
Consecutive Available Count : 4
Measurement Window Size    : 1200
Class of Service            : 6
Total Configured MEPs      : 1
Total Active MEPs          : 1
MEP ID                      : 10
LMM Transmission           : In Progress
Transmission Mode          : On Demand
Total Frames to be sent     : 45
Frames Transmitted          : 1
Pending Frames              : 33
Frames Received             : 1
Availability Status         : Available
Unavailability Start Time   : None
Unavailability End Time     : None
Tx Start Time               : Mon Mar 12 10:28:29 2012
-----
Total Configured Segments   : 11
Total Active Segments       : 11
E4G-200.63 #
E4G-200.63 #
E4G-200.63 #
E4G-200.63 #
E4G-200.63 #

```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms.



clear counters cfm segment frame-loss mep

```
clear counters cfm segment segment_name frame-loss mep mep_id
```

Description

This command clears only frame-loss information for the given MEP in segment with given segment name.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
---------------------	---

Default

N/A.

Usage Guidelines

Use this command to clear only frame-loss information for the given MEP in segment with given segment name.

Example

```
E4G-200.24 # clear counters cfm segment "cs2" frame-loss mep 3
E4G-200.25 #
E4G-200.25 #
E4G-200.25 #
E4G-200.25 # sho cfm segment
CFM Segment Name           : cs2
Domain Name                 : dom2
Association                  : a2
MD Level                    : 2
Destination MAC             : 00:04:96:52:a7:64
Frame Delay:
DMM Transmission            : Disabled
Frames Transmitted          : 0
Frames Received             : 0
DMM Tx Interval             : 10 secs
DMR Rx Timeout              : 50 msec
Alarm Threshold              : 10 %
Clear Threshold              : 95 %
Measurement Window Size     : 60
Class of Service            : 6
Tx Start Time               : None
Min Delay                   : None
Max Delay                   : None
Last Alarm Time             : None
Alarm State                 : None
Lost Frames                 : 0
```



```

Frame Loss:
LMM Tx Interval           : 10 secs
SES Threshold             : 1.000000e-02
Consecutive Available Count : 4
Measurement Window Size   : 1200
Class of Service          : 6
Total Configured MEPs     : 1
Total Active MEPs        : 1
MEP ID                    : 3
LMM Transmission          : In Progress
Transmission Mode         : Continuous
Frames Transmitted        : 0
Frames Received           : 0
Availability Status       : Idle
Unavailability Start Time : None
Unavailability End Time   : None
Tx Start Time             : None
-----
Total Configured Segments : 1
Total Active Segments     : 1
E4G-200.26 #
E4G-200.26 #

```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms.

configure bfd vlan

```
configure bfd vlan vlan_name [{detection-multiplier multiplier} {receive-interval
rx_interval} {transmit-interval tx_interval}]
```

Description

Configures BFD transmit (TX) and receive (RX) intervals and multipliers on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN.
<i>multiplier</i>	Specifies the detection multiplier. The range is 1 to 255.
<i>rx_interval</i>	Specifies the receive interval for control packets in milliseconds. The range is 100 to 4294967 ms.
<i>tx_interval</i>	Specifies the transmit interval for control packets in milliseconds. The range is 100 to 4294967 ms.



Default

The default value for RX and TX intervals is 1000 ms.

The default value for the detection-multiplier is 3.

Usage Guidelines

Use this command to configure BFD.

Use the `show bfd vlan` command to display the current settings.

Example

The following command configures a transmit and receive interval of 2000 ms and a detection multiplier of 2 on the VLAN `vlan1`:

```
configure bfd vlan vlan1 detection-multiplier 2 receive-interval 2000
transmit-interval 2000
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

configure bfd vlan authentication

```
configure bfd vlan vlan_name authentication [none | simple-password {encrypted}
password]]
```

Description

Configures authentication for BFD on a VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN name.
none	Specifies that no authentication is to be used. (Default)
encrypted	Indicates that the password is already encrypted.
<i>password</i>	Specifies a simple password to use to authenticate.



Default

The authentication default is none.

Usage Guidelines

Use this command to configure authentication for BFD on a VLAN using a password or specify that none is required.

Use the `show bfd vlan` command to display the authentication setting.

The encrypted keyword is primarily for the output of the show configuration command, so that the password is not revealed in the command output. Do not use it to set the password

Example

The following command configures authentication using the password password:

```
configure bfd vlan vlan1 authentication simple-password password
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

configure cfm domain add association integer

```
configure cfm domain domain_name add association integer int [vlan vlan_name | vman vman_name ]
```

Description

Creates a maintenance association (MA) related to a specified maintenance domain (MD). This command supports the 2-octet integer MA format.

Syntax Description

<i>domain_name</i>	Specifies the domain you want to associate with this MA.
<i>int</i>	Enter an integer to name the MA. The range is 0 to 65535.
<i>vlan_name</i>	Specifies the VLAN you want to assign to this MA. Each MA contains only one VLAN, VMAN, BVLAN or SVLAN.
<i>vman_name</i>	Specifies the VMAN you want to assign to this MA.



Default

N/A.

Usage Guidelines

All ports configured on the specified VLAN are now CFM ports in the specified MA.

You add the MA, or association, to the domain, and the MA uses the MD level assigned to the domain. Each MA can belong to only one domain, but several MAs can belong to a given domain. The MA is unique within a given domain.

Example

The following command creates a 2-octet integer MA (350) that associates the domain brazil and the VLAN admin:

```
configure cfm domain brazil add association integer 350 vlan admin
```

History

This command was first available in ExtremeXOS 11.4.

The SVLAN option was added in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

configure cfm domain add association string

```
configure cfm domain domain_name add association string name [vlan vlan_name | vman vman_name ]
```

Description

Creates a maintenance association (MA) related to a specified maintenance domain (MD). This command supports the character string MA format.

Syntax Description

<i>domain_name</i>	Specifies the domain you want to associate with this MA.
string	Enter up to 45 alphanumeric characters to name the MA.
<i>vlan_name</i>	Specifies the VLAN you want to assign to this MA. Each MA contains only one VLAN, VMAN, or BVLAN.
<i>vman_name</i>	Specifies the VMAN you want to assign to this MA.



Default

N/A.

Usage Guidelines

All ports configured on the specified VLAN are now CFM ports in the specified MA.

You add the MA, or association, to the domain, and the MA uses the MD level assigned to the domain. Each MA can belong to only one domain, but several MAs can belong to a given domain. The MA is unique within a given domain.

Example

The following command creates an MA named service that associates the MD spain and the VLAN finance:

```
configure cfm domain service add association string spain vlan finance
```

History

This command was first available in ExtremeXOS 11.4.

The SVLAN option was added in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

configure cfm domain add association vlan-id

```
configure cfm domain domain_name add association vlan-id vlanid [vlan vlan_name | vman vman_name ]
```

Description

Creates a maintenance association (MA) related to a specified maintenance domain (MD). This command supports the VLAN ID MA format.

Syntax Description

<i>domain_name</i>	Specifies the domain you want to associate with this MA.
<i>vlanid</i>	Specifies the VLAN ID.
<i>vlan_name</i>	Specifies the VLAN you want to assign to this MA. Each MA contains only one VLAN, VMAN, or BVLAN.
<i>vman_name</i>	Specifies the VMAN you want to assign to this MA.



Default

N/A.

Usage Guidelines

All ports configured on the specified VLAN are now CFM ports in the specified MA.

You add the MA, or association, to the domain, and the MA uses the MD level assigned to the domain. Each MA can belong to only one domain, but several MAs can belong to a given domain. The MA is unique within a given domain.

History

This command was first available in ExtremeXOS 12.1.

The SVLAN option was added in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

configure cfm domain add association vpn-id oui index

```
configure cfm domain domain_name add association vpn-id oui oui index index [vlan vlan_name | vman vman_name ]
```

Description

Creates a maintenance association (MA) related to a specified maintenance domain (MD). This command supports the RFC 2685 VPN ID MA format.

Syntax Description

<i>domain_name</i>	Specifies the domain you want to associate with this MA.
<i>oui</i>	Enter a virtual private network (VPN) Organizational Unique Identifier (OUI) in the format XX:XX:XX as part of the name for the MA.
<i>index</i>	Enter the 32-bit VPN index you want to append to the OUI to name the MA. The range is 0 to 4294967295.
<i>vlan_name</i>	Specifies the VLAN you want to assign to this MA. Each MA contains only one VLAN, VMAN, or BVLAN.
<i>vman_name</i>	Specifies the VMAN you want to assign to this MA.

Default

N/A.



Usage Guidelines

All ports configured on the specified VLAN are now CFM ports in the specified MA. You add the MA, or association, to the domain, and the MA uses the MD level assigned to the domain. Each MA can belong to only one domain, but several MAs can belong to a given domain. The MA is unique within a given domain.

Example

The following command creates an MA with the VPN ID of 11:22:33 50 that associates the domain spain and the VLAN accounting:

```
configure cfm domain spain add association vpn-id oui 11:22:33 index 50 vlan
accounting
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure cfm domain association add

```
configure cfm domain domain_name association association_name [ports port_list
add [[end-point [up|down] mepid { group group_name } ] | [intermediate-point]]
```

Description

This command allows you to create an up MEP, down MEP, intermediate-point (MIP) on a maintenance association, a group. You can also combine different maintenance points.

Combining different Maintenance points is restricted per the following:

- Up MEP and Down MEP in a single association is not allowed.
- Down MEP and MIP in a single association is not allowed.
- More than one Up MEP in a single association is not allowed.
- Up MEP and MIP in a single association is allowed.
- More than one Down MEP in a single association is allowed.
- A group can be created while creating a MEP.
- With CFM Support over VPLS, this command is used to associate pseudo wires of a VPLS service instance to an association & domain.
- Portlist can have only one port configured for a MEP configuration but can have multiple ports in MIP configuration, when Hwaoam is supported on the system.



Syntax Description

<i>domain_name</i>	Enter the domain associated with the MA you are configuring.
<i>association_name</i>	Enter the name of the MA for which you are configuring a MIP.
<i>port_list</i>	Specifies the port number(s).
up	Enter the port to be the UP port of the MA; this MEP sends CCM messages to all ports—other than the sending switch port—in this MA on this switch.
down	Enter the port to be the DOWN port of the MA; this MEP sends CCM messages out of the configured physical port.
<i>mepid</i>	Specifies a value for this MEP. The range is 1 to 8191. NOTE: On each MA, each MEPID must be unique.
group	CFM group that binds an LMEP to RMEPs. If not specified, the client does not receive events from the respective RMEPs.
<i>group_name</i>	Group name, maximum of 31 characters.

Default

N/A.

Usage Guidelines

These ports must already be in the MA (VLAN or VMAN) prior to assigning a MEP function to them. If you try to assign a port not in the MA as an end-point, the system returns the following message:

```
The following port(s) <portlist> are not part of the associations VLAN.
```



Note

Ensure that you assigned the port number correctly to the UP MEP and to the DOWN MEP, or the CCM messages go in the wrong direction.

Each MA needs at least two MEPs that can reach each other to exchange CCM messages.

You can also combine different maintenance points. The following are CLI restrictions on MP combinations:

- DOWN and UP MEP cannot be present on the same association
- DOWN MEP and MIP cannot be present on the same association
- UP MEP and MIP can be present on the same association
- Only one UP MEP is allowed in an association
- Multiple DOWN MEPs are allowed in an association

You can configure a total of 32 MIPs on a single switch.

Use the `show cfm` command to verify your configuration.



Example

The following command configures port 1:20 as a MIP on the 350 association in the spain domain:

```
configure cfm domain spain association 350 ports 1:20 add intermediate-point
```

The following command configures port 5:10 to be the UP MEP on the test association in the brazil domain, with a mepid of 500:

```
configure cfm domain brazil association test ports 5:10 add end-point up 500
```

History

This command was first available in ExtremeXOS 11.4.

This command was updated in ExtremeXOS 15.2 to include the optional group parameter.

Platform Availability

This command is available on all platforms.

configure cfm domain association add remote-mep

```
configure cfm domain domain-name association association_name add remote-mep  
mepid { mac-address mac_address }
```

Description

Allows you to add a remote MEP with the given MEP ID and MAC address to an existing association.

Syntax Description

<i>domain_name</i>	Enter the domain associated with the MA you are configuring.
<i>association_name</i>	Enter the name of the MA for which you are adding a remote MEP.
<i>mepid</i>	Enter the MEP ID of the remote MEP being added. The range is 1 to 8191.
<i>mac_address</i>	Specifies the MAC address for the remote MEP being added.

Default

N/A.



Usage Guidelines

Use this command to add a remote MEP with given MEP ID and MAC address to an existing association. Use the `show cfm detail` command to verify your configuration.



Note

Because the Summit X460 does not support unicast CCM generation, creating an RMEP with unicast MAC address is not meaningful. Therefore, it is an optional parameter on E4G400, E4G200, and Summit X460 platforms.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

configure cfm domain association delete

```
configure cfm domain domain_name association association_name [ports port_list
delete [[end-point [up|down]] | [intermediate-point] ] ]
```

Description

Deletes a maintenance end point (MEP) or maintenance intermediate point (MIP) from that MA.

Syntax Description

<i>domain_name</i>	Enter the domain associated with the MA you are configuring.
<i>association_name</i>	Enter the name of the MA for which you are deleting an MIP.
<i>port_list</i>	Specifies the port number(s).
up	Specifies that an UP MEP is to be deleted.
down	Specifies that a DOWN MEP is to be deleted.

Default

N/A.

Usage Guidelines

Use this command to delete an MEP or MIP.

If the VPLS option is chosen then the CFM deletes all the VPLS-based MIPs.

Use the `show cfm` command to verify your configuration.



Example

The following command deletes port 5:12 as an MIP on the test association in the brazil domain:

```
configure cfm domain brazil association test ports 5:12 delete intermediate-point
```

The following command deletes an UP MEP on port 5:10 on the test association in the brazil domain:

```
configure cfm domain brazil association test ports 5:10 delete end-point up
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure cfm domain association delete remote-mep

```
configure cfm domain domain-name association association_name delete remote-mep mepid
```

Description

Allows you to delete a remote MEP for a specific MEP ID and MAC address.

Syntax Description

<i>domain_name</i>	Enter the domain associated with the MA you are configuring.
<i>association_name</i>	Enter the name of the MA for which you are changing an MEP ID.
<i>mepid</i>	Enter the MEP ID of the remote MEP that is to be deleted.

Default

N/A.

Usage Guidelines

Use this command to delete a remote MEP of an MA for a specific MEP ID.

Use the `show cfm detail` command to verify your configuration.



History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

configure cfm domain association destination-mac-type

```
configure cfm domain domain-name association association_name destination-mac-type [unicast | multicast]
```

Description

Allows you to choose the destination MAC type for sending CFM PDUs for an MA.

Syntax Description

<i>domain_name</i>	Enter the domain associated with the MA you are configuring.
<i>association_name</i>	Enter the name of the MA for which you are changing the MAC type.
unicast	CFM PDUs are sent to the unicast MAC address configured in static remote MEP creation.
multicast	CFM PDUs are sent to the standard multicast destination address.

Default

Multicast.

Usage Guidelines

Use this command to change the MAC type on a previously configured MA. If multicast is selected, CFM PDUs are sent to the standard multicast destination. If unicast is selected, CFM PDUs are sent to the unicast MAC address configured in static remote MEP creation.

Use the `show cfm` command to verify your configuration.

E4G400, E4G200, and Summit X460 do not support unicast CCM (Continuity Check Message) generation. When the user configures the destination MAC type as unicast, the following message appears:

```
Error: IEEE 802.1ag PDUs can be sent only to standard multicast address on this platform
```

History

This command was first available in ExtremeXOS 12.1.



Platform Availability

This command is available on all platforms.

configure cfm domain association end-point add group

```
configure cfm domain domain-name association association-name ports port-list
end-point [up | down] add group group_name
```

Description

This command allows you to create a group for an existing local end-point.

Syntax Description

<i>domain_name</i>	Enter the domain associated with the MA you are configuring.
<i>association_name</i>	Enter the name of the MA for which you are configuring an MEP.
<i>port_list</i>	Enter the port number you want to configure as either an UP or DOWN MEP.
delete	Delete configuration from the association.

Default

N/A.

Usage Guidelines

Use this command to add a group to the association.

Example

```
configure cfm domain "MD1" association "MD1v1" ports 17 end-point down add
group "eapsCfmGrp"
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on all platforms.



configure cfm domain association ports end-point ccm

```
configure cfm domain domain_name association association_name ports port_list
end-point [up | down ] ccm [disable | enable]
```

Description

This command is used to enable or disable sending CCMs on a given MEP.

Syntax Description

<i>domain_name</i>	Enter the domain associated with the MA you are configuring.
<i>association_name</i>	Enter the name of the MA for which you are configuring an MEP.
<i>port_list</i>	Enter the port number you want to configure as either an UP or DOWN MEP.

Default

Enabled.

Usage Guidelines

Each MA needs at least two MEPs that can reach each other to exchange CCM messages.



Note

Ensure that you assigned the port number correctly to the UP MEP and to the DOWN MEP, or the CCM messages go in the wrong direction.

These ports must already be in the MA (VLAN or VMAN) prior to assigning a MEP function to them. If you try to assign a port not in the MA as an end-point, the system returns the following message:

```
The following port(s) <portlist> are not part of the associations VLAN.
```

Use the `show cfm` command to verify your configuration.

Example

```
configure cfm domain "MD1" association "MD1v1" ports 17 end-point down delete group "eapsCfmGrp"
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.



configure cfm domain association end-point delete group

```
configure cfm domain domain_name association association_name ports port_list
end-point [up|down] delete group [group_name | all ]
```

Description

This command allows you to delete one or all groups.

Syntax Description

<i>domain_name</i>	Enter the domain associated with the MA you are configuring.
<i>association_name</i>	Enter the name of the MA for which you are configuring an MEP.
<i>port_list</i>	Enter the port number you want to configure as either an UP or DOWN MEP.
delete	Delete configuration from the association

Default

N/A.

Usage Guidelines

Use this command to delete one or all groups from the association.

Example

```
configure cfm domain "MD1" association "MD1v1" ports 17 end-point down delete
group "eapsCfmGrp"
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on all platforms.

configure cfm domain association ports end-point mepid

```
configure cfm domain domain-name association association_name ports port_list
end-point [up | down] mepid mepid
```



Description

Allows you to change the MEP ID for a previously configured MEP. Each MEP within a single MA must have a unique MEP ID.

Syntax Description

<i>domain_name</i>	Enter the domain associated with the MA you are configuring.
<i>association_name</i>	Enter the name of the MA for which you are changing an MEP ID.
<i>port_list</i>	Enter the port number you want to change the MEP ID.
up	Enter this variable if you are changing the MEP ID on an UP MEP.
down	Enter this variable if you are changing the MEP ID on a DOWN MEP.
<i>mepid</i>	Enter the new value for this MEP. The range is 1 to 8191. NOTE: On each MA, each MEPID must be unique.

Default

N/A.

Usage Guidelines

Use this command to change the MEPID on a previously configured UP or DOWN MEP. If you attempt to change the MEPID on a port that is either not an MEP or having wrong MEP type, the system returns an error message.

Use the `show cfm` command to verify your configuration.

Example

The following command changes the MEP ID to 75 on the previously configured port 2:4 UP MEP on the 350 association in the finance domain:

```
configure cfm domain finance association 350 ports 2:4 end-point up mepid 75
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.



configure cfm domain association ports end-point sender-id-ipaddress

```
configure cfm domain domain_name association association_name ports port_list
end-point [up | down ] sender-id-ipaddress [disable | enable ip-address]
```

Description

This command is used to disable or enable configuring the sender-id-ipaddress on a given MEP.

Syntax Description

<i>domain_name</i>	Enter the domain associated with the MA you are configuring.
<i>association_name</i>	Enter the name of the MA for which you are configuring an MEP.
<i>port_list</i>	Enter the port number.
<i>ip-address</i>	Specifies the IP address that is sent in the sender-id TLV of the CFM PDUs.

Default

Disable.

Usage Guidelines

Each MA needs at least two MEPs that can reach each other to exchange CCM messages.



Note

Ensure that you assigned the port number correctly to the UP MEP and to the DOWN MEP, or the CCM messages go in the wrong direction.

You must create the MEP for which the configuration is being made before changing the configuration. Otherwise, the following error message is displayed:

```
The following port(s) <portlist> are not part of the associations VLAN.
```

Use the `show cfm` command to verify your configuration.



Note

E4G400, E4G200, and Summit X460 do not support this option. When the user configures a sender-id-ipaddress on an end-point, the following message appears:

```
Error: Sender ID IP Address configuration is not supported on this platform.
```

History

This command was first available in ExtremeXOS 12.1.



Platform Availability

This command is available on all platforms except E4G400, E4G200, and Summit X460.

configure cfm domain association end-point transmit-interval

```
configure cfm domain domain_name association association_name [ports port_list
end-point [up | down]] transmit-interval [3|10|100|1000|10000|60000|600000]
```

Description

Allows you to change time interval for an MEP to send out a CCM. We recommend configuring this value as at least 1 second.

Syntax Description

<i>domain_name</i>	Enter the domain associated with the MA you are configuring.
<i>association_name</i>	Enter the name of the MA for which you are changing the time interval an MEP sends out CCM.
<i>port_list</i>	Enter the port number of the MEP on which you are changing the time interval it sends out a CCM.
up	Enter this variable if you are changing the time interval for sending a CCM on an UP MEP.
down	Enter this variable if you are changing the time interval for sending a CCM on a DOWN MEP.

Default

1000 ms.

Usage Guidelines

Use this command to change the time interval between sending out CCMs on a previously configured UP or DOWN MEP. If you attempt to change the interval on a port that is either not an MEP or having wrong MEP type, the system returns an error message.



Note

We recommend that you use a transmit interval of at least 1 second (1000 ms).

The receiving system also uses this value multiplied by 3.5 to determine when the MEP is no longer alive.



Use the `show cfm` command to verify your configuration and the `show cfm detail` command to display the configured lifetime.



Note

The transmit interval value “3” is 3.3 msec. The values 3 and 10 are supported on platforms x460, E4G400 and E4G200 only for down MEPS. Also, the values 60000 and 600000 are supported in hardware.

Example

The following command changes the interval the UP MEP (previously configured on port 2:4) uses to send CCM messages on the 350 association in the finance domain to 10 seconds:

```
configure cfm domain finance association 350 ports 2:4 end-point up transmit-
interval 10000
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure cfm domain association ports end-point

```
configure cfm domain domain_name association association_name ports port_list
end-point [up | down] [enable | disable]
```

Description

Enables or disables an MEP.

Syntax Description

<i>domain_name</i>	Specifies the domain name.
<i>association_name</i>	Specifies the MA name.
<i>port_list</i>	Specifies the ports to configure.
up	Specifies that the end point is up.
down	Specifies that the end point is down.

Default

MEP is enabled by default.



Usage Guidelines

Use this command to enable or disable an MEP.

Use the `show cfm` command to verify your configuration.

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.

configure cfm domain association remote-mep mac-address

```
configure cfm domain domain-name association association_name remote-mep mepid
mac-address mac_address
```

Description

Allows you to modify the MAC address of an existing MEP.

Syntax Description

<i>domain_name</i>	Enter the domain associated with the MA you are configuring.
<i>association_name</i>	Specifies the name of the MA for which you are modifying a remote MEP.
<i>mepid</i>	Specifies the MEP ID of the remote MEP being modified. The range is 1 to 8191.
<i>mac_address</i>	Specifies the MAC address for the remote MEP being modified.

Default

N/A.

Usage Guidelines

Use this command to modify a remote MEP with given MEP ID and MAC address in an existing association. Use the `show cfm detail` command to verify your configuration.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.



configure cfm domain delete association

```
configure cfm domain domain_name delete association association_name
```

Description

Deletes a maintenance association (MA), including all its configured values, from the switch.

Syntax Description

<i>domain_name</i>	Enter the domain associated with the MA you are deleting.
<i>association_name</i>	Enter the name of the MA you are deleting.

Default

N/A.

Usage Guidelines

When you delete an association, or MA, you also remove all its configured values from the switch. These values include all configured MEPs, MIPs, and static remote MEPs.

Example

The following command deletes the MA test, in the domain of brazil, from the switch, along with all its configured MIPs, MEPs, and static remote MEPs:

```
configure cfm domain brazil delete association test
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure cfm domain md-level

```
configure cfm domain domain_name md-level level
```

Description

Changes a previously configured MD level for the specified domain.



Syntax Description

<i>domain_name</i>	Enter the name of the domain for which you want to change the MD level.
<i>level</i>	Specifies the new MD level you are assigning to this domain. Enter a value between 0 and 7.

Default

N/A.

Usage Guidelines

You can have up to 8 domains on a switch, and each one must have a unique MD level. Thus, a given MD level exists only once on a switch.

The IEEE standard 801.2ag specifies different levels for different network users, as follows:

- 5 to 7 for end users
- 3 and 4 for Internet service providers (ISPs)
- 0 to 2 for operators (entities carrying the information for the ISPs)

Example

The following command changes the MD level of a previously created domain extreme to 2:

```
configure cfm domain extreme md-level 2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure cfm group add rmp

```
configure cfm group group_name add rmp mepid
```

Description

This command allows you to create and associate an RMEP to a group.

Syntax Description

<i>mepid</i>	Specifies the MEP ID of the remote MEP being created. The range is 1 to 8191.
--------------	---



Default

N/A.

Usage Guidelines

Use this command to create and associate an RMEP to a group.

Example

```
configure cfm group "eapsCfmGroup" add rmep 2
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on all platforms.

configure cfm group delete rmep

```
configure cfm group group_name delete rmep [mepid | all]
```

Description

This command allows you to delete one or all RMEPs from a group.

Syntax Description

<i>mepid</i>	Specifies the MEP ID of the remote MEP being created. The range is 1 to 8191.
--------------	---

Default

N/A.

Usage Guidelines

Use this command to delete one or all RMEPs from a group.

Example

```
configure cfm group "eapsCfmGroup" delete rmep 2
```



History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on all platforms.

configure cfm segment add domain association

```
configure cfm segment segment_name add domain domain_name association
association_name
```

Description

Adds a CFM domain and association to a CFM segment.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
<i>domain_name</i>	Specifies the IEEE 802.lag maintenance domain.
<i>association_name</i>	Specifies the IEEE 802.lag association name.

Default

N/A.

Usage Guidelines

Use this command to add a CFM domain and an association to a CFM segment. It is used to enable DMM/DMR in the association that is configured in the CFM domain.

Example

The following command adds the domain cfm3 and the association as3 to the segment s2.

```
configure cfm segment s2 add domain cfm3 association as3
```

To delete the domain and/or association, use the command, `configure cfm segment delete domain association`.

History

This command was first available in ExtremeXOS 12.3.



Platform Availability

This command is available on all platforms.

configure cfm segment delete domain association

```
configure cfm segment segment_name delete domain association
```

Description

Deletes a CFM domain from a CFM segment.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
---------------------	---

Default

N/A.

Usage Guidelines

Use this command to delete a CFM domain from a CFM segment.

Example

The following command deletes the domain and association from the segment s2.

```
configure cfm segment s2 delete domain association
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.

configure cfm segment dot1p

```
configure cfm segment segment_name dot1p dot1p_priority
```

Description

Configures the priority for the segment.



Syntax Description

<code>segment-name</code>	An alpha numeric string identifying the segment name.
<code>dot1p_priority</code>	Priority value that is set in the DMM/DMR. The range is 0 to 7.

Default

The default is "6."

Usage Guidelines

Use this command to configure the dot1p priority that a DMM/DMR frame can get.

Example

The following command configures a dot1p priority of 3 for segment s2.

```
configure cfm segment s2 dot1p 3
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.

configure cfm segment frame-delay dot1p

```
configure cfm segment segment_name frame-delay dot1p dot1p_priority
```

Description

This command configures the class of service for a particular cfm segment. This value is used to fill the dot1p priority bit in the Ethernet header during transmission.

If the optional keyword **frame-delay** is not specified, the same value of Dot1p will be used for both DMM and LMM. The optional keyword allows configuring different values for DMM and LMM.

Syntax Description

<code>segment_name</code>	An alpha numeric string identifying the segment name.
<code>dot1p_priority</code>	Priority value that is set in the DMM/DMR. The range is 0 to 7.



Default

N/A.

Usage Guidelines

Use this command to configure the class of service for a particular cfm segment.

Example

```
configure cfm segment frame-delay dot1p 4
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.

configure cfm segment frame-delay/frame-loss transmit interval

```
configure cfm segment segment_name {frame-delay | frame-loss} transmit-interval
interval
```

Description

Configures the delay between two consecutive DMM/LMM frames.

Syntax Description

<i>segment_name</i>	Alphanumeric string identifying the segment name.
frame-delay	Y.1731 Ethernet frame delay measurement.
frame-loss	Y.1731 Ethernet frame loss measurement.
<i>interval</i>	Transmit interval in seconds, with a range of 1 to 90.

Default

N/A.



Usage Guidelines

Configures the delay between two consecutive DMM/LMM frames. The configured delay would be for both continuous and on-demand transmission. This command is optional, and if not configured, the default interval would be 10 seconds.

If the optional keyword `frame-delay` or `frame-loss` is not specified, the same value of `transmit-interval` will be used for both DMM and LMM. The optional keyword allows configuring different values for DMM and LMM.

Example

```
configure cfm segment cs2 frame-delay transmit-interval 10
configure cfm segment cs2 frame-loss transmit-interval 10
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms.

configure cfm segment frame-delay window

```
configure cfm segment segment_name frame-delay window window_size
```

Description

This command is used to configure the window size for calculating the alarm/clear threshold values for DMM and Severely Errored Second (SES) threshold for LMM. This window size denotes the total number of recent frames for which the threshold values will be measured.

If the optional keyword `frame-delay` or `frame-loss` is not specified, the same value of window size will be used for both DMM and LMM. The optional keyword allows configuring values for DMM and LMM.

Syntax Description

<i>segment_name</i>	Alphanumeric string identifying the segment name.
frame-delay	Y.1731 Ethernet frame delay measurement.
<i>window_size</i>	Window size for delay measurement; number of frames 1-1800 to be used.

Default

60.



Usage Guidelines

Use this command to configure the window size for calculating the alarm/clear threshold values for DMM and Severely Errored Second (SES) threshold for LMM.

Example

```
configure cfm segment cs2 frame-delay window 1000
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms.

configure cfm segment frame-loss dot1p

```
configure cfm segment segment_name frame-loss dot1p dot1p_priority
```

Description

This command configures the class of service for a particular cfm segment. This value is used to fill the dot1p priority bit in the Ethernet header during transmission.

If the optional keyword `frame-loss` is not specified, the same value of Dot1p will be used for both DMM and LMM. The optional keyword allows configuring different values for DMM and LMM.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
<i>dot1p_priority</i>	Priority value that is set in the DMM/DMR. The range is 0 to 7.

Default

N/A.

Usage Guidelines

Use this command to configure the class of service for a particular cfm segment.



Example

```
configure cfm segment frame-loss dot1p 4
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.

configure cfm segment frame-loss window

```
configure cfm segment segment_name frame-loss window window_size
```

Description

This command is used to configure the window size for calculating the alarm/clear threshold values for DMM and Severely Errored Second (SES) threshold for LMM. This window size denotes the total number of recent frames for which the threshold values will be measured.

If the optional keyword `frame-delay` or `frame-loss` is not specified, the same value of window size will be used for both DMM and LMM. The optional keyword allows configuring values for DMM and LMM.

Syntax Description

<i>segment_name</i>	Alphanumeric string identifying the segment name.
frame-loss	Y.1731 Ethernet frame loss measurement.
<i>window_size</i>	Window size for loss measurement; number of frames 1-1800 to be used.

Default

1200.

Usage Guidelines

Use this command to configure the window size for calculating the alarm/clear threshold values for DMM and Severely Errored Second (SES) threshold for LMM.

Example

```
configure cfm segment cs2 frame-loss window 900
```



History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms.

configure cfm segment frame-loss mep

```
configure cfm segment segment_name frame-loss [add|delete] mep mep_id
```

Description

This command is used to add/delete the local MEP for a given CFM segment.

Syntax Description

<i>segment_name</i>	Alphanumeric string identifying the segment name.
frame-loss	Y.1731 Ethernet frame loss measurement.

Default

N/A.

Usage Guidelines

The MEP with the given MEP ID should already be created in the system. The domain and association for the segment should be configured before executing this command. If the domain and association are not configured, the command throws an error.

Configuring of local MEP is mandatory to start the Frame Loss measurements.

Example

```
configure cfm segment cs2 add mep 3
configure cfm segment cs2 delete mep 3
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms.



configure cfm segment frame-loss consecutive

```
configure cfm segment segment_name frame-loss consecutive frames
```

Description

This command is used to configure the number of consecutive measurements to be used to determine the availability status of a CFM segment.

Syntax Description

<i>segment_name</i>	Alphanumeric string identifying the segment name.
frame-loss	Y.1731 Ethernet frame loss measurement.

Default

10.

Usage Guidelines

This configuration is optional.

Example

```
configure cfm segment cs2 frame-loss consecutive 10
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms.

configure cfm segment frame-loss ses-threshold

```
configure cfm segment segment_name frame-loss ses-threshold percent
```

Description

This command is used to configure the percentage of frames lost in a measurement period for it to be marked as SES (Severely Errored Second).



Syntax Description

<i>segment_name</i>	Alphanumeric string identifying the segment name.
ses	Severely errored second.
frame-loss	Y.1731 Ethernet frame loss measurement.

Default

30%.

Usage Guidelines

This configuration is optional.

Example

```
configure cfm segment cs2 frame-loss ses-threshold .02
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms.

configure cfm segment threshold

```
configure cfm segment segment_name [alarm-threshold | clear-threshold] value
```

Description

Configures the alarm threshold and clear threshold.

Syntax Description

alarm-threshold	Specifies the minimum threshold percentage.
clear-threshold	Specifies the maximum threshold percentage.
<i>value</i>	Specified the threshold percentage in a range of 1-99%.

Default

Alarm threshold is 10% of the total frames received during the current window.



Clear-threshold is 95% of the total frames received during the current window.

Usage Guidelines

Use this command to configure the alarm and clear threshold value for a CFM segment. Upon reaching the alarm threshold, an error message is generated and displayed once, and the state is maintained until the threshold reaches the clear threshold value.

This command is optional, and if not configured the default intervals are used.

Example

The following commands configure an alarm threshold of 15% and a clear-threshold of 90% for segment-first.

```
configure cfm segment segment-first alarm-threshold 15
configure cfm segment segment-first clear-threshold 90
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.

configure cfm segment timeout

```
configure cfm segment segment_name timeout msec
```

Description

Configures the timeout for a segment.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
<i>msec</i>	Specifies the number of milliseconds. The range is 1 to 65535.

Default

50 milliseconds.



Usage Guidelines

Use this command to configure the timeout value for the reception of a DMR frame. If a DMR frame is not received within this specified time, that frame is considered as an errored frame, and if the number of errored frames reaches the alarm threshold of the current window size, an alarm is generated.

This command is optional, and if not configured, timeout is set to the default.

Example

The following command configures a timeout value of 45 milliseconds for the s4 segment:

```
configure cfm segment s4 timeout 45
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.

configure cfm segment transmit-interval

```
configure cfm segment segment_name { frame-delay | frame loss } transmit-interval
interval
```

Description

Configures the transmission interval of DMM frames.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
frame-delay	Y.1731 Ethernet Frame Delay Measurement.
frame loss	Y.1731 Ethernet Frame Loss Measurement.
<i>interval</i>	Specifies the transmit interval in seconds. The range is 1 to 90.

Default

10 seconds.



Usage Guidelines

Use this command to configure the delay between two consecutive DMM frames. The configured delay is for both continuous and on-demand transmission. This command is optional, and if not configured the default interval is used.

Example

The following command configures a transmission interval of 5 seconds for segment s2.

```
configure cfm segment s2 transmit-interval 5
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.

configure cfm segment window

```
configure cfm segment segment_name window size
```

Description

Configures the measurement window size.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
<i>size</i>	Specifies the number of frames to be used for delay measurement. The range is 1 to 1800.

Default

60 frames.

Usage Guidelines

Use this command to configure the window size to be used for calculating the threshold values. This window size denotes the total number of recent frames for which the threshold values are to be measured.



This is an optional command and if not configured, the lower of either the default value or the total number of frames sent is used.



Note

MEPs with intervals 3 and 10 cannot be created in this domain as the domain name format is of dns type.

Example

The following command configures the measurement window size for the CFM segment segment-first at 55:

```
configure cfm segment segment-first window 55
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.

create cfm domain dns md-level

```
create cfm domain dns name md-level level
```

Description

Creates a maintenance domain (MD) in the DNS name format and assigns an MD level to that domain.

Syntax Description

<i>name</i>	Assigns the name you want for this domain, using the DNS name format. Enter alphanumeric characters for this format; the maximum is 43 characters.
<i>level</i>	Specifies the MD level you are assigning to this domain. Enter a value between 0 and 7.

Default

N/A.

Usage Guidelines

You can have up to 8 domains on a switch, and each one must have a unique MD level.



You assign each domain a maintenance domain (MD) level, which function in a hierarchy for forwarding CFM messages. The levels are from 0 to 7; with the highest number being superior in the hierarchy.

The IEEE standard 801.2ag specifies different levels for different network users, as follows:

- 5 to 7 for end users
- 3 and 4 for Internet service providers (ISPs)
- 0 to 2 for operators (entities carrying the information for the ISPs)



Note

MEPs with intervals 3 and 10 cannot be created in this domain as the domain name format is of dns type.

Example

The following command creates a domain, using the DNS name format, named extreme and assigns that domain an MD level of 2:

```
create cfm domain dns extreme md-level 2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

create cfm domain mac md-level

```
create cfm domain mac mac-addr int md-level level
```

Description

Creates a maintenance domain (MD) in the MAC address + 2-octet integer format and assigns an MD level to that domain.

Syntax Description

<i>mac-addr</i>	Enter a MAC address in the format XX:XX:XX:XX:XX:XX to specify part of the domain name.
<i>int</i>	Enter the 2-octet integer you want to append to the MAC address to specify the domain name.
<i>level</i>	Specifies the MD level you are assigning to this domain. Enter a value between 0 and 7.



Default

N/A.

Usage Guidelines

You can have up to 8 domains on a switch, and each one must have a unique MD level.

You assign each domain a maintenance domain (MD) level, which function in a hierarchy for forwarding CFM messages. The levels are from 0 to 7; with the highest number being superior in the hierarchy.

The IEEE standard 801.2ag specifies different levels for different network users, as follows:

- 5 to 7 for end users
- 3 and 4 for Internet service providers (ISPs)
- 0 to 2 for operators (entities carrying the information for the ISPs)

Example

The following command creates a domain, using the MAC + 2-octet integer format, with the MAC address of 11:22:33:44:55:66 and an integer value of 63; it also assigns that domain an MD level of 2:

```
create cfm domain mac 11:22:33:44:55:66 63 md-level 2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

create cfm domain string md-level

```
create cfm domain string str_name md-level level
```

Description

Creates a maintenance domain (MD) in the string name format and assigns an MD level to that domain.

Syntax Description

<i>str_name</i>	Enter a character string to specify part of the domain name. The maximum length is 43 characters.
<i>level</i>	Specifies the MD level you are assigning to this domain. Enter a value between 0 and 7.



Default

N/A.

Usage Guidelines

You can have up to 8 domains on a switch, and each one must have a unique MD level.

You assign each domain a maintenance domain (MD) level, which function in a hierarchy for forwarding CFM messages. The levels are from 0 to 7; with the highest number being superior in the hierarchy.

The IEEE standard 801.2ag specifies different levels for different network users, as follows:

- 5 to 7 for end users
- 3 and 4 for Internet service providers (ISPs)
- 0 to 2 for operators (entities carrying the information for the ISPs)

Example

The following command creates a domain, using the string format having a value of extreme; it also assigns that domain an MD level of 2:

```
create cfm domain string extreme md-level 2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

create cfm segment destination

```
create cfm segment segment_name destination mac_addr {copy segment_name_to_copy}
```

Description

Creates a CFM segment.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
<i>mac_addr</i>	Specifies the MAC address.
<i>segment_name_to_copy</i>	Specifies the CFM segment whose configuration is to be copied.



Default

N/A.

Usage Guidelines

Use this command to explicitly create a CFM segment where the segment name is a 32-byte long alpha-numeric character string.

Example

The following command creates a CFM segment named segment-new using MAC address 00:11:22:11:33:11 and copying segment-old:

```
create cfm segment segment-new destination 00:11:22:11:33:11 copy segment-old
```

Here, the *copy existing cfm segment* is an optional parameter, and if used, the following configurations from the existing CFM segment are copied to the newly created segment:

- DMM transmission interval
- Class of service
- Threshold values
- Measurement window size
- Timeout value



Note

The copy option is not shown in "show config" as it is used only for copying the existing values when creating a segment.

If you later configure any of the above mentioned information in segment-new, the old value(s) which were copied from segment-old will be overwritten with the new one in segment-new, as is done for any other commands. The same will not be true on the reverse case. If you modify the values of segment-old, the modified value will NOT be propagated to the CFM segments which use segment-old's configurations. In other words, the configurations of segment-old that are at the time of creating segment-new will alone be copied and not any other changes that are made to segment-old later on.

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.

delete cfm domain

```
delete cfm domain domain
```



Description

Deletes the specified maintenance domain (MD) from the switch, as well as all configuration setting related to this MD.

Syntax Description

<i>domain</i>	Enter the name of the domain you want to delete.
---------------	--

Default

N/A.

Usage Guidelines

This command deletes all configuration settings related to the domain—for example, all MAs, MIPs, and MEPs—as well as the domain itself.

Example

The following command deletes the domain atlanta (as well as all settings related to this domain):

```
delete cfm domain atlanta
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

delete cfm segment

```
delete cfm segment [segment_name | all]
```

Description

Deletes one or all CFM segments.

Syntax Description

<i>segment_name</i>	An alpha-numeric string identifying the segment name.
all	Specifies all CFM segments.



Default

N/A.

Usage Guidelines

Use this command to delete one or all CFM segments.

Example

The following command deletes the CFM segment named segment-new:

```
delete cfm segment segment-new
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.

disable cfm segment frame-delay measurement

```
disable cfm segment frame-delay measurement segment_name
```

Description

Stops DMM frame transmission.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
---------------------	---

Default

N/A.

Usage Guidelines

Use this command to stop transmission of DMM frames for a selected CFM segment. This command stops transmission that has been triggered using the command `enable cfm segment frame-delay measurement`.

This stops the transmission for both continuous and on-demand mode.



Example

The following command stops frame transmission on the CFM segment segment-first:

```
disable cfm frame-delay measurement segment-first
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.

disable cfm segment frame-loss measurement mep

This stops the transmission for both continuous and on-demand mode.

```
disable cfm segment frame-loss measurement segment_name mep mep_id
```

Description

This command stops the transmission of the LMM frames for a particular cfm segment.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
---------------------	---

Default

N/A.

Usage Guidelines

This below command stops the transmission of the LMM frames for a particular cfm segment. This stops the transmission for both continuous and on-demand mode.

Example

```
disable cfm segment cs2 frame-loss measurement mep 3
```

History

This command was first available in ExtremeXOS 15.1.



Platform Availability

This command is available on all platforms.

enable/disable bfd vlan

```
[enable | disable] bfd vlan vlan_name
```

Description

Enables or disables BFD on a VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN name.
------------------	--------------------------

Default

N/A.

Usage Guidelines

Use this command to enable or disable BFD on a VLAN.

Example

The following command enables the bfd on the VLAN named finance:

```
enable bfd vlan finance
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

enable cfm segment frame-delay measurement

```
enable cfm segment frame-delay measurement segment_name { mep med_id }  
[continuous | count value]
```



Description

Triggers DMM frame transmission.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
mep	Specifies the maintenance association End Point that helps trigger a particular MEP level session on that segment.
<i>mep_id</i>	Specifies the MEP-ID. The range is 1-8191]. The default is all MEPs on the segment.
continuous	Specifies that frames are to be sent continuously until stopped.
count	Specifies that a number of frames are to be sent.
<i>value</i>	Specifies the number of frames to send. The range is 1 to 4294967295.

Default

N/A.

Usage Guidelines

Use this command to trigger DMM frames at the specified transmit interval configured using the command `configure cfm segment transmit-interval`.

Continuous transmission continues until it is stopped with the command `disable cfm segment frame-delay measurement` or `delete cfm segment`.

Note



If you try to trigger the DMM frames for a segment that is not completely configured, the frames are not transmitted for that segment, and an error message is displayed on the console.

Example

The following command triggers continuous frame transmission on the CFM segment `segment-first`:

```
enable cfm frame-delay measurement segment-first continuous
```

History

This command was first available in ExtremeXOS 12.3.

The `mep` keyword was added in ExtremeXOS 15.4.



Platform Availability

This command is available on all platforms.

enable cfm segment frame-loss measurement mep

If the user specifies the mode as continuous, the LMM transmission will continue till it is stopped by the user.

```
enable cfm segment frame-loss measurement segment_name mep mep_id [continuous | count frames]
```

Description

This command is used to trigger LMM frames at the configured transmit-interval.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
continuous	Specifies that frames are to be sent continuously until stopped.
count	Specifies that a number of frames are to be sent.
value	Specifies the number of frames to send. The range is 1 to 4294967295.

Default

N/A.

Usage Guidelines

This command is used to trigger LMM frames at the configured transmit-interval. If the user specifies the mode as continuous, the LMM transmission will continue till it is stopped by the user.



Note

If the user tries to trigger the LMM frames for a segment which is not completely configured, the frames will not be transmitted for that segment, and an error message will be thrown.

Example

```
enable cfm segment cs2 frame-loss measurement mep 3 count 10
enable cfm segment cs2 frame-loss measurement mep 3 continuous
```

History

This command was first available in ExtremeXOS 15.1.



Platform Availability

This command is available on all platforms.

ping mac port

The ping, or loopback message (LBM), goes from the MEP configured on the port toward the given MAC address.

```
ping mac mac port port {domain} domain_name {association} association_name
```

Description

Allows you to ping on the Layer 2 level throughout the specified domain and MA.

Syntax Description

<i>mac</i>	Enter the unique system MAC address on the device you want to reach. Enter this value in the format XX:XX:XX:XX:XX:XX.
<i>port</i>	Enter the port number of the MEP from which you are issuing the ping.
domain	Enter this keyword.
<i>domain_name</i>	Enter the name of the domain from which you are issuing the ping.
association	Enter this keyword.
<i>association_name</i>	Enter the name of the association from which you are issuing the ping.

Default

N/A.

Usage Guidelines

You must have CFM parameters configured prior to issuing a Layer 2 ping.

In order to send a Layer 2 ping, you must specify the port (MEP), the domain, and the MA from which you are issuing the ping. An UP MEP sends the ping to all ports (except the sending port) on the VLAN that is assigned to the specified MA, and a DOWN MEP sends the ping out from that port from that MA toward the specified MAC address.

All MIPs along the way forward the LBM to the destination. The destination MP responds back to the originator with a loopback reply (LBR).

This command sends out a ping from the MEP configured on the specified port toward the specified MAC address. If you attempt to send a ping message from a port that is not configured as a MEP, the system returns an error message. If the specified MAC address is not present in the Layer 2 forwarding table (FDB), the system cannot send the ping (applies to UpMEP, not DownMEP).



Example

The following command sends a Layer 2 ping to the unique system MAC address 00:04:96:1F:A4:31 from the previously configured UP MEP (port 2:4) in the speed association in the atlanta domain:

```
ping mac 00:04:96:1F:A4:31 port 2:4 atlanta speed
```

The following is sample output from the Layer 2 ping command:

```
BD-12802.48 # ping mac 00:04:96:1e:14:70 port 2:12 "extreme" 100
Send L2 Ping from Down MEP on 2:12, waiting for responses [press Ctrl-C to
abort].
42 bytes from 00:04:96:1e:14:70, seq=4 time=17 ms
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

show bfd

show bfd

Description

Displays information on existing BFD sessions.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to show the status of the current BFD sessions.

The following session states are displayed:

- Init—The state when BFD is establishing the session.
- Down—The state when BFD detects that the session is down.



- Admin Down—The state when the user disables BFD on that interface.
- Up—The state when the BFD session is established.

Example

The following command displays information on current BFD sessions:

```
show bfd
```

Following is sample output from this command:

```
Number of sessions           : 2
Sessions in Init State       : 0
Sessions in Down State       : 0
Sessions in Admin Down State : 1
Sessions in Up State         : 1
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

show bfd counters

show bfd counters

Description

Displays the readings of the global BFD counters.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to display global BFD counters.

To clear the counters, use the `clear counters bfd` command.



Example

The following command displays BFD global counters:

```
show bfd counters
```

Following is sample output from this command:

```
Valid Tx Pkt           : 177      Valid Rx Pkt           : 177
Rx Invalid TTL         : 0        Rx Invalid UDP SrcPort : 0
Interface Not found   : 0        Rx Invalid Version     : 0
Rx Invalid Length Pkt : 0        Rx Invalid Multiplier  : 0
Rx Invalid Demand Mode : 0      Rx Poll & Final set    : 0
Rx Invalid My Discriminator : 0    Rx Invalid Your Discriminator : 0
Rx Invalid Auth Length : 0        Rx session Not Found   : 6
Auth Type Fails       : 0        Authentication Fails   : 0
Tx Fails              : 0        Rx Discarded Pkt      : 0
```

Note



The Rx session Not Found counter is incremented when the BFD session corresponding to the received BFD packet is not found. The Rx Discarded Pkt counter is incremented when the neighbor state indicated in the BFD packet is not one of the expected/allowed states.

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

show bfd session client

```
show bfd session client [mpls | ospf {ipv4 | ipv6} | static {ipv4 | ipv6}] {vr
[vrname | all]}
```

Description

Displays the BFD session information for a specified client.

Syntax Description

mpls	Specifies an MPLS client.
ospf	OSPF Protocol.
ipv4	Displays sessions requested by IPv4 version client, e.g. OSPFv2 (Default).
ipv6	Displays sessions requested by IPv6 version client, e.g. OSPFv3.



static	Specifies a static route.
vrname	Specifies the name of the virtual router.

Default

IPv4.

Usage Guidelines

Use this command to display session information for a specified client.

Example

The following command displays the BFD sessions for an MPLS client on all VRs:

```
show bfd session client mpls vr all
```

Following is sample output from this command:

```
Neighbor      Interface      Detection      Status
-----      -
10.10.10.2    vlan10         3000           Up
=====
```

NOTE: All timers in milliseconds.

History

This command was first available in ExtremeXOS 12.4.

Support for BFD protected static route was added in ExtremeXOS 12.5.3.

The **ospf** keyword was added in ExtremeXOS 15.3.2.

Platform Availability

This command is available on all platforms.

show bfd session counters vr all

```
show bfd session {ipv4 | ipv6} {ipaddress} counters {vr [vrname | all]}
```

Syntax Description

ipv4	Displays all IPv4 sessions.
ipv6	Displays all IPv6 sessions.
<i>ipaddress</i>	Displays sessions in specified VR.



Default

Displays all IPv4 sessions counters by default if IPv4 or IPv6 is not specified.

Usage Guidelines

Use this command to display BFD session counters.

To clear the counters, use the `clear counters bfd` command.

Example

The following command displays the session counters:

```
show bfd session counters vr all
```

Following is sample output from this command:

```
Neighbor : 10.10.10.1      Interface : vlan10Vr-Name :   bfd_vr10
Valid Rx Pkt           : 87
Total Tx Pkt           : 87
Auth Type Fails        : 0
Authentication Fails   : 0
Discarded Pkt          : 0
```

History

This command was first available in ExtremeXOS 12.4.

IPv6 version of this command was added in ExtremeXOS 15.3.2.

Platform Availability

This command is available on all platforms.

show bfd session detail vr all

```
show bfd session {ipv4 | ipv6} {ipaddress } detail {vr [vrname | all]}
```

Description

Displays detailed information about a BFD session.

Syntax Description

ipv4	Displays all IPv4 sessions.
ipv6	Displays all IPv6 sessions.



<i>ipaddress</i>	Displays sessions in specified VR.
<i>vrname</i>	Displays sessions in specified VR.

Default

Displays all IPv4 sessions by default if ipv4 or ipv6 is not specified.

Usage Guidelines

Use this command to display BFD session information in detail.

Example

The following command displays the BFD session information in detail:

```
show bfd session detail vr all
```

Following is sample output from this command:

```
Neighbor      : 10.10.10.1          Local      : 10.10.10.2
Vr-Name       : bfd_vr10      Interface   : vlan10
Session Type  : Single Hop    State      : Up
Detect Time   : 3000 mc       Age        : 250 ms
Discriminator (local/remote) : 1 / 1
Demand Mode (local/remote)  : 0 / 0
Poll (local/remote)         : 0 / 0
Tx Interval (local/remote)  : 1000 / 1000 ms
Rx Interval (local/remote)  : 1000 / 1000 ms
oper Tx Interval            : 1000 ms
oper Rx Interval            : 1000 ms
Multiplier (local/remote)  : 3 / 3
Local Diag                  : 0 (No Diagnostic)
Remote Diag                  : 0 (No Diagnostic)
Authentication               : None
Clients                      : MPLS,
Uptime                       : 00 days 00 hours 00 minutes 41 seconds
Up Count                     : 1
Last Valid Packet Rx         : 00:51:49.300000
Last Packet Tx               : 00:51:48.820000
```

The following command displays a specified IPv6 BFD session in detail:

```
sh bfd session fe80::204:96ff:fe1f:a800%v2 detail
```

```
Neighbor      : fe80::204:96ff:fe1f:a800
Local        : fe80::204:96ff:fe27:2c6a
VR-Name      : VR-Default      Interface   : v2
Session Type  : Single Hop    State      : Up
Detect Time   : 60000 ms      Age        : 460 ms
Discriminator (local/remote) : 1 / 1
Demand Mode (local/remote)  : Off / Off
Poll (local/remote)         : Off / Off
Tx Interval (local/remote)  : 20000 / 1000 ms
```




```

Rx Interval (local/remote)      : 20000 / 1000 ms
Oper Tx Interval                : 20000 ms
Oper Rx Interval                : 20000 ms
Multiplier (local/remote)     : 3 / 3
Local Diag                     : 0 (No Diagnostic)
Remote Diag                    : 0 (No Diagnostic)
Authentication                  : None
Clients                        : OSPFv3
Uptime                          : 00 days 01 hours 35 minutes 43 seconds
Up Count                       : 9
Last Valid Packet Rx           : 12:27:36.464105
Last Packet Tx                 : 12:27:19.34236

```

History

This command was first available in ExtremeXOS 12.4.

IPv6 version was added in ExtremeXOS 15.3.2.

Platform Availability

This command is available on all platforms.

show bfd session vr all

```
show bfd session { ipv4 | ipv6 } { ipaddress } { vr [vrname | all ] }
```

Description

Displays general information about a BFD session.

Syntax Description

ipv4	Displays all IPv4 sessions.
ipv6	Displays all IPv6 sessions.
<i>ipaddress</i>	Displays session that has specified address as destination address.
vrname	Displays sessions in specified VR.

Default

Displays all IPv4 sessions by default if **ipv4** or **ipv6** keyword is not specified.

Usage Guidelines

Use this command to display general information about a BFD session.



Example

The following command displays general information about the BFD session:

```
show bfd session vr all
```

Following is sample output from this command:

```
Neighbor      Interface      Clients  Detection  Status      VR
=====
30.30.30.2    bfdVlan        ----s    0           Down        VR-Default
=====
Clients Flag: m - MPLS, o - OSPF, s - Static
NOTE: All timers in milliseconds.
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

show bfd vlan

```
show bfd vlan {vlan_name}
```

Description

Displays the BFD settings for the specified VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN name.
------------------	--------------------------

Default

N/A.

Usage Guidelines

Use this command to display the BFD settings on a specified VLAN.



Example

The following command displays the BFD settings for the VLAN vlan10:

```
show bfd vlan vlan10
```

Following is sample output from this command:

```
VLAN           : vlan10
BFD            : Enabled
Tx Interval    : 1000
Rx Interval    : 1000
Detection Multiplier : 3
Authentication : None
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

show bfd vlan counters

```
show bfd vlan {vlan_name} counters
```

Description

Displays BFD counters on a specified VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN name.
------------------	--------------------------

Default

N/A.

Usage Guidelines

Use this command to display counter readings for a specified VLAN.



Example

The following command displays the counter readings for the VLAN vlan10:

```
show bfd vlan vlan10 counters
```

Following is sample output from this command:

```
VLAN                               : vlan10
Valid Rx Pkt                       : 144
Total Tx Pkt                       : 144
Auth Type Fails                    : 0
Authentication Fails               : 0
Discarded Pkt                      : 0
Rx session Not Found               : 6
```

Note



The Discarded Pkt counter is incremented when the neighbor state indicated in the BFD packet is not one of the expected/allowed states. The Rx session Not Found counter is incremented when the BFD session corresponding to the received BFD packet is not found.

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

show cfm

```
show cfm { domain_name { association_name {{ports port_list {[intermediate-point
| [end-point [up|down]]}}}}
```

Description

Displays the current CFM configuration on the switch.

Syntax Description

<i>domain_name</i>	Enter the name of the domain you want to display.
<i>association_name</i>	Enter the name of the association you want to display.
<i>port_list</i>	Enter the ports in the domain and association you want to display.
up	Enter this to display the UP MEP for the specified MA.



down	Enter this to display the DOWN MEP for the specified MA.
intermediate-point	Enter this to display the MIPs for the specified MA.

Default

N/A.

Usage Guidelines

This command displays the following information:

- Domain names
- MA levels
- Association names
- VLAN names
- Transmit Interval
- UP MEPs
- MEPIDs
- MEP transmit intervals
- MEP State
- DOWN MEPs
- Intermediate points (MIPs)
- Total number of CFM ports on the switch
- Destination MAC Type
- VPLS-based MPs
- Sender ID information
- ISID Intermediate Point

See [Supported Instances for CFM](#) for the number of domains, ports, MEPs, MIPs, and associations supported on the switch.

Example

The `show cfm` command displays the current CFM configuration on the switch:

```
* (debug) switch # show cfm
Domain: "sVlanDom5", MD Level: 5
Association: "sVlanAssoc", Destination MAC Type: Multicast, SVLAN "s1" with 2
cfm ports
Transmit Interval: 60000 ms
port 1:5; Intermediate Point ( Dynamic )
port 1:8; Up End Point, mepid: 1, transmit-interval: 60000 ms (from
association),
MEP State: Enabled, CCM Message: Enabled, Send SenderId TLV: Disabled
ISID Intermediate Point;
Domain: "vplsDom6", MD Level: 6
Association: "vplsAssoc1", Destination MAC Type: Multicast, VLAN "v1" with 1
cfm ports
and VPLS MIP; Transmit Interval: 1000 ms
```



```

port 1:3; Up End Point, mepid: 1, transmit-interval: 1000 ms (from
association),
MEP State: Enabled, CCM Message: Enabled, Send SenderId TLV: Disabled
vpls name: vp101; Intermediate Point; Link Trace Encoding : VPLS-Name:System-
Name
Association: "vplsAssoc2", Destination MAC Type: Multicast, VLAN "v2" with 1
cfm ports
and VPLS MIP; Transmit Interval: 1000 ms
port 1:2; Up End Point, mepid: 1, transmit-interval: 1000 ms (from
association),
MEP State: Enabled, CCM Message: Enabled, Send SenderId TLV: Disabled
vpls name: vp100; Intermediate Point; Link Trace Encoding :VPLS-Name:Private-
IP
Total Number of Domain                : 2
Total Number of Association            : 3
Total Number of Up MEP                 : 3
Total Number of Down MEP               : 0
Total Number of MIP                    : 4
Total Number of Number of CFM port    : 8
Total Number of VPLS MIP(Static/Up): 2 / 4

```

History

This command was first available in ExtremeXOS 11.4.

Transmit Interval and MEP State were added in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.

show cfm detail

```

show cfm {domain_name {association_name {ports port_list} {[end-point [up |
down]]}}}} detail

```

Description

Displays the MEP CCM database.

Syntax Description

<i>domain_name</i>	Enter the name of the domain for which you want to display the MEP CCM databases.
<i>association_name</i>	Enter the name of the association for which you want to display the MEP CCM databases.
<i>port_list</i>	Enter the ports in the domain/association for which you want to display the CCM databases.



up	Enter this to display the CCM database on the UP MEP for the specified MA.
down	Enter this to display the CCM database on the DOWN MEP for the specified MA.

Default

N/A.

Usage Guidelines

If you do not specify any parameters or variables, the system displays information on all CCM databases on the switch.

This command displays the following items of the CCM database:

- The name of the domain and association
- Port number
- MP and type
- MAC address of remote end points
- MEP IDs
- Lifetime for CCM messages from each remote end point
- Actual age of CCM messages



Note

The TTL for the CCM messages from the MP you are working on is 3.5 times the transmission interval.

Example

The following command displays the CCM databases on the switch:

```
show cfm detail
```

The following is sample output from this command:

```
BD-12802.48 # sh cfm detail
Domain/      Port    MP  Remote End-Point  Remote End-Point MEP
Life        Flags
Association          MAC Address      IP Address      ID    time  Age
=====
=====
extreme
100          2:1      UE  00:04:96:10:e5:f0  0.0.0.0          2    3500
950      DMA
2:3      UE  00:04:96:10:e5:f0  0.0.0.0          1    3500  50    DMA
=====
=====
Maintenance Point: (UE) Up End-Point, (DE) Down End-Point
```



Flags: (S) Static Entry (D) Dynamic Entry
 CCM Destination MAC: (U) Unicast (M) Multicast
 Status: (A) Active, (I) Inactive
 NOTE: The Domain and Association names are truncated to 13 characters,
 Lifetime
 and Age are in milliseconds.

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

show cfm groups

```
show cfm groups {group_name}
```

Description

This command displays the details of specified or all groups. The information contains group name, group status, LMEP id, the physical port of the LMEP, RMEP ids, registered clients, domain and association names.

Syntax Description

<i>group_name</i>	Group name, maximum of 31 characters.
-------------------	---------------------------------------

Default

N/A.

Usage Guidelines

Use this command to display the details of specified or all groups. The information contains group name, group status, LMEP id, the physical port of the LMEP, RMEP ids, registered clients, domain and association names.

Example

```
X480-48t.1 # sh cfm groups
Group : eapsCfmGrp1      Status : UP
Local MEP      : 11      port   : 41
Remote MEPs    : 10
Client(s)      : eaps
Domain         : MD1
```




```

Association      : MD1v2
Group : eapsCfmGrp2      Status : UP
Local MEP       : 12      port   : 31
Remote MEPs    : 13
Client(s)      : eaps
Domain         : MD1
Association    : MD1v2

```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on all platforms.

show cfm segment

```
show cfm segment {segment_name}
```

Description

Displays information for CFM segments.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
---------------------	---

Default

N/A.

Usage Guidelines

Use this command to display information for the selected CFM segment.

If a segment name is not specified, the information for all of the segments that are currently configured are displayed.

Example

The following command displays information for an active CFM segment that is configured to transmit with a specific count:

```

show cfm segment s2
CFM Segment Name           : s2

```



```

Domain Name           : pbt-d2
Association           : pbt-d2-protecting
MD Level              : 2
Destination MAC       : 00:04:96:1e:14:70
DMM Transmission      : In Progress
Transmission mode     : Continuous
Frames Transmitted    : 2
Frames Received       : 2
DMM TX Interval       : 2secs
DMR RX Timeout        : 10 msec
Alarm Threshold        : 10 %
Clear Threshold        : 95 %
Measurement Window Size : 60
Class of Service      : 0
Tx Start Time         : Sun Apr 19 21:18:58 2009
Min Delay              : Sun Apr 19 21:18:58 2009
Max Delay              : Sun Apr 19 21:19:00 2009
Last Alarm Time       : None
Alarm State           : Not Set
Lost Frames in Current window : 0
-----
Total Configured Segments : 2
Total Active Segments    : 1

```

The following command displays information for a disabled segment:

```

BD-12804.1 # sh cfm seg s2
CFM Segment Name      : s2
Domain Name           : pbt-d2
Association           : pbt-d2-protecting
MD Level              : 2
Destination MAC       : 00:04:96:1e:14:70
DMM Transmission      : Disabled
Frames Transmitted    : 10
Frames Received       : 10
DMM TX Interval       : 2secs
DMR RX Timeout        : 10 msec
Alarm Threshold        : 10 %
Clear Threshold        : 95 %
Measurement Window Size : 60
Class of Service      : 0
Tx Start Time         : Sat Apr 18 05:39:54 2000
Min Delay              : Sat Apr 18 05:40:12 2000
Max Delay              : Sat Apr 18 05:39:56 2000
Last Alarm Time       : None
Alarm State           : Not Set
Lost Frames in Current window : 1
-----
Total Configured Segments : 2
Total Active Segments    : 0

```

History

This command was first available in ExtremeXOS 12.3.



Platform Availability

This command is available on all platforms.

show cfm segment frame-delay

```
show cfm segment frame-delay {segment_name}
```

Description

This command displays frame-delay information for the given CFM segment.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
---------------------	---

Default

N/A.

Usage Guidelines

Use this command to display frame-delay information for the given CFM segment.

Example

Example output not yet available and will be provided in a future release.

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.

show cfm segment frame-delay/frame-loss mep id

```
show cfm segment {{segment_name} | {frame-delay {segment_name}} | {frame-loss  
{segment_name {mep mep_id}}}}
```

Description

This command is used to display the current status and configured values of a cfm segment.



Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
---------------------	---

Default

N/A.

Usage Guidelines

Use this command to display the current status and configured values of a cfm segment.



Note

In this command, the row "pending frames" will be displayed only for on-demand mode of transmission.

A segment is considered as active if any of the MEPs in the segment is enabled for Frame Loss measurement. Active Segment count will be incremented by one only even if there are multiple MEPs enabled for Frame Loss. For example, assume that there are 3 segments created - seg1, seg2 and seg3. Segment "seg1" is enabled for Frame Delay measurement. Segment "seg3" has 10 MEPs added with 4 enabled for Frame Loss measurement, the following are the valid counts. Switch wide "Total Configured Segments" will be 3 and "Total Active Segments" will be 2. For Segments "seg1" and "seg2", "Total Configured MEPs" and "Total Active MEPs" will be 0. For segment "seg3", "Total Configured MEPs" will be 10 and "Total Active MEPs" will be 4.

By default, both the Frame Delay and Frame Loss sections are displayed for all the CFM segments. The user has option to filter out based on Segment Name or Frame Delay / Frame Loss.

The behavior for each of the optional parameters is explained below:

- Show cfm segment: Displays frame-delay and frame-loss information for all the CFM segments.
- Show cfm segment *segment_name*: Displays frame-delay and frame-loss information for the given CFM segment.
- Show cfm segment frame-delay: Displays frame-delay information for all the CFM segments.
- Show cfm segment frame-delay *segment_name*: Displays frame-delay information for the given CFM segment.
- Show cfm segment frame-loss: Displays frame-loss information for all the CFM segments (and all the MEPs under each of the segment).
- Show cfm segment frame-loss *segment_name*: Displays frame-loss information for the given CFM segment (and all the MEPs under the given segment).
- Show cfm segment frame-loss *segment_name* mep *mep_id*: Displays frame-loss information for the given CFM segment - MEP ID combination.

Example

```
Switch#show cfm segment sc-rtp
CFM Segment Name      : sc-rtp
Domain Name           : pbt-d2
```



```

Association                : pbt-d2-protecting
MD Level                   : 2
Destination MAC           : 00:04:96:1e:14:70
Frame Delay:
DMM Transmission          : In Progress
Transmission mode         : Continuous
Frames Transmitted        : 24
Frames Received           : 15
DMM Tx Interval           : 2 secs
DMR Rx Timeout            : 10 msec
Alarm Threshold           : 10 %
Clear Threshold           : 95 %
Measurement Window Size   : 60
Class of Service          : 0
Tx Start Time             : Fri Apr 17 01:29:45 2009
Min Delay                 : Fri Apr 17 01:30:29 2009
Max Delay                 : Fri Apr 17 01:30:03 2009
Last Alarm Time           : Fri Apr 17 01:29:59 2009
Alarm State               : Set
Lost Frames in Current Window : 9
Frame Loss:
LMM Tx Interval           : 2 secs
LMR Rx Timeout            : 10 msec
SES Threshold             : 30 %
Consecutive Available Count : 10
Measurement Window Size   : 60
Class of Service          : 0
Total Configured MEPs     : 2
Total Active MEPs         : 2
MEP ID                    : 100
LMM Transmission          : In Progress
Transmission mode         : Continuous
Frames Transmitted        : 24
Frames Received           : 15
Availability Status        : Available/Unavailable
Unavailability Start Time : Fri Apr 17 01:10:45 2011
Unavailability End Time   : Fri Apr 17 01:20:45 2011
Tx Start Time             : Fri Apr 17 01:10:45 2011
Min Near-End Frame Loss   : Fri Apr 17 01:29:45 2009
Max Near-End Frame Loss   : Fri Apr 17 01:39:45 2009
Min Far-End Frame Loss    : Fri Apr 17 01:49:45 2009
Max Far-End Frame Loss    : Fri Apr 17 01:59:45 2009
MEP ID                    : 200
LMM Transmission          : In Progress
Transmission mode         : Continuous
Frames Transmitted        : 24
Frames Received           : 15
Availability Status        : Available/Unavailable
Unavailability Start Time : Fri Apr 17 01:10:45 2011
Unavailability End Time   : Fri Apr 17 01:20:45 2011
Tx Start Time             : Fri Apr 17 01:10:45 2011
Min Near-End Frame Loss   : Fri Apr 17 01:29:45 2009
Max Near-End Frame Loss   : Fri Apr 17 01:39:45 2009
Min Far-End Frame Loss    : Fri Apr 17 01:49:45 2009
Max Far-End Frame Loss    : Fri Apr 17 01:59:45 2009
-----
Total Configured Segments : 1

```



Total Active Segments : 1

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.

show cfm segment frame-delay statistics

```
show cfm segment frame-delay statistics {segment-name} {mep mep_id}
```

Description

This command displays frame-delay information for the given CFM segment.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
mep	Maintenance association End Point.
<i>mep_id</i>	MEP-ID. The range is 1-8191.

Default

N/A.

Usage Guidelines

Use this command to display the delay for the last received frame, the minimum, maximum and average delay, and the delay variance during the current transmission. When the segment name is not specified, only the segments which have valid statistics alone are displayed. When the segment name is specified, that particular segment's information, although not present, is displayed.

Example

The following command displays the frame delay statistics for the CFM segment:

```
show cfm segment frame-delay statistics
```



Following is sample output for this command:

```

-----
Segment Name      Mep      Recent   Min      Max      Mean      Jitter    Errored
                  ID       Delay    Delay    Delay    Delay    Delay     Frames*
                  (ms)    (ms)    (ms)    (ms)    (ms)    (ms)
-----
segment1          ----    0.000   0.000   0.000   0.000   0.000     0
segment2          100    0.000   0.000   0.000   0.000   0.000
0
segment3          200    0.000   0.000   0.000   0.000   0.000     0
                  100    0.000   0.000   0.000   0.000   0.000     0
                  300    0.000   0.000   0.000   0.000   0.000     0
-----
Flags: (*) % of frames beyond alarm threshold in the current measurement
window

```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.

show cfm segment frame-loss

```
show cfm segment frame-loss {segment_name}
```

Description

This command displays frame-loss information for the given CFM segment.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
---------------------	---

Default

N/A.

Usage Guidelines

Use this command to display frame-delay information for the given CFM segment.



Example

```

sho cfm seg frame-loss
CFM Segment Name      : cs2
Domain Name           : dom2
Association            : a2
MD Level              : 2
Destination MAC       : 00:04:96:52:a7:64
Frame Loss:
LMM Tx Interval       : 10 secs
SES Threshold         : 1.000000e-02
Consecutive Available Count : 4
Measurement Window Size : 1200
Class of Service      : 6
Total Configured MEPs : 1
Total Active MEPs     : 1
MEP ID                : 3
LMM Transmission      : In Progress
Transmission Mode     : Continuous
Frames Transmitted    : 483
Frames Received       : 483
Availability Status    : Available
Unavailability Start Time : None
Unavailability End Time : None
Tx Start Time         : Mon Apr 23 12:28:28 2012
-----
Total Configured Segments : 1
Total Active Segments     : 1
E4G-200.31 #
E4G-200.31 #

```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.

show cfm segment frame-loss statistics

```
show cfm segment frame-loss statistics {segment-name}
```

Description

Displays shows frame-loss statistics.



Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
---------------------	---

Default

N/A.

Usage Guidelines

The below output is an example for displaying the frame-loss stats for the cfm segments. This command shows the recent, minimum, maximum and average near-end and far-end frame loss ratios during the current transmission. The stats for a particular segment will be preserved till the user triggers the next LMM transmission or until it does a clear counter.

Example

The following command displays the frame loss statistics for the CFM segment:

```

LEFT.93 # show cfm segment frame-loss statistics
-----
Segment Name      MEP      Last      Last      Min      Max      Min      Max      Mean
Mean
                  ID        NE        FE        NE        NE        FE        FE        NE
FE
                  FLR      FLRFLRFLRFLRFLRFLRFLR      NLR
-----
seg1               111       10       10       10       10       10       10       10      10
seg1               222       10       10       10       10       10       10       10      10
seg2               333       10       10       10       10       10       10       10      10
-----
Legend: FE - Far End, NE - Near End, FLR - Frame Loss Ratio

Window FE FLR  Last FE Tx  Last FE Rx
-----
cs2                3  0.000000e+00  509467221  526672689
0.000000e+00  501936465  544907407
-----
Legend: FE - Far End, NE - Near End, FLR - Frame Loss Ratio
    
```

History

This command was first available in ExtremeXOS 15.1.



Platform Availability

This command is available on all platforms.

show cfm segment mep

```
show cfm segment {segment_name} {mep mep_id }
```

Description

This command displays frame-delay information for the given CFM segment – MEP ID combination.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
---------------------	---

Default

N/A.

Usage Guidelines

Use this command to display frame-delay information for the given CFM segment – MEP ID combination.

Example

```
Switch#showcfm segment sc-rtp
CFM Segment Name      : sc-rtp
Domain Name           : pbt-d2
Association            : pbt-d2-protectingMD
Level                 : 2
Destination MAC       : 00:04:96:1e:14:70
Frame Delay:
MEP ID                : 100

-----
DMM Transmission      : In Progress
  Transmission mode   : Continuous
  Frames Transmitted  : 24
  Frames Received     : 15
  DMM Tx Interval    : 2 secs
  DMR Rx Timeout     : 10 msec
  Alarm Threshold     : 10 %
  Clear Threshold     : 95 %
  Measurement Window Size : 60
  Class of Service    : 0
  Tx Start Time      : Fri Apr 17 01:29:45 2009
  Min Delay          : Fri Apr 17 01:30:29 2009
  Max Delay          : Fri Apr 17 01:30:03 2009
  Last Alarm Time    : Fri Apr 17 01:29:59 2009
  Alarm State        : Set
```



```

Lost Frames in Current Window : 9

MEP ID : 200
DMM Transmission : In Progress
Transmission mode : Continuous
Frames Transmitted : 24
Frames Received : 15
DMM Tx Interval : 2 secs
DMR Rx Timeout : 10 msec
Alarm Threshold : 10 %
Clear Threshold : 95 %
Measurement Window Size : 60
Class of Service : 0
Tx Start Time : Fri Apr 17 01:29:45 2009
Min Delay : Fri Apr 17 01:30:29 2009
Max Delay : Fri Apr 17 01:30:03 2009
Last Alarm Time : Fri Apr 17 01:29:59 2009
Alarm State : Set
Lost Frames in Current Window : 9
Frame Loss:
LMM Tx Interval : 2 secs
LMR Rx Timeout : 10 msec
SES Threshold : 30 %
Consecutive Available Count : 10
Measurement Window Size : 60
Class of Service : 0
Total Configured MEPS : 2
Total Active MEPS : 2

MEP ID : 100
LMM Transmission : In Progress
Transmission mode : Continuous
Frames Transmitted : 24
Frames Received : 15
Availability Status : Available/Unavailable
Unavailability Start Time : Fri Apr 17 01:10:45 2011
Unavailability End Time : Fri Apr 17 01:20:45 2011
    Start Time : Fri Apr 17 01:10:45 2011
        Near-End Frame Loss : Fri Apr 17 01:29:45 2009 Tx
        Near-End Frame Loss : Fri Apr 17 01:39:45 2009 Min
        Far-End Frame Loss : Fri Apr 17 01:49:45 2009 Max
        Far-End Frame Loss : Fri Apr 17 01:59:45 2009 Max

MEP ID : 200
LMM Transmission : In Progress
Transmission mode : Continuous
Frames Transmitted : 24
Frames Received : 15
Availability Status : Available/Unavailable
Unavailability Start Time : Fri Apr 17 01:10:45 2011
Unavailability End Time : Fri Apr 17 01:20:45 2011
Tx Start Time : Fri Apr 17 01:10:45 2011
Min Near-End Frame Loss : Fri Apr 17 01:29:45 2009
Max Near-End Frame Loss : Fri Apr 17 01:39:45 2009
Min Far-End Frame Loss : Fri Apr 17 01:49:45 2009
Max Far-End Frame Loss : Fri Apr 17 01:59:45 2009

```



```
Total Configured Segments      : 1
Total Active Segments          : 1
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

traceroute mac port

```
traceroute mac mac {up-end-point} port port {domain} domain_name {association}
association_name {t1 t1}
```

Description

Allows you to send out a Link Trace Message (LTM) for the specified MA from the MEP configured on the port for the specified MAC address to the end of the MA.

Syntax Description

<i>mac</i>	Enter the unique system MAC address on the port configured as a MEP for the specified MA. Enter this value in the format XX:XX:XX:XX:XX:XX.
up-end-point	Use this keyword to force the LTM to be send from an UP MEP if both a DOWN MEP and an UP MEP are configured on the same port.
<i>port</i>	Enter the port number of the MEP from which you are issuing the LTM.
domain	Enter this keyword.
<i>domain_name</i>	Enter the name of the domain from which you are issuing the ping.
association	Enter this keyword.
<i>association_name</i>	Enter the name of the association from which you are issuing the ping.
t1	Enter this keyword.
<i>t1</i>	Enter the upper limit of MIPs the LTM can pass prior to reaching its destination.

Default

TTL default value is 64.

Usage Guidelines

Use this command to send an LTM from the MEP on the port for the given MAC address. If no MEP is configured on the port, the system returns an error message.



If both an UP and DOWN MEP are configured on the same port, the system uses the DOWN MEP. If you want to use the UP MEP in this situation, enter the up-end-point keyword. After you issue the command, the system prints out the route the LTM message took.

Each MIP along the route passes the LTM along only in the direction of the path and sends a packet back to the originating MAC notifying that it passed the LTM. If the destination MAC type is configured as unicast on the association to which this MEP belongs to, link trace replies will not be received from any of the MIPs configured on the intermediate switches. If there is a MIP on the switch that originated the trace route, the MIP sends a link trace reply.

Example

The following commands send an LTM:

1. A trace route invoked from a customer device CE1 to another customer device CE3 connected through an MPLS cloud (MTU1 -' PE1 'PE3), where a VPLS MIP is configured to encode a system-name, will have a response as follows:

```
(debug) Switch # traceroute mac 00:04:96:28:02:15 port 1 "extr_cfm5" "extr_ma"
Send out Link Trace Message(LTM), collecting responses [press Ctrl-C to
abort].
TTL  CFM Source MAC      Reply  Reply Mac              Port ID
=====
=
63   00:04:96:1e:6d:40  I F-f- 00:04:96:1e:6d:40  o--  1:8
62   00:04:96:1e:6d:40  E F-f- 00:04:96:1e:6d:40  o--  vp100:MTU-1
61   00:04:96:1e:16:10  I F-f- 00:04:96:1e:16:10  o--  vp100:PE-1
60   00:04:96:1e:16:10  E F-f- 00:04:96:1e:16:10  o--  vp100:PE-1
59   00:04:96:1e:14:90  I F-f- 00:04:96:1e:14:90  o--  vp100:PE-3
58   00:04:96:1e:14:90  E F-f- 00:04:96:1e:14:90  o--  1:8
57   00:04:96:28:02:15  I -h-- 00:04:96:28:02:15  o--  1
=====
=
Reply Flags: (I) Ingress, (E) Egress, (F) FwdYes, (h) RlyHit, (f) RlyFDB
Flags: (o) Ok, (d) Down, (b) Blocked
```

2. A trace route Invoked within an MPLS Cloud from MTU1 to PE3 (MTU1 -' PE1 'PE3), where a VPLS MIP is configured to encode a private-ip, will have a response as follows:

```
(debug) Switch # traceroute mac 00:04:96:1e:14:90 port 1:8 extr_cfm2 "extr_ma"
Send out Link Trace Message(LTM), collecting responses [press Ctrl-C to
abort].
TTL  CFM Source MAC      Reply  Reply Mac              Port ID
=====
=
63   00:04:96:1e:6d:40  E F-f- 00:04:96:1e:6d:40  o--  vp100:3.3.3.3
62   00:04:96:1e:16:10  I F-f- 00:04:96:1e:16:10  o--  vp100:1.1.1.1
61   00:04:96:1e:16:10  E F-f- 00:04:96:1e:16:10  o--  vp100:5.5.5.5
60   00:04:96:1e:14:90  I F-f- 00:04:96:1e:14:90  o--  vp100:3.3.3.3
59   00:04:96:1e:14:90  E -h-- 00:04:96:1e:14:90  o--  1:8
=====
=
```



Reply Flags: (I) Ingress, (E) Egress, (F) FwdYes, (h) RlyHit, (f) RlyFDB
 Flags: (o) Ok, (d) Down, (b) Blocked

If in PE1 alone, a VPLS MIP is configured to encode a system name, the response will be as follows:

```
(debug) Switch # traceroute mac 00:04:96:1e:14:90 port 1:8 extr_cfm2 "extr_ma"
Send out Link Trace Message(LTM), collecting responses [press Ctrl-C to
abort].
TTL  CFM Source MAC      Reply  Reply Mac              Port ID
=====
=
63   00:04:96:1e:6d:40  E F-f- 00:04:96:1e:6d:40  o--  vp100:3.3.3.3
62   00:04:96:1e:16:10  I F-f- 00:04:96:1e:16:10  o--  vp100:PE1
61   00:04:96:1e:16:10  E F-f- 00:04:96:1e:16:10  o--  vp100:PE1
60   00:04:96:1e:14:90  I F-f- 00:04:96:1e:14:90  o--  vp100:3.3.3.3
59   00:04:96:1e:14:90  E -h-- 00:04:96:1e:14:90  o--  1:8
=====
=
Reply Flags: (I) Ingress, (E) Egress, (F) FwdYes, (h) RlyHit, (f) RlyFDB
Flags: (o) Ok, (d) Down, (b) Blocked
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

unconfigure bfd vlan

```
unconfigure bfd vlan vlan_name
```

Description

Unconfigures BFD settings from a specified VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN name.
------------------	--------------------------

Default

N/A.

Usage Guidelines

Use this command to unconfigure BFD settings from a specified VLAN.



Example

The following command unconfigures the BFD settings on the VLAN named vlan1:

```
unconfigure bfd vlan vlan1
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

unconfigure cfm domain association end-point transmit-interval

```
unconfigure cfm domain domain_name association association_name {ports port_list  
end-point [up | down]} transmit-interval
```

Description

Unconfigures the CCM interval of the association or MEP to the default interval.

Syntax Description

<i>domain_name</i>	Specifies the domain associated with the MA.
<i>association_name</i>	Specifies the name of the MA.
<i>port_list</i>	Specifies the ports to unconfigure.
up	Enter this variable if you are changing the time interval for sending a CCM on an UP MEP.
down	Enter this variable if you are changing the time interval for sending a CCM on a DOWN MEP.

Default

1000 ms.

Usage Guidelines

Use this command to revert the CCM interval of either the association or the MEP back to the default CCM interval.



Example

The following command changes the interval the UP MEP (previously configured on port 2:4) uses to send CCM messages on the 350 association in the finance domain to the default of 1000 ms:

```
unconfigure cfm domain finance association 350 ports 2:4 end-point up  
transmit-interval
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.



12 PoE Commands

Extreme Networks PoE Devices
Summary of PoE Software Features
clear inline-power stats ports
configure inline-power budget
configure inline-power disconnect-precedence
configure inline-power label ports
configure inline-power operator-limit ports
configure inline-power priority ports
configure inline-power usage-threshold
disable inline-power
disable inline-power legacy
disable inline-power legacy slot
disable inline-power ports
disable inline-power slot
enable inline-power
enable inline-power legacy
enable inline-power legacy slot
enable inline-power ports
enable inline-power slot
reset inline-power ports
show inline-power
show inline-power configuration ports
show inline-power info ports
show inline-power slot
show inline-power stats
show inline-power stats ports
show inline-power stats slot
unconfigure inline-power budget slot
unconfigure inline-power disconnect-precedence
unconfigure inline-power operator-limit ports
unconfigure inline-power priority ports
unconfigure inline-power usage-threshold

Power over Ethernet (PoE) is an effective method of supplying 48 VDC power to certain types of powered devices (PDs) through Category 5 or Category 3 twisted pair Ethernet cables. PDs include wireless access points, IP telephones, laptop computers, web cameras, and other devices. With PoE, a single Ethernet cable supplies power and the data connection, reducing costs associated with separate power cabling and supply. PoE for ExtremeXOS includes a method of detection to assure that power is

delivered to devices that meet the IEEE 802.3af specification for PoE, as well as to many legacy devices.

Extreme Networks PoE Devices

Following is a list of the Extreme Networks devices that support PoE and the minimum required software:

- 8500-G48T-e module (with daughter card) for the BlackDiamond 8800 series switch—ExtremeXOS 12.3 and higher.
- G48Tc module (with daughter card) for the BlackDiamond 8800 series switch—ExtremeXOS 12.1 and higher.
- G48Te2 module (with daughter card) for the BlackDiamond 8800 series switch—ExtremeXOS 12.1 and higher.
- 8900-G48T-xl module (with daughter card) for the BlackDiamond 8800 series switch—ExtremeXOS 12.4 and higher.

Following is a list the Extreme Networks devices that support PoE+ and the minimum required software:

- Summit X460-24p switch—ExtremeXOS 12.5 and later.
- Summit X460-48p switch—ExtremeXOS 12.5 and later.
- Summit X440-24p switch—ExtremeXOS 15.1.1 and later.
- Summit X440-L2-24t—ExtremeXOS 15.2.1 and later.
- Summit X440-L2-48t—ExtremeXOS 15.2.1 and later.

Summary of PoE Software Features

The Extreme Networks PoE devices support the following PoE software features:

- Configuration and control of the power distribution for PoE at the system, slot, and port levels.
- Real-time discovery and classification of 802.3af-compliant PDs and many legacy (non-standard) devices.
- Monitor and control of PoE fault conditions.
- Support for configuring and monitoring PoE status at the system, slot, and port levels.
- LED control for indicating the port's PoE inline power state.
- Management of an over-subscribed power budget.
- Support for hitless failover in a chassis with two MSMs.
- Support for failover in a SummitStack.

For more information about configuring and managing PoE, see the *ExtremeXOS Concepts Guide*.



clear inline-power stats ports

```
clear inline-power stats ports [all | port_list]
```

Description

Clears the inline statistics for the selected port to zero.

Syntax Description

all	Specifies all ports.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

Use this command to clear all the information displayed by the `show inline-power stats ports port_list` command.

Example

The following command clears the inline statistics for ports 1-8 on slot 3 on a modular switch:

```
clear inline-power stats ports 3:1-3:8
```

The following command displays cleared inline power configuration information for ports 1-8 in slot 3:

```
show inline-power stats ports 3:1-3:8
```

Following is sample output from this command:

```
STATISTICS COUNTERS
Port  State      Class      Absent  InvSig  Denied  OverCurrent  Short
3:1   delivering  class3     0       0       0       0           0
3:2   delivering  class3     0       0       0       0           0
3:3   searching   class0     0       0       0       0           0
3:4   searching   class0     0       0       0       0           0
3:5   searching   class0     0       0       0       0           0
3:6   searching   class0     0       0       0       0           0
3:7   searching   class0     0       0       0       0           0
3:8   searching   class0     0       0       0       0           0
```



History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the PoE devices listed in [Extreme Networks PoE Devices](#).

configure inline-power budget

```
configure inline-power budget num_watts {slot slot}
```

Description

Sets the reserved power on the switch or specified slot to the specified watts.

Syntax Description

<i>num_watts</i>	Specifies the number of watts to reserve for specified switch or slot for inline power. Enter an integer. The minimum value is 37, or 0 if the slot is disabled; the maximum is 768; and the default value is 50.
<i>slot</i>	Specifies a slot. The slot must be configured to hold a PoE module.

Default

50 W.

Usage Guidelines

This command sets the budgeted power reserved for all PDs connected to the switch or specified slot in Watts. On a modular switch, none of the power budget on a specified slot can be used to power other slots or PDs on other slots.

On a modular switch, if you specify a slot that is not configured to hold a PoE module, the system returns the following error message:

```
Error: Slot 2 is not capable of inline-power.
```

You can modify the power budget without disabling the switch or slot.

If the power consumption of the PDs on the switch or a specified slot exceeds this configured power budget, the system disconnects the lowest priority ports. (Refer to [configure inline-power priority ports](#) for information on configuring this parameter.)



If you attempt to configure this power budget for a value that the system cannot safely provide, the system returns an error message. To display inline power settings, use the command `show inline-power`; to display the power for the entire switch, use the command `show power budget`.



Note

You must disable inline power for the switch or the specified slot using the `disable inline-power slot` command prior to setting the budget to 0.

To reduce the chances of ports fluctuating between powered and non-powered states, newly inserted PDs are not powered when the actual delivered power for the module is within approximately 19 W of the configured inline power budget for that switch or slot. However, actual aggregate power can be delivered up to the configured inline power budget for the switch or slot (for example, when delivered power from ports increases or when the configured inline power budget for the switch or slot is reduced).

Each Summit family switch has its own PSU and the power budget for each Summit switch is determined by the internal/external PSUs connected to that Summit switch. For this reason, `configure inline-power budget num_watts {slotslot}` is not applicable to Summit family switches or SummitStack.

Example

The following command sets the power for slot 4 to 150 W on a modular switch:

```
configure inline-power budget 150 slot 4
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the BlackDiamond 8000 series modules listed in Extreme Networks PoE Devices.

configure inline-power disconnect-precedence

```
configure inline-power disconnect-precedence [deny-port | lowest-priority]
```

Description

Configures the disconnect precedence priority for the switch when a new PD is detected and the measured inline power for that switch or specified slot is within 19 W of the switch's or slot's PoE power budget.



Syntax Description

deny-port	Specifies power be denied to PD requesting power, regardless of priority.
lowest-priority	Specifies power be withdrawn from lowest-priority port(s) when next PD requesting power connects.

Default

Deny-port.

Usage Guidelines

You configure this parameter for the switch and for the entire modular switch; you cannot configure this per slot or per port.

If the power supplied to the PDs on a switch or specified slot exceeds the power that was budgeted for that switch or specified slot, the system disconnects power to one or more ports to prevent power overload. Refer to `configure inline-power budget` for information on configuring and modifying the power budgeted for each switch or specified slot.

You configure the switch to either deny power to the next PD that requests power on that switch or slot, regardless of the priority, or to disconnect those PDs on ports with lower priorities until there is enough power for the new PD. If you select this last argument and you did not configure port priorities or if several ports have the same priority, the switch withdraws power (or disconnects) those ports with the highest port number (s). Refer to `configure inline-power priority ports` for information on configuring the PoE priority for the ports.

The default value is deny-port. So, if you do not change the default value and the switch's or slot's power is exceeded, the next PD requesting power will not be connected.

When the setting is lowest priority, the switch continues dropping ports with the lowest configured PoE port priorities, or the highest port number in the case of equal PoE port priorities, until there is enough power for the requesting PD.

Example

The following command sets the switch to withdraw power from the lowest-priority port(s):

```
configure inline-power disconnect-precedence lowest-priority
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on:

- The following modules on BlackDiamond 8800 series switches



- 8500-G48T-e module (with daughter card)—ExtremeXOS 12.3 and higher.
- G48P module—ExtremeXOS 11.1 and higher.
- G48Pe module—ExtremeXOS 11.5 and higher.
- G48Tc module (with daughter card)—ExtremeXOS 12.1 and higher.
- G48Te2 module (with daughter card)—ExtremeXOS 12.1 and higher.
- 8900-G48T-xl module (with daughter card)—ExtremeXOS 12.4.2 and higher.
- Summit X460-24p and X460-48p switches—ExtremeXOS 12.5 and higher.

configure inline-power label ports

```
configure inline-power label string ports port_list
```

Description

Lets you create your own label for a specified PoE port or group of PoE ports.

Syntax Description

<i>string</i>	Specifies a name up to 15 characters in length to identify the specified power port(s).
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

No label.

Usage Guidelines

Use the `show inline-power configuration ports` command, as shown in the following example, to display inline power configuration information, including the label (if any) for each port:

```
show inline-power configuration port 3:1-10
```

Following is sample output from this command on a modular switch:

Port	Config	Operator	Limit	Priority	Label
3:1	Enabled	16000	mW	Low	finance
3:2	Enabled	15000	mW	Low	finance
3:3	Enabled	15000	mW	Low	
3:4	Enabled	15000	mW	Low	
3:5	Enabled	15000	mW	Low	
3:6	Enabled	15000	mW	Low	marketing
3:7	Enabled	15000	mW	Low	marketing
3:8	Enabled	15000	mW	Low	marketing
3:9	Enabled	15000	mW	Low	
3:10	Enabled	15000	mW	Low	



Example

The following command assigns the name “alpha-test_1” to port 1 on slot 4:

```
config inline-power label alpha-test_1 ports 4:1
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the PoE devices listed in [Extreme Networks PoE Devices](#).

configure inline-power operator-limit ports

```
configure inline-power operator-limit milliwatts ports [all |port_list]
```

Description

Sets the power limit allowed for PDs connected to the specified ports.

Syntax Description

<i>milliwatts</i>	An integer specifying the maximum allowed power in milliwatts.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

PoE—15400 mW.

PoE+—30000 mW.

Usage Guidelines

This command sets the power limit that a PD can draw on the specified ports. For PoE, the range is 3000 to 16800mW and the default value is 15400 mW. For PoE+, the range is 3000 to 30000 mW and the default value is 30000 mW.

If the measured power for a specified port exceeds the port’s operator limit, the power is withdrawn from that port and the port moves into a fault state.

If you try to set an operator-limit outside the accepted range, the system returns the following error message:

```
Error: Invalid operator-limit value. Must be in the range of 3000-16800 mW
```



Example

The following command sets the limit for legacy PDs on ports 3 – 6 of slot 5 on a modular switch to 10000 mW:

```
configure inline-power operator-limit 10000 ports 5:3-5:6
```

History

This command was first available in ExtremeXOS 11.1.

PoE+ was added in ExtremeXOS 12.5.

Platform Availability

This command is available on the PoE devices listed in [Extreme Networks PoE Devices](#).

configure inline-power priority ports

```
configure inline-power priority [critical | high | low] ports port_list
```

Description

Sets the PoE priority on the specified ports.

Syntax Description

critical high low	Sets the PoE priority for the specified ports.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Low.

Usage Guidelines

The system allocates power to those ports with the highest priorities first. This command can also be used in conjunction with the `configure inline-power disconnect-precedence` command. If you configure the disconnect precedence as lowest priority, then newly detected PDs will be powered if that port has higher priority than the existing powered ports.

If there are multiple ports at the same priority level (either configured or by default) and one of the ports must have power withdrawn because of excessive power demands, those ports with the lower port number are powered first. The higher port numbers have power withdrawn first in the case of equal PoE port priorities.



Example

The following command assigns a critical PoE priority on ports 4 – 6 on slot 3 on a modular switch:

```
configure inline-power priority critical ports 3:4-3:6
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on:

- The following modules on BlackDiamond 8800 series switches:
 - 8500-G48T-e module (with daughter card)—ExtremeXOS 12.3 and higher.
 - G48P module—ExtremeXOS 11.1 and higher.
 - G48Pe module—ExtremeXOS 11.5 and higher.
 - G48Tc module (with daughter card) —ExtremeXOS 12.1 and higher.
 - G48Te2 module (with daughter card)—ExtremeXOS 12.1 and higher.
 - 8900-G48T-xl module (with daughter card)—ExtremeXOS 12.4.2 and higher.
- Summit X460-24p and X460-48p switches—ExtremeXOS 12.5 and higher.

configure inline-power usage-threshold

```
configure inline-power usage-threshold threshold
```

Description

Sets the inline power usage SNMP event threshold.

Syntax Description

<i>threshold</i>	Specifies the percentage of budgeted power used on any PoE module or stand-alone switch that causes the system to send an SNMP event and create a log message. The range 1 to 99; the default value is 70.
------------------	--

Default

70.

Usage Guidelines

This command sets the threshold for generating an SNMP event and an Event Management System (EMS) message. On a modular switch, this threshold is when the measured power for a PoE module compared to the budgeted power for that slot exceeds a certain value. On stand-alone switches, this



threshold applies to the total power available to the entire switch. The configured threshold value initiates the event and message once that percentage of the budgeted power is being used.

On a modular switch, the PoE threshold applies only to the percentage per slot of measured to budgeted power use; it does not apply systemwide.

The system generates an additional SNMP event and EMS message once the power usage falls below the threshold again; once the condition clears.

Example

The following command sets the inline power usage alarm threshold at 75%:

```
configure inline-power usage-threshold 75
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the PoE devices listed in [Extreme Networks PoE Devices](#).

disable inline-power

disable inline-power

Description

Shuts down PoE power currently provided on all ports on all slots.

Syntax Description

This command has no arguments or variables.

Default

Enable.

Usage Guidelines

You can control whether inline power is provided to the system by using the `disable inline-power` command and the `enable inline-power` command. Using the `disable inline-power` command shuts down inline power currently provided on the entire switch or to specified ports



and slots. Disabling inline power to a switch, port, or slot immediately removes power to any connected PDs. By default, inline power provided to all ports is enabled.

Note

Disabling inline power using the `disable inline-power` command does not affect the data traffic traversing the port. And, disabling the port using the `disable port` command does not affect the inline power supplied to the port.

On modular switches, disabling inline power does not allow PoE power reserved for slots to be allocated to other slots that may be needing more power to become operational. However, when you issue the command `disable slot` on a slot holding a PoE module, the inline power is also disabled and that slot is totally offline.

Note

Inline power cannot be delivered to connected PDs unless the Summit family switch or BlackDiamond 8800 chassis and module are powered on.

Example

The following command shuts down inline power currently provided to all ports and all slots:

```
disable inline-power
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the PoE devices listed in [Extreme Networks PoE Devices](#).

disable inline-power legacy

disable inline-power legacy

Description

Disables the non-standard (or capacitance) power detection mechanism for the switch.

Syntax Description

This command has no arguments or variables.



Default

Disable.

Usage Guidelines

This command disables the non-standard power-detection mechanism on the switch. Legacy PDs do not conform to the IEEE 802.3af standard but may be detected by the switch through a capacitance measurement.

However, measuring the power through capacitance is used only if this parameter is enabled and after an unsuccessful attempt to discover the PD using the standard resistance measurement method. The default for legacy is disabled.

The reason legacy detection is configurable is that it is possible for a normal (non-PoE) device to have a capacitance signature that causes the device to be detected as a legacy PoE device and have power delivered to it, potentially causing damage to the device.

Example

The following command disables capacitance detection of PDs on the switch:

```
disable inline-power legacy
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on the Summit family switches listed in [Extreme Networks PoE Devices](#).

disable inline-power legacy slot

```
disable inline-power legacy slot slot
```

Description

Disables the non-standard (or capacitance) power detection mechanism for the specified slot.

Syntax Description

<i>slot</i>	Disables non-standard power detection for specified slot on a modular switch.
-------------	---

Default

Disable.



Usage Guidelines

This command disables the non-standard power-detection mechanism on the switch or specified slot. Legacy PDs do not conform to the IEEE 802.3af standard but may be detected by the switch through a capacitance measurement.

However, measuring the power through capacitance is used only if this parameter is enabled and after an unsuccessful attempt to discover the PD using the standard resistance measurement method. The default for legacy is disabled.

The reason legacy detection is configurable is that it is possible for a normal (non-PoE) device to have a capacitance signature that causes the device to be detected as a legacy PoE device and have power delivered to it, potentially causing damage to the device.

On a stack if you do not specify a slot number, the command operates on all active nodes. This command operates only on nodes in the active topology.

Example

The following command disables capacitance detection of PDs on slot 3 of a modular switch:

```
disable inline-power legacy slot 3
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the BlackDiamond 8000 series modules listed in [Extreme Networks PoE Devices](#). It is available on SummitStack when the stack contains Summit family switches listed in [Extreme Networks PoE Devices](#).

disable inline-power ports

```
disable inline-power ports [all | port_list]
```

Description

Shuts down PoE power currently provided to all ports or to specified ports.

Syntax Description

all	Disables inline power to all ports on the switch.
<i>port_list</i>	Disables inline power to the specified ports.



Default

Enable.

Usage Guidelines

Disabling inline power to ports immediately removes power to any connected PDs. By default, the capability to provide inline power to all ports is enabled.



Note

Disabling inline power using the `disable inline-power` command does not affect the data traffic traversing the port. And, disabling the port using the `disable port` command does not affect the inline power supplied to the port.

Disabling inline power to a port providing power to a PD immediately removes power to the PD.



Note

On a modular switch, PoE power removed from ports using this command can be used by other ports on the same module.

Example

The following command shuts down inline power currently provided to ports 4 and 5 on slot 3 on a modular switch:

```
disable inline-power ports 3:4-5
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the PoE devices listed in [Extreme Networks PoE Devices](#).

disable inline-power slot

```
disable inline-power slot slot
```

Description

Shuts down PoE power currently provided to the specified slot.

Syntax Description

<code>slot</code>	Disables inline power to specified slot.
-------------------	--



Default

Enable.

Usage Guidelines

Disabling inline power to a slot immediately removes power to any connected PDs. By default, the capability to provide inline power to a slot is enabled.

Disabling a slot using this command does not change the power budgeted to a specified slot using the `configure inline-power budget` command; nor can that power be used by PDs connected to any other slot.



Note

You can set the reserved power budget to 0 for a slot if, and only if, you first issue this command.

On a stack if you do not specify a slot number, the command operates on all active nodes. This command operates only on nodes in the active topology.

Example

The following command removes power to all PDs on slot 3:

```
disable inline-power slot 3
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the PoE devices listed in [Extreme Networks PoE Devices](#).

enable inline-power

enable inline-power

Description

Enables PoE power to all ports; on a modular switch, this is all ports on all slots.

Syntax Description

This command has no arguments or variables.



Default

Enable.

Usage Guidelines

You can control whether inline power is provided to the system by using the `disable inline-power` command and the `enable inline-power` command. By default, inline power provided to all ports is enabled.

Enabling inline power starts the PoE detection process used to discover, classify, and power remote PDs.

Note



When you are working on a modular switch, if your chassis has an inline power module and there is not enough power to supply a slot, that slot will not be powered on; the slot will not function in data-only mode without enough power for inline power.

Disabling inline power using the `disable inline-power` command does not affect the data traffic traversing the port. And, disabling the port using the `disable port` command does not affect the inline power supplied to the port.

However, when you issue the command `disable slot` for a modular switch on a slot holding a PoE module, the inline power is also disabled; that slot is totally offline.

Note



Inline power cannot be delivered to connected PDs unless the Summit family switch or BlackDiamond 8800 chassis and module are powered on.

Example

The following command enables inline power currently provided to all ports and all slots:

```
enable inline-power
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the PoE devices listed in [Extreme Networks PoE Devices](#).

enable inline-power legacy

```
enable inline-power legacy
```



Description

Enables the non-standard (or capacitance) power detection mechanism for the switch.

Syntax Description

This command has no arguments or variables

Default

Disable.

Usage Guidelines

This command disables the non-standard power-detection mechanism on the switch. Legacy PDs do not conform to the IEEE 802.3af standard but may be detected by the switch through a capacitance measurement.

However, measuring the power through capacitance is used only if this parameter is enabled and after an unsuccessful attempt to discover the PD using the standard resistance measurement method. The default for legacy is disabled.



Caution

A normal (non-PoE) device may have a capacitance signature that causes the device to be detected as a legacy PoE device (and have power supplied), potentially causing damage to the device.

Example

The following command enables capacitance detection of PDs on the switch:

```
enable inline-power legacy
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on the Summit family switches listed in [Extreme Networks PoE Devices](#).

enable inline-power legacy slot

```
enable inline-power legacy slot slot
```



Description

Enables non-standard (or capacitance) power detection mechanism for the specified slot on a modular switch.

Syntax Description

<code>slot</code>	Enables non-standard power detection for specified slot on a modular switch.
-------------------	--

Default

Disable.

Usage Guidelines

This command enables the non-standard power-detection mechanism on the specified slot. Legacy PDs do not conform to the IEEE 802.3af standard but may be detected by the switch through a capacitance measurement.

However, measuring the power through capacitance is used only if this parameter is enabled and after an unsuccessful attempt to discover the PD using the standard resistance measurement method. The default for legacy is disabled.



Caution

A normal (non-PoE) device may have a capacitance signature that causes the device to be detected as a legacy PoE device (and have power supplied), potentially causing damage to the device.

On stack, if you do not specify a slot number, the command operates on all active nodes. The command operates only on nodes in the active topology.

Example

The following command enables capacitance detection of PDs on slot 3 on a modular switch:

```
enable inline-power legacy slot 3
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the BlackDiamond 8000 series modules listed in [Extreme Networks PoE Devices](#). It is available on SummitStack when the stack contains Summit family switches listed in [Extreme Networks PoE Devices](#).



enable inline-power ports

```
enable inline-power ports [all | port_list]
```

Description

Enables PoE power currently provided to all ports or to specified ports.

Syntax Description

all	Enables inline power to all ports on the switch.
<i>port_list</i>	Enables inline power to the specified ports.

Default

Enable.

Usage Guidelines

Disabling inline power to a port immediately removes power to any connected PD. By default, inline power provided to all ports is enabled.

On modular switches, to deliver inline power to ports with connected PDs, you must also reserve power for the slot with the PDs using the `configure inline-power budget` command. If you do not have enough reserved power for the port, that port moves into a Denied state.



Note

On modular switches, if your chassis has an inline power module and there is not enough power to supply a slot, that slot will not be powered on; the slot will not function in data-only mode without enough power for inline power.

Disabling inline power using the `disable inline-power` command does not affect the data traffic traversing the port. And, disabling the port using the `disable port` command does not affect the inline power supplied to the port.

Example

The following command enables inline power to ports 4 and 5 on slot 3 on a modular switch:

```
enable inline-power ports 3:4-5
```

History

This command was first available in ExtremeXOS 11.1.



Platform Availability

This command is available on the PoE devices listed in [Extreme Networks PoE Devices](#).

enable inline-power slot

```
enable inline-power slot slot
```

Description

Enables PoE power to the specified slot on modular switches.

Syntax Description

<i>slot</i>	Enables inline power to specified slot.
-------------	---

Default

Enable.

Usage Guidelines

Disabling inline power to a slot immediately removes power to any connected PDs. By default, inline power provided to all slots is enabled.

To deliver inline power to slots, you must reserve power for that slot using the `configure inline-power budget` command. By default, each PoE module has 50 W of power reserved for inline power.

Note



If your chassis has an inline power module and there is not enough power to supply a slot, that slot will not be powered on; the slot will not function in data-only mode without enough power for inline power.

Disabling inline power using the `disable inline-power` command does not affect the data traffic traversing the slot. And, disabling the slot using the `disable slot` command does not affect the inline power supplied to the slot.

On a stack, if you do not specify a slot number, the command operates on all active nodes. This command operates only on nodes in the active topology.

Example

The following command makes inline power available to slot 3:

```
enable inline-power slot 3
```



History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the BlackDiamond 8000 series modules listed in [Extreme Networks PoE Devices](#). It is available on SummitStack when the stack contains Summit family switches listed in [Extreme Networks PoE Devices](#).

reset inline-power ports

```
reset inline-power ports port_list
```

Description

Power cycles the specified ports.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports for which power is to be reset.
------------------	--

Default

N/A.

Usage Guidelines

This command power cycles the specified ports. Ports are immediately disabled and then re-enabled, allowing remote PDs to be power-cycled.

This command affects only inline power; it does not affect network connectivity for the port(s).

Example

The following command resets power for port 4 on slot 3 on a modular switch:

```
reset inline-power ports 3:4
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the PoE devices listed in [Extreme Networks PoE Devices](#).



show inline-power

show inline-power

Description

Displays inline power status information for the specified PoE switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The output varies depending on the PoE device you are using.



Note

For additional information on inline power parameters, refer to the `show power budget` command.

Example

The following command displays inline power status for the switch:

```
show inline-power
```

Following is sample output from this command for the Summit X460-24p switch switch:

```
Inline Power System Information
Configured           : Enabled
Power Usage Threshold : 70 percent
Firmware Status     Power (Watts) Power (Watts) Legacy
Operational         405 W          0 W          Disabled
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the Summit family switches listed in Extreme Networks PoE Devices.



show inline-power configuration ports

```
show inline-power configuration ports port_list
```

Description

Displays inline power configuration information for the specified ports.

Syntax Description

<code>port_list</code>	Specifies one or more ports.
------------------------	------------------------------

Default

N/A.

Usage Guidelines

The output displays the following inline power configuration information for the specified ports:

- Config—Indicates whether the port is enabled to provide inline power:
 - Enabled: The port can provide inline power.
 - Disabled: The port cannot provide inline power.
- Operator Limit—Displays the configured limit, in milliwatts, for inline power on the port.
- Label—Displays a text string, if any, associated with the port.

The following also displays for this command on modular PoE devices:

- Priority—Displays inline power priority of the port, which is used when the disconnect precedence is set to lowest priority:
 - Low.
 - High.
 - Critical.

Example

The following command displays inline power configuration information for ports 1 to 10 in slot 3 on a modular switch:

```
show inline-power configuration port 3:1-10
```

Following is sample output from this command:

Port	Config	Operator Limit	Priority	Label
3:1	Enabled	15000 mW	Low	
3:2	Enabled	15000 mW	Low	
3:3	Enabled	15000 mW	Low	
3:4	Enabled	15000 mW	Low	



3:5	Enabled	15000 mW	Low
3:6	Enabled	15000 mW	Low
3:7	Enabled	15000 mW	Low
3:8	Enabled	15000 mW	Low
3:9	Enabled	15000 mW	Low
3:10	Enabled	15000 mW	Low

History

This command was first available in ExtremeXOS 11.1

Platform Availability

This command is available on the PoE devices listed in [Extreme Networks PoE Devices](#).

show inline-power info ports

```
show inline-power info {detail} ports port_list
```

Description

Displays inline power information for the specified ports.

Syntax Description

<i>port_list</i>	Specifies one or more ports.
------------------	------------------------------

Default

N/A.

Usage Guidelines



Note

Ports in the denied or faulted state periodically display the searching state as the hardware retests the PD state.

You can use this command to generate a summary report or a detailed report.

Summary output displays the following inline power information for the specified ports:

- State—Displays the port power state:
 - Disabled.
 - Searching.
 - Delivering.
 - Faulted.
 - Disconnected.



- Other.
- Denied.
- PD's power class—Displays the class type of the connected PD:
 - “----”: disabled or searching.
 - “class0”: class 0 device.
 - “class1”: class 1 device.
 - “class2”: class 2 device.
 - “class3”: class 3 device.
 - “class4”: class 4 device.
- Volts—Displays the measured voltage. A value from 0 to 2 is valid for ports that are in a searching state.
- Curr—Displays the measured current, in milliamperes, drawn by the PD.
- Power—Displays the measured power, in watts, supplied to the PD.
- Fault—Displays the fault value:
 - None.
 - UV/OV fault.
 - UV/OV spike.
 - Over current.
 - Overload.
 - Undefined.
 - Underload.
 - HW fault.
 - Discovery resistance fail.
 - Operator limit violation.
 - Disconnect.
 - Discovery resistance, A2D failure.
 - Classify, A2D failure.
 - Sample, A2D failure.
 - Device fault, A2D failure.
 - Force on error.

The detail command lists all inline power information for the selected ports.

Detail output displays the following information:

- Configured Admin State—Displays the port's configured state; Enabled or Disabled.
- Inline Power State—Displays the port power state.
- MIB Detect Status—Displays the port state as reported by SNMP; valid values are as follows:
 - disabled
 - searching
 - delivering
 - fault
 - test
 - otherFault
 - denyLowPriority



- Label—Displays the port's configured label.
- Operator Limit—Displays the port's configured operator limit value.
- PD Class—Displays the class type of connected PD:
- Max Allowed Power—Displays the amount of maximum allowed power for a device of this class.
- Measured Power—Displays the measured power, in watts, supplied to the PD.
- Line Voltage—Displays the measured voltage. A value from 0 to 2 is valid for ports in a searching state.
- Current—Displays the measured current, in milliamperes, drawn by the PD.
- Fault Status—Displays the fault value.
- Detailed Status

Example

The following command displays summary inline power information for ports 1 to 3 on slot 3 on a modular switch:

```
show inline-power info ports 3:1-3
```

Following is sample output from this command:

Port (mA)	State (Watts)	Class	Volts	Curr	Power	Fault
3:1	delivering	class3	48.3	192	9.300	None
3:2	delivering	class3	48.3	192	9.300	None
3:3	searching	-----	0.0	0	0.0	None

The following command displays detail inline power information for port 1 on slot 3:

```
show inline-power info detail port 3:1
```

Following is sample output from this command:

```
Port 3:1
Configured Admin State: enabled
Inline Power State      : delivering
MIB Detect Status      : delivering
Label                  :
Operator Limit         : 16800 milliwatts
PD Class               : class3
Max Allowed Power      : 15.400 W
Measured Power         : 9.400 W
Line Voltage           : 48.3 Volts
Current                : 193 mA
Fault Status           : None
Detailed Status        :
```



The following command displays detail inline power information for port 3 on a Summit X460 switch:

```
show inline-power info detail ports
```

Following is sample output from this command:

```
Port 3
Configured Admin State: enabled
Inline Power State      : delivering
MIB Detect Status      : delivering
Label                  :
Operator Limit         : 30000 milliwatts
PD Class               : class4
Max Allowed Power     : 30.0 W
Measured Power        : 28.400 W
Line Voltage          : 54.3 Volts
Current               : 523 mA
Fault Status          : None
Detailed Status       : Delivering power to IEEE PD
Priority              : low
The Detailed Status field, when available, will be one of the following
strings for PoE+.
Delivering power to IEEE PD
Delivering power to Pre-standard PD
Fault - Maintain Power Signature (MPS) absent
Fault - Short
Fault - Overload
Fault - Thermal shutdown
Fault - Startup failure
Fault - Classification failure
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the PoE devices listed in [Extreme Networks PoE Devices](#).

show inline-power slot

```
show inline-power slot slot
```

Description

Displays inline power information for the specified slot on modular switches.



Syntax Description

<i>slot</i>	Specifies the slot.
-------------	---------------------

Default

N/A.

Usage Guidelines

On modular switches, the output indicates the following inline power status for each system:

- Configured power:
 - Enabled.
 - Disabled.
- System power surplus.
- Redundant power surplus.
- Power usage threshold.
- Disconnect precedence.
- Legacy—The status of the legacy mode, which allows detection of many non-standard PDs.

On modular switches, the output indicates the following inline power status information for each slot:

- Inline power status—The status of inline power. The status conditions are:
 - Enabled.
 - Disabled.
- Firmware status—The operational status of the slot. The status conditions are:
 - Operational.
 - Not operational.
 - Disabled.
 - Subsystem failure.
 - Card not present.
 - Slot disabled.
- Budgeted power—The amount of power, in watts, that is available to the slot.
- Measured power—The amount of power, in watts, that currently being used by the slot.

On a stack, if you do not specify a slot number, the command operates on all active nodes. This command operates only on nodes in the active topology.

Example

The following command displays inline power information for slot 3 on a modular switch:

```
show inline-power slot 3
```



Following is sample output from this command on a modular switch:

```

Inline Power System Information
Configured           : Enabled
System Power Surplus : 1500 Watts available for budgeting
Redundant Power Surplus : 465 Watts available for budgeting to maintain N
+1
Power Usage Threshold : 70 percent (per slot)
Disconnect Precedence : lowest-priority
Legacy Mode          : Disabled
Budgeted           Measured
Slot  Inline-Power  Firmware Status    Power (Watts)  Power (Watts)
3     Enabled      Operational        50 W           9 W
4     Enabled      Card Not Present   ( 50 W)        n/a
7     Enabled      Operational        50 W           0 W
Note: A budget value in parentheses is not allocated from the system power

```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the BlackDiamond 8000 series modules listed in [Extreme Networks PoE Devices](#). It is available on SummitStack when the stack contains Summit family switches listed in [Extreme Networks PoE Devices](#).

show inline-power stats

```
show inline-power stats
```

Description

Displays inline power statistics for the specified switch.

Syntax Description

There are no variables or parameters for this command.

Default

N/A.

Usage Guidelines

Use this command to produce a report that shows the firmware status and version plus how many ports are currently faulted, powered, and waiting for power for the switch. Unlike the values displayed



with the `show inline-power stats ports` command, these values are current readings, not cumulative counters.

Example

The following command displays inline power statistics information for the Summit X460-48p switch:

```
show inline-power stats
```

Following is sample output from this command:

```
Inline-Power Slot Statistics
Firmware status           : Operational
Firmware revision        : 292b1
Total ports powered      : 7
Total ports awaiting power : 17
Total ports faulted      : 0
Total ports disabled     : 0
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on the Summit family switches listed in Extreme Networks PoE Devices.

show inline-power stats ports

```
show inline-power stats ports port_list
```

Description

Displays inline power statistics for the specified ports.

Syntax Description

<code>port_list</code>	Specifies one or more slots and ports.
------------------------	--

Default

N/A.

Usage Guidelines

The output displays the following inline power statistics for the specified ports:

- State—Displays the port power state:



- Disabled.
- Searching.
- Delivering.
- Faulted.
- Disconnected.
- Other.
- Denied.
- PD's power class—Displays the class type of the connected PD:
 - “----”: disabled or searching.
 - “class0”: class 0 device.
 - “class1”: class 1 device.
 - “class2”: class 2 device.
 - “class3”: class 3 device.
 - “class4”: class 4 device.
- Absent—Displays the number of times the port was disconnected.
- InvSig—Displays the number of times the port had an invalid signature.
- Denied—Displays the number of times the port was denied.
- Over-current—Displays the number of times the port entered an overcurrent state.
- Short—Displays the number of times the port entered undercurrent state.

Example

The following command displays inline power configuration information for ports 1 to 10 in slot 3 on a modular switch:

```
show inline-power stats ports 3:1-10
```

Following is sample output from this command:

```

STATISTICS COUNTERS
Port  State      Class      Absent  InvSig  Denied  OverCurrent  Short
3:1   delivering  class3     0       0       0       18           0
3:2   delivering  class3     0       0       0       0            0
3:3   searching   class0     0       0       0       0            0
3:4   searching   class0     0       0       0       0            0
3:5   searching   class0     0       0       0       0            0
3:6   searching   class0     0       0       0       0            0
3:7   searching   class0     0       0       0       0            0
3:8   searching   class0     0       0       0       0            0
3:9   searching   class0     0       0       0       0            0
3:10  searching   class0     0       0       0       0            0

```

History

This command was first available in ExtremeXOS 11.1.



Platform Availability

This command is available on the PoE devices listed in [Extreme Networks PoE Devices](#).

show inline-power stats slot

```
show inline-power stats slot slot
```

Description

Displays inline power statistics for the specified slot on modular switches.

Syntax Description

<code>slot</code>	Specifies the slot.
-------------------	---------------------

Default

N/A.

Usage Guidelines

Use this command to produce a report that shows the firmware status and version plus how many ports are currently faulted, powered, and waiting for power for the selected slots. Unlike the values displayed with the `show inline-power stats ports` command, these values (displayed with the `show inline-power stats slot` command) are current readings, not cumulative counters.

On a stack, if you do not specify a slot number, the command operates on all active nodes. This command operates only on nodes in the active topology.

Example

The following command displays inline power statistics information for slot 3 on a modular switch:

```
show inline-power stats slot 3
```

Following is sample output from this command:

```
Inline-Power Slot Statistics
Slot: 3
Firmware status           : Operational
Firmware revision         : 292b1
Total ports powered       : 7
Total ports awaiting power : 41
Total ports faulted       : 0
Total ports disabled      : 0
```



History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the BlackDiamond 8000 series modules listed in [Extreme Networks PoE Devices](#). It is available on SummitStack when the stack contains Summit family switches listed in [Extreme Networks PoE Devices](#).

unconfigure inline-power budget slot

```
unconfigure inline-power budget slot slot
```

Description

Unconfigures the inline reserved power on modular switches on the specified slot and returns the power budget on that slot to the default value of 50 W.

Syntax Description

<code>slot</code>	Specifies the slot.
-------------------	---------------------

Default

50 W.

Usage Guidelines

This command unconfigures any previously configured power budget for the specified slot and resets the budgeted power reserved for all PDs connected to this slot to 50 W. The rest of the previously configured power budget on this slot cannot be used to power other slots or PDs on other slots (unless you explicitly reconfigure the power budget for other slots).

If you specify a slot that does not have a PoE module, the system returns the following error message:

```
Error: Slot 2 is not capable of inline-power.
```

Example

The following command resets the power for slot 4 to 50 W:

```
unconfigure inline-power budget slot 4
```

History

This command was first available in ExtremeXOS 11.1.



Platform Availability

This command is available on the BlackDiamond 8000 series modules listed in [Extreme Networks PoE Devices](#).

unconfigure inline-power disconnect-precedence

unconfigure inline-power disconnect-precedence

Description

On a modular switch, unconfigures the disconnect precedence setting and returns the switch to the default disconnect precedence value of deny port.

Syntax Description

This command has no arguments or variables.

Default

Deny-port.

Usage Guidelines

You configure this parameter for the entire switch; you cannot configure this per slot or per port.

Unconfigures the PoE disconnect precedence previously set for the modular switch and returns the disconnect precedence to the default value of deny port. Deny port denies power to the next PD that requests inline power from the slot when the inline power budget for the switch or slot is reached, regardless of the inline power port priority.

Example

The following command resets the switch to the PoE disconnect precedence value, which is deny port:

```
unconfigure inline-power disconnect-precedence
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the PoE devices listed in [Extreme Networks PoE Devices](#).



unconfigure inline-power operator-limit ports

```
unconfigure inline-power operator-limit ports [all | port_list]
```

Description

Unconfigures the PoE operator limit setting and resets the power limit allowed for PDs connected to the specified ports to the default values.

Syntax Description

all	Specifies all ports.
<i>port_list</i>	Specifies one or more slots and ports.

Default

PoE—15400 mW.

PoE+—30000 mW.

Usage Guidelines

This command unconfigures any previously configured operator limit for the specified ports. It resets the maximum power that any PD can draw to 15400 mW for PoE and 30000 mW for PoE+.

Example

The following command resets the limit on ports 3 to 6 of slot 5 on a modular switch to the default value of 15400 mW:

```
unconfigure inline-power operator-limit ports 5:3-5:6
```

History

This command was first available in ExtremeXOS 11.1.

PoE+ was added in ExtremeXOS 12.5.

Platform Availability

This command is available on the PoE devices listed in [Extreme Networks PoE Devices](#).

unconfigure inline-power priority ports

```
unconfigure inline-power priority ports [all | port_list]
```



Description

On modular switches, unconfigures the PoE priority on the specified ports, and returns the ports to the default PoE port priority value of low.

Syntax Description

all	Specifies all ports.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Low.

Usage Guidelines

Use this to reset the PoE port priority on specified ports on modular switches to the default value of low.

If there are multiple ports on the modular switch at the same priority level (either configured, or by default) and one of the ports must have power withdrawn because of excessive power demands, those ports with the lower port number are powered first. The higher port numbers have power withdrawn first in the case of equal PoE port priorities.

Example

The following command resets the PoE priority on ports 4 – 6 on slot 3 to low:

```
unconfigure inline-power priority ports 3:4-3:6
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the PoE devices listed in [Extreme Networks PoE Devices](#).

unconfigure inline-power usage-threshold

```
unconfigure inline-power usage-threshold
```

Description

Unconfigures the inline power usage alarm threshold and returns threshold to the default value of 70%.



Syntax Description

This command has no arguments or variables.

Default

70.

Usage Guidelines

This command unconfigures the PoE usage threshold setting for initiating SNMP event and EMS messages and returns the switch's inline power usage threshold for to 70%. The system initiates an event and message once that percentage of the budgeted power is being used.

On stand-alone switches, this PoE threshold applies to the entire switch. On modular switches, the threshold applies only to the percentage per slot of measured to budgeted power use; the threshold does not apply to the entire switch.

The system generates an additional SNMP event and EMS message once the power usage falls below the threshold again; once the condition clears.

Example

The following command resets the inline power usage alarm threshold to 70%:

```
unconfigure inline-power usage-threshold
```

History

This command was first available in ExtremeXOS 11.1

Platform Availability

This command is available on the PoE devices listed in [Extreme Networks PoE Devices](#).



13 Commands for Status Monitoring and Statistics

```
Event Management System
Extreme Link Status Monitoring
sFlow Statistics
RMON
clear counters
clear counters xml-notification
clear cpu-monitoring
clear elsm ports auto-restart
clear elsm ports counters
clear log
clear log counters
clear sys-recovery-level
configure elsm ports hellotime
configure elsm ports hold-threshold
configure elsm ports uptimer-threshold
configure log display
configure log filter events
configure log filter events match
configure log target filter
configure log target format
configure log target match
configure log target severity
configure log target syslog
configure log target xml-notification filter
configure ports monitor vlan
configure sflow agent ipaddress
configure sflow collector ipaddress
configure sflow max-cpu-sample-limit
configure sflow poll-interval
configure sflow ports sample-rate
configure sflow sample-rate
configure sys-health-check all level
configure sys-health-check interval
configure sys-recovery-level
configure sys-recovery-level slot
```

```
configure sys-recovery-level switch
configure syslog add
configure syslog delete
configure xml-notification target
configure xml-notification target add/delete
create log filter
create log target xml-notification
create xml-notification target url
delete log filter
delete log target xml-notification
delete xml-notification target
disable cli-config-logging
disable cpu-monitoring
disable elsm ports
disable elsm ports auto-restart
disable log display
disable log target
disable log target xml-notification
disable rmon
disable sflow
disable sflow ports
disable sys-health-check
disable syslog
enable cli-config-logging
enable cpu-monitoring
enable elsm ports
enable elsm ports auto-restart
enable log display
enable log target
enable log target xml-notification
enable rmon
enable sflow
enable sflow ports
enable sys-health-check
enable syslog
enable/disable xml-notification
show configuration "xmlc"
show cpu-monitoring
show elsm
show elsm ports
show fans
show log
```




```
show log components
show log configuration
show log configuration filter
show log configuration target
show log configuration target xml-notification
show log counters
show log events
show ports rxerrors
show ports statistics
show ports txerrors
show ports vlan statistics
show rmon memory
show sflow configuration
show sflow statistics
show temperature
show version
show vlan statistics
show xml-notification configuration
show xml-notification statistics
unconfigure log filter
unconfigure log target format
unconfigure ports monitor vlan
unconfigure sflow
unconfigure sflow agent
unconfigure sflow collector
unconfigure sflow ports
unconfigure xml-notification
upload log
```

This chapter describes commands for:

- Configuring and managing the Event Management System/Logging.
- Configuring and monitoring system health and statistics.
- Enabling, disabling, and configuring the Extreme Link Status Monitoring (ELSM) protocol.
- Enabling and disabling the collection of remote monitoring (RMON) statistics on the switch.
- Enabling, disabling, and configuring sFlow™ statistics collection.
- Monitoring CPU utilization.

Event Management System

When an event occurs on a switch, the Event Management System (EMS) allows you to send messages generated by these events to a specified log target. You can send messages to the memory buffer, NVRAM, the console display, the current session, to a syslog host, or to the other Management Switch



Fabric Module (MSM) or Management Module (MM). The log messages contain configuration and fault information pertaining to the device. You can format the log messages to contain various items of information, but typically a message consists of:

- **Timestamp**—The timestamp records when the event occurred.
- **Severity level**:
 - **Critical**—A desired switch function is inoperable. The switch may need to be reset.
 - **Error**—A problem is interfering with normal operation.
 - **Warning**—An abnormal condition exists that may lead to a function failure.
 - **Notice**—A normal but significant condition has been detected; the system is functioning as expected.
 - **Info**—Actions and events that are consistent with expected behavior.
 - **Debug-Summary, Debug-Verbose, and Debug-Data**—Information that is useful when performing detailed trouble shooting procedures.

By default, log entries that are assigned a critical, error, or warning level are considered static entries and remain in the NVRAM log target after a switch reboot.

- **Component**—The component refers to the specific functional area to which the error refers.
- **Message**—The message contains the log information with text that is specific to the problem.

The switch maintains a configurable number of messages in its internal (memory-buffer) log (1000 by default). You can display a snapshot of the log at any time. In addition to viewing a snapshot of the log, you can configure the system to maintain a running real-time display of log messages on the console display or telnet session. In addition to maintaining an internal log, the switch supports remote logging by way of the UNIX syslog host facility.

Beginning with ExtremeXOS 11.2, EMS supports IPv6 as a parameter for filtering events.

Extreme Link Status Monitoring

ExtremeXOS 11.4 introduces support for the Extreme Link Status Monitoring (ELSM) protocol. ELSM is an Extreme Networks proprietary protocol that monitors network health by detecting CPU and remote link failures. ELSM is available only on Extreme Networks devices and operates on a point-to-point basis. You configure ELSM on the ports that connect to other network devices and on both sides of the peer connection.

ELSM monitors network health by exchanging various hello messages between two ELSM peers. ELSM uses an open-ended protocol, which means that an ELSM-enabled port expects to send and receive hello messages from its peer. The Layer 2 connection between ports determines the peer connection. Peers can be either directly connected or separated by one or more hubs. If there is a direct connection between peers, they are considered neighbors.

If ELSM detects a failure, the ELSM-enabled port responds by blocking traffic on that port. For example, if a peer stops receiving messages from its peer, ELSM brings down that connection. ELSM does this by blocking all incoming and outgoing data traffic on the port and notifying applications that the link is down.

In some situations, a software or hardware fault may prevent the CPU from transmitting or receiving packets, thereby leading to the sudden failure of the CPU. If the CPU is unable to process or send



packets, ELSM isolates the connections to the faulty switch from the rest of the network. If the switch fabric sends packets during a CPU failure, the switch may appear healthy when it is not. For example, if hardware forwarding is active and software forwarding experiences a failure, traffic forwarding may continue. Such failures can trigger control protocols such as Extreme Standby Router Protocol (ESRP) or Ethernet Automatic Protection Switching (EAPS) to select different devices to resume forwarding. This recovery action, combined with the CPU failure, can lead to loops in a Layer 2 network.

Configuring ELSM on Extreme Networks devices running ExtremeXOS is backward compatible with Extreme Networks devices running ExtremeWare.

sFlow Statistics

sFlow® is a technology for monitoring traffic in data networks containing switches and routers.

It relies on statistical sampling of packets from high-speed networks, plus periodic gathering of the statistics. A User Datagram Protocol (UDP) datagram format is defined to send the information to an external entity for analysis. sFlow consists of a (Management Information Base) MIB and a specification of the packet format for forwarding information to a remote agent. Details of sFlow specifications can be found in RFC 3176 and at the following website: www.sflow.org

ExtremeXOS allows you to collect sFlow statistics on a per port basis. An agent, residing locally on the switch, sends data to a collector that resides on another machine. You configure the local agent, the address of the remote collector, and the ports of interest for sFlow statistics gathering. You can also modify default values for how frequently on average a sample is taken, how often the data is sent to the collector, and the maximum load allowed on the CPU before throttling the statistics gathering.

Licensing

For information about software licensing, including how to obtain and upgrade your license, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

RMON

RMON is the common abbreviation for the Remote Monitoring Management Information Base (MIB) system defined by the Internet Engineering Task Force (IETF) documents RFC 1757 and RFC2021, which allows you to monitor LANs remotely. Using the RMON capabilities of the switch allows network administrators to improve system efficiency and reduce the load on the network.

The IETF defines nine groups of Ethernet RMON statistics. The switch supports the following four of these groups, as defined in RFC1757:

- Statistics.
- History.
- Alarms.
- Events.

The switch also supports the following parameters for configuring the RMON probe and the trap destination table, as defined in RFC2021:



- probeCapabilities.
- probeSoftwareRev.
- probeHardwareRev.
- probeDateTime.
- probeResetControl.
- trapDestTable.

clear counters

clear counters

Description

Clears all switch statistics and port counters, including port packet statistics, bridging statistics, IP statistics, and log event counters.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

You should view the switch statistics and port counters before you clear them. Use the `show ports` command to view port statistics. Use the `show log counters` command to show event statistics.

The CLI also provides a number of options that you can specify with the `clear counters` command. If you specify an option, the switch only clears the statistics for that option. For example, if you want to clear, reset only the STP statistics and counters, use the `clear counters stp` command. Please refer to the specific chapter in the *ExtremeXOS Command Reference Guide* for more detailed information about those commands.

Viewing and maintaining statistics on a regular basis allows you to see how well your network is performing. If you keep simple daily records, you will see trends emerging and notice problems arising before they cause major network faults. By clearing the counters, you can see fresh statistics for the time period you are monitoring.

Example

The following command clears all switch statistics and port counters:

```
clear counters
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

clear counters xml-notification

```
clear counters xml-notification {all | target}
```

Description

Clears the statistics counters.

Syntax Description

<i>target</i>	Specifies an alpha numeric string that identifies the configured target.
---------------	--

Default

N/A.

Usage Guidelines

Use this command to unconfigure and reset all statistics counters.

Example

The following command clears all of the xml-notification statistics counters:

```
clear counters xml-notification all
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on BlackDiamondX8, BlackDiamond 8000 series modules, and Summit Family switches.




clear cpu-monitoring

```
clear cpu-monitoring {process name} {slot slotid}
```

Description

Clears, resets the CPU utilization history and statistics stored in the switch.

Syntax Description

<i>name</i>	Specifies the name of the process.
<i>slotid</i>	Specifies the slot number of the MSM/MM module: A specifies the MSM installed in slot A. B specifies the MSM installed in slot B.
	 Note This parameter is available only on modular switches.

Default

N/A.

Usage Guidelines

When you do not specify any keywords, this command clears the CPU utilization history for the entire switch, including processes, and resets the statistics to zero (0). On modular switches, this command also clears the CPU utilization history of the installed MSMs/MMs.

When you specify process, the switch clears and resets the CPU utilization history for the specified process.

Modular Switches Only

When you specify slot, the switch clears and resets the CPU utilization history for the specified MSM/MM.

Example

The following command resets the CPU history and resets the statistics to 0 for the TFTP process running on the MSM/MM installed in slot A of a modular switch:

```
clear cpu-monitoring process tftpd slot A
```

The following command resets the CPU history and resets statistics to 0 for the TFTP process running on a Summit family switch:

```
clear cpu-monitoring process tftpd
```



History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

clear elsm ports auto-restart

```
clear elsm ports port_list auto-restart
```

Description

Clears one or more ELSM-enabled ports that are in the Down-Stuck state.

Syntax Description

<i>port_list</i>	Specifies the ELSM-enabled ports that are permanently in the Down-Stuck state.
------------------	--

Default

N/A.

Usage Guidelines

If you do not have automatic restart enabled, use this command to transition ELSM-enabled ports that are permanently in the Down-Stuck state to the Down state. You can also use the `enable elsm ports port_list auto-restart` command to transition a port from the Down-Stuck state to the Down state.

For information about the ELSM-enabled ports states, see the command `show elsm ports`.

If automatic restart is enabled (this is the default behavior), automatic restart automatically transitions the ports from the Down-Stuck state to the Down state. For more information, see the command `enable elsm ports auto-restart`.

Example

The following command transitions the ports from the Down-Stuck state to the Down state:

```
clear elsm ports 2:1-2:2 auto-restart
```

History

This command was first available in ExtremeXOS 11.4.



Platform Availability

This command is available on all platforms.

clear elsm ports counters

```
clear elsm {ports port_list} counters
```

Description

Clears the statistics gathered by ELSM for the specified ports or for all ports.

Syntax Description

<i>port_list</i>	Specifies the ELSM-enabled ports for which ELSM statistics are being cleared.
------------------	---

Default

N/A.

Usage Guidelines

You should view the ELSM statistics and counters before you clear them. To view ELSM-specific counter information, use the `show elsm ports all | port_list` command. To view summary ELSM information, including the ports configured for ELSM, use the `show elsm` command.

Use this command to clear only the ELSM-related counters. To clear all of the counters on the switch, including those related to ELSM, use the `clear counters` command.

Viewing and maintaining statistics on a regular basis allows you to see how well your network is performing. If you keep simple daily records, you will see trends emerging and notice problems arising before they cause major network faults. By clearing the counter, you can see fresh statistics for the time period you are monitoring.

Example

The following command clears the statistics gathered by ELSM for slot 2, ports 1-2:

```
clear elsm ports 2:1-2:2 counters
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.



clear log

```
clear log {error-led | static | messages [memory-buffer | nvram]}
```

Description

Clears the log messages in memory and NVRAM, and clears the ERR LED on the MSM/MM.

Syntax Description

error-led	Clears the ERR LED on the MSM/MM.
static	Specifies that the messages in the NVRAM and memory-buffer targets are cleared, and the ERR LED on the MSM/MM is cleared.
memory-buffer	Clears entries from the memory buffer.
nvram	Clears entries from NVRAM.

Default

N/A.

Usage Guidelines

The switch log tracks configuration and fault information pertaining to the device.

By default, log entries that are sent to the NVRAM remain in the log after a switch reboot. The `clear log` and `clear log messages memory-buffer` commands remove entries in the memory buffer target; the `clear log static` and `clear log messages nvram` commands remove messages from the NVRAM target. In addition, the `clear log static` command will also clear the memory buffer target.

On modular switches and SummitStack, there are three ways to clear the ERR LED: clear the log, reboot the switch, or use the `clear log error-led` command. To clear the ERR LED without rebooting the switch or clearing the log messages, use the `clear log error-led` command.

Execution of these commands on a backup or standby node results in the clearing of that node's information only. Execution of these commands on the master node results in the clearing of information on all nodes in the system.

Example

The following command clears all log messages, from the NVRAM:

```
clear log static
```

History

This command was first available in ExtremeXOS 10.1.



Platform Availability

This command is available on all platforms.

clear log counters

```
clear log counters [event-condition | [all | event-component] {severity severity
{only}}]
```

Description

Clears the incident counters for events.

Syntax Description

<i>event-condition</i>	Specifies the event condition counter to clear.
all	Specifies that all events counters are to be cleared.
<i>event-component</i>	Specifies that all the event counters associated with a particular component should be cleared.
<i>severity</i>	Specifies the minimum severity level of event counters to clear (if the keyword only is omitted).
only	Specifies that only event counters of the specified severity level are to be cleared.

Default

If severity is not specified, then the event counters of any severity are cleared in the specified component.

Usage Guidelines

This command sets the incident counters to zero for each event specified. To display event counters, use the following command:

```
show log counters
```

See the command `show log` for more information about severity levels.

To get a listing of the event conditions in the system, use the following command:

```
show log events {detail}
```

To get a listing of the components present in the system, use the following command:

```
show log components
```

In a SummitStack, execution of these commands on a backup or standby node results in the clearing of that node's information only. Execution of these commands on the master node results in the clearing of information on all nodes in the system.



Example

The following command clears the event counters for event conditions of severity error or greater in the component BGP:

```
clear log counters "BGP" severity error
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

clear sys-recovery-level

clear sys-recovery-level

Description

If configured and the switch detects a hardware fault and enters the shutdown state, this command clears the shutdown state and renders the switch, I/O, or MSM/MM module(s) operational.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

If you configure the switch or one or more modules to shutdown upon detecting a hardware fault, and the switch or module enters the shutdown state, you must explicitly clear the shutdown state and reset the switch or the affected modules for the switch to become operational.

To clear the shutdown state, use the following command:

```
clear sys-recovery-level
```

The switch prompts you to confirm this action. The following is a sample confirmation message:

```
Are you sure you want to clear sys-recovery-level? (y/n)
```

Enter y to confirm this action and clear the shutdown state. Enter n or press [Enter] to cancel this action.



Modular Switches Only

On a modular switch, after using the `clear sys-recovery-level` command, you must reset each affected module.

If you configured only a few I/O modules to shutdown, reset each affected I/O module as follows:

Disable the slot using the `disable slot slot` command.

Re-enable the slot using the `enable slot slot` command.



Note

You must complete this procedure for each module that enters the shutdown state.

If you configured all I/O modules or one or more MSMs/MMs to shut down, use the `reboot` command to reboot the switch and reset all affected modules.

After you clear the shutdown state and reset the affected module, each port is brought offline and then back online before the module and the entire system is operational.

Summit Family Switches Only

On a Summit family switch, after you clear the shutdown state, use the `reboot` command to bring the switch and ports back online. After you use the `reboot` command, the switch is operational.

Differences in the Command Line Prompt

When an exclamation point (!) appears in front of the command line prompt, it indicates that one or more slots of a modular switch or the entire stand-alone switch is shut down as a result of your hardware recovery configuration and a switch error.

The following is truncated sample output of the command line prompt for a modular switch:

```
The I/O modules in the following slots are shut down: 1,3
Use the "clear sys-recovery-level" command to restore I/O modules
! BD-8810.1 #
```

The following is sample output for a Summit switch:

```
All switch ports have been shut down.
Use the "clear sys-recovery-level" command to restore all ports.
! SummitX #
```

Example

The following command clears the shutdown state:

```
clear sys-recovery-level
```



History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

configure elsm ports hellotime

```
configure elsm ports port_list hellotime hello_time
```

Description

Configures the ELSM hello timer by specifying the time between consecutive hello messages for the specified ports.

Syntax Description

<i>port_list</i>	Specifies the port or ports for which the ELSM hello timer should be configured.
<i>hello_time</i>	Specifies the time in seconds between consecutive hello messages. Use the same value for the hello interval on peer ports. The default value is 1 second, and the range is 1 to 128 seconds.

Default

The default is 1 second.

Usage Guidelines

ELSM works between two connected ports, and each ELSM instance is based on a single port.

When you enable ELSM on the specified ports, the ports participate in ELSM with their peers and begin exchanging ELSM hello messages.

ELSM uses two types of hello messages to communicate the health of the network to other ELSM ports:

- Hello+ — The ELSM-enabled port receives a hello message from its peer and no problem is detected.
- Hello- — The ELSM-enabled port does not receive a hello message from its peer.

ELSM also has hello transmit states. The hello transmit states display the current state of transmitted ELSM hello messages. For more information about the hello transmit states, see the `show elsm ports` command.

A high hello timer value can increase the time it takes for the ELSM-enabled port to enter the Up state. The down timer is $(2 + \text{hold threshold}) * \text{hello timer}$. Assuming the default value of 2 for the hold



threshold, configuring a hello timer of 128 seconds creates a down timer of $(2 + 2) 128$, or 512 seconds. In this scenario it would take 512 seconds for the port to transition from the Down to the Up state.

If you modify the hello timer on one port, we recommend that you use the same hello timer value on its peer port.

Example

The following command specifies 5 seconds between consecutive ELSM hello messages for slot 2, ports 1-2 on the switch:

```
configure elsm ports 2:1-2:2 hellotime 5
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure elsm ports hold-threshold

```
configure elsm ports port_list hold-threshold hold_threshold
```

Description

Configures the number of Hello+ messages required by the specified ELSM-enabled ports to transition from the Down-Wait state to the Up state.

Syntax Description

<i>port_list</i>	Specifies the port or ports for which the ELSM hold threshold should be configured.
<i>hold_threshold</i>	Specifies the number of Hello+ messages required to transition from the Down-Wait state to the Up state. The default is 2 messages, and the range is 1 to 40 messages.

Default

The default is 2 Hello+ messages.

Usage Guidelines

The port begins in the Down state, so the first received Hello+ message transitions the ELSM-enabled port from the Down state to the Down-Wait state. After that transition, the configured hold-threshold



value determines the number of Hello+ messages required to transition from Down-Wait state to the Up state.

The ELSM hold threshold determines the number of Hello+ messages the ELSM peer port must receive to transition from the Down-Wait state to the Up state. For example, a threshold of 1 means the ELSM port must receive at least one Hello+ message to transition from the Down-Wait state to the Up state.

After the down timer expires, the port checks the number of Hello+ messages against the hold threshold. If the number of Hello+ messages received is greater than or equal to the configured hold threshold, the ELSM receive port moves from the Down-Wait state to the Up state.

If the number of Hello+ messages received is less than the configured hold threshold, the ELSM receive port moves from the Down-Wait state back to the Down state and begins the process again.

If you modify the hold threshold on one port, we recommend that you use the same hold threshold value on its peer port.

You configure the hold threshold on a per-port basis, not on a per-switch basis.

Example

The following command specifies that two Hello+ messages are required for the ELSM receive ports configured on slot 2, ports 1-2, to transition from the Down-Wait state to the Up state:

```
configure elsm hold-threshold 2 ports 2:1-2:2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure elsm ports uptimer-threshold

```
configure elsm ports port_list uptimer-threshold uptimer_threshold
```

Description

Configures the number of Hello+ messages required by the specified ELSM-enabled ports to transition from the Up state to the Down state.



Syntax Description

<i>port_list</i>	Specifies the port or ports for which the ELSM hold threshold should be configured.
<i>uptimer_threshold</i>	Specifies the number of Hello+ messages required to transition from the Up-state to the Down state. The default is 6 messages, and the range is 3 to 60 messages.

Default

The default is 6 Hello+ messages.

Usage Guidelines

The ELSM up timer begins when the ELSM-enabled port enters the UP state. Each time the port receives a Hello+ message, the timer restarts. Up timer is `Uptimer_threshold * hello timer`. When the Up timer expires, it transits from UP state to DOWN state.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

configure log display

configure log display severity {only}

Description

Configures the real-time log-level message to display.

Syntax Description

<i>severity</i>	Specifies a message severity. Severities include critical, error, warning, notice, info, debug-summary, debug-verbose, and debug-data.
only	Specifies only log messages of the specified severity level.

Default

If not specified, messages of all severities are displayed on the console display.



Usage Guidelines

You must enable the log display before messages are displayed on the log display. Use the `enable log display` command to enable the log display. This allows you to configure the system to maintain a running real-time display of log messages on the console.

Severity filters the log to display messages with the selected severity or higher (more critical). Severities include critical, error, warning, info, notice, debug-summary, debug-verbose, and debug-data.

You can also control log data to different targets. The command equivalent to `configure log display` is the following:

```
configure log target console-display severity severity
```

To display the current configuration of the log display, use the following command:

```
show log configuration target console-display
```

In a stack, this command is applicable only to Master and Backup nodes and not applicable to the standby nodes.

Example

The following command configures the system log to maintain a running real-time display of log messages of critical severity or higher:

```
configure log display critical
```

The following command configures the system log to maintain a running real-time display of only log messages of critical severity:

```
configure log display critical only
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure log filter events

```
configure log filter name [add | delete] {exclude} events [event-condition | [all  
| event-component] {severity severity {only}}]
```



Description

Configures a log filter to add or delete detailed feature messages based on a specified set of events.

In a stack, this command is applicable only to Master and Backup nodes and not applicable to the standby nodes.

Syntax Description

<i>name</i>	Specifies the filter to configure.
add	Add the specified events to the filter.
delete	Remove the specified events from the filter.
exclude	Events matching the specified events will be excluded.
<i>event-condition</i>	Specifies an individual event.
all	Specifies all components and subcomponents.
<i>event-component</i>	Specifies all the events associated with a particular component.
<i>severity</i>	Specifies the minimum severity level of events (if the keyword only is omitted).
only	Specifies only events of the specified severity level.

Default

If the exclude keyword is not used, the events will be included by the filter. If severity is not specified, then the filter will use the component default severity threshold (see the note: [If no severity is specified when delete or exclude is specified, severity all is used when delete or exclude is specified](#)).

Usage Guidelines

This command controls the incidents that pass a filter by adding, or deleting, a specified set of events. If you want to configure a filter to include or exclude incidents based on event parameter values (for example, MAC address or BGP Neighbor) see the command `configure log filter events match`.

When the add keyword is used, the specified event name is added to the beginning of the filter item list maintained for this filter. The new filter item either includes the events specified, or if the exclude keyword is present, excludes the events specified.

The delete keyword is used to remove events from the filter item list that were previously added using the add command. All filter items currently in the filter item list that are identical to, or a subset of, the set of events specified in the delete command will be removed.

Event Filtering Process

From a logical standpoint, the filter associated with each enabled log target is examined to determine whether a message should be logged to that particular target. The determination is made for a given filter by comparing the incident with the most recently configured filter item first. If the incident matches this filter item, the incident is either included or excluded, depending on whether the exclude



keyword was used. Subsequent filter items on the list are compared if necessary. If the list of filter items has been exhausted with no match, the incident is excluded.

Events, Components, and Subcomponents

As mentioned, a single event can be included or excluded by specifying the event's name. Multiple events can be added or removed by specifying an ExtremeXOS component name plus an optional severity. Some components, such as BGP, contain subcomponents, such as Keepalive, which is specified as BGP.Keepalive. Either components or subcomponents can be specified. The keyword all in place of a component name can be used to indicate all ExtremeXOS components.

Severity Levels

When an individual event name is specified following the events keyword, no severity value is needed since each event has pre-assigned severity. When a component, subcomponent, or the all keyword is specified following the events keyword, a severity value is optional. If no severity is specified, the severity used for each applicable subcomponent is obtained from the pre-assigned severity threshold levels for those subcomponents. For example, if STP were specified as the component, and no severity is specified for the add of an include item, then only messages with severity of error and greater would be passed, since the threshold severity for the STP component is error. If STP.InBPDU were specified as the component, and no severity is specified, then only messages with severity of warning and greater would be passed, since the threshold severity for the STP.InBPDU subcomponent is warning. Use the `show log components` command to see this information.

The severity keyword all can be used as a convenience when delete or exclude is specified. The use of delete (or exclude) with severity all deletes (or excludes) previously added events of the same component of all severity values.



Note

If no severity is specified when delete or exclude is specified, severity all is used.

If the only keyword is present following the severity value, then only the events in the specified component at that exact severity are included. Without the only keyword, events in the specified component at that severity or more urgent are included. For example, using the option severity warning implies critical, error, or warning events, whereas the option severity warning only implies warning events only. Severity all only is not a valid choice.

Any EMS events with severity debug-summary, debug-verbose, or debug-data will not be logged unless debug mode is enabled. See the command `enable log debug-mode`.

Filter Optimization

Each time a configure log filter command is issued for a given filter name, the events specified are compared against the current configuration of the filter to try to logically simplify the configuration.

For example, if the command:

```
configure log filter bgpFilter1 add events bgp.keepalive severity error only
```

were to be followed by the command:



```
configure log filter bgpFilter1 add events bgp severity info
```

the filter item in the first command is automatically deleted since all events in the BGP.Keepalive subcomponent at severity error would be also included as part of the second command, making the first command redundant.

More Information

See the command `show log` for more information about severity levels.

To get a listing of the components present in the system, use the following command:

```
show log components
```

To get a listing of event condition definitions, use the following command:

```
show log events
```

To see the current configuration of a filter, use the following command:

```
show log configuration filter {filter name}
```

Example

The following command adds all STP component events at severity info to the filter mySTPFilter:

```
configure log filter myStpFilter add events stp severity info
```

The following command adds the STP.OutBPDU subcomponent, at the pre-defined severity level for that component, to the filter myStpFilter:

```
configure log filter myStpFilter add events stp.outbpdu
```

The following command excludes one particular event, STP.InBPDU.Drop, from the filter:

```
configure log filter myStpFilter add exclude events stp.inbpdu.drop
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure log filter events match



```
configure log filter name [add | delete] {exclude} events [event-condition | [all
| event-component] {severity severity {only}}] [match | strict-match] type value
```

Description

Configures a log filter to add or delete detailed feature messages based on a specified set of events and match parameter values.

In a stack, this command is applicable only to Master and Backup nodes and not applicable to the standby nodes.

Syntax Description

<i>name</i>	Specifies the filter to configure.
add	Add the specified events to the filter.
delete	Remove the specified events from the filter.
exclude	Events matching the filter will be excluded.
<i>event-condition</i>	Specifies the event condition.
all	Specifies all events.
<i>event-component</i>	Specifies all the events associated with a particular component.
<i>severity</i>	Specifies the minimum severity level of events (if the keyword only is omitted).
only	Specifies only events of the specified severity level.
match	Specifies events whose parameter values match the <i>type value</i> pair.
strict-match	Specifies events whose parameter values match the <i>type value</i> pair, and possess all the parameters specified.
<i>type</i>	Specifies the type of parameter to match. For more information about types and values see Types and Values .
<i>value</i>	Specifies the value of the parameter to match. For more information about types and values see Types and Values .

Default

If the **exclude** keyword is not used, the events will be included by the filter. If severity is not specified, then the filter will use the component default severity threshold (see the note on [note: If no severity is specified when delete or exclude is specified, severity all is used when delete or exclude is specified](#)).

Usage Guidelines

This command controls the incidents that pass a filter by adding or deleting a specified set of events that match a list of *type value* pairs. This command is an extension of the command `configure log filter events`, and adds the ability to filter incidents based on matching specified event parameter values to the event.



See the `configure log filter events` command `configure log filter events` for more information on specifying and using filters, on event conditions and components, and on the details of the filtering process. The discussion here is about the concepts of matching *type value* pairs to more narrowly define filters.

Types and Values

Each event in ExtremeXOS is defined with a message format and zero or more parameter types. The `show log events` command `show log events` can be used to display event definitions (the event text and parameter types). The syntax for the parameter types (represented by *type* in the command syntax above) is:

```
[address-family [ipv4-multicast | ipv4-unicast | ipv6-multicast | ipv6-unicast] |
bgp-neighbor ip address | bgp-routeridip address | eaps eaps domain name |
{destination | source} [ipaddress ip address | L4-port | mac-address ] | esrpesrp
domain name | {egress | ingress} [slotslot number | portsport_list] | ipaddressip
address | L4-portL4-port | mac-addressmac_address | netmask netmask | number
number | portport_list | processprocess name | slot slotid | stringexact string
to be matched | vlanvlan name | vlan tagvlan tag]
```



Note

The slot parameters are available only on modular switches.

Beginning with ExtremeXOS 11.2, you can specify the `ipaddress` type as IPv4 or IPv6, depending on the IP version. The following examples show how to configure IPv4 addresses and IPv6 addresses:

- IPv4 address.

To configure an IP address, with a mask of 32 assumed, use the following command:

```
configure log filter myFilter add events all match ipaddress 12.0.0.1
```

To configure a range of IP addresses with a mask of 8, use the following command:

```
configure log filter myFilter add events all match ipaddress 12.0.0.0/8
```

- IPv6 address.

To configure an IPv6 address, with a mask of 128 assumed, use the following command:

- `configure log filter myFilter add events all match ipaddress 3ffe::1`
- To configure a range of IPv6 addresses with a mask of 16, use the following command:
- `configure log filter myFilter add events all match ipaddress 3ffe::/16`

- IPv6 scoped address.

IPv6 scoped addresses consist of an IPv6 address and a VLAN. The following examples identify a link local IPv6 address.

To configure a scoped IPv6 address, with a mask of 128 assumed, use the following command:



Note

In the previous example, if you specify the VLAN name, it must be a full match; wild cards are not allowed.



The *value* depends on the parameter type specified. As an example, an event may contain a physical port number, a source MAC address, and a destination MAC address. To allow only those incidents with a specific source MAC address, use the following in the command:

```
configure log filter myFilter add events aaa.radius.requestInit severity notice
match source mac-address 00:01:30:23:C1:00 configure log filter myFilter add
events bridge severity notice match source mac-address 00:01:30:23:C1:00
```

The string type is used to match a specific string value of an event parameter, such as a user name. The exact string is matched with the given parameter and no regular expression is supported.

Match Versus Strict-Match

The match and strict-match keywords control the filter behavior for incidents whose event definition does not contain all the parameters specified in a `configure log filter events match` command. This is best explained with an example. Suppose an event in the XYZ component, named XYZ.event5, contains a physical port number, a source MAC address, but no destination MAC address. If you configure a filter to match a source MAC address and a destination MAC address, XYZ.event5 will match the filter when the source MAC address matches regardless of the destination MAC address, since the event contains no destination MAC address. If you specify the strict-match keyword, then the filter will never match, since XYZ.event5 does not contain the destination MAC address.

In other words, if the match keyword is specified, an incident will pass a filter so long as all parameter values in the incident match those in the match criteria, but all parameter types in the match criteria need not be present in the event definition.

More Information

See the command `show log` for more information about severity levels.

To get a listing of the components present in the system, use the following command:

```
show log components
```

To get a listing of event condition definitions, use the following command:

```
show log events
```

To see the current configuration of a filter, use the following command:

```
show log configuration filter {filter name}
```

Example

By default, all log targets are associated with the built-in filter, DefaultFilter. Therefore, the most straightforward way to send additional messages to a log target is to modify DefaultFilter. In the following example, the command modifies the built-in filter to allow incidents in the STP component, and all subcomponents of STP, of severity critical, error, warning, notice and info. For any of these



events containing a physical port number as a match parameter, limit the incidents to only those occurring on physical ports 3, 4 and 5 on slot 1, and all ports on slot 2:

```
configure log filter DefaultFilter add events stp severity info match ports
1:3-1:5, 2:*
```

If desired, issue the `unconfigure log DefaultFilter` command to restore the DefaultFilter back to its original configuration.

History

This command was first available in ExtremeXOS 10.1.

New parameter *type* values, including `esrp` and `eaps` were added in ExtremeXOS 11.0 and 11.1.

Support for IPv6 addresses was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure log target filter


```
configure log target [console | memory-buffer | primary-msm | primary-node |
backup-msm | backup-node | nvram | session | syslog [all | ipaddress | ipPort {vr
vr_name} [local0...local7]]] filter filter-name{severity severity {only}}
```

Description


Associates a filter to a target.

In a stack, this command is applicable only to Master and Backup nodes. This command is not applicable to standby nodes.

Syntax Description

target	Specifies the device to send the log entries.
console	Specifies the console display.
memory-buffer	Specifies the switch memory buffer.
primary-msm	Specifies the primary MSM.
	 Note This parameter is available only on modular switches.
primary-node	Specifies the primary node in a stack.



backup-msm	Specifies the backup MSM.
	 Note This parameter is available only on modular switches.
backup-node	Specifies the backup node in a stack.
nvr	Specifies the switch NVRAM.
session	Specifies the current session (including console display).
syslog	Specifies a syslog remote server.
all	Specifies all of the syslog remote servers.
<i>ipaddress</i>	Specifies the syslog IP address.
<i>ipPort</i>	Specifies the UDP port number for the syslog target.
<i>vr_name</i>	Specifies the virtual router that can reach the server IP address.
	 Note User-created VRs are supported only on the platforms listed for this feature in Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i> .
local10 ... local17	Specifies the local syslog facility.
<i>filter-name</i>	Specifies the filter to associate with the target.
<i>severity</i>	Specifies the minimum severity level to send (if the keyword only is omitted).
only	Specifies that only the specified severity level is to be sent.

Default

If severity is not specified, the severity level for the target is left unchanged. If a virtual router is not specified, VR-Mgmt is used.

Usage Guidelines

This command associates the specified filter and severity with the specified target. A filter limits messages sent to a target.

Although each target can be configured with its own filter, by default, all targets are associated with the built-in filter, DefaultFilter. Each target can also be configured with its own severity level. This provides the ability to associate multiple targets with the same filter, while having a configurable severity level for each target.

A message is sent to a target if the target has been enabled, the message passes the associated filter, the message is at least as severe as the configured severity level, and the message output matches the regular expression specified. By default, the memory buffer, NVRAM, primary MSM/MM, and backup MSM/MM targets are enabled. For other targets, use the command `enable log target`. The following table describes the default characteristics of each type of target.



Table 17: Default target log characteristics

Target	Enabled	Severity Level
console display	no	info
memory buffer	yes	debug-data
NVRAM	yes	warning
primary MSM/MM	yes	warning
backup MSM/MM	yes	warning
session	no	info
syslog	no	debug-data

The built-in filter, DefaultFilter, and a severity level of info are used for each new telnet session. These values may be overridden on a per-session basis using the `configure log target filter` command and specify the target as session. Use the following form of the command for per-session configuration changes:

```
configure log target session filter filtername {severity severity {only}}
```

Configuration changes to the current session target are in effect only for the duration of the session, and are not saved in FLASH memory. The session option can also be used on the console display, if the changes are desired to be temporary. If changes to the console-display are to be permanent (saved to FLASH memory), use the following form of the command:

```
configure log target console filter filtername {severity severity {only}}
```

Modular Switches Only

If the condition for the backup-msm target is met by a message generated on the primary, the event is sent to the backup MSM/MM. When the backup MSM/MM receives the event, it will see if any of the local targets (nvram, memory, or console) are matched. If so it gets processed. The session and syslog targets are disabled on the backup MSM/MM, as they are handled on the primary. If the condition for the primary-msm target is met by a message generated on the backup, the event is sent to the primary MSM.

Note that the backup-msm target is only active on the primary MSM/MM, and the primary-msm target is only active on the backup MSM/MM.

SummitStack only

The backup-node target is only active on the primary-node, and the primary-node target is active on backup-node and standby-nodes.



Example

The following command sends log messages to the previously syslog host at 10.31.8.25, port 8993, and facility local3, that pass the filter myFilter and are of severity warning and above:

```
configure log target syslog 10.31.8.25:8993 local3 filter myFilter severity
warning
```

The following command sends log messages to the current session, that pass the filter myFilter and are of severity warning and above:

```
configure log target session filter myFilter severity warning
```

History

This command was first available in ExtremeXOS 10.1.

The primary-msm and backup-msm options were first available in ExtremeXOS 11.0.

The ipPort parameter was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure log target format

For console display, session, memory buffer, and NVRAM targets:

```
configure log target [ console | session | memory-buffer | nvrnm ] format
[timestamp [ seconds | hundredths | none]] [date [ dd-Mmm-yyyy | yyyy-mm-dd |
Mmm-dd | mm-dd-yyyy | mm/dd/yyyy | dd-mm-yyyy | none]] {event-name [component |
condition | none]} {process-name} {process-slot} {severity} {source-line} {host-
name}
```

For syslog targets:

```
configure log target syslog [[all | [ipaddress|ipPort]] {vr vr_name} {local}]
format [timestamp [ seconds | hundredths | none]] [date [ dd-Mmm-yyyy | yyyy-mm-
dd | Mmm-dd | mm-dd-yyyy | mm/dd/yyyy | dd-mm-yyyy | none]] {event-name
[component | condition | none]} {process-slot} {severity} {priority} {host-name}
{source-line} {tag-id} {tag-name}
```



Description

Configures the formats of the displayed message, on a per-target basis.

In a stack, this command is applicable only to Master and Backup nodes and not applicable to the standby nodes.



Syntax Description

console	Specifies the console display.
session	Specifies the current session (including console display).
memory-buffer	Specifies the switch memory buffer.
nvrAm	Specifies the switch NVRAM.
syslog	Specifies a syslog target.
all	Specifies all remote syslog servers.
<i>ipaddress</i>	Specifies the syslog IP address.
<i>ipPort</i>	Specifies the UDP port number for the syslog target.
<i>vr_name</i>	Specifies the virtual router that can reach the server IP address.
	<div style="display: flex; align-items: center;">  <div> <p>Note User-created VRs are supported only on the platforms listed for this feature in Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i>.</p> </div> </div>
local0 ... local7	Specifies the local syslog facility.
timestamp	Specifies a timestamp formatted to display seconds, hundredths, or none.
date	Specifies a date formatted as specified, or none.
event-name	Specifies how detailed the event description will be. Choose from none, component or condition.
host-name	Specifies whether to include the syslog host name.
priority	Specifies whether to include the priority.
process-name	Specifies whether to include the internal process name.
process-slot	Specifies which slot number the message was generated.
	<div style="display: flex; align-items: center;">  <div> <p>Note This parameter is available only on modular switches.</p> </div> </div>
severity	Specifies whether to include the severity.
source-line	Specifies whether to include the source file name and line number.
tag-id	Specifies whether to include the tag ID.
tag-name	Specifies whether to include the tag name.

Default

The following defaults apply to console display, memory buffer, NVRAM, and session targets:

- timestamp—hundredths.
- date—mm-dd-yyyy.
- event-name—condition.
- process-name—off.
- process-slot—off (modular switches only).
- severity—on.



- source-line—off.
- host-name—off.

The following defaults apply to syslog targets (per RFC 3164):

- timestamp—seconds.
- date—mmm-dd.
- event-name—none.
- process-slot—off (modular switches only).
- severity—on.
- priority—on.
- host-name—off.
- source-line—off.
- tag-id—off.
- tag-name—on.

If a virtual router is not specified, VR-Mgmt is used.

Usage Guidelines

This command configures the format of the items that make up log messages. You can choose to include or exclude items and set the format for those items, but you cannot vary the order in which the items are assembled.

When applied to the targets console or session, the format specified is used for the messages sent to the console display or telnet session. Configuration changes to the session target, be it either a telnet or console display target session, are in effect only for the duration of the session, and are not saved in FLASH.

When this command is applied to the target memory-buffer, the format specified is used in subsequent `show log` and `upload log` commands. The format configured for the internal memory buffer can be overridden by specifying a format on the `show log` and `upload log` commands.

When this command is applied to the target syslog, the format specified is used for the messages sent to the specified syslog host.

Timestamps

Timestamps refer to the time an event occurred, and can be output in either seconds as described in RFC 3164 (for example, “13:42:56”), hundredths of a second (for example, “13:42:56.98”), or suppressed altogether. To display timestamps as hh:mm:ss, use the seconds keyword, to display as hh:mm:ss.HH, use the hundredths keyword, or to suppress timestamps altogether, use the none keyword. Timestamps are displayed in hundredths by default.

Date

The date an event occurred can be output as described in RFC 3164. Dates are output in different formats, depending on the keyword chosen. The following lists the date keyword options, and how the date “March 26, 2005” would be output:



- Mmm-dd—Mar 26.
- mm-dd-yyyy—03/26/2005.
- dd-mm-yyyy—26-03-2005.
- yyyy-mm-dd—2005-03-26.
- dd-Mmm-yyyy—26-Mar-2005.

Dates are suppressed altogether by specifying none. Dates are displayed as mm-dd-yyyy by default.

Event Names

Event names can be output as the component name only by specifying event-name component and as component name with condition mnemonic by specifying event-name condition, or suppressed by specifying event-name none. The default setting is event-name condition to specify the complete name of the events.

Host Name

The configured SNMP name of the switch can be output as HOSTNAME described in RFC 3164 by specifying host-name. The default setting is off.

Process Name

For providing detailed information to technical support, the (internal) ExtremeXOS task names of the applications detecting the events can be displayed by specifying process-name. The default setting is off.

Process Slot

For providing detailed information to technical support, the slot from which the logged message was generated can be displayed by specifying process-slot. The default setting is off. This is available only on modular switches.

Severity

A four-letter abbreviation of the severity of the event can be output by specifying severity on or suppressed by specifying severity off. The default setting is severity on. The abbreviations are: Crit, Erro, Warn, Noti, Info, Summ, Verb, and Data. These correspond to: Critical, Error, Warning, Notice, Informational, Debug-Summary, Debug-Verbose, and Debug-Data.

Source Line

For providing detailed information to technical support, the application source file names and line numbers detecting the events can be displayed by specifying source-line. The default setting is off. You must enable debug mode using the `enable log debug-mode` command to view the source line information. For messages generated prior to enabling debug mode, the source line information is not displayed.



Tag ID

The process-id of the (internal) ExtremeXOS process that generated the event that resulted in the log message can be displayed by specifying tag-id. The default setting is off.

Tag Name

The name of the log component to which the generated event belongs can be displayed by specifying tag-name. The default setting is on. The tag name would be the same as the output of event-name component.

Example

In the following example, the switch generates the identical event from the component Sntp, using three different formats.

Using the default format for the session target, an example log message might appear as:

```
05/29/2005 12:15:25.00 <Warn:Sntp.RslvSrvrFail> The Sntp server parameter
value (TheWrongServer.example.com) can not be resolved.
```

If you set the current session format using the following command:

```
configure log target session format timestamp seconds date mm-dd-yyyy event-
name component
```

The same example would appear as:

```
05/29/2005 12:16:36 <Warn:Sntp> The Sntp server parameter value
(TheWrongServer.example.com) can not be resolved.
```

To provide some detailed information to technical support, you set the current session format using the following command:

```
configure log target session format timestamp hundredths date mmm-dd event-
name condition source-line process-name
```

The same example would appear as:

```
May 29 12:17:20.11 Sntp: <Warn:Sntp.RslvSrvrFail> tSntpc: (sntpcLib.c:606)
The Sntp server parameter value (TheWrongServer.example.com) can not be
resolved.
```

History

This command was first available in ExtremeXOS 10.1.



The ipPort and host-name parameters were first introduced in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure log target match



```
configure log target [console | memory-buffer | nvram | primary-msm | primary-
node| backup-msm | backp-node | session | syslog [all | ipaddress | ipPort {vr
vr_name}{local0 ... local7}]] match [any |match-expression]
```

Description


Associates a match expression to a target.

In a stack, this command is applicable only on a Master and Backup nodes. This command is not applicable for standby nodes.

Syntax Description

console	Specifies the console display.
memory-buffer	Specifies the switch memory buffer.
nvram	Specifies the switch NVRAM.
primary-msm	Specifies the primary MSM.  Note This parameter is available only on modular switches.
primary-node	Specifies the primary node in a stack.
backup-msm	Specifies the backup MSM.  Note This parameter is available only on modular switches.
backup-node	Specifies the backup-node in a stack.
session	Specifies the current session (including console display).
syslog	Specifies a syslog target.
all	Specifies all of the remote syslog servers.
<i>ipaddress</i>	Specifies the syslog IP address.
<i>ipPort</i>	Specifies the UDP port number for the syslog target.



<i>vr_name</i>	Specifies the virtual router that can reach the server IP address.
	 <p>Note User-created VRs are supported only on the platforms listed for this feature in <i>Feature License Requirements in the ExtremeXOS Concepts Guide</i>.</p>
local0 ... local7	Specifies the local syslog facility.
any	Specifies that any messages will match. This effectively removes a previously configured match expression.
<i>match-expression</i>	Specifies a regular expression. Only messages that match the regular expression will be sent.

Default

By default, targets do not have a match expression. If a virtual router is not specified, VR-Mgmt is used.

Usage Guidelines

This command configures the specified target with a match expression. The filter associated with the target is not affected. A message is sent to a target if the target has been enabled, the message passes the associated filter, the message is at least as severe as the configured severity level, and the message output matches the regular expression specified.

See the command `show log` for a detailed description of simple regular expressions. By default, targets do not have a match expression.

Specifying `any` instead of `match-expression` effectively removes a match expression that had been previously configured, causing any message to be sent that has satisfied all of the other requirements.

To see the configuration of a target, use the following command:

```
show log configuration target {console | memory-buffer | nvrnm | primary-msm |
primary-node | backup-msm | backup-node | session | syslog {ipaddress | ipPort |
vrvr_name} {[local0 ... local7]}}
```

To see the current configuration of a filter, use the following command:

```
show log configuration filter {filter name}
```

Example

The following command sends log messages to the current session, that pass the current filter and severity level, and contain the string `user5`:

```
configure log target session match user5
```

History

This command was first available in ExtremeXOS 10.1.



The primary-msm and backup-msm options were first available in ExtremeXOS 11.0.

The ipPort parameter was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure log target severity



```
configure log target [console | memory-buffer | nvram | primary-msm | primary-node | backup-msm | backup-node | session | syslog [all | ipaddress | ipPort {vr_name} [local0...local7 ]]] {severity severity {only}}
```

Description


Sets the severity level of messages sent to the target.

In a stack, this command is applicable only to Master and Backup nodes. You cannot run this command on standby nodes.

Syntax Description

console	Specifies the console display.
memory-buffer	Specifies the switch memory buffer.
nvr am	Specifies the switch NVRAM.
primary-msm	Specifies the primary MSM.  Note This parameter is available only on modular switches.
primary-node	Specifies the primary node in a stack.
backup-msm	Specifies the backup MSM.  Note This parameter is available only on modular switches.
backup-node	Specifies the backup node in a stack.
session	Specifies the current session (including console display).
syslog	Specifies a syslog target.
all	Specifies all of the remote syslog servers.
<i>ipaddress</i>	Specifies the syslog IP address.
<i>ipPort</i>	Specifies the UDP port number for the syslog target.



<i>vr_name</i>	Specifies the virtual router that can reach the server IP address.
	 <p>Note User-created VRs are supported only on the platforms listed for this feature in Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i>.</p>
local0 ... local7	Specifies the local syslog facility.
<i>severity</i>	Specifies the least severe level to send (if the keyword only is omitted).
only	Specifies that only the specified severity level is to be sent.

Default

By default, targets are sent messages of the following severity level and above:

- console display—info.
- memory buffer—debug-data.
- NVRAM—warning.
- session—info.
- syslog—debug-data.
- primary MSM/MM—warning (modular switches only).
- backup MSM/MM—warning (modular switches only).
- primary node—warning (stack only).
- backup node—warning (stack only).

If a virtual router is not specified, VR-Mgmt is used.

Usage Guidelines

This command configures the specified target with a severity level. The filter associated with the target is not affected. A message is sent to a target if the target has been enabled, the message passes the associated filter, the message is at least as severe as the configured severity level, and the message output matches the regular expression specified.

See the command `show log` for a detailed description of severity levels.

To see the current configuration of a target, use the following command:

```
show log configuration target {console | memory-buffer | nvrn | primary-msm |
primary-node | backup-msm | backup-node | session | syslog {ipaddress | ipPort |
vrvr_name} {[local0 ... local7]}}
```

To see the current configuration of a filter, use the following command:

```
show log configuration filter {filter name}
```



Example

The following command sends log messages to the current session, that pass the current filter at a severity level of info or greater, and contain the string user5:

```
configure log target session severity info
```

History

This command was first available in ExtremeXOS 10.1.

The primary-msm and backup-msm options were first available in ExtremeXOS 11.0.

The ipPort parameter was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.


configure log target syslog

```
configure log target syslog [all | ipaddress | ipPort] {vr vr_name}  
{local10...local17} from source-ip-address
```

Description

Configures the syslog server's IP address for one or all syslog targets.

Syntax Description

syslog	Specifies a syslog target.
all	Specifies all of the remote syslog servers.
<i>ipaddress</i>	Specifies the syslog server's IP address.
<i>ipPort</i>	Specifies the UDP port number for the syslog target.
<i>vr_name</i>	Specifies the virtual router that can reach the server IP address.
	<div style="display: flex; align-items: center;">  <div> <p>Note User-created VRs are supported only on the platforms listed for this feature in Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i>.</p> </div> </div>
local10 ... local17	Specifies the local syslog facility.
<i>source-ip-address</i>	Specifies the local source IP address to use.



Default

If a virtual router is not specified, the following virtual routers are used:

- ExtremeXOS 10.1—VR-0.
- ExtremeXOS 11.0 and later—VR-Mgmt.

Usage Guidelines

Use this command to identify and configure the syslog server's IP address. By configuring a source IP address, the syslog server can identify from which switch it received the log message.

Options for configuring the remote syslog server include:

- `all`—Specifies all of the remote syslog server hosts.
- `ipaddress`—The IP address of the remote syslog server host.
- `ipPort`—The UDP port.
- `vr_name`—The virtual router that can reach the syslog host.
- `local0-local7`—The syslog facility level for local use.
- `from`—The local source IP address.

If you do not configure a source IP address for the syslog target, the switch uses the IP address in the configured VR that has the closed route to the destination.

Example

The following command configures the IP address for the specified syslog target named orange:

```
configure log target syslog orange from 10.234.56.78
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure log target xml-notification filter

```
configure log target xml-notification xml_target_name filter filter-name
{severity [[severity] {only}]}
```

Description

Configures a Web server target with an EMS filter.



Syntax Description

<i>xml_target_name</i>	Specifies the name of the xml notification target.
<i>filter-name</i>	Specifies the name of the EMS filter.
<i>severity</i>	Specifies the least severe level to send (if the keyword only is omitted).

Default

N/A.

Usage Guidelines

Use this command to configure a Web server target with an EMS filter. All EMS filters can be applied.

Example

The following command configures the Web server target test2 with EMS filter filtertest2:

```
configure log target xml-notification test filter filtertest2
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit family switches.

configure ports monitor vlan

```
configure ports [port_list|all] monitor vlan vlan_name {rx-only | tx-only}
```

Description

Starts counting VLAN statistics on a port or a group of ports.

Syntax Description

<i>port_list</i>	Specifies one or more ports. May be in the form: 1, 2, 3-5, 2:5, 2:6-2:8.
all	Specifies all ports.
<i>vlan_name</i>	Specifies a VLAN name.



rx-only	Specifies receive statistics. (BlackDiamond X8 series switches, BlackDiamond 8000 series modules and Summit family switches only.)
tx-only	Specifies transmit statistics. (BlackDiamond X8 series switches, BlackDiamond 8900 series modules and Summit X460, X480, and X670 switches only.)

Default

N/A.

Usage Guidelines

Use this command to configure access to VLAN statistics per port.

The rx-only and tx-only parameters are intended for, but not restricted to, use on ports that support both receive and transmit statistics. Ports on slots that do not support transmit statistics do not require explicit use of the rx-only keyword. In the absence of specifying either rx-only or tx-only, both RX and TX VLAN statistics are gathered if both are supported on the configured port.

When both receive and transmit statistics are configured and resources for either receive or transmit are not available, neither receive nor transmit statistics will be configured.

On BlackDiamond 8000 series modules and Summit family switches, the number of VLANs that can be monitored is dependent on filtering resources on the involved module or switch.

When per-port monitoring is configured, the following commands display the latest statistics directly from the hardware in real time. This information is not logged.

To display VLAN statistics at the port level, use the following command:

```
show ports {port_list} vlan statistics {no-refresh}
```

To display VLAN statistics at the VLAN level, use the following command:

```
show vlan {vlan_name} statistics {no-refresh}
```

Example

The following command configures per-port monitoring of transmit statistics for a set of ports for the VLAN named finance on a Summit X480 switch:

```
configure ports 2,3 monitor vlan finance tx-only
```

History

This command was first available in ExtremeXOS 12.0.

Support for BlackDiamond 8000 series modules, SummitStack, and Summit family switches was added in ExtremeXOS 12.5.



Platform Availability

This command is available only on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, SummitStack and Summit family switches.

configure sflow agent ipaddress

```
configure sflow agent {ipaddress} ipaddress
```

Description

Configures the sFlow agent's IP address.

Syntax Description

<i>ipaddress</i>	Specifies the IP address from which sFlow data is sent on the switch.
------------------	---

Default

The default configured IP address is 0.0.0.0, but the effective IP address is the management port IP address.

Usage Guidelines

This command allows you to configure the IP address of the sFlow agent. Typically, you would set this to the IP address used to identify the switch in the network management tools that you use. The agent address is stored in the payload of the sFlow data, and is used by the sFlow collector to identify each agent uniquely. The default configured value is 0.0.0.0, but the switch will use the management port IP address if it exists.

Both the commands `unconfigure ports monitor vlan` and `unconfigure sflow agent` will reset the agent parameter to the default.

Example

The following command sets the sFlow agent's IP address to 10.2.0.1:

```
configure sflow agent ipaddress 10.2.0.1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.




configure sflow collector ipaddress

```
configure sflow collector {ipaddress} ipaddress {port udp-port-number} {vr
vr_name}
```

Description

Configures the sFlow collector IP address.

Syntax Description

<i>ipaddress</i>	Specifies the IP address to send the sFlow data.
<i>udp-port-number</i>	Specifies the UDP port to send the sFlow data.
<i>vr_name</i>	Specifies from which virtual router to send the sFlow data.
<div style="display: flex; align-items: center;">  <div> <p>Note</p> <p>User-created VRs are supported only on the platforms listed for this feature in Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i>.</p> </div> </div>	

Default

The following values are the defaults for this command:

- UDP port number—6343.
- Virtual router—VR-Mgmt (previously called VR-0).

Usage Guidelines

This command allows you to configure where to send the sFlow data. You must specify an IP address for the sFlow data collector, and you may specify a particular UDP port, if your collector uses a non-standard port. You may also need to specify from which virtual router to send the data.

You can configure up to four sFlow collectors. Each unique IP address/UDP port/virtual router combination identifies a collector.

Both the commands `unconfigure ports monitor vlan` and `unconfigure sflow collector` will reset the collector parameters to the default.

Example

The following command specifies that sFlow data should be sent to port 6343 at IP address 192.168.57.1 using the virtual router VR-Mgmt:

```
configure sflow collector ipaddress 192.168.57.1
```



History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure sflow max-cpu-sample-limit

```
configure sflow max-cpu-sample-limit rate
```

Description

Configures the maximum number of sFlow samples handled by the CPU per second.

Syntax Description

<i>rate</i>	Specifies the maximum sFlow samples per second.
-------------	---

Default

The default value is 2000 samples per second.

Usage Guidelines

This command configures the maximum number of samples sent to the CPU per second. If this rate is exceeded, the internal sFlow CPU throttling mechanism kicks in to limit the load on the CPU.

Every time the limit is reached, the sample rate is halved (the value of number in the `configure sflow sample-rate number` or `configure sflow ports port_list sample-ratenum` command is doubled) on the slot (modular switch) or ports (stand-alone switch) on which maximum number of packets were received during the last snapshot.

This effectively halves the sampling frequency of all the ports on that slot or stand-alone switch with a sub-sampling factor of 1. The sampling frequency of ports on that slot or stand-alone switch with a sub-sampling factor greater than 1 will not change; the sub-sampling factor is also halved so that the same rate of samples are sent from that port.

The maximum CPU sample rate is based on the total number of samples received from all the sources. The valid range is 100 to 200000 samples per second.



Example

The following command specifies that the sFlow maximum CPU sample rate should be set to 4000 samples per second:

```
configure sflow max-cpu-sample-limit 4000
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure sflow poll-interval

```
configure sflow poll-interval seconds
```

Description

Configures the sFlow counter polling interval.

Syntax Description

<i>seconds</i>	Specifies the number of seconds between polling each counter. The value can range from 0 to 3600 seconds.
----------------	---

Default

The default polling interval is 20 seconds.

Usage Guidelines

Each sFlow statistics counter is polled at regular intervals, and this data is then sent to the sFlow collector. This command is used to set the polling interval. To manage CPU load, polling for sFlow enabled ports are distributed over the polling interval, so that all ports are not polled at the same instant. For example, if the polling interval is 20 seconds and there are twenty counters, data is collected successively every second.

Specifying a poll interval of 0 (zero) seconds disables polling.



Example

The following command sets the polling interval to 60 seconds:

```
configure sflow poll-interval 60
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure sflow ports sample-rate

```
configure sflow ports port_list sample-rate number
```

Description

Configures the sFlow per-port sampling rate.

Syntax Description

<i>port_list</i>	Specifies a list of ports.
<i>number</i>	Specifies the fraction (1/number) of packets to be sampled.

Default

The default number is 8192, unless modified by the `configure sflow sample-rate` command.

Usage Guidelines

This command configures the sampling rate on a particular set of ports, and overrides the system-wide value set in the `configure sflow sample-rate` command. The rate is rounded off to the next power of two, so if 400 is specified, the sample rate is configured as 512. The valid range is 256 to 536870912.

Summit Family Switches and BlackDiamond 8000 c-, e-, xl-, and xm-Series Modules, BlackDiamond X8 Series Switches

All ports on the switch or same I/O module are sampled individually.



Example

The following command sets the sample rate for the ports 4:6 to 4:10 to one packet out of every 16384:

```
configure sflow ports 4:6-4:10 sample-rate 16384
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure sflow sample-rate

```
configure sflow sample-rate number
```

Description

Configures the sFlow default sampling rate.

Syntax Description

<i>number</i>	Specifies the fraction (1/number) of packets to be sampled.
---------------	---

Default

The default number is 8192.

Usage Guidelines

This command configures the default sampling rate. This is the rate that newly enabled sFlow ports will have their sample rate set to. Changing this rate will not affect currently enabled sFlow ports. The rate is rounded off to the next power of two, so if 400 is specified, the sample rate is configured as 512. The valid range is 256 to 536870912.

Configuring a lower number for the sample rate means that more samples will be taken, increasing the load on the switch. Do not configure the sample rate to a number lower than the default unless you are sure that the traffic rate on the source is low.

Summit Family Switches, BlackDiamond X8 Series Switches, BlackDiamond 8000 c-, e-, xl-, and xm-Series Modules Only

The minimum rate that these platforms sample is 1 out of every 256 packets. If you configure a rate to be less than 256, the switch automatically rounds up the sample rate to 256.



Example

The following command sets the sample rate to one packet out of every 16384:

```
configure sflow sample-rate 16384
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure sys-health-check all level

```
configure sys-health-check all level [normal | strict]
```

Description

Configures how the ExtremeXOS software handles faults for BlackDiamond X8 and BlackDiamond 8800 series switches and Summit family switches.

Syntax Description

normal	Upon a fault detection, the switch only sends a message to the syslog. This is the default setting.
strict	Upon a fault detection, the switch takes the action configured by the <code>configure sys-recovery-level slot</code> or the command.

Default

The default setting is normal.

Usage Guidelines

On a BlackDiamond X8 series switch or a BlackDiamond 8800 series switch, use this command in conjunction with the `configure sys-recovery-level slot [all | slot_number] [none | reset | shutdown]` command to implement your network's fault handling strategy.

On the Summit family switches, use this command in conjunction with the `configure sys-recovery-level switch [none | reset | shutdown]` command to implement your network's fault handling strategy.

ExtremeXOS 11.5 enhances the number of switch-fabric tests completed and monitored by the polling module of the system health checker. Additionally with ExtremeXOS 11.5, you can now configure how ExtremeXOS handles a detected fault based on the configuration of the `configure sys-recovery-`



`level slot [all | slot_number] [none | reset | shutdown]` or the `configure sys-recovery-level switch [none | reset | shutdown]` command.

If you configure the `strict` parameter, the switch takes the action configured by the `configure sys-recovery-level slot` or the `configure sys-recovery-level switch` command, which can include logging only or restarting, rebooting, or shutting down the suspect device.

To maintain a smooth upgrade for devices running ExtremeXOS 11.4 and earlier, the switch-fabric tests introduced in ExtremeXOS 11.5 are set to only log error messages ('normal mode') by default. However, we recommend that you configure 'strict mode' so the system can attempt to recover by utilizing the action configured in the `configure sys-recovery-level slot` or the `configure sys-recovery-level switch` command (which by default is reset).

System Behavior for the BlackDiamond X8 Series Switches and BlackDiamond 8800 Series Switches

Depending on your switch configuration, the following table shows how the BlackDiamond X8 series switches and BlackDiamond 8800 series switches behave when the ExtremeXOS software detects a fault:

Table 18: System behavior for the BlackDiamond X8 and 8800 series switches

Fault Handling Configuration	Module Recovery Configuration	Behavior
<code>configure sys-health-check all level normal</code>	<code>configure sys-recovery-level slot none</code>	The switch sends messages to the syslog.
Same as above.	<code>configure sys-recovery-level slot reset</code>	Same as above.
Same as above.	<code>configure sys-recovery-level slot shutdown</code>	Same as above.
<code>configure sys-health-check all level strict</code>	<code>configure sys-recovery-level slot none</code>	Same as above.
Same as above.	<code>configure sys-recovery-level slot reset</code>	ExtremeXOS reboots the affected switch or module.
Same as above.	<code>configure sys-recovery-level slot shutdown</code>	ExtremeXOS shuts down the affected switch or module.

System Behavior for Summit Family Switches

Depending on your switch configuration, the following table shows how Summit family switches behave when the ExtremeXOS software detects a fault:

Table 19: System behavior for Summit family switches

Fault Handling Configuration	Hardware Recovery Configuration	Behavior
<code>configure sys-health-check all level normal</code>	<code>configure sys-recovery-level switch none</code>	The switch sends messages to the syslog.
Same as above.	<code>configure sys-recovery-level switch reset</code>	Same as above.



Table 19: System behavior for Summit family switches (continued)

Fault Handling Configuration	Hardware Recovery Configuration	Behavior
Same as above.	<code>configure sys-recovery-level switch shutdown</code>	Same as above.
<code>configure sys-health-check all level strict</code>	<code>configure sys-recovery-level switch none</code>	Same as above.
Same as above.	<code>configure sys-recovery-level switch reset</code>	ExtremeXOS reboots the affected switch.
Same as above.	<code>configure sys-recovery-level switch shutdown</code>	ExtremeXOS shuts down the affected switch.

Displaying the System Health Check Setting

To display the system health check setting, including polling and how ExtremeXOS handles faults on the switch, use the following command:

```
show switch
```

The system health check setting, displayed as SysHealth check, shows the polling setting and how ExtremeXOS handles faults. The polling setting appears as Enabled, and the fault handling setting appears in parenthesis next to the polling setting. In the following truncated output from a BlackDiamond 8810 switch, the system health check setting appears as SysHealth check: Enabled (Normal):

```
SysName:          TechPubs Lab
SysName:          BD-8810Rack3
SysLocation:
SysContact:      support@extremenetworks.com, +1 888 257 3000
System MAC:      00:04:96:1F:A2:60
SysHealth check: Enabled (Normal)
Recovery Mode:   None
System Watchdog: Enabled
```

If you use the strict parameter, which configures the switch to take the action configured by the `configure sys-recovery-level slot` or the `configure sys-recovery-level switch` command, (Strict) would appear next to Enabled.

Example

On a BlackDiamond 8800 series switch, the following command configures the switch to forward faults to be handled by the level set by the `configure sys-recovery-level slot` command:

```
configure sys-health-check all level strict
```



On Summit family switches, the following command configures the switch to forward faults to be handled by the level set by the `configure sys-recovery-level switch` command:

```
configure sys-health-check all level strict
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available only on the BlackDiamond X8 and 8800 series switches, SummitStack, and Summit family switches.

configure sys-health-check interval

```
configure sys-health-check interval interval
```

Description

Configures the frequency of sending backplane diagnostic packets.

Syntax Description

<i>interval</i>	BlackDiamond X8 and BlackDiamond 8800 series switches —Specifies the frequency of sending backplane diagnostic packets.If backplane diagnostic packets are enabled on a particular slot, the default value for sending diagnostic packets is 5 seconds on that slot.If only polling occurs (this is the system default), the default value is 5seconds. (The polling interval is not a user-configured parameter, and polling always occurs.)
-----------------	--

Default

Depending upon your platform, the following defaults apply:

- BlackDiamond X8 and BlackDiamond 8800 series switches:
 - If backplane diagnostics are enabled on a particular slot, the default for sending packets is 5 seconds on that slot.
 - The polling interval is always 5 seconds (this is a not a user-configured parameter).

Usage Guidelines

Use this command with the guidance of Extreme Networks Technical Support personnel.

The system health checker tests I/O modules and the backplane by forwarding backplane diagnostic packets. Use this command to configure the amount of time it takes for the packets to be forwarded and returned to the MSM.



To enable backplane diagnostic packets, use the `enable sys-health-check slot slot` command. With backplane diagnostic packets enabled on a specific slot, the interval option of the `configure sys-health-check interval` command specifies the frequency of sending backplane diagnostic packets. For example, if you specify an interval of 9, backplane diagnostic packets are sent every 9 seconds on only the enabled slot.

**Note**

Extreme Networks does not recommend configuring an interval of less than the default interval. Doing this can cause excessive CPU utilization.

BlackDiamond X8 and BlackDiamond 8800 Series Switches Only

By default, the system health checker always polls the control plane health between MSMs/MMs and I/O modules, monitors memory levels on the I/O module, monitors the health of the I/O module, and checks the health of applications and processes running on the I/O module. If the system health checker detects an error, the health checker notifies the MSM/MM.

You must enable the backplane diagnostic packets feature to send backplane diagnostic packets. If you enable this feature, the system health checker tests the data link for a specific I/O module every 5 seconds by default. The MSM/MM sends and receives diagnostic packets from the I/O module to determine the state and connectivity. If you disable backplane diagnostics, the system health checker stops sending backplane diagnostic packets.

Example

The following examples assume that you enabled backplane diagnostic packets on a specific I/O slot.

On the BlackDiamond 8800 series switches, the following command configures the backplane diagnostic packet interval to 8 seconds:

```
configure sys-health-check interval 8
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on the BlackDiamond X8 and BlackDiamond 8800 series switches.

configure sys-recovery-level


```
configure sys-recovery-level [all | none]
```



Description

Configures a recovery option for instances where a software exception occurs in ExtremeXOS.

Syntax Description

all	Configures ExtremeXOS to log an error into the syslog and reboot the system after any software task exception occurs.
none	Configures the recovery level to none. No action is taken when a software task exception occurs; there is no system reboot, which can cause unexpected switch behavior.
	 Note Use this parameter only under the guidance of Extreme Networks Technical Support personnel.

Default

The default setting is all.

Usage Guidelines

If the software fails, the switch automatically reboots or leaves the system in its current state. You must specify one of the following parameters for the system to respond to software failures:

- **all**—The system will send error messages to the syslog and reboot if any software task exception occurs.

On modular switches, this command sets the recovery level only for the MSMs/MMs. The MSM/MM should reboot only if there is a software exception that occurs on the MSM/MM. The MSM/MM should not reboot if a software exception occurs on an I/O module.

To set the recovery level for all slots (MSM/MM and I/O) use the `configure sys-recovery-level slot` command.

- **none**—No action is taken when a software task exception occurs. The system does not reboot, which can cause unexpected switch behavior.



Note

Use the none parameter only under the guidance of Extreme Networks Technical Support personnel.

The default setting and behavior is all. Extreme Networks strongly recommends using the default setting.

Displaying the System Recovery Setting

To display the software recovery setting on the switch, use the following command:

```
show switch
```



This command displays general switch information, including the software recovery level. The following truncated output from a Summit switch displays the software recovery setting (displayed as Recovery Mode):

```
SysName:          TechPubs Lab
SysLocation:
SysContact:       support@extremenetworks.com, +1 888 257 3000
System MAC:       00:04:96:20:B4:13
SysHealth check:  Enabled (Normal)
Recovery Mode:    All
System Watchdog:  Enabled
```



Note

All platforms display the software recovery setting as Recovery Mode.

Example

The following command configures a switch to not take an action when any software task exception occurs:

```
configure sys-recovery-level none
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure sys-recovery-level slot

```
configure sys-recovery-level slot [all | slot_number] [none | reset | shutdown]
```

Description

Configures a recovery option for instances where an exception occurs on the specified MSM/MM or I/O module.

Syntax Description

all	Specifies all slots of the MSM/MM and I/O module.
<i>slot_number</i>	Specifies the slot of the MSM/MM or I/O module. A and B—Indicate an MSM/MM1 through 10—Indicate an I/O module



none	Configures the MSM/MM or I/O module to maintain its current state regardless of the detected hardware fault. The offending MSM/MM or I/O module is not reset. For more information about the states of an MSM/MM or I/O module see the <code>show slot</code> command.
reset	Configures the offending MSM/MM or I/O module to reset upon a hardware fault detection. For more detailed information, see the Usage Guidelines described below.
shutdown	Configures the switch to shut down all slots/modules configured for shutdown upon fault detection. On the modules configured for shutdown, all ports in the slot are taken offline in response to the reported errors; however, the MSMs/MMs remain operational for debugging purposes only. ExtremeXOS logs fault, error, system reset, system reboot, and system shutdown messages to the syslog.

Default

The default setting is reset.

Usage Guidelines

Use this command for system auto-recovery upon detection of hardware problems. You can configure the MSMs/MMs or I/O modules installed in a modular switch to take no action, automatically reset, shutdown, or if dual MSMs/MMs are installed, failover to the other MSM/MM, if the switch detects a faulty MSM/MM or I/O module. This enhanced level of recovery detects faults in the ASICs as well as packet buses.

You must specify one of the following parameters for the system to respond to MSM/MM or I/O module failures:

- **none**—Configures the MSM/MM or I/O module to maintain its current state regardless of the detected fault. The offending MSM/MM or I/O module is not reset. ExtremeXOS logs fault and error messages to the syslog and notifies you that the errors are ignored. This does not guarantee that the module remains operational; however, the switch does not reboot the module.
- **reset**—Configures the offending MSM/MM or I/O module to reset upon fault detection. ExtremeXOS logs fault, error, system reset, and system reboot messages to the syslog.
- **shutdown**—Configures the switch to shut down all slots/modules configured for shutdown upon fault detection. On the modules configured for shutdown, all ports in the slot are taken offline in response to the reported errors; however, the MSMs/MMs remain operational for debugging purposes only. You must save the configuration, using the `save configuration` command, for it to take effect. ExtremeXOS logs fault, error, system reset, system reboot, and system shutdown messages to the syslog.

Depending on your configuration, the switch resets the offending MSM/MM or I/O module if fault detection occurs. An offending MSM/MM is reset any number of times, and the MSM/MM is not permanently taken offline. On the BlackDiamond 8800 series switches, an offending I/O module is reset a maximum of five times. After the maximum number of resets, the I/O module is permanently taken offline.



Messages Displayed

If you configure the hardware recovery setting to either none (ignore) or shutdown, the switch prompts you to confirm this action. The following is a sample shutdown message:

```
Are you sure you want to shutdown on errors? (y/n)
```

Enter y to confirm this action and configure the hardware recovery level. Enter n or press [Enter] to cancel this action.

Taking Ports Offline

Beginning with ExtremeXOS 11.5, you can configure the switch to shut down one or more modules upon fault detection by specifying the shutdown option. If you configure one or more slots to shut down and the switch detects a hardware fault, all ports in all of the configured shut down slots are taken offline in response to the reported errors. (MSMs are available for debugging purposes only.)

The affected module remains in the shutdown state across additional reboots or power cycles until you explicitly clear the shutdown state. If a module enters the shutdown state, the module actually reboots and the show slot command displays the state of the slot as Initialized; however, the ports are shut down and taken offline. For more information about clearing the shutdown state, see the `clear sys-recovery-level` command.

Module Recovery Actions--BlackDiamond 8800 Series Switches Only

The following table describes the actions module recovery takes based on your module recovery setting. For example, if you configure a module recovery setting of reset for an I/O module, the module is reset a maximum of five times before it is taken permanently offline.

From left to right, the columns display the following information:

- Module Recovery Setting—This is the parameter used by the `configure sys-recovery-level slot` command to distinguish the module recovery behavior.
- Hardware—This indicates the hardware that you may have in your switch.
- Action Taken—This describes the action the hardware takes based on the module recovery setting.

Table 20: Module Recovery Actions for the BlackDiamond X8 Series Switches and BlackDiamond 8800 Series Switches

Module Recovery Setting	Hardware	Action Taken
none		
	Single MSM	The MSM remains powered on in its current state. This does not guarantee that the module remains operational; however, the switch does not reboot the module.
	Dual MSM	The MSM remains powered on in its current state. This does not guarantee that the module remains operational; however, the switch does not reboot the module.
	I/O Module	The I/O module remains powered on in its current state. The switch sends error messages to the log and notifies you that the errors are ignored. This does not guarantee that the module remains operational; however, the switch does not reboot the module.



Table 20: Module Recovery Actions for the BlackDiamond X8 Series Switches and BlackDiamond 8800 Series Switches (continued)

Module Recovery Setting	Hardware	Action Taken
reset		
	Single MSM	Resets the MSM.
	Dual MSM	Resets the primary MSM and fails over to the backup MSM.
	I/O Module	Resets the I/O module a maximum of five times. After the fifth time, the I/O module is permanently taken offline.
shutdown		
	Single MSM	The MSM is available for debugging purposes only (the I/O ports also go down); however, you must clear the shutdown state using the <code>clear sys-recovery-level</code> command for the MSM to become operational. After you clear the shutdown state, you must reboot the switch. For more information see the <code>clear sys-recovery-level</code> command.
	Dual MSM	The MSM is available for debugging purposes only (the I/O ports also go down); however, you must clear the shutdown state using the <code>clear sys-recovery-level</code> command for the MSM to become operational. After you clear the shutdown state, you must reboot the switch. For more information see the <code>clear sys-recovery-level</code> command.
	I/O Module	Reboots the I/O module. When the module comes up, the ports remain inactive because you must clear the shutdown state using the <code>clear sys-recovery-level</code> command for the I/O module to become operational. After you clear the shutdown state, you must reset each affected I/O module or reboot the switch. For more information see the <code>clear sys-recovery-level</code> command.

Displaying the Module Recovery Setting

To display the module recovery setting, use the following command:

```
show slot
```

Beginning with ExtremeXOS 11.5, the `show slot` output has been modified to include the shutdown configuration. If you configure the module recovery setting to shutdown, the output displays an “E” flag that indicates any errors detected on the slot disables all ports on the slot. The “E” flag appears only if you configure the module recovery setting to shutdown.

Note



If you configure one or more slots for shut down and the switch detects a hardware fault on one of those slots, all of the configured slots enter the shutdown state and remain in that state until explicitly cleared.



If you configure the module recovery setting to none, the output displays an “e” flag that indicates no corrective actions will occur for the specified MSM/MM or I/O module. The “e” flag appears only if you configure the module recovery setting to none.

The following sample output displays the module recovery action. In this example, notice the flags identified for slot 2:

Slots	Type	Configured	State	Ports	Flags
-					
Slot-1	8900-G96T-c	8900-G96T-c	Operational	96	MB
Slot-2	8900-10G24X-c	8900-10G24X-c	Operational	24	MB E
Slot-3	8900-40G6X-xm	8900-40G6X-xm	Operational	24	MB
Slot-4	G48Xc	G48Xc	Operational	48	MB
Slot-5	G8Xc	G8Xc	Operational	8	MB
Slot-6			Empty	0	
Slot-7	G48Te2(PoE)	G48Te2(PoE)	Operational	48	MB
Slot-8	G48Tc	G48Tc	Operational	48	MB
Slot-9	10G4Xc	10G4Xc	Operational	4	MB
Slot-10			Empty	0	
MSM-A	8900-MSM128		Operational	0	
MSM-B	8900-MSM128		Operational	0	

Flags : M - Backplane link to Master is Active
 B - Backplane link to Backup is also Active
 D - Slot Disabled
 I - Insufficient Power (refer to "show power budget")
 e - Errors on slot will be ignored (no corrective action initiated)
 E - Errors on slot will disable all ports on slot

Note



In ExtremeXOS 11.4 and earlier, if you configure the module recovery setting to none, the output displays an “E” flag that indicates no corrective actions will occur for the specified MSM or I/O module. The “E” flag appears only if you configure the module recovery setting to none.

Displaying Detailed Module Recovery Information

To display the module recovery setting for a specific port on a module, including the current recovery mode, use the following command:

```
show slot slot
```

In addition to the information displayed with show slot, this command displays the module recovery setting configured on the slot. The following truncated output displays the module recovery setting (displayed as Recovery Mode) for the specified slot:

```
Slot-2 information:
State:                Operational
Download %:           100
Flags:                MB  E
Restart count:        0 (limit 5)
Serial number:        800264-00-01 0907G-00166
```




```

Hw Module Type:      8900-10G24X-c
SW Version:          15.2.0.26
SW Build:            v1520b26
Configured Type:     8900-10G24X-c
Ports available:     24
Recovery Mode:       Shutdown
Debug Data:          Peer=Operational
Flags : M - Backplane link to Master is Active
      B - Backplane link to Backup is also Active
      D - Slot Disabled
      I - Insufficient Power (refer to "show power budget")
      e - Errors on slot will be ignored (no corrective action initiated)
      E - Errors on slot will disable all ports on slot

```

Troubleshooting Module Failures

If you experience an I/O module failure, use the following troubleshooting methods when you can bring the switch offline to solve or learn more about the problem:

- Restarting the I/O module—Use the `disable slot slot` command followed by the `enable slot slot` command to restart the offending I/O module. By issuing these commands, the I/O module and its associated fail counter is reset. If the module does not restart, or you continue to experience I/O module failure, please contact Extreme Networks Technical Support.
- Running diagnostics—Use the `run diagnostics normal slot` command to run operational diagnostics on the offending I/O module to ensure that you are not experiencing a hardware issue. If the module continues to enter the failed state, please contact Extreme Networks Technical Support.

If you experience an MSM/MM failure, please contact Extreme Networks Technical Support.

Example

The following command configures a switch to not take an action if a hardware fault occurs:

```
configure sys-recovery-level slot none
```

History

This command was first available in ExtremeXOS 11.3.

The shutdown parameter was added in ExtremeXOS 11.5.

Platform Availability

This command is available only on modular switches.

configure sys-recovery-level switch

```
configure sys-recovery-level switch [none | reset | shutdown]
```



Description

Configures a recovery option for instances where a hardware exception occurs on Summit family switches.

Syntax Description

none	Configures the switch to maintain its current state regardless of the detected fault. The switch does not reboot or shutdown. ExtremeXOS logs fault and error messages to the syslog.
reset	Configures the switch to reboot upon detecting a hardware fault. ExtremeXOS logs fault, error, system reset, and system reboot messages to the syslog.
shutdown	Configures the switch to shut down upon detecting a hardware fault. All ports are taken offline in response to the reported errors; however, the management port remains operational for debugging purposes only. If the switch shuts down, it remains in this state across additional reboots or power cycles until you explicitly clear the shutdown state.

Default

The default setting is reset.

Usage Guidelines

Use this command for system auto-recovery upon detection of hardware problems. You can configure Summit family switches to take no action, automatically reboot, or shutdown if the switch detects a hardware fault. This enhanced level of recovery detects faults in the CPU.

You must specify one of the following parameters for the switch to respond to hardware failures:

- **none**—Configures the switch to maintain its current state regardless of the detected fault. The switch does not reboot or shutdown.
- **reset**—Configures the switch to reboot upon detecting a hardware fault.
- **shutdown**—Configures the switch to shutdown upon fault detection. All ports are taken offline in response to the reported errors; however, the management port remains operational for debugging purposes only.

Messages Displayed

If you configure the hardware recovery setting to either none (ignore) or shutdown, the switch prompts you to confirm this action by displaying a message similar to the following:

```
Are you sure you want to shutdown on errors? (y/n)
```

Enter y to confirm this action and configure the hardware recovery level. Enter n or press [Enter] to cancel this action.

Displaying the Hardware Recovery Setting

To display the hardware recovery setting, use the following command:



```
show switch
```

If you change the hardware recovery setting from the default (reset) to either none (ignore) or shutdown, the Recovery Mode output is expanded to include a description of the hardware recovery mode. If you keep the default behavior or return to reset, the Recovery Mode output lists only the software recovery setting.

The following truncated output from a Summit switch displays the software recovery and hardware recovery settings (displayed as Recovery Mode):

```
SysName:          TechPubs Lab
SysLocation:
SysContact:       support@extremenetworks.com, +1 888 257 3000
System MAC:       00:04:96:1F:A5:71
Recovery Mode:    All, Ignore
System Watchdog:  Enabled
```

If you configure the hardware recovery setting to none, the output displays “Ignore” to indicate that no corrective actions will occur on the switch. “Ignore” appears only if you configure the hardware recovery setting to none.

If you configure the hardware recovery setting to shutdown, the output displays “Shutdown” to indicate that the switch will shutdown if fault detection occurs. “Shutdown” appears only if you configure the hardware recovery setting to shutdown.

If you configure the hardware recovery setting to reset, the output displays only the software recovery mode.

Example

The following command configures the switch to not take an action if a hardware fault occurs:

```
configure sys-recovery-level switch none
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available only on the Summit family switches.

configure syslog add


```
configure syslog add [ipaddress | ipPort] {vr vr_name} [local0...local17]
```



Description

Configures the remote syslog server host address, and filters messages to be sent to the remote syslog target.

Syntax Description

<code>ipaddress</code>	Specifies the remote syslog server IP address.
<code>ipPort</code>	Specifies the UDP port number for the syslog target.
<code>vr_name</code>	Specifies the virtual router that can reach the server IP address. <div style="display: flex; align-items: center;">  <div> <p>Note User-created VRs are supported only on the platforms listed for this feature in Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i>.</p> </div> </div>
<code>local0 ... local7</code>	Specifies the local syslog facility.

Default

If a virtual router is not specified, VR-Mgmt is used. If UDP port is not specified, 514 is used.

Usage Guidelines

Options for configuring the remote syslog server include:

- `ipaddress`—The IP address of the remote syslog server host.
- `ipPort`—The UDP port.
- `vr_name`—The virtual router that can reach the syslog host.
- `local0-local7`—The syslog facility level for local use.

The switch log overwrites existing log messages in a wrap-around memory buffer, which may cause you to lose valuable information once the buffer becomes full. The remote syslog server does not overwrite log information, and can store messages in non-volatile files (disks, for example).

The `enable syslog` command must be issued in order for messages to be sent to the remote syslog server(s). Syslog is disabled by default. A total of four syslog servers can be configured at one time.

When a syslog server is added, it is associated with the filter `DefaultFilter`. Use the `configure log target filter` command to associate a different filter.

The syslog facility level is defined as `local0 - local7`. The facility level is used to group syslog data.

Example

The following command configures the remote syslog server target:

```
configure syslog 123.45.67.78 local1
```



History

This command was first available in ExtremeXOS 10.1.

The `ipPort` parameter was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure syslog delete

```
configure syslog delete [all | ipaddress | ipPort] {vr vr_name}
{local10...local17 }
```

```
configure syslog delete host name/ip {:udp-port} [local10...local17]
```

Description

Deletes a remote syslog server address.

Syntax Description

all	Specifies all remote syslog servers.
<i>ipaddress</i>	Specifies the remote syslog server IP address.
<i>ipPort</i>	Specifies the UDP port number for the syslog target.
<i>vr_name</i>	Specifies the virtual router that can reach the server IP address.
	<div data-bbox="636 1199 721 1283"></div> <div data-bbox="743 1184 1398 1299"> <p>Note User-created VRs are supported only on the platforms listed for this feature in Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i>.</p> </div>
local10 ... local17	Specifies the local syslog facility.

Default

If a virtual router is not specified, VR-Mgmt is used.

If a UDP port number is not specified, 514 is used.

Usage Guidelines

This command is used to delete a remote syslog server target.



Example

The following command deletes the remote syslog server with an IP address of 10.0.0.1:

```
configure syslog delete 10.0.0.1 local1
```

History

This command was first available in ExtremeXOS 10.1.

The ipPort parameter was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure xml-notification target

```
configure xml-notification target target [url url {vr vr_name} | user [none | user] | [encrypted-auth encrypted-auth] | [queue-size queue-size]]
```

Description

Configures the Web server target in the XML client.

Syntax Description

<i>target</i>	Specifies an alpha numeric string that identifies the configured target.
<i>url</i>	Specifies the Web server URL.
<i>vr_name</i>	Specifies the virtual router over which the XML client process can connect to a Web server to send push notifications.
<i>user</i>	Specifies the alpha numeric string identifying the Web server user.
<i>encrypted-auth</i>	Specifies the encrypted user authentication string.
<i>queue-size</i>	Specifies in numeric format, the size of the buffer that stores incoming events from ExtremeXOS software.

Default

N/A.

Usage Guidelines

Use this command to configure the Web server target in XML client process.



Example

The following command configures the target `target2` for the user `admin`:

```
configure xml-notification target target2 user admin
```

History

This command was first available in ExtremeXOS 12.4.

The virtual router option was added in ExtremeXOS 12.4.2.

Platform Availability

This command is available on BlackDiamond X8, BlackDiamond 8000 series modules, and Summit Family switches.

configure xml-notification target add/delete

```
configure xml-notification target target [add | delete] module
```

Description

Adds or deletes an EXOS module to or from the Web server target.

Syntax Description

<i>target</i>	Specifies an alpha numeric string that identifies the configured target.
<i>module</i>	Specifies the name of the ExtremeXOS module.

Default

N/A.

Usage Guidelines

Use the `add` option to attach a module to the Web server target in order to receive events from that application and send them to the targeted Web server. There is no limitation to the number of modules that can be attached.

Only Identity Management and EMS are supported targets.

Use the `delete` option to detach ExtremeXOS modules from the Web server target in order to stop receiving events from that module.



Example

The following command deleted the target test2 from EMS:

```
configure xml-notification target test2 ems
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on BlackDiamond X8, BlackDiamond 8000 series modules, and Summit Family switches.

create log filter

```
create log filter name {copy filter_name}
```

Description

Creates a log filter with the specified name.

Syntax Description

<i>name</i>	Specifies the name of the filter to create.
copy	Specifies that the new filter is to be copied from an existing one.
<i>filter_name</i>	Specifies the existing filter to copy.

Default

N/A.

Usage Guidelines

This command creates a filter with the name specified. A filter is a customizable list of events to include or exclude, and optional parameter values. The list of events can be configured by component or subcomponent with optional severity, or individual condition, each with optional parameter values. See the commands `configure log filter events` and `configure log filter events match` for details on how to add items to the filter.

The filter can be associated with one or more targets using the `configure log target filter` command to control the messages sent to those targets. The system has one built-in filter named DefaultFilter, which itself may be customized. Therefore, the `create log filter` command can be used if a filter other than DefaultFilter is desired. As its name implies, DefaultFilter initially contains the



default level of logging in which every ExtremeXOS component and subcomponent has a pre-assigned severity level.

If another filter needs to be created that will be similar to an existing filter, use the copy option to populate the new filter with the configuration of the existing filter. If the copy option is not specified, the new filter will have no events configured and therefore no incidents will pass through it.

The total number of supported filters, including DefaultFilter, is 20.

Example

The following command creates the filter named fdb2, copying its configuration from the filter DefaultFilter:

```
create log filter fdb2 copy DefaultFilter
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

create log target xml-notification

```
create log target xml-notification [ target_name | xml_target_name ]
```

Description

Creates a Web server XML-notification target name.

Syntax Description

<i>target_name</i>	Specifies the name of a non-existing XML notification target.
<i>xml_target_name</i>	Specifies the name of an already existing XML notification target.

Default

N/A.

Usage Guidelines

Use this command to create a Web server XML-notification target name for EMS.



Example

The following command creates the target name test2:

```
create log target xml-notification test2
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on BlackDiamond X8, BlackDiamond 8000 series modules, and Summit Family switches.

create xml-notification target url

```
create xml-notification target new-target url url {vr vr_name} {user [none | user]} {encrypted-auth encrypted-auth} {queue-size queue-size}
```

Description

Creates the Web server target in the XML client.

Syntax Description

<i>new-target</i>	Specifies a name for the target being created.
<i>url</i>	Specifies the Web server URL.
<i>vr_name</i>	Specifies the name of the virtual router over which the XML client process can connect to the Web server.
<i>user</i>	Specifies the name of the user.
<i>encrypted-auth</i>	Specifies the encrypted user authentication string.
<i>queue-size</i>	Specifies, in numeric format, the size of the buffer that stores incoming events from ExtremeXOS.

Default

N/A.



Usage Guidelines

Use this command to create the Web server target in the XML client process.



Note

You cannot enter a password in the CLI directly. It is a two-step process similar to creating a user account in ExtremeXOS.

Example

The following command creates a target target2 on `http://10.255.129.22:8080/xos/webservice` with a queue size of 100:

```
create xml-notification target target2 url http://10.255.129.22:8080/xos/webservice queue-size 100
```

History

This command was first available in ExtremeXOS 12.4.

The virtual router option was added in ExtremeXOS 12.4.2.

Platform Availability

This command is available on BlackDiamond X8, BlackDiamond 8000 series modules, and Summit Family switches.

delete log filter

Deletes a log filter with the specified name.

```
delete log filter [filter_name | all]
```

Syntax Description

<i>filter_name</i>	Specifies the filter to delete.
all	Specifies that all filters, except DefaultFilter, are to be deleted

Default

N/A.

Usage Guidelines

This command deletes the specified filter, or all filters except for the filter DefaultFilter. The specified filter must not be associated with a target. To remove that association, associate the target with DefaultFilter instead of the filter to be deleted, using the following command:

```
configure log target target filter DefaultFilter
```



Example

The following command deletes the filter named fdb2:

```
delete log filter fdb2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

delete log target xml-notification

```
delete log target xml-notification xml_target_name
```

Description

Deletes a Web server target.

Syntax Description

<i>xml_target_name</i>	Specifies the name of the xml notification target.
------------------------	--

Default

N/A.

Usage Guidelines

Use this command to delete a Web server target.

Example

The following command deleted the Web server target target2:

```
delete log target xml-notification target2
```

History

This command was first available in ExtremeXOS 12.4.



Platform Availability

This command is available on BlackDiamond X8, BlackDiamond 8000 series modules, and Summit family switches.

delete xml-notification target

```
delete xml-notification target target
```

Description

Deletes the Web server target on the XML client process.

Syntax Description

<i>target</i>	Specifies the configured target.
---------------	----------------------------------

Default

N/A.

Usage Guidelines

Use this command to delete the Web server target on the XML client process.

Example

The following command deletes the target test2:

```
delete xml-notification target test2
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on BlackDiamond X8, BlackDiamond 8000 series modules, and Summit family switches.

disable cli-config-logging

```
disable cli-config-logging
```



Description

Disables the logging of CLI configuration commands to the switch Syslog.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Every command is displayed in the log window which allows you to view every command executed on the switch.

The `disable cli-config-logging` command discontinues the recording of all switch configuration changes and their sources that are made using the CLI via Telnet or the local console. After you disable configuration logging, no further changes are logged to the system log.

To view the status of configuration logging on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for configuration logging.

Example

The following command disables the logging of CLI configuration command to the Syslog:

```
disable cli-config-logging
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

disable cpu-monitoring

```
disable cpu-monitoring
```

Description

Disables CPU monitoring on the switch.



Syntax Description

This command has no arguments or variables.

Default

CPU monitoring is enabled and occurs every 5 seconds.

Usage Guidelines

Use this command to disable CPU monitoring on the switch.

This command does not clear the monitoring interval. Therefore, if you altered the CPU monitoring interval, this command does not return the CPU monitoring interval to 5 seconds. To return to the default frequency level, use the `enable cpu-monitoring {interval seconds} {thresholdpercent}` and specify 5 for the interval.

Example

The following command disables CPU monitoring on the switch:

```
disable cpu-monitoring
```

History

This command was first available in an ExtremeXOS 11.2.

The default value shown began in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

disable elsm ports

```
disable elsm ports port_list
```

Description

Disables the ELSM protocol for the specified ports.

Syntax Description

<code>port_list</code>	Specifies the port or ports for which ELSM should be disabled.
------------------------	--



Default

The default is disabled.

Usage Guidelines

ELSM works between two connected ports, and each ELSM instance is based on a single port. When you disable ELSM on the specified ports, the ports no longer send ELSM hello messages to their peers and no longer maintain ELSM states.

When you enable ELSM on the specified ports, the ports participate in ELSM with their peers and begin exchanging ELSM hello messages. To enable ELSM, use the following command:

```
enable elsm ports port_list
```

For more information about ELSM, see the command `enable elsm ports`.

Example

The following command disables ELSM for slot 2, ports 1-2 on the switch:

```
disable elsm ports 2:1-2:2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

disable elsm ports auto-restart

```
disable elsm ports port_list auto-restart
```

Description

Disable ELSM automatic restart for the specified ports.

Syntax Description

<i>port_list</i>	Specifies the port or ports for which ELSM auto-restart is being disabled.
------------------	--

Default

The default is enabled.



Usage Guidelines

If you disable ELSM automatic restart, the ELSM-enabled port can transition between the following states multiple times: Up, Down, and Down-Wait. When the number of state transitions is greater than or equal to the sticky threshold, the port enters and remains in the Down-Stuck state.

The ELSM sticky threshold specifies the number of times a port can transition between the Up and Down states. The sticky threshold is not user-configurable and has a default value of 1. That means a port can transition only one time from the Up state to the Down state. If the port attempts a subsequent transition from the Up state to the Down state, the port enters the Down-Stuck state.

If the port enters the Down-Stuck state, you can clear the stuck state and have the port enter the Down state by using one of the following commands:

```
clear elsm ports port_list auto-restart enable elsm ports port_list auto-restart
```

If you use the `enable elsm ports port_list auto-restart` command, automatic restart is always enabled; you do not have to use the `clear elsm ports port_list auto-restart` command to clear the stuck state.

Enabling Automatic Restart

To enable ELSM automatic restart, you must explicitly configure this behavior on each ELSM-enabled port. If you enable ELSM automatic restart and an ELSM-enabled port goes down, ELSM bypasses the Down-Stuck state and automatically transitions the down port to the Down state, regardless of the number of times the port goes up and down.

To enable automatic restart, use the following command:

```
enable elsm ports port_list auto-restart
```

If you configure automatic restart on one port, Extreme Networks recommends that you use the same configuration on its peer port.

Example

The following command disables ELSM automatic restart for slot 2, ports 1-2 on the switch:

```
disable elsm ports 2:1-2:2 auto-restart
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.



disable log display

`disable log display`

Description

Disables the sending of messages to the console display.

In a stack, this command is applicable only to Master and Backup nodes and not applicable to the standby nodes.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

If the log display is disabled, log information is no longer written to the serial console.

This command setting is saved to FLASH and determines the initial setting of the console display at boot up.

You can also use the following command to control logging to different targets:

```
disable log display
```

The `disable log display` command is equivalent to `disable log target console-display` command.

Example

The following command disables the log display:

```
disable log display
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.



disable log target


```
disable log target [console | memory-buffer | nvram | primary-msm | primary-node
| backup-msm | backup-node | session | syslog [all | ipaddress | ipPort] {vr
vr_name} [local0 ... local7]]
```

Description

Stops sending log messages to the specified target.

In a stack, this command is applicable only to Master and Backup nodes and not applicable to the standby nodes.

Syntax Description

console	Specifies the console display.
memory-buffer	Specifies the switch memory buffer.
nvram	Specifies the switch NVRAM.
primary-msm	Specifies the primary MSM. NOTE: This parameter is available only on modular switches.
primary-node	Specifies the primary node in a stack.
backup-msm	Specifies the backup MSM. NOTE: This parameter is available only on modular switches.
backup-node	Specifies the backup node in a stack.
session	Specifies the current session (including console display).
syslog	Specifies a syslog target.
all	Specifies all of the remote syslog servers.
<i>ipaddress</i>	Specifies the syslog host name or IP address.
<i>ipPort</i>	Specifies the UDP port number for the syslog target.
<i>vr_name</i>	Specifies the virtual router that can reach the server IP address. <div style="display: flex; align-items: center;">  <div> <p>Note User-created VRs are supported only on the platforms listed for this feature in Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i>.</p> </div> </div>
local0 ... local7	Specifies the local syslog facility.

Default

Enabled, for memory buffer, NVRAM, primary MSM, and backup MSM/MM; all other targets are disabled by default.



Usage Guidelines

This command stops sending messages to the specified target. By default, the memory buffer, NVRAM, primary MSM/MM, and backup MSM/MM targets are enabled. Other targets must be enabled before messages are sent to those targets.

Configuration changes to the session target are in effect only for the duration of the console display or telnet session, and are not saved in FLASH. Changes to the other targets are saved to FLASH.

You can also use the following command to disable displaying the log on the console:

```
disable log display
```

The `disable log display` command is equivalent to `disable log target console-display` command.

Modular Switches Only

Note that the `backup-msm` target is only active on the primary MSM/MM, and the `primary-msm` target is only active on the backup MSM/MM.

Example

The following command disables log messages to the current session:

```
disable log target session
```

History

This command was first available in ExtremeXOS 10.1.

The `primary-msm` and `backup-msm` options were first available in ExtremeXOS 11.0.

The `ipPort` parameter was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

disable log target xml-notification

```
disable log target xml-notification xml_target_name
```

Description

Disables a Web server target.



Syntax Description

<code>xml_target_name</code>	Specifies the name of the xml-notification target.
------------------------------	--

Default

N/A.

Usage Guidelines

Use this command to disable a Web server EMS target.

Example

The following command disables the Web server target target2:

```
disable log target xml-notification target2
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on BlackDiamond X8, BlackDiamond 8000 series modules, and Summit family switches.

disable rmon

disable rmon

Description

Disables the collection of RMON statistics on the switch.

Syntax Description

This command has no arguments or variables.

Default

By default, RMON is disabled. However, even in the disabled state, the switch responds to RMON queries and sets for alarms and events.



Usage Guidelines

The switch supports four out of nine groups of Ethernet RMON statistics. In a disabled state, the switch continues to respond queries of statistics. Collecting of history, alarms, and events is stopped; however, the switch still queries old data.

To view the status of RMON polling on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for RMON polling.

To view the RMON memory usage statistics for a specific memory type (for example, statistics, events, logs, history, or alarms) or for all memory types, use the following command:

```
show rmon memory {detail | memoryType}
```

Example

The following command disables the collection of RMON statistics on the switch:

```
disable rmon
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

disable sflow

```
disable sflow
```

Description

Globally disables sFlow statistical packet sampling.

Syntax Description

This command has no arguments or variables.

Default

Disabled.



Usage Guidelines

This command disables sFlow globally on the switch. When you disable sFlow globally, the individual ports are also put into the disabled state. If you later enable the global sFlow state, individual ports return to their previous state.

Example

The following command disables sFlow sampling globally:

```
disable sflow
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

disable sflow ports

```
disable sflow ports port_list
```

Description

Disables sFlow statistical packet sampling and statistics gathering on a particular list of ports.

Syntax Description

<i>port_list</i>	Specifies a list of ports.
------------------	----------------------------

Default

Disabled.

Usage Guidelines

This command disables sFlow on a particular list of ports. Once sFlow is disabled on a port, sampling and polling will stop. If sFlow is disabled globally, all sampling and polling stops.

Use the following command to disable sFlow globally:

```
disable sflow
```



Example

The following command disables sFlow sampling on port 3:1:

```
disable sflow ports 3:1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

disable sys-health-check

```
disable sys-health-check slot slot
```

Description

Discontinues sending backplane diagnostic packets on modular switches.

Syntax Description

<i>slot</i>	Specifies the slot to disable sending backplane diagnostic packets.
-------------	---

Default

Polling is enabled, backplane diagnostic packets are disabled.

Depending upon your platform, when disabling backplane diagnostic packets, the following defaults apply:

- BlackDiamond 8800 series switches—By default, the system health checker discontinues sending backplane diagnostic packets to the specified slot. Only polling is enabled.

Usage Guidelines

When you use this command, backplane diagnostic packets are disabled and no longer sent by the system health checker.

BlackDiamond 8800 Series Switches Only

If you modify the interval in the `configure sys-health-check interval interval` command and later disable backplane diagnostics, the configured interval for sending backplane diagnostic packets remains. The next time you enable backplane diagnostic packets, the health checker sends backplane



diagnostics packets at the configured interval. For example, if you configure an interval of 8seconds, the system health checker sends backplane diagnostic packets every 8seconds.

To return to the "default" interval of 5seconds, configure the frequency of sending backplane diagnostic packets to 5 seconds using the following command:

```
configure sys-health-check interval 5
```

Example

On the BlackDiamond 8800 series switches, the following example assumes that you did not modify the interval option in the `configure sys-health-check interval interval` command.

The following command disables backplane diagnostics on slot 3, polling is always enabled and occurs every 5 seconds.

```
disable sys-health-check slot 3
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available only on modular switches.

disable syslog

disable syslog

Description

Disables logging to all remote syslog server targets.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Disables logging to all remote syslog server targets, not to the switch targets. This setting is saved in FLASH, and will be in effect upon boot up.



Example

The following command disables logging to all remote syslog server targets:

```
disable syslog
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable cli-config-logging

enable cli-config-logging

Description

Enables the logging of CLI configuration commands to the Syslog for auditing purposes.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

ExtremeXOS allows you to record all configuration changes and their sources that are made using the CLI by way of Telnet or the local console. The changes are logged to the system log. Each log entry includes the user account name that performed the changes and the source IP address of the client (if Telnet was used). Configuration logging applies only to commands that result in a configuration change.

To view the status of configuration logging on the switch, use the `show management` command. This command displays information about the switch including the enable/disable state for configuration logging.



Example

The following command enables the logging of CLI configuration commands to the Syslog:

```
enable cli-config-logging
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

enable cpu-monitoring

```
enable cpu-monitoring {interval seconds} {threshold percent}
```

Description

Enables CPU monitoring on the switch.

Syntax Description

<i>seconds</i>	Specifies the monitoring interval, in seconds. The default is 5 seconds, and the range is 5 to 60 seconds.
threshold	Specifies the CPU threshold value. CPU usage is measured in percentages. The default is 90%, and the range is 0% to 100%.

Default

CPU monitoring is enabled and occurs every 5 seconds. The default CPU threshold value is 90%.

Usage Guidelines

CPU monitoring allows you to monitor the CPU utilization and history for all of the processes running on the switch. By viewing this history on a regular basis, you can see trends emerging and identify processes with peak utilization. Monitoring the workload of the CPU allows you to troubleshoot and identify suspect processes before they become a problem.

To specify the frequency of CPU monitoring, use the `interval` keyword. We recommend the default setting for most network environments.

CPU usage is measured in percentages. By default, the CPU threshold value is 90%. When CPU utilization of a process exceeds 90% of the regular operating basis, the switch logs an error message specifying the process name and the current CPU utilization for the process. To modify the CPU threshold level, use the `threshold` keyword. The range is 0% to 100%.



Example

The following command enables CPU monitoring every 30 seconds:

```
enable cpu-monitoring interval 30
```

History

This command was first available in ExtremeXOS 11.2.

The default values shown began in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

enable elsm ports

```
enable elsm ports port_list
```

Description

Enables the ELSM protocol for the specified ports.

Syntax Description

<i>port_list</i>	Specifies the port or ports for which ELSM should be enabled.
------------------	---

Default

The default is disabled.

Usage Guidelines

The ELSM protocol allows you to detect CPU and remote link failures in the network. ELSM operates on a point-to-point basis; you only configure ELSM on the ports that connect to other devices within the network, but you must configure ELSM on both sides of the peer connections.

The Layer 2 connection between the ports determines the peer. You can have a direct connection between the peers or hubs that separate peer ports. In the first instance, the peers are also considered neighbors. In the second instance, the peer is not considered a neighbor.

An Extreme Networks device with ELSM enabled detects CPU and remote link failures by exchanging hello messages between two ELSM peers. If ELSM detects a failure, the ELSM-enabled port responds by blocking traffic on that port. For example, if a peer stops receiving messages from its peer, ELSM brings down that connection by blocking all incoming and outgoing data traffic on the port and notifying applications that the link is down.



Configuring and enabling ELSM on Extreme Networks devices running ExtremeXOS is backward compatible with Extreme Networks devices running ExtremeWare.

When you enable ELSM on a port, ELSM immediately blocks the port and it enters the Down state. When the port detects an ELSM-enabled peer, the peer ports exchange ELSM hello messages. At this point, the ports enter the transitional Down-Wait state. If the port receives Hello+ messages from its peer and does not detect a problem, the peers enter the Up state. If a peer detects a problem or there is no peer port configured, the port enters the Down state.

For more information about the types of ELSM hello messages, see the `configure elsm ports hellotime` command.



Note

ELSM and mirroring are mutually exclusive. You can enable either ELSM, or mirroring, but not both.

If you try to enable ELSM on a port that is already configured as a mirrored port, the switch displays messages similar to the following:

- Stand-alone switch

```
Cannot enable ELSM on port 1. Port is configured as mirror monitor
port
```

- Modular switch

```
Cannot enable ELSM on port 1:1. Port is configured as mirror monitor
port
```

```
Configuration failed on backup MSM, command execution aborted!
```

Configuring the Hello Timer Interval

ELSM ports use hello messages to communicate information about the health of the network to its peer port. You can also configure the interval by which the ELSM-enabled ports sends hello messages. For more information about configuring the hello interval, see the command `configure elsm ports hellotime`.

Disabling ELSM

ELSM works between two connected ports, and each ELSM instance is based on a single port. When you disable ELSM on the specified ports, the ports no longer send ELSM hello messages to their peers and no longer maintain ELSM states. To disable ELSM, use the following command:

```
disable elsm ports port_list
```

Example

The following command enables ELSM for slot 2, ports 1-2 on the switch:

```
enable elsm ports 2:1-2:2
```



History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

enable elsm ports auto-restart

```
enable elsm ports port_list auto-restart
```

Description

Enables ELSM automatic restart for the specified ports.

Syntax Description

<i>port_list</i>	Specifies the port or ports for which ELSM auto-restart is being enabled.
------------------	---

Default

The default is enabled.

Usage Guidelines

You must explicitly configure this behavior on each ELSM-enabled port; this is not a global command.

By default, ELSM automatic restart is enabled. If an ELSM-enabled port goes down, ELSM bypasses the Down-Stuck state and automatically transitions the down port to the Down state, regardless of the number of times the port goes up and down.

If you disable ELSM automatic restart, the ELSM-enabled port can transition between the following states multiple times: Up, Down, and Down-Wait. When the number of state transitions is greater than or equal to the sticky threshold, the port enters the Down-Stuck state.

The ELSM sticky threshold specifies the number of times a port can transition between the Up and Down states. The sticky threshold is not user-configurable and has a default value of 1. That means a port can transition only one time from the Up state to the Down state. If the port attempts a subsequent transition from the Up state to the Down state, the port enters the Down-Stuck state.

If the port enters the Down-Stuck state, you can clear the stuck state and enter the Down state by using one of the following commands:

```
clear elsm ports port_list auto-restart
```

OR

```
enable elsm ports port_list auto-restart
```



If you use the `enable elsm ports port_list auto-restart` command, automatic restart is always enabled; you do not have to use the `clear elsm ports port_list auto-restart` command to clear the stuck state.

To disable automatic restart, use the following command:

```
disable elsm ports port_list auto-restart
```

If you configure automatic restart on one port, Extreme Networks recommends that you use the same configuration on its peer port.

Example

The following command enables ELSM automatic restart for slot 2, ports 1-2 on the switch:

```
enable elsm ports 2:1-2:2 auto-restart
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

enable log display

enable log display

Description

Enables a running real-time display of log messages on the console display. In a stack, this command is applicable only to Master and Backup nodes. You cannot run this command on standby nodes.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

If you enable the log display on a terminal connected to the console port, your settings will remain in effect even after your console session is ended (unless you explicitly disable the log display).



You configure the messages displayed in the log using the `configure log display`, or configure log target console-display commands.

You can also use the following command to control logging to different targets:

```
enable log display
```

The `enable log display` command is equivalent to enable log target console-display command.

To change the log filter association, severity threshold, or match expression for messages sent to the console display, use the `configure log target console-display` command

Example

The following command enables a real-time display of log messages:

```
enable log display
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.


enable log target

```
enable log target [console | memory-buffer | nvram | primary-msm | primary-node |
backup-msm | backup-node | session | syslog [all | ipAddress | ipPort] {vr
vr_name} [local0...local7]]]
```



Description

Starts sending log messages to the specified target.

Syntax Description

console	Specifies the console display.
memory-buffer	Specifies the switch memory buffer.
nvram	Specifies the switch NVRAM.
primary-msm	Specifies the primary MSM.
	 Note This parameter is available only on modular switches.
primary-node	Specifies the primary node of a stack.



backup-msm	Specifies the backup MSM.  Note This parameter is available only on modular switches.
backup-node	Specifies the backup node of a stack.
session	Specifies the current session (including console display).
syslog	Specifies a syslog target.
all	Specifies all of the remote syslog servers.
<i>ipaddress</i>	Specifies the syslog IP address.
<i>ipPort</i>	Specifies the UDP port number for the syslog target.
<i>vr_name</i>	Specifies the virtual router that can reach the server IP address.  Note User-created VRs are supported only on the platforms listed for this feature in <i>Feature License Requirements in the ExtremeXOS Concepts Guide</i> .
local10 ... local17	Specifies the local syslog facility.

Default

Enabled, for memory buffer and NVRAM; all other targets are disabled by default.

Usage Guidelines

This command starts sending messages to the specified target. By default, the memory-buffer, NVRAM, primary MSM/MM, and backup MSM/MM targets are enabled. Other targets must be enabled before messages are sent to those targets.

Configuration changes to the session target are in effect only for the duration of the console display or Telnet session, and are not saved in FLASH. Others are saved in FLASH.

You can also use the following command to enable displaying the log on the console:

```
enable log display
```

This command is equivalent to the enable log target console-display command.

Modular Switches Only

Note that the backup-msm target is only active on the primary MSM/MM, and the primary-msm target is only active on the backup MSM/MM.



Example

The following command enables log messages on the current session:

```
enable log target session
```

History

This command was first available in ExtremeXOS 10.1.

The primary-msm and backup-msm options were first available in ExtremeXOS 11.0.

The ipPort parameter was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

enable log target xml-notification

```
enable log target xml-notification xml_target_name
```

Description

Enables a Web server target.

Syntax Description

<i>xml_target_name</i>	Specifies the name of the xml-notification target.
------------------------	--

Default

N/A.

Usage Guidelines

Use this command to enable a Web server target for EMS.

Example

The following command enables the Web server target target2:

```
enable log target xml-notification target2
```



History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on BlackDiamond X8, BlackDiamond 8000 series modules, and Summit family switches.

enable rmon

enable rmon

Description

Enables the collection of RMON statistics on the switch.

Syntax Description

This command has no arguments or variables.

Default

By default, RMON is disabled. However, even in the disabled state, the switch responds to RMON queries and sets for alarms and events. By enabling RMON, the switch begins the processes necessary for collecting switch statistics.

Usage Guidelines

The switch supports four out of nine groups of Ethernet RMON statistics. In an enabled state, the switch responds to the following four groups:

- **Statistics**—The RMON Ethernet Statistics group provides traffic and error statistics showing packets, bytes, broadcasts, multicasts, and errors on a LAN segment or VLAN.
- **History**—The History group provides historical views of network performance by taking periodic samples of the counters supplied by the Statistics group. The group features user-defined sample intervals and bucket counters for complete customization of trend analysis.
- **Alarms**—The Alarms group provides a versatile, general mechanism for setting threshold and sampling intervals to generate events on any RMON variable. Both rising and falling thresholds are supported, and thresholds can be on the absolute value of a variable or its delta value. In addition, alarm thresholds may be auto calibrated or set manually.
- **Events**—The Events group creates entries in an event log and/or sends SNMP traps to the management workstation. An event is triggered by an RMON alarm. The action taken can be configured to ignore it, to log the event, to send an SNMP trap to the receivers listed in the trap receiver table, or to both log and send a trap. The RMON traps are defined in RFC 1757 for rising and falling thresholds.

The switch also supports the following parameters for configuring the RMON agent, as defined in RFC2021:



- `probeCapabilities`—If you configure the `probeCapabilities` object, you can view the RMON MIB groups supported on at least one interface by the probe.
- `probeSoftwareRev`—If you configure the `probeSoftwareRev` object, you can view the current software version of the monitored device.
- `probeHardwareRev`—If you configure the `probeHardwareRev` object, you can view the current hardware version of the monitored device.
- `probeDateTime`—If you configure the `probeDateTime` object, you can view the current date and time of the probe.
- `probeResetControl`—If you configure the `probeResetControl` object, you can restart a managed device that is not running normally. Depending on your configuration, you can do one of the following:
 - Warm boot—A warm boot restarts the device using the current configuration saved in non-volatile memory.
 - Cold boot—A cold boot causes the device to reset the configuration parameters stored in non-volatile memory to the factory defaults and then restarts the device using the restored factory default configuration.

**Note**

You can only use the RMON features of the system if you have an RMON management application and have enabled RMON on the switch.

RMON requires one probe per LAN segment, and stand-alone RMON probes have traditionally been expensive. Therefore, the approach taken by Extreme Networks has been to build an inexpensive RMON probe into the agent of each system. This allows RMON to be widely deployed around the network without costing more than traditional network management. The switch accurately maintains RMON statistics at the maximum line rate of all of its ports.

For example, statistics can be related to individual ports. Also, because a probe must be able to see all traffic, a stand-alone probe must be attached to a nonsecure port. Implementing RMON in the switch means that all ports can have security features enabled.

To view the status of RMON polling on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for RMON polling.

To view the RMON memory usage statistics for a specific memory type (for example, statistics, events, logs, history, or alarms) or for all memory types, use the following command:

```
show rmon memory {detail | memoryType}
```

Example

The following command enables the collection of RMON statistics on the switch:

```
enable rmon
```



History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

enable sflow

enable sflow

Description

Globally enables sFlow statistical packet sampling.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables sFlow globally on the switch.



Note

On BlackDiamond 8000 c-, e-, xl-, and xm-series modules, and Summit X460, X480, X670 and X770 series switches (whether or not included in a SummitStack), sFlow and mirroring are not mutually exclusive. You can enable sFlow and mirroring at the same time.

Any traffic grouping using QP2 may encounter unexpected results when sFlow is enabled. For more information about QoS, see QoS in the *ExtremeXOS Concepts Guide*.

Example

The following command enables sFlow sampling globally:

```
enable sflow
```

History

This command was first available in ExtremeXOS 11.0.



Platform Availability

This command is available on all platforms.

enable sflow ports

```
enable sflow ports port_list {ingress | egress | both }
```

Description

Enables sFlow statistical packet sampling on a particular list of ports.

Syntax Description

<i>port_list</i>	Specifies a list of ports.
ingress	Enables ingress sFlow on a per-port basis.
egress	Enables egress sFlow on a per-port basis.
both	Enables both ingress and egress sFlow on a per-port basis.

Default

Ingress.

Usage Guidelines

This command enables sFlow on a particular list of ports. Ingress, egress, or a combination of both types of sampling can be enabled on a port. You also need to enable sFlow globally in order to gather statistics and send the data to the collector. Once sFlow is enabled globally, and on the ports of interest, sampling and polling begins.

Use the following command to enable sFlow globally:

```
enable sflow
```

Note



On BlackDiamond 8000 c-, e-, xl-, and xm-series modules, and Summit family switches (whether or not included in a SummitStack), sFlow and mirroring are not mutually exclusive. You can enable sFlow and mirroring at the same time.

For more information about mirroring, see [Configuring Slots and Ports on a Switch](#) in the *ExtremeXOS Concepts Guide*.



Example

The following command enables egress sFlow sampling on the port 3:1:

```
enable sflow ports 3:1 egress
```

History

This command was first available in ExtremeXOS 11.0.

The **ingress** keyword was added in ExtremeXOS 15.3.

The **egress** and **both** keywords were added in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

enable sys-health-check

```
enable sys-health-check slot slot
```

Description

Enables backplane diagnostic packets on the specified slot.

Syntax Description

<i>slot</i>	Specifies the slot to participate in sending backplane diagnostic packets.
-------------	--

Default

Polling is enabled, backplane diagnostic packets are disabled.

Depending upon your platform, when enabling diagnostic packets, the following defaults apply:

- BlackDiamond 8800 series switches—The system health checker tests the data link every 5 seconds for the specified slot.

Usage Guidelines

Configure the system health checker with guidance from Extreme Networks Technical Support personnel.

The system health checker tests I/O modules and the backplane by sending diagnostic packets. By isolating faults to a specific module or backplane connection, the system health checker notifies you of a possible hardware failure.



System health check errors are reported to the syslog. Syslog output includes the slot number where the problem occurred, the loopback packet ID number, and a notification that the MSM/MM did not receive the last packet. If you see an error, please contact Extreme Networks Technical Support.



Note

Enabling backplane diagnostic packets increases CPU utilization and competes with network traffic for resources.

The system health checker continues to periodically forward test packets to failed components.

To configure the frequency of the backplane diagnostic packets on the BlackDiamond 8800 series switches, use the `configure sys-health-check interval` command.

Displaying the System Health Check Setting

To display the system health check polling setting on the switch, use the following command:

```
show switch
```

As previously described, polling is always enabled on the switch, which is why you see the system health check setting as Enabled. The following truncated output from a BlackDiamond 8810 switch displays the system health check setting (displayed as SysHealth check):

```
SysName:          TechPubs Lab
SysName:          BD-8810Rack3
SysLocation:
SysContact:       support@extremenetworks.com, +1 888 257 3000
System MAC:       00:04:96:1F:A2:60
SysHealth check:  Enabled
Recovery Mode:    None
System Watchdog:  Enabled
```

Example

The following command enables backplane diagnostic packets on slot 6:

```
enable sys-health-check slot 6
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available only on modular switches.



enable syslog

enable syslog

Description

Enables logging to all remote syslog host targets.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

To enable remote logging, you must do the following:

- Configure the syslog host to accept and log messages.
- Enable remote logging by using the `enable syslog` command.
- Configure remote logging by using the `configure syslog` command.

When you use the `enable syslog` command, the exporting process of the syslog begins. This command also determines the initial state of an added remote syslog target.

Example

The following command enables logging to all remote syslog hosts:

```
enable syslog
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable/disable xml-notification

```
[enable|disable] xml-notification [all | target]
```



Description

Enables or disables Web server target(s).

Syntax Description

<i>target</i>	wSpecifies the configured target.
---------------	-----------------------------------

Default

By default, the target Web server is not enabled for xml-notifications. You have to explicitly enable it, and the display value is “no.”

Usage Guidelines

Use the enable option to enable Web server target(s) in order to receive events from ExtremeXOS modules and to send out events to the targeted Web server(s).

Use the disable option to disable the Web server target(s).

Example

The following command enables all of the configured targets:

```
enable xml-notification all
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on BlackDiamond X8, BlackDiamond 8000 series modules, and Summit family switches.

show configuration “xmlc”

```
show configuration "xmlc" {detail | non-persistent {detail}}
```

Description

Displays the configuration of an XMLC module.

Syntax Description

This command has no arguments or variables.



Default

N/A.

Usage Guidelines

Displays the configuration of an XMLC module.

When the detail option is chosen, all configuration data including the default is displayed. Otherwise the default would not be displayed.

When the non-persistent option is chosen, UPM non-persistent configuration data is displayed.

Example

The following command displays the xmlc configuration:

```
show configuration "xmlc" detail
```

Following is sample output from this command:

```
Module xmlc configuration.
#
create xml-notification target test url http://10.255.42.73:9080/axis/
services/eventPort
configure xml-notification test queue-size 100
disable xml-notification test
create xml-notification target test2 url http://10.255.42.48:9080/axis/
services/eventPort
configure xml-notification test2 queue-size 100
enable xml-notification test2
create xml-notification target epicenter-target url http://
10.255.42.48:8080/xos/webservice
configure xml-notification target epicenter-target user admin encrypted-auth
YWRtaW46ZXBpY2VudGVy
configure xml-notification epicenter-target queue-size 100
enable xml-notification epicenter-target
configure xml-notification target test add idMgr
configure xml-notification target test2 add idMgr
configure xml-notification target epicenter-target add idMgr
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on BlackDiamond X8, BlackDiamond 8000 series modules, and Summit family switches.




show cpu-monitoring

```
show cpu-monitoring {process name} {slot slotid}
```

Description

Displays the CPU utilization history of one or more processes.

Syntax Description

<i>name</i>	Specifies the name of the process.
<i>slotid</i>	Specifies the slot number of the MSM/MM module: "A" specifies the MSM installed in slot A. "B" specifies the MSM installed in slot B.
	 Note This parameter is available only on modular switches.

Default

N/A.

Usage Guidelines

Viewing statistics on a regular basis allows you to see how well your network is performing. If you keep simple daily records, you will see trends emerging and notice problems arising before they cause major network faults. This way, statistics can help you get the best out of your network.

By default, CPU monitoring is enabled and occurs every 5 seconds. The default CPU threshold value is 90%.

This information may be useful for your technical support representative if you experience a problem.

Depending on the software version running on your switch or your switch model, additional or different CPU and process information might be displayed.

SummitStack Only

When you issue the command with out any parameters:

- From the stack manager or backup node, the stack displays CPU utilization history for all the processes running on the master node and the backup node in the Active Topology.
- From the stack manager or a standby node, the stack displays CPU utilization history for all the processes running on the master node and the standby node in the Active Topology.



Modular Switches Only

When you issue the command without any parameters, the switch displays CPU utilization history for all of the processes running on the MSMs/MMs installed in your system.

Reading the Output

The `show cpu-monitoring` command is helpful for understanding the behavior of a process over an extended period of time. The following information appears in a tabular format:

- **Card**—The location (MSM A or MSM B) where the process is running on a modular switch.
- **Process**—The name of the process.
- **Range of time (5 seconds, 10 seconds, and so forth)**—The CPU utilization history of the process or the system. The CPU utilization history goes back only 1 hour.
- **Total User/System CPU Usage**—The amount of time recorded in seconds that the process spends occupying CPU resources. The values are cumulative meaning that the values are displayed as long as the system is running. You can use this information for debugging purposes to see where the process spends the most amount of time: user context or system context.

Example

The following command displays CPU utilization on the switch:

```
show cpu-monitoring
```

The following is sample truncated output from a modular switch:

```

CPU Utilization Statistics - Monitored every 5 seconds
-----
-
Card   Process           5   10   30   1   5   30   1   Max   Total
secs  secs  secs min  mins mins hour      User/System
util  util  util util util util util util  CPU Usage
(%)  (%)  (%) (%) (%) (%) (%) (%)  (secs)
-----
-
MSM-A  System             0.0  0.0  0.1  0.0  0.0  0.0  0.0  0.9
MSM-B  System             0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
MSM-A  GNSS_cpuif         0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
MSM-A  GNSS_ctrlif        0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
MSM-A  GNSS_esmi          0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
MSM-A  GNSS_fabric        0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
MSM-A  GNSS_mac_10g       0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
MSM-A  GNSS_pbusmux       0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
MSM-A  GNSS_pktengine     0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
MSM-A  GNSS_pktif         0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
MSM-A  GNSS_switch        0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
MSM-A  aaa                 0.0  0.0  0.0  0.0  0.0  0.0  0.0  8.4  0.82  0.56
MSM-A  acl                 0.0  0.0  0.0  0.0  0.0  0.0  0.0  7.5  0.37  0.33
MSM-A  bgp                 0.0  0.0  0.0  0.0  0.0  0.0  0.0  5.2  0.27  0.42
MSM-A  cfgmgr              0.0  0.9  0.3  3.7  1.2  1.2  1.3  27.3  7.70  7.84
MSM-A  cli                 0.0  0.0  0.0  48.3  9.6  2.5  2.1  48.3  0.51  0.37
    
```



MSM-A	devmgr	0.0	0.0	0.0	0.9	0.3	0.2	0.2	17.1	2.22	2.50
MSM-A	dirser	0.0	0.0	0.0	0.0	0.0	0.0	0.0	9.5	0.0	0.0
MSM-A	dosprotect	0.0	0.0	0.0	0.0	0.0	0.0	0.0	3.8	0.20	0.26
MSM-A	eaps	1.9	0.9	0.4	0.0	0.0	0.0	0.0	8.4	2.40	1.40
MSM-A	edp	0.0	0.0	0.0	0.0	0.0	0.0	0.0	10.2	0.99	0.47
MSM-A	elrp	0.0	0.0	0.0	0.0	0.0	0.0	0.0	8.4	0.44	0.28
MSM-A	ems	0.0	0.0	0.0	0.0	0.0	0.0	0.0	12.2	1.1	1.16
MSM-A	epm	0.0	0.0	0.0	0.9	0.1	0.2	0.2	4.7	2.6	4.18
MSM-A	esrp	0.0	0.0	0.0	0.0	0.0	0.0	0.0	7.5	0.44	0.36
MSM-A	etmon	0.9	0.4	0.6	1.2	1.1	1.0	1.0	23.3	21.84	7.24
...											

The following is sample output from a Summit switch:

```
CPU Utilization Statistics - Monitored every 25 seconds
```

Process	5 secs	10 secs	30 secs	1 hour	5 mins	30 mins	1 hour	Max	Total	
secs	secs	secs	min	mins	mins	hour	User/System	CPU Usage		
util	util	util	util	util	util	util	util	(secs)		
(%)	(%)	(%)	(%)	(%)	(%)	(%)	(%)			
System	n/a	n/a	0.0	0.9	0.1	0.2	0.5	34.6		
aaa	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.8	1.72	0.78
acl	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.40	0.24
bgp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	12.6	11.18	2.21
cfgmgr	n/a	n/a	0.0	0.0	0.0	0.0	0.8	39.8	4743.92	3575.79
cli	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.59	0.42
devmgr	n/a	n/a	0.0	0.0	0.0	0.0	0.0	19.5	74.44	24.52
dirser	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
dosprotect	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.8	0.12
eaps	n/a	n/a	0.0	0.0	0.0	0.0	0.1	5.5	36.40	15.41
edp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	11.1	10.92	3.97
elrp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.49	0.44
ems	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	1.19	1.29
epm	n/a	n/a	0.0	0.0	0.0	0.0	0.0	30.7	48.74	32.93
esrp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	2.7	0.82	0.45
etmon	n/a	n/a	0.0	0.0	0.0	0.0	0.5	30.5	4865.78	873.87
...										

The following is sample truncated output from a stack:

```
Slot-1 stack.3 # sh cpu-monitoring
```

```
CPU Utilization Statistics - Monitored every 20 seconds
```

Card	Process	5 secs	10 secs	30 secs	1 hour	5 mins	30 mins	1 hour	Max	Total	
secs	secs	secs	min	mins	mins	hour	User/System	CPU Usage			
util	util	util	util	util	util	util	util	(secs)			
(%)	(%)	(%)	(%)	(%)	(%)	(%)	(%)				
Slot-1	System	n/a	n/a	0.0	1.6	0.8	0.5	0.5	2.5		
Slot-6	System	n/a	n/a	0.3	0.9	0.7	0.4	0.5	4.6		
Slot-1	aaa	n/a	n/a	0.0	0.0	0.0	0.0	0.0	3.6	1.22	0.75
Slot-1	acl	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.8	1.8	0.52



Slot-1 bgp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Slot-1 brm	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.8	0.53	0.17	
Slot-1 cfgmgr	n/a	n/a	0.1	0.0	0.0	0.0	0.0	0.8	3.18	0.65	
Slot-1 cli	n/a	n/a	0.9	0.8	0.1	0.0	0.3	97.2	13.7	2.12	
Slot-1 devmgr	n/a	n/a	0.0	0.0	0.0	0.0	0.0	5.0	1.1	1.24	
Slot-1 dirser	n/a	n/a	0.0	0.0	0.0	0.0	0.0	5.9	0.0	0.0	
Slot-1 dosprotect	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.18	0.12	
Slot-1 eaps	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.3	0.92	0.45	
Slot-1 edp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.3	0.68	0.20	
Slot-1 elrp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.9	0.49	0.21	
Slot-1 elsm	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.3	0.38	0.34	
Slot-1 ems	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.3	1.0	0.41	
Slot-1 epm	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.3	1.63	1.28	
Slot-1 esrp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.9	0.50	0.21	
Slot-1 etmon	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.8	4.0	0.65	
...											
...											
Slot-1 stp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.9	0.67	0.27	
Slot-1 telnetd	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.5	0.23	0.6	
Slot-1 tftpd	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.19	0.10	
Slot-1 thttpd	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.8	0.21	0.13	
Slot-1 upm	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.8	0.43	0.22	
Slot-1 vlan	n/a	n/a	0.0	0.0	0.0	0.0	0.1	4.3	4.28	1.56	
Slot-1 vrrp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.8	0.38	0.13	
Slot-1 xmld	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.48	0.25	
Slot-6 aaa	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.42	0.26	
Slot-6 acl	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.40	0.26	
Slot-6 bgp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	
Slot-6 brm	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.18	0.7	
Slot-6 cfgmgr	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.81	0.28	
Slot-6 cli	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.9	7.17	1.2	
Slot-6 devmgr	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.35	0.88	
Slot-6 dirser	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	
Slot-6 dosprotect	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.6	0.2	
Slot-6 eaps	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.60	0.20	
Slot-6 edp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.23	0.11	
Slot-6 elrp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.9	0.20	0.4	
Slot-6 elsm	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.21	0.9	
Slot-6 ems	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.44	0.22	
Slot-6 epm	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	1.78	1.29	
Slot-6 esrp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.24	0.8	
Slot-6 etmon	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.8	1.11	0.28	
...											

History

This command was first available in an ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.



show elsm

show elsm

Description

Displays summary information for all of the ELSM-enabled ports on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to display the operational state of ELSM on the configured ports.

If no ports are configured for ELSM, the switch does not display any information.

For ELSM-enabled ports, this command displays the following information in a tabular format:

- Port—The port number of the ELSM-enabled port.
- ELSM State—The current state of ELSM on the port. The ELSM state can be one of the following:
 - Up—Indicates a healthy remote system and this port is receiving Hello+ messages from its peer.

If an ELSM-enabled port enters the Up state, the up timer begins. Each time the port receives a Hello+ message from its peer, the up timer restarts and the port remains in the Up state.

- Down—Indicates that the port is down, blocked, or has not received Hello+ messages from its peer.

If an ELSM-enabled port does not receive a hello message from its peer before the up timer expires, the port transitions to the Down state. When ELSM is down, data packets are neither forwarded nor transmitted out of that port.

- Down-Wait—Indicates a transitional state.

If the port enters the Down state and later receives a Hello+ message from its peer, the port enters the Down-Wait state. If the number of Hello+ messages received is greater than or equal to the hello threshold (by default 2 messages), the port transitions to the Up state. If the number of Hello+ messages received is less than the hold threshold, the port enters the Down state.

- Down-Stuck—Indicates that the port is down and requires user intervention.

If the port repeatedly flaps between the Up and Down states, the port enters the Down-Stuck state. Depending on your configuration, there are two ways for a port to transition out of this state:

By default, automatic restart is enabled, and the port automatically transitions out of this state. See the command `enable elsm ports auto-restart` for more information.



If you disabled automatic restart, and the port enters the Down-Stuck state, you can clear the stuck state and enter the Down state by using one of the following commands:

```
clear elsm ports port_list auto-restart
```

```
enable elsm ports port_list auto-restart
```

- Hello time —The current value of the hello timer, which by default is 1 second. The hello timer indicates the number of seconds between consecutive hello messages.

Additional show Command

You can also use the `show ports {port_list} information {detail}` command to display ELSM information.

If you do not specify the detail parameter, the following columns display the current state of ELSM on the switch:

- Flags:
 - L—Indicates that ELSM is enabled on the switch.
 - - —Indicates that ELSM is disabled on the switch.
- ELSM:
 - up—Indicates that ELSM is enabled and the ELSM peer ports are up and communicating; the ELSM link state is up. In the up state, the ELSM-enabled port sends and receives hello messages from its peer.
 - dn—Indicates that ELSM is enabled, but the ELSM peers are not communicating; the ELSM link state is down. In the down state, ELSM transitions the peer port on this device to the down state. ELSM blocks all incoming and outgoing switching traffic and all control traffic except ELSM PDUs.
 - - —Indicates that ELSM is disabled on the switch.

If you specify the optional detail parameter, the following ELSM output is called out in written explanations versus displayed in a tabular format:

- ELSM Link State (displayed only if ELSM is enabled on the switch):
 - Up—Indicates that ELSM is enabled and the ELSM peer ports are up and communicating; the ELSM link state is up. In the up state, the ELSM-enabled port sends and receives hello messages from its peer.
 - Down—Indicates that ELSM is enabled, but the ELSM peers are not communicating; the ELSM link state is down. In the down state, ELSM transitions the peer port on this device to the down state. ELSM blocks all incoming and outgoing switching traffic and all control traffic except ELSM PDUs.
- ELSM:
 - Enabled—Indicates that ELSM is enabled on the switch.
 - Disabled—Indicates that ELSM is disabled on the switch.



Example

The following command displays summary configuration information for all of the ELSM-enabled ports on the switch:

```
show elsm
```

The following is sample output from this command:

```
Port      ELSM State   Hello Time
====      =
5:14     Up 1 (second)
5:18     Down         1 (second)
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

show elsm ports

```
show elsm ports all | port_list
```

Description

Displays detailed information for one or more ELSM-enabled ports.

Syntax Description

all	Displays detailed ELSM information for all ports.
<i>port_list</i>	Displays detailed ELSM information for one or more ports.

Default

N/A.

Usage Guidelines

Use this command to display detailed information about the operational state of ELSM on the configured ports.

This command displays in a tabular format the following ELSM data for one or more ELSM-enabled ports on the switch:



- Port—The port number of the ELSM-enabled port.
- Link State—The state of the link between ELSM-enabled (peer) ports. The link state can be one of the following:
 - Ready—Indicates that the port is enabled but there is no link.
 - Active—Indicates that the port is enabled and the physical link is up.
- ELSM Link State—The current state of the ELSM logical link on the switch. The ELSM link state can be one of the following:
 - Up—Indicates that ELSM is enabled and the ELSM peer ports are up and communicating; the ELSM link state is up. In the up state, the ELSM-enabled port sends and receives hello messages from its peer.
 - Down—Indicates that ELSM is enabled, but the ELSM peers are not communicating; the ELSM link state is down. In the down state, ELSM transitions the peer port on this device to the down state. ELSM blocks all incoming and outgoing switching traffic and all control traffic except ELSM PDUs.
- ELSM State—The current state of ELSM on the port. The ELSM state can be one of the following:
 - Up—Indicates a healthy remote system and this port is receiving Hello+ messages from its peer.

If an ELSM-enabled port enters the Up state, the up timer begins. Each time the port receives a Hello+ message from its peer, the up timer restarts and the port remains in the Up state. The up timer is 6 * hello timer, which by default is 6 seconds.

- Down—Indicates that the port is down, blocked, or has not received Hello+ messages from its peer.

If an ELSM-enabled port does not receive a hello message from its peer before the up timer expires, the port transitions to the Down state. When ELSM is down, data packets are neither forwarded nor transmitted out of that port.

- Down-Wait—Indicates a transitional state.

If the port enters the Down state and later receives a Hello+ message from its peer, the port enters the Down-Wait state. If the number of Hello+ messages received is greater than or equal to the hold threshold, the port transitions to the Up state. If the number of Hello+ messages received is less than the hold threshold, the port enters the Down state.

- Down-Stuck—Indicates that the port is down and requires user intervention.

If the port repeatedly flaps between the Up and Down states, the port enters the Down-Stuck state. Depending on your configuration, there are two ways for a port to transition out of this state:

By default, automatic restart is enabled, and the port automatically transitions out of this state. See the command `enable elsm ports auto-restart` for more information.

If you disabled automatic restart, and the port enters the Down-Stuck state, you can clear the stuck state and enter the Down state by using one of the following commands:

```
clear elsm ports port_list auto-restart
```

```
enable elsm ports port_list auto-restart
```

- Hello Transmit State—The current state of ELSM hello messages being transmitted. The transmit state can be one of the following:



- HelloRx(+)—Specifies that the ELSM-enabled port is up and receiving Hello+ messages from its peer. The port remains in the HelloRx+ state and restarts the HelloRx timer each time it receives a Hello+ message. If the HelloRx timer expires, the hello transmit state enters HelloRX(-). The HelloRx timer is $6 * \text{hello timer}$, which by default is 6 seconds.
- HelloRx(-)—Specifies that the ELSM-enabled port either transitions from the initial ELSM state or is up and not receiving hello messages because there is a problem with the link or the peer is missing.
- Hello time—The current value of the hello timer, which by default is 1 second. The hello timer indicates the number of seconds between consecutive hello messages.
- Hold Threshold—The number of Hello+ messages required by the ELSM-enabled port to transition from the Down-Wait state to the Up state within the hold threshold.
- UpTimer Threshold—The number of hello times that span without receiving Hello+ packets before a port changes its ELSM state from Up to Down.
- Auto Restart—The current state of ELSM automatic restart on the port. The state of Auto Restart can be one of the following:
 - Enabled—If an ELSM-enabled port goes down, ELSM automatically brings up the down port. This is the default behavior.
 - Disabled If an ELSM-enabled port goes down, the port enters and remains in the Down-Stuck state until you clear the stuck state.

For more information about automatic restart, see the command `enable elsm ports auto-restart`.

- Sticky Threshold—Specifies the number of times a port can transition between the Up and Down states. The sticky threshold is not user-configurable and has a default value of 1. That means a port can transition only one time from the Up state to the Down state. If the port attempts a subsequent transition from the Up state to the Down state, the port enters the Down-Stuck state.
- Sticky Threshold Counter—The number of times the port transitions from the Up state to the Down state.
- Down Timeout—The actual waiting time (msecs or secs) before a port changes its ELSM state from Down to Up. When ELSM is enabled on a port and it is in a Down state, before it changes its ELSM state from Down to Up, it expects to receive at least a “Hold Threshold” number of Hello+ packets during the Down Timeout period after it receives the first Hello+ packet from its peer. It is equal to $[\text{Hello Time} * (\text{Hold Threshold} + 2)]$.
- Up Timeout—The actual waiting time (msecs or secs) before a port changes its ELSM state from Up to Down after receiving the last Hello+ packets. When a port is in an Up state, it expects to receive a Hello+ packet from its peer every “Hello Time” period to maintain its Up state. When it does not receive a Hello+ packet after an “Up Timeout” period, it changes its ELSM state from Up to Down. It is equal to $[\text{Hello Time} * \text{UpTimer Threshold}]$.

The remaining output displays counter information. Use the counter information to determine the health of the ELSM peers and how often ELSM has gone up or down. The counters are cumulative.

- RX Hello+—The number of Hello+ messages received by the port.
- Rx Hello- —The number of Hello- messages received by the port.
- Tx Hello+—The number of Hello+ messages sent by the port.
- Tx Hello- —The number of Hello- messages sent by the port.
- ELSM Up/Down Count—The number of times ELSM has been up or down.



To clear, reset the counters, use either the `clear elsm {ports port_list} counters` or the `clear counters` command.

Additional show Command

You can also use the `show ports {port_list} information {detail}` command to display ELSM information.

If you do not specify the detail parameter, the following columns display the current state of ELSM on the switch:

- Flags:
 - L—Indicates that ELSM is enabled on the switch.
 - - —Indicates that ELSM is disabled on the switch.
- ELSM:
 - up—Indicates that ELSM is enabled and the ELSM peer ports are up and communicating; the ELSM link state is up. In the up state, the ELSM-enabled port sends and receives hello messages from its peer.
 - dn—Indicates that ELSM is enabled, but the ELSM peers are not communicating; the ELSM link state is down. In the down state, ELSM transitions the peer port on this device to the down state. ELSM blocks all incoming and outgoing switching traffic and all control traffic except ELSM PDUs.
 - - —Indicates that ELSM is disabled on the switch.

If you specify the optional detail parameter, the following ELSM output is called out in written explanations versus displayed in a tabular format:

- ELSM Link State (displayed only if ELSM is enabled on the switch).
 - Up—Indicates that ELSM is enabled and the ELSM peer ports are up and communicating; the ELSM link state is up. In the up state, the ELSM-enabled port sends and receives hello messages from its peer.
 - Down—Indicates that ELSM is enabled, but the ELSM peers are not communicating; the ELSM link state is down. In the down state, ELSM transitions the peer port on this device to the down state. ELSM blocks all incoming and outgoing switching traffic and all control traffic except ELSM PDUs.
- ELSM:
 - Enabled—Indicates that ELSM is enabled on the switch.
 - Disabled—Indicates that ELSM is disabled on the switch.

Example

The following command displays detailed ELSM information for all configured ports on the switch:

```
show elsm ports all
```

The following is sample output from this command:

```
ELSM Info Port 4:4
```



```

Link State           : Active
ELSM Link State     : Up
ELSM State          : Up
Hello Transmit State : HelloRx(+)
Hello Time          : 100 msec
Hold Threshold      : 2
UpTimer Threshold   : 6
Auto Restart        : Enabled
Down Timeout        : 400 msec
Up Timeout          : 600 msec
Rx Hello+           : 667960
Rx Hello-           : 0
Tx Hello+           : 667958
Tx Hello-           : 0
ELSM Up/Down Count  : UP: 0    DOWN: 0
ELSM Info Port 4:4
Link State           : Active
ELSM Link State     : Up
ELSM State          : Up
Hello Transmit State : HelloRx(+)
Hello Time          : 100 msec
Hold Threshold      : 2
UpTimer Threshold   : 6
Auto Restart        : Disabled
Sticky Threshold    : 1
Sticky Threshold Counter : 0
Down Timeout        : 400 msec
Up Timeout          : 600 msec
Rx Hello+           : 708204
Rx Hello-           : 0
Tx Hello+           : 708201
Tx Hello-           : 0
ELSM Up/Down Count  : UP: 0    DOWN: 0

```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

show fans

```
show fans {detail}
```

Description

Displays the status of the fans in the system.



Syntax Description

detail	The detail option is reserved for future use.
---------------	---

Default

N/A.

Usage Guidelines

Use this command to view detailed information about the health of the fans.

This status information may be useful for your technical support representative if you have a network problem.

The switch collects and displays the following fan information:

- **State**—The current state of the fan. Options are:
 - Empty: There is no fan installed.
 - Failed: The fan failed.
 - Operational: The fan is installed and working normally.
- **NumFan**—The number of fans in the fan tray.
- **Fan Name**, displayed as Fan-1, Fan-2, and so on (modular switches also include a description of the location, for example, Upper or Upper-Right)—Specifies the individual state for each fan in a fan tray and its current speed in revolutions per minute (rpm).

On modular switches, the output also includes the following information:

- **PartInfo**—Information about the fan tray, including the:
 - **Serial number**—A collection of numbers and letters, that make up the serial number of the fan. This is the first series of numbers and letters in the display.
 - **Part number**—A collection of numbers and letters, that make up the part number of the fan. This is the second series of numbers and letters in the display.
- **Revision**—The revision number of the fan.
- **Odometer**—Specifies the power-on date and how long the fan tray has been operating since it was first powered-on.

Example

The following command displays the status of the installed fans. If a fan is not installed, the state of the fan is Empty.

```
show fans
```

The following is sample output from a BlackDiamond 8800 series switch:

```
FanTray information:
State:                Operational
NumFan:              9
```



```

PartInfo:          0404X-00015 450102-00-01
Revision:          1.0
Odometer:          111 days 16 hours 30 minutes since Oct-13-2004
Upper-Left Fan-1:  Operational at 2880 RPM
Middle-Left Fan-2: Operational at 2820 RPM
Lower-Left Fan-3:  Operational at 2820 RPM
Upper-Center Fan-4: Operational at 2820 RPM
Center Fan-5:      Operational at 2820 RPM
Lower-Center Fan-6: Operational at 2880 RPM
Upper-Right Fan-7: Operational at 2880 RPM
Middle-Right Fan-8: Operational at 2820 RPM
Lower-Right Fan-9: Operational at 2880 RPM

```

The following is sample output from a BlackDiamond X8 switch:

```

BD-X8.8 # show fans
FanTray-1 information:
State:          Operational
NumFan:         6
PartInfo:       1135G-02213 450350-00-01
Revision:       1.0
Odometer:       19 hours since Mar-01-2010
Top Fan-1:      Operational at 3300 RPM
Fan-2:          Operational at 3360 RPM
Fan-3:          Operational at 3240 RPM
Fan-4:          Operational at 3240 RPM
Fan-5:          Operational at 3240 RPM
Bottom Fan-6:   Operational at 3240 RPM
FanTray-2 information:
State:          Operational
NumFan:         6
PartInfo:       1135G-02244 450350-00-01
Revision:       1.0
Odometer:       19 hours since Mar-01-2010
Top Fan-1:      Operational at 3360 RPM
Fan-2:          Operational at 3300 RPM
Fan-3:          Operational at 3360 RPM
Fan-4:          Operational at 3300 RPM
Fan-5:          Operational at 3300 RPM
Bottom Fan-6:   Operational at 3360 RPM
FanTray-3 information:
State:          Operational
NumFan:         6
PartInfo:       1135G-02306 450350-00-01
Revision:       1.0
Odometer:       41 days 19 hours 30 minutes since Oct-14-2011
Top Fan-1:      Operational at 3240 RPM
Fan-2:          Operational at 3240 RPM
Fan-3:          Operational at 3360 RPM
Fan-4:          Operational at 3360 RPM
Fan-5:          Operational at 3360 RPM
Bottom Fan-6:   Operational at 3360 RPM
FanTray-4 information:
State:          Operational
NumFan:         6
PartInfo:       1135G-02275 450350-00-01
Revision:       1.0

```




```

Odometer:                19 hours  since Mar-01-2010
Top      Fan-1:           Operational at 3300 RPM
Fan-2:                Operational at 3300 RPM
Fan-3:                Operational at 3300 RPM
Fan-4:                Operational at 3240 RPM
Fan-5:                Operational at 3300 RPM
Bottom   Fan-6:           Operational at 3300 RPM
FanTray-5 information:
State:                Operational
NumFan:                6
PartInfo:             1135G-02337 450350-00-01
Revision:             1.0
Odometer:             41 days 18 hours  since Oct-14-2011
Top      Fan-1:           Operational at 3300 RPM
Fan-2:                Operational at 3360 RPM
Fan-3:                Operational at 3240 RPM
Fan-4:                Operational at 3300 RPM
Fan-5:                Operational at 3240 RPM
Bottom   Fan-6:           Operational at 3240 RPM

```

The following is sample output from a Summit switch:

```

FanTray information:
State:                Operational
NumFan:                6
PartInfo:             0931G-00064 450237-00-05
Revision:             5.0
Fan-1:                Operational at 14894 RPM
Fan-2:                Operational at 15360 RPM
Fan-3:                Operational at 15360 RPM
Fan-4:                Operational at 9637 RPM
Fan-5:                Operational at 9637 RPM
Fan-6:                Operational at 9637 RPM

```

The following is a sample output from a SummitStack:

```

FanTray-1 information:
State:                Operational
NumFan:                1
Fan-1:                Operational at 1000 RPM
FanTray-2 information:
State:                Operational
NumFan:                1
Fan-1:                Operational at 1000 RPM
FanTray-3 information:
State:                Operational
NumFan:                1
Fan-1:                Operational at 1000 RPM
FanTray-4 information:
State:                Operational
NumFan:                1
Fan-1:                Operational at 1000 RPM
FanTray-5 information:
State:                Empty
FanTray-6 information:

```



```

State:                               Empty
FanTray-7 information:
State:                               Empty
FanTray-8 information:
State:                               Empty

```

History

This command was first available in an ExtremeXOS 10.1.

Information about the location of the fan tray for the BlackDiamond 8810 switch (upper-left, middle left, lower-left, upper- lower-center, upper-right, middle-right, and lower-right) was added to the `show fans` output in ExtremeXOS 11.1.

Information about the current speed in rpm for the Summit family switches was added to the `show fans` output in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

show log

```

show log {messages [memory-buffer | nvram]} {events {event-condition | event-
component}} {severity severity {only}} {starting [date date time time | date date
| time time]} {ending [date date time time | date date | time time]} {match
regex} {chronological}

```

Description

Displays the current log messages.

Syntax Description

messages	Specifies the target location from which to display the log messages.
memory-buffer	Show messages stored in volatile memory (default).
nvram	Show messages stored in NVRAM.
events	Show event messages.
<i>event-condition</i>	Specifies the event condition to display.
<i>event-component</i>	Specifies the event component to display.
<i>severity</i>	Specifies the minimum severity level to display (if the keyword only is omitted).
only	Specifies that only the specified severity level is to be displayed
starting	Show messages with timestamps equal to or greater than that specified



<i>date</i>	Specifies the date, where date is <i>month</i> (1-12) / <i>day</i> (1-31) {/ <i>year</i> (yyyy)}.
<i>time</i>	Specifies the time, where time is <i>hour</i> (0-23) {: <i>minute</i> (0-59) {: <i>seconds</i> (0-59) {. <i>hundredths</i> }}.
ending	Show messages with timestamps equal to or less than that specified.
<i>regex</i>	Specifies a regular expression. Only messages that match the regular expression will be displayed.
chronological	Specifies displaying log messages in ascending chronological order (oldest to newest).

Default

The following defaults apply:

- messages—memory buffer.
- event—no restriction (displays user-specified event).
- severity—none (displays everything stored in the target).
- starting, ending—if not specified, no timestamp restriction.
- match—no restriction.
- chronological—if not specified, show messages in order from newest to oldest.

Usage Guidelines

Switch configuration and fault information is filtered and saved to target logs, in a memory buffer, and in NVRAM. Each entry in the log contains the following information:

- Timestamp—records the month and day of the event, along with the time (hours, minutes, seconds, and hundredths).
- Severity Level—indicates the urgency of a condition reported in the log. [Table 21: Severity Levels Assigned by the Switch](#) on page 1038 Table describes the severity levels assigned to events.
- Component, Subcomponent, and Condition Name—describes the subsystem in the software that generates the event. This provides a good indication of where a fault might lie.
- Message—a description of the event occurrence. If the event was caused by a user, the user name is also provided.

This command displays the messages stored in either the internal memory buffer or in NVRAM. The messages shown can be limited by specifying a severity level, a time range, or a match expression. Messages stored in the target have already been filtered as events occurred, and specifying a severity or match expression on the show log command can only further limit the messages shown.

If the messages keyword is not present, the messages stored in the memory-buffer target are displayed. Otherwise, the messages stored in the specified target are displayed.

If the only keyword is present following the severity value, then only the events at that exact severity are included. Without the only keyword, events at that severity or more urgent are displayed. For example, severity warning implies critical, error, or warning, whereas severity warning only implies only warning.



Messages whose timestamps are equal or later than the starting time and are equal or earlier than the specified ending time will be shown if they also pass the severity requirements and match expression, if specified.

If a match phrase is specified, the formatted message must match the simple regular expression specified by match-expression for it to be shown.

A simple regular expression is a string of single characters including the dot character (.), which are optionally combined with quantifiers and constraints. A dot matches any single character while other characters match only themselves (case is significant). Quantifiers include the star character (*) that matches zero or more occurrences of the immediately preceding character or dot. Constraints include the caret character (^) that matches at the beginning of a message, and the currency character (\$) that matches at the end of a message. Bracket expressions are not supported. There are a number of sources available on the Internet and in various language references describing the operation of regular expressions.

If the chronological keyword is specified, messages are shown from oldest to newest; otherwise, messages are displayed newest to oldest.

Severity Level

The severity levels are critical, error, warning, notice, and info, plus three severity levels for extended debugging, debug-summary, debug-verbose, and debug-data. In log messages, the severity levels are shown by four letter abbreviations. The abbreviated forms are:

- Critical—Crit.
- Error—Erro.
- Warning—Warn.
- Notice—Noti.
- Info—Info.
- Debug-Summary—Summ.
- Debug-Verbose—Verb.
- Debug-Data—Data.

The three severity levels for extended debugging, debug-summary, debug-verbose, and debug-data, require that debug mode be enabled (which may cause a performance degradation). See the command `enable log debug-mode`. The following table describes the security levels:

Table 21: Severity Levels Assigned by the Switch

Level	Description
Critical	A serious problem has been detected that is compromising the operation of the system and that the system cannot function as expected unless the situation is remedied. The switch may need to be reset.
Error	A problem has been detected that is interfering with the normal operation of the system and that the system is not functioning as expected.
Warning	An abnormal condition, not interfering with the normal operation of the system, has been detected that may indicate that the system or the network in general may not be functioning as expected.



Table 21: Severity Levels Assigned by the Switch (continued)

Level	Description
Notice	A normal but significant condition has been detected, which signals that the system is functioning as expected.
Info (Informational)	A normal but potentially interesting condition has been detected, which signals that the system is functioning as expected and simply provides information or confirmation about the condition.
Debug-Summary	A condition has been detected that may interest a developer determining the reason underlying some system behavior.
Debug-Verbose	A condition has been detected that may interest a developer analyzing some system behavior at a more verbose level than provided by the debug summary information.
Debug-Data	A condition has been detected that may interest a developer inspecting the data underlying some system behavior.

Messages stored in NVRAM are in encoded format. To restore the ASCII text of a message, the version of ExtremeXOS loaded must be able to interpret the data written prior to reboot. When the encoded format for a particular message cannot be interpreted by the version of ExtremeXOS currently loaded, the messages are displayed in the following format:

```
03/21/2005 17:15:37.36 : NO MESSAGE DECODE; Missing component "epm" v24.2
DUMP-10: 00 14 C3 C1 00 11 00 1C 01 FF 00 08 65 70 6D 00
'.....epm.' DUMP-20: 08 FF 00 0C 00 18 00 02 65 70 6D 00
'.....epm.'
```

Log entries remain in the NVRAM log after a switch reboot. Issuing a clear log command does not remove these static entries. To remove log entries from NVRAM, use the following command:

```
clear log messages nvram
```

Example

The following command displays messages with a critical severity:

```
show log severity critical
```

The following command displays messages with warning, error, or critical severity:

```
show log severity warning
```

The following is sample output from a modular switch:

```
11/12/2004 00:38:10.30 <Warn:dm.Warn> MSM-A: Insufficient Power to power-on
Slot-7
11/12/2004 00:38:08.77 <Warn:dm.Warn> MSM-A: Slot-7 being Powered OFF due to
insuf
ficient power
11/12/2004 00:36:23.77 <Warn:dm.Warn> MSM-A: Slot-7 being Powered OFF due to
insuf
```



```

efficient power
...
A total of 83 log messages were displayed.

```

The following command displays messages containing the string "slot 2":

```
show log match "slot 2"
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show log components

```
show log components {event component } {version}
```

Description

Displays the name, description and default severity for all components.

Syntax Description

<i>event component</i>	Specifies the component to display.
version	Specifies the version number of the component.

Default

N/A.

Usage Guidelines

This command displays the name, description, and default severity defined for the specified components or subcomponents.

Depending on the software version running on your switch or your switch model, additional or different component information might be displayed.



Example

The following command displays the log components:

```
show log components
```

The following is sample output from this command:

Severity Component	Title	Threshold
-----	-----	-----
AAA	Authentication, Authorization, Accounting	Info
RADIUS	Remote Authentication Dial In User Service	Error
TACACS	Terminal Access Controller Access Control Syst	Info
ACL	ACL	Info
CLEARFlow	CLEARFlow	Info
Policy	Policy actions	Info
bgp	Border Gateway Protocol	Info
damp	BGP Route Flap Dampening related debug message	Error
event	BGP FSM related events	Error
inUpdt	Incoming Update related debug msgs	Warning
keepalive	BGP keepalive message	Warning
misc	Miscellaneous debug (Import, Aggregate, NextHop	Warning
msgs	Debug for BGP messages (OPEN, Update, Notifica	Warning
outUpdt	Transmit Update related debug	Warning
bootp	BOOTP, DHCP Component	Error
relay	BOOTP Relay trace component	Error
server	DHCP Server subcomponent	Info
cli	Command Line Interface	Info
shell	CLI configuration shell.	Error
subagent	CLI application subagent	Error
cm	Configuration Manager	Warning
file	CM file operation events	Warning
sys	CM system events	Warning
DM	Device Manager	Info
Card	Device Manager Card State Machine	Info
dosprot	dosprot	Info
ds	Directory Services	Error
EAPS	Ethernet Automatic Protection Switching	Info
SharedPort	EAPS SharedPort Domain	Info
EDP	Extreme Discovery Protocol (EDP)	Error
ELRP		
Report	Extreme Loop Recognition Protocol	Warning
EPM	Extreme Process Manager	Info
KLM	Kernel Loadable Module Manager	Notice
Msg	Message Handler	Info
Upgrade	Upgrade Manager	Info
Version	Version Manager	Critical
ESRP	Extreme Standby Router Protocol	Error
Aware	Subsystem description	Info
InPdu	Subsystem description	Info
Nbr	Subsystem description	Info
OutPdu	Subsystem description	Info
State	ESRP State Transitions	Warning
System	Subsystem description	Warning



Track	Subsystem description	Warning	
Vlan	Extreme Standby Router Protocol	Info	
fdb	fdb module event		Error
HAL	Hardware Abstraction Layer		Error
Card	Card State Driver	Info	
FDB	Forwarding Database Driver	Info	
IPv4ACL	IPv4 Access Control List Driver	Info	
IPv4Adj	IPv4 Adjacency Driver	Info	
IPv4FIB	IPv4 FIB Driver	Info	
IPv4Mc	IPv4 Multicast Driver	Info	
Mirror	Mirroring Driver	Error	
Msg	Message Handler	Info	
Port	I/O Port Driver	Info	
SM	Switch Manager	Info	
Sys	System Driver	Info	
VLAN	VLAN Driver	Info	
IPMC	IP Multicast Main Module		Info
Snoop	IP Multicast Snooping Module	Error	
VLAN	IP Multicast VLAN Module	Error	
ISIS	Intermediate-to-Intermediate		Error
Export	Route Redistribution into ISIS	Error	
IFSM	ISIS Interface Finite State Machine (IFSM)	Warning	
IIH	ISIS Hello (IIH) PDU	Warning	
LSP	ISIS Link State PDU	Notice	
NFSM	ISIS Neighbor Finite State Machine (NFSM)	Warning	
PDU	ISIS General PDU	Warning	
Restart	ISIS Restart	Notice	
SPF	ISIS Shortest Path First (SPF)	Warning	
VLAN	ISIS VLAN-Related Events	Error	
Kern	Kernel messages		Error
LACP	Link Aggregation Control Protocol	Info	
lldp	Link Layer Discovery Protocol (IEEE 802.1AB)	Warning	
log	Log server messages	Warning	
netTool	netTools framework		Error
dnsclient	Dns Client	Error	
dnsproxy	Dns Proxy	Error	
routeradv	IPv6 Router Advertisements	Warning	
sntp	Sntp client	Warning	
nl	Network Login		Info
dot1x	802.1x-based Network Login	Warning	
mac	MAC-based Network Login	Warning	
web	Web-based Network Login	Warning	
NM	Node Manager		Info
ospf	open shortest path first		Error
event	ospf events	Info	
hello	ospf hello	Error	
lsa	ospf link-state advertisement	Error	
neighbor	ospf neighbor	Error	
spf	ospf shortest path first	Error	
ospfv3	OSPFv3 related EMS messages		Warning
events	OSPF6 events related messages	Error	
lsa	LSA related messages	Warning	
nbr	OSPF6 neighbor related EMS messages	Warning	
pkt	OSPF6 Packet receive/transmit/processing relat	Warning	
route	OSPF6 route add/delete related messages	Warning	
spf	SPF computation related messages	Error	
pim	Pim Protocol Events		Warning
cache	PIM cache maintenance.	Warning	



debug	PIM debug messages	Notice	
hello	Hello messages	Warning	
mcdcb	multicast forwarding engine	Warning	
msg	Trace for pim control packtes	Notice	
nbr	Neighbor creation/deletion etc	Warning	
rpm	RP message exchange.	Warning	
pm	Policy Manager		Error
config	Policy file events	Info	
POE	Inline Power		Notice
rip	RIP routing		Error
cfg	rip configuration	Warning	
event	rip events	Warning	
inUpdt	rip - inbound route updates	Warning	
msgs	rip - socket messages in and out	Warning	
outUpdt	rip - outbound route updates	Warning	
sys	rip - exos kernel interface	Warning	
ripng	RIPng Protocol Events		Warning
debug	RIPng debug messages	Notice	
external	RIPng external interface related messages	Warning	
message	RIPng control messages	Warning	
route	Hello messages	Warning	
rmon	RMON general info		Error
alarm	RMON alarm info	Error	
estat	RMON statistics info	Error	
event	RMON event info	Error	
history	RMON history	Error	
RtMgr	Route Manager		Info
VLAN	rtmgr vlan interface	Info	
sflow	Sflow Protocol Events		Warning
debug	SFLOW debug messages	Notice	
extended	SFLOW extended data collection	Notice	
msg	SFLOW process initializaion related message	Warning	
sample	SFLOW sample collection related messages	Warning	
statistics	SFLOW port statistics related message	Warning	
STP	Spanning-Tree Protocol		Error
InBPDU	STP In Bridge Protocol Data Unit	Warning	
OutBPDU	STP Out Bridge Protocol Data Unit	Warning	
System	STP System	Error	
System	XOS system related log messages		Info
telnetd	telnet server		Info
tftpd	tftp server		Info
thttpd	thttp server		Info
trace	Debug trace messages		Warning
vlan	Vlan mgr		Info
ack	vlan ack	Error	
dbg	Debug information	Info	
err	errors	Error	
mac	Virtual MAC Debugging	Info	
msgs	Messages	Info	
VRRP	Config/State messages		Warning
Advert	Subsystem description	Warning	
System	System/Library messages	Warning	

A total of 143 component(s) were displayed.



The following command displays the version number of the VRRP component:

```
show log components vrrp version
```

The following is sample output from this command:

Component	Title	Version
VRRP	Config/State messages	2.4
Advert	Subsystem description	3.1
System	System/Library messages	3.2

A total of 3 component(s) were displayed.

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show log configuration

show log configuration

Description

Displays the log configuration for switch log settings, and for certain targets.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the log configuration for all targets. The state of the target, enabled or disabled is displayed. For the enabled targets, the associated filter, severity, match expression, and format is displayed. The debug mode state of the switch is also displayed.



Example

The following command displays the configuration of all the log targets and all existing filters:

```
show log configuration
```

The following is sample output from this command:

```

Debug-Mode: Enabled
Log Target      : memory-buffer
Enabled ?      : yes
Filter Name     : DefaultFilter
Match regex    : Any
Severity       : Debug-Data (through Critical)
Format         : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condit
ion>
Buffer size    : 1000 messages
Log Target     : nvrnm
Enabled ?      : yes
Filter Name     : DefaultFilter
Match regex    : Any
Severity       : Warning (through Critical)
Format         : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condit
ion>
Log Target     : console
Enabled ?      : no
Filter Name     : DefaultFilter
Match regex    : Any
Severity       : Info (through Critical)
Format         : MM/DD/YYYY HH:MM:SS.hh
<Severity:Component.SubComponent.Condition>
Log Filter Name: DefaultFilter
I/
E Comp.      Sub-comp.      Condition      Severity
- -----
I All
Log Filter Name: myFilter
I/
E Comp.      Sub-comp.      Condition      Severity
- -----
I STP
Include/Exclude: I - Include, E - Exclude
Component Unreg: * - Component/Subcomponent is not currently registered
Severity Values: C - Critical, E - Error, W - Warning, N - Notice, I -
Info
Debug Severity : S - Debug-Summary, V - Debug-Verbose, D - Debug-Data
+ - Debug Severities, but log debug-mode not enabled
If Match parameters present:
Parameter Flags: S - Source, D - Destination, (as applicable)
I - Ingress, E - Egress, B - BGP
Parameter Types: Port - Physical Port list, Slot - Physical Slot #
MAC - MAC address, IP - IP Address/netmask, Mask - Netmask
VID - Virtual LAN ID (tag), VLAN - Virtual LAN name
L4 - Layer-4 Port #, Num - Number, Str - String
Nbr - Neighbor, Rtr - Routerid, EAPS - EAPS Domain
Proc - Process Name

```



Include/Exclude: I - Include, E - Exclude
 Component Unreg: * - Component/Subcomponent is not currently registered
 Severity Values: C - Critical, E - Error, W - Warning, N - Notice, I - Info
 * - Pre-assigned severities in effect for specified component
 Debug Severity : S - Debug-Summary, V - Debug-Verbose, D - Debug-Data
 + - Debug Severities, but log debug-mode not enabled
 If Match parameters present:
 Parameter Flags: S - Source, D - Destination, (as applicable)
 I - Ingress, E - Egress, B - BGP
 Parameter Types: Port - Physical Port list, Slot - Physical Slot #
 MAC - MAC address, IP - IP Address/netmask, Mask - Netmask
 VID - Virtual LAN ID (tag), VLAN - Virtual LAN name
 L4 - Layer-4 Port #, Num - Number, Str - String
 Nbr - Neighbor, Rtr - Routerid, EAPS - EAPS Domain
 Proc - Process Name
 Strict Match : Y - every match parameter entered must be present in the event
 N - match parameters need not be present in the event

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show log configuration target


```

show log configuration target {console | memory-buffer | nvram | primary-msm |
primary-node | backup-msm | backup-node | session | syslog {ipaddress | ipPort |
vr vr_name} {[local0...local7 ]}}
  
```



Description

Displays the log configuration for the specified target.

Syntax Description

console	Show the log configuration for the console display.
memory-buffer	Show the log configuration for volatile memory.
nvram	Show the log configuration for NVRAM.
primary-msm	Specifies the primary MSM.
	 Note This parameter is available only on modular switches.
primary-node	Specifies the primary node in a stack.



backup-msm	Specifies the backup MSM.  Note This parameter is available only on modular switches.
backup-node	Specifies the backup-node in a stack.
session	Show the log configuration for the current session (including console display).
syslog	Show the configuration for the specified syslog target.
<i>ipaddress</i>	Specifies the syslog IP address.
<i>ipPort</i>	Specifies the UDP port number for the syslog target.
<i>vr_name</i>	Specifies the virtual router that can reach the server IP address.  Note User-created VRs are supported only on the platforms listed for this feature in <i>Feature License Requirements</i> in the <i>ExtremeXOS Concepts Guide</i> .
local10 ... local17	Specifies the local syslog facility.

Default

If no options are specified, the command displays the configuration for the current session and console display.

If a virtual router is not specified, VR-Mgmt is used.

Usage Guidelines

This command displays the log configuration for the specified target. The associated filter, severity, match expression, and format is displayed.

Example

The following command displays the log configuration:

```
show log configuration target
```

The following is sample output from this command:

```
Log Target      : memory-buffer
Enabled ?      : yes
Filter Name     : DefaultFilter
Match regex    : Any
Severity       : Debug-Data (through Critical)
Format         : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>
Buffer size    : 1000 messages
Log Target     : nvram
```



```

Enabled ?      : yes
Filter Name   : DefaultFilter
Match regex   : Any
Severity      : Warning (through Critical)
Format        : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>
Log Target    : console
Enabled ?     : no
Filter Name   : DefaultFilter
Match regex   : Any
Severity      : Info (through Critical)
Format        : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>
Log Target    : primary-msm
Enabled       : yes
Filter Name   : DefaultFilter
Match regex   : Any
Severity      : Warning (through Critical)
Log Target    : backup-msm
Enabled       : yes
Filter Name   : DefaultFilter
Match regex   : Any
Severity      : Warning (through Critical)

```

History

This command was first available in ExtremeXOS 10.1.

The primary-msm and backup-msm options were first available in ExtremeXOS 11.0.

The ipPort parameter was first available in ExtremeXOS 11.0.

The local0 ... local7 keywords were made optional in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

show log configuration target xml-notification

```
show log configuration target xml-notification {xml_target_name}
```

Description

Displays XML target information.

Syntax Description

<i>xml_target_name</i>	Specifies the configured xml notification target.
------------------------	---



Default

N/A.

Usage Guidelines

Use this command to display XML target information.

Example

The following command displays XML target information for all targets:

```
show log configuration target xml-notification
```

Following is sample output from the command:

```
Log Target      : xml-notification (sqa)
Enabled        : yes
Filter Name     : DefaultFilter
Match regex    : Any
Severity       : Info (through Critical)
Log Target      : xml-notification (epi)
Enabled        : yes
Filter Name     : xmlc_filter_epi
Match regex    : Any
Severity       : Info (through Critical)
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on BlackDiamond X8, BlackDiamond 8000 series modules, and Summit Family switches.

show log counters

```
show log counters {event condition | [all | event component]} {include | notified | occurred} {severity severity {only}}
```

Description

Displays the incident counters for events.



Syntax Description

<i>event condition</i>	Specifies the event condition to display.
all	Specifies that all events are to be displayed.
<i>event component</i>	Specifies that all the events associated with a particular component or subcomponent should be displayed.
include	Specifies if one or more targets should be included in this event.
notified	Specifies the number of times this event has occurred.
occurred	Specifies the number of times this event has occurred since the last clear or reboot.
<i>severity</i>	Specifies the minimum severity level of events to display (if the keyword only is omitted).
only	Specifies that only events of the specified severity level are to be displayed.

Default

If severity is not specified, then events of all severity are displayed.

Usage Guidelines

This command displays the incident counters for each event specified. Two incident counters are displayed. One counter displays the number of times an event has occurred, and the other displays the number of times that notification for the event was made to the system (an incident record was injected into the system for further processing). Both incident counters reflect totals accumulated since reboot or since the counters were cleared using the `clear log counters` or `clear counters` command, regardless of whether it was filtered or not.

The keywords `include`, `notified`, and `occurred` only display events with non-zero counter values for the corresponding counter.

This command also displays a reference count (the column titled Rf in the output). The reference count is the number of enabled targets receiving notifications of this event.

See the command `show log` for more information about severity levels.

To get a listing of the event conditions in the system, use the following command:

```
show log events
```

To get a listing of the components present in the system, use the following command:

```
show log components
```



Example

The following command displays the event counters for event conditions of severity debug-summary or greater in the component STP.InBPDU:

```
show log counters stp.inbpdu severity debug-summary
```

The following is sample output from this command:

Comp Notified	SubComp	Condition	Severity	Occurred	In
STP 0	InBPDU	Drop	Error	0	Y
STP 0	InBPDU	Ign	Debug-Summary	0	N
STP 0	InBPDU	Mismatch	Warning	0	Y

Occurred : # of times this event has occurred since last clear or reboot
 Flags : (*) Not all applications responded in time with there count values
 In(cluded): Set to Y(es) if one or more targets filter includes this event
 Notified : # of times this event has occurred when 'Included' was Y(es)

The following command displays the event counters for the event condition PDUDrop in the component STP.InBPDU:

```
show log counters "STP.InBPDU.Drop"
```

The following is sample output from this command:

Comp Notified	SubComp	Condition	Severity	Occurred	In
STP 0	InBPDU	Drop	Error	0	Y

Occurred : # of times this event has occurred since last clear or reboot
 Flags : (*) Not all applications responded in time with there count values
 In(cluded): Set to Y(es) if one or more targets filter includes this event
 Notified : # of times this event has occurred when 'Included' was Y(es)

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.



show log events

```
show log events [event condition | [all | event component] {severity severity
{only}}] {details}
```

Description

Displays information about the individual events (conditions) that can be logged.

Syntax Description

<i>event condition</i>	Specifies the event condition to display.
all	Specifies that all events are to be displayed.
<i>event component</i>	Specifies that all the events associated with a particular component should be displayed.
<i>severity</i>	Specifies the minimum severity level of events to display (if the keyword only is omitted).
only	Specifies that only events of the specified severity level are to be displayed.
details	Specifies that detailed information, including the message format and parameter types, be displayed.

Default

If severity is not specified, then events of all severity are displayed. If detail is not specified, then summary only information is displayed.

Usage Guidelines

This command displays the mnemonic, message format, severity, and parameter types defined for each condition in the event set specified.

See the command `show log` for more information about severity levels.

When the detail option is specified, the message format is displayed for the event conditions specified. The message format parameters are replaced by the value of the parameters when the message is generated.

To get a listing of the components present in the system, use the following command:

```
show log components
```

Example

The following command displays the event conditions of severity debug-summary or greater in the component STP.InBPDU:

```
show log events stp.inbpdu severity debug-summary
```



The following is sample output from this command:

Comp	SubComp	Condition	Severity	Parameters
STP	InBPDU	Drop	Error	2 total
STP	InBPDU	Ign	Debug-Summary	2 total
STP	InBPDU	Mismatch	Warning	2 total

The following command displays the details of the event condition PDUTrace in the component STP.InBPDU:

```
show log events stp.inbpdu.pdutrace details
```

The following is sample output from this command:

Comp	SubComp	Condition	Severity	Parameters
STP	InBPDU	Trace	Debug-Verbose	2 total
0 - string				
1 - string (printf)				
Port=%0%: %1%				

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show ports rxerrors

```
show ports {port_list | stack-ports stacking-port-list} rxerrors {no-refresh}
```

Description

Displays real-time receive error statistics. The switch automatically refreshes the output unless otherwise specified.



Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>stacking-port-list</i>	Specifies one or more stacking ports or slots. Applies to SummitStack and the Summit family switches only.
no-refresh	Specifies that auto-refresh is disabled. The output provides a real-time snapshot of the receive errors at the time the command is issued. This setting is not saved.

Default

The switch automatically refreshes the output.

Usage Guidelines

If you do not specify a port number or range of ports, receive error statistics are displayed for all ports.

If you do not specify the no-refresh parameter, the switch automatically refreshes the output (this is the default behavior).

If you specify the no-refresh parameter, the output provides a snapshot of the real-time receive error statistics at the time you issue the command and displays the output in page-by-page mode (this was the default behavior in ExtremeXOS 11.2 and earlier). This setting is not saved; therefore, you must specify the no-refresh parameter each time you want a snapshot of the port receive errors.

This status information may be useful for your technical support representative if you have a network problem.

Collected Port Receive Error Information

The switch collects the following port receive error information:

- Port Number.
- Link State—The current state of the link. Options are:
 - Active (A)—The link is present at this port.
 - Ready (R)—The port is ready to accept a link.
 - Not Present (NP)—The port is configured, but the module is not installed in the slot (modular switches only).
 - Loopback (L)—The port is in Loopback mode.
- Receive Bad CRC Frames (RX CRC)—The total number of frames received by the port that were of the correct length, but contained a bad FCS value.
- Receive Oversize Frames (RX Over)—The total number of good frames received by the port greater than the supported maximum length of 1,522 bytes.
- Receive Undersize Frames (RX Under)—The total number of frames received by the port that were less than 64 bytes long.
- Receive Fragmented Frames (RX Frag)—The total number of frames received by the port were of incorrect length and contained a bad FCS value.
- Receive Jabber Frames (RX Jabber)—The total number of frames received by the port that was of greater than the support maximum length and had a Cyclic Redundancy Check (CRC) error.



- Receive Alignment Errors (RX Align)—The total number of frames received by the port that occurs if a frame has a CRC error and does not contain an integral number of octets.
- Receive Frames Lost (RX Lost)—The total number of packets dropped due to the memory queue being full.

Port Monitoring Display Keys

For information about the available port monitoring display keys, see the `show ports statistics` command.

Example

The following command displays receive error statistics for slot 5, ports 4 through 7 on a modular switch with auto-refresh disabled:

```
show ports 5:4-5:7 rxerrors no-refresh
```

The following is sample output from this command:

```
Port Rx Error monitor
Port      Link  Rx      Rx      Rx      Rx      Rx      Rx      Rx
State    Crc   Over   Under   Frag   Jabber   Align   Lost
=====
5:4      R      0      0      0      0      0      0      0
5:5      R      0      0      0      0      0      0      0
5:6      R      0      0      0      0      0      0      0
5:7      R      0      0      0      0      0      0      0
=====
Link State: A-Active, R-Ready, NP-Port not present, L-Loopback
```

The following command displays receive error statistics for all ports on the Summit family switches with auto-refresh enabled (the default behavior):

```
show ports rxerrors
```

The following is sample truncated output from this command:

```
Port Rx Error Monitor Tue Jul 5 15:07:13 UTC 2005
Port  Link  Rx      Rx      Rx      Rx      Rx      Rx      Rx
State Crc   Over   Under   Frag   Jabber   Align   Lost
=====
1      R      0      0      0      0      0      0
0      0
2      R      0      0      0      0      0      0
0      0
3      R      0      0      0      0      0      0
```



```

0          0
4          R          0          0          0          0          0
0          0
5          R          0          0          0          0          0
0          0
6          R          0          0          0          0          0
0          0
7          R          0          0          0          0          0
0          0
8          R          0          0          0          0          0
0          0
9          R          0          0          0          0          0
0          0
10         R          0          0          0          0          0
0          0
11         R          0          0          0          0          0
0          0
12         R          0          0          0          0          0
0          0
13         R          0          0          0          0          0
0          0
14         R          0          0          0          0          0
0          0
15         R          0          0          0          0          0
0          0
16         R          0          0          0          0          0
0          0
17         R          0          0          0          0          0
0          0
=====
==
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
0->Clear Counters U->page up D->page down ESC->exit

```

History

This command was first available in ExtremeXOS 10.1.

The link state information was updated to include NP-Port not present in ExtremeXOS 11.2.

Support for the auto-refresh functionality and the no-refresh parameter were added in ExtremeXOS 11.3. Auto-refresh continually updates the display. The no-refresh parameter takes a real-time snapshot of the display at the time you issue the command.

Platform Availability

This command is available on all platforms.

The stack-ports option is available only on SummitStack and the Summit family switches.

show ports statistics

```
show ports {port_list | stack-ports stacking-port-list} statistics {no-refresh}
```



Description

Displays real-time port statistic information. The switch automatically refreshes the output unless otherwise specified.

Syntax Description

<i>stacking-port-list</i>	Specifies one or more stacking slots and ports. Applies to SummitStack and the Summit family switches only.
<i>port_list</i>	Specifies one or more ports or slots and ports.
no-refresh	Specifies that auto-refresh is disabled. The output provides a real-time snapshot of the port statistics at the time the command is issued. This setting is not saved.

Default

The switch automatically refreshes the output.

Usage Guidelines

If you do not specify a port number or range of ports, statistics are displayed for all ports.

If you do not specify the no-refresh parameter, the switch automatically refreshes the output (this is the default behavior).

If you specify the no-refresh parameter, the output provides a snapshot of the real-time port statistics at the time you issue the command and displays the output in page-by-page mode (this was the default behavior in ExtremeXOS 11.2 and earlier). This setting is not saved; therefore, you must specify the no-refresh parameter each time you want a snapshot of the port statistics.

Jumbo frame statistics are displayed for switches only that are configured for jumbo frame support.

This status information may be useful for your technical support representative if you have a network problem.

Collected Port Statistics

The switch collects the following port statistic information:

- Port Number.
- Link State—The current state of the link. Options are:
 - Active (A)—The link is present at this port.
 - Ready (R)—The port is ready to accept a link.
 - Not Present (NP)—The port is configured, but the module is not installed in the slot (modular switches only).
 - Loopback (L)—The port is in Loopback mode.
- Transmitted Packet Count (Tx Pkt Count)—The number of packets that have been successfully transmitted by the port.



- Transmitted Byte Count (Tx Byte Count)—The total number of data bytes successfully transmitted by the port.
- Received Packet Count (RX Pkt Count)—The total number of good packets that have been received by the port.
- Received Byte Count (RX Byte Count)—The total number of bytes that were received by the port, including bad or lost frames. This number includes bytes contained in the Frame Check Sequence (FCS), but excludes bytes in the preamble.
- Received Broadcast (RX Bcast)—The total number of frames received by the port that are addressed to a broadcast address.



Note

On Summit switches, when a broadcast jumbo frame is sent, the RX Bcast counter is not updated. The RX Pkt counter is updated to reflect the received broadcast jumbo frames.

- Received Multicast (RX Mcast)—The total number of frames received by the port that are addressed to a multicast address.

Port Monitoring Display Keys

The following table describes the keys used to control the display that appears if auto-refresh is enabled (the default behavior).

Table 22: Port Monitoring Display Keys with Auto-Refresh Enabled

Key(s)	Description
U	Displays the previous page of ports.
D	Displays the next page of ports.
[Esc]	Exits from the screen.
0	Clears all counters.

The following table describes the keys used to control the display that appears if you auto-refresh is disabled.

Table 23: Port Monitoring Displays Keys with Auto-Refresh Disabled

Key	Description
Q	Exits from the screen.
[Space]	Displays the next page of ports.

Example

The following command displays port statistics for slot 1, ports 1 through 2 on a modular switch with auto-refresh disabled:

```
show ports 1:1-1:2 statistics no-refresh
```



The following is sample output from this command:

```

Port Statistics
Port      Link      Tx Pkt    Tx Byte    Rx Pkt    Rx Byte    Rx      Rx
State     Count    Count     Count      Count     Bcast     Mcast
=====
==
1:1      A        7241     2722608    14482    3968068    0
0
1:2      R         0         0          0         0          0
0
=====
==
Link State: A-Active, R-Ready, NP-Port not present, L-Loopback
    
```

The following command displays port statistics for all ports on the Summit family switches with auto-refresh enabled (the default behavior):

```
show ports statistics
```

The following is truncated sample output from this command:

```

Port Statistics
Port      Link      Tx Pkt    Tx Byte    Rx Pkt    Rx Byte    Rx      Rx
State     Count    Count     Count      Count     Bcast     Mcast
=====
==
1         R         0         0          0          0          0
0         0
2         R         0         0          0          0          0
0         0
3         R         0         0          0          0          0
0         0
4         R         0         0          0          0          0
0         0
5         R         0         0          0          0          0
0         0
6         R         0         0          0          0          0
0         0
7         R         0         0          0          0          0
0         0
8         R         0         0          0          0          0
0         0
9         R         0         0          0          0          0
0         0
10        R         0         0          0          0          0
0         0
11        R         0         0          0          0          0
0         0
12        R         0         0          0          0          0
0         0
13        R         0         0          0          0          0
0         0
14        R         0         0          0          0          0
0         0
=====
Tue Jul  5 14:18:34 UTC 2005
    
```



```

0      0
15     R      0      0      0      0
0      0
16     R      0      0      0      0
0      0
17     R      0      0      0      0
0      0
=====
==
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
0->Clear Counters  U->page up  D->page down  ESC->exitPort Statistics

```

History

This command was first available in ExtremeXOS 10.1.

The link state information was updated to include NP-Port not present in ExtremeXOS 11.2.

Support for the auto-refresh functionality and the no-refresh parameter were added in ExtremeXOS 11.3. Auto-refresh continually updates the display. The no-refresh parameter takes a real-time snapshot of the display at the time you issue the command.

Platform Availability

This command is available on all platforms.

The stack-ports option is available on SummitStack and the Summit Family switches only.

show ports txerrors

```
show ports {port_list | stack-ports stacking-port-list} txerrors {no-refresh}
```

Description

Displays real-time transmit error statistics. The switch automatically refreshes the output unless otherwise specified.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>stacking-port-list</i>	Specifies one or more stacking slot ports for display. Applies to SummitStack and Summit family switches only.
no-refresh	Specifies that auto-refresh is disabled. The output provides a real-time snapshot of the transmit errors at the time the command is issued. This setting is not saved.

Default

The switch automatically refreshes the output.



Usage Guidelines

If you do not specify a port number or range of ports, error statistics are displayed for all ports.

If you do not specify the no-refresh parameter, the switch automatically refreshes the output (this is the default behavior).

If you specify the no-refresh parameter, the output provides a snapshot of the real-time transmit error statistics at the time you issue the command and displays the output in page-by-page mode (this was the default behavior in ExtremeXOS 11.2 and earlier). This setting is not saved; therefore, you must specify the no-refresh parameter each time you want a snapshot of the port transmit errors.

This status information may be useful for your technical support representative if you have a network problem.

Collected Port Transmit Error Information

The switch collects the following port transmit error information:

- Port Number.
- Link State—The current state of the link. Options are:
 - Active (A)—The link is present at this port.
 - Ready (R)—The port is ready to accept a link.
 - Not Present (NP)—The port is configured, but the module is not installed in the slot (modular switches only).
 - Loopback (L)—The port is in Loopback mode.
- Transmit Collisions (TX Coll)—The total number of collisions seen by the port, regardless of whether a device connected to the port participated in any of the collisions.
- Transmit Late Collisions (TX Late Coll)—The total number of collisions that have occurred after the port's transmit window has expired.
- Transmit Deferred Frames (TX Deferred)—The total number of frames that were transmitted by the port after the first transmission attempt was deferred by other network traffic.
- Transmit Errored Frames (TX Errors)—The total number of frames that were not completely transmitted by the port because of network errors (such as late collisions or excessive collisions).
- Transmit Lost Frames (TX Lost)—The total number of transmit frames that do not get completely transmitted because of buffer problems (FIFO underflow).
- Transmit Parity Frames (TX Parity)—The bit summation has a parity mismatch.

Port Monitoring Display Keys

For information about the available port monitoring display keys, see the `show ports statistics` command.



Example

The following command displays transmit error statistics for slot 5, ports 4 through 7 on a modular switch with auto-refresh disabled:

```
show ports 5:4-5:7 txerrors no-refresh
```

The following is sample output from this command:

```
Port Transmission errors
Port      Link   Tx      Tx      Tx      Tx      Tx
State    Coll  Late coll  Deferred  Errors  Lost  Parity
=====
==
5:4      R      0      0      0      0      0
5:5      R      0      0      0      0      0
5:6      R      0      0      0      0      0
5:7      R      0      0      0      0      0
=====
==
Link State: A-Active, R-Ready, NP-Port not present, L-Loopback
```

The following command displays transmit error statistics for all ports on a Summit switch:

```
show ports txerrors
```

The following is a truncated sample output from this command:

```
Port Tx Error Monitor                               Tue Jul  5 15:07:13 UTC
2005
Port      Link   Tx      Tx      Tx      Tx      Tx      Tx
State    Coll  Late coll  Deferred  Errors  Lost  Parity
=====
==
1      0      R      0      0      0      0
2      0      R      0      0      0      0
3      0      R      0      0      0      0
4      0      R      0      0      0      0
5      0      R      0      0      0      0
6      0      R      0      0      0      0
7      0      R      0      0      0      0
8      0      R      0      0      0      0
9      0      R      0      0      0      0
10     0      R      0      0      0      0
```



```

0          0
11         R          0          0          0          0
0          0
12         R          0          0          0          0
0          0
13         R          0          0          0          0
0          0
14         R          0          0          0          0
0          0
15         R          0          0          0          0
0          0
16         R          0          0          0          0
0          0
17         R          0          0          0          0
0          0
=====
==
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
0->Clear Counters U->page up D->page down ESC->exitPort Tx Error

```

History

This command was first available in ExtremeXOS 10.1.

The link state information was updated to include NP-Port not present in ExtremeXOS 11.2.

Support for the auto-refresh functionality and the no-refresh parameter were added in ExtremeXOS 11.3. Auto-refresh continually updates the display. The no-refresh parameter takes a real-time snapshot of the display at the time you issue the command.

Platform Availability

This command is available on all platforms.

show ports vlan statistics

```
show ports {port_list} vlan statistics {no-refresh}
```

Description

Displays VLAN statistics at the port level.

Syntax Description

port_list	Specifies one or more ports or slots and ports. Can be one or more port numbers. May be in the form: 1, 2, 3-5, 1*, 1:5, 1:6-1:8.
no-refresh	Specifies that there is no continuous refresh. The prompt comes back to the user after fetching statistics once.



Default

N/A.

Usage Guidelines

This command is used in conjunction with the `configure ports [port_list|all] monitor vlanvlan_name {rx-only | tx-only}` command.

Example

The following command displays statistics for the ports 1-2 on slot 5 of a BlackDiamond 12804 switch:

```

show ports 5:1-2 vlan stats
Displays the vlan statistics in a real time continuous refresh mode or no-
refresh mode.
* (debug) BD-12804.31 # show ports 5:1-2 vlan stat
Port VLAN Statistics                               Wed Mar 28 10:52:59 2007
Port   Vlan      Rx Frames      Rx Byte      Tx Frame      Tx Byte
Count   Count      Count          Count        Count
=====
==
5:1     Default  318750522     20400046784   318750588     20400051672
5:2     Default  292811491     18739948736   292811975     18739980504
0->Clear Counters  U->page up  D->page down ESC->exit

```

For ports that do not support transmit statistics, a '-' will be displayed. For ports that do not support transmit byte counters, a '-' will be displayed for that row and column. Similarly, configuration using rx-only or tx-only will result in the display of "-"s in the appropriate rows and columns.

History

This command was first available in ExtremeXOS 12.0.

Support for BlackDiamond 8000 series modules, SummitStack, and Summit family switches was added in ExtremeXOS 12.5.

Support for BlackDiamond X8 series switches was added in ExtremeXOS 15.1.

Platform Availability

This command is available only on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, SummitStack, and Summit family switches.

show rmon memory

```
show rmon memory {detail | memoryType}
```



Description

Displays RMON specific memory usage and statistics.

Syntax Description

detail	Displays detailed information.
<i>memoryType</i>	Specifies the type of memory usage and statistics to display.

Default

N/A.

Usage Guidelines

If you do not specify the detailed keyword or a enter a specific RMON memory type, the output contains usage information for all memory types.

Example

The following command displays RMON memory statistics:

```
show rmon memory
```

The following is sample output from this command:

```

RMON Memory Information
-----
Bytes Allocated: 14298032 AllocFailed: 0
Current Memory Utilization Level: GREEN
Memory Utilization Statistics
-----
Size      16      32      48      64      80      96      112     128     144     176
208      256     384      5
12      768     1024    2048    4096    8192   16384   18432   40960   64000
-----
-----
Used Blocks 1558      3  2490      1      0      0      0      1
1      0 63444      1  1869
0      311      0      0      0      0      0      0      0      0
rmonEstat  0      0      0      0      0      0      0      0      0      0
0      0      0      311
0      0      0      0      0      0      0      0      0      0
rmonOwner 1555      0      0      0      0      0      0      0      0      0
0      0      0      0
0      0      0      0      0      0      0      0      0      0
rmonHisc  0      0      0      0      0      0      0      0      0      0
0      0      0  1244
0      0      0      0      0      0      0      0      0      0
rmonHist  0      0      0      0      0      0      0      0      0      0

```




```

0 63444      0      0
0      0      0      0      0      0      0      0      0      0
rmonAlarm    0      0      0      0      0      0      0      0      0      0
0      0      0      3
0      0      0      0      0      0      0      0      0      0
rmonLogDescription 0      0      0      0      0      0      0      0      0
0      1      0      0      0      0      0      0      0      0
0      0      0      0      0      0      0      0      0      0
rmonLog      0      1      0      0      0      0      0      0      0      0
0      0      0      0
0      0      0      0      0      0      0      0      0      0
rmonEvent    0      0      0      0      0      0      0      0      1      0
0      0      0      0
0      0      0      0      0      0      0      0      0      0
rmonEventDescription 0      1      0      0      0      0      0      0
0      0      0      0      0
0      0      0      0      0      0      0      0      0      0
rmonEventCommunity 0      1      0      0      0      0      0      0
0      0      0      0      0
0      0      0      0      0      0      0      0      0      0
rmonCommunity 1      0      0      0      0      0      0      0
0      0      0      0      0
0      0      0      0      0      0      0      0      0      0
rmonDs       0      0      0      0      0      0      0      0
0      0      0      0
0      311     0      0      0      0      0      0      0      0
rmonDbx      0      0      2490  0      0      0      0      0
0      0      0      0
0      0      0      0      0      0      0      0      0      0
rmonOid      0      0      0      0      0      0      0      0
0      0      0      311
0      0      0      0      0      0      0      0      0      0
rmonMdbIndexOid 2      0      0      1      0      0      0      0
0      0      0      0      0      0      0      0      0      0
0      0      0      0      0      0      0      0      0      0
rmonMdbString 0      0      0      0      0      0      0      0
0      0      0      1      0
0      0      0      0      0      0      0      0      0      0

```

The following command displays RMON event statistics:

```
show rmon memory rmonEvent
```

The following is sample output from this command:

```

RMON Memory Information
-----
Bytes Allocated: 14298032 AllocFailed: 0
Current Memory Utilization Level: GREEN
Memory Utilization Statistics
-----
Memory Statistics for rmonEvent
-----
Size      16      32      48      64      80      96      112     128     144     176
208     256     384     512     768     1024    2048    4096    8192   16384   18432

```



```

40960 64000
-----
-----
-----
Alloced      0      0      0      0      0      0      0      0      1      0
0      0      0      0
0      0      0      0      0      0      0      0      0      0
AllocedPeak  0      0      0      0      0      0      0      0      0      1
0      0      0      0      0      0      0      0      0      0
0      0      0      0      0      0      0      0      0      0
AllocSuccess 0      0      0      0      0      0      0      0      0      1
0      0      0      0      0      0      0      0      0      0
0      0      0      0      0      0      0      0      0      0
FreeSuccess  0      0      0      0      0      0      0      0      0      0
0      0      0      0      0
0      0      0      0      0      0      0      0      0      0
AllocFail    0      0      0      0      0      0      0      0      0      0
0      0      0      0
0      0      0      0      0      0      0      0      0      0
FreeFail     0      0      0      0      0      0      0      0      0      0
0      0      0      0
0      0      0      0      0      0      0      0      0      0

```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

show sflow configuration

```
show sflow {configuration} {egress}
```

Description

Displays the current sFlow configuration.

Syntax Description

configuration	Displays the sFlow configuration.
ingress	Displays type of sFlow and the sampling rate.
egress	Displays type of sFlow and the sampling rate.

Default

N/A.



Usage Guidelines

This command displays the sFlow configuration of your system.

The following fields are displayed:

- Global Status—sFlow is globally enabled or disabled.
- Polling interval—How often the hardware is polled for statistics, in seconds.
- Sampling rate—Packets are sampled, on average, once for every rate-number of packets.
- Maximum cpu sample limit—Maximum number of packets per second sampled before sample throttling takes effect.
- Agent IP—IP address inserted into the sFlow data packets to identify the sFlow switch.
- Collectors—To which IP address and port, and from which virtual router, the sFlow packets are sent.
- Port Status—Enabled or disabled for statistics gathering.
- Port Sample-rate—Shows the sampling rate configured for the port and the actual rate if CPU throttling has taken effect.
- Port Subsampling factor—See the command `configure sflow ports sample-rate` for details.

Example

To display the sFlow configuration on your system, use the following command:

```
show sflow
```

The following is an example of the show sflow configuration command :

```
SFLOW Global Configuration
Global Status: enabled
Polling interval: 20
Sampling rate: 8192
Maximum cpu sample limit: 200000
SFLOW Configured Agent IP: 0.0.0.0
Operational Agent IP: 10.127.11.88
Collectors

SFLOW Port Configuration
Port  Status                Sample-rate                Subsampling                Sflow-type
      Config / Actual        factor                    Ingress / Egress
5:21  enabled                 8192 / 8192              1                          Disabled / Enabled
```

The following example displays the type of sFlow configured on the physical interface, and the flow rate:

```
# show sflow configuration egress
SFLOW Global ConfigurationGlobal
Status: enabled
Polling interval: 20
Sampling rate: 8192
Maximum cpu sample limit: 200000
SFLOW Configured Agent IP: 0.0.0.0
Operational Agent IP: 10.127.11.88
Collectors
```



```

SFLOW Port Configuration
Port      Status          Sample-rate          Subsampling          Sflow-
type
                                     Config / Actual      factor
Ingress / Egress
5:21     enabled          8192 / 8192         1
Disabled / Enabled

```

History

This command was first available in an ExtremeXOS 11.0.

The output for the **ingress** and **egress** keywords was added in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

show sflow statistics

```
show sflow statistics
```

Description

Displays sFlow statistics.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays sFlow statistics for your system.

The following fields are displayed:

- Received frames—Number of frames received on sFlow enabled ports.
- Sampled Frames—Number of packets that have been sampled by sFlow.
- Transmitted Frames—Number of UDP packets sent to remote collector(s).
- Broadcast Frames—Number of broadcast frames received on sFlow enabled ports.
- Multicast Frames—Number of multicast frames received on sFlow enabled ports.
- Packet Drops—Number of samples dropped.



Example

To display sFlow statistics for your system, use the following command:

```
show sflow statistics
```

The output from this command is similar to the following:

```
SFLOW Statistics

Received frames      : 1159044921
Sampled Frames      : 104944
Transmitted Frames  : 10518
Broadcast Frames    : 0
Multicast Frames    : 1055652
Packet Drops        : 0
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show temperature

```
show temperature
```

Description

Depending on the platform, this command displays the current temperature of the I/O modules, management modules, power supply controllers, XGM-2xn card, and the switch.

On a stack, the command displays the current temperature of the modules in each slot.

Syntax Description

This command has no arguments or variables

Default

N/A.



Usage Guidelines

Depending on the software version running on your switch or your switch model, additional or different temperature information might be displayed.

Modular Switches Only

Use this command to display the temperature in Celsius and the current status of the following installed components in the switch:

- Management modules (MSM/MM).
- I/O modules.
- Power controllers.

The switch monitors the temperature of each component and generates a warning if the temperature exceeds the normal operating range. If the temperature exceeds the minimum/maximum limits, the switch shuts down the overheated module.

Summit Family Switches Only

Use this command to display the temperature in Celsius and the current status of the following components:

- Switch.
- XGM-2xn card.

SummitStack Only.

Use this command to display the temperature in Celsius and the current status of the following components:

- All switches in the stack.
- XGM-2xn cards if present.

The switch monitors its temperature and generates a warning if the temperature exceeds the normal operating range. If the temperature exceeds the maximum limit, the `show switch` output indicates the switch in an OPERATIONAL (Overheat) mode, and the `show temperature` output indicates an error state due to overheat.

Displaying the Temperature of Other Installed Components--Modular Switches Only

You can also view the temperature of the power supplies and the fan trays in the switch.

To view the temperature of the power supplies installed in a modular switch, use the following command:

```
show power {ps_num} {detail}
```



Example

Depending on the platform, the following command displays the temperature of various switch components:

```
show temperature
```

In the BlackDiamond X8 switch, the temperature shown is the adjusted maximum of all temperatures from all monitored devices on the card. This is due to the fact that individual device temperature sensors are monitored within the device's operating temperature range. The following is sample output from a BlackDiamond X8 switch:

```
BD-X8.3 # show temperature
Field Replaceable Units           Temp (C)   Status   Min   Normal   Max
-----
Slot-1       : BDXA-10G48X           71.00    Normal    0   25-100   107
Slot-2       : BDXA-10G48X           71.00    Normal    0   25-100   107
Slot-3       :
Slot-4       :
Slot-5       :
Slot-6       : BDXA-40G24X           86.00    Normal    0   25-100   107
Slot-7       : BDXA-40G24X           86.00    Normal    0   25-100   107
Slot-8       : BDXA-40G24X           90.00    Normal    0   25-100   107
FM-1         : BDXA-FM20T             85.00    Normal    0   25-100   107
FM-2         : BDXA-FM20T             83.00    Normal    0   25-100   107
FM-3         : BDXA-FM20T             88.00    Normal    0   25-100   107
FM-4         : BDXA-FM20T             88.00    Normal    0   25-100   107
MM-A         : BDX-MM1                79.00    Normal    0   25-100   105
MM-B         : BDX-MM1                88.00    Normal    0   25-100   105
```

The following is sample output from a BlackDiamond 8810 switch:

```
BD-8810.7 # show temperature
Field Replaceable Units           Temp (C)   Status   Min   Normal   Max
-----
Slot-1       : 8900-10G24X-c           37.00    Normal  -10   0-55    65
Slot-2       : 10G8Xc                   33.00    Normal  -10   0-50    60
Slot-3       :
Slot-4       :
Slot-5       : G8Xc                      31.50    Normal  -10   0-50    60
Slot-6       :
Slot-7       :
Slot-8       : 10G4Xa                    26.00    Normal  -10   0-50    60
Slot-9       :
Slot-10      : 8900-G96T-c              38.50    Normal  -10   0-55    65
MSM-A        : 8900-MSM128              31.00    Normal  -10   0-55    65
MSM-B        : 8900-MSM128              31.00    Normal  -10   0-55    65
PSUCTRL-1    :                          34.04    Normal  -10   0-50    60
PSUCTRL-2    :                          35.79    Normal  -10   0-50    60
```



The following is sample output from a SummitStack of 8 nodes:

```
Slot-3 Stack.1 # show temperature
Field Replaceable Units          Temp (C)   Status   Min   Normal  Max
-----
Slot-1          :
Slot-2          : SummitX      34.50    Normal  -10   0-54   59
Slot-3          : SummitX      36.50    Normal  -10   0-66   67
Slot-4          :
Slot-5          :
Slot-6          :
Slot-7          :
Slot-8          :
Slot-3 Stack.2 #
```

History

This command was first available in an ExtremeXOS 10.1.

Information about the power controller(s), a component status column, and the minimum, normal, and maximum temperature ranges of the components was added to the `show temperature` output in ExtremeXOS 11.0.

Information about the XGM-2xn card was added to the `show temperature` output in ExtremeXOS 11.2.

Support for stacking was added to `show temperature` output in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

show version

```
show version {detail | process name | images {partition partition} {slot
slotid} }
```


Description

Displays the hardware serial and version numbers, the software version currently running on the switch, and (if applicable) the software version running on the modules and power controllers.

Syntax Description

detail	Specifies display of slot board name and chassis or platform name.
process	Specifies display of all of the processes on the switch.
<i>name</i>	Specifies display of a specific process on the switch.
images	Specifies the display of installed images.



<i>partition</i>	Specifies display of a specific partition (primary or secondary).
<i>slotid</i>	Specifies display of an MSM/MM in a specific slot (A or B).
 Note This parameter is available only on modular switches.	

Default

N/A.

Usage Guidelines

The following describes the information displayed when you execute the `show version` or `show version detail` commands:

- **Part Number**—A collection of numbers and letters that make up the part number of the switch and when applicable the hardware components installed in a modular switch.
- **Serial Number**—A collection of numbers and letters that make up the serial number of the switch and when applicable the hardware components installed in a modular switch.



Note

For information about the physical location of the serial number on your switch, refer to the section that describes your specific switch model in the hardware documentation.

- **Image**—The ExtremeXOS software version currently running on the switch. If you have two software images downloaded on the switch, only the currently running ExtremeXOS version information is displayed. The information displayed includes the major version number, minor version number, a specific patch release, and the build number. The software build date is also displayed.
- **BootROM**—The BootROM version currently running on the switch.
- **Diagnostics**—For BlackDiamond 8800 series switches, this is the version number of operational diagnostics software that runs on the I/O module included in the particular version of ExtremeXOS. For BlackDiamond X8 switches, this is the version number of operational diagnostics software that runs either on the Management Module (MM) or the I/O and Fabric Modules (FM), respectively, included in the particular version of ExtremeXOS.
- **FPGA**—The field-programmable gate array firmware version currently running on the module (BlackDiamond X8 switches only).

Depending on the model of your switch and the software running on your switch, different version information may be displayed.



Note

The information displayed does not include the I/O version number on the BlackDiamond 8800 series switch. The I/O version number includes the major, minor, and I/O version number, not the patch and build numbers.

If you use the `process` option, you will see the following information about the processes running on the switch:



- Card—The location (MSM/MM) where the process is running on a modular switch.
- Process Name—The name of the process.
- Version—The version number of the process.
- BuiltBy—The name of the software build manager.
- Link Date—The date the executable was linked.

Example

The following command displays the hardware and software versions currently running on the switch:

```
show version
```

The following is sample output from a BlackDiamond X8 switch:

```
BD-X8.5 # show version
Chassis      : 800427-00-01 1135G-02182 Rev 1.0
Slot-1       : 800435-00-01 1138G-00252 Rev 1.0 BootROM: 1.0.0.6   IMG:
15.1.0.30
Slot-2       : 800435-00-01 1138G-00250 Rev 1.0 BootROM: 1.0.0.6   IMG:
15.1.0.30
Slot-3       :
Slot-4       :
Slot-5       :
Slot-6       : 800439-00-01 1138G-00342 Rev 1.0 BootROM: 1.0.0.6   IMG:
15.1.0.30
Slot-7       : 800439-00-01 1138G-00335 Rev 1.0 BootROM: 1.0.0.6   IMG:
15.1.0.30
Slot-8       : 800439-00-01 1138G-00346 Rev 1.0 BootROM: 1.0.0.6   IMG:
15.1.0.30
FM-1         : 800433-00-01 1138G-00179 Rev 1.0 BootROM: 1.0.0.6   IMG:
15.1.0.30
FM-2         : 800433-00-01 1138G-00225 Rev 1.0 BootROM: 1.0.0.6   IMG:
15.1.0.30
FM-3         : 800433-00-01 1138G-00220 Rev 1.0 BootROM: 1.0.0.6   IMG:
15.1.0.30
FM-4         : 800433-00-01 1138G-00224 Rev 1.0 BootROM: 1.0.0.6   IMG:
15.1.0.30
MM-A         : 800432-00-01 1136G-00468 Rev 1.0 BootROM: 1.0.0.2   IMG:
15.1.0.30
MM-B         : 800432-00-01 1136G-00473 Rev 1.0 BootROM: 1.0.0.2   IMG:
15.1.0.30
PSUCTRL-1   : 450357-00-01 1135G-02368 Rev 1.0
PSUCTRL-2   : 450357-00-01 1135G-02399 Rev 1.0
FanTray-1   : 450350-00-01 1135G-02213 Rev 1.0 BootROM: 1.0.2.5
FanTray-2   : 450350-00-01 1135G-02244 Rev 1.0 BootROM: 1.0.2.5
FanTray-3   : 450350-00-01 1135G-02306 Rev 1.0 BootROM: 1.0.2.5
FanTray-4   : 450350-00-01 1135G-02275 Rev 1.0 BootROM: 1.0.2.5
FanTray-5   : 450350-00-01 1135G-02337 Rev 1.0 BootROM: 1.0.2.5
PSU-1       :
PSU-2       :
PSU-3       :
PSU-4       :
PSU-5       : H2500A2-EX 800429-00 1121X-88454 Rev 1.0
```



```

PSU-6      : H2500A2-EX  800429-00 1121X-88598 Rev 1.0
PSU-7      : H2500A2-EX  800429-00 1109X-88765 Rev 1.0
PSU-8      :
Image      : ExtremeXOS version 15.1.0.30 v1510b30 by release-manager
on Thu Jan 12 12:57:57 EST 2012
BootROM    : 1.0.0.2
Diagnostics : 1.8 (MM), 1.6 (I/O and FM)

```

The following is sample output from a BlackDiamond 8810 switch (the output from the BlackDiamond 8806 is similar):

```

Chassis    : 800129-00-02 04344-00039 Rev 2.0
Slot-1     : 800114-00-04 04364-00021 Rev 4.0 BootROM: 1.0.4.0   IMG:
12.4.0.11
Slot-2     : 800115-00-02 04344-00006 Rev 2.0
Slot-3     : 800113-00-04 04354-00031 Rev 4.0 BootROM: 1.0.4.0   IMG:
12.4.0.11
Slot-4     :
Slot-5     : 800112-00-03 04334-00040 Rev 3.0 BootROM: 1.0.4.0   IMG:
12.4.0.11
Slot-6     : 800112-00-03 04334-00004 Rev 3.0 BootROM: 1.0.3.7   IMG:
12.4.0.11
Slot-7     :
Slot-8     : 800157-00-02 06034-00015 Rev 2.0 BootROM: 1.0.4.0   IMG:
12.4.0.11
Slot-9     : 800159-00-02 06044-00037 Rev 2.0 BootROM: 1.0.4.0   IMG:
12.4.0.11
Slot-10    :
MSM-A      : 800112-00-03 04334-00040 Rev 3.0 BootROM: 1.0.4.2   IMG:
12.4.0.11
MSM-B      : 800112-00-03 04334-00004 Rev 3.0 BootROM: 1.0.3.8   IMG:
12.4.0.11
PSUCTRL-1  : 450117-00-01 04334-00021 Rev 1.0 BootROM: 2.16
PSUCTRL-2  : 450117-00-01 04334-00068 Rev 1.0 BootROM: 2.16
PSU-1      : PS 2336 4300-00137 0441J-01807 Rev 5.0
PSU-2      : PS 2336 4300-00137 0536J-06779 Rev 7.0
PSU-3      :
PSU-4      :
PSU-5      :
PSU-6      :
Image      : ExtremeXOS version 12.4.0.11 v1240b11 by release-manager
on Wed Nov 18 19:39:10 PST 2009
BootROM    : 1.0.4.2
Diagnostics : 1.10

```

The following is a sample from a SummitStack:

```

Slot-1     : 800152-00-04 0624G-01015 Rev 4.0 BootROM: 1.0.2.0   IMG: 12.0.0.4
Slot-2     : 800192-00-01 0620G-00020 Rev 1.0 BootROM: 1.0.1.9   IMG: 12.0.0.4
Slot-3     : 800163-00-04 0630G-00804 Rev 4.0 BootROM: 1.0.2.0   IMG: 12.0.0.4
Slot-4     : 800153-00-04 0630G-00672 Rev 4.0 BootROM: 1.0.2.0   IMG: 12.0.0.4
Slot-5     :
Slot-6     :
Slot-7     :
Slot-8     :

```



```
XGM2-1      :
Image       : ExtremeXOS version 12.0.0.4 branch-fixes_v1200b4 by mmroz
on Thu Mar 15 10:55:58 EDT 2007
BootROM    : 1.0.2.0
```

Using the process option of the show version command produces output similar to the following on a modular switch:

Card	Process Name	Version	BuiltBy	Link Date
MSM-A	aaa	3.0.0.2	release-manager	Thu Mar 31 09:23:54 PST 2005
MSM-A	acl	3.0.0.2	release-manager	Thu Mar 31 09:26:46 PST 2005
MSM-A	bgp	3.0.0.2	release-manager	Thu Mar 31 09:27:54 PST 2005
MSM-A	cfgmgr	3.0.0.21	release-manager	Thu Mar 31 09:23:42 PST 2005
MSM-A	cli	3.0.0.22	release-manager	Thu Mar 31 09:23:34 PST 2005
MSM-A	devmgr	3.0.0.2	release-manager	Thu Mar 31 09:23:22 PST 2005
MSM-A	dirser	3.0.0.2	release-manager	Thu Mar 31 09:24:02 PST 2005
MSM-A	eaps	3.0.0.8	release-manager	Thu Mar 31 09:26:34 PST 2005
MSM-A	edp	3.0.0.2	release-manager	Thu Mar 31 09:25:56 PST 2005
MSM-A	elrp	3.0.0.1	release-manager	Thu Mar 31 09:25:14 PST 2005
MSM-A	ems	3.0.0.2	release-manager	Thu Mar 31 09:35:08 PST 2005
MSM-A	epm	3.0.0.3	release-manager	Thu Mar 31 09:23:11 PST 2005
MSM-A	esrp	3.0.0.4	release-manager	Thu Mar 31 09:26:23 PST 2005
			

The following is sample output from the Summit switch:

Process Name	Version	BuiltBy	Link Date
aaa	3.0.0.2	release-manager	Thu Mar 31 09:34:17 PST 2005
acl	3.0.0.2	release-manager	Thu Mar 31 09:38:44 PST 2005
bgp	Not Started	Unknown	Unknown
cfgmgr	3.0.0.21	release-manager	Thu Mar 31 09:33:58 PST 2005
cli	3.0.0.22	release-manager	Thu Mar 31 09:33:45 PST 2005
cna	3.1.0.1	release-manager	Thu Mar 31 09:49:28 PST 2005
devmgr	3.0.0.2	release-manager	Thu Mar 31 09:33:26 PST 2005
dirser	3.0.0.2	release-manager	Thu Mar 31 09:34:31 PST 2005
dosprotect	3.0.0.1	release-manager	Thu Mar 31 09:48:58 PST 2005
eaps	3.0.0.8	release-manager	Thu Mar 31 09:38:25 PST 2005
edp	3.0.0.2	release-manager	Thu Mar 31 09:37:24 PST 2005
elrp	3.0.0.1	release-manager	Thu Mar 31 09:36:27 PST 2005
ems	3.0.0.2	release-manager	Thu Mar 31 09:50:40 PST 2005



```

epm          3.0.0.3   release-manager Thu Mar 31 09:33:08 PST 2005
esrp         3.0.0.4   release-manager Thu Mar 31 09:38:07 PST 2005
etmon        1.0.0.1   release-manager Thu Mar 31 09:47:16 PST 2005
. . . .

```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show vlan statistics

```
show vlan {vlan_name} statistics {no-refresh}
```

Description

Displays VLAN statistics at the VLAN level.

Syntax Description

vlan_name	Specifies to display VLAN statistics from the VLAN with this name.
no-refresh	Specifies that there is no continuous refresh. The prompt comes back to the user after fetching statistics once.

Default

N/A.

Usage Guidelines

This command displays statistics based on the sum of the statistics for individual ports. Use it to display the VLAN statistics monitored using the `configure ports [port_list|all] monitor vlanvlan_name {rx-only | tx-only}` command.

Example

The following command displays VLAN statistics:

```

* (debug) BD-12804.17 # show vlan statistics no-refresh
Vlan          Rx Total          Rx
Byte          Tx Total          Tx Byte
              Frames
Count        Frames          Count
=====

```



```

=
Default                30251013                22034
7326296
901840
=====
=

```

If the VLAN contains ports that do not support a certain type of VLAN statistic, such as transmit statistics or byte counters, then a dash character ('-') will be displayed in that column.

History

This command was first available in ExtremeXOS 12.0.

Support for BlackDiamond 8000 series modules, SummitStack, and Summit family switches was added in ExtremeXOS 12.5.

Platform Availability

This command is available only on BlackDiamond X8, BlackDiamond 8000 series modules, SummitStack, and Summit family switches.

show xml-notification configuration

```
show xml-notification configuration {target}
```

Description

Displays the configuration of the Web server target.

Syntax Description

<i>target</i>	Specifies an alpha numeric string that identifies the configured target.
---------------	--

Default

N/A.

Usage Guidelines

Use this command to display information about the configuration of the Web server target. If a target is not specified, all configured targets are displayed.



Example

The following command displays the configuration of the configured targets:

```
show xml-notification configuration
```

The following is sample output from this command:

```
Target Name           : sqa
**Server URL          : http://10.255.129.22:8080/xos/webservice (VR-Mgmt)
Server User Name      : admin
Enabled               : yes
Queue Size            : 100
Connection Status     : connected
Configured Modules    : ems,idmgr
Target Name           : epi
**Server URL          : http://10.255.59.6:8080/xos/webservice (VR-Finance)
Server User Name      : admin
Enabled               : yes
Queue Size            : 100
Connection Status     : connected
Configured Modules    : ems
Target Name           : test3
**Server URL          : https://10.120.91.64:8443/xos/webservice (VR-Mgmt)
Server User Name      : admin
Enabled               : yes
Queue Size            : 100
Connection Status     : not connected
Configured Modules    : ems
Target Name           : testingcorrect
**Server URL          : http://10.66.254.211:8080/xos/webservice (VR-Mgmt)
Server User Name      : admin
Enabled               : no
Queue Size            : 100
Connection Status     : not connected
Configured Modules    : idMgr,ems
```

Note



When a particular VR has been specified in the configuration process, that VR is displayed next to the URL. When no VR is specified since the parameter is optional, the default VR supplied by the XML client is VR-Mgmt. When you are using a version released before the virtual router option was added, VR-Mgmt is displayed.

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on BlackDiamond X8, BlackDiamond 8000 series modules, and Summit Family switches.



show xml-notification statistics

```
show xml-notification statistics {target}
```

Description

Displays statistics for of the Web server target.

Syntax Description

<i>target</i>	Specifies an alpha numeric string that identifies the configured target.
---------------	--

Default

N/A.

Usage Guidelines

Use this command to display the connection status, enable status and event statistics of the Web server target. If a target is not specified, all configured targets are displayed.

Example

The following command displays statistics for all of the configured targets:

```
show xml-notification statistics
```

The following is sample output from this command:

```
Target Name           : epi
Server URL            : http://10.255.129.22:8080/xos/webservice
Server Queue Size     : 100
Enabled               : yes
Connection Status     : connected
Events Received       : 450
Connection Failures   : 0
Events Sent Success   : 450
Events Sent Fail      : 0
Events Dropped        : 0
Target Name           : epi
Server URL            : http://10.255.59.6:8080/xos/webservice
Server Queue Size     : 100
Enabled               : yes
Connection Status     : fail
Events Received       : 31
Connection Failures   : 3
Events Sent Success   : 2
Events Sent Fail      : 29
Events Dropped        : 0
```



History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on BlackDiamond X8, BlackDiamond 8000 series modules, and Summit Family switches.

unconfigure log filter

```
unconfigure log filter filter_name
```

Description

Resets the log filter to its default values; removes all filter items.

Syntax Description

<i>filter_name</i>	Specifies the log filter to unconfigure.
--------------------	--

Default

N/A.

Usage Guidelines

If the filter name specified is DefaultFilter, this command restores the configuration of DefaultFilter back to its original settings.

If the filter name specified is not DefaultFilter, this command sets the filter to have no events configured and therefore, no incidents will pass. This is the configuration of a newly created filter that was not copied from an existing one.

See the `delete log filter` command for information about deleting a filter.

Example

The following command sets the log filter myFilter to stop passing any events:

```
unconfigure log filter myFilter
```

History

This command was first available in ExtremeXOS 10.1.



Platform Availability

This command is available all platforms.


unconfigure log target format

```
unconfigure log target [console | memory-buffer | nvram | session | syslog [all | ipaddress | ipPort {vr vr_name} [local0...local7]]] format
```

Description

Resets the log target format to its default values.

Syntax Description

console	Specifies the console display format.
memory-buffer	Specifies the switch memory buffer format.
nvr am	Specifies the switch NVRAM format.
session	Specifies the current session (including console display) format.
syslog	Specifies a syslog target format.
all	Specifies all remote syslog servers.
<i>ipaddress</i>	Specifies the syslog IP address.
<i>ipPort</i>	Specifies the UDP port number for the syslog target.
<i>vr_name</i>	Specifies the virtual router that can reach the server IP address.
	<div style="border: 1px solid black; padding: 5px;">  <p>Note User-created VRs are supported only on the platforms listed for this feature in Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i>.</p> </div>
local0 ... local7	Specifies the local syslog facility.
format	Specifies that the format for the target will be reset to the default value.

Default

When a target format is unconfigured, it is reset to the default values.

The following defaults apply to console display, memory buffer, NVRAM, and session targets:

- timestamp—hundredths.
- date—mm-dd-yyyy.
- severity—on.
- event-name—condition.
- host-name—off.
- sequence-number—off.
- process-name—off.



- process-slot—on (modular switches only).
- process-id—off.
- source-line—off.

The following defaults apply to syslog targets (per RFC 3164):

- timestamp—seconds.
- date—mmm-dd.
- severity—on.
- event-name—none.
- host-name—off.
- sequence-number—off.
- process-name—off.
- process-slot—on (modular switches only).
- process-id—off.
- source-line—off.

Usage Guidelines

Use this command to reset the target format to the default format.

Example

The following command sets the log format for the target session (the current session) to the default:

```
unconfigure log target session format
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

unconfigure ports monitor vlan

```
unconfigure ports [port_list |all] monitor vlan vlan_name
```

Description

Stops counting VLAN statistics on a port or group of ports.



Syntax Description

<i>port_list</i>	Specifies one or more ports. May be in the form: 1, 2, 3-5, 2:5, 2:6-2:8.
<i>vlan_name</i>	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

None.

Example

The following command removes monitoring for ports on a supported modular switch on the VLAN named accounting:

```
unconfigure ports 8:1-8:6 monitor vlan accounting
```

History

This command was first available in ExtremeXOS 12.0.

Support for BlackDiamond 8000 series modules, SummitStack, and Summit family switches was added in ExtremeXOS 12.5.

Platform Availability

This command is available only on BlackDiamond X8, BlackDiamond 8000 series modules, SummitStack and Summit family switches.

unconfigure sflow

unconfigure sflow

Description

Resets all the sFlow values to the default values.

Syntax Description

This command has no arguments or variables.



Default

The default values for sFlow are as follows:

- sFlow agent IP address—0.0.0.0.
- sampling frequency—sample one every 8196 packets.
- polling interval—20 seconds.
- maximum CPU sample limit—2000 samples per second.

sFlow is unconfigured and disabled on all ports.

Usage Guidelines

This command resets sFlow values to the default values, and removes any port configurations, and any sFlow collectors configured on the switch.

Example

The following command unconfigures sFlow:

```
unconfigure sflow
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

unconfigure sflow agent

```
unconfigure sflow agent
```

Description

Resets the sFlow agent's IP address to the default value.

Syntax Description

This command has no arguments or variables.

Default

The default IP address is 0.0.0.0.



Usage Guidelines

This command resets the sFlow agent IP address to its default value.

Example

The following command resets the agent IP back to the management IP address:

```
unconfigure sflow agent
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.


unconfigure sflow collector

```
unconfigure sflow collector {ipaddress} ipaddress {port udp-port-number} {vr  
vr_name}
```

Description

Unconfigures the sFlow collector.

Syntax Description

<i>ipaddress</i>	Specifies the IP address of the collector to reset.
<i>udp-port-number</i>	Specifies the UDP port.
<i>vr_name</i>	Specifies which virtual router.
 Note User-created VRs are supported only on the platforms listed for this feature in Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i> .	

Default

The following values are the defaults for this command:

- UDP port number—6343.
- Virtual router—VR-Mgmt (previously called VR-0).



Usage Guidelines

This command allows you to reset the specified sFlow collector parameters to the default values.

Both the commands `unconfigure ports monitor vlan` and `unconfigure sflow collector` will reset the collector parameters to the default.

Example

The following command removes the collector at IP address 192.168.57.1:

```
unconfigure sflow collector ipaddress 192.168.57.1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

unconfigure sflow ports

```
unconfigure sflow ports port_list
```

Description

Removes the specified ports from the sFlow configuration, and stops sampling them.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
------------------	---

Default

N/A.

Usage Guidelines

This command removes the specified ports from the sFlow configuration, and stops sampling them.



Example

The following command unconfigures sFlow on the ports 2:5-2:7:

```
unconfigure sflow ports 2:5-2:7
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

unconfigure xml-notification

unconfigure xml-notification

Description

Unconfigures the XML notification client.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to unconfigure the XML client process including the associated log target configuration.

Example

The following command unconfigures the xml-notification client:

```
unconfigure xml-notification
```

History

This command was first available in ExtremeXOS 12.4.



Platform Availability

This command is available on BlackDiamond X8, BlackDiamond 8000 series modules, and Summit Family switches.


upload log

```
upload log ipaddress {vr vr_name} filename {messages [memory-buffer | nvram]
{events {event-condition | event_component}}} {severity severity {only}} {match
regex} {chronological}
```

Description

Uploads the current log messages to a TFTP server.

Syntax Description

<i>ipaddress</i>	Specifies the ipaddress of the TFTP server.
<i>vr_name</i>	Specifies the virtual router that can reach the TFTP server.
	 <p>Note User-created VRs are supported only on the platforms listed for this feature in Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i>.</p>
<i>filename</i>	Specifies the file name for the log stored on the TFTP server.
messages	Specifies the location from which to display the log messages.
memory-buffer	Show messages stored in volatile memory.
nvram	Show messages stored in NVRAM.
events	Show event messages.
<i>event-condition</i>	Specifies the event condition to display.
<i>event-component</i>	Specifies the event component to display.
<i>severity</i>	Specifies the minimum severity level to display (if the keyword only is omitted).
only	Specifies that only the specified severity level is to be displayed.
<i>regex</i>	Specifies a regular expression. Only messages that match the regular expression will be displayed.
chronological	Specifies uploading log messages in ascending chronological order (oldest to newest).

Default

The following defaults apply:

- messages—memory buffer.
- severity—none (displays everything stored in the target).



- match—no restriction.
- chronological—if not specified, show messages in order from newest to oldest.

Usage Guidelines

This command is similar to the `show log` command, but instead of displaying the log contents on the command line, this command saves the log to a file on the TFTP server you specify. For more details on most of the options of this command, see the command `show log`.

Host Name and Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z).
- Numerals (0-9).
- Period (.).
- Dash (-) Permitted only for host names.
- Underscore (_) Permitted only for host names.
- Colon (:).

When naming or configuring an IP address for your network server, remember the requirements listed above.

Remote Filename Character Restrictions

This section provides information about the characters supported by the switch for remote filenames.

When specifying a remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z).
- Numerals (0-9).
- Period (.).
- Dash (-).
- Underscore (_).
- Slash (/).

When naming a local or remote file, remember the requirements listed above.

Example

The following command uploads messages with a critical severity to the filename `switch4critical.log` on TFTP server at `10.31.8.25`:

```
upload log 10.31.8.25 switch4critical.log critical
```



The following command uploads messages with warning, error, or critical severity to the filename switch4warn.log on TFTP server at 10.31.8.25:

```
upload log 10.31.8.25 switch4warn.log warning
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.



14 VLAN Commands

```
configure private-vlan add network
configure private-vlan add subscriber
configure private-vlan delete
configure protocol add
configure protocol delete
configure vlan add ports
configure vlan add ports private-vlan translated
configure vlan delete ports
configure vlan description
configure vlan ipaddress
configure vlan name
configure vlan protocol
configure vlan tag
configure vlan-translation add loopback-port
configure vlan-translation add member-vlan
configure vlan-translation delete loopback-port
configure vlan-translation delete member-vlan
create private-vlan
create protocol
create vlan
delete private-vlan
delete protocol
delete vlan
disable loopback-mode vlan
disable vlan
enable loopback-mode vlan
enable vlan
show private-vlan
show private-vlan <name>
show protocol
show vlan
show vlan description
unconfigure vlan description
unconfigure vlan ipaddress
```

This chapter describes commands for configuring and managing:

- VLANs.
- Private VLANs (PVLANS).

- VLAN translation.

For an introduction to VLAN features, see the *ExtremeXOS Concepts Guide*.

configure private-vlan add network

```
configure private-vlan name add network vlan_name
```

Description

Adds the specified VLAN as the network VLAN on the specified PVLAN.

Syntax Description

<i>name</i>	Specifies the name of the PVLAN to which the VLAN is added.
<i>vlan_name</i>	Specifies a VLAN to add to the PVLAN.

Default

N/A.

Usage Guidelines

The VLAN must be created and configured with a tag before it is added to the PVLAN.

Example

The following command adds VLAN sharednet as the network VLAN for the PVLAN named companyx:

```
configure private-vlan companyx add network sharednet
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms that support the Private VLAN feature. The features and the platforms that support them are listed in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure private-vlan add subscriber

```
configure private-vlan name add subscriber vlan_name {non-isolated} {loopback-  
port port}
```



Description

Adds the specified VLAN as a subscriber VLAN on the specified PVLAN.

Syntax Description

<i>name</i>	Specifies the name of the PVLAN to which the VLAN is added.
<i>vlan_name</i>	Specifies a VLAN to add to the PVLAN.
non-isolated	Configures the subscriber VLAN as a non-isolated subscriber VLAN.
<i>port</i>	Specifies the port that serves as the loopback port.

Default

If the non-isolated option is omitted, this command adds the specified VLAN as an isolated subscriber VLAN.

Usage Guidelines

The VLAN must be created and configured with a tag before it is added to the PVLAN. If the non-isolated option is omitted, the VLAN is added as an isolated subscriber VLAN. If the non-isolated option is included, the VLAN is added as a non-isolated subscriber VLAN.

The `loopback-port port` option is available only on BlackDiamond 8000 series modules and Summit family switches, whether or not included in a SummitStack. If two or more subscriber VLANs have overlapping ports (where the same ports are assigned to both VLANs), each of the subscriber VLANs with overlapping ports must have a dedicated loopback port.

Example

The following command adds VLAN restricted as a subscriber VLAN for the PVLAN named companyx:

```
configure private-vlan companyx add subscriber restricted isolated
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms that support the Private VLAN feature. For features and the platforms that support them, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure private-vlan delete

```
configure private-vlan name delete [network | subscriber] vlan_name
```



Description

Deletes the specified VLAN from the specified PVLAN.

Syntax Description

<i>name</i>	Specifies the name of the PVLAN from which the VLAN is deleted.
network	Specifies that the VLAN to be deleted is a network VLAN.
subscriber	Specifies that the VLAN to be deleted is a subscriber VLAN.
<i>vlan_name</i>	Specifies the VLAN to delete from the PVLAN.

Default

N/A.

Usage Guidelines

This command deletes a VLAN from a PVLAN, but it does not delete the VLAN from the system—it just breaks the link between the VLAN and the PVLAN. You can use this command to delete both network and subscriber VLANs.

Example

The following command deletes network VLAN sharednet from the PVLAN named companyx:

```
configure private-vlan companyx delete network sharednet
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms that support the Private VLAN feature. For features and the platforms that support them, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure protocol add

```
configure protocol name add [etype | llc | snap] hex {[etype | llc | snap] hex}
```

Description

Configures a user-defined protocol filter.



Syntax Description

<i>name</i>	Specifies a protocol filter name.
<i>hex</i>	Specifies a four-digit hexadecimal number between 0 and FFFF that represents: The Ethernet protocol type taken from a list maintained by the IEEE. The DSAP/SSAP combination created by concatenating a two-digit LLC Destination SAP (DSAP) and a two-digit LLC Source SAP (SSAP). The SNAP-encoded Ethernet protocol type.

Default

N/A.

Usage Guidelines

Supported protocol types include:

- `etype` - IEEE Ethertype.
- `llc` - LLC Service Advertising Protocol.
- `snap` - Ethertype inside an IEEE SNAP packet encapsulation.

A maximum of 15 protocol filters, each containing a maximum of six protocols, can be defined.

The protocol filter must already exist before you can use this command. Use the `create protocol` command to create the protocol filter.

No more than seven protocols can be active and configured for use.



Note

Protocol based VLAN for Etype from 0x0000 to 0x05ff are not classifying as per filter. When traffic arrive with these Etypes, it is classified to native VLAN rather protocol based vlan.

Example

The following command configures a protocol named Fred by adding protocol type LLC SAP with a value of FFEF:

```
configure protocol fred add llc 0xfeff
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.



configure protocol delete

```
configure protocol name delete [etype | llc | snap] hex {[etype | llc | snap]
hex} ...
```

Description

Deletes the specified protocol type from a protocol filter.

Syntax Description

<i>name</i>	Specifies a protocol filter name.
<i>hex</i>	Specifies a four-digit hexadecimal number between 0 and FFFF that represents: The Ethernet protocol type taken from a list maintained by the IEEE. The DSAP/SSAP combination created by concatenating a two-digit LLC Destination SAP (DSAP) and a two-digit LLC Source SAP (SSAP). The SNAP-encoded Ethernet protocol type.

Default

N/A.

Usage Guidelines

Supported protocol types include:

- **etype** - IEEE Ethertype.
- **llc** - LLC Service Advertising Protocol.
- **snap** - Ethertype inside an IEEE SNAP packet encapsulation.

Example

The following command deletes protocol type LLC SAP with a value of FEFF from protocol fred:

```
configure protocol fred delete llc feff
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.



configure vlan add ports

```
configure {vlan} vlan_name add ports [port_list | all] {tagged tag | untagged}
{{stp} stp_name} {dot1d | emistp | pvst-plus}}
```

Description

Adds one or more ports in a VLAN.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>port_list</i>	Specifies a list of ports or slots and ports.
all	Specifies all ports.
tagged tag	Specifies the ports should be configured as tagged.
untagged	Specifies the ports should be configured as untagged.
<i>stp_name</i>	Specifies an STP domain name.
dot1d emistp pvst-plus	Specifies the BPDU encapsulation mode for these STP ports.

Default

Untagged.

Usage Guidelines

The VLAN must already exist before you can add (or delete) ports: use the `create vlan` command to create the VLAN.

If the VLAN uses 802.1Q tagging, you can specify tagged or untagged port(s). If the VLAN is untagged, the ports cannot be tagged.

Untagged ports can only be a member of a single VLAN. By default, they are members of the default VLAN (named Default). In order to add untagged ports to a different VLAN, you must first remove them from the default VLAN. You do not need to do this to add them to another VLAN as tagged ports. If you attempt to add an untagged port to a VLAN prior to removing it from the default VLAN, you see the following error message:

```
Error: Protocol conflict when adding untagged port 1:2. Either add this port as tagged or assign another protocol to this VLAN.
```



Note

This print is not displayed if keyword "all" is used as `port_list`.



The ports that you add to a VLAN and the VLAN itself cannot be explicitly assigned to different virtual routers (VRs). When multiple VRs are defined, consider the following guidelines while adding ports to a VLAN:

- A VLAN can belong (either through explicit or implicit assignment) to only one VR.
- If a VLAN is not explicitly assigned to a VR, then the ports added to the VLAN must be explicitly assigned to a single VR.
- If a VLAN is explicitly assigned to a VR, then the ports added to the VLAN must be explicitly assigned to the same VR or to no VR.
- If a port is added to VLANs that are explicitly assigned to different VRs, the port must be explicitly assigned to no VR.

**Note**

User-created VRs are supported only on the platforms listed for this feature in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*. On switches that do not support user-created VRs, all VLANs are created in VR-Default and cannot be moved.

Refer to [STP Commands](#) for more information on configuring Spanning Tree Domains.

**Note**

If you use the same name across categories (for example, STPD and EAPS names), we recommend that you specify the identifying keyword as well as the actual name. If you do not use the keyword, the system may return an error message.

Beginning with ExtremeXOS version 11.4, the system returns the following message if the ports you are adding are already EAPS primary or EAPS secondary ports:

```
WARNING: Make sure Vlan1 is protected by EAPS, Adding EAPS ring ports to a VLAN could cause a loop in the network. Do you really want to add these ports? (y/n)
```

Example

The following command assigns tagged ports 1:1, 1:2, 1:3, and 1:6 to a VLAN named accounting:

```
configure vlan accounting add ports 1:1, 1:2, 1:3, 1:6 tagged
```

History

This command was first available in ExtremeXOS 10.1.

The **tagged** keyword was added in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.



configure vlan add ports private-vlan translated

Translation from network VLAN tag to each subscriber VLAN tag is done by default in a private VLAN.

```
configure {vlan} vlan_name add ports port_list private-vlan translated
```

Description

Adds the specified ports to the specified network VLAN and enables tag translation for all subscriber VLAN tags to the network VLAN tag.

Syntax Description

<i>vlan_name</i>	Specifies the network VLAN to which the ports are added.
<i>port_list</i>	Specifies the ports to be added to the network VLAN.

Default

N/A.

Usage Guidelines

This command is allowed only when the specified VLAN is configured as a network VLAN on a PVLAN.

Example

The following command adds port 2:1 to VLAN sharednet and enables VLAN translation on that port:

```
configure sharednet add ports 2:1 private-vlan translated
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms that support the Private VLAN feature. For features and the platforms that support them, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vlan delete ports

```
configure {vlan} vlan_name delete ports [all | port_list {tagged tag}]
```



Description

Deletes one or more ports in a VLAN.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
all	Specifies all ports.
<i>port_list</i>	A list of ports or slots and ports.
tagged tag	Specifies the port-specific VLAN tag. When there are multiple ports specified in the <i>port_list</i> , the same tag is used for all of them.

Default

When unspecified, the port tag is equal to the VLAN tag.

Usage Guidelines

Specify port tag to delete a VLAN port that has a different tag from the VLAN tag.

Example

The following command removes ports 1:1, 1:2, 4:3, and 5:6 on a modular switch from a VLAN named accounting:

```
configure accounting delete port 1:1, 1:2, 4:3, 5:6
```

The following example deletes a VLAN port with tag 10:

```
create vlan exchange tag 100
config vlan exchange del ports 3 tag 10
```

The following example deletes a VLAN port tag of 10 on two ports:

```
create vlan exchange tag 100
config vlan exchange d ports 3,4 tag 10
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.



configure vlan description

```
configure {vlan} vlan_name description [vlan-description | none]
```

Description

Configures a description for the specified VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN name.
<i>vlan-description</i>	Specifies a VLAN description (up to 64 characters) that appears in show vlan commands and can be read from the ifAlias MIB object for the VLAN.
none	This keyword removes the configured VLAN description.

Default

No description.

Usage Guidelines

The VLAN description must be in quotes if the string contains any space characters. If a VLAN description is configured for a VLAN that already has a description, the new description replaces the old description.

Example

The following command assigns the description Campus A to VLAN vlan1:

```
configure vlan vlan1 description "Campus A"
```

History

This command was first available in ExtremeXOS 12.4.4.

Platform Availability

This command is available on all platforms.

configure vlan ipaddress

```
configure {vlan} vlan_name ipaddress [ipaddress {ipNetmask} | ipv6-link-local |  
{eui64} ipv6_address_mask]
```



Description

Assigns an IPv4 address and an optional subnet mask or an IPv6 address to the VLAN. Beginning with ExtremeXOS version 11.2, you can specify IPv6 addresses. You can assign either an IPv4 address, and IPv6 address, or both to the VLAN. Beginning with ExtremeXOS software version 11.3, you can use this command to assign an IP address to a specified VMAN and enable multicasting on that VMAN.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>ipaddress</i>	Specifies an IPv4 address.
<i>ipNetmask</i>	Specifies an IPv4 subnet mask in dotted-quad notation (for example, 255.255.255.0).
ipv6-link-local	Specifies IPv6 and configures a link-local address generated by combining the standard link-local prefix with the automatically generated interface in the EUI-64 format. Using this option automatically generates an entire IPv6 address; this address is only a link-local, or VLAN-based, IPv6 address, that is, ports on the same segment can communicate using this IP address and do not have to pass through a gateway.
eui64	Specifies IPv6 and automatically generates the interface ID in the EUI-64 format using the interface's MAC address. Once you enter this parameter, you must add the following variables: <i>ipv6_address_mask</i> . Use this option when you want to enter the 64-bit prefix and use a EUI-64 address for the rest of the IPv6 address.
<i>ipv6_address_mask</i>	Specify the IPv6 address in the following format: x:x:x:x:x:x/prefix length, where each x is the hexadecimal value of one of the 8 16-bit pieces of the 128-bit wide address.

Default

N/A.

Usage Guidelines

Note



You can also use this command to assign an IP address to a VMAN on all platforms that support the VMAN feature. For information on which software licenses and platforms support the VMAN feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

The VLAN must already exist before you can assign an IP address: use the `create vlan` command to create the VLAN (also the VMAN must already exist).

Note



If you plan to use the VLAN as a control VLAN for an EAPS domain, do NOT configure the VLAN with an IP address.

See [IP Unicast Commands](#) for information on adding secondary IP addresses to VLANs.

Beginning with ExtremeXOS software version 11.2, you can specify IPv6 addresses. See [IPv6 Unicast Routing](#) for information on IPv6 addresses.



Beginning with ExtremeXOS software version 11.3, you can assign an IP address (including IPv6 addresses) to a VLAN. Beginning with version 11.4, you can enable multicasting on that VLAN.

To enable multicasting on the specified VLAN once you assigned an IP address, take the following steps:

- Enable IP multicast forwarding.
- Enable and configure multicasting.

Example

The following commands are equivalent; both assign an IPv4 address of 10.12.123.1 to a VLAN named accounting:

```
configure vlan accounting ipaddress 10.12.123.1/24
```

```
configure vlan accounting ipaddress 10.12.123.1 255.255.255.0
```

The following command assigns a link local IPv6 address to a VLAN named management:

```
configure vlan accounting ipaddress ipv6-link-local
```

History

This command was first available in ExtremeXOS 10.1.

The IPv6 parameters were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure vlan name

```
configure {vlan} vlan_name name name
```

Description

Renames a previously configured VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the current (old) VLAN name.
<i>name</i>	Specifies a new name for the VLAN.



Default

N/A.

Usage Guidelines

You cannot change the name of the default VLAN “Default.”

For information on VLAN name requirements and a list of reserved keywords, see [Object Names](#) in the *ExtremeXOS Concepts Guide*.



Note

If you use the same name across categories (for example, STPD and EAPS names), we recommend that you specify the identifying keyword as well as the actual name. If you do not use the keyword, the system may return an error message.

Example

The following command renames VLAN vlan1 to engineering:

```
configure vlan vlan1 name engineering
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure vlan protocol

```
configure {vlan} vlan_name protocol protocol_name
```

Description

Configures a VLAN to use a specific protocol filter.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>protocol_name</i>	Specifies a protocol filter name. This can be the name of a predefined protocol filter, or one you have defined. The following protocol filters are predefined: IPIPv6IPXNetBIOSDECNetIPX_8022IPX_SNAPAppleTalk any indicates that this VLAN should act as the default VLAN for its member ports.



Default

Protocol any.

Usage Guidelines

If the keyword any is specified, all packets that cannot be classified into another protocol-based VLAN are assigned to this VLAN as the default for its member ports.

Use the configure protocol command to define your own protocol filter.

Protocol Filters on BlackDiamond 8800 Series Switches, SummitStack, and the Summit Family Switches Only

These devices do not forward packets with a protocol-based VLAN set to AppleTalk. To ensure that AppleTalk packets are forwarded on the device, create a protocol-based VLAN set to “any” and define other protocol-based VLANs for other traffic, such as IP traffic. The AppleTalk packets pass on the “any” VLAN, and the other protocols pass traffic on their specific protocol-based VLANs.

Example

The following command configures a VLAN named accounting as an IP protocol-based VLAN:

```
configure accounting protocol ip
```

History

This command was first available in ExtremeXOS 10.1.

The IPv6 parameter was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure vlan tag

```
configure {vlan} vlan_name tag tag {remote-mirroring}
```

Description

Assigns a unique 802.1Q tag to the VLAN.



Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>tag</i>	Specifies a value to use as an 802.1Q tag. The valid range is from 2 to 4095.
remote-mirroring	Specifies that the tagged VLAN is for remote mirroring.

Default

The default VLAN uses an 802.1Q tag (and an internal VLANid) of 1.

Usage Guidelines

If any of the ports in the VLAN use an 802.1Q tag, a tag must be assigned to the VLAN. The valid range is from 2 to 4094 (tag 1 is assigned to the default VLAN, and tag 4095 is assigned to the management VLAN).

The 802.1Q tag is also used as the internal VLANid by the switch.

You can specify a value that is currently used as an internal VLANid on another VLAN; it becomes the VLANid for the VLAN you specify, and a new VLANid is automatically assigned to the other untagged VLAN.

Example

The following command assigns a tag (and internal VLANid) of 120 to a VLAN named accounting:

```
configure accounting tag 120
```

History

This command was first available in ExtremeXOS 10.1.

The remote-mirroring option was added in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

configure vlan-translation add loopback-port

```
configure {vlan} vlan_name vlan-translation add loopback-port port
```

Description

Adds the specified port as a loopback port for the specified member VLAN.



Syntax Description

<code>vlan_name</code>	Specifies the name of the member VLAN to which you want to add the loopback port.
<code>port</code>	Specifies the port that serves as the loopback port.

Default

N/A.

Usage Guidelines

The `loopback-port port` option is available only on BlackDiamond 8000 series modules and Summit family switches, whether or not included in a SummitStack. If two or more member VLANs have overlapping ports (where the same ports are assigned to both VLANs), each of the member VLANs with overlapping ports must have a dedicated loopback port.

The loopback port can be added to the member VLAN when the member VLAN is created, or you can use this command to add the loopback port at a later time.

Example

The following command adds port 2:1 as a loopback port for the member VLAN leafvlan:

```
configure leafvlan vlan-translation add loopback-port 2:1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms that support the VLAN Translation feature. For features and the platforms that support them, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vlan-translation add member-vlan

```
configure {vlan} vlan_name vlan-translation add member-vlan member_vlan_name
{loopback-port port}
```

Description

Adds a member VLAN to a translation VLAN.



Syntax Description

<i>vlan_name</i>	Specifies the name of the translation VLAN to which you want to add the member VLAN.
<i>member_vlan_name</i>	Specifies the member VLAN to be added to the translation VLAN.
<i>port</i>	Specifies the port that serves as the loopback port.

Default

N/A.

Usage Guidelines

This command configures VLAN tag translation between the two VLANs specified. The member VLAN is added to the list maintained by translation VLAN. A translation VLAN can have multiple member VLANs added to it.

The `loopback-port port` option is available only on BlackDiamond 8000 series modules and Summit family switches, whether or not included in a SummitStack. If two or more member VLANs have overlapping ports (where the same ports are assigned to both VLANs), each of the member VLANs with overlapping ports must have a dedicated loopback port.

Example

The following command adds member VLAN leafvlan to the translation VLAN branchvlan:

```
configure branchvlan vlan-translation add member-vlan leafvlan
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms that support the VLAN Translation feature. For features and the platforms that support them, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vlan-translation delete loopback-port

```
configure {vlan} vlan_name vlan-translation delete loopback-port
```

Description

Deletes the loopback port from the specified member VLAN.



Syntax Description

<i>vlan_name</i>	Specifies the name of the member VLAN from which you want to delete the loopback port.
------------------	--

Default

N/A.

Usage Guidelines

This command disables and deletes the loopback port from the specified member VLAN. This command does not delete the member VLAN.

Example

The following command deletes the loopback port from the member VLAN leafvlan:

```
configure leafvlan vlan-translation delete loopback-port
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms that support the VLAN Translation feature. For features and the platforms that support them, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vlan-translation delete member-vlan

```
configure {vlan} vlan_name vlan-translation delete member-vlan [member_vlan_name | all]
```

Description

Deletes one or all member VLANs from a translation VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the name of the translation VLAN from which you want to delete the member VLAN.
<i>member_vlan_name</i>	Specifies the member VLAN to be deleted from the translation VLAN.
all	Deletes all member VLANs from the specified translation VLAN.



Default

N/A.

Usage Guidelines

This command removes the link between the translation VLAN and the specified member VLANs, but it does not remove the VLANs from the switch.

Example

The following command deletes member VLAN leafvlan from the translation VLAN branchvlan:

```
configure branchvlan vlan-translation delete member-vlan leafvlan
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms that support the VLAN Translation feature. For features and the platforms that support them, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

create private-vlan

```
create private-vlan name {vr vr_name}
```

Description

Creates a PVLAN framework with the specified name.

Syntax Description

<i>name</i>	Specifies a name for the new PVLAN.
<i>vr_name</i>	Specifies the VR in which the PVLAN is created.

Default

N/A.

Usage Guidelines

The PVLAN is a framework that links network and subscriber VLANs; it is not an actual VLAN.



A private VLAN name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For private VLAN naming guidelines and a list of reserved names, see [Object Names](#) in the *ExtremeXOS Concepts Guide*.

If no VR is specified, the PVLAN is created in the default VR context.

Example

The following command creates a PVLAN named companyx:

```
create private-vlan companyx
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms that support the Private VLAN feature. For features and the platforms that support them, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

create protocol

```
create protocol name
```

Description

Creates a user-defined protocol filter.

Syntax Description

<i>name</i>	Specifies a protocol filter name. The protocol filter name can have a maximum of 31 characters.
-------------	---

Default

N/A.

Usage Guidelines

Protocol-based VLANs enable you to define packet filters that the switch can use as the matching criteria to determine if a particular packet belongs to a particular VLAN.

After you create the protocol, you must configure it using the `configure protocol` command. To assign it to a VLAN, use the `configure {vlan} vlan_name protocolprotocol_name` command.



Example

The following command creates a protocol named fred:

```
create protocol fred
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

create vlan

```
create vlan vlan_name {description vlan-description} {vr name}
```

Description

Creates a named VLAN.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name (up to 32 characters).
<i>vlan-description</i>	Specifies a VLAN description (up to 64 characters) that appears in show vlan commands and can be read from the ifAlias MIB object for the VLAN.
<i>name</i>	Specifies a VR or virtual routing and forwarding (VRF) instance in which to create the VLAN. NOTE: User-created VRs are supported only on the platforms listed for this feature in Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i> . On switches that do not support user-created VRs, all VLANs are created in VR-Default and cannot be moved.

Default

A VLAN named Default exists on all new or initialized Extreme switches:

- It initially contains all ports on a new or initialized switch, except for the management port(s), if there are any.
- It has an 802.1Q tag of 1.
- The default VLAN is untagged on all ports.
- It uses protocol filter any.

A VLAN named Mgmt exists on switches that have management modules or management ports:

- It initially contains the management port(s) the switch.
- It is assigned the next available internal VLANid as an 802.1Q tag.



If you do not specify the VR, the VLAN is created in the current VR.

If the VLAN description contains one or more space characters, you must enclose the complete name in double quotation marks.

Usage Guidelines

A newly-created VLAN has no member ports, is untagged, and uses protocol filter any until you configure it otherwise. Use the various configure vlan commands to configure the VLAN to your needs.

Internal VLANids are assigned automatically using the next available VLANid starting from the high end (4094) of the range.

The VLAN name can include up to 32 characters. VLAN names must begin with an alphabetical letter, and only alphanumeric, underscore (_), and hyphen (-) characters are allowed in the remainder of the name. VLAN names cannot match reserved keywords. For more information on VLAN name requirements and a list of reserved keywords, see [Object Names](#) in the *ExtremeXOS Concepts Guide*.

Note



If you use the same name across categories (for example, STPD and EAPS names), we recommend that you specify the identifying keyword as well as the actual name. If you do not use the keyword, the system may return an error message.

VLAN names are locally significant. That is, VLAN names used on one switch are only meaningful to that switch. If another switch is connected to it, the VLAN names have no significance to the other switch.

You must use mutually exclusive names for:

- VLANs.
- VMANs.
- Ipv6 tunnels.
- BVLANS.
- SVLANS.
- CVLANS



Note

The VLAN description is stored in the ifAlias MIB object.

If you do not specify a VR when you create a VLAN, the system creates that VLAN in the default VR (VR-Default). The management VLAN is always in the management VR (VR-Mgmt).

Once you create VRs, ExtremeXOS software allows you to designate one of these as the domain in which all your subsequent configuration commands, including VLAN commands, are applied. If you create VRs, ensure that you are creating the VLANs in the desired virtual-router domain.

Note



User-created VRs are supported only on the platforms listed for this feature in [ExtremeXOS Software Licenses](#) in the *ExtremeXOS Concepts Guide*. On switches that do not support user-created VRs, all VLANs are created in VR-Default and cannot be moved.



Example

The following command creates a VLAN named accounting on the current VR:

```
create vlan accounting description "Accounting Dept"
```

History

This command was first available in ExtremeXOS 10.1.

The VR option `vr`, was added in ExtremeXOS 11.0.

The VLAN description option was added in ExtremeXOS 12.4.4.

Platform Availability

This command is available on all platforms.

delete private-vlan

```
delete private-vlan name
```

Description

Deletes the PVLAN framework with the specified name.

Syntax Description

<i>name</i>	Specifies the name of the PVLAN to be deleted.
-------------	--

Default

N/A.

Usage Guidelines

The PVLAN is a framework that links network and subscriber VLANs; it is not an actual VLAN.

This command deletes the PVLAN framework, but it does not delete the associated VLANs. If the ports in the network VLAN were set to translate, they are changed to tagged.

Example

The following command deletes the PVLAN named companyx:

```
delete private-vlan companyx
```



History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms that support the Private VLAN feature. For features and the platforms that support them, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

delete protocol

```
delete protocol name
```

Description

Deletes a user-defined protocol.

Syntax Description

<i>name</i>	Specifies a protocol name.
-------------	----------------------------

Default

N/A.

Usage Guidelines

If you delete a protocol that is in use by a VLAN, the protocol associated with that VLAN becomes none.

Example

The following command deletes a protocol named fred:

```
delete protocol fred
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.



delete vlan

delete vlan *vlan_name*

Description

Deletes a VLAN.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
------------------	------------------------

Default

N/A.

Usage Guidelines

If you delete a VLAN that has untagged port members and you want those ports to be returned to the default VLAN, you must add them back explicitly using the `configure svlan delete ports` command.



Note

The default VLAN cannot be deleted. Before deleting an ISC VLAN, you must delete the MLAG peer.

Example

The following command deletes the VLAN accounting:

```
delete accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable loopback-mode vlan

disable loopback-mode vlan *vlan_name*



Description

Disallows a VLAN to be placed in the UP state without an external active port. This allows (disallows) the VLANs routing interface to become active.

Syntax Description

<code>vlan_name</code>	Specifies a VLAN name.
------------------------	------------------------

Default

N/A.

Usage Guidelines

Use this command to specify a stable interface as a source interface for routing protocols. This decreases the possibility of route flapping, which can disrupt connectivity.

Example

The following command disallows the VLAN accounting to be placed in the UP state without an external active port:

```
disable loopback-mode vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable vlan

```
disable vlan vlan_name
```

Description

Use this command to disable the specified VLAN.

Syntax Description

<code>vlan_name</code>	Specifies the VLAN you want to disable.
------------------------	---



Default

Enabled.

Usage Guidelines

This command allows you to administratively disable specified VLANs. The following guidelines apply to working with disabling VLANs:

- Disabling a VLAN stops all traffic on all ports associated with the specified VLAN.
- You cannot disable a VLAN that is running Layer 2 protocol control traffic for protocols such as EAPS, STP, and ESRP.

When you attempt to disable a VLAN running Layer 2 protocol control traffic, the system returns a message similar to the following:

```
VLAN accounting cannot be disabled because it is actively use by an L2 Protocol
```

- You can disable the default VLAN; ensure that this is necessary prior to disabling the default VLAN.
- You cannot disable the management VLAN.
- You cannot bind Layer 2 protocols to a disabled VLAN.
- You can add ports to and delete ports from a disabled VLAN.

Example

The following command disables the VLAN named accounting:

```
disable vlan accounting
```

History

This command was first available in ExtremeXOS 11.4.

The ability to add ports to a disabled VLAN was added in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

enable loopback-mode vlan

```
enable loopback-mode vlan vlan_name
```

Description

Allows a VLAN to be placed in the UP state without an external active port. This allows (disallows) the VLANs routing interface to become active.



Syntax Description

<code>vlan_name</code>	Specifies a VLAN name.
------------------------	------------------------

Default

N/A.

Usage Guidelines

Use this command to specify a stable interface as a source interface for routing protocols. This decreases the possibility of route flapping, which can disrupt connectivity.

Example

The following command allows the VLAN accounting to be placed in the UP state without an external active port:

```
enable loopback-mode vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable vlan

```
enable vlan vlan_name
```

Description

Use this command to re-enable a VLAN that you previously disabled.

Syntax Description

<code>vlan_name</code>	Specifies the VLAN you want to disable.
------------------------	---

Default

Enabled.



Usage Guidelines

This command allows you to administratively enable specified VLANs that you previously disabled.

Example

The following command enables the VLAN named accounting:

```
enable vlan accounting
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

show private-vlan

show private-vlan

Description

Displays information about all the PVLANS on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

If the PVLAN is incomplete because it does not have a network or any subscriber VLAN configured, [INCOMPLETE] appears next to the PVLAN name.

Example

The following command displays all the PVLANS on the switch:

```
* (debug) BD-8808.1 # show private-vlan
```

```
-----  
-----
```



```

Name          VID  Protocol Addr      Flags          Proto  Ports
Virtual
Active router
/Total
-----
-----
Engineering
Network VLAN:
-Engr1        10  -----          ANY    4 /5  VR-
Default
Non-Isolated Subscriber VLAN:
-ni1          400 -----          ANY    1 /1  VR-
Default
-ni2          401 -----          ANY    1 /1  VR-
Default
Isolated Subscriber VLAN:
-il           500 -----          ANY    1 /1  VR-
Default
Ops
Network VLAN:
-Ops          20  -----          ANY    2 /2  VR-
Default
Non-Isolated Subscriber VLAN:
-OpsNi1       901 -----          ANY    1 /1  VR-
Default
-OpsNi2       902 -----          ANY    1 /1  VR-
Default
-OpsNi3       903 -----          ANY    1 /1  VR-
Default
-OpsNi4       904 -----          ANY    1 /1  VR-
Default
Isolated Subscriber VLAN:
-OpsI0        600 -----          ANY    1 /1  VR-
Default
-OpsI1        601 -----          ANY    1 /1  VR-
Default
-OpsI2        602 -----          ANY    1 /1  VR-
Default
-OpsI3        603 -----          ANY    1 /1  VR-
Default
-OpsI4        604 -----          ANY    1 /1  VR-
Default
Sales [INCOMPLETE]
Network VLAN:
-NONE
Non-Isolated Subscriber VLAN:
-SalesNi1     701 -----          ANY    1 /1  VR-
Default
-SalesNi2     702 -----          ANY    1 /1  VR-
Default
Isolated Subscriber VLAN:
-SalesI0      800 -----          ANY    1 /1  VR-
Default
-----
-----

```

```

Flags : (C) EAPS Control vlan, (d) NetLogin Dynamically created VLAN,
(D) VLAN Admin Disabled, (E) ESRP Enabled, (f) IP Forwarding Enabled,
(i) ISIS Enabled, (I) IP Forwarding lpm-routing Enabled, (L) Loopback Enabled,

```



```
(l) MPLS Enabled, (m) IPmc Forwarding Enabled, (n) IP Multinetting Enabled,
(N) Network LogIn vlan, (o) OSPF Enabled, (p) PIM Enabled,
(P) EAPS protected vlan, (r) RIP Enabled,
(T) Member of STP Domain, (V) VPLS Enabled, (v) VRRP Enabled
Total number of PVLAN(s) : 3
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms that support the Private VLAN feature. For features and the platforms that support them, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show private-vlan <name>

```
show {private-vlan} name
```

Description

Displays information about the specified PVLAN.

Syntax Description

<i>name</i>	Specifies the name of the PVLAN to display.
-------------	---

Default

N/A.

Usage Guidelines

If the PVLAN is incomplete because it does not have a network or any subscriber VLAN configured, [INCOMPLETE] appears next to the PVLAN name.

Example

The following command displays information for the companyx PVLAN:

```
* (debug) BD-8808.1 # show private-vlan "Engineering"
-----
-----
Name          VID  Protocol Addr          Flags          Proto  Ports
Virtual
Active router
/Total
```



```

-----
-----
Engineering
Network VLAN:
-Engr1      10  ----- ANY      4 /5  VR-
Default
Non-Isolated Subscriber VLAN:
-ni1       400 ----- ANY      1 /1  VR-
Default
-ni2       401 ----- ANY      1 /1  VR-
Default
Isolated Subscriber VLAN:
-il        500 ----- ANY      1 /1  VR-
Default
-----
-----
Flags : (C) EAPS Control vlan, (d) NetLogin Dynamically created VLAN,
(D) VLAN Admin Disabled, (E) ESRP Enabled, (f) IP Forwarding Enabled,
(i) ISIS Enabled, (I) IP Forwarding lpm-routing Enabled, (L) Loopback Enabled,
(l) MPLS Enabled, (m) IPmc Forwarding Enabled, (n) IP Multinetting Enabled,
(N) Network LogIn vlan, (o) OSPF Enabled, (p) PIM Enabled,
(P) EAPS protected vlan, (r) RIP Enabled,
(T) Member of STP Domain, (V) VPLS Enabled, (v) VRRP Enabled

```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms that support the Private VLAN feature. For features and the platforms that support them, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show protocol

```
show protocol {name}
```

Description

Displays protocol filter definitions.

Syntax Description

<i>name</i>	Specifies a protocol filter name.
-------------	-----------------------------------

Default

Displays all protocol filters.



Usage Guidelines

Displays the defined protocol filter(s) with the types and values of its component protocols.

Example

The following is an example of the show protocol command:

Protocol Name	Type	Value
IP	etype	0x0800
etype 0x0806		
ANY	ANY	0xffff
ipx	etype	0x8137
decnet	etype	0x6003
etype 0x6004		
netbios	llc	0xf0f0
llc 0xf0f1		
ipx_8022	llc	0xe0e0
ipx_snap	snap	0x8137
appletalk	snap	0x809b
snap 0x80f3		

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show vlan

```
show vlan {virtual-router vr-name}show {vlan} vlan_name {ipv4 | ipv6}show vlan
[tag tag | detail] {ipv4 | ipv6}show vlan ports
```

Description

Displays information about one or all VLANs.



Syntax Description

vr-name	Specifies a VR name for which to display summary information for all VLANs. If no VR name is specified, the software displays summary information for all VLANs in the current VR context. NOTE: User-created VRs are supported only on the platforms listed for this feature in Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i> . On switches that do not support user-created VRs, all VLANs are created in VR-Default and cannot be moved.
vlan_name	Specifies a VLAN name for which to display detailed VLAN information.
tag	Specifies the 802.1Q tag of a VLAN for which to display detailed VLAN information.
detail	Specifies that detailed information should be displayed for all VLANs.
ipv4	Specifies IPv4.
ipv6	Specifies IPv6.
ports	Displays VLAN ports information.

Default

Summary information for all VLANs on the device.

Usage Guidelines



Note

To display IPv6 information, you must issue either the `show vlan detail` command or `show vlan` command with the name of the specified VLAN.

Unlike many other VLAN-related commands, the keyword `vlan` is required in all forms of this command except when requesting information for a specific VLAN.

Use the command `show vlan` to display summary information for all VLANs. It shows various configuration options as a series of flags (see the example below). VLAN names, descriptions, and protocol names may be abbreviated in this display.

Use the command `show vlan detail` to display detailed information for all VLANs. This displays the same information as for an individual VLAN, but shows every VLAN, one-by-one. After each VLAN display you can elect to continue or quit.

Protocol `none` indicates that this VLAN was configured with a user-defined protocol that has subsequently been deleted.



Note

The BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches display the Mgmt VLAN in VR-Mgmt.

When an IPv6 address is configured for the VLAN, the system may display one of the following two address types in parentheses after the IPv6 address:

- Tentative.



- Duplicate.



Note

See the ExtremeXOS Concepts Guide for information on IPv6 address types.

You can display additional useful information on VLANs configured with IPv6 addresses by issuing the `show ipconfig ipv6 vlan vlan_name` command.

When a displayed VLAN is part of a PVLAN, the display includes the PVLAN name and type (which is network, non-isolated subscriber, or isolated subscriber).

When the displayed VLAN is configured for VLAN translation, the display provides translation VLAN information. If the displayed VLAN is a translation VLAN, a list of translation VLAN members appears. If the displayed VLAN is a member VLAN, the display indicates the translation VLAN to which the member VLAN belongs.

Example

The following is an example of the `show vlan` command on a switch where PTP and CES are configured (for example, an E4G-200 or E4G-400):

```
E4G-400.15 # sh vlan
-----
Name                VID  Protocol Addr          Flags                Proto  Ports
Virtual
Active router
/Total
-----
Default             1    -----T----- ANY    1 /34
VR-Default
Mgmt                4095 ----- ANY    1 /1
VR-Mgmt
v1                  40   1.1.1.51      /24 -fL-----ek ANY    1 /10
VR-Default
v2                  20   1.1.2.52      /24 -f-----e- ANY    0 /1
VR-Default
-----
Flags : (B) BFD Enabled, (c) 802.1ad customer VLAN, (C) EAPS Control VLAN,
(d) NetLogin Dynamically created VLAN, (D) VLAN Admin Disabled,
(e) CES Configured, (E) ESRP Enabled, (f) IP Forwarding Enabled,
(F) Learning Disabled, (i) ISIS Enabled, (I) Inter-Switch Connection VLAN for
MLAG,
(k) PTP Configured, (l) MPLS Enabled, (L) Loopback Enabled,
(m) IPmc Forwarding Enabled, (M) Translation Member VLAN or Subscriber VLAN,
(n) IP Multinetting Enabled, (N) Network Login VLAN, (o) OSPF Enabled,
(O) Flooding Disabled, (p) PIM Enabled, (P) EAPS protected VLAN,
(r) RIP Enabled, (R) Sub-VLAN IP Range Configured,
(s) Sub-VLAN, (S) Super-VLAN, (t) Translation VLAN or Network VLAN,
(T) Member of STP Domain, (v) VRRP Enabled, (V) VPLS Enabled, (W) VPWS Enabled
Total number of VLAN(s) : 4
```



The following sample output shows OpenFlow status:

```
E4G-200.5 # show vlan
-----
Name          VID  Protocol Addr      Flags          Proto
Ports  Virtual Active router  /Total
-----
Default          1  -----
ANY                                0/0
VR-Default  ext  4094 -----
ANY                                0 /12
VR-Default  Mgmt 4095 -----
ANY                                1/1
VR-Mgmt
-----

Flags : (B) BFD Enabled, (c) 802.1ad customer VLAN, (C) EAPS Control VLAN,
        (d) Dynamically created VLAN, (D) VLAN Admin Disabled,
        (e) CES Configured, (E) ESRP Enabled, (f) IP Forwarding Enabled,

        (F) Learning Disabled, (i) ISIS Enabled, (I) Inter-Switch Connection
VLAN for MLAG,
        (k) PTP Configured, (l) MPLS Enabled, (L) Loopback Enabled,
        (m) IPmc Forwarding Enabled, (M) Translation Member VLAN or
Subscriber VLAN,
        (n) IP Multinetting Enabled, (N) Network Login VLAN, (o) OSPF
Enabled,
        (O) Flooding Disabled, (p) PIM Enabled, (P) EAPS protected VLAN,

        (r) RIP Enabled, (R) Sub-VLAN IP Range Configured,
        (s) Sub-VLAN, (S) Super-VLAN, (t) Translation VLAN or Network VLAN,

        (T) Member of STP Domain, (v) VRRP Enabled, (V) VPLS Enabled, (W)
VPWS Enabled
        (Z) Openflow Enabled

Total number of VLAN(s) : 3
```

The following sample output shows detailed OpenFlow status:

```
E4G-200.7 # show vlan detail
VLAN Interface with name Default created by user
Admin State:      Enabled          Tagging:      802.1Q Tag 1
Description:      None
Virtual router:  VR-Default
IPv4 Forwarding: Disabled
IPv4 MC Forwarding: Disabled
IPv6 Forwarding: Disabled
IPv6 MC Forwarding: Disabled
IPv6:             None
STPD:             s0(Disabled,Auto-bind)
Protocol:         Match all unfiltered protocols
Loopback:         Disabled
NetLogin:         Disabled
QosProfile:       None configured
Egress Rate Limit Designated Port: None configured
```




```

Flood Rate Limit QosProfile:      None configured
Ports:  0.      (Number of active ports=0)

#
#
VLAN Interface with name ext created by user
Admin State:      Enabled
Tagging:Untagged (Internal tag 4094)
Description:      None
Virtual router:   VR-Default
IPv6 Forwarding: Disabled
IPv6:             None
STPD:            None
Protocol:         Match all unfiltered protocols
Loopback:         Disabled
NetLogin:         Disabled
QosProfile:       None configured
Openflow:         Enabled
Egress Rate Limit Designated Port: None configured
Flood Rate Limit QosProfile:      None configured
Ports:  12.      (Number of active ports=0)
      Untag:      1,      2,      3,      4,      5,      6,      7,
                  8,      9,      10,     11,     12
Flags:  (*) Active, (!) Disabled, (g) Load Sharing port
        (b) Port blocked on the vlan, (m) Mac-Based port
        (a) Egress traffic allowed for NetLogin
        (u) Egress traffic unallowed for NetLogin
        (t) Translate VLAN tag for Private-VLAN
        (s) Private-VLAN System Port, (L) Loopback port
        (x) VMAN Tag Translated port
        (G) Multi-switch LAG Group port
        (H) Dynamically added by MVRP

#
#
VLAN Interface with name Mgmt created by user
Admin State:      Enabled
Tagging:          802.1Q Tag 4095
Description:      Management VLAN
Virtual router:   VR-Mgmt
IPv4 Forwarding: Disabled
IPv6 Forwarding: Disabled
IPv6:             None
STPD:            None
Protocol:         Match all unfiltered protocols
Loopback:         Disabled
NetLogin:         Disabled
QosProfile:       None configured
Flood Rate Limit QosProfile:      None configured
Ports:  1.      (Number of active ports=1)
Untag: Mgmt-port on Mgmt is active

```

The following example displays VLAN ports information:

```

show vlan ports 1,2,3,4,5,6,7,8,9,10,11,12
-----
Name          VID      Protocol Addr      Flags
Proto Ports  Virtual  Active router  /Total

```



```

-----
ext          4094
-----
ANY    0 /12  VR-Default
-----
Flags : (B) BFD Enabled, (c) 802.1ad customer VLAN, (C) EAPS Control VLAN,
(d) Dynamically created VLAN, (D) VLAN Admin Disabled,
(e) CES Configured, (E) ESRP Enabled, (f) IP Forwarding Enabled,
(F) Learning Disabled, (i) ISIS Enabled, (I) Inter-Switch Connection VLAN for
MLAG,
(k) PTP Configured, (l) MPLS Enabled, (L) Loopback Enabled,
(m) IPmc Forwarding Enabled, (M) Translation Member VLAN or Subscriber VLAN,

(n) IP Multinetting Enabled, (N) Network Login VLAN, (o) OSPF Enabled,

(O) Flooding Disabled, (p) PIM Enabled, (P) EAPS protected VLAN,
(r) RIP Enabled, (R) Sub-VLAN IP Range Configured,
(s) Sub-VLAN, (S) Super-VLAN, (t) Translation VLAN or Network VLAN,
(T) Member of STP Domain, (v) VRRP Enabled, (V) VPLS Enabled, (W) VPWS
Enabled
(Z) Openflow Enabled

Total number of VLAN(s) : 3 (1 displayed)

show vlan ports 1 detail
VLAN Interface with name ext created by user
Admin State:      Enabled
Tagging:Untagged (Internal tag 4094)
Description:      None
Virtual router:  VR-Default
IPv4 Forwarding: Disabled
IPv4 MC Forwarding: Disabled
IPv6 Forwarding: Disabled
IPv6 MC Forwarding: Disabled
IPv6:             None
STPD:             None
Protocol:         Match all unfiltered protocols
Loopback:         Disabled
NetLogin:         Disabled
QosProfile:       None configured
Openflow:         Enabled Egress Rate Limit Designated Port: None configured

Flood Rate Limit QosProfile:      None configured
Ports:  12.      (Number of active ports=0)
  Untag:    1,    2,    3,    4,    5,    6,    7,
           8,    9,   10,   11,   12
Flags:    (*) Active, (!) Disabled, (g) Load Sharing port

(b) Port blocked on the vlan, (m) Mac-Based port
(a) Egress traffic allowed for NetLogin
(u) Egress traffic unallowed for NetLogin
(t) Translate VLAN tag for Private-VLAN
(s) Private-VLAN System Port, (L) Loopback port
(x) VMAN Tag Translated port

```



- (G) Multi-switch LAG Group port
- (H) Dynamically added by MVRP

The following example is the show output of a vlan that has port-specific tag. The tag is displayed in parentheses.

```
VLAN Interface with name v11 created by user
  Admin State:      Enabled      Tagging:   802.1Q Tag 100
  Description:      None
  Virtual router:   VR-Default
  IPv4 Forwarding:  Disabled
  IPv4 MC Forwarding: Disabled
  IPv6 Forwarding:  Disabled
  IPv6 MC Forwarding: Disabled
  IPv6:             None
  STPD:             None
  Protocol:         Match all unfiltered protocols
  Loopback:         Disabled
  NetLogin:         Disabled
  OpenFlow:         Disabled
  QosProfile:       None configured
  Egress Rate Limit Designated Port: None configured
  Flood Rate Limit QosProfile:       None configured
  Ports:           8.             (Number of active ports=2)
    Untag:         5
    Tag:           1,           10,           11
    Port-specific Tag:           1(0010),           1(0011),
*3(0101),
                               *4(0102)

  Flags:           (*) Active, (!) Disabled, (g) Load Sharing
port
                  (b) Port blocked on the vlan, (m) Mac-Based
port
                  (a) Egress traffic allowed for NetLogin
NetLogin          (u) Egress traffic unallowed for
                  (t) Translate VLAN tag for Private-VLAN
port              (s) Private-VLAN System Port, (L) Loopback
                  (e) Private-VLAN End Point Port
                  (x) VMAN Tag Translated port
                  (G) Multi-switch LAG Group port
                  (H) Dynamically added by MVRP
                  (U) Dynamically added uplink port
                  (V) Dynamically added by VM Tracking
```

History

This command was first available in ExtremeXOS 10.1.

The IPv6 information was added in ExtremeXOS 11.2.

The netlogin information was added in ExtremeXOS 11.3.

The VR and administratively enabled/disabled information was added in ExtremeXOS 11.4.



The tag option was added in ExtremeXOS 12.4.4.

The OpenFlow status feature was added in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

Information on MAC-based ports is available only on the Summit family of switches, SummitStack, and the BlackDiamond 8800 series switch.

show vlan description

show vlan description

Description

Displays a list of VLANs and VLAN descriptions.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the descriptions for all VLANs:

```
* Switch.4 # show vlan description
-----
Name          VID  Description
-----
ctrl1         11   Control Vlan
ctrl2         102  Control Vlan 2
Default       1
v1            60   vlan 1
vplsVlan      3296 L2 VPN to home office
-----
Total number of VLAN(s) : 5
```



History

This command was first available in ExtremeXOS 12.4.4.

Platform Availability

This command is available on all platforms.

unconfigure vlan description

```
unconfigure {vlan} vlan_name description
```

Description

Removes the description for the specified VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN name.
------------------	--------------------------

Default

N/A.

Usage Guidelines

None.

Example

The following command removes the description from VLAN vlan1:

```
unconfigure vlan vlan1 description
```

History

This command was first available in ExtremeXOS 12.4.4.

Platform Availability

This command is available on all platforms.

unconfigure vlan ipaddress

```
unconfigure {vlan} vlan_name ipaddress {ipv6_address_mask}
```



Description

Removes the IP address of the VLAN or a VMAN. With no parameters, the command removes the primary IPv4 address on the specified VLAN. Using the IPv6 parameters, you can remove specified IPv6 addresses from the specified VLAN.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN or VMAN name.
<i>ipv6_address_mask</i>	Specifies an IPv6 address using the format of IPv6-address/prefix-length, where IPv6 is the 128-bit address and the prefix length specifies the number of leftmost bits that comprise the prefix.

Default

Removes the primary IPv4 address from the specified VLAN or VMAN.

Usage Guidelines



Note

With IPv6, you cannot remove the last link local IPv6 address until all global IPv6 addresses are removed. For MLAG configurations, you cannot remove an IP address from a VLAN until after you delete the MLAG peer.

Example

The following command removes the primary IPv4 address from the VLAN accounting:

```
unconfigure vlan accounting ipaddress
```

The following command removes an IPv6 addresses from the VLAN finance:

```
unconfigure vlan finance ipaddress 3ffe::1
```

History

This command was first available in ExtremeXOS 10.1.

The IPv6 parameters were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.



15 VMAN (PBN) Commands

```
configure port ethertype
configure vman add ports
configure vman add ports cep
configure vman delete ports
configure vman ethertype
configure vman ports add cvid
configure vman ports delete cvid
configure vman tag
create vman
delete vman
disable dot1p examination inner-tag ports
disable vman cep egress filtering ports
enable dot1p examination inner-tag port
enable vman cep egress filtering ports
show vman
show vman eaps
show vman ethertype
unconfigure vman ethertype
```

This chapter describes commands for managing the following Layer 2 feature:

- Provider bridge networks (PBNs—also known as VMANs).

Note



The “VMAN” term is an Extreme Networks term that became familiar to Extreme Networks customers before the PBN standard was complete. The VMAN term is used in the ExtremeXOS software and also in this book to support customers who are familiar with this term. The PBN term is also used in this guide to establish the relationship between this industry standard technology and the Extreme Networks VMAN feature.

For an introduction to these features, see the *ExtremeXOS Concepts Guide*.

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on the platforms listed for the PBN feature in [Feature License Requirements](#).

configure port ethertype

```
configure port port_list ethertype {primary | secondary}
```

Description

Assigns the primary or secondary ethertype value to the specified ports.

Syntax Description

<i>port_list</i>	Specifies the list of ports to be configured.
primary	Assigns the primary ethertype value to the specified ports.
secondary	Assigns the secondary ethertype value to the specified ports.

Default

N/A.

Usage Guidelines

None.

Example

The following command configures port 2:1 to use the secondary ethertype:

```
configure port 2:1 ethertype secondary
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

configure vman add ports

```
configure vman vman-name add ports [ all | port_list ] {untagged { port-cvid  
port_cvid} | tagged}
```

Description

Adds one or more ports to a VMAN.



Syntax Description

<i>vman-name</i>	Specifies the VMAN to configure.
all	Specifies all switch ports.
<i>port_list</i>	Specifies a list of ports.
untagged	Configures the specified ports as Customer Network Ports (CNPs).
tagged	Configures the specified ports as Provider Network Ports (PNPs), which are also called VMAN network ports.
<i>port_cvid</i>	Port's CVID used for untagged packets. If unspecified, untagged packets will be single tagged with the VMAN's SVID. If specified, untagged packets will be double tagged with the VMAN's SVID and the port's CVID.

Default

If you do not specify a parameter, the default value is untagged, which creates a CNP.

Usage Guidelines

This command adds ports as either CNPs or PNPs. To add a port to a VMAN as a CEP, use the following command:

```
configure vman vman_name add portsport_list cep cvidcvid_first {-cvid_last}
{translatecvid_first_xlate {-cvid_last_xlate}}
```

The VMAN must already exist before you can add (or delete) ports. VMAN ports can belong to load-sharing groups.

When a port is configured serve as a CNP for one VMAN and A PNP for another VMAN, it inspects the VMAN ethertype in received packets. Packets with a matching ethertype are treated as tagged and switched across the associated PNP VMAN. Packets with a non-matching ethertype are treated as untagged and forwarded into the associated CNP VMAN.

When a port is configured only as a CNP (an untagged VMAN member), whether the VMAN ethertype is 0x8100 or otherwise, all received packets ingress the associated VMAN regardless of the packet's tagging.

Note



If you use the same name across categories (for example, STPD and EAPS names), we recommend that you specify the identifying keyword as well as the actual name. If you do not use the keyword, the system may return an error message.

The following guidelines apply to all platforms:

- You must enable or disable jumbo frames before configuring VMANs. You can enable or disable jumbo frames on individual ports or modules, or on the entire switch. See [Configuring Slots and Ports on a Switch](#) in the *ExtremeXOS Concepts Guide* for more information on configuring jumbo frames.
- Each port can serve in only one VMAN role per VMAN. When multiple roles are configured on a port, each role must be configured for a different VMAN.



- Multiple VMAN roles can be combined on one port with certain VLAN types as shown in the following table.

Table 24: Port Support for Combined VMAN Roles and VLANs

Platform	Combined CNP, CEP, and Tagged VLAN ^{5, 6}	Combined PNP, CNP, and CEP ^{a, b, 7}	Combined PNP and Tagged VLAN	Combined PNP and Untagged VLAN
Summit X440, X460, X480, X670 and X770	X	X	X ⁸	X
BlackDiamond 8500 and 8800 c-, and e-series modules	X	X	Xd	X
BlackDiamond X8 series switches and BlackDiamond 8900 c-, xl-, and xm-series modules	X	X	Xe	X

Note

If you already configured VLANs and VMANs on the same module or stand-alone switch using ExtremeXOS 11.4, you cannot change the VMAN ethertype from 0X8100 without first removing either the VLAN or VMAN configuration.

Example

The following command assigns ports 1:1, 1:2, 1:3, and 1:6 to a VMAN named accounting:

```
configure vman accounting add ports 1:1, 1:2, 1:3, 1:6 tag 100
```

History

This command was first available in ExtremeXOS 11.0.

The `svid` keyword was added in ExtremeXOS 12.2.

The `cvid` keyword was added in ExtremeXOS 15.3.2.

⁵ Subsets of this group are also supported. That is, any two of these items are supported.

⁶ When a CNP is combined with a CEP or tagged VLAN, any CVIDs not explicitly configured for a CEP or tagged VLAN are associated with the CNP.

⁷ A PNP (tagged VMAN) and a CNP (untagged VMAN) or CEP cannot be combined on a port for which the selected VMAN ethertype is 0x8100.

⁸ If the secondary VMAN ethertype is selected for the port, it must be set to 0x8100.

⁵ Subsets of this group are also supported. That is, any two of these items are supported.

⁶ When a CNP is combined with a CEP or tagged VLAN, any CVIDs not explicitly configured for a CEP or tagged VLAN are associated with the CNP.

⁷ A PNP (tagged VMAN) and a CNP (untagged VMAN) or CEP cannot be combined on a port for which the selected VMAN ethertype is 0x8100.



Platform Availability

This command is available on all platforms.

configure vman add ports cep

```
configure vman vman_name add ports port_list cep cvid cvid_first {- cvid_last }
{translate cvid_first_xlate {- cvid_last_xlate }
```

Description

Adds one or more switch ports to the specified VMAN as Customer Edge Ports (CEPs), and configures the CVIDs on those ports to map to the VMAN.

Syntax Description

<i>vman_name</i>	Specifies the VMAN to configure.
<i>port_list</i>	Specifies a list of ports.
<i>cvid_first</i>	Specifies a CVLAN ID (CVID) or the first in a range of CVIDs that the CEP will accept and map to the specified VMAN. Valid values are 1-4095.
<i>cvid_last</i>	Specifies the last in a range of CVIDs that the CEP will accept and map to the VMAN. Valid values are 1-4095.
translate	Enables translation of the specified CEP CVID range to the specified VMAN CVID range.
<i>cvid_first_xlate</i>	Specifies a VMAN CVID or the first in a range of VMAN CVIDs to which the CEP CVIDs will map. Valid values are 1-4095.
<i>cvid_last_xlate</i>	Specifies the last in a range of VMAN CVIDs to which the CEP CVIDs will map. Valid values are 1-4095. The number of VMAN CVIDs in this range must equal the number of CEP CVIDs specified in this command.

Default

N/A.

Usage Guidelines

If you specify only one CVID or a range of CVIDs without translation, the specified CVIDs are mapped to the specified VMAN and appear unchanged in the VMAN.

If you specify CVID translation, the CEP CVIDs map to different VMAN CVIDs. The number of CEP CVIDs specified must equal the number of VMAN CVIDs specified. The first CEP CVID in the specified range maps to the first CVID in the range specified for the VMAN. The difference between *cvid_first* and *cvid_first_xlate* establishes an offset N that maps CEP CVIDs to VMAN CVIDs. (Offset N = *cvid_first_xlate* - *cvid_first*.) The translated VMAN CVID that corresponds to a CEP CVID can be determined as follows:



VMAN CVID = CEP CVID + N



Note

CVID translation can reduce the number of CVIDs that can be mapped to VMANs.

After you enable and configure a CEP with this command, you can use the following command to map additional CVIDs on the port to the VMAN:

```
configure vman vman_name ports port_list add cvidcvid_first {-cvid_last}
{translatecvid_first_xlate {-cvid_last_xlate}}
```

When this command specifies multiple ports, each port gets an independent CVID map; the ports do not share a common map. Changes to the CVID map affect only the ports specified in the configuration command. For example, consider the following commands:

```
configure vman vman1 add port 1-2 cep cvid 10
configure vman vman1 port 1 add cvid 11
```

After these commands are entered, port 1 maps CVIDs 10 and 11 to VMAN vman1, and port 2 maps only CVID 10 to vman1.

You can add the same port as a CEP to multiple VMANs. A port can also support multiple VMANs in different roles as shown in [Table 24: Port Support for Combined VMAN Roles and VLANs](#) on page 1140.

To view the CEP CVID configuration for a port, use the `show vman` command.

Example

The following command configures port 1 as a CEP for VMAN vman1 and specifies that CEP CVID 5 maps to CVID 5 on the VMAN:

```
configure vman vman1 add port 1 cep cvid 5
```

The following command configures port 1 as a CEP for VMAN vman1 and enables the port to translate CEP CVIDs 10-19 to VMAN CVIDs 20-29:

```
configure vman vman1 add port 1 cep cvid 10 - 19 translate 20 - 29
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on BlackDiamond X8, BlackDiamond 8800 series switches and Summit family switches

The CVID translation feature is available only on BlackDiamond X8, BlackDiamond 8900 c-, xl-, and xm-series modules and Summit X440, X460, X480, X670 and X770 series switches.



configure vman delete ports

```
configure vman vman-name delete ports [all | port_list]
```

Description

Deletes one or more ports from a VMAN.

Syntax Description

<i>vman_name</i>	Specifies a VMAN name.
all	Specifies all ports in the VMAN.
<i>port_list</i>	Specifies a list of ports.

Default

N/A.

Usage Guidelines

The VMAN must already exist before you can delete ports.

Example

The following command deletes ports 1:1, 1:2, 1:3, and 1:6 on a modular switch for a VMAN named accounting:

```
configure vman accounting delete ports 1:1, 1:2, 1:3, 1:6
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure vman ethertype

```
configure vman ethertype value [primary | secondary]
```

Description

Changes the default ethertype for the VMAN header.



Syntax Description

<i>value</i>	Specifies an ethertype value in the format of 0xffff.
primary	Assigns the ethertype as the primary Ethernet value.
secondary	Assigns the ethertype as the secondary Ethernet value.

Default

Ethertype value of 0x88a8 and type primary.

Usage Guidelines

The software supports two VMAN ethertype values, a primary value and a secondary value. By default, the primary ethertype applies to all VMANs. To use the secondary ethertype, define the ethertype with this command, and then assign the secondary ethertype to ports with the following command:

```
configure port port_list ethertype {primary | secondary}
```

If your VMAN transits a third-party device (other than an Extreme Networks device), you must configure the ethertype for the VMAN tag as the ethertype that the third-party device uses. If you configure both primary and secondary ethertypes, you can connect to devices that use either of the two values assigned.

The system supports all VMAN ethertypes, including the standard ethertype of 0x8100.

Example

The following command changes the VMAN ethertype value to 8100:

```
configure vman ethertype 0x8100
```

History

This command was first available in ExtremeXOS 11.0.

Support for a secondary ethertype was added in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

configure vman ports add cvid

```
configure vman vman_name ports port_list add cvid cvid_first {- cvid_last}
{translate cvid_first_xlate {- cvid_last_xlate}}
```



Description

Adds one or more CVIDs to a CEP.

Syntax Description

<i>vman_name</i>	Specifies the VMAN to configure.
<i>port_list</i>	Specifies a list of ports.
<i>cvid_first</i>	Specifies a Customer VLAN ID (CVID) or the first in a range of CVIDs that the CEP will accept and map to the specified VMAN. Valid values are 1-4095.
<i>cvid_last</i>	Specifies the last in a range of CVIDs that the CEP will accept and map to the VMAN. Valid values are 1-4095.
translate	Enables translation of the specified CEP CVID range to the specified VMAN CVID range.
<i>cvid_first_xlate</i>	Specifies a VMAN CVID or the first in a range of VMAN CVIDs to which the CEP CVIDs will map. Valid values are 1-4095.
<i>cvid_last_xlate</i>	Specifies the last in a range of VMAN CVIDs to which the CEP CVIDs will map. Valid values are 1-4095. The number of VMAN CVIDs in this range must equal the number of CEP CVIDs specified in this command.

Default

N/A.

Usage Guidelines

Before you can add CVIDs to CEPs, you must configure the target physical ports as CEPs using the following command:

```
configure vman vman_name add portsport_list cep cvidcvid_first {- cvid_last}
{translatecvid_first_xlate {-cvid_last_xlate}}
```

If you specify only one CVID or a range of CVIDs without translation, the specified CVIDs are mapped to the specified VMAN and appear unchanged in the VMAN.

If you specify CVID translation, the CEP CVIDs map to different VMAN CVIDs. The number of CEP CVIDs specified must equal the number of VMAN CVIDs specified. The first CEP CVID in the specified range maps to the first CVID in the range specified for the VMAN. The difference between *cvid_first* and *cvid_first_xlate* establishes an offset N that maps CEP CVIDs to VMAN CVIDs. (Offset N = *cvid_first_xlate* - *cvid_first*.) The translated VMAN CVID that corresponds to a CEP CVID can be determined as follows:

VMAN CVID = CEP CVID + N



Note

CVID translation can reduce the number of CVIDs that can be mapped to VMANs.



When this command specifies multiple ports, each port gets an independent CVID map; the ports do not share a common map. Changes to the CVID map affect only the ports specified in the configuration command. For example, consider the following commands:

```
configure vman vman1 add port 1-2 cep cvid 10
configure vman vman1 port 1 add cvid 11
```

After these commands are entered, port 1 maps CVIDs 10 and 11 to VMAN vman1, and port 2 maps only CVID 10 to vman1.

To view the CEP CVID configuration for a port, use the `show vman` command.

Example

The following command adds CVIDs 20-29 to port 1 and VMAN vman1 and enables translation to CVIDs 30-39:

```
configure vman vman1 port 1 add cvid 20 - 29 translate 30 - 99
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8800 series switches and Summit family switches.

The CVID translation feature is available only on BlackDiamond X8 series switches, BlackDiamond 8900 c-, xl-, and xm-series modules and Summit X440, X460, X480, X670 and X770 series switches.

configure vman ports delete cvid

```
configure vman vman_name ports port_list delete cvid cvid_first {- cvid_last }
```

Description

Deletes one or more CVIDs from a CEP.

Syntax Description

<i>vman_name</i>	Specifies the VMAN to configure.
<i>port_list</i>	Specifies a list of ports.
<i>cvid_first</i>	Specifies a CVID or the first in a range of CVIDs that are to be deleted. Valid values are 1-4095.
<i>cvid_last</i>	Specifies the last in a range of CVIDs that are to be deleted. Valid values are 1-4095.



Default

N/A.

Usage Guidelines

Each CEP has its own CVID map, and this command deletes CVIDs only from the ports specified with this command.

If all the CVIDs are deleted from a CEP, the CEP is deleted from the VMAN.

To view the CEP CVID configuration for a port, use the `show vman` command.

Example

The following command deletes CVID 15 on port 1 from VMAN vman1:

```
configure vman vman1 port 1 delete cvid 15
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8800 series switches and Summit family switches.

configure vman tag

```
configure vman vman_name tag tag
```

Description

Assigns a tag to a VMAN.

Syntax Description

<i>vman_name</i>	Specifies a VMAN name.
<i>tag</i>	Specifies a value to use as the VMAN tag. The valid range is from 2 to 4094.

Default

N/A.



Usage Guidelines

Every VMAN requires a unique tag.

You can specify a value that is currently used as an internal VLAN ID on another VLAN; it becomes the VLAN ID for the VLAN you specify, and a new VLAN ID is automatically assigned to the other untagged VLAN.

Example

The following command assigns a tag of 120 to a VMAN named accounting:

```
configure vman accounting tag 120
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

create vman

```
create vman vman-name {learning-domain} {vr vr_name}
```

Description

Creates a VMAN.

Syntax Description

<i>vman-name</i>	Specifies a VMAN name using up to 32 characters.
learning-domain	Specifies that this VMAN is a learning domain, which supports inter-VMAN forwarding.
vr	Specifies a virtual router.
<i>vr_name</i>	Specifies a virtual router name. NOTE: User-created VRs are supported only on the platforms listed for this feature in Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i> . On switches that do not support user-created VRs, all VMANs are created in VR-Default and cannot be moved.

Default

N/A.



Usage Guidelines

For information on VMAN name requirements and a list of reserved keywords, see [Object Names](#) of the *ExtremeXOS Concepts Guide*. You must use mutually exclusive names for:

- VLANs.
- VMANs.
- IPv6 tunnels.

The keyword `learning-domain` enables you to create a VMAN that serves as a learning domain for inter-VMAN forwarding.

If you do not specify the virtual router, the VMAN is created in the current virtual router.

After you create the VMAN, you must configure the VMAN tag and add the ports that you want.

Example

The following command creates a VMAN named `fred`:

```
create vman fred
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

delete vman

```
delete vman vman-name
```

Description

Deletes a previously created VMAN.

Syntax Description

<i>vman-name</i>	Specifies a VMAN name.
------------------	------------------------

Default

N/A.



Usage Guidelines

None.

Example

The following command deletes the VMAN accounting:

```
delete vman accounting
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

disable dot1p examination inner-tag ports

```
disable dot1p examination inner-tag ports [all | port_list]
```

Description

Used with VMANs, and instructs the switch to examine the 802.1p value of the outer tag, or added VMAN header, to determine the correct egress queue on the egress port.

Syntax Description

all	Specifies all ports.
<i>port_list</i>	Specifies a list of ports or slots and ports.

Default

Disabled.

Usage Guidelines

Use this command to instruct the system to refer to the 802.1p value contained in the outer tag, or VMAN encapsulation tag, when assigning the packet to an egress queue at the egress port of the VMAN.



Note

See [QoS Commands](#) for information on configuring and displaying the current 802.1p and DiffServ configuration for the inner, or original header, 802.1p value.



Example

The following command uses the 802.1p value on the outer tag, or VMAN encapsulation, to put the packet in the egress queue on the VMAN egress port:

```
disable dot1p examination inner-tag port 3:2
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available only on the BlackDiamond X8, BlackDiamond 8800 series switches, SummitStack, and Summit family switches.

disable vman cep egress filtering ports

```
disable vman cep egress filtering ports {port_list | all}
```

Description

Disables the egress filtering of CVIDs that are not configured in the CVID map for a CEP.

Syntax Description

<i>port_list</i>	Specifies a list of ports.
all	Specifies all switch ports.

Default

Egress CVID filtering is disabled.

Usage Guidelines

To view the configuration setting for the egress CVID filtering feature, use the `show ports information` command.



Note

When CVID egress filtering is enabled, it reduces the maximum number of CVIDs supported on a port. The control of CVID egress filtering applies to fast-path forwarding. When frames are forwarded through software, CVID egress filtering is always enabled.



Example

The following command disables egress CVID filtering on port 1:

```
disable vman cep egress filtering port 1
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on the BlackDiamond X8, BlackDiamond 8900 c-, xl-, and xm-series modules. This command is also available on Summit X440, X460, X480, X670 and X770 series switches.

enable dot1p examination inner-tag port

```
enable dot1p examination inner-tag port [all | port_list]
```

Description

Used with VMANs, and instructs the switch to examine the 802.1p value of the inner tag, or header of the original packet, to determine the correct egress queue on the egress port.

Syntax Description

all	Specifies all ports.
<i>port_list</i>	Specifies a list of ports or slots and ports.

Default

Disabled.

Usage Guidelines

Use this command to instruct the system to refer to the 802.1p value contained in the inner, or original, tag when assigning the packet to an egress queue at the egress port of the VMAN.



Note

See [QoS Commands](#) for information on configuring and displaying the current 802.1p and DiffServ configuration for the inner, or original header, 802.1p value.



Example

The following command puts the packets in the egress queue of the VMAN egress port according to the 802.1p value on the inner tag:

```
enable dot1p examination inner-tag port 3:2
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available only on the BlackDiamond X8, BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches.

enable vman cep egress filtering ports

```
enable vman cep egress filtering ports {port_list | all}
```

Description

Enables the egress filtering of frames based on their CVIDs on ports configured as CEPs.

Syntax Description

<i>port_list</i>	Specifies a list of ports.
all	Specifies all switch ports.

Default

Egress CVID filtering is disabled.

Usage Guidelines

For a given VMAN and a port configured as a CEP for that VMAN, only frames with CVIDs that have been mapped from the CEP to the VMAN are forwarded from the VMAN and out the CEP.

To view the configuration setting for the egress CVID filtering feature, use the `show ports information` command.

Note



CVID egress filtering is available only on switches that support this feature, and when this feature is enabled, it reduces the maximum number of CVIDs supported on a port. The control of CVID egress filtering applies to fast-path forwarding. When frames are forwarded through software, CVID egress filtering is always enabled.



Example

The following command enables egress CVID filtering on port 1:

```
enable vman cep egress filtering port 1
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on the BlackDiamond X8 series switches and the BlackDiamond 8900 c-, xl-, and xm-series modules. This command is also available on Summit X440, X460, X480, X670 and X770 series switches.

show vman

```
show vman show {vman} vman_name {ipv4 | ipv6} show vman [tag tag | detail] {ipv4 | ipv6}
```

Description

Displays information about one or all VMANs.



Note

The information displayed for this command depends on the platform and configuration you are using.

Syntax Description

<i>vman_name</i>	Specifies that information is displayed for the specified VMAN.
<i>tag</i>	Specifies a VMAN using the 802.1Q tag.
detail	Specifies that all information is displayed for each VMAN.
ipv4	Specifies IPv4.
ipv6	Specifies IPv6.

Default

Summary information for all VMANs on the switch.

Usage Guidelines

The information displayed with this command depends on the platform and configuration you are using.



Example

The following example displays a list of all the VMANs on the switch:

```
* BD-12804.17 # show vman
-----
-----
Name          VID  Protocol Addr          Flags          Proto  Ports
Virtual
Active router
/Total
-----
-----
le1           4091 -----a      ANY    2 /2   VR-
Default
le2           4090 -----a      ANY    0 /0   VR-
Default
vm1           4089 -----      ANY    0 /0   VR-
Default
-----
-----
Flags : (a) Learning Domain (C) EAPS Control vlan, (E) ESRP Enabled,
(f) IP Forwarding Enabled, (i) ISIS Enabled, (I) IP Forwarding lpm-routing
Enabled,
(L) Loopback Enabled, (m) IPmc Forwarding Enabled,
(n) IP Multinetting Enabled, (N) Network LogIn vlan,
(o) OSPF Enabled, (p) PIM Enabled,
(P) EAPS protected vlan, (r) RIP Enabled, (T) Member of STP Domain,
(v) VRRP Enabled, (B) 802.1ah Backbone VMAN, (S) 802.1ah Service VMAN
Total number of vman(s) : 3
```

The following example displays information on a single VMAN named vman1:

```
# show vman blue
VMAN Interface with name vman1 created by user
Admin State:      Enabled          Tagging:      802.1Q Tag 100
Virtual router:  VR-Default
IPv4 Forwarding: Disabled
IPv6 Forwarding: Disabled
IPv6:            None
STPD:           None
Protocol:       Match all unfiltered protocols
Loopback:      Disabled
NetLogin:      Disabled
QosProfile:     None configured
Egress Rate Limit Designated Port: None configured
Flood Rate Limit QosProfile:      None configured
Ports:        2.          (Number of active ports=0)
Tag:          *1,          *2
CEP:          *3: CVID 20-29
              *4: CVID 10-19 translate 20-29
              *5: CVID 10-19 translate 20-29,CVID 30
Flags:        (*) Active, (!) Disabled, (g) Load Sharing port
              (b) Port blocked on the vlan, (m) Mac-Based port
              (a) Egress traffic allowed for NetLogin
              (u) Egress traffic unallowed for NetLogin
```



- (t) Translate VLAN tag for Private-VLAN
- (s) Private-VLAN System Port, (L) Loopback port
- (e) Private-VLAN End Point Port
- (x) VMAN Tag Translated port
- (G) Multi-switch LAG Group port

The Port CVID output was added in the display of `show vman vlan_name | detail` in ExtremeXOS 15.3.2:

```
VMAN Interface with name vml created by user
  Admin State: Enabled           Tagging:           802.1Q Tag
1000
  Description: None
  Virtual router: VR-Default
  IPv4 Forwarding: Disabled
  IPv6 Forwarding: Disabled
  IPv6:           None
  STPD:           None
  Protocol:       Match all unfiltered protocols
  Loopback:       Disabled
  NetLogin:       Disabled
  QosProfile:     None configured
  Egress Rate Limit Designated Port: None configured
  Flood Rate Limit QosProfile:     None configured
  Ports:         3.           (Number of active ports=3)
    Untag:        *21: Port CVID 5,
                 *24: Port CVID 7,
  Tag:           *22
  Flags:         (*) Active, (!) Disabled, (g) Load Sharing port
                 (b) Port blocked on the vlan, (m) Mac-Based port
                 (a) Egress traffic allowed for NetLogin
                 (u) Egress traffic unallowed for NetLogin
                 (t) Translate VLAN tag for Private-VLAN
                 (s) Private-VLAN System Port, (L) Loopback port
                 (x) VMAN Tag Translated port
                 (G) Multi-switch LAG Group port
```

The `show vman detail` command shows all the information shown in the `show vman vlan_name` command, but displays information for all configured VMANs.

History

This command was first available in ExtremeXOS 11.0.

Information on IEEE 802.1ah was added in ExtremeXOS 11.4.

The tag option was added in ExtremeXOS 12.4.4.

Port CVID output was added in ExtremeXOS 15.3.2.

Platform Availability

This command is available on all platforms.



CEP information is displayed only on BlackDiamond X8, BlackDiamond 8800 series switches and Summit family switches.

show vman eaps

```
show {vman} vman_name eaps
```

Description

Displays the EAPS domains to which the VMAN belongs.

Syntax Description

<i>vman_name</i>	Specifies the name of the VMAN for which EAPS information is to be displayed.
------------------	---

Default

N/A.

Usage Guidelines

None.

Example

The following example displays a list of EAPS domains for the campus1 VMAN:

```
* BD-12804.17 # show vman campus1 eaps
```

History

This command was first available in ExtremeXOS 11.0.

Information on IEE 802.1ah was added in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

show vman ethertype

```
show vman etherType
```



Description

Displays the etherType information for VLANs, VMANs, and PBBNs.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following example shows the display from the show vman etherType command on switches that support only VMANs:

```
vMan EtherType: 0x88a8
```

The following example shows the display from the show vman etherType command on switches that support PBBNs:

```
BlackDiamond 12804.41 # show vman etherType
vman EtherType : 0x88a8
bvlan EtherType: 0x88b5
```

The following example shows the display from the show vman etherType command when a secondary etherType is configured:

```
BD-12804.3 # show vman etherType
Vman Primary EtherType   : 0x9100
Vman Secondary EtherType : 0x8100
BVlan EtherType          : 0x88b5
Secondary EtherType ports : 6:2g 6:3
```

The letter g in the port list indicates that the port is a LAG/Trunk port, the details of which can be seen using the show port sharing command.

History

This command was first available in ExtremeXOS 11.0.

Information on IEE 802.1ah was added in ExtremeXOS 11.4.



Platform Availability

This command is available on all platforms.

unconfigure vman ethertype

```
unconfigure vman ethertype {secondary}
```

Description

Restores the default primary VMAN ethertype value of 0x88A8 or deletes the secondary ethertype value.

Syntax Description

secondary	Deletes the secondary ethertype value.
------------------	--

Default

N/A.

Usage Guidelines

When you enter this command without the secondary option, the primary VMAN ethertype returns to the default value of 0x88A8. If you specify the secondary option, the secondary VMAN ethertype value is deleted (no value is assigned).



Note

Before unconfiguring the secondary VMAN ethertype, any secondary VMAN port must be changed to the primary VMAN ethertype; otherwise this command fails.

Example

The following command restores the primary VMAN ethertype to the default value:

```
unconfigure vman ethertype
```

The following command restores the secondary VMAN ethertype to the default value:

```
unconfigure vman ethertype secondary
```

History

This command was first available in ExtremeXOS 11.0.



Platform Availability

This command is available on all platforms.



16 FDB Commands

```
clear fdb
configure fdb mac-tracking ports
configure fdb static-mac-move packets
create fdbentry vlan ports
delete fdb mac-tracking entry
delete fdbentry
disable fdb static-mac-move
disable flooding ports
disable learning iparp sender-mac
disable learning port
disable snmp traps fdb mac-tracking
enable fdb static-mac-move
enable flooding ports
enable learning iparp sender-mac
enable learning port
enable snmp traps fdb mac-tracking
show fdb
show fdb mac-tracking configuration
show fdb mac-tracking statistics
show fdb static-mac-move configuration
show fdb stats
```

This chapter describes commands for:

- Configuring FDB entries.
- Displaying FDB entries.
- Managing the MAC tracking feature.

For an introduction to FDB features, see the *ExtremeXOS Concepts Guide*.

clear fdb

```
clear fdb {mac_addr | ports port_list | vlan vlan_name | blackhole}
```

Description

Clears dynamic FDB entries that match the filter.

Syntax Description

<i>mac_addr</i>	Specifies a MAC address, using colon-separated bytes.
<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>vlan_name</i>	Specifies a VLAN name.
blackhole	Specifies the blackhole entries.

Default

Clears all dynamic FDB entries.

Usage Guidelines

This command clears FDB entries based on the specified criteria. When no options are specified, the command clears all dynamic FDB entries.

Example

The following command clears any FDB entries associated with ports 4:3-4:5 on a modular switch:

```
clear fdb ports 4:3-4:5
```

The following command clears any FDB entries associated with VLAN corporate:

```
clear fdb vlan corporate
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure fdb mac-tracking ports

```
configure fdb mac-tracking {[add|delete]} ports [port_list|all]
```

Description

Enables or disables MAC address tracking for all MAC addresses on the specified ports.



Syntax Description

add	Enables MAC address tracking for the specified ports.
delete	Disables MAC address tracking for the specified ports.
<i>port_list</i>	Specifies a list of ports on which MAC address tracking is to be enabled or disabled.
all	Specifies that MAC address tracking is to be enabled or disabled on all ports.

Default

No ports are enabled for MAC address tracking.

Usage Guidelines

MAC address tracking events on enabled ports generate EMS messages and can optionally generate SNMP traps.



Note

When a MAC address is configured in the tracking table, but detected on a MAC tracking enabled port, the per MAC address statistical counters are not updated.

Example

The following command enables MAC address tracking for all MAC addresses on port 2:1:

```
configure fdb mac-tracking add ports 2:1
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

configure fdb static-mac-move packets

```
configure fdb static-mac-move packets count
```

Description

Configures the number of EMS and SNMP reports that can be generated each second for MAC addresses that are duplicates of statically configured MAC addresses.



Syntax Description

<i>count</i>	Specifies the number of duplicate MAC address events that are reported each second. The range is 1 to 25.
--------------	---

Default

2.

Usage Guidelines

None.

Example

The following command configures the switch to report up to 5 duplicate MAC address events per second:

```
configure fdb static-mac-move packets 5
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on the Summit family switches.

create fdbentry vlan ports

```
create fdbentry mac_addr vlan vlan_name [ports port_list {tagged tag} | blackhole]
```

Description

Creates a permanent static FDB entry.

Syntax Description

<i>mac_addr</i>	Specifies a device MAC address, using colon-separated bytes.
<i>vlan_name</i>	Specifies a VLAN name associated with a MAC address.
<i>port_list</i>	Specifies one or more ports or slots and ports associated with the MAC address.



tagged tag	Specifies the port-specific VLAN tag. When there are multiple ports specified in the <i>port_list</i> , the same tag is used for all of them.
blackhole	Enables the blackhole option. Any packets with either a source MAC address or a destination MAC address matching the FDB entry are dropped.

Default

N/A.

Usage Guidelines

Permanent entries are retained in the database if the switch is reset or a power off/on cycle occurs. A permanent static entry can either be a unicast or multicast MAC address. After they have been created, permanent static entries stay the same as when they were created. If the same MAC address and VLAN is encountered on another virtual port that is not included in the permanent MAC entry, it is handled as a blackhole entry. The static entry is not updated when any of the following take place:

- A VLAN identifier (VLANid) is changed.
- A port is disabled.
- A port enters blocking state.
- A port goes down (link down).

A permanent static FDB entry is deleted when any of the following take place:

- A VLAN is deleted.
- A port mode is changed (tagged/untagged).
- A port is deleted from a VLAN.

Permanent static entries are designated by `spm` in the flags field of the `show fdb` output. You can use the `show fdb` command to display permanent FDB entries.

If the static entry is for a PVLAN VLAN that requires more than one underlying entry, the system automatically adds the required entries. For example, if the static entry is for a PVLAN network VLAN, the system automatically adds all required extra entries for the subscriber VLANs.

You can create FDB entries to multicast MAC addresses and list one or more ports. If more than one port number is associated with a permanent MAC entry, packets are multicast to the multiple destinations.

IGMP snooping rules take precedence over static multicast MAC addresses in the IP multicast range (01:00:5e:xx:xx:xx) unless IGMP snooping is disabled.



Note

When a multiport list is assigned to a unicast MAC address, load sharing is not supported on the ports in the multiport list.



Example

The following command adds a permanent, static entry to the FDB for MAC address 00 E0 2B 12 34 56, in VLAN marketing on slot 2, port 4 on a modular switch:

```
create fdbentry 00:E0:2B:12:34:56 vlan marketing port 2:4
```

The following example creates a multiport unicast FDB entry, in VLAN black, on slot 1, ports 1, 2, and 4, on the BlackDiamond 8800 family of switches:

```
create fdbentry 01:00:00:00:00:01 vlan black port 1:1, 1:2, 1:4
```

The following command adds a permanent, static entry to the FDB for MAC address 00:01:02:03:04:05, in VLAN marketing, on a vlan port that has tag 100 on port 3 on a switch:

```
create fdbentry 00:01:02:03:04:05 vlan msk ports 3 tag 100
```

History

This command was first available in ExtremeXOS 10.1.

The ability to create a multicast FDB with multiple entry ports was added in ExtremeXOS 11.3.

The blackhole option was first available for all platforms in ExtremeXOS 12.1.

The ability to create a unicast FDB with multiple entry ports was available for the BlackDiamond 8000 c-, and e-series modules in ExtremeXOS 12.1. This feature is supported on all later platforms when introduced.

The **tag** keyword and example was added in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

delete fdb mac-tracking entry

```
delete fdb mac-tracking entry [mac_addr | all]
```

Description

Deletes a MAC address from the MAC address tracking table.



Syntax Description

<i>mac_addr</i>	Specifies a device MAC address, using colon-separated bytes.
all	Specifies that all MAC addresses are to be deleted from the MAC address tracking table.

Default

The MAC address tracking table is empty.

Usage Guidelines

None.

Example

The following command deletes a MAC address from the MAC address tracking table:

```
delete fdb mac-tracking entry 00:E0:2B:12:34:56
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.

delete fdbentry

```
delete fdbentry [all | mac_address [vlan vlan_name ]
```

Description

Deletes one or all permanent FDB entries.

Syntax Description

all	Specifies all FDB entries.
<i>mac_address</i>	Specifies a device MAC address, using colon-separated bytes.
<i>vlan_name</i>	Specifies the specific VLAN name.

Default

N/A.



Usage Guidelines

None.

Example

The following example deletes a permanent entry from the FDB:

```
delete fdbentry 00:E0:2B:12:34:56 vlan marketing
```

The following example deletes all permanent entries from the FDB:

```
delete fdbentry all
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

disable fdb static-mac-move

disable fdb static-mac-move

Description

Disables EMS and SNMP reporting of discovered MAC addresses that are duplicates of statically configured MAC addresses.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.



Example

The following command disables this feature:

```
disable fdb static-mac-move
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on Summit family switches.

disable flooding ports

With the BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches, you can further identify the type of packets for which to block flooding.

```
disable flooding [all_cast | broadcast | multicast | unicast] ports [port_list | all]
```

Description

Disables Layer 2 egress flooding on one or more ports.

Syntax Description

all_cast	Specifies disabling egress flooding for all packets on specified ports.
broadcast	Specifies disabling egress flooding only for broadcast packets.
multicast	Specifies disabling egress flooding only for multicast packets.
unicast	Specifies disabling egress flooding only for unknown unicast packets.
<i>port_list</i>	Specifies one or more ports or slots and ports.
all	Specifies all ports on the switch.

Default

Enabled for all packet types.

Usage Guidelines



Note

If an application requests specific packets on a specific port, those packets are not affected by the `disable flooding ports` command.



You might want to disable egress flooding to do the following:

- enhance security
- enhance privacy
- improve network performance

This is particularly useful when you are working on an edge device in the network. The practice of limiting flooded egress packets to selected interfaces is also known as upstream forwarding.



Note

If you disable egress flooding with static MAC addresses, this can affect many protocols, such as IP and ARP.

The following guidelines apply to enabling and disabling egress flooding:

- Disabling multicasting egress flooding does not affect those packets within an IGMP membership group at all; those packets are still forwarded out. If IGMP snooping is disabled, multicast packets are not flooded.
- Egress flooding can be disabled on ports that are in a load-sharing group. In a load-sharing group, the ports in the group take on the egress flooding state of the master port; each member port of the load-sharing group has the same state as the master port.
- On all platforms FDB learning takes place on ingress ports and is independent of egress flooding; either can be enabled or disabled independently.
- Disabling unicast or all egress flooding to a port also stops packets with unknown MAC addresses to be flooded to that port.
- Disabling broadcast or all egress flooding to a port also stops broadcast packets to be flooded to that port.

BlackDiamond X8 Series switches, BlackDiamond 8800 family of switches, SummitStack, and the Summit switch only

You can disable egress flooding for unicast, multicast, or broadcast MAC addresses, as well as for all packets on the ports of the BlackDiamond 8800 family of switches, SummitStack, and the Summit switch. The default behavior for the BlackDiamond 8800 family of switches, SummitStack, and the Summit is enabled egress flooding for all packet types.

The following command disables unicast flooding on ports 10-12 on a Summit series switch:

```
disable flooding unicast port 10-27
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on BlackDiamond X8 and 8800 series switches, SummitStack, and Summit family switches.



disable learning iparp sender-mac

```
disable learning iparp {vr vr_name} sender-mac
```

Description

Disables MAC address learning from the payload of IP ARP packets.

Syntax Description

<i>vr_name</i>	Specifies a virtual router.
----------------	-----------------------------

Default

Disabled.

Usage Guidelines

To view the configuration for this feature, use the following command:

```
show iparp
```

Example

The following command disables MAC address learning from the payload of IP ARP packets:

```
disable learning iparp sender-mac
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all Summit family switches, SummitStack, and BlackDiamond 8800 series switches.

disable learning port

```
disable learning {drop-packets | forward-packets} port [port_list | all]
```

Description

Disables MAC address learning on one or more ports for security purposes.



Syntax Description

port	Specifies the port.
<i>port_list</i>	Specifies one or more ports or slots and ports.
all	Specifies all ports and slots.
drop-packets	Specifies that packets with unknown source MAC addresses be dropped. If you do not specify the forward-packets option, this option is used.
forward-packets	Specifies that packets with unknown source MAC addresses be forwarded.

Default

Enabled.

Usage Guidelines

Use this command in a secure environment where access is granted via permanent forwarding database (FDB) entries per port.

Example

The following command disables MAC address learning on port 4:3:

```
disable learning ports 4:3
```

History

This command was first available in ExtremeXOS 10.1.

The drop packets and forward packets options were added in ExtremeXOS 12.1.

Platform Availability

This command is available on all Summit family switches, SummitStack, and BlackDiamond X8 and 8800 series switches.

disable snmp traps fdb mac-tracking

```
disable snmp traps fdb mac-tracking
```

Description

Disables SNMP trap generation when MAC-tracking events occur for a tracked MAC address.



Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables SNMP traps for MAC-tracking events:

```
disable snmp traps fdb mac-tracking
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.

enable fdb static-mac-move

```
enable fdb static-mac-move
```

Description

Enables EMS and SNMP reporting of discovered MAC addresses that are duplicates of statically configured MAC addresses.

Syntax Description

This command has no arguments or variables.

Default

Disabled.



Usage Guidelines

This command enables reporting only. All packets that arrive from a duplicate MAC address on another port (other than the statically configured port) are dropped.

The switch reports the source MAC address, port, and VLAN for each duplicate MAC address.

Example

The following command enables this feature:

```
enable fdb static-mac-move
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on Summit family switches.

enable flooding ports

```
enable flooding [all_cast | broadcast | multicast | unicast] ports [port_list | all]
```

Description

Enables egress flooding on one or more ports. With the BlackDiamond X8, BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches, you can further identify the type of packets to flood on the specified ports.

Syntax Description

all_cast	Specifies enabling egress flooding for all packets on specified ports.
broadcast	Specifies enabling egress flooding only for broadcast packets. NOTE: This parameter is available only on the BlackDiamond X8, BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches.
multicast	Specifies enabling egress flooding only for multicast packets. NOTE: This parameter is available only on the BlackDiamond X8, BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches.
unicast	Specifies enabling egress flooding only for unknown unicast packets. NOTE: This parameter is available only on the BlackDiamond X8, BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches.
<i>port_list</i>	Specifies one or more ports or slots and ports.
all	Specifies all ports on the switch.



Default

Enabled for all packet types.

Usage Guidelines

Use this command to re-enable egress flooding that you previously disabled using the `disable flooding ports` command.

The following guidelines apply to enabling and disabling egress flooding:

- Disabling multicasting egress flooding does not affect those packets within an IGMP membership group at all; those packets are still forwarded out. If IGMP snooping is disabled, multicast packets are not flooded.
- Egress flooding can be disabled on ports that are in a load-sharing group. If that is the situation, the ports in the group take on the egress flooding state of the master port; each member port of the load-sharing group has the same state as the master port.
- FDB learning is independent of egress flooding. FDB learning and egress flooding can be enabled or disabled independently.
- Disabling unicast or all egress flooding to a port also stops packets with unknown MAC addresses to be flooded to that port.
- Disabling broadcast or all egress flooding to a port also stops broadcast packets to be flooded to that port.

BlackDiamond X8, 8800 series switches, SummitStack, and the Summit family of switches only

You can disable egress flooding for unicast, multicast, or broadcast MAC addresses, as well as for all packets on the ports of the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches. The default behavior for the BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches is enabled egress flooding for all packet types.

Example

The following command enables unicast flooding on ports 13-17 on a Summit series switch:

```
enable flooding unicast port 13-17
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on BlackDiamond X8 and BlackDiamond 8800 series switches, SummitStack, and Summit family switches.



enable learning iparp sender-mac

```
enable learning iparp {request | reply | both-request-and-reply} {vr vr_name}
sender-mac
```

Description

Enables MAC address learning from the payload of IP ARP packets.

Syntax Description

request	Enables learning only for IP ARP request packets.
reply	Enables learning only for IP ARP reply packets.
both-request-and-reply	Enables learning for both request and reply packets.
<i>vr_name</i>	Specifies a virtual router.

Default

Disabled.

Usage Guidelines

To view the configuration for this feature, use the following command:

```
show iparp
```

Example

The following command enables MAC address learning from the payload of reply IP ARP packets:

```
enable learning iparp reply sender-mac
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all Summit family switches, SummitStack, and BlackDiamond X8 and BlackDiamond 8800 series switches.

enable learning port

```
enable learning {drop-packets} ports [all | port_list]
```



Description

Enables MAC address learning on one or more ports.

Syntax Description

drop-packets	Forwards EDP packets, and drops all unicast, multicast, and broadcast packets from a source address not in the FDB. No further processing occurs for dropped packets. NOTE: This parameter is available only on BlackDiamond X8, 8800 series switches, SummitStack, and the Summit family switches.
all	Specifies all ports.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Enabled.

Usage Guidelines

Use this command to enable MAC address learning on one or more ports.

Example

The following command enables MAC address learning on slot 1, ports 7 and 8 on a modular switch:

```
enable learning ports 1:7-8
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all Summit family switches, SummitStack, and BlackDiamond X8 and BlackDiamond 8800 series switches.

enable snmp traps fdb mac-tracking

```
enable snmp traps fdb mac-tracking
```

Description

Enables SNMP trap generation when MAC-tracking events occur for a tracked MAC address.



Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables SNMP traps for MAC-tracking events:

```
enable snmp traps fdb mac-tracking
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.

show fdb

```
show fdb {blackhole {netlogin [all | mac-based-vlans]} | netlogin [all | mac-
based-vlans] | permanent {netlogin [all | mac-based-vlans]} | mac_addr {netlogin
[all | mac-based-vlans]} | ports port_list {netlogin [all | mac-based-vlans]} |
vlan vlan_name {netlogin [all | mac-based-vlans]} | {{vpls} {vpls_name}}}
```

Description

Displays FDB entries.

Syntax Description

blackhole	Displays the blackhole entries. (All packets addressed to these entries are dropped.)
slot	Specifies a slot in the switch.
num_entries	Specifies the maximum number of hardware entries to display. The range is 1 to 25.
netlogin all	Displays all FDBs created as a result of the netlogin process.



netlogin mac-based-vlans	Displays all netlogin MAC-based VLAN FDB entries. NOTE: This parameter is supported only for Summit family switches, SummitStack, and the BlackDiamond 8800 series switches. See Network Login Commands for more information on netlogin.
permanent	Displays all permanent entries, including the ingress and egress QoS profiles.
<i>mac_addr</i>	Specifies a MAC address, using colon-separated bytes, for which FDB entries should be displayed.
<i>port_list</i>	Displays the entries for one or more ports or ports and slots.
<i>vlan_name</i>	Displays the entries for a specific VLAN.
<i>vpls_name</i>	Specifies a specific VPLS for which to display entries.

Default

All.

Usage Guidelines

The pulling of MAC addresses for display purposes is given a lower priority to the actual data path learning. Eventually all the MAC addresses are learned in a quiescent system.

The show fdb command output displays the following information:

Mac	The MAC address that defines the entry.
Vlan	The PVLAN or VLAN for the entry.
Age	The age of the entry, in seconds (does not appear if the keyword permanent is specified). The age parameter does not display for the backup MSM/MM on modular switches. On BlackDiamond 8900 xl-series and Summit X480 switches, the Age is always 000 and the h flag is set for entries that are hardware aged.
Flags	Flags that define the type of entry: b - Ingress Blackhole B - Egress Blackhole D - Drop entry for an isolated subscriber VLAN d - Dynamich - Aged in hardware (Applies to BlackDiamond 8900 xl-series and Summit X480 switches) i - an entry also exists in the IP FDB l - lockdown MAC L - lockdown-timeout MAC m - MACM - Mirror n - NetLogino - IEEE 802.1ah backbone MAC P - PVLAN created entry p - Permanent s - Static v - NetLogin MAC-Based VLAN (only supported on the Summit switch, SummitStack, and the BlackDiamond 8800 family of switches) x - an entry also exists in the IPX FDBs.
Port List	The ports on which the MAC address has been learned.

Example

The following command example shows how the FDB entries appear for all options except the hardware option:

```
# show fdb
Mac                Vlan    Age  Flags          Port / Virtual Port List
-----
```



```

00:0c:29:4b:34:cf      v101(0101) 0041 d m      D 1:2
00:0c:29:4b:34:cf      v100(0100) 0041 d m      P 1:2
00:0c:29:d2:2d:48      v102(0102) 0045 d m      1:3
00:0c:29:d2:2d:48      v100(0100) 0045 d m      P 1:3
00:0c:29:f1:f2:f5      v100(0100) 0045 d m      1:1
00:0c:29:f1:f2:f5      v102(0102) 0045 d m      P 1:1
00:0c:29:f1:f2:f5      v101(0101) 0045 d m      P 1:1
Flags : d - Dynamic, s - Static, p - Permanent, n - NetLogin, m - MAC, i - IP,
x - IPX, l - lockdown MAC, L - lockdown-timeout MAC, M- Mirror, B - Egress
      Blackhole,
b - Ingress Blackhole, v - MAC-Based VLAN, P - Private VLAN, T - VLAN
      translation,
D - drop packet, h - Hardware Aging, o - IEEE 802.1ah Backbone MAC.
Total: 3 Static: 0 Perm: 0 Dyn: 3 Dropped: 0 Locked: 0 Locked with
Timeout: 0
FDB Aging time: 300
FDB VPLS Aging time: 300

```

The following example displays output where the port tag is displayed in parentheses:

```

# show fdb
Mac                Vlan      Age  Flags      Port / Virtual Port List
-----
00:00:00:00:04:0a  test(0200) 0057 d m      3(0010)
00:00:00:00:04:0b  test(0200) 0300 d m      3(0011)
00:01:02:03:04:05  test(0200) 0000 spm     3(0010)

Flags : d - Dynamic, s - Static, p - Permanent, n - NetLogin, m - MAC, i - IP,
      x - IPX, l - lockdown MAC, L - lockdown-timeout MAC, M- Mirror, B -
      Egress
      Blackhole,
      b - Ingress Blackhole, v - MAC-Based VLAN, P - Private VLAN, T - VLAN
      translation,
      D - drop packet, h - Hardware Aging, o - IEEE 802.1ah Backbone MAC.
Total: 3 Static: 0 Perm: 0 Dyn: 3 Dropped: 0 Locked: 0 Locked
with Timeout: 0
FDB Aging time: 300
FDB VPLS Aging time: 300

```

History

This command was first available in ExtremeXOS 10.1.

The stats and netlogin parameters were first available in ExtremeXOS 11.3.

The blackhole output under the b and B flags was first available for all platforms in ExtremeXOS 12.1.

The o flag was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.



show fdb mac-tracking configuration

show fdb mac-tracking configuration

Description

Displays configuration information for the MAC address tracking feature.

Syntax Description

This command has no arguments or variables.

Default

The MAC address tracking table is empty.

Usage Guidelines

None.

Example

The following command example displays the contents of the MAC address tracking table:

```
Switch.8 # show fdb mac-tracking configuration
MAC-Tracking enabled ports: 1-3,10,20
SNMP trap notification      : Enabled
MAC address tracking table (4 entries):
00:30:48:72:ee:88
00:21:9b:0e:ca:32
00:12:48:82:9c:56
00:30:48:84:d4:16
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.

show fdb mac-tracking statistics

show fdb mac-tracking statistics {*mac_addr*} {**no-refresh**}



Description

Displays statistics for the MAC addresses that are being tracked.

Syntax Description

<i>mac_addr</i>	Specifies a MAC address, using colon-separated bytes, for which FDB entries should be displayed.
no-refresh	Specifies a static snapshot of data instead of the default dynamic display.

Default

N/A.

Usage Guidelines

Use the keys listed below the display to clear the statistics counters or page up or down through the table entries.

Example

The following command example displays statistics for the entries in the MAC address tracking table:

```
Switch.1 # show fdb mac-tracking statistics
MAC Tracking Statistics      Fri Mar 20 15:25:01 2009
Add      Move      Delete
MAC Address      events      events      events
=====
00:00:00:00:00:01      0           0           0
00:00:00:00:00:02      0           0           0
00:00:00:00:00:03      0           0           0
00:00:00:00:00:04      0           0           0
00:00:00:00:00:05      0           0           0
00:00:00:00:00:06      0           0           0
00:00:00:00:00:07      0           0           0
00:00:00:00:00:08      0           0           0
00:00:00:00:00:09      0           0           0
00:00:00:00:00:10      0           0           0
00:00:00:00:00:11      0           0           0
00:00:00:00:00:12      0           0           0
00:00:00:00:00:13      0           0           0
00:00:00:00:00:14      0           0           0
00:00:00:00:00:15      0           0           0
00:00:00:00:00:16      0           0           0
00:00:00:00:00:17      0           0           0
00:00:00:00:00:18      0           0           0
=====
0->Clear Counters  U->page up  D->page down  ESC->exit
```



History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.

show fdb static-mac-move configuration

show fdb static-mac-move configuration

Description

Displays the configuration for the feature that reports the discovery of MAC addresses that are duplicates of statically configured MAC addresses.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following example shows the command display:

```
Switch.37 # show fdb static-mac-movement configuration
Static MAC Movement Notification: Enabled
MAC learning Packets Count       : 5
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on Summit family switches.



show fdb stats

```
show fdb stats {{ports {all | port_list} | vlan {all} | {vlan} vlan_name } {no-refresh}}
```

Description

Displays FDB entry statistics for the specified ports or VLANs in either a dynamic or a static report.

Syntax Description

all	Requests statistics for all ports or all VLANs.
<i>port_list</i>	Specifies which ports are to be included in the statistics display.
<i>vlan_name</i>	Specifies a single VLAN to be included in the statistics display.
no-refresh	Specifies a static display, which is not automatically updated.

Default

Summary FDB statistics for the switch.

Usage Guidelines

The dynamic display remains visible and continues to update until you press [Esc].

The show fdb stats command output displays the following information:

Port	When you chose to display statistics for ports, this column displays port numbers.
Link State	When you chose to display statistics for ports, this column displays the link states, which are described at the bottom of the display.
VLAN	When you chose to display statistics for VLANs, this column displays VLAN names.
MAC Addresses	This column displays the total number of MAC addresses for each port or VLAN.
Dynamic	This column displays the total number of MAC addresses that were learned dynamically for each port or VLAN.
Static	This column displays the total number of MAC addresses that are configured on this switch for each port or VLAN.
Dropped	This column displays the total number of dynamic MAC addresses that were discovered, but not stored in the FDB. Discovered MAC addresses might be dropped because a configured learning limit is reached, the FDB is in lockdown, or a port forwarding state is in transition. Some conditions that lead to dropped MAC addresses can produce log messages or SNMP traps.



Example

The following command example displays summary FDB statistics for the switch:

```
torino1.1 # show fdb stats
Total: 4 Static: 3 Perm: 3 Dyn: 1 Dropped: 0
FDB Aging time: 300
FDB VPLS Aging time: 300
torino1.2 #
```

The following command example displays FDB statistics for ports 1 to 16 on slot 1:

```
# show fdb stats ports 1:1-1:16
FDB Stats                               Mon Mar 15 15:30:49 2010
Port      Link  MAC
State     Addresses  Dynamic      Static      Dropped
=====
1:1       A      2394        2389        5           2
1:2       A      37          37          0           0
1:3       A      122         121         1           452
1:4       R      0           0           0           0
1:5       R      0           0           0           0
1:6       A      43          43          0           0
1:7       A      118         118         0           0
1:8       R      0           0           0           0
1:9       R      0           0           0           0
1:10      A      8           8           0           0
1:11      A      2998        2990        8           1
1:12      A      486         486         0           0
1:13      R      0           0           0           0
1:14      A      42          42          0           0
1:15      A      795         795         0           0
1:16      A      23          23          0           2
=====
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
U->page up  D->page down  ESC->exit
```

The following command example displays FDB statistics for all VLANs:

```
# show fdb stats vlan all
FDB Stats                               Mon Mar 15 15:30:49 2010
VLAN      MAC Addresses  Dynamic      Static      Dropped
=====
SV_PPPOE      2394        2389        5           2
NV_PPPOE      122         121         1           452
=====
U->page up  D->page down  ESC->exit
```

History

The dynamic display for this command was first available in ExtremeXOS 12.4.2.



Platform Availability

This command is available on all platforms.



17 Data Center Solution Commands

```
show vm-tracking repository
configure fip snooping add vlan
configure fip snooping add fcf
configure fip snooping delete vlan
configure fip snooping delete fcf
configure fip snooping fcf-update
configure fip snooping fcmmap
configure fip snooping port location
configure lldp ports dcbx add application
configure lldp ports dcbx delete application
configure lldp ports vendor-specific dcbx
configure port reflective-relay
configure vlan dynamic-vlan uplink-ports
configure vm-tracking authentication database-order
configure vm-tracking blackhole
configure vm-tracking local-vm
configure vm-tracking nms
configure vm-tracking nms timeout
configure vm-tracking repository
configure vm-tracking timers
configure vm-tracking vpp add
configure vm-tracking vpp delete
configure vm-tracking vpp vlan-tag
configure vm-tracking vpp counters
create vm-tracking local-vm
create vm-tracking vpp
delete vm-tracking local-vm
delete vm-tracking vpp
disable fip snooping
disable vm-tracking
disable vm-tracking dynamic-vlan ports
disable vm-tracking ports
enable fip snooping
enable vm-tracking
enable vm-tracking dynamic-vlan ports
enable vm-tracking ports
run vm-tracking repository
```

```

show fip snooping access-list
show fip snooping counters
show fip snooping enode
show fip snooping fcf
show fip snooping virtual-link
show fip snooping vlan
show lldp dcbx
show vlan dynamic-vlan
show vm-tracking
show vm-tracking local-vm
show vm-tracking network-vm
show vm-tracking nms
show vm-tracking port
show vm-tracking repository
show vm-tracking vpp
unconfigure vm-tracking local-vm
unconfigure vm-tracking repository
unconfigure vm-tracking vpp vlan-tag
unconfigure vm-tracking vpp
unconfigure vm-tracking nms

```

This chapter describes commands for:

- Managing The Extreme Network Virtualization (XNV) feature
- Managing the Direct Attach feature
- Using FIP Snooping

For an introduction to Data Center Solutions, see the ExtremeXOS Concepts Guide.

show vm-tracking repository

```
show vm-tracking repository {primary | secondary}
```

Description

Displays the FTP file synchronization configuration for NVPP and VM MAP files.

Syntax Description

primary secondary	Specifies whether you are displaying the primary or secondary FTP server configuration.
----------------------------	---



Default

If you do not specify primary or secondary, the default action is to display both the primary and secondary FTP server configurations.

Usage Guidelines

None.

Example

The following command displays the configuration for the primary and secondary FTP servers:

```
show vm-tracking repository
Primary VM-Map FTP server:
Server name:
IP address      : 10.100.1.200
VR Name        : VR-Mgmt
Refresh-interval: 600 seconds
Path Name      : /pub (default)
User Name      : anonymous (default)
Secondary vm-map FTP server: Unconfigured
Last sync      : 16:35:15      Last sync server : Primary
Last sync status : Successful
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

configure fip snooping add vlan

```
configure fip snooping add {vlan} vlan_name
```

Description

FIP Snooping must be configured to operate.

When a VLAN is added to FIP Snooping using this command, a record containing all FIP configuration information is created for that VLAN with default settings for all configuration elements. If the configuration is saved, the record persists across reboots. The user can see the record when using the “show fip snooping vlan” command to see FIP Snooping information for a VLAN. If the record does not exist, no information appears.



Syntax Description

add	Allows use of FIP Snooping on the VLAN.
vlan	Optional VLAN keyword.

Default

- Feature is disabled on the specified VLAN.
- Port locations default to “perimeter”.
- FCF-update mode is “automatic”.
- No FCFs exist in the configuration.
- The FC-MAP prefix is 0e:fc:00.
- There are no ACLs.
- There are no ENodes.
- There are no virtual links.
- All counters contain zero.

Usage Guidelines

This command creates the FIP Snooping configuration record for the specified VLAN. All default settings are in effect.

Example

```
configure fip snooping add vlan v3
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the following platforms:

- BlackDiamond X8
- BlackDiamond 8800 series BD8900-40G6X-c
- Summit X670
- Summit X770

configure fip snooping add fcf

```
configure fip snooping {vlan} vlan_name add fcf mac_addr port port
```



Description

This command is used to add an FCF to a FIP Snooping VLAN port when in manual fcf-update mode.

If the fcf-update mode is manual, this command adds a new FCF MAC to the list of FCFs. The command does not allow the same FCF MAC to be added to multiple ports in the same VLAN.

When a new FCF is added, ACLs are added to accept FIP frames from the new FCF.

An FCF can only be configured on a FIP Snooping VLAN port that has port location FCF-to-Enode or All configured.

If the fcf-update mode is automatic and this command is executed, the add is not allowed and the user is informed.

Syntax Description

fip	FCoE Initialization Protocol.
snooping	Snooping on FIP frames.
<i>vlan_name</i>	Name of a FIP Snooping VLAN where fcf-update is configured to be in manual mode.
add	Add to the list of FCoE forwarders.
<i>mac_addr</i>	MAC address of the FCoE Forwarder specified in the format of hh:hh:hh:hh:hh:hh.
<i>port</i>	Port through which the FCF is reachable.

Default

N/A.

Usage Guidelines

This command is used to add an FCF to a FIP Snooping VLAN port when in manual fcf-update mode. The command does not allow the same FCF MAC to be added to multiple ports in the same VLAN.

Example

```
configure fip snooping v3 add fcf aa:bb:cc:dd:00:00 port 1:2
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the following platforms:



- BlackDiamond X8
- BlackDiamond 8800 series BD8900-40G6X-c
- Summit X670
- Summit X770

configure fip snooping delete vlan

```
configure fip snooping delete [{vlan} vlan_name | all]
```

Description

This command deletes the FIP Snooping configuration record for the specified VLAN. If “all” is specified, all FIP Snooping configuration information is removed from the system after the user confirms this request:

Warning: This command will remove all FIP Snooping configuration for all VLANs. Do you want to continue? (y/N)

If FIP Snooping is enabled on an affected VLAN it is first disabled causing the removal of related ACL and FDB information from the system. Also removed are any virtual links, Enodes, and FCFs.

Note



A VLAN cannot be deleted when FIP Snooping is configured. For example: * BDx8.60 # delete vlan v1 Error: Failed to delete VLAN v1; FIP Snooping is configured on this VLAN. Configuration failed on backup MM, command execution aborted! * BDx8.61 # configure fip snooping delete vlan v1 * BDx8.62 # delete vlan v1 * BDx8.63 #

Syntax Description

delete	Remove use of FIP Snooping from the VLAN.
vlan	Optional VLAN keyword

Default

N/A.

Usage Guidelines

Use this command to delete the FIP Snooping configuration record for the specified VLAN.

Example

```
configure fip snooping delete vlan v3
```



History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the following platforms:

- BlackDiamond X8
- BlackDiamond 8800 series BD8900-40G6X-c
- Summit X670
- Summit X770

configure fip snooping delete fcf

```
configure fip snooping {vlan} vlan_name delete fcf mac_addr port port
```

Description

This command is used to remove an FCF from a FIP Snooping VLAN port when in manual fcf-update mode.

If the fcf-update mode is manual, this command removes the FCF MAC from the list of FCFs configured on the FIP Snooping VLAN. When an FCF is removed from the list, the ACLs referencing the FCF (including virtual links) are removed.

If the fcf-update mode is automatic and this command is executed, the remove is not allowed and the user is informed.

Syntax Description

fip	FCoE Initialization Protocol.
snooping	Snooping FIP frames.
<i>vlan_name</i>	Name of a FIP Snooping VLAN where the specified FCF MAC address has been configured.
delete	Delete from the list of FCoE forwarders.
fcf	The list of FCoE forwarders in the VLAN.
<i>mac_addr</i>	MAC address of the FCoE Forwarder specified in the format of hh:hh:hh:hh:hh:hh.
<i>port</i>	Port through which the FCF is reachable.

Default

N/A.



Usage Guidelines

This command is used to remove an FCF from a FIP Snooping VLAN port when in manual fcf-update mode. The fcf-update mode must be “manual”.

Example

```
configure fip snooping v3 delete fcf aa:bb:cc:dd:00:00 port 1:2
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the following platforms:

- BlackDiamond X8
- BlackDiamond 8800 series BD8900-40G6X-c
- Summit X670
- Summit X770

configure fip snooping fcf-update

```
configure fip snooping {vlan} vlan_name fcf-update [auto | manual]
```

Description

This command configures the update mode of the list of FCFs per FIP Snooped VLAN. The default mode is auto.

The list of FCFs to which ENodes establish FCoE virtual links is updated either administratively or dynamically via snooped FIP frames. This command selects the method of updating the list of FCFs per VLAN. When the updating method changes, the following events occur.

- FDB entries of FCFs' MACs are removed.
- ACLs checking the FCFs' MACs are removed.

In automatic mode, the list of FCFs is automatically constructed through observation of FCF discovery advertisement packets. An attempt to configure an FCF while in automatic mode is rejected.

In manual mode the list of FCFs is configured by the user. Use the following commands to configure the list of FCFs:

- `configure fip snooping add fcf`
- `configure fip snooping delete fcf`
- `configure fip snooping fcf-update`

When the fcf-update mode is changed from manual to automatic, all configured FCFs are removed.



Syntax Description

<i>vlan_name</i>	Name of a FIP Snooping VLAN where the fcf-update mode is to be configured.
auto	Learn the list of FCoE forwarders from snooped FIP frames.
manual	FCoE forwarders are configured manually using the “configure fip snooping <i>vlan</i> add fcf” command.

Default

Auto.

Usage Guidelines

This command configures the update mode of the list of FCFs per FIP Snooped VLAN. In automatic mode, the list of FCFs is automatically constructed through observation of FCF discovery advertisement packets. An attempt to configure an FCF while in automatic mode is rejected. In manual mode the list of FCFs is configured by the user. When the fcf-update mode is changed, all FCFs are removed.

Example

```
configure fip snooping vlan v3 fcf-update manual
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the following platforms:

- BlackDiamond X8
- BlackDiamond 8800 series BD8900-40G6X-c
- Summit X670
- Summit X770

configure fip snooping fcmmap

```
configure fip snooping {vlan} vlan_name fcmmap mac_prefix
```

Description

An FCF and an Enode negotiate whether the FCF or the Enode will provide a VN_Port MAC address for each virtual link. The Enode (also called a server) can provide a Server Provided MAC Address (SPMA), or the FCF can provide a Fabric Provided MAC Address (FPMA). An individual FPMA is assigned by the



FibreChannel fabric to the VN_Port during fabric login. An FPMA address begins with the 24-bit FC-MAP prefix. The default value of the FC-MAP prefix is 0E:FC:00 but can be changed. The low order three octets of the FPMA will contain the FibreChannel fabric-assigned FibreChannel ID (also called a VN_Port_ID) for the virtual link.

This command configures the expected MAC address prefix (used when in FPMA mode) of all FPMA used on the FIP Snooping VLAN. The FPMA for a VN_Port is assigned by the FCF using its configured FC-MAP prefix to construct the VN_Port FPMA. Therefore the FC-MAP prefix configured on the switch must be the same as that configured on the FCF for the VLAN. The default value of *mac_prefix* is 0E:FC:00:00:00:00.

The *mac_prefix* value must be between 0e:fc:00 and 0e:fc:ff and the lower three MAC octets must be specified as zero or the following message will be displayed:

```
Error: Invalid FC-MAP, use 0e:fc:xx:00:00:00 where xx is a two-digit hexadecimal value.
```

The user should not use the same FIP Snooping VLAN for connection to more than one FibreChannel fabric or storage area network. Duplicate FPMA could be assigned by the different fabrics causing connectivity issues.

Syntax Description

fip	FCoE Initialization Protocol.
snooping	Snooping FIP frames.
<i>vlan_name</i>	Name of the VLAN for which the FC-MAP prefix is to be changed.
fcmap	24 prefix for MAC address assigned to VN Port in an FPMA mode.
<i>mac_prefix</i>	24bit prefix of MAC followed by 24 zeros formatted as 0e:fc:xx:00:00:00 where xx is a two-digit hexadecimal number.

Default

The default value of *mac_prefix* is 0E:FC:00:00:00:00.

Usage Guidelines

This command configures the expected MAC address prefix (used when in FPMA mode) of all FPMA used on the FIP Snooping VLAN.

Example

```
configure fip snooping vlan v3 fcmap 0e:fc:01:00:00:00
```

History

This command was first available in ExtremeXOS 15.1.



Platform Availability

This command is available on the following platforms:

- BlackDiamond X8
- BlackDiamond 8800 series BD8900-40G6X-c
- Summit X670
- Summit X770

configure fip snooping port location

```
configure fip snooping {vlan} vlan_name ports port_list {location [perimeter | enode-to-fcf | fcf-to-enode | all]}
```

The default ACLs on the port are changed to be consistent with the new location.

Description

This command configures the port location for a member of a VLAN that is configured to perform FIP Snooping. The default port location type is perimeter. If no FIP Snooping configuration record was previously created for the VLAN, this command causes its creation with defaults (except for the particular port's location as specified) set.

The acceptable FIP frames differ per port location. The command specifies the port location and guides the switch to install different ACLs. The default port location, i.e. port type, is perimeter, where the port is expected to be connected to ENodes. The change of the port type triggers the following events.

If FIP Snooping is enabled:

- All FDB entries previously stored for the VLAN on the specified port are removed, except for those related to manually configured FCFs.
- All virtual links are removed.
- All knowledge of ENodes (if any) learned on this port is removed.
- All knowledge of discovered FCFs (if any) learned on this port is removed.

Syntax Description

<i>vlan-name</i>	Name of the VLAN whose port(s) will have the location changed.
<i>port_list</i>	A port or a list of ports.
perimeter	Port is directly connected to ENodes. Per virtual link ACLs are installed providing the most security.
enode-to-fcf	Port sees packets from FCoE nodes to FCoE forwarders only.
fcf-to-enode	Port sees packets from FCoE forwarders to FCoE nodes only.
all	Port sees packets both from FCoE forwarders and FCoE nodes.

Default

Perimeter.



Usage Guidelines

This command configures the port location for a member of a VLAN that is to perform FIP Snooping.

Example

```
configure fip snooping vlan "v3" port 1:1 location fcf-to-encode
```

History

This command is available on the following platforms:

- BlackDiamond X8
- BlackDiamond 8800 series BD8900-40G6X-c
- Summit X670
- Summit X770

configure lldp ports dcbx add application

```
configure lldp ports [all | port_list] dcbx add application [name  
application_name | ethertype ethertype_value | L4-port port_number | tcp-port  
port_number | udp-port port_number] priority priority_value
```

Description

Configures an application priority to be advertised to DCBX end stations.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>application_name</i>	Specifies an application. Supported values are: <ul style="list-style-type: none"> • fcoe—Fiber Channel Over Ethernet (FCoE). • fip—FCoE Initiation Protocol (FIP). • iscsi—Internet Small Computer System Interface (iSCSI).
<i>ethertype_value</i>	Specifies an ethertype value in the range of 1536 to 65535.
L4-port <i>port_number</i>	Specifies a Layer 4 port number in the range of 0 to 65535. Supported Layer4 protocols include TCP, SCTP, UDP, and DCCP.
tcp-port <i>port_number</i>	Specifies a TCP port number in the range of 0 to 65535.
udp-port <i>port_number</i>	Specifies a UDP port number in the range of 0 to 65535.
<i>priority_value</i>	Specifies a priority in the range of 0 to 7.



Default

N/A.

Usage Guidelines

This command configures the switch to advertise the priority that an end station should use for the specified application or port number. The priority number is mapped to an 802.1p value, which determines how the switch manages traffic from that application or port.

The switch supports a maximum of 8 DCBX applications per port. If an application configuration already exists on the specified port or ports, the priority is updated to the new value. If the maximum number of applications for a port is exceeded, the switch logs an error message.

Example

The following command configures the switch to advertise priority 4 for the iSCSI application on ports 1 to 24:

```
configure lldp ports 1-24 dcbx add application name iscsi priority 4
```

The following command configures the switch to advertise priority 3 for ethertype value 34525 on port 1:

```
configure lldp ports 1 dcbx add application ethertype 34525 priority 3
```

The following command configures the switch to advertise priority 6 for Layer 4 port 992 on port 1:

```
configure lldp ports 1 dcbx add application L4-port 992 priority 6
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on all platforms.

configure lldp ports dcbx delete application

```
configure lldp ports [all | port_list] dcbx delete application [all-applications | name application_name | ethertype ethertype_value | L4-port port_number | tcp-port port_number | udp-port port_number]
```



Description

Removes the priority configuration for one or all applications from the specified ports.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>application_name</i>	Specifies an application. Supported values are: <ul style="list-style-type: none"> • fcoe—Fiber Channel Over Ethernet (FCoE). • fip—FCoE Initiation Protocol (FIP). • iscsi—Internet Small Computer System Interface (iSCSI).
<i>ethertype_value</i>	Specifies an ethertype value in the range of 1536 to 65535.
L4-port <i>port_number</i>	Specifies a Layer 4 port number in the range of 0 to 65535. Supported Layer4 protocols include TCP, SCTP, UDP, and DCCP.
tcp-port <i>port_number</i>	Specifies a TCP port number in the range of 0 to 65535.
udp-port <i>port_number</i>	Specifies a UDP port number in the range of 0 to 65535.

Default

N/A.

Usage Guidelines

This command configures the switch to advertise the priority that an end station should use for the specified application or port number. The priority number is mapped to an 802.1p value, which determines how the switch manages traffic from that application or port.

If an application configuration already exists on the specified port or ports, the priority is updated to the new value.

Example

The following command removes the priority configuration for Layer 4 port 30 on port 23:

```
configure lldp ports 23 dcbx delete application L4-port 30
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on all platforms.



configure lldp ports vendor-specific dcbx

```
configure lldp ports [all | port_list] [advertise | no-advertise] vendor-specific dcbx {ieee|baseline}
```

Description

Configures the LLDP port to advertise or not to advertise Data Center Bridging Exchange (DCBX) information to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.
ieee	Specifies the DCBX protocol defined in IEEE 802.1Qaz.
baseline	Specifies the DCBX protocol known as Baseline Version 1.01, which was defined before IEEE 802.1Qaz.

Default

No advertisement for both DCBX protocols.

Usage Guidelines

If you do not specify a protocol with this command, the advertise option enables advertisement for the IEEE 802.1Qaz protocol, and the no-advertise option disables advertisement for both protocols.

Example

The following command advertises DCBX information according to IEEE 802.1Qaz for port 1:5:

```
configure lldp ports 1:5 advertise vendor specific dcbx
```

The following command advertises DCBX information according to Baseline Version 1.01 for port 2:1:

```
configure lldp ports 2:1 advertise vendor specific dcbx baseline
```

The following command disables advertisement of DCBX information on all ports:

```
configure lldp ports all no-advertise vendor specific dcbx
```

History

This command was first available in ExtremeXOS 12.6.



Platform Availability

This command is available on all platforms.

configure port reflective-relay

```
configure port port reflective-relay [on | off]
```

Description

Enables the direct attach feature on the specified port.

Syntax Description

<i>port</i>	Specifies a single port on which to enable the direct attach feature.
-------------	---

Default

Off.

Usage Guidelines

You should only enable the direct attach feature on ports that directly connect to a VM server running VEPA software.

This feature requires installation of the Direct Attach feature pack. For more information, see [Feature Pack Features](#).

Example

The following command enables the direct attach feature on port 2:1:

```
configure port 2:1 reflective-relay on
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all Summit family switches and BlackDiamond X8, BlackDiamond 8000 series modules.

configure vlan dynamic-vlan uplink-ports




```
configure vlan dynamic-vlan uplink-ports [ add {ports} port_list | delete {ports}
[port_list | all] ]
```

Description

Statically provisions uplink ports for all dynamically created VLANs.

Syntax Description

dynamic-vlan	Configuration options for dynamically created VLANs.
uplink-ports	Tagged uplink ports for VLANs created by EXOS.
add	Add ports to dynamic VLAN uplink port list.
delete	Remove ports from dynamic VLAN uplink port list.
ports	Ports to be configured as uplink ports.
<i>port_list</i>	List of ports separated by a comma or "-" ;type=portlist_t";
all	Clear the dynamic VLAN uplink port list.

Default

N/A.

Usage Guidelines

Use this command to statically provision uplink ports for dynamically created VLANs.

Example

```
X460-48p.3 # conf vlan dynamic-vlan uplink-ports add ports 16-18X460-48p.4 #
conf vlan dynamic-vlan uplink-ports add 20,22,24X460-48p.5 # configure vlan
dynamic-vlan uplink-ports delete ports 22X460-48p.7 # configure vlan dynamic-
vlan uplink-ports delete 16-18X460-48p.8 # configure vlan dynamic-vlan uplink-
ports delete all
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

configure vm-tracking authentication database-order



```
configure vm-tracking authentication database-order [[nms] | [vm-map] | [local] | [nms local] | [local nms] | [nms vm-map] | [vm-maplocal] | [local vm-map] | [nms vm-map local] | [localnmsvm-map]]
```

Description

Configures the authentication database options and sequence for VM authentication.

Syntax Description

nms	Specifies the configured Network Management System (NMS).
vm-map	Specifies the configured VMMAP file.
local	Specifies the configured local database.

Default

nms vm-map local.

Usage Guidelines

The switch attempts VM authentication in the sequence specified. For example, in the default configuration, the switch attempts NMS authentication first, VMMAP authentication second, and local authentication third. If nms is specified, the switch always attempts NMS authentication before attempting VMMAP file authentication.

Example

The following command configures the database authentication order:

```
configure vm-tracking authentication database-order local nms vm-map
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

configure vm-tracking blackhole

```
configure vm-tracking blackhole [policy policy_name | dynamic-rule rule_name | none]
```



Description

Specifies a policy file or dynamic ACL rule to apply to VMs during periods that are outside of the approved time slot for that VM.

Syntax Description

<i>policy_name</i>	Specifies the name of a policy file to apply to the VM authentication request.
<i>rule_name</i>	Specifies the name of an ACL rule to apply to the VM authentication request.

Default

N/A.

Usage Guidelines

This command is not supported in this software release. It will be supported in a future release.

The none option applies no policy name or ACL rule during periods that are outside of the approved time slot for that VM.



Note

This command is provided to support future identity management features. It serves no practical purpose in this release.

Example

The following command applies no policy name or ACL rule during periods that are outside of the authorized authentication period:

```
configure vm-tracking blackhole none
```

History

This command was first visible in ExtremeXOS 12.5, but it is not supported in this release. This command will be supported in a future release.

Platform Availability

This command is available on all platforms.

configure vm-tracking local-vm

```
configure vm-tracking local-vm mac-address mac [name name | ip-address ipaddress  
| vpp vpp_name] | vlan-tag tag {vr vr_name}]
```



Description

Configures the parameters associated with a local VM database entry to be used for VM MAC local authentication.

Syntax Description

<i>mac</i>	Specifies the MAC address for the VM database entry you want to configure.
<i>name</i>	Specifies a name to represent this VM in show vm-tracking command display.
<i>ipaddress</i>	Specifies the IP address for the VM. This must match the IP address configured on the VM.
<i>vpp_name</i>	Specifies the name of a VPP to apply to the local VM.
<i>tag</i>	VLAN tag between 1 and 4094.
<i>vr_name</i>	Virtual router name.

Default

N/A.

Usage Guidelines

Before you configure a VM entry in the local VM database, you must create the entry with the `create vm-tracking local-vm` command.

Before you assign an VPP to a VM entry in the local VM database, you must create the VPP with the `create vm-tracking vpp` command.

Example

The following command configures an IP address for the VM entry specified by the MAC address:

```
configure vm-tracking local-vm mac-address 00:E0:2B:12:34:56 ip-address
10.10.10.1
```

History

This command was first available in ExtremeXOS 12.5.

The **ingress-vpp** and **egress-vpp** options were replaced with the **vpp** option in ExtremeXOS 12.6.

The **vlan-tag** and **vr-name** options were added in 15.3.

Platform Availability

This command is available on all platforms.



configure vm-tracking nms

```
configure vm-tracking nms [primary | secondary] server [ipaddress | hostname]
{udp_port} client-ip client_ip shared-secret {encrypted} secret {vr vr_name}
```

Description

Configures the switch RADIUS client to an NMS for VM authentication.

Syntax Description

primary secondary	Specifies the whether you are configuring the primary or secondary NMS.
<i>ipaddress</i>	Specifies the NMS IP address.
<i>hostname</i>	Specifies the NMS DNS hostname.
<i>udp_port</i>	Specifies the UDP port number of the NMS application.
<i>client_ip</i>	Specifies the client IP address, which is the switch IP address on the interface leading to the NMS.
encrypted	Specifies that the secret key for communications with the NMS is encrypted.
<i>secret</i>	Specifies a key or password for communications with the NMS.
<i>vr_name</i>	Specifies the VR that is used to access the NMS.

Default

N/A.

Usage Guidelines

The NMS is a RADIUS server such as the one provided with Ridgeline.

Example

The following command configures the switch to authenticate VMs through the primary NMS server Ridgeline using the password password:

```
configure vm-tracking nms primary server Ridgeline client-ip 10.10.3.3 shared-
secret password
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.



configure vm-tracking nms timeout

```
configure vm-tracking nms timeout seconds
```

Description

Configures the timeout period for authentication attempts with the configured NMS servers.

Syntax Description

<i>seconds</i>	Specifies the timeout period in seconds.
----------------	--

Default

3 seconds.

Usage Guidelines

None.

Example

The following command configures the switch to allow 1 minute for successful authentication of a VM with the NMS server:

```
configure vm-tracking nms timeout 60
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

configure vm-tracking repository

```
configure vm-tracking repository [primary | secondary] server [ipaddress |  
hostname] {vr vr_name} {refresh-interval seconds} {path-name path_name} {user  
user_name {encrypted} password}
```

Description

Configures FTP file synchronization for NVPP and VMMAF files.



Syntax Description

primary secondary	Specifies the whether you are configuring the primary or secondary FTP server.
<i>ipaddress</i>	Specifies the FTP server IP address.
<i>vr_name</i>	Specifies the VR that is used to access the FTP server.
<i>seconds</i>	Specifies how often the switch updates the local files that are synchronized with the FTP server. The range is 40 to 3600 seconds.
<i>path_name</i>	Specifies the path to the repository server files from the FTP server root directory. The default directory for repository server files is: pub.
<i>user_name</i>	Specifies a user name for FTP server access. If no username is specified, the switch uses user name anonymous.
encrypted	This keyword indicates that the specified password is encrypted.
<i>password</i>	Specifies the password for the specified user name.

Default

Refresh interval: 600 seconds.

Usage Guidelines

Some jitter is added to the refresh interval period to prevent all switches from downloading files at the same time.

Example

The following command configures the switch to refresh the VM MAP and NVPP files from primary FTP server ftp1 every 5 minutes:

```
configure vm-tracking repository primary server ftp1 refresh-interval 300
```

History

This command was first available in ExtremeXOS 12.5.

Support for specifying an FTP user name was added in ExtremeXOS 12.6.

Platform Availability

This command is available on all platforms.

configure vm-tracking timers

```
configure vm-tracking timers reauth-period reauth_period
```



Description

Configures the RADIUS reauthentication period for VM MAC addresses.

Syntax Description

<i>reauth_period</i>	Specifies the reauthentication period in seconds. The ranges are 0 and 30-7200 seconds.
----------------------	---

Default

0 seconds.

Usage Guidelines

One way to periodically apply Virtual Port Profiles (VPPs) to VM MAC addresses is to configure a reauthentication period. At the end of each reauthentication period, the switch reauthenticates each VM MAC address and applies any updated VPPs.

This command applies to only those VMs that authenticate through RADIUS. Reauthentication is disabled when the reauthentication period is set to 0 seconds. When reauthentication is disabled, the VM MAC address remains authenticated until the FDB entry for that VM expires.

Example

The following command enables RADIUS server reauthentication at 2 minute intervals:

```
configure vm-tracking timers reauth-period 120
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

configure vm-tracking vpp add

```
configure vm-tracking vpp vpp_name add [ingress | egress] [policy policy_name |
dynamic-rule rule_name] {policy-order policy_order}
```

Description

Configures an LVPP to use the specified policy or ACL rule.



Syntax Description

<i>vpp_name</i>	Specifies the name of an existing LVPP.
add	Specifies whether the LVPP should start using the specified policy or rule.
ingress	Specifies that the policy mapped to the LVPP is for ingress traffic.
egress	Specifies that the policy mapped to the LVPP is for egress traffic.
<i>policy_name</i>	Specifies a policy to add to or delete from the LVPP.
<i>rule_name</i>	Specifies a dynamic ACL rule to add to or delete from the LVPP.

Default

N/A.

Usage Guidelines

Multiple ACL or policy files can be mapped to each LVPP. A maximum of 8 ingress and 4 egress ACL or policies are available to be mapped to each LVPP. If the policy file or dynamic rule specified in this command fails to bind, then the CLI command is rejected.

Before you can configure an LVPP, you must first create it with the `create vm-tracking vpp` command.

Example

The following command configures LVPP vpp1 to use the dynamic ACL rule named rule1 for ingress traffic:

```
configure vm-tracking vpp vpp1 add ingress dynamic-rule rule1
```

History

This command was first available in ExtremeXOS 12.5.

The ingress and egress keywords were added in ExtremeXOS 12.6.

Platform Availability

This command is available on all platforms.

configure vm-tracking vpp delete

```
configure vm-tracking vpp vpp_name delete [ingress | egress] [policy policy_name
| dynamic-rule rule_name] {policy-order policy_order}
```



Description

Specifies that the LVPP should stop using the specified policy or rule.

Syntax Description

<i>vpp_name</i>	Specifies the name of an existing LVPP.
add	Specifies whether the LVPP should stop using the specified policy or rule.
ingress	Specifies that the policy mapped to the LVPP is for ingress traffic.
egress	Specifies that the policy mapped to the LVPP is for egress traffic.
<i>policy_name</i>	Specifies a policy to add to or delete from the LVPP.
<i>rule_name</i>	Specifies a dynamic ACL rule to add to or delete from the LVPP.

Default

N/A.

Usage Guidelines

Multiple ACL or policy files can be mapped to each LVPP. A maximum of 8 ingress and 4 egress ACL or policies are available to be mapped to each LVPP. If the policy file or dynamic rule specified in this command fails to bind, then the CLI command is rejected.

Before you can configure an LVPP, you must first create it with the `create vm-tracking vpp` command.

Example

The following command configures LVPP vpp1 to use the dynamic ACL rule named rule1 for ingress traffic:

```
configure vm-tracking vpp vpp1 add ingress dynamic-rule rule1
```

History

This command was first available in ExtremeXOS 12.5.

The ingress and egress keywords were added in ExtremeXOS 12.6.

Platform Availability

This command is available on all platforms.

configure vm-tracking vpp vlan-tag



```
configure vm-tracking vpp vpp_name vlan-tag tag {vr vr_name}
```

Description

This command configures the VLAN tag and VR name for VPP. If the detected VM MAC uses this VPP, then the port in which the VM MAC is detected will be placed on this VR/VLAN.

Syntax Description

<i>vpp_name</i>	Specifies a name for the LVPP to delete.
<i>tag</i>	
<i>vr_name</i>	

Default

N/A.

Usage Guidelines

Use this command to configure the VLAN tag and VR name for VPP. If the detected VM MAC uses this VPP, then the port in which the VM MAC is detected will be placed on this VR/VLAN.

Example

Example output not yet available and will be provided in a future release.

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

configure vm-tracking vpp counters

```
configure vm-tracking vpp vpp_name counters [ingress-only | egress-only | both | none]
```

Description

Configures whether counters need to be installed for Virtual Machine MAC which receives this VPP mapping.



Syntax Description

ingress-only	Only counts packets ingressing the switch whose source MAC address matches VM MAC.
egress-only	Only counts packets egressing the switch whose source MAC address matches VM MAC.
both	Counts packets ingressing and egressing the switch whose source MAC address matches VM MAC.
none	No packets will be counted.

Default

N/A.

Usage Guidelines

Use this command to configure whether counters need to be installed for Virtual Machine MAC which receives this VPP mapping.

Example

Example output not yet available and will be provided in a future release.

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms that support egress ACLs.

create vm-tracking local-vm

```
create vm-tracking local-vm mac-address mac {name name | ipaddress ipaddress vpp
vpp_name | vlan-tag tag {vr vr_name}}
```

Description

Creates a local VM database entry to be used for VM MAC local authentication, with optional parameters.



Syntax Description

<i>mac</i>	Specifies the MAC address for the VM. This must match the MAC address configured on the VM and be unique among the locally configure VM addresses.
<i>name</i>	Specifies a name to represent this VM in show vm-tracking command display.
<i>ipaddress</i>	Specifies the IP address for the VM. This must match the IP address configured on the VM.
<i>vpp_name</i>	Specifies the virtual port profile to apply for the local VM.
<i>tag</i>	VLAN tag between 1 and 4094.
<i>vr_name</i>	Virtual router name.

Default

N/A.

Usage Guidelines

A VM name can include up to 32 characters. VM names must begin with an alphabetical letter, and only alphanumeric, underscore (_), and hyphen (-) characters are allowed in the remainder of the name. VM names cannot match reserved keywords. For more information on VM name requirements and a list of reserved keywords, see [Object Names](#) of the *ExtremeXOS Concepts Guide*.

The following command creates a VM entry named VM1 in the local VM database:

```
create vm-tracking local-vm mac-address 00:E0:2B:12:34:56 name VM1
```

The following command creates a VM entry and assigns IP address 10.10.2.2 to the entry:

```
create vm-tracking local-vm mac-address 00:E0:2B:12:34:57 ip-address 10.10.2.2
```

The following command creates a VM entry and assigns VPP vpp1 to it:

```
create vm-tracking local-vm mac-address 00:E0:2B:12:34:58 vpp vpp1
```

History

This command was first available in ExtremeXOS 12.5.

The ingress-vpp and egress-vpp options were replaced with the vpp option in ExtremeXOS 12.6.

The vlan-tag and vr-name options were added in 15.3.

Platform Availability

This command is available on all platforms.



create vm-tracking vpp

```
create vm-tracking vpp vpp_name
```

Description

Creates a Local VPP (LVPP).

Syntax Description

<i>vpp_name</i>	Specifies a name for the new VPP.
-----------------	-----------------------------------

Default

N/A.

Usage Guidelines

A VPP name can include up to 32 characters. VPP names must begin with an alphabetical letter, and only alphanumeric, underscore (_), and hyphen (-) characters are allowed in the remainder of the name. VPP names cannot match reserved keywords. For more information on VPP name requirements and a list of reserved keywords, see [Object Names](#) of the ExtremeXOS Concepts Guide.

Example

The following command creates a VPP named vpp1:

```
create vm-tracking vpp vpp1
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

delete vm-tracking local-vm

```
delete vm-tracking local-vm {mac-address mac}
```

Description

Deletes the specified VM entry in the local VM database.



Syntax Description

<i>mac</i>	Specifies the MAC address for a VM entry to delete.
------------	---

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the VM entry for MAC address 00:E0:2B:12:34:56 in the local VM database:

```
delete vm-tracking local-vm mac-address 00:E0:2B:12:34:56
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

delete vm-tracking vpp

```
delete vm-tracking vpp {vpp_name}
```

Description

Deletes the specified LVPP.

Syntax Description

<i>vpp_name</i>	Specifies a name for the LVPP to delete.
-----------------	--

Default

N/A.



Usage Guidelines

None.

Example

The following command deletes the VPP named vpp1:

```
delete vm-tracking vpp vpp1
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

disable fip snooping

```
disable fip snooping {{vlan} vlan_name}
```

Description

This command disables FIP Snooping on one VLAN, or on all VLANs on which FIP Snooping is currently enabled.

Disabling FIP Snooping on a VLAN causes the following changes on that VLAN:

- All ACLs installed for the VLAN for FIP Snooping operation are removed.
- All FDB entries for the VLAN are removed.



Note

Depending on the activity of connected devices, some dynamic FDB entries may appear.

- All Enodes and virtual links learned on the VLAN are removed.
- If the fcf-update mode is automatic, all FCFs learned on the VLAN are removed.
- FDB learning is turned on for the VLAN.

Syntax Description

fip	FCoE Initialization Protocol.
snooping	Snooping FIP frames.
<i>vlan_name</i>	Name of the VLAN on which FIP Snooping is to be disabled.



Default

Disabled.

Usage Guidelines

Use this command to disable FIP Snooping in the VLAN. This command has no effect if executed on a VLAN for which no configuration record has been created. If a *vlan_name* is not specified, the command disables FIP Snooping on all VLANs on which it is enabled.

Example

```
disable fip snooping vlan v3
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the following platforms:

- BlackDiamond X8
- BlackDiamond 8800 series BD8900-40G6X-c
- Summit X670
- Summit X770

disable vm-tracking

```
disable vm-tracking
```

Description

Disables the Extreme Network Virtualization (XNV) feature on the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.



Usage Guidelines

This command disables the XNV feature, which tracks virtual machines (VMs) that connect to the switch.



Note

When the VM tracking feature is disabled, file synchronization with the FTP server stops.

Example

The following command disables the XNV feature:

```
disable vm-tracking
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

disable vm-tracking dynamic-vlan ports

```
disable vm-tracking dynamic-vlan ports port_list
```

Description

This command disables VM-tracking dynamic VLAN on specific ports.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Use this command to disable VM-tracking dynamic VLAN on specific ports. The ALL option is not supported because VM-tracking dynamic VLAN should not be enabled on a switch's uplink port.

Example

Example output not yet available and will be provided in a future release.



History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

disable vm-tracking ports

```
disable vm-tracking ports port_list
```

Description

Disables the XNV feature on the specified ports.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
------------------	---

Default

Disabled.

Usage Guidelines

This command disables VM tracking on the specified ports.

Example

The following command disables VM tracking on port 2:1:

```
disable vm-tracking ports 2:1
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

enable fip snooping

```
enable fip snooping {{vlan} vlan_name}
```



Description

This command enables FIP Snooping in the VLAN. If no VLAN is specified, FIP Snooping is enabled on all VLANs that have been added using the `configure fip snooping add the {vlan} vlan_name` command.

A FIP Snooping VLAN is disabled by default.

Once FIP Snooping is enabled on a VLAN, the following events occur:

- FDB learning is turned off for the VLAN.
- All FDB entries of the VLAN are removed. If FCFs are manually configured FDB entries are added for each such FCF.
- ACLs are installed to block most FIP and FCoE frames.

Syntax Description

fip	FCoE Initialization Protocol
snooping	Snooping FIP frames
<i>vlan_name</i>	Name of the vlan on which FIP Snooping is to be enabled.

Default

Disabled.

Usage Guidelines

This command enables FIP Snooping in the VLAN.

Example

```
enable fip snooping vlan v3
```

History

This command was first available in ExtremeXOS 15.1

Platform Availability

This command is available on the following platforms:

- BlackDiamond X8
- BlackDiamond 8800 series BD8900-40G6X-c
- Summit X670
- Summit X770



enable vm-tracking

enable vm-tracking

Description

Enables the XNV feature on the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables the XNV feature, which tracks VMs that connect to the switch.

This command does not enable XNV on any ports. To start tracking VMs, you must enable VM tracking on one or more ports using the `enable vm-tracking ports` command.

Example

The following command enables the XNV feature:

```
enable vm-tracking
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

enable vm-tracking dynamic-vlan ports

enable vm-tracking dynamic-vlan ports *port_list*

Description

This command enables VM-tracking dynamic VLAN on specific ports. The ALL option is not supported because VM-tracking dynamic VLAN should never be enabled on a switch's uplink port.



Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Use this command to enable VM-tracking dynamic VLAN on specific ports. The ALL option is not supported because VM-tracking dynamic VLAN should not be enabled on a switch's uplink port.

Example

Example output not yet available and will be provided in a future release.

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

enable vm-tracking ports

```
enable vm-tracking ports port_list
```

Description

Enables the XNV feature on the specified ports.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
------------------	---

Default

Disabled.

Usage Guidelines

You must enable VM tracking on the switch with the `enable vm-tracking` command before you can use this command. This command enables VM tracking on the specified ports. You should enable VM tracking only on ports that connect directly to a server that hosts VMs that you want to track. You should never enable VM tracking on a switch uplink port.



Example

The following command enables VM tracking on port 2:1:

```
enable vm-tracking ports 2:1
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

run vm-tracking repository

```
run vm-tracking repository sync-now
```

Description

Manually starts FTP file synchronization for NVPP and VMMAP files.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Before you can manually start FTP file synchronization, you must configure FTP servers using the `configure vm-tracking repository` command.

Example

The following command starts file synchronization with the configured FTP server:

```
run vm-tracking repository sync-now
```

History

This command was first available in ExtremeXOS 12.5.



Platform Availability

This command is available on all platforms.

show fip snooping access-list

```
show fip snooping {vlan} vlan_name access-list {[fcf mac_addr | virtual-link
mac_addr | all]}
```

Description

The command lists all the FCoE ACLs meeting the criteria.

The list can be shortened by specifying the MAC of an FCF or the VN_Port MAC assigned to a virtual link in the VLAN. The ACL with higher priority appears first.

By default, the command lists all the ACLs installed by the VLAN. The example below shows the output of the command followed by the default ACLs installed when fip-snooping is enabled on the VLAN.

Syntax Description

fip	FCoE Initialization Protocol.
snooping	Snooping FIP frames.
<i>vlan_name</i>	Name of the VLAN for which the access-list is shown.
fcf	List FCoE access-lists matching the FCoE forwarder's MAC.
<i>mac_addr</i>	MAC address of the FCoE forwarder.
virtual-link	List FCoE access-lists matching FCoE virtual link's MAC.
<i>mac_addr</i>	MAC address assigned to a VN-Port in the form xx:xx:xx:xx:xx:xx where xx is a pair of hexadecimal digits.
all	All FCoE access-lists in the VLAN.

Default

N/A.

Usage Guidelines

The command lists all the FCoE ACLs meeting the criteria.

Example

```
BDXA.112 # show fip snooping vlan v3

VLAN          : v3
FIP Snooping : Enabled
```




```
FCF Update    : Auto
FC-MAP       : 0e:fc:00:00:00:00
```

```
Port  Location
-----
1:1   Perimeter
1:2   FCF-to-Enode
1:3   Enode-to-FCF
1:4   All
-----
```

```
?
BDXA.113 # show fip snooping vlan v3 access-list
```

```
VLAN : v3
```

```
entry f424c0TffffS0efc00000000 { if match all {
    ethernet-type 0x0;
    ethernet-destination-address 0e:fc:00:00:00:00;
} then {
    deny ;
    do-not-learn ;
}}
```

```
entry f424c1T8914D011018010002 { if match all {
    ethernet-type 0x8914;
    ethernet-destination-address 01:10:18:01:00:02;
} then {
    permit ;
    mirror-cpu ;
}}
```

```
entry f424c2T8914D011018010001 { if match all {
    ethernet-type 0x8914;
    ethernet-destination-address 01:10:18:01:00:01;
} then {
    permit ;
    mirror-cpu ;
}}
```

```
entry f424c3T8906 { if match all {
    ethernet-type 0x8906;
} then {
    deny ;
    do-not-learn ;
}}
```

```
entry f424c3T8914 { if match all {
    ethernet-type 0x8914;
} then {
    deny ;
    do-not-learn ;
}}
```

```
Total number of ACL : 5
```

```
BDXA.114 #
```



History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the following platforms:

- BlackDiamond X8
- BlackDiamond 8800 series BD8900-40G6X-c
- Summit X670
- Summit X770

show fip snooping counters

```
show fip snooping {vlan} vlan_name counters
```

Description

This command shows the number of FIP frames snooped per type.

Syntax Description

fip	FCoE Initialization Protocol.
snooping	Snooping FIP frames.
<i>vlan_name</i>	Name of the VLAN for which the counters are shown.
counters	Number and types of FIP frames snooped on the VLAN.

Default

N/A.

Usage Guidelines

The command shows the number of FIP frames snooped per type.

Example

```
BDX8.62 # show fip snooping vlan v1 counters
```

```
VLAN : v1
```

FIP Frame type	Snooped
Solicited Discovery Request	0
Unsolicited Discovery Request	1
Solicited Discovery Advertisement	1



Unsolicited Discovery Advertisement	12
Fabric Login (FLOGI)	1
FLOGI Accept	1
FLOGI Reject	0
NPortID Virtualization Fabric Discovery (NPIV FDISC)	5
NPIV FDISC Accept	5
NPIV FDISC Reject	0
Fabric Logout (FLOGO)	0
FLOGO Accept	0
FLOGO Reject	0
Exchange Link Parameters (ELP)	0
ELP Accept	0
ELP Reject	0
ENode Keep-alive	11
VN_Port Keep-alive	6
Clear Virtual-link	0
VLAN Request	0
VLAN Notify	0
Unknown FIP Frame Type	0

BDX8.63 #

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the following platforms:

- BlackDiamond X8
- BlackDiamond 8800 series BD8900-40G6X-c
- Summit X670
- Summit X770

show fip snooping enode

```
show fip snooping {vlan} vlan_name enode
```

Description

This command shows the list of ENodes that are learned from FIP protocol packets on the specified VLAN.

The maximum FCoE size is in the snooped FIP discovery request from the ENode.



Syntax Description

fip	FCoE Initialization Protocol.
snooping	Snooping FIP frames.
<i>vlan_name</i>	Name of the VLAN for which the ENodes are shown.
enode	FCoE node. A Fiber Channel node that is able to transmit FCoE frames using one or more ENode MACs.

Default

N/A.

Usage Guidelines

This command shows the list of ENodes that are learned from FIP protocol packets on the specified VLAN.

Example

```
BDX8.92 # show fip snooping vlan v2 enode
VLAN : v2
Max
FCoE
ENode MAC          Port  Location          Age  Size
-----
00:00:00:A2:10:25  1:1  Perimeter         23   2098
00:00:01:C9:64:32  1:1  Perimeter         11   2098
00:00:05:A2:03:53  1:3  ENode to FCF      11   2098
00:00:00:9A:12:32  1:3  ENode to FCF      19   2098
Age      :The time in seconds since last FIP frame from the FCoE forwarder.
Total number of Enode MAC : 0
BDX8.93 #
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the following platforms:

- BlackDiamond X8
- BlackDiamond 8800 series BD8900-40G6X-c
- Summit X670
- Summit X770



show fip snooping fcf

```
show fip snooping {vlan} vlan_name fcf
```

Description

This command shows the list of FCFs in a VLAN on each member port. If the FCFs are added manually, the age is set to 0.

Syntax Description

fip	FCoE Initializaton Protocol
snooping	Snooping on FIP frames
<i>vlan_name</i>	Name of the VLAN for which the FCFs are shown.
fcf	FCoE forwarder. A Fiber Channel switching element that is able to forward FCoE frames.

Default

N/A.

Usage Guidelines

This command shows the list of FCFs in a VLAN on each member port. If the FCFs are added manually, the age is set to 0.

Example

```
BDX8.74 # show fip snooping vlan v2 fcf
VLAN          : v2
FCF Update    : Manual
FCF MAC       Port  Location      Age
-----
e2:ee:00:00:00:01  1:2  FCF-to-Enode    0
e2:ee:00:00:00:02  1:2  FCF-to-Enode    0
e2:ee:00:00:00:03  1:4  All              0
e2:ee:00:00:00:04  1:4  All              0
-----
Age          :The time in seconds since last FIP frame from the FCoE forwarder.
Total number of FCF MAC : 4
BDX8.75 #
```

History

This command was first available in ExtremeXOS 15.1.



Platform Availability

This command is available on the following platforms:

- BlackDiamond X8
- BlackDiamond 8800 series BD8900-40G6X-c
- Summit X670
- Summit X770

show fip snooping virtual-link

```
show fip snooping {vlan} vlan_name virtual-link {[enode mac_addr | fcf mac_addr]}
```

Description

This command lists the virtual links established in the VLAN. The list can be narrowed down to per ENode or per FCF where the ending point of the virtual link resides. The display shows all virtual links on the VLAN (as limited by the specification of enode or fcf) regardless of whether they are using SPMA or FPMA. Virtual links are differentiated within a VLAN by the VN_Port_ID (which is also contained in the low-order three octets of an FPMA MAC address, but not that of an SPMA MAC address).

Syntax Description

fip	FCoE Initialization Protocol.
snooping	Snooping FIP frames.
<i>vlan_name</i>	Name of the VLAN for which the FCFs are shown.
virtual-link	FCoE virtual link.
enode	Show virtual links related to the specified ENode only.
<i>mac_addr</i>	MAC of FCoE node originating the virtual link.
fcf	Show virtual links related to the specified FCF only.
<i>mac_addr</i>	MAC address of FCoE forwarder ending the virtual link in the form xx:xx:xx:xx:xx:xx where xx is a pair of hexadecimal digits.

Default

N/A.

Usage Guidelines

This command lists the virtual links established in the VLAN.

Example

```
BDX8.93 # show fip snooping v1 virtual-link
```



```

VLAN : v1
Port  ENode MAC          VN_Port MAC          VNPortId  FCF MAC          Age
-----
-----
1:1   aa:bb:cc:00:00:00    0e:fc:00:01:00:01   01:00:01   aa:bb:cd:00:00:00
2856
1:1   aa:bb:cc:00:00:00    0e:fc:00:01:00:02   01:00:02   aa:bb:cd:00:00:00
3106
1:1   aa:bb:cc:00:00:00    0e:fc:00:01:00:03   01:00:03   aa:bb:cd:00:00:00
3106
1:1   aa:bb:cc:00:00:00    0e:fc:00:01:00:04   01:00:04   aa:bb:cd:00:00:00
3106
1:1   aa:bb:cc:00:00:00    0e:fc:00:01:00:05   01:00:05   aa:bb:cd:00:00:00
3106
1:1   aa:bb:cc:00:00:00    0e:fc:00:01:00:06   01:00:06   aa:bb:cd:00:00:00
3106
VN_Port : Virtual N_Port instantiated on successful completion of
FIP FLOGI or FIP NPIV FDISC Exchange
Age      : The time in seconds since last FIP frame from the VN_Port
Total number of Virtual Link : 6
BDX8.94 #

```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the following platforms:

- BlackDiamond X8
- BlackDiamond 8800 series BD8900-40G6X-c
- Summit X670
- Summit X770

show fip snooping vlan

```
show fip snooping {vlan} vlan_name
```

Description

This command shows the FIP-snooping configuration status in the VLAN.

Syntax Description

<i>vlan_name</i>	Name of the VLAN for which the FIP Snooping configuration is shown.
------------------	---

Default

N/A.



Usage Guidelines

Use this command to show the FIP-snooping configuration status in the VLAN.

Example

```
BDX8.73 # show fip snooping vlan v2
VLAN          : v2
FIP Snooping  : Disabled
FCF Update    : Manual
FC-MAP        : 0e:fc:00:00:00:00
Port   Location
-----
1:1    Perimeter
1:2    FCF-to-Enode
1:3    Enode-to-FCF
1:4    All
-----
BDX8.74 #
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the following platforms:

BlackDiamond X8

BlackDiamond 8800 series BD8900-40G6X-c

Summit X670

Summit X770

show lldp dcbx

```
show lldp {port [all | port_list]} dcbx {ieee|baseline} {detailed}
```

Description

Displays DCBX configuration and statistics information for one or all ports.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.



ieee	Specifies IEEE 802.1Qaz information only.
baseline	Specifies Baseline v1.01 information only.
detailed	Shows information on the configured VLANs on the port.

Default

N/A.

Usage Guidelines

The summary display (without the detailed option) displays the status for each DCBX TLV on each port. For each TLV, the status is reported as shown in the following table.

DISABLED	DCBX is disabled on the port. This port status appears only in the summary display when DCBX is enabled for one version and disabled for the other. In the detailed display, ports on which DCBX are disabled are not shown.
OK	This TLV has been received by the peer, and either the configuration matches, or the peer is reporting that it is in willing mode and is not reporting an explicit error.
UNKNOWN	This TLV has not been received by the peer since the port has been active.
EXPIRED	This TLV has been received by the peer, but the time to live has expired.
ERROR	Either a mismatch exists between the local and remote configuration and the peer is not willing, or the peer is reporting an error.
MULTIPLE PEERS	More than one LLDP peer has been detected on the link.

When you specify a port or the detailed option, local TLV information includes the information that will be contained in the next TLV that is sent, and if the configuration hasn't changed, this is the same information that was sent in the last TLV. Peer TLV information displays the information from the last TLV that has been received. For each TLV, statistics are reported as follows:

- Sent: Total number of TLVs sent since port has been operational.
- Received: Total number of TLVs received since port has been operational.
- Errors: Total number of mal-formed TLVs received since port has been operational.

You can clear the statistics using the clear counters command.

Table 25: IEEE 802.1Qaz DCBX TLVs on page 1236 describes the IEEE 802.1Qaz DCBX TLVs that can be displayed. Table 26: Baseline v1.01 DCBX TLVs on page 1238 describes the Baseline v1.01 DCBX TLVs.



Table 25: IEEE 802.1Qaz DCBX TLVs



TLV/Description	Contents/Description
ETS TLV Advertises the ETS configuration of the local port and the configuration recommended to/by the peer for the specified port, respectively.	<p>Willing—Whether or not the device is willing to accept configuration from its DCBX peer. Zero (0) means No, and one (1) means Yes. The Willing bit does not apply to the ETS Recommendation TLV, and should always be zero.</p> <p>CBS—Whether the device supports the credit-based shaper algorithm. Zero (0) means No, and one (1) means Yes.</p> <p>Max TCs— Maximum number of traffic classes that the node can support.</p> <p>Priority Assgn—Priority Assignment Table. A priority group (PG) table describing how 802.1p priorities are assigned to PGs. The table is laid out as follows: Priority-0 : Priority-1 : Priority-2 : Priority-3 : Priority-4 : Priority-5 : Priority-6 : Priority-7</p> <p>The value in the Priority-N position indicates the TC ID to which packets with an 802.1p priority of N are mapped.</p> <hr/> <p>Note</p> <p> For Extreme Networks products, a traffic class (TC) is synonymous with a QoS Profile (QP), except that TCs are zero-based, and QPs are one-based, so TC 1 maps to QP 0.</p> <hr/> <p>TC Bwdth—TC Bandwidth Table. Indicates the percentage of bandwidth allocated for each traffic class. The table is laid out as follows: TC%-0 : TC%-1 : TC%-2 : TC%-3 : TC%-4 : TC%-5 : TC%-6 : TC%-7</p> <p>The value in the TC%-N position indicates the percentage of the link bandwidth allocated to TC N. The total of all positions must equal 100.</p> <p>TSA—Transmission Selection Algorithm (TSA) Assignment Table. The table is laid out as follows: TC-0 : TC-1 : TC-2 : TC-3 : TC-4 : TC-5 : TC-6 : TC-7</p> <p>The value in the TC-N position indicates the TSA used by TC N, which is one of the following: S - Strict priority (TSA 0)C - Credit-based shaper (TSA 1)E - Enhanced Transmission Selection (TSA 2)V - Vendor-specific Transmission Selection algorithm (TSA 255)</p> <hr/> <p>Note</p> <p> TSA values 3 to 254 are reserved for future standardization.</p>
Common Feature TLVs TLVs common to the Priority Group, PFC, and Application TLVs	<p>Oper Vers—Operating version of the feature.</p> <p>Max Vers—Highest feature version supported by the system.</p> <p>Enabled—Locally administered parameter that indicates whether the DCB feature is enabled. Zero (0) means No, and one (1) means Yes.</p> <p>Willing—Indicates whether the device is willing to accept configuration from its DCBX peer. Zero (0) means No, and one (1) means Yes.</p> <p>Error—Indicates whether an error has occurred during the configuration exchange with the peer. Zero (0) means No, and one (1) means Yes.</p>



Table 25: IEEE 802.1Qaz DCBX TLVs (continued)



TLV/Description	Contents/Description
<p>Priority Group TLV Advertises priority to priority group mapping, priority group bandwidth and the scheduling algorithm.</p>	<p>PG IDs—Priority Allocation Table. A priority group (PG) table describing how 802.1p priorities are assigned to PGs. The table is laid out as follows: Priority-0 : Priority-1 : Priority-2 : Priority-3 : Priority-4 : Priority-5 : Priority-6 : Priority-7</p> <p>The value in the Priority-N position indicates the PG ID to which packets with an 802.1p priority of N are mapped. If the value is in the range of 0 to 7, this is the actual PG. If the value is equal to 15, this priority is mapped to a non-ETS group. In the case of Extreme Networks products, this would be a strict priority group.</p> <hr/> <p> Note For Extreme Networks products, a priority group (PG) is synonymous with a QoS Profile (QP), except that PGs are zero-based, and QPs are one-based, so PG1 maps to QP 0.</p> <hr/> <p>PG%—Priority Group Allocation Table. Indicates the percentage of bandwidth allocated for each priority group. The table is laid out as follows: PG%-0 : PG%-1 : PG%-2 : PG%-3 : PG%-4 : PG%-5 : PG%-6 : PG%-7</p> <p>The value in the PG%-N position indicates the percentage of the link bandwidth allocated to PG N. The total of all positions must equal 100.</p> <p>Num TCs—Maximum number of priority groups that the node can support.</p>
<p>PFC TLV Describes the PFC configuration for the given port.</p>	<p>Willing—Whether or not the device is willing to accept configuration from its DCBX peer. Zero (0) means No, and one (1) means Yes.</p> <p>MBC—MACsec Bypass Capability. If set to zero (0), the device is capable of bypassing MACsec processing when MACsec is disabled. If set to one (1), the sending station is not capable of bypassing MACsec processing when MACsec is disabled.</p> <p>PFC Cap—PFC Capability. The maximum number of classes on which the device may simultaneously support PFC.</p> <p>PFC Enable—List of priorities on which PFC is enabled.</p>
<p>Application TLV Displays the priority the device expects to be used for the specified application.</p>	<p>Priority—The priority to be used for the given protocol.</p> <p>Application—Specifies one of the following:</p> <ul style="list-style-type: none"> • FCoE • FIP • iSCSI • EtherType: <i>ethertype</i> • TCP/UDP Port: <i>port number</i> • TCP Port: <i>port number</i> • TCP Port: <i>port number</i>



Table 26: Baseline v1.01 DCBX TLVs

TLV/Description	Contents/Description
Control TLV Contains general information about the DCBX session.	Oper Vers—Operating version of the DCBX protocol. Max Vers—Highest DCBX protocol version supported by the system. Seq No—A value that changes each time an exchanged parameter in one or more of the DCB feature TLVs changes. Ack No—The SeqNo value from the most recent peer DCBX TLV that has been handled. This value acknowledges to the peer that a specific SeqNo has been received.
Common Feature TLVs TLVs common to the Priority Group, PFC, and Application TLVs	Oper Vers—Operating version of the feature. Max Vers—Highest feature version supported by the system. Enabled—Locally administered parameter that indicates whether the DCB feature is enabled. Zero (0) means No, and one (1) means Yes. Willing—Indicates whether the device is willing to accept configuration from its DCBX peer. Zero (0) means No, and one (1) means Yes. Error—Indicates whether an error has occurred during the configuration exchange with the peer. Zero (0) means No, and one (1) means Yes.
Priority Group TLV Advertises priority to priority group mapping, priority group bandwidth and the scheduling algorithm.	PG IDs—Priority Allocation Table. A priority group (PG) table describing how 802.1p priorities are assigned to PGs. The table is laid out as follows: Priority-0 : Priority-1 : Priority-2 : Priority-3 : Priority-4 : Priority-5 : Priority-6 : Priority-7 The value in the Priority-N position indicates the PG ID to which packets with an 802.1p priority of N are mapped. If the value is in the range of 0 to 7, this is the actual PG. If the value is equal to 15, this priority is mapped to a non-ETS group. In the case of Extreme Networks products, this would be a strict priority group.  Note For Extreme Networks products, a priority group (PG) is synonymous with a QoS Profile (QP), except that PGs are zero-based, and QPs are one-based, so PG1 maps to QP 0. PG%—Priority Group Allocation Table. Indicates the percentage of bandwidth allocated for each priority group. The table is laid out as follows: PG%-0 : PG%-1 : PG%-2 : PG%-3 : PG%-4 : PG%-5 : PG%-6 : PG%-7 The value in the PG%-N position indicates the percentage of the link bandwidth allocated to PG N. The total of all slots must equal 100. Num TCs—Maximum number of priority groups that the node can support.
PFC TLV Describes the PFC configuration for the given port.	PFC Enable—List of priorities on which PFC is enabled. Num TC PFCs—The maximum number of classes on which the device may simultaneously support PFC.
Application TLV Displays the priority the device expects to be used for the specified application.	Priority—The priority to be used for the given protocol. Application—Specifies one of the following: <ul style="list-style-type: none"> • FCoE • FIP • iSCSI • EtherType: <i>ethertype</i> • TCP/UDP Port: <i>port number</i>



Example

The following example displays the summary DCBX configuration and statistics:

```
# show lldp dcbx
=====
==
Baseline DCBX TLV Status:          IEEE DCBX TLV Status:
Port      Control  PG      PFC      App      ETS-Conf ETS-Rec  PFC      App
=====
==
1         OK        OK      OK       OK       OK       OK       OK       OK
2         OK        OK      OK       OK       OK       OK       OK       OK
3         OK        OK      OK       OK       OK       OK       OK       OK
4         OK        OK      OK       OK       OK       OK       OK       OK
5         UNKNOWN UNKNOWN UNKNOWN UNKNOWN UNKNOWN UNKNOWN UNKNOWN UNKNOWN
UNKNOWN
9         UNKNOWN UNKNOWN UNKNOWN UNKNOWN UNKNOWN UNKNOWN UNKNOWN UNKNOWN
UNKNOWN
10        UNKNOWN UNKNOWN UNKNOWN UNKNOWN UNKNOWN DISABLED DISABLED DISABLED
DISABLED
16        DISABLED DISABLED DISABLED DISABLED UNKNOWN UNKNOWN UNKNOWN
UNKNOWN
23        UNKNOWN UNKNOWN UNKNOWN UNKNOWN UNKNOWN UNKNOWN UNKNOWN UNKNOWN
UNKNOWN
24        UNKNOWN UNKNOWN UNKNOWN UNKNOWN UNKNOWN UNKNOWN UNKNOWN UNKNOWN
UNKNOWN
=====
==
Control  - Control TLV
PG       - Priority Group TLV
PFC     - Priority-Based Flow Control TLV
App     - Application Configuration TLV
ETS-Conf - ETS Configuration TLV
ETS-Rec - ETS Recommendation TLV
```

The following example displays detailed IEEE 802.1Qaz DCBX configuration and statistics information for port 1:

```
# show lldp ports 1 dcbx ieee
Port number : 1
IEEE 802.1Qaz DCBX Information:
-----
ETS Configuration TLV: Sent: 5996, Received: 5997, Errors: 0, Status: OK
Local TLV : Willing: 0, CBS: 1, Max TCs: 8
Priority Assgn: 0:0:0:0:0:0:7, TC Bwdth: 33:0:0:33:34:0:0:0, TSA:
E:S:S:E:E:S:S:S
Peer TLV : Willing: 0, CBS: 1, Max TCs: 8
Priority Assgn: 0:0:0:0:0:0:7, TC Bwdth: 33:0:0:33:34:0:0:0, TSA:
E:S:S:E:E:S:S:S
ETS Recommendation TLV: Sent: 5996, Received: 5997, Errors: 0, Status: OK
Local TLV : Willing: 0, CBS: 0, Max TCs: 8
Priority Assgn: 0:0:0:0:0:0:7, TC Bwdth: 33:0:0:33:34:0:0:0, TSA:
E:S:S:E:E:S:S:S
Peer TLV : Willing: 0, CBS: 0, Max TCs: 8
Priority Assgn: 0:0:0:0:0:0:7, TC Bwdth: 33:0:0:33:34:0:0:0, TSA:
```



```

E:S:S:E:E:S:S:S
PFC TLV: Sent: 5996, Received: 5997, Errors: 0, Status: OK
Local TLV : Willing: 0, MBC: 0, Max PFCs: 8, PFC Enable: 3,4
Peer TLV  : Willing: 0, MBC: 0, Max PFCs: 8, PFC Enable: 3,4
Application TLV: Sent: 5987, Received: 5988, Errors: 0, Status: OK
Local TLV : Priority: 4, iSCSI
Priority: 3, FCoE
Priority: 3, FIP
Peer TLV  : Priority: 4, iSCSI
Priority: 3, FCoE
Priority: 3, FIP

```

The following example displays detailed Baseline v1.01 DCBX configuration and statistics information for port 1:

```

# show lldp ports 1 dcbx baseline
Port number : 1
Baseline v1.01 DCBX Information:
-----
Control TLV: Sent: 5999, Received: 6000, Errors: 0, Status: OK
Local TLV : Oper Vers: 0, Max Vers: 0, Seq No: 17, Ack No: 17
Peer TLV  : Oper Vers: 0, Max Vers: 0, Seq No: 17, Ack No: 17
Priority Group TLV: Sent: 5999, Received: 6000, Errors: 0, Status: OK
Local TLV : Oper Vers: 0, Max Vers: 0, Enabled: 1, Willing: 0 Error: 0
PG IDs: 0:0:0:0:0:0:0:15, PG%: 33:0:0:33:34:0:0:0, Num TCs: 8
Peer TLV  : Oper Vers: 0, Max Vers: 0, Enabled: 1, Willing: 0 Error: 0
PG IDs: 0:0:0:0:0:0:0:15, PG%: 33:0:0:33:34:0:0:0, Num TCs: 8
PFC TLV: Sent: 5999, Received: 6000, Errors: 0, Status: OK
Local TLV : Oper Vers: 0, Max Vers: 0, Enabled: 1, Willing: 0 Error: 0
Max PFCs: 8, PFC Enable: 3,4
Peer TLV  : Oper Vers: 0, Max Vers: 0, Enabled: 1, Willing: 0 Error: 0
Max PFCs: 8, PFC Enable: 3,4
App TLV: Sent: 5990, Received: 5991, Errors: 0, Status: OK
Local TLV : Oper Vers: 0, Max Vers: 0, Enabled: 1, Willing: 0 Error: 0
Priority: 4, iSCSI
Priority: 3, FCoE
Priority: 3, FIP
Peer TLV  : Oper Vers: 0, Max Vers: 0, Enabled: 1, Willing: 0 Error: 0
Priority: 4, iSCSI
Priority: 3, FCoE
Priority: 3, FIP

```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on all platforms.

show vlan dynamic-vlan

```
show vlan dynamic-vlan
```



Description

Displays the configuration related to dynamically created VLANs.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays configuration related to dynamically created VLANs.

```
X460-48p.7 # sh vlan dynamic-vlan
Uplink Ports      : 12-15, 18-20
X460-48p.8 #
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

show vm-tracking

show vm-tracking

Description

Displays the XNV feature configuration and the authenticated VM information.

Syntax Description

This command has no arguments or variables.

Default

N/A.



Usage Guidelines

None.

Example

The following command displays the XNV configuration and the authenticated VMs:

```
* Switch.51 # sh vm-tracking

-----
VM Tracking Global Configuration
-----
VM Tracking                : Enabled
VM Tracking authentication order: nms vm-map local
VM Tracking nms reauth period  : 0 (Re-authentication disabled)
VM Tracking blackhole policy   : none
-----

Port                : 1:20
VM TRACKING         : ENABLED

-----
MAC                Flags
APC    IP Address  Type    Value
-----
-----
00:00:00:00:00:11 LBI    255.255.255.255  VM
VPP    lvpp1
IEP
EEP
00:00:00:00:00:12 ---
VM
VPP
IEP
EEP
00:00:00:00:00:13 V---  30.30.30.30     VM    VMware-VM#2
VPP    nvpp1
IEP    a1.pol
EEP    a2.pol
-----
-----

Flags :
(A)uthenticated : L - Local, N - NMS, V - VMMAP
(P)olicy Applied : B - All Ingress and Egress, E - All Egress, I - All Ingress
(C)ounter Installed: B - Both Ingress and Egress, E - Egress, I - Ingress

Type :
IEP - Ingress Error Policies
EEP - Egress Error Policies

Number of Network VMs Authenticated: 1
Number of Local VMs Authenticated  : 1
Number of VMs Authenticated        : 2
```



History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

show vm-tracking local-vm

```
show vm-tracking local-vm {mac-address mac}
```

Description

Displays one or all of the VM entries in the local VM database.

Syntax Description

<i>mac</i>	Specifies the MAC address of a VM database entry that you want to display.
------------	--

Default

N/A.

Usage Guidelines

If you do not enter a MAC address with this command, the command displays all entries in the local VM database.

Example

The following command displays the local database VMs:

```
* Switch.52 # show vm-tracking local-vm
```

MAC Address	IP Address	Type	Value
00:00:00:00:00:21		VM	
		VLAN Tag	100
		VR Name	VR-Default
		VPP	vpp1

```
-----
Number of Local VMs: 1
```

History

This command was first available in ExtremeXOS 12.5.



Platform Availability

This command is available on all platforms.

show vm-tracking network-vm

show vm-tracking network-vm

Description

Displays all of the VM entries in the network VM database.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the configuration for all entries in the network VM database:

```
* Switch.52 # show vm-tracking network-vm
```

MAC Address	IP Address	Type	Value
00:00:00:00:00:11	192.168.100.200	VM VPP	KVM-VM-#101 vpp300
00:01:02:03:04:06	192.168.100.201	VM VPP	VM #200 vpp201

```
Number of Network VMs: 2
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.



show vm-tracking nms

```
show vm-tracking nms server {primary | secondary}
```

Description

Displays the RADIUS client configuration and operating statistics for one or both NMS servers.

Syntax Description

primary secondary	Specifies whether you are displaying the primary or secondary NMS server information.
-----------------------------------	---

Default

If you do not specify primary or secondary, the default action is to display both the primary and secondary NMS server configurations.

Usage Guidelines

None.

Example

The following command displays the RADIUS client information for the primary and secondary NMS servers:

```
show vm-tracking nms server
VM Tracking NMS (RADIUS): enabled
VM Tracking Radius server connect time out: 3 seconds
Primary VM Tracking NMS server:
Server name      :
IP address       : 10.127.5.221
Server IP Port   : 1812
Client address   : 10.127.10.173 (VR-Mgmt)
Shared secret    : pmckmtpq
Access Requests  : 0                Access Accepts    : 0
Access Rejects   : 0                Access Challenges  : 0
Access Retransmits: 0                Client timeouts    : 0
Bad authenticators: 0                Unknown types      : 0
Round Trip Time  : 0
Secondary VM Tracking NMS server:
Server name      :
IP address       : 10.127.5.223
Server IP Port   : 1812
Client address   : 10.127.10.173 (VR-Mgmt)
Shared secret    : rjgueogu
Access Requests  : 0                Access Accepts    : 0
Access Rejects   : 0                Access Challenges  : 0
Access Retransmits: 0                Client timeouts    : 0
```



```
Bad authenticators: 0           Unknown types      : 0
Round Trip Time   : 0
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

show vm-tracking port

```
show vm-tracking port port_list
```

Description

Displays the XNV feature configuration for the specified port and information for all VMs authenticated on the port.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
------------------	---

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the XNV configuration for port 1:20 and the authenticated VMs:

```
* (Private) Slot-1 Access3.14 # sh vm-tracking port 1:20
-----
      VM Tracking Global Configuration
-----
VM Tracking                : Enabled
VM Tracking authentication order : nms vm-map local
VM Tracking nms reauth period  : 0 (Re-authentication disabled)
VM Tracking blackhole policy   : none
-----
Port                : 1:20
VM Tracking         : Enabled
```



```
VM Tracking Dynamic VLAN      : Enabled
```

MAC	Flags AP	IP Address	Type	Value
00:00:00:00:00:11	LBI	255.255.255.255	VM	
			VLAN Tag	100
			VR Name	VR-Default
			VPP	lvpp1
			IEP	
			EEP	

```
-----
Flags :
  (A)uthenticated : L - Local, N - NMS, V - VM MAP
  (P)olicy Applied : B - All Ingress and Egress, E - All Egress, I - All
Ingress
  (C)ounter Installed : B - Both Ingress and Egress, E - Egress, I -
Ingress
  All Ingress and Egress, E - All Egress, I - All Ingress
Type :
  IEP - Ingress Error Policies      EEP - Egress Error Policies
```

```
Number of Network VMs Authenticated: 0
Number of Local VMs Authenticated  : 1
Number of VMs Authenticated        : 1
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

show vm-tracking repository

```
show vm-tracking repository {primary | secondary}
```

Description

Displays the FTP file synchronization configuration for NVPP and VM MAP files.

Syntax Description

primary secondary	Specifies whether you are displaying the primary or secondary FTP server configuration.
-----------------------------------	---

Default

If you do not specify primary or secondary, the default action is to display both the primary and secondary FTP server configurations.



Usage Guidelines

None.

Example

The following command displays the configuration for the primary and secondary FTP servers:

```
show vm-tracking repository
Primary VM-Map FTP server:
Server name:
IP address      : 10.100.1.200
VR Name        : VR-Mgmt
Refresh-interval: 600 seconds
Path Name       : /pub (default)
User Name       : anonymous (default)
Secondary vm-map FTP server: Unconfigured
Last sync      : 16:35:15      Last sync server : Primary
Last sync status : Successful
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

show vm-tracking vpp

```
show vm-tracking vpp {vpp_name}
```

Description

Displays the configuration of one or all VPPs.

Syntax Description

<i>vpp_name</i>	Specifies the name of an existing local VPP.
-----------------	--

Default

All.

Usage Guidelines

You can only specify local VPPs with this command. If you do not enter a VPP name with this command, the command displays all local and network VPPs.



Example

The following command displays the configuration of all VPPs:

```
* (Private) Slot-1 Access3.14 # sh vm-tracking vpp
VPP Name                               Type                               Value
-----
nvpp1                                  origin                            network
                                       counters                          ingress-only
                                       VLAN Tag                          200
                                       VR Name                           VR-Default
                                       ingress                            ingLocal1.pol(1)
                                                                       ingLocal2.pol(2)
                                       egress                            egrLocal1.pol(1)
                                                                       egrLocal2.pol(2)

lvpp1                                  origin                            local
                                       counters                          egress-only
                                       VLAN Tag                          100
                                       VR Name                           VR-Default
                                       ingress                            ingl.pol(1)
                                       egress                            egr1.pol(1)
                                                                       egr2.pol(2)

Number of Local VPPs   : 1
Number of Network VPPs: 1
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

unconfigure vm-tracking local-vm

```
unconfigure vm-tracking local-vm mac-address mac [name | ip-address | vpp | vlan-tag]
```

Description

Unconfigures the parameters associated with a local VM database entry to be used for VM MAC local authentication.

Syntax Description

mac	Specifies the MAC address for the local VM database entry you want to unconfigure.
name	Removes the name configured for the VM database entry.



ip-address	Removes the IP address configured for the VM database entry.
vpp	Removes the VPP configured for the VM database entry.
vlan-tag	Removes the VLAN tag configured for the VM database entry.

Default

N/A.

Usage Guidelines

None.

Example

The following command removes the IP address configuration for the VM entry specified by the MAC address:

```
unconfigure vm-tracking local-vm mac-address 00:E0:2B:12:34:56 ip-address
```

History

This command was first available in ExtremeXOS 12.5.

The ingress-vpp and egress-vpp options were replaced with the vpp option in ExtremeXOS 12.6.

The VLAN-tag option was added in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

unconfigure vm-tracking repository

```
unconfigure vm-tracking repository {primary | secondary}
```

Description

Removes the configuration for FTP file synchronization for NVPP and VM MAP files.

Syntax Description

primary secondary	Specifies the whether you are unconfiguring the primary or secondary FTP server.
-----------------------------------	--



Default

If you do not specify primary or secondary, the default action is to remove both the primary and secondary FTP server configurations.

Usage Guidelines

None.

Example

The following command removes the configuration for the primary and secondary FTP servers:

```
unconfigure vm-tracking repository
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

unconfigure vm-tracking vpp vlan-tag

```
unconfigure vm-tracking vpp vpp_name vlan-tag
```

Description

Unconfigures the VLAN tag of VPP.

Syntax Description

<i>vpp_name</i>	Specifies a name of the VPP.
-----------------	------------------------------

Default

N/A.

Usage Guidelines

Use this command to unconfigure the VLAN tag of VPP.

Example

Example output not yet available and will be provided in a future release.



History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

unconfigure vm-tracking vpp

unconfigure vm-tracking vpp *vpp_name*

Description

Removes the association of a policy or ACL rule to an LVPP.

Syntax Description

<i>vpp_name</i>	Specifies the name of an existing LVPP.
-----------------	---

Default

N/A.

Usage Guidelines

None.

Example

The following command removes the configuration of LVPP vpp1:

```
unconfigure vm-tracking vpp vpp1
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

unconfigure vm-tracking nms

unconfigure vm-tracking nms {**server** [**primary** | **secondary**]}



Description

Removes the configuration for one or both NMS servers.

Syntax Description

primary secondary	Specifies the whether you are unconfiguring the primary or secondary NMS.
-----------------------------------	---

Default

N/A.

Usage Guidelines

If you do not specify primary or secondary, this command removes the configuration for both NMS servers.

Example

The following command removes the configuration for the secondary NMS server:

```
unconfigure vm-tracking nms server secondary
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.



18 AVB Commands

```
clear msrp counters
clear mvrp counters
clear network-clock gtp counters
configure mrp ports timers
configure msrp latency-max-frame-size
configure msrp ports sr-pvid
configure msrp ports traffic-class delta-bandwidth
configure msrp timers first-value-change-recovery
configure mvrp stpd
configure mvrp tag ports registration
configure mvrp tag ports transmit
configure mvrp vlan auto-creation
configure mvrp vlan registration
configure network-clock gtp default-set
configure network-clock gtp ports announce
configure network-clock gtp ports peer-delay
configure network-clock gtp ports sync
disable avb
disable avb ports
disable msrp
disable msrp ports
disable mvrp
disable mvrp ports
disable network-clock gtp
disable network-clock gtp ports
enable avb
enable avb ports
enable msrp
enable msrp ports
enable mvrp
enable mvrp ports
enable network-clock gtp
enable network-clock gtp ports
show avb
show mrp ports
show msrp
show msrp listeners
```

```

show msrp ports
show msrp ports bandwidth
show msrp ports counters
show msrp streams
show msrp talkers
show mvrp
show mvrp ports counters
show mvrp tag
show network-clock gtp
show network-clock gtp ports
unconfigure avb
unconfigure mrp ports timers
unconfigure msrp
unconfigure mvrp
unconfigure mvrp stpd
unconfigure mvrp tag
unconfigure network-clock gtp ports

```

This chapter describes commands for managing the Audio Video Bridging (AVB) feature and its associated protocols including Multiple Registration Protocol (MRP), Multiple VLAN Registration Protocol (MVRP), Multiple Stream Registration Protocol (MSRP) and Generalized Precision Time Protocol (gPTP).

clear msrp counters

```
clear msrp counters {ports [port_list | all]}
```

Description

Clears both the PDU and attribute event counters per port.

Syntax Description

msrp	Multiple Stream Registration Protocol.
counters	MSRP packet and attribute event counters.
<i>port_list</i>	Port list separated by a comma or "-".
all	All ports.

Default

N/A.



Usage Guidelines

Use this command to clear both the PDU and attribute event counters per port.

Example

```
clear msrp counters
clear msrp counters ports 1-5
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

clear mvrp counters

```
clear mvrp counters {event | packet} {ports [port_list | all]}
```

Description

Clears MVRP statistics.

Syntax Description

mvrp	Multiple VLAN Registration Protocol.
event	MVRP event counters.
packet	MVRP packet counters.

Default

Clears both event and packet counters if none of the options are specified.

Usage Guidelines

Use this command to clear MVRP statistics. The default behavior clears both event and packet counters if none of the options are specified. The statistics that are reset are the number of failed registrations on that port, number of MVRPDUs sent, number of MVRPDUs received with error and without error for packet counters and different MVRP events rx/tx counters for event counters. If no port is specified, MVRP statistics of all ports are reset.



Example

The following command clears event counters:

```
clear mvrp event counters
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

clear network-clock gptp counters

```
clear network-clock gptp ports counters {ports [port_list | all]}
```

Description

Clears gPTP port counters.

Syntax Description

gptp	IEEE 802.1AS Generalized Precision Time Protocol.
counters	gPTP port counters.
<i>port_list</i>	Specifies one or more of the switch's physical ports.
all	Specifies all of the switch's physical ports.

Default

N/A.

Usage Guidelines

Use this command to clear gPTP port counters. The command `clear counters` also clears the gPTP port counters (along with all other counters).

Example

```
clear network-clock gptp counters
clear network-clock gptp counters ports 2-4
clear network-clock gptp counters ports all
```



History

This command was first available in ExtremeXOS 15.3

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

configure mrp ports timers

```
configure mrp ports [port_list | all] timers [ {join join_msec} {leave
leave_msec} {leave-all leave_all_msec} {periodic [periodic_msec | off]}{refresh
[ auto-refresh | refresh_msec | off ]} ]
```

Description

This command sets the join, leave, leave all, periodic, and extended-refresh timer values for a list of ports. The unit value is in milliseconds. The join timer, leave all timer, and periodic timer are started for each MRP application per port. The leave timer is started for each state machine that is in LV (leave) state. The default values for join, leave, leave-all, periodic and extended-refresh timers are 200, 600, 10000, 1000, and 0 milliseconds, respectively.

Syntax Description

mrp	Multiple Registration Protocol.
ports	Ports.
<i>port_list</i>	Port list separated by a comma or "-" type="portlist_t".
all	All ports.
timers	Multiple Registration Protocol timers.
join	The time interval to delay sending MRP advertisements.
<i>join_msec</i>	Join timer value in milliseconds (range is 0 ms to 500 ms, default is 200 ms).
leave	The time interval to wait in the leaving state before transitioning to the empty state.
<i>leave_msec</i>	Leave timer value in milliseconds (range is 600 ms to 3000 ms, default is 600 ms).
leave-all	The time interval used to control the frequency of "leave all" messages.
<i>leave_all_msec</i>	Leave All timer value in milliseconds (range is 5000 ms to 20000 ms, default is 10000 ms).
periodic	The time interval between two periodic events.
<i>periodic_msec</i>	Periodic timer value in milliseconds (range is 1000ms to 300000 ms, default is 1000 ms); type="uint32_t".
off	Turn off timer.



refresh	Timer value to use in place of regular timer, only in cases when leave-all is received or sent.
auto-refresh	Automatically calculate timer values based on number of talkers and listeners.
<i>refresh_msec</i>	Refresh timer value in milliseconds (range is 600ms to 300000ms, default is 0ms (off)).

Default

The default values for join, leave, leave-all, periodic and extended-refresh timers are 200, 600, 10000, 1000, and 0 milliseconds, respectively.

Usage Guidelines

This command is used to set the join, leave, and leave-all timer values for a list of ports. The unit value is in milliseconds. The join timer and leave all timer are started for each MRP application per port. The leave timer is started for each state machine that is in LV (leave) state. The default values for these timers are 200, 600, and 10000, respectively.

```
configure mrp ports 4 timers join 300
configure mrp ports all timers leave-all 15000
configure mrp ports all timers join 300 leave-all 15000
```

History

The extended-refresh and period timer options were added in 15.3.2.

Platform Availability

This command is available on all platforms.

configure msrp latency-max-frame-size

```
configure msrp [ latency-max-frame-size frame_size | [ ignore-latency-changes | talker-vlan-pruning ] [ on | off ] ]
```

Description

This command configures the system-wide MSRP variables.

Syntax Description

msrp	Multiple Stream Registration Protocol.
latency-max-frame-size	Maximum size of interfering frame (used in latency calculations).
<i>frame_size</i>	The maximum frame size in bytes (range 64 to 2000, default is 1522).
ignore-latency-changes	Ignore accumulated latency changes when evaluating first value change.



talker-vlan-pruning	Talker propagation is filtered on ports where VLAN does not exist.
on	Turn on.
off	Turn off.

Default

1522.

Usage Guidelines

Use this command to configure the system-wide MSRP variables.

Example

```
configure msrp latency-max-frame-size 100
```

History

The ignore-latency-changes, talker-vlan-pruning, and on | off options were added in 15.3.2.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

configure msrp ports sr-pvid

```
configure msrp ports [port_list | all] sr-pvid vlan_tag
```

Description

Specifies the default VLAN ID on the port for MSRP data stream. The sr-pvid serves as a recommendation to connected AVB devices; AVB devices may still use other VLAN IDs if they are configured to do so.

Syntax Description

msrp	Multiple Stream Registration Protocol
<i>port_list</i>	List of ports in the switch.
all	All the ports in the switch.
sr-pvid	Default VLAN Identifier for stream-related traffic.
<i>vlan_tag</i>	VLAN ID ranging from 1 to 4094 (default is 2).



Default

2.

Usage Guidelines

Use this command to specify the default VLAN ID on the port for MSRP data streams. The `sr-pvid` serves as a recommendation to connected AVB devices; AVB devices may still use other VLAN IDs if they are configured to do so.

Example

```
configure msrp ports 1,2,3 sr-pvid 2
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

configure msrp ports traffic-class delta-bandwidth

```
configure msrp ports [port_list | all] traffic-class [A | B] delta-bandwidth  
percentage
```

Description

Configures delta-bandwidth value per traffic class per MSRP port.

Syntax Description

msrp	Multiple Stream Registration Protocol.
<i>port_list</i>	List of ports in the switch.
traffic-class	Traffic class.
A	Traffic class A.
B	Traffic class B.
delta-bandwidth	Delta-bandwidth percentage (range 0 to 100, default 75 for class A, 0 for class B).

Default

Class A: 75, Class B: 0.



Usage Guidelines

The delta bandwidth configuration limits the amount of bandwidth that can be used by the given stream reservation class. Each class is allowed to use a maximum of its delta bandwidth plus the delta bandwidth configured for each of the higher classes. For example, if the delta bandwidth for classes A and B are configured to 10 and 10 respectively, class A streams can use up to 10 percent of the link bandwidth, and class B streams can use up to 20 percent of the link bandwidth. The sum of the class A and B delta bandwidth values must be less than 100 percent.

Example

```
configure msrp ports all traffic-class A delta-bandwidth 50
configure msrp ports 1-5 traffic-class B delta-bandwidth 0
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

configure msrp timers first-value-change-recovery

```
configure msrp timers first-value-change-recovery [first_value_change_msec | off]
```

Description

This command configures MSRP first value change recovery timer, or disables the timer. If configured, the system waits until the configured timer value before allowing recovery of streams from first value change failure. If disabled, the system does not recover from first value change failure.

Syntax Description

msrp	Multiple Stream Registration Protocol.
timers	Multiple Stream Registration Protocol timers.
first-value-change-recovery	The time interval to wait to allow recovery of stream from first value change failure.
first_value_change_msec	First Value Change Recovery time in milliseconds (range is 10000 ms to 5400000 ms, default is 30000 ms); type="uint32_t"; range="[10000, 5400000]".
off	Turn off first value change recovery timer, and do not recover from first value change failure.



Default

30000 ms.

Usage Guidelines

Use this command to allow streams to recover from first value change failure.

Example

```
configure msrp timers first-value-change recovery 20000
configure msrp timers first-value-change recovery off
```

History

This command was first available in ExtremeXOS 15.3.2.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

configure mvrp stpd

```
configure mvrp stpd stpd_name
```

Description

Configures the STP domain to use for dynamically created VLANs.

Syntax Description

mvrp	Multiple VLAN Registration Protocol.
stpd	The STP domain used for MVRP.
<i>stpd_name</i>	The STP domain the VLAN is to be associated. All ports of the domain will be advertised, when this VLAN gets registered.

Default

s0.

Usage Guidelines

Use this command to configure the STP domain used for MVRP.



Example

The following example configures the default STP domain for MVRP to "stpd2":

```
configure mvrp stpd stpd2
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

configure mvrp tag ports registration

```
configure mvrp tag vlan_tag ports [port_list | all] registration [forbidden | normal ]
```

Description

This command is used to configure MVRP VLAN registration as forbidden or normal for specific VLANs on specific ports. If registration for a VLAN is configured forbidden on a specific port, then when MVRP PDU is received on the port with the particular forbidden VLAN Id, the VLAN is not created and if the VLAN is already there, the port is not added to the VLAN. For normal MVRP behavior, the registration needs to be normal. The forbidden / normal setting is only for dynamic addition of ports to VLANs. Any static addition of ports to the VLANs will override this setting and will mark the status as fixed. The forbidden setting can be used to control MSRP advertisements, in typical scaling scenarios.

Syntax Description

mvrp	Multiple VLAN Registration Protocol.
tag	The 802.1Q VLAN ID.
<i>vlan_tag</i>	VLAN ID ranging from 1 to 4094; type=uint16_t"; range="[1,4094]".
ports	Ports.
<i>port_list</i>	Port list separated by a comma or "-"; type="portlist_t";
all	All ports.
registration	Whether port can be added dynamically to the VLAN.
forbidden	Port cannot be added dynamically to the VLAN.
normal	Port can be added dynamically to the VLAN.

Default

Normal.



Usage Guidelines

Use this command to control dynamic addition of ports to VLANs.

Example

```
configure mvrp tag 2 ports 2,3,4 registration forbidden
configure mvrp tag 2 ports all registration normal
```

History

This command was first available in ExtremeXOS 15.3.

The registration option, and forbidden and normal keywords were added in 15.3.2.

Platform Availability

This command is available on all commands.

configure mvrp tag ports transmit

```
configure mvrp tag vlan_tag ports [port_list | all] transmit [on | off ]
```

Description

Controls whether the given VLAN ID may be advertised in MVRP messages transmitted on the given set of ports.

Syntax Description

mvrp	Multiple VLAN Registration Protocol.
tag	The 802.1Q VLAN ID.
transmit	When enabled, MVRP message are sent on the ports.
on	Transmission of MVRP messages are enabled on the port(s) for the given tag.
off	Transmission of the MVRP messages are disabled on the port(s) for the given tag.

Default

Transmit on.

Usage Guidelines

Use this command to control whether the given VLAN ID may be advertised in MVRP messages transmitted on the given set of ports.



Example

The following command configures transmit off for VLAN ID 100 on all MVRP ports:

```
configure mvrp tag 100 ports all transmit off
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

configure mvrp vlan auto-creation

```
configure mvrp vlan auto-creation [on | off]
```

Description

Enables or disables the dynamic VLAN creation feature of MVRP.

Syntax Description

mvrp	Multiple VLAN Registration Protocol.
auto-creation	When enabled, results in VLANs added dynamically on the switch through MVRP.
on	Enable auto-creation.
off	Disable auto-creation.

Default

Enabled.

Usage Guidelines

Use this command to enable or disable the dynamic VLAN creation of MVRP. By default, auto-creation is enabled. If disabled, the switch may participate in the MVRP protocol, and advertised static VLANs, but will not dynamically create VLANs.

Example

The following command enables MVRP VLAN auto creation:

```
configure mvrp vlan auto-creation on
```



History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

configure mvrp vlan registration

configure mvrp vlan registration forbidden | normal

Description

This command is a global system setting. If global registration is forbidden, ports cannot be added to any VLAN dynamically.

Syntax Description

mvrp	Multiple VLAN Registration Protocol.
vlan	VLAN.
registration	Whether all ports can be added to new dynamic VLANs. This can be overridden by static port addition to VLAN.
forbidden	Ports cannot be added dynamically to the VLAN. This can be overridden by static port addition.
normal	Ports can be added dynamically to the VLAN (default).

Default

Normal.

Usage Guidelines

Use this command to set global registration. If global registration is forbidden, ports cannot be added to any VLAN dynamically.

Example

The following command allows ports to be added dynamically to the VLAN:

```
configure mvrp vlan registration normal
```

History

This command was first available in ExtremeXOS 15.3.

The **registration** keyword was first available in ExtremeXOS 15.3.2.



Platform Availability

This command is available on all platforms.

configure network-clock gtp default-set

```
configure network-clock gtp default-set [{priority1 priority1_value} {priority2
priority2_value}]
```

Description

This command configures the switch's default-set parameters, specifically its grandmaster clock priority values that are used to elect the grandmaster clock in the network.

Syntax Description

<i>priority1_value</i>	The switch's grandmaster clock priority1 value. This is the most significant parameter used to select the grandmaster clock in the network. Lower values indicate higher priority, and 255 prevents the switch from becoming the grandmaster clock.
<i>priority2_value</i>	The switch's grandmaster clock priority2 value. This is one of the least significant parameters used to select the grandmaster clock in the network. Lower values indicate higher priority.

Default

- Priority1_value = 246 (from 802.1AS 8.6.2.1)
- Priority2_value = 248 (from 802.1AS 8.6.2.5)

Usage Guidelines

Use this command to configure the switch's default-set parameters, specifically its grandmaster clock priority values that are used to elect the grandmaster clock in the network. The Best Master Clock Algorithm uses six parameters from each time-aware system in the network to select the grandmaster clock in the network. Priority1 is the highest precedence value; it allows users to preemptively configure which systems they prefer to be the grandmaster clock. Priority2 is a lower precedence value; it allows users to configure tiebreaker priorities.

The default priority1 values defined by IEEE 802.1AS-2011 clause 8.6.2.1 give preference to network infrastructure systems such as Extreme switches.

Example

```
configure network-clock gtp default-set priority1 248
configure network-clock gtp default-set priority2 100
configure network-clock gtp default-set priority1 248 priority2 100
```



History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

configure network-clock gtp ports announce

```
configure network-clock gtp ports [port_list | all] announce [initial-interval
log_2_interval | receipt-timeout timeout_count]
```

Description

Configures gPTP Announce parameters on the specified ports. Announce messages are used to elect the grandmaster clock and determine the time-synchronous spanning tree.

Syntax Description

<i>port_list</i>	Specifies one or more of the switch's physical ports.
all	Specifies all of the switch's physical ports.
<i>log_2_interval</i>	The interval between Announce messages used by the switch on the port when the port is initialized or when the switch receives a message interval request TLV with announceInterval value 126. This value is in \log_2 seconds. The valid range of values is -3 ($2^{-3} = 0.125$ seconds) to 17 ($2^{17} = 131072$ seconds).
<i>timeout_count</i>	On a gPTP slave port, the number of announce intervals to wait without receiving an Announce message before assuming the master is no longer sending Announce messages.

Default

- *log_2_interval* = 0 (1 second; 802.1AS-2011 10.6.2.2)
- *timeout_count* = 3 (802.1AS-2011 10.6.3.2)

Usage Guidelines

Use this command to configure gPTP Announce parameters on the specified ports. Announce messages are used to elect the grandmaster clock and determine the time-synchronous spanning tree. Announce selects the grandmaster in the network and establishes the tree from the grandmaster to all other time-aware systems in the network.

initial-interval corresponds to 802.1AS parameter `initialLogAnnounceInterval`.

receipt-timeout corresponds to 802.1AS parameter `announceReceiptTimeout`.



Example

```
configure network-clock gtp ports 1-2 announce initial-interval 127
configure network-clock gtp ports all announce receipt-timeout 5
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

configure network-clock gtp ports peer-delay

```
configure network-clock gtp ports [port_list | all] peer-delay [{allowed-lost-responses lost_responses_value} {initial-req-interval log_2_interval} {asymmetry_time [nanoseconds | microseconds | milliseconds | seconds] | neighbor-thresh [auto | neighbor_thresh_time [nanoseconds | microseconds | milliseconds | seconds]]}]
```


Description

Configures gPTP peer delay parameters on the specified ports.

Syntax Description

<i>port_list</i>	Specifies one or more of the switch's physical ports.
all	Specifies all of the switch's physical ports.
<i>lost_responses_value</i>	The number of consecutive Peer Delay RequestPdelay_Req messages that the switch must send on a port without receiving a valid response before it considers the port not to be exchanging Peer Delay messages with its neighbor.
<i>log_2_interval</i>	The interval between Peer Delay RequestPdelay_Req messages sent by the switch on the port when the port is initialized or when the switch receives on the port a message interval request TLV with linkDelayInterval value of 126. This value is in log2 seconds. The valid range of values is -3 (2 ⁻³ = 0.125 seconds) to 17 (2 ¹⁷ = 131072 seconds).



<i>asymmetry_time</i>	<p>The time that the propagation delay from this switch to the neighbor is less than the estimated one-way propagation delay between the switch and its neighbor (which is also the time that the propagation delay from the neighbor to this switch is greater than the estimate). This value is negative if the propagation delay to the neighbor is greater than the estimate. It can be in nanoseconds, microseconds, milliseconds, or seconds. The maximum value is 4,294,967,295 nanoseconds (approximately 4.3 seconds). Let tIR be the propagation delay from this switch (initiator) to the neighbor (responder), tRI be the propagation delay from the neighbor to this switch, and meanPathDelay be the estimated one-way propagation delay. Then:</p> <ul style="list-style-type: none"> • $\text{meanPathDelay} = (tIR + tRI) / 2$ • $tIR = \text{meanPathDelay} - \text{asymmetry_time}$ • $tRI = \text{meanPathDelay} + \text{asymmetry_time}$
<i>neighbor_thresh_time</i>	<p>The maximum measured mean of the propagation delay between this switch and the neighbor above which the switch considers the port unable to run gPTP. This value can be in nanoseconds, microseconds, milliseconds, or seconds.</p>
auto	<p>Use a media specific default value for the neighbor_thresh_time:</p> <ul style="list-style-type: none"> • Copper: 800 nanoseconds. This category includes short range copper cables such as SFP+ Direct Attach and QSRP+ Passive Copper. • Multi-mode fiber: 11 microseconds. This category includes the QSFP+ Active Optical cables. 11 microseconds allows 10 microseconds for 100BASE-FX 2 km plus 10% tolerance.) • Single-mode fiber: 550 microseconds. This allows 500 microseconds for our "LX100" transceiver plus 10% tolerance. <hr/> <p> Note These values may change. A draft of the 802.1AS corrigendum (P802.1AS-Cor-1/D1.1) specifies 800 ns for 100BASE-TX and 1000BASE-T.</p>

Default

- `Lost_responses_value` = 3 (802.1AS 11.5.3)
- `Log_2_interval` = 0 (1 second; not specified in 802.1AS)
- `Asymmetry_time` = 0 (802.1AS 10.2.4.8)
- `Neighbor_thresh_time` = Copper media: 800 nanoseconds, fiber media: 4,294,967,295 nanoseconds

Usage Guidelines

Peer Delay messages determine whether a neighboring system is gPTP capable and measure the propagation delay on the link between the switch and a neighboring gPTP capable system.

- **allowed-lost-responses** corresponds to 802.1AS parameter `allowedLostResponses`.
- **initial-req-interval** corresponds to 802.1AS parameter `initialLogPdelayReqInterval`.
- **asymmetry** corresponds to 802.1AS parameter `delayAsymmetry`.
- **neighbor-thresh** corresponds to 802.1AS parameter `neighborPropDelayThresh`.



Example

```
configure network-clock gtp ports 1-3 peer-delay allowed-lost-responses 5
configure network-clock gtp ports 1-2 peer-delay initial-log-interval -3
configure network-clock gtp ports 1-2 peer-delay neighbor-thresh 3
nanoseconds
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

configure network-clock gtp ports sync

```
configure network-clock gtp ports [port_list | all] sync [initial-interval
log_2_interval receipt-timeout timeout_count]
```

Description

Configures gPTP synchronization parameters on the specified ports.

Syntax Description

<i>port_list</i>	Specifies one or more of the switch's physical ports.
all	Specifies all of the switch's physical ports.
<i>log_2_interval</i>	The interval between Sync messages used by the switch for the port when the port is initialized or when the switch receives a message interval request TLV with timeSyncInterval value of 126. This value is in log2 seconds. The valid range of values is -3 (2 ⁻³ = 0.125 seconds) to 17 (2 ¹⁷ = 131072 seconds).
<i>timeout_count</i>	On a gPTP slave port, the number of sync intervals to wait without receiving a Sync message before assuming the adjacent master port is no longer sending Sync messages.

Default

- *log_2_interval* = -3 (0.125 second; 802.1AS 11.5.2.3)
- *timeout_count* = 3 (802.1AS 10.6.3.1)

Usage Guidelines

Synchronization distributes the time from the grandmaster to all other time-aware systems in the networks.



initial-interval corresponds to 802.1AS parameter `initialLogSyncInterval`.

receipt-timeout corresponds to 802.1AS parameter `syncReceiptTimeout`.

Example

```
configure network-clock gptp ports 1-2 sync initial-interval -1
configure network-clock gptp ports all sync receipt-timeout 5
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

disable avb

disable avb

Description

This command is a macro command that can be used to disable all AVB protocols globally on the switch. It is equivalent to issuing the following three commands:

```
disable mvrp
```

```
disable msrp
```

```
disable network-clock gptp
```

Syntax Description

avb	Audio Video Bridging
------------	----------------------

Default

Disabled.

Usage Guidelines

Use this command to disable all AVB protocols globally on the switch.



Example

```
disable avb
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

disable avb ports

```
disable avb ports [port_list | all]
```

Description

This command is a macro command that can be used to disable all AVB protocols on the given ports. It is equivalent to issuing the following three commands:

```
disable mvrp ports [port_list | all]
```

```
disable msrp ports [port_list | all]
```

```
disable network-clock gptp ports [port_list | all]
```

Syntax Description

avb	Audio Video Bridging.
<i>port_list</i>	Port list separated by a comma or "-".
all	All ports.

Default

Disabled.

Usage Guidelines

Use this command to disable all AVB protocols on the given ports.

Example

```
disable avb ports all
```



History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

disable msrp

disable msrp

Description

Disables MSRP on the switch.

Syntax Description

msrp	Multiple Stream Registration Protocol.
-------------	--

Default

Disabled.

Usage Guidelines

Use this command to disable MSRP on a switch.

Example

The following command disables MSRP:

```
disable msrp
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.



disable msrp ports

```
disable msrp ports [port_list | all]
```

Description

Disables MSRP on the ports listed in the command after the keyword **ports**.

Syntax Description

msrp	Multiple Stream Registration Protocol.
<i>port_list</i>	Port list separated by a comma or "-".
all	All ports.

Default

Disabled.

Usage Guidelines

Use this command to disable MSRP in the ports listed or all ports.

Example

```
disable msrp ports all
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

disable mvrp

```
disable mvrp
```

Description

Disables MVRP globally on a switch.



Syntax Description

mvrp	Multiple VLAN Registration Protocol.
-------------	--------------------------------------

Default

Disabled.

Usage Guidelines

Use this command to disable MVRP globally on a switch. MVRP is run on the MVRP enabled ports only if the global setting is enabled. By default, MVRP is disabled globally and on individual ports. When MVRP is disabled globally, all MVRP packets are forwarded transparently.

Example

The following command disables MVRP:

```
disable mvrp
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

disable mvrp ports

```
disable mvrp ports [port_list | all]
```

Description

Disable MVRP on a given set of ports.

Syntax Description

mvrp	Multiple VLAN Registration Protocol.
<i>port_list</i>	Port(s) on which MVRP is to be enabled.
all	All ports.

Default

Disabled.



Usage Guidelines

Use this command to disable MVRP on given set of ports. MVRP is run on the MVRP enabled ports only if the global setting is enabled. By default MVRP is disabled globally and on individual ports. When MVRP is disabled globally, all MVRP packets will be forwarded transparently.

Example

The following command disables MVRP on ports 4 and 5:

```
disable mvrp ports 4-5
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all commands.

disable network-clock gtp

```
disable network-clock gtp
```

Description

Disables gPTP on the switch.

Syntax Description

network-clock	Network clock.
gtp	IEEE 802.1AS Generalized Precision Time Protocol (gPTP).

Default

Disabled.

Usage Guidelines

Use this command to disable gPTP after having enabled it.

Example

```
disable network-clock gtp
```



History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

disable network-clock gptp ports

```
disable network-clock gptp ports [port_list | all]
```

Description

Disables gPTP on one or more ports.

Syntax Description

<i>port_list</i>	Specifies one or more the the switch's physical ports.
all	Specifies all of the switch's physical ports.

Default

Disabled.

Usage Guidelines

Use this command to configure on which ports gPTP runs. gPTP runs on no ports if it is not enabled in the switch by `enable network-clock gptp`.

Example

```
disable network-clock gptp ports 1-3
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.



enable avb

enable avb

Description

This command is a macro command that can be used to enable all AVB protocols globally on the switch. It is equivalent to issuing the following three commands:

```
enable mvrp
```

```
enable msrp
```

```
enable network-clock gtp
```

Syntax Description

avb	Audio Video Bridging.
------------	-----------------------

Default

Disabled.

Usage Guidelines

Use this command to globally enable all AVB protocols globally on the switch.

Example

```
enable avb
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

enable avb ports

enable avb ports [*port_list* | **all**]



Description

This command is a macro command that can be used to enable all AVB protocols on the switch. It is equivalent to issuing the following three commands:

```
enable mvrp
```

```
enable msrp
```

```
enable network-clock gptp
```

Syntax Description

avb	Audio Video Bridging.
<i>port_list</i>	Port list separated by a comma or "-".
all	All ports.

Default

Disabled.

Usage Guidelines

Use this command to enable all AVB protocols on the given ports.

Example

```
enable avb ports 1-5
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

enable msrp

```
enable msrp
```

Description

Enables MSRP globally on the switch.



Syntax Description

msrp	Multiple Stream Registration Protocol.
-------------	--

Default

Disabled.

Usage Guidelines

Use this command to enable MSRP globally on a switch.

Example

The following command enables MSRP:

```
enable msrp
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

enable msrp ports

```
enable msrp ports [port_list | all]
```

Description

Enables MSRP in the ports listed in the command after the keyword **ports**.

Syntax Description

msrp	Multiple Stream Registration Protocol.
<i>port_list</i>	Port list separated by a comma or "-".
all	All ports.

Default

Disabled.



Usage Guidelines

Use this command to enable MSRP in the ports listed or all ports.

Example

```
enable msrp ports 1-3
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

enable mvrp

```
enable mvrp
```

Description

Enables MVRP globally on a switch.

Syntax Description

mvrp	Multiple VLAN Registration Protocol.
-------------	--------------------------------------

Default

Disabled.

Usage Guidelines

Use this command to enable MVRP globally on a switch. MVRP is run on the MVRP enabled ports only if the global setting is enabled. By default, MVRP is disabled globally and on individual ports. When MVRP is disabled globally, all MVRP packets are forwarded transparently.

Example

The following command enables MVRP globally on the switch:

```
enable mvrp
```



History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

enable mvrp ports

```
enable mvrp ports [port_list | all]
```

Description

Enables MVRP on a given set of ports.

Syntax Description

mvrp	Multiple VLAN Registration Protocol
<i>port_list</i>	Port(s) on which MVRP is to be enabled.
all	All ports.

Default

Disabled.

Usage Guidelines

Use this command to enable MVRP on given set of ports. MVRP is run on the MVRP enabled ports only if the global setting is enabled. By default, MVRP is disabled globally and on individual ports. When MVRP is disabled globally, all MVRP packets will be forwarded transparently.

Example

The following command enables MVRP on ports 4 and 5:

```
enable mvrp ports 4-5
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.



enable network-clock gptp

enable network-clock gptp

Description

Enables gPTP on the switch.

Syntax Description

network-clock	Network clock.
gptp	IEEE 802.1AS Generalized Precision Time Protocol (gPTP).

Default

Disabled.

Usage Guidelines

Use this command to enable gPTP.

Example

```
enable network-clock gptp
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

enable network-clock gptp ports

enable network-clock gptp ports [*port_list* | **all**]

Description

Enables gPTP on one or more ports.



Syntax Description

<i>port_list</i>	Specifies one or more of the switch's physical ports.
all	Specifies all of the switch's physical ports.

Default

Disabled.

Usage Guidelines

Use this command to configure on which ports gPTP runs. gPTP does not run on any ports if it is not first enabled in the switch by the `enable network-clock gptp` command.

Example

```
enable network-clock gptp
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

show avb

show avb

Description

Displays a summary of MSRP, MVRP, and gPTP configuration on the switch.

Syntax Description

avb	Audio Video Bridging.
------------	-----------------------

Default

N/A.



Usage Guidelines

Use this command to display a summary of MSRP, MVRP, and gPTP configuration and status on the switch.

Example

```
#show avb
gPTP status           : Enabled
gPTP enabled ports   : *17d *19d

MSRP status          : Enabled
MSRP enabled ports   : !3 *17ab *19a

MVRP status          : Enabled
MVRP enabled ports   : *17 *19

Flags:                (*) Active, (!) Administratively disabled,
                      (a) SR Class A allowed, (b) SR Class B allowed,
                      (d) Disabled gPTP port role, (m) Master gPTP port role,
                      (p) Passive gPTP port role, (s) Slave gPTP port role.
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

show mrp ports

```
show mrp ports {port_list}
```

Description

Shows the MRP timers configured on the given list of ports on the switch.

Syntax Description

mrp	Multiple Registration Protocol.
<i>port_list</i>	Ports on which MRP timers are configured or unconfigured.

Default

N/A.



Usage Guidelines

Use this command to view MRP timers configured on the given list of ports on the switch.

Example

```
# show mrp ports 1, 4, 5
-----
-----
Ports      Join Time (ms)      Leave Time (ms)      Leave All
Time (ms)  Periodoc (ms)      Extended
-----
-----
Refresh (ms)
-----
-----
1          200                600
10000     1000              10000
4          300                800
10000     1000              10000
5          200                600
10000     1000              10000
-----
-----
```

History

This command was first available in ExtremeXOS 15.3.

Output for periodic and extended refresh timers added in 15.3.2.

Platform Availability

This command is available on all commands.

show msrp

show msrp

Description

Displays the MSRP configuration on the switch.

Syntax Description

msrp	Multiple Stream Registration Protocol.
-------------	--



Default

N/A.

Usage Guidelines

Use this command to display MSRP configuration on the switch.

Example

```
# show msrp
MSRP Status                : Enabled
MSRP Max Latency Frame Size : 1522
MSRP Max Fan-in Ports      : No limit
MSRP First Value Change Recovery Time : 10000 (ms)
MSRP Ignore Latency Changes : On
MSRP Talker VLAN Pruning   : On
MSRP Enabled Ports         : *17ab      *19a    !5
Total MSRP streams         : 4
Total MSRP reservations    : 2
Flags:      (*) Active,          (!) Administratively disabled,
            (a) SR Class A allowed, (b) SR Class B allowed.
```

History

This command was first available in ExtremeXOS 15.3.

The MSRP First Value Change Recovery Time, MSRP Ignore Latency Change, and MSRP Talker VLAN Pruning example outputs were added in 15.3.2.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

show msrp listeners

```
show msrp listeners {egress | ingress | ingress-and-egress} {port port_num}
{source-mac-addr source_mac_addr | stream-id stream_id}
```

Description

Shows MSRP listener information.

Syntax Description

msrp	Multiple Stream Registration Protocol.
listeners	Listener attributes.



egress	Display egress listeners only.
ingress	Display ingress listeners only (default).
ingress-and-egress	Display all listeners.
port_num	Filter based on ingress port number of the stream.
source-mac-addr	Filter based on source MAC address of a data stream.
stream-id	Filter based on stream ID of a data stream.

Default

N/A.

Usage Guidelines

Use this command to show MSRP listener information. The output can be filtered based on the stream id, source MAC, or port number on which the listener is registered.

Example

```
X460-24t.1 # show msrp listeners
      Stream Id          Port  Dec      Dir      State      Stream Age
                                     App  Reg  (days, hr:mm:ss)
-----
00:50:c2:4e:d3:2d:00:00    19  Ready  Ingress  VO  IN      0, 00:58:12
00:50:c2:4e:d3:2d:00:01    19  Ready  Ingress  VO  IN      0, 00:58:12
00:50:c2:4e:d3:2d:00:02    19  Ready  Ingress  VO  IN      0, 00:58:12
-----
-
App      : Applicant State,          Dec      : MSRP Declaration
Types,
Dir      : Direction of MSRP attribute,  Reg      : Registrar State

MSRP Declaration Types:
AskFail : Listener Asking Failed,    RdyFail : Listener Ready Failed,
Ready   : Listener Ready

Applicant States:
AA      : Anxious active,           AN      : Anxious new,
AO      : Anxious observer,        AP      : Anxious passive,
LA      : Leaving active,          LO      : Leaving observer,
QA      : Quiet active,            QO      : Quiet observer,
QP      : Quiet passive,           VN      : Very anxious new,
VO      : Very anxious observer,    VP      : Very anxious passive

Registrar States:
IN      : In - Registered,          LV      : Leaving - Timing out,
MT      : Empty - Not Registered
```



History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

show msrp ports

```
show msrp ports {port_list}{detail}
```

Description

Displays the MSRP configured port information.

Syntax Description

msrp	Multiple Stream Registration Protocol.
ports	Ports.
<i>port_list</i>	Port list separated by a comma or "-".
detail	Port information with more detail.

Default

N/A.

Usage Guidelines

Displays the MSRP configured port information. Specifying **detail** displays port information with more detail.

Example

```
# show msrp ports
      Port Enabled  Oper   Port  Dplx  Jumbo  Jumbo  Cls  Bndry
State Sr-Pvid
      Speed                Size                App/
Reg
-----
      IN      5  Y      Up/dbg      N      9216  A  N  QA/
      IN      2
      IN      2
      *21  Y      Up      1000 M Full  N      9216  A  N  QA/
```



```
IN          2
                                                    B   N   QA/
IN          2
```

```
-----
--
Flags       : (*) Active,                (!) Administratively disabled

App        : Applicant State,           Bndry   : Boundary,
Cls        : Traffic Class,            Dplx    : Duplex,
Oper       : MSRP Operational State,   Prop    : Propagated,
Reg        : Registrar State

MSRP Declaration Types:
  Adv      : Talker Advertise,          AskFail  : Listener Asking Failed,
  Fail     : Talker Fail,              RdyFail  : Listener Ready Failed,
  Ready    : Listener Ready

Applicant States:
  AA       : Anxious active,           AN        : Anxious new,
  AO       : Anxious observer,        AP        : Anxious passive,
  LA       : Leaving active,          LO        : Leaving observer,
  QA       : Quiet active,            QO        : Quiet observer,
  QP       : Quiet passive,          VN        : Very anxious new,
  VO       : Very anxious observer,   VP        : Very anxious passive

Registrar States:
  IN       : In - Registered,          LV        : Leaving - Timing out,
  MT       : Empty - Not Registered
```

```
#show msrp ports
      21 detail Port Enabled
      Oper Port Dplx Jumbo Jumbo Cls Bndry
      State Sr-Pvid
Size   App/Reg  ----  -----
-----
----- *21  Y
Up     1000 M Full N      9216
A   N   QA/IN
2
B   N   QA/IN
2
Talkers:
State                               Stream Id           Declaration
Prop App Reg                       ----- Rx -----
-----
VO                                     00:50:c2:4e:d3:2d:00:00 Adv Adv
IN                                     00:50:c2:4e:d3:2d:00:01 Adv Adv VO
IN
Listeners:
State                               Stream Id           Declaration
Prop App Reg                       ----- Rx -----
-----
VO                                     00:50:c2:4e:d3:3d:00:00 Ready Ready
VO
IN                                     00:50:c2:4e:d3:3d:00:01 Ready Ready VO
IN
```

```
-----
--  Flags       : (*) Active,                (!) Administratively disabled
App        : Applicant State,           Bndry
```



```

: Boundary Cls : Traffic Class, Dplx :
Duplex Oper : MSRP Operational State, Prop : Propagated
Reg : Registrar State MSRP Declaration
Types: Adv :
Talker Advertise, AskFail : Listener Asking Failed,
Fail : Talker
Fail, RdyFail : Listener Ready Failed, Ready :
Listener Ready Applicant
States: AA :
Anxious active, AN : Anxious new, AO :
Anxious observer, AP : Anxious passive, LA :
Leaving active, LO : Leaving observer, QA :
Quiet active, QO : Quiet observer, QP :
Quiet passive, VN : Very anxious new, VO : Very
anxious observer, VP : Very anxious passive Registrar
States: IN : In -
Registered, LV : Leaving - Timing out MT :
Empty - Not Registered
    
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

show msrp ports bandwidth

```
show msrp ports {port_list} bandwidth
```

Description

Displays bandwidth information of an MSRP port.

Syntax Description

msrp	Multiple Stream Registration Protocol.
ports	Ports.
<i>port_list</i>	Port list separated by a comma or "-".
bandwidth	Bandwidth information per port per traffic-class.

Default

N/A.



Usage Guidelines

Use this command to display bandwidth information of an MSRP port.

Example

```
# show msrp ports bandwidth
Port          Port      Class  Delta      Maximum      Reserved      Available
              Speed
-----
      5ab          0 M   A          75.00%      0.00%      0.00%      0.00%
                                   B          0.00%      0.00%      0.00%      0.00%
*21ab        1000 M   A          75.00%      75.00%      0.00%      75.00%
                                   B          0.00%      75.00%      0.00%      75.00%
```

Flags: (*) Active, (!) Administratively disabled,
(a) SR Class A allowed, (b) SR Class B allowed.

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

show msrp ports counters

```
show msrp ports {port_list} counters {event | packet}
```

Description

Shows PDU or event counters per port.

Syntax Description

msrp	Multiple Stream Registration Protocol.
ports	Ports.
<i>port_list</i>	Port list separated by a comma or "-".
counters	MSRP packet and attribute event counters.
event	MSRP attribute event counters.
packet	MSRP packet counters (default).



Default

N/A.

Usage Guidelines

Use this command to display PDU or event counters per port. The counters count the received attributes from talkers and listeners per attribute event, or the number of PDUs received. `show msrp counters` by itself displays PDU counters.

Example

```
#show msrp ports 17 counters packet
Port      Streams      Reservations  Rx Pkt      Rx Error     Tx Pkt
-----
17         0              0             2            0             2
```

```
#show msrp ports 17 counters event
Port : 17
      MRP Attribute Events      Rx      Tx
-----
In           250          56
JoinIn       0              0
JoinMt      224          386
Lv           0              0
Mt           0            152
New          0              0
```

```
      MSRP Declarations
-----
Listener Asking Failed      0          0
Listener Ready              56          8
Listener Ready Failed       0          0
Talker Advertise            8          56
Talker Failed               0          3
```

```
-----
In      : Not declared, but registered
JoinIn  : Declared and Registered
JoinMt  : Declared, but not registered
Lv      : Previously registered, but now withdrawn
Mt      : Not declared, and not registered
New     : Newly declared, and possibly not previously
         registered
```

History

This command was first available in ExtremeXOS 15.3.



Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

show msrp streams

```
show msrp streams {detail | propagation} {port port_num} {source-mac-addr
source_mac_addr | stream-id stream_id}{destination-mac-addr destination_mac_addr}
```

Description

Shows the MSRP stream information collected from the talker's attributes.

Syntax Description

msrp	Multiple Stream Registration Protocol.
streams	Data streams advertising QoS specification using MSRP.
detail	Show stream information with more detail.
propagation	Show stream propagation through switch.
port	Filter based on ingress port number of the stream.
source-mac-addr	Filter based on source MAC address of a data stream.
stream-id	Filter based on stream ID of a data stream.
destination-mac-addr	Filter based on destination MAC address of a data stream.

Default

N/A.

Usage Guidelines

Use this command to show the MSRP stream information collected from the talker's attributes. The output can be filtered based on the stream id, source MAC, destination MAC, or port number on which the stream is registered.

Example

```
# show msrp streams
  Stream Id           Destination      Port  Dec  VID  Cls/Rn  BW
-----
00:50:c2:4e:d3:2d:00:00  91:e0:f0:00:35:80   17  Adv   2  A/1    6.336
Mb
00:50:c2:4e:d3:2d:00:01  91:e0:f0:00:35:81   17  Adv   2  A/1    6.336
Mb
```



```

00:50:c2:4e:d3:2d:00:02  91:e0:f0:00:35:82    17  Adv    2  A/1    6.336
Mb
00:50:c2:4e:d3:2d:00:03  91:e0:f0:00:35:83    17  Adv    2  A/1    6.336
Mb
00:50:c2:4e:d3:2d:00:04  91:e0:f0:00:35:84    17  Adv    2  A/1    6.336
Mb
Total Streams: 5

```

```

-----
-
BW   : Bandwidth,           Cls   : Traffic Class,
Dec  : Prop Declaration Types, Rn     : Rank
(!)  : Talker pruned or forbidden
MSRP Declaration Types:
  Adv  : Talker Advertise,      AskFail : Listener Asking Failed,
  Fail : Talker Fail,          RdyFail : Listener Ready Failed,
  Ready : Listener Ready

```

```

#show msrp streams detail
  Stream Id          Destination      Port  Dec  VID  Cls/Rn  BW
-----
00:50:c2:4e:d3:2d:00:00  91:e0:f0:00:0e:80    17  Adv    2  A/1    6.336
Mb
      Accumulated Latency(nSec) : 0
      Max Frame Size             : 56
      Max Interval Frames        : 1
      Frame Rate (fps)           : 8000
00:50:c2:4e:d3:2d:00:01  91:e0:f0:00:0e:81    17  Fail    2  A/1    6.336
Mb
      Failure Code                : (10) Out of MSRP resrc
      Fail Bridge                  :
08:00:e0:e0:e0:e0:e0:e0
      Accumulated Latency(nSec) : 0
      Max Frame Size             : 56
      Max Interval Frames        : 1
      Frame Rate (fps)           : 8000
Total Streams: 2

```

```

-----
-
BW   : Bandwidth,           Cls   : Traffic Class,
Dec  : Prop. Declaration Types, Rn     : Rank
(!)  : Talker pruned or forbidden
MSRP Declaration Types:
  Adv  : Talker Advertise,      AskFail : Listener Asking Failed,
  Fail : Talker Fail,          RdyFail : Listener Ready Failed,
  Ready : Listener Ready

```

```

# show msrp streams propagation
  Stream Id          Destination      Port  Prop  VID  Cls/Rn  BW
-----
                                Dec
-----
00:50:c2:4e:d3:2d:00:00  91:e0:f0:00:35:80    17  Adv    2  A/1    6.336

```



Mb

```

Talker Propagation:
  Ingress      Ingress      Propagated      Propagated      Egress
  DecType      Port          DecType          Ports           DecType
  -----      -
  Adv    -->    17 --> Adv    -->    19 --> Adv
                                     21 --> Adv

Listener Propagation:
  Egress      Egress      Propagated      Listener      Ingress
  DecType      Port          DecType          Ports           DecType
  -----      -
  RdyFail <--    17 <-- Ready    <--    19 <-- Ready
                                     <--    21 <-- AskFail
  
```

Total Streams: 1

```

-----
-
BW      : Bandwidth,          Cls      : Traffic Class,
Dec     : Prop. Declaration Types, Rn       : Rank
(!)    : Talker pruned or forbidden
  
```

```

MSRP Declaration Types:
  Adv    : Talker Advertise,          AskFail : Listener Asking Failed,
  Fail   : Talker Fail,              RdyFail : Listener Ready Failed,
  Ready  : Listener Ready
  
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

show msrp talkers

```

show msrp talkers {egress | ingress | ingress-and-egress} {port port_num}{source-
mac-addr source_mac_addr | stream-id stream_id}
  
```

Description

Shows MSRP talker attributes.



Syntax Description

msrp	Multiple Stream Registration Protocol.
talkers	Talker attributes.
egress	Display egress talkers only (default).
ingress	Display ingress talkers only.
port	Filter based on ingress port number of the stream.
source-mac-addr	Filter based on source MAC address of a data stream.
stream-id	Filter based on stream ID of a data stream.

Default

N/A.

Usage Guidelines

Use this command to shows MSRP talker attributes. The output can be filtered based on the stream id, source MAC, or port number on which the talker is registered.

Example

```

X460-24t.1 # show msrp talkers
      Stream Id          Port  Dec   Dir   State          Failure Code
      -----
      -----
      App  Reg
      ---  ---
00:50:c2:4e:d3:2d:00:00    19  Adv   Egress  QA  MT  -
                                21  Fail  Egress  QA  MT  AVB incapbl
port(8)
00:50:c2:4e:d3:2d:00:01    19  Adv   Egress  QA  MT  -
                                21  Fail  Egress  QA  MT  AVB incapbl port(8)
-----
-
App      : Applicant State,          Dec   : MSRP Declaration Types,
Dir      : Direction of MSRP attribute,  Reg   : Registrar State

MSRP Declaration Types:
  Adv    : Talker Advertise,          Fail  : Talker Fail

Applicant States:
  AA     : Anxious active,            AN     : Anxious new,
  AO     : Anxious observer,          AP     : Anxious passive,
  LA     : Leaving active,            LO     : Leaving observer,
  QA     : Quiet active,              QO     : Quiet observer,
  QP     : Quiet passive,             VN     : Very anxious new,
  VO     : Very anxious observer,     VP     : Very anxious passive
    
```



Registrar States:

```

IN      : In - Registered,           LV      : Leaving - Timing out,
MT      : Empty - Not Registered

```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

show mvrp

show mvrp

Description

Shows MVRP settings.

Syntax Description

mvrp	Multiple VLAN Registration Protocol.
-------------	--------------------------------------

Default

N/A.

Usage Guidelines

Use this command to show the MVRP settings.

Example

```

X460-24t.3 # show mvrp
MVRP enabled           : Enabled
MVRP dynamic VLAN creation : Enabled
MVRP VLAN registration   : Forbidden
MVRP default STP domain   : s0
MVRP enabled ports      : 9    *11    *13
Flags: (*) Active, (!) Administratively disabled.

```

History

This command was first available in ExtremeXOS 15.3.



MRVP VLAN registration output was added in 15.3.2.

Platform Availability

This command is available on all commands.

show mvrp ports counters

```
show mvrp ports {port_list} counters {event | packet}
```

Description

Shows the port MVRP statistics. The statistics for packet or event counters are displayed as per input.

Syntax Description

mvrp	Multiple VLAN Registration Protocol.
ports	Ports.
<i>port_list</i>	List of ports in the switch separated by a comma or "-".
counters	MVRP packet and attribute event counters.
event	MVRP attribute event counters.
packet	MVRP packet counters (default).

Default

Packet counters.

Usage Guidelines

This command is to show the port MVRP statistics. The statistics for packet or event counters will be displayed as per input. The default is packet counters. The packet counters include Number of VLANs registered on the port, Number of Failed Registrations, Number of MVRPDUs received, Number of MVRPDUs sent, Number of erroneous MVRPDUs received, and the source address of the MVRP message last received by the port. The event counters include the number of different events received/transmitted.

Example

```
X460-24t.5 # show mvrp ports 9,11,13 counters packet
Port      VLANs    Failed   Rx Pkt  Rx Error  Tx Pkt   Last Source
      Regs  Regs    Count   Count    Count    Address
-----
          9      2       0       0         0       64  00:00:00:00:00:00
```



```

11      2      0      806836      0      433754      00:22:97:00:41:e7
13      2      0      784176      0      404794      00:22:97:00:41:e8

```

 Regs: Registrations

X460-24t.7 # show mvrp ports 9 counters event

Port : 17

MRP Attribute	Events	Rx	Tx
In		250	56
JoinIn		0	0
JoinMt		224	386
LeaveAll		5	0
Lv		0	0
Mt		0	152
New			

 In : Not declared, but registered
 JoinIn : Declared and Registered
 JoinMt : Declared, but not registered
 LeaveAll : All registrations will shortly be deregistered
 Lv : Previously registered, but now withdrawn
 Mt : Not declared, and not registered
 New : Newly declared, and possibly not previously registered

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

show mvrp tag

```
show mvrp tag vlan_tag {ports {port_list}}
```

Description

Shows the port specific applicant and registrar states and the configured control values for all MVRP enabled ports.

Syntax Description

mvrp	Multiple VLAN Registration Protocol.
tag	The 802.1Q VLAN ID.



<code>vlan_tag</code>	VLAN ID ranging from 1 to 4094 (default is 2).
<code>port_list</code>	Port list separated by comma or "-".

Default

N/A.

Usage Guidelines

Use this command to show the port specific applicant and registrar states and the configured control values for all MVRP enabled ports. The registrar control value is derived as follows:

- Normal = Dynamically ordered port.
- Fixed = Statically added port.
- Forbidden = VLAN is configured to be forbidden on the port.

Example

```
X460-24t.4 # show mvrp tag 2
```

Port	Applicant State	Applicant Control	Registrar State	Registrar Control	
-----	-----	-----	-----	-----	-----
	9	VN	On	MT	Normal
	11	QA	On	IN	Normal
	13	QA	On	IN	Normal

Applicant States:

```

AA      : Anxious active,      AN      : Anxious new,
AO      : Anxious observer,   AP      : Anxious passive,
LA      : Leaving active,     LO      : Leaving observer,
QA      : Quiet active,       QO      : Quiet
observer,
QP      : Quiet passive,      VN      : Very anxious
new,
VO      : Very anxious observer, VP      : Very anxious passive

```

Registrar States:

```

IN      : In - Registered,     LV      : Leaving - Timing out,
MT      : Empty - Not Registered

```

Applicant Control:

```

On      : Transmit On,       Off     : Transmit Off

```

Registrant Control:

```

Fixed   : Statically added,   Forbidden : Forbidden VLAN,
Normal  : Dynamically added

```



History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

show network-clock gtp

```
show network-clock gtp {default-set | current-set | parent-set | time-
properties-set}
```

Description

Displays global gPTP configuration and data.

Syntax Description

default-set	Displays this switch's native time capabilities.
current-set	Displays this switch's state relative to the grandmaster system.
parent-set	Displays the upstream (i.e., toward the grandmaster) system's parameters.
time-properties-set	Displays the grandmaster's parameters.

Default

N/A.

Usage Guidelines

Use this command to display whether gPTP is enabled in the switch and the ports on which gPTP is enabled.

Example

```
# show network-clock gtp
gPTP status:      Enabled
gPTP enabled ports: *1m,      *2s,      *3p,      *4d,      !5d,
                   11d      12d      13d      14d      *15d
                   *16d     17d      18d     19d     20d
                   21d     22d     23d     24d     25d
                   26d     27d     28d     29d     30d
                   31d     32d     33d     34d
Flags:            (*) Active, (!) Administratively disabled,
                  (d) Disabled gPTP port role, (m) Master gPTP port role,
                  (p) Passive gPTP port role, (s) Slave gPTP port role

# show network-clock gtp default-set
```



```

Local Clock Identity           : 00:04:96:FF:FE:52:2C:BE
Number of gPTP ports          : 24
Local Clock Class              : 255 (slave only clock)
Local Clock Accuracy           : 254 (unknown)
Local Offset Scaled Log Variance : 65535
GM Capable                     : No
Local Priority1                 : 255
Local Priority2                 : 248
Current UTC Offset             : unknown
Leap 59                         : No
Leap 61                         : No
Time Traceable                 : No
Frequency Traceable            : No
Time Source                     : 160 (Internal Oscillator)

# show network-clock gtp current-set
Steps Removed                  : 1
Offset from GM                 : 10 nanoseconds
Last GM Phase Change           : 548 nanoseconds
Last GM Frequency Change       : 100
GM Time Base Indicator         : 2
GM Change Count                : 1
Last GM Change Event           : Tue Nov 22 03:32:07 2011
Last GM Frequency Change Event : Tue Nov 22 03:32:07 2011
Last GM Phase Change Event     : Tue Nov 22 03:32:07 2011

# show network-clock gtp parent-set
Parent Clock Identity          : 00:04:96:FF:FE:52:34:5F
Parent port number             : 21
Cumulative Rate Ratio          : 10000
GM Clock Identity              : 00:12:34:FF:FE:56:78:9A
GM Clock Accuracy              : 32 (25 ns)
GM Offset Scaled Log Variance  : 32767
GM Priority1                   : 245
GM Priority2                   : 248

# show network-clock gtp time-properties-set
Current UTC Offset             : 33 seconds
Leap 59                         : No
Leap 61                         : No
Time Traceable                 : Yes
Frequency Traceable            : Yes
Time Source                     : 32 (GPS)

```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.



show network-clock gptp ports

```
show network-clock gptp ports [port_list | all] {counters}
```

Description

Displays gPTP port parameters and counters.

Syntax Description

<i>port_list</i>	Specifies one or more of the switch's physical ports.
all	Specifies all of the switch's physical ports.

Default

N/A.

Usage Guidelines

The command `show network-clock gptp port` displays the specified port's gPTP parameters:

Physical port number	The switch's number for this physical port.
gPTP port status	Indicates whether gPTP is enabled on this port.
Clock Identity	This switch's gPTP Clock Identity.
gPTP Port Number	gPTP number for this physical port.
IEEE 802.1AS Capable	Indicates whether this switch and the neighboring systemdevice connected via this port can interoperate via gPTP.
Port Role	The port's gPTP role: <ul style="list-style-type: none"> • Disabled (3) • Master (6) • Passive (7) • Slave (9)
Announce Initial Interval	The initial announce interval on this port. The interval is represented as the log base 2 of the interval in seconds; for example, 0 = 1 second.
Announce Current Interval	The current announce interval on this port. The interval is represented as the log base 2 of the interval in seconds; for example, 1 = 2 seconds.
Announce Receipt Timeout	The number of announce intervals a slave port waits without receiving an Announce message before it assumes the master port is no longer sending Announce messages and the BMCA needs to be run.
Sync Initial Interval	The initial time-synchronization transmission interval on this port. The interval is represented as the log base 2 of the interval in seconds; for example, -1 = 500 milliseconds.
Sync Current Interval	The current time-synchronization transmission sync interval on this port. The interval is represented as the log base 2 of the interval in seconds; for example, -2 = 250 milliseconds.



Sync Receipt Timeout	The number of time-synchronization transmission intervals a slave port waits without receiving a Sync message before it assumes the master port is no longer sending Sync messages and the BMCA needs to be run.
Sync Receipt Timeout Interval	Sync Receipt Timeout in time units.
Peer Delay Initial Interval	The initial Peer Delay Request interval on this port. The interval is represented as the log base 2 of the interval in seconds; for example, 2 = 4 seconds.
Peer Delay Current Interval	The current Peer Delay Request interval on this port. The interval is represented as the log base 2 of the interval in seconds; for example, 3 = 8 seconds.
Peer Delay Allowed Lost Responses	The number of consecutive Peer Delay Request messages that the switch must send on this port without receiving a valid response before it considers the port not to be exchanging Peer Delay messages with its neighbor.
Measuring Propagation Delay	Indicates whether this port is measuring its link's propagation delay.
Mean Propagation Delay	The link's estimated one-way propagation delay. The peer delay protocol measures the sum of the link's propagation delays in each direction, and this is that sum divided by two, which is accurate only if the link is symmetrical.
Mean Propagation Delay Threshold	The propagation delay above which the switch considers this port unable to run gPTP.
Propagation Delay Asymmetry	The configured time that the propagation delay from this switch to the neighbor is less than the estimated one-way propagation delay between the switch and its neighbor (which is also the time that the propagation delay from the neighbor to this switch is greater than the estimate). This value is negative if the propagation delay to the neighbor is greater than the estimate. Let t_{IR} be the propagation delay from this switch (initiator) to the neighbor (responder), t_{RI} be the propagation delay from the neighbor to this switch, and $meanPathDelay$ be the estimated one-way propagation delay. Then: $meanPathDelay = (t_{IR} + t_{RI}) / 2$ $t_{IR} = meanPathDelay - asymmetry_time$ $t_{RI} = meanPathDelay + asymmetry_time$
Neighbor Rate Ratio	The estimated ratio of the frequency of the local clock in the neighboring systemdevice connected via this port, to this switch's local clock's frequency. The ratio is represented as the ratio minus 1, multiplied by 2^{41} : $(ratio - 1) * 2^{41}$
PTP Version	The PTP version number used on this port. Always 2.

The command `show network-clock gptp port counters` displays the specified port's gPTP counters:

Physical port number	The switch's number for this physical port.
gPTP port status	Indicates whether gPTP is enabled on this port.
Announce	The number of Announce messages received and sent.
Sync	The number of Sync messages received and sent.
Follow Up	The number of Follow Up messages received and sent.
Peer Delay Request	The number of Peer Delay Request messages received and sent.
Peer Delay Response	The number of Peer Delay Response messages received and sent.
Peer Delay Response Followup	The number of Peer Delay Response Follow Up messages received and sent



gPTP packet discards	The number of received gPTP packets discarded or lost for one of the following reasons (from 802.1AS-2011 14.7.8): <ul style="list-style-type: none"> • Announce message from this switch • Announce message with stepsRemoved >= 255 • Announce message with a Path Trace TLV that includes this switch • Follow Up message not received following Sync message received • Peer Delay Response message not received following Peer Delay Request message sent • Peer Delay Response Follow Up message not received following Peer Delay Request message sent
Announce Receipt Timeout Count	The number of Announce Receipt timeouts.
Sync Receipt Timeout Count	The number of Sync Receipt timeouts.
Peer Delay Allowed Lost Responses Exceeded Count	The number of times the number of consecutive Peer Delay Request messages sent without receiving a valid response exceeded the Peer Delay Allowed Lost Responses.

Example

```
# show network-clock gtp ports 2
Physical port number      : 2
gPTP port status         : Enabled
Clock Identity           : 00:04:96:FF:FE:52:2C:BE
gPTP Port Number         : 2
IEEE 802.1AS Capable     : Yes
Port Role                 : 9 (Slave)
Announce Initial Interval : 0 (1 second)
Announce Current Interval : 1 (2 seconds)
Announce Receipt Timeout : 3
Sync Initial Interval     : -3 (125 milliseconds)
Sync Current Interval     : -2 (250 milliseconds)
Sync Receipt Timeout      : 3
Sync Receipt Timeout Interval : 750 milliseconds
Peer Delay Initial Interval : 2 (4 seconds)
Peer Delay Current Interval : 4 (8 seconds)
Peer Delay Allowed Lost Responses : 3
Measuring Propagation Delay : Yes
Mean Propagation Delay    : 1000 nanoseconds
Mean Propagation Delay Threshold : 10000 nanoseconds
Propagation Delay Asymmetry : 0
Neighbor Rate Ratio       : 200
PTP Version               : 2
```

```
# show network-clock gtp ports 3 counters
Physical port number      : 3
gPTP port status         : Enabled
-----
Parameter                Receive      Transmit
-----
Announce                  1000          2000
Sync                     1000           500
Follow Up                 2000          2500
Peer Delay Request        3000          1000
```



```

Peer Delay Response                500          1500
Peer Delay Response Follow Up      200          1000
gPTP packet discards              2000         -
-----
Announce Receipt Timeout Count      : 1000
Sync Receipt Timeout Count          : 500
Peer Delay Allowed Lost Responses Exceeded Count : 2000

```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

unconfigure avb

unconfigure avb

Description

This command is a macro command that can be used to unconfigure all AVB protocols globally on the switch. It is equivalent to issuing the following four commands:

```
unconfigure mvrp
```

```
unconfigure msrp
```

```
unconfigure network-clock gptp
```

```
unconfigure mrp ports all
```

Syntax Description

avb	Audio Video Bridging
------------	----------------------

Default

N/A.

Usage Guidelines

Use this command to unconfigure all AVB protocols globally on the switch.



Example

```
unconfigure avb
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

unconfigure mrp ports timers

```
unconfigure mrp ports [port_list | all] {timers {refresh} {join} {leave} {leave-all} {periodic}}
```

Description

Unconfigure MRP timers, or only reset the MRP timer values to default if the **timer** keyword is specified.

Syntax Description

mrp	Multiple Registration Protocol.
ports	Ports on which MRP timers are to be configured.
all	All ports.
timers	Multiple Registration Protocol timers.
refresh	Timer value to use in place of regular leave timer, only in cases when leave-all is received or sent.
join	The time interval to delay sending MRP advertisements.
leave	The time interval to wait in the leaving state before transitioning to the empty state.
leave-all	The time interval used to control the frequency of "leave all" messages.
periodic	The time interval between two periodic events.

Default

The default values for join, leave, leave-all, periodic and extended-refresh timers are 200, 600, 10000, 1000, and 0 milliseconds, respectively.



Usage Guidelines

Use this command to unconfigure MRP timers, or only reset the MRP timer values to default if the **timer** keyword is specified. If none of the timers are specified, this command resets all three timers to the default values. The default values for the join, leave, and leave-all timers are 200, 600, and 10000 ms respectively.

Example

```
unconfigure mrp ports all
unconfigure mrp ports all timers
unconfigure mrp ports all timers join
```

History

This command was first available in ExtremeXOS 15.3.

The **extended-refresh** and **periodic timer** options were added in 15.3.2.

Platform Availability

This command is available on all platforms.

unconfigure msrp

```
unconfigure msrp {ports [port_list | all]}
```

Description

Disables MSRP and removes all configuration. If a list of ports is specified, MSRP is disabled and the related configuration is removed only on the ports and the system-wide MSRP configuration stays intact.

Syntax Description

msrp	Multiple Stream Registration Protocol.
<i>port_list</i>	List of ports in the switch.
all	All the ports in the switch.

Default

N/A.



Usage Guidelines

Use this command to disable MSRP and remove all configuration. If a list of ports is specified, MSRP is disabled and the related configuration is removed only on the ports and the system-wide MSRP configuration stays intact.

Example

```
unconfigure msrp
unconfigure msrp ports all
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.

unconfigure mvrp

unconfigure mvrp

Description

Unconfigures MVRP on a switch and all MVRP port and bridge settings.

Syntax Description

mvrp	Multiple VLAN Registration Protocol.
-------------	--------------------------------------

Default

N/A.

Usage Guidelines

Use this command to unconfigure MVRP on a switch. This command unconfigures all MVRP port and bridge settings.

Example

The following command unconfigures MVRP:

```
unconfigure mvrp
```



History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

unconfigure mvrp stpd

unconfigure mvrp stpd

Description

Resets the MVRP STP domain.

Syntax Description

mvrp	Multiple VLAN Registration Protocol.
stpd	The STP domain the VLAN is to be associated with. All ports of the domain will be advertised when this VLAN is registered.

Default

s0.

Usage Guidelines

Use this command to reset the STP domain associated with a particular VLAN or all VLANs to default. If a VLAN is specified, the specific VLAN will be associated to the default STP, which is configured using the `configure mvrp stpd stpd_name default` command. If VLAN is not specified, all VLANs are associated to STP domain s0.

Example

The following are examples of unconfigure mvrp stpd commands:

```
unconfigure mvrp stpd
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.



unconfigure mvrp tag

```
unconfigure mvrp tag vlan_tag
```

Description

Resets all MVRP settings for the given VLAN id. The STP domain, the registrar state machine settings, applicant state machine settings for the given VLAN are reset to default values.

Syntax Description

mvrp	Multiple VLAN Registration Protocol.
<i>vlan_tag</i>	The 802.1Q VLAN ID.

Default

N/A.

Usage Guidelines

Use this command to reset all MVRP settings for the given VLAN id. The STP domain, the registrar state machine settings, and the applicant state machine settings for the given VLAN are reset to default values. All dynamically added ports of the VLAN are removed. If the VLAN was created dynamically, it is removed. If VLAN is not specified, MVRP settings for all VLANs are reset and the dynamic VLAN creation feature is reset to “enabled”.

Example

The following example shows unconfiguring an MVRP:

```
unconfigure mvrp tag 100
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

unconfigure network-clock gptp ports

```
unconfigure network-clock gptp ports [port_list | all]
```



Description

Restores all configuration parameters on the specified ports to their default values. This command does not disable gPTP on the ports.

Syntax Description

<i>port_list</i>	Specifies one or more of the switch's physical ports.
all	Specifies all of the switch's physical ports.

Default

N/A.

Usage Guidelines

Use this command to restore all configuration parameters on the specified ports to their default values.

Example

```
unconfigure network-clock gtp ports all
unconfigure network-clock gtp ports 1,2
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X430, X440, X460, and X670 switches if the AVB feature pack license is installed on the switch.



19 TRILL Commands

```
clear trill counters
configure trill add access tag
configure trill add network vlan
configure trill appointed-forwarder
configure trill delete access tag
configure trill delete network vlan
configure trill designated-vlan
configure trill inhibit-time
configure trill mtu probe
configure trill mtu probe fail-count
configure trill mtu size
configure trill nickname
configure trill ports
configure trill ports protocol
configure trill pseudonode
configure trill system-id
configure trill timers csnp
configure trill timers hello
configure trill timers lsp
configure trill timers spf
configure trill timers spf backoff-delay
configure trill tree prune vlan
create trill nickname
delete trill nickname
disable trill
enable trill
show trill
show trill distribution-tree
show trill lsdb
show trill neighbor
show trill ports
show trill rbridges
traceroute trill
```



clear trill counters

```
clear trill counters
```

Description

This command clears all protocol and port counters associated with TRILL.

Syntax Description

trill	Transparent Interconnection of Lots of Links
counters	Reset all TRILL counters to zero.

Default

N/A.

Usage Guidelines

Use this command to clear all protocol and port counters associated with TRILL.

Example

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



configure trill add access tag

```
configure trill add access tag first_tag {-last_tag}
```

Description

This command administratively adds a VLAN tag or range of VLAN tags that represent edge (or access) VLANs attached to the TRILL network. The VLAN can be created using any method, for example, via the CLI, XNV, ID Manager, or MVRP.



Syntax Description

trill	Transparent Interconnection of Lots of Links.
add	Add entities to TRILL.
access	Access VLAN.
tag	VLAN ID values.
<i>first_tag</i>	Start VLAN ID between 1 and 4094";type="int32_t";range="[1,4094]
-	Range separator.
<i>last_tag</i>	End VLAN ID between 1 and 4094";type="int32_t";range="[1,4094]

Default

N/A.

Usage Guidelines

Use this command to add a VLAN tag or range of VLAN tags that represent edge (or access) VLANs attached to the TRILL network. The VLAN can be created using any method, for example, via the CLI, XNV, ID Manager, or MVRP. If the VLAN is already associated with a TRILL network, the operation is ignored and the command continues configuring the remaining VLAN IDs in the specified set. The *last_tag* is optional. When specified, all VLANs including the *first_tag* and *last_tag* are configured as TRILL Access VLANs. If the *last_tag* is not specified, only the *first_tag* is configured as a TRILL access VLAN. By default, no VLANs are configured as access VLANs when TRILL is enabled. VLAN one (also called the Default VLAN) is also the default TRILL Network VLAN used to carry TRILL data traffic. TRILL Hellos are sent on all TRILL access VLANs. TRILL Hellos sent on access VLANs are sent with the AC flag set, indicating that the VLAN link should *not* be used in the TRILL shortest best path and distribution tree topology calculations. The valid tag range is one to 4094.

Example

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



configure trill add network vlan

```
configure trill add network vlan vlan_name
```



Description

This command configures the TRILL protocol on the specified VLAN. If the specified VLAN name does not exist, an error is returned. Once TRILL is added to a VLAN, the VLAN may become a designated TRILL VLAN and carry TRILL data traffic. If TRILL is already configured on *vlan_name*, the command completes with no errors.

Syntax Description

trill	Transparent Interconnection of Lots of Links.
add	Add entities to TRILL.
network	Trunk network.
vlan	VLAN.
<i>vlan_name</i>	VLAN name.

Default

N/A.

Usage Guidelines

Use this command to configure the TRILL protocol on the specified VLAN. If the specified VLAN name does not exist, an error is returned. Once TRILL is added to a VLAN, the VLAN may become a designated TRILL VLAN and carry TRILL data traffic. If TRILL is already configured on *vlan_name*, the command completes with no errors.

There are three types of TRILL VLANs: access, network, and hybrid. An access VLAN carries only native Ethernet data traffic. Traffic associated with TRILL edge VLANs is either encapsulated for transmit over the TRILL network or decapsulated from the TRILL network and transmitted as native tagged or untagged traffic on the local VLAN. TRILL Network VLANs carry only TRILL encapsulated VLAN traffic. Any native Ethernet traffic received on a TRILL Network VLAN is discarded and TRILL encapsulated packets are never decapsulated and transmitted as native Ethernet traffic on a TRILL Network VLAN. Link state information associated with Network TRILL VLANs is advertised and the associated links to each RBridge neighbor are used in SPF calculations when building broadcast/multicast forwarding trees. By default, when TRILL is enabled, the Default VLAN is a TRILL Network VLAN.

If the switch receives a TRILL Hello packet on a VLAN that is not configured for TRILL Network or Access VLAN, an error message is logged. If the DRB specifies a designated VLAN ID and there is no TRILL configured Network VLAN with that VLAN ID, then an error message is logged and the TRILL peer session is not established. It is recommended that the TRILL Network VLAN ID be one (the *Default* VLAN), which is the default TRILL Network VID stated in the TRILL IETF RFCs. This should allow networks to be automatically configured and easily setup for network operation.

TRILL hybrid VLANs are VLANs that are configured as TRILL Network VLANs and a TRILL Access VLAN and thus operate as both an Access and Network TRILL VLAN. TRILL Hybrid VLANs are not supported.



Example

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



configure trill appointed-forwarder

```
configure trill appointed-forwarder {distribution [shared {reappoint-mode [sticky
| balanced]} | priority]} {timer hold-time seconds }
```

Description

This command configures the algorithm the RBridge will use to appoint forwarders if the RBridge is the Designated RBridge (DRB). By default, the DRB will assign VLAN IDs from the common set of access VLANs configured by the RBridges on the shared link based on RBridge priority.

Syntax Description

trill	Transparent Interconnection of Lots of Links.
appointed-forwarder	Appointed Forwarder control. Only valid if the RBridge is the Designated RBridge.
distribution	Algorithm for assigning Appointed Forwarders.
shared	Evenly distribute the VLANs across all RBridges (default).
reappoint-mode	Method for managing Appointed Forwarders when new RBridges become active or inactive.
sticky	Newly active RBridges don't change current VLAN assignments. Newly inactive RBridge VLANs are reappointed (default).
balanced	Forces a recalculation of VLAN assignments for all VLANs when an RBridge goes active or inactive.
priority	The RBridge with the highest priority becomes the Appointed Forwarder for the VLAN.
timer	Timers for Appointed Forwarders.
hold-time	Amount of time Designated RBridge will remember VLAN appointments for an RBridge that has gone inactive.
<i>seconds</i>	Hold time in seconds (0-600). (Default 180).



Default

Shared, 180 seconds.

Usage Guidelines

Use this command to configure the algorithm the RBridge will use to appoint forwarders if the RBridge is the Designated RBridge (DRB). By default, the DRB will assign VLAN IDs from the common set of access VLANs configured by the RBridges on the shared link based on RBridge priority. This provides the simplest algorithm for assigning Access VLANs. The shared keyword assigns each access VLAN to the set of connected RBridges such that each RBridge has the same number of forwarding VLANs. If an RBridge becomes unavailable, the DRB reassigns VLAN IDs that were appointed to the down RBridge to one of the remaining active RBridges. If the DRB fails, then the next highest priority RBridge assumes the role of the DRB and takes over the responsibility of assigning appointed forwarders for each VLAN. The default mode of appointing RBridge forwarding status is priority.

If the distribution algorithm is shared, the optional reappoint-mode keyword configures how VLAN forwarding appointments are managed as RBridges become active or go inactive on a shared VLAN distribution link. The default mode is sticky. In sticky reappoint mode, VLAN forwarding appointments are maintained so long as the appointed forwarder RBridge is active. Only the VLANs that are appointed to a failed RBridge are reappointed. This mode provides the most consistent and least packet forwarding disruption when an RBridge fails. The downside to the sticky mode is that over time, VLAN appointments between RBridges may become unbalanced. The balanced mode maintains VLAN appointments based on the VLAN appointed forwarder algorithm. If an RBridge fails, The DRB recalculates VLAN appointments after the appointed-forwarder hold time expires and communicates the new assignments to the remaining active RBridges.

Example

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



configure trill delete access tag

```
configure trill delete access tag first_tag {-last_tag}
```



Description

This command administratively deletes a VLAN tag or range of VLAN tags from the configured access VLAN set. VLANs that match the specified VLAN ID to be deleted are detached (disassociated) from the TRILL network.

Syntax Description

trill	Transparent Interconnection of Lots of Links.
delete	Delete entities from TRILL.
access	Access VLAN.
tag	VLAN ID values.
<i>first_tag</i>	Start VLAN ID between 1 and 4094";type="int32_t";range="[1,4094]
-	Range separator.
<i>last_tag</i>	End VLAN ID between 1 and 4094";type="int32_t";range="[1,4094]

Default

N/A.

Usage Guidelines

Use this command to delete a VLAN tag or range of VLAN tags from the configured access VLAN set. VLANs that match the specified VLAN ID to be deleted are detached (disassociated) from the TRILL network. If the VLAN has already been deleted from the TRILL network, the operation is ignored and the command continues disassociating the remaining specified VLAN tags. The *last_tag* is optional. When specified, all VLANs including the *first_tag* and *last_tag* are disassociated from the TRILL network as access VLANs. If the *last_tag* is not specified, then only the VLAN identified by the *first_tag* is disassociated with the TRILL network. By default, no VLANs are configured as access VLANs when TRILL is enabled. The VLAN of one (also called the Default VLAN) also defaults to the TRILL Network (or native) VLAN and is used to carry TRILL data traffic. If a VLAN is disassociated from the TRILL network, the RBridge will cease sending TRILL Hellos on the specified VLAN. If the RBridge was the appointed forwarder for the VLAN, one of the remaining RBridges (if any) is assigned the appointed forwarding status for the VLAN. If there are no other RBridges connected to the VLAN, traffic for this VLAN is no longer transported across the TRILL network and any received TRILL traffic destined for devices attached to the VLAN is discarded.

Example

History

This command was first available in ExtremeXOS 15.4.



Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



configure trill delete network vlan

```
configure trill delete network vlan [ vlan_name | all ]
```

Description

This command removes TRILL from the specified network VLAN from the TRILL configuration. Once deleted, no TRILL data packets are sent and any TRILL packets received are discarded. Any RBridge links associated with the Network VLAN are removed from the database and the updated information is communicated to remaining known RBridge peers. Deleting TRILL from a VLAN or deleting the VLAN entirely may trigger RBridges to recalculate their dtrees. There must be at least one TRILL network VLAN configured for the RBridge to participate in the TRILL network.

Syntax Description

trill	Transparent Interconnection of Lots of Links.
delete	Delete entities to TRILL.
network	Trunk network.
vlan	VLAN.
<i>vlan_name</i>	VLAN name.
all	All VLANs.

Default

N/A.

Usage Guidelines

Use this command to remove TRILL from the specified network VLAN from the TRILL configuration. Once deleted, no TRILL data packets are sent and any TRILL packets received are discarded. Any RBridge links associated with the Network VLAN are removed from the database and the updated information is communicated to remaining known RBridge peers. Deleting TRILL from a VLAN or deleting the VLAN entirely may trigger RBridges to recalculate their dtrees. There must be at least one TRILL network VLAN configured for the RBridge to participate in the TRILL network.

Example



History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



configure trill designated-vlan

```
configure trill designated-vlan vlan_name
```

Description

This command configures the desired VLAN to use as the designated TRILL VLAN. The designated VLAN is the VLAN that is used to carry TRILL data traffic between RBridges.

Syntax Description

trill	Transparent Interconnection of Lots of Links.
designated-vlan	Desired designated VLAN.
<i>vlan_name</i>	Desired designated VLAN name.

Default

N/A.

Usage Guidelines

Use this command to configure the desired VLAN to use as the designated TRILL VLAN. The designated VLAN is the VLAN that is used to carry TRILL data traffic between RBridges. By default, the designated-vlan is the default VLAN, VLAN ID = 1. The designated VLAN may use any VLAN tag from one through 4094. Only one VLAN per VR can be configured as the designated VLAN. In the event that two RBridges are configured with different Designated VLAN IDs, the lowest value VLAN ID is used as the designated VLAN tag. In this event, a informational message is logged indicating that another RBridge was detected attempting to communicate using a different designated VLAN ID.

Example

History

This command was first available in ExtremeXOS 15.4.



Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



configure trill inhibit-time

configure trill inhibit-time *seconds*

Description

This command configures the amount of time the RBridge does not forward traffic to end stations after detecting a root RBridge topology change or under certain conditions when the appointed forwarder status changes.

Syntax Description

trill	Transparent Interconnection of Lots of Links.
inhibit-time	Amount of time the RBridge does not forward traffic after detecting an RBridge topology change.
<i>seconds</i>	Inhibit time in seconds (0-30). (Default 15).

Default

15.

Usage Guidelines

Use this command to configure the amount of time the RBridge does not forward traffic to end stations after detecting a root RBridge topology change or under certain conditions when the appointed forwarder status changes. The inhibit time is set to avoid packet loops when there are network topology changes. The parameter *seconds* has a valid range of zero to 30 seconds and defaults to 15 seconds. A value of zero means that no packets will be inhibited from being forwarded.

Example

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.





configure trill mtu probe

```
configure trill mtu probe [enable | disable]
```

Description

This command configures the TRILL MTU probe protocol. When enabled, the RBridge will transmit an MTU probe to known peers to validate that the RBridge peers are reachable using the configured MTU size.

Syntax Description

trill	Transparent Interconnection of Lots of Links.
mtu	Maximum Transmission Unit.
probe	Send MTU probes to determine and verify the configured MTU size supported on TRILL links.
enable	Enable MTU probing (default).
disable	Disable MTU probing.

Default

Enabled.

Usage Guidelines

Use this command to configure the TRILL MTU probe protocol. When enabled, the RBridge will transmit an MTU probe to known peers to validate that the RBridge peers are reachable using the configured MTU size. MTU probes are only transmitted and acknowledged on the each link's designated VLAN. The MTU probe packet is padded to the MTU size tested. When MTU probes are disabled, no MTU probe packets are transmitted and the network administrator must ensure that all TRILL links support the configured MTU size. The default MTU probe setting is enable.

Example

History

This command was first available in ExtremeXOS 15.4.



Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



configure trill mtu probe fail-count

```
configure trill mtu probe fail-count probes_sent
```

Description

This command configures the MTU probe failure count. RBridges send MTU probes to known peers to determine and verify that the configured MTU size is supported on TRILL links. MTU probes are only transmitted and acknowledged on the each link's designated VLAN.

Syntax Description

trill	Transparent Interconnection of Lots of Links.
mtu	Maximum Transmission Unit.
probe	Send MTU probes to determine and verify the configured MTU size supported on TRILL links.
fail-count	Number of probe requests sent out before declaring failure.
<i>probes_sent</i>	Number of probes sent to determine the MTU size (1-10). (Default 2.)

Default

2.

Usage Guidelines

Use this command to configure the MTU probe failure count. RBridges send MTU probes to known peers to determine and verify that the configured MTU size is supported on TRILL links. MTU probes are only transmitted and acknowledged on the each link's Designated VLAN. The MTU probe packet is padded to the MTU size tested. If the RBridge fails to receive an MTU acknowledgement after fail-count attempts, the MTU size test fails. When the RBridge's configured MTU size is not supported, the RBridge attempts to determine the peer's MTU setting by lowering the MTU size and resending the MTU probe. A modified binary algorithm of lowering and increasing the MTU size continues until the exact MTU size is determined. MTU probe packets are sent at 500ms intervals until either an acknowledgement is received or the configured probe fail-count is reached. The default fail-count setting is 2 and the valid range is 1 to 10. Incoming Ethernet packets that are larger than the supported MTU size and cannot be entirely encapsulated into a TRILL packet are discarded.

Example



History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



configure trill mtu size

```
configure trill mtu size mtu_size
```

Description

This command configures the TRILL MTU size. Jumbo frames must be enabled on TRILL Network ports. Jumbo frames allow fully formed non-jumbo Ethernet packet to be encapsulated and transported across the TRILL network without fragmentation.

Syntax Description

trill	Transparent Interconnection of Lots of Links.
mtu	Maximum Transmission Unit.
<i>mtu_size</i>	Size of Maximum Transmission Unit in octets (1470-9188). (Default 1518).

Default

1518 octets.

Usage Guidelines

Use this command to configure the TRILL MTU size. Jumbo frames must be enabled on TRILL Network ports. Jumbo frames allow fully formed non-jumbo Ethernet packet to be encapsulated and transported across the TRILL network without fragmentation. The default TRILL MTU size is 1518 octets. Thus the minimum jumbo frame size is 1518 plus the size of the TRILL header plus any TRILL options. The minimum TRILL header size with no options is 28 octets.

Example

History

This command was first available in ExtremeXOS 15.4.



Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



configure trill nickname

```
configure trill nickname [nickname_id | nickname_string ] {new-nickname
new_nickname_id} {nickname-priority new_id_priority} {root-priority
new_root_priority} {name new_nickname_string}
```

Description

This command configures nickname parameters. By default, the switch randomly chooses a nickname and verifies that the nickname is not already being used by another RBridge in the TRILL network.

Syntax Description

trill	Transparent Interconnection of Lots of Links
nickname	Nickname
<i>nickname_id</i>	Identifier between 1 and 0xFFBF in hex.
<i>nickname_string</i>	Name string up to 32 characters.
new-nickname	Priority value between 128 and 255. Lower numbers represent lower priority. Default is 192.
<i>new_nickname_id</i>	Identifier between 1 and 0xFFBF in hex";type="hex_t";range="[1,65471].
nickname-priority	Nickname priority.
<i>new_id_priority</i>	Priority value between 128 and 255. Lower numbers represent lower priority.
root-priority	Root priority.
<i>new_root_priority</i>	Priority value between 0 and 65535. Lower numbers represent lower priority.";type="uint16_t";range="[0,65535].
name	Human readable name associated with nickname.
<i>new_nickname_string</i>	Name string up to 32 characters";type="string";range="[1,32].

Default

N/A.

Usage Guidelines

Use this command to configure the nickname parameters. By default, the switch randomly chooses a nickname and verifies that the nickname is not already being used by another RBridge in the TRILL network. Each created nickname is advertised to the other RBridges in the TRILL network. If the nickname is randomly generated, the nickname default priority is 0x00 (lowest priority). Optionally, nickname string name may be configured.



Example

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



configure trill ports

```
configure trill ports [port_list | all] {link-type [broadcast | point-to-point] }
{metric [ metric | automatic] } {drb-election priority}
```

Description

Syntax Description

trill	Transparent Interconnection of Lots of Links.
ports	Port options.
<i>port_list</i>	Port list to configure.
all	Modify all ports in the system.
link-type	Interface type for link.
broadcast	Indicates there may be multiple R Bridges on the link. Does not run Hello Protocol (default).
point-to-point	Assumes only one R Bridge on the link. Does not run Hello Protocol.
metric	Link cost metric.
<i>metric</i>	Manual link cost metric. Range is 1 - 16777214.
automatic	Automatically determine link cost metric (default).
drb-election	Designated R Bridge election.
<i>priority</i>	Priority value between 0 and 127. Lower numbers represent lower priority. (Default is 64).

Default

Defaults are broadcast, automatic, and priority = 64.



Usage Guidelines

Example

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



configure trill ports protocol

```
configure trill ports [port_list | all] protocol [enable | disable]
```

Description

This command enables and disables the TRILL Hello protocol per port. By default, the Hello protocol is enabled on all ports. When enabled the RBridge is able to discover its neighbors and to exchange topology information.

Syntax Description

trill	Transparent Interconnection of Lots of Links.
ports	Port options.
<i>port_list</i>	Port list to configure.
all	Modify all ports in the system.
protocol	Enable or disable the TRILL protocol on port.
enable	Enable (default).
disable	Disable.

Default

Enabled.

Usage Guidelines

Use this command to enable and disable the TRILL Hello protocol per port. By default, the Hello protocol is enabled on all ports. When enabled the RBridge is able to discover its neighbors and to



exchange topology information. This information is used to dynamically program the data plane for TRILL packet forwarding. If the TRILL Hello protocol is disabled, the RBridge is not able to peer with any other RBridges in the TRILL network on the disabled port. This command can be used to reduce the number of Hello packets transmitted on TRILL Access VLANs, by disabling TRILL on ports that are not connected to a peer RBridge.

Example

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



configure trill pseudonode

```
configure trill pseudonode [enable | disable]
```

Description

This command configures the RBridge to represent multiple devices on a shared link as being connected to a TRILL pseudonode. By default, TRILL will automatically allocate a pseudonode ID when two or more devices are connected to the RBridge on the same interface.

Syntax Description

trill	Transparent Interconnection of Lots of Links.
pseudonode	Allow bypassing of pseudonodes on links.
enable	Enable bypass pseudonodes.
disable	Disable bypass pseudonodes.

Default

Enabled.

Usage Guidelines

Use this command to configure the RBridge to represent multiple devices on a shared link as being connected to a TRILL pseudonode. By default, TRILL will automatically allocate a pseudonode ID when



two or more devices are connected to the RBridge on the same interface. Enabling pseudonode support simplifies the LSDB topology graph. Specifying the `disable` keyword disables the capability. When disabled, three or RBridges on the same interface is represented as a full mesh topology in the LSDB.

The hardware data plane supports only one neighbor RBridge per port. If a second RBridge is detected on the same port, a warning message is logged. The control plane will continue to exchange control packets with all RBridges on the port, but only one RBridge per port will go to active state. The second RBridge may transition to active state if the peer session to the active RBridge terminates.

Example

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



configure trill system-id

```
configure trill system-id [switch-mac | system_id]
```

Description

This command configures the RBridge's TRILL System ID. If not specified, the RBridge automatically generates the TRILL ID by using internal system MAC address as the `system-id`.

Syntax Description

trill	Transparent Interconnection of Lots of Links.
system-id	System ID.
switch-mac	Use the System MAC address as the <code>system-id</code> (default).
<i>system_id</i>	System ID. A six octet value in the format of <code>xxxx.xxxx.xxxx</code> ; type=string; range=[14,14]

Default

System MAC address.



Usage Guidelines

This command configures the RBridge's TRILL System ID. The system-id is a six octet value. If not specified, the RBridge automatically generates the TRILL ID by using internal system MAC address as the system-id. An additional octet is appended to the end of the system ID to form the TRILL ID. The default TRILL ID is formed by appending 0x00 to the end of the system ID. If the appended octet is non-zero, then the TRILL ID represents a pseudonode.

The TRILL ID is used to identify the RBridge node in the Links State Database. When multiple RBridge nodes are attached to the same link, the RBridge dynamically creates a pseudonode by assigning a non-zero octet to the system ID. All of the TRILL neighbors on the shared link are represented in the LSDB as being connected to the RBridge's pseudonode. This simplifies the TRILL topology representation. Up to 254 pseudonodes are supported.

Specifying the keyword `switch-mac` instructs the RBridge to use the switch's universally administered system MAC address to generate the TRILL ID. This is the default setting for the system-id.

Example

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



configure trill timers csnp

```
configure trill timers csnp interval seconds
```

Description

This command sets the minimum time between consecutive complete sequence number packet (CSNP) transmissions on the specified interface.

Syntax Description

trill	Transparent Interconnection of Lots of Links.
timers	Timers.
csnp	Complete Sequence Number Packet.
interval	Minimum time between CSNP transmissions on the specified interface.
<i>seconds</i>	Interval in seconds (1-65535). (Default 10).



Default

N/A.

Usage Guidelines

Use this command to set the minimum time between consecutive complete sequence number packet (CSNP) transmissions on the specified interface. Periodic CSNPs are only sent on broadcast interfaces and only by the DRB. Point-to-point adjacencies also use CSNPs also, but they are not periodic and are unaffected by this command. The valid range for this parameter is 1-65535 seconds with a default value of 10 seconds.

Example

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



configure trill timers hello

```
configure trill timers hello {multiplier number} {interval [seconds | minimal]}
```

Description

This command configures the TRILL Hello protocol timers. The multiplier keyword specifies the hello multiplier. The hello multiplier is used in conjunction with the hello interval to compute the hold time.

Syntax Description

trill	Transparent Interconnection of Lots of Links.
timers	Timers.
hello	Hello protocol.
multiplier	Hello multiplier.
<i>number</i>	Hello multiplier value from 2 to 100. (Default 3).
interval	Interval between two consecutive Hello transmissions.
<i>seconds</i>	Interval in seconds (1-65535). (Default 10).
minimal	Set interval based on 1 sec hold time. Hold time = interval * multiplier.



Default

N/A.

Usage Guidelines

Use this command to configure the TRILL Hello protocol timers. The multiplier keyword specifies the hello multiplier. The hello multiplier is used in conjunction with the hello interval to compute the hold time. The hold time is included in Hello PDUs and is calculated by multiplying the hello multiplier by the hello transmit interval. The hold time tells the neighboring RBridge how long to wait before declaring the sending RBridge inactive. The valid range for this parameter is from two to 100 with a default value of 3. Interval specifies the interval between two consecutive Hello transmissions. The valid range for this parameter is from one to 65535 seconds with a default value of 10 seconds.

Example

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



configure trill timers lsp

```
configure trill timers lsp {generation-interval generation_seconds} {refresh-interval refresh_seconds} {lifetime lifetime_seconds} {transmit-interval milliseconds} {retransmit-interval retransmit_seconds} {checksum [enable | disable]}
```

Description

This command configures the TRILL Link State Protocol timers.

Syntax Description

trill	Transparent Interconnection of Lots of Links.
timers	Timers.
lsp	Link State Protocol.
generation-interval	Minimum time required to wait before regenerating the same LSP.
<i>generation_seconds</i>	Generation interval in seconds (1-120). (Default 30).



refresh-interval	Refresh rate for locally originated LSPs.
<i>refresh_seconds</i>	Refresh interval in seconds (1-65535). (Default 90).
lifetime	Remaining time for locally originated LSPs.
<i>lifetime_seconds</i>	Lifetime in seconds (350-65535). (Default 1200).
transmit-interval	Minimum time between LSP transmissions.
<i>milliseconds</i>	Transmit interval in milliseconds (1-4294967295). (Default 33).
retransmit-interval	Time to wait for an acknowledgement of the transmitted LSP on a point-to-point interface.
<i>retransmit_seconds</i>	Retransmit interval in seconds (1-65535). (Default 5).
checksum	Checksum calculation for Link State PDUs.
enable	Enable checksum.
disable	Disable checksum.

Default

N/A.

Usage Guidelines

Use this command to configure the TRILL Link State Protocol timers.

Example

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



configure trill timers spf

```
configure trill timers spf {restart restart_interval} {interval seconds}
```

Description

This command configures the TRILL SPF restart time and periodic calculation interval. The restart keyword specifies the restart timer for the RBridge LSP database.



Syntax Description

trill	Transparent Interconnection of Lots of Links.
timers	Timers.
spf	Shortest Path First.
restart	Restart timer for the RBridge LSP database.
<i>restart_interval</i>	Restart interval in seconds (5-65535). (Default 60).
interval	Minimum time to wait between SPF calculations.
<i>seconds</i>	Interval in seconds (1-120). (Default 10).

Default

N/A.

Usage Guidelines

Use this command to configure the TRILL SPF restart time and periodic calculation interval. The restart keyword specifies the restart timer for the RBridge LSP database. If timer expires before the LSDB has been resynchronized, the SPF is run. The valid range for this parameter is from five to 65535 seconds with a default value of 60 seconds.

This interval keyword specifies the minimum time to wait between SPF calculations. This helps prevent overloading the switch CPU in the event a link flap causes several back-to-back SPF calculations. The valid range for this parameter is from one to 120 seconds with a default value of 10 seconds.

Example

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



configure trill timers spf backoff-delay

```
configure trill timers spf backoff-delay {minimum minimum_delay} {maximum
maximum_delay}
```



Description

This command configures the TRILL SPF back-off timer delay. The backoff-delay keyword specifies the exponential back-off delays between SPF calculations.

Syntax Description

trill	Transparent Interconnection of Lots of Links.
timers	Timers.
spf	Shortest Path First.
backoff-delay	Exponential back-off delays between SPF calculations.
minimum	Minimum value for backoff-delay.
<i>minimum_delay</i>	Minimum backoff delay in milliseconds (1-2147483647). (Default 500).
maximum	Maximum value for backoff-delay.
<i>maximum_delay</i>	Maximum backoff delay in milliseconds (1-2147483647). (Default 50000).

Default

N/A.

Usage Guidelines

Use this command to configure the TRILL SPF back-off timer delay. The backoff-delay keyword specifies the exponential back-off delays between SPF calculations. The minimum keyword specifies the minimum duration between SPF calculations. The maximum keyword specifies the maximum duration between SPF calculations. The valid range for these parameters is from one to 2,147,483,647 milliseconds with a default minimum value of 500 milliseconds and a default maximum value of 50000 milliseconds.

Example

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



configure trill tree prune vlan



```
configure trill tree prune vlan [enable | disable]
```

Description

This command configures tree pruning for all trees in-use as specified by the use count. VLAN tags can be pruned from trees to reduce unnecessary traffic flooding.

Syntax Description

trill	Transparent Interconnection of Lots of Links.
tree	Distribution Tree.
prune	Prune unreachable destinations.
vlan	Prune unreachable VLAN destinations.
enable	Enable pruning (default).
disable	Disable pruning.

Default

Enabled.

Usage Guidelines

Use this command to configure tree pruning for all trees in-use as specified by the use count. VLAN tags can be pruned from trees to reduce unnecessary traffic flooding. VLANs cannot be pruned when R Bridges are configured to perform VLAN translation. This is indicated by the VLAN Mapping (VM) flag in the R Bridge TRILL Hellos. By default, tree pruning is disabled. When VLAN pruning is enabled, leaves of trees may be filtered for Access VLANs that are not configured at the edge of an R Bridge. Thus, flood traffic is only transported to remote R Bridges that have devices connected to the encapsulated packet's VLAN.

Example

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.





create trill nickname

```
create trill nickname nickname_id {nickname-priority id_priority } { root-priority root_priority } {name nickname_string }
```

Description

This command allocates a nickname for use by the local RBridge. The nickname is a 16-bit number that is unique within the TRILL network.

Syntax Description

trill	Transparent Interconnection of Lots of Links.
nickname	Nickname.
<i>nickname_id</i>	Identifier between 1 and 0xFFBF in hex";type="hex_t";range="[1,65471]
nickname-priority	Nickname priority
<i>id_priority</i>	Priority value between 128 and 255. Lower numbers represent lower priority. Default is 192.
root-priority	Root priority
<i>root_priority</i>	Priority value between 0 and 65535. Lower numbers represent lower priority. Default is 32768.
name	Human readable name associated with nickname.
<i>nickname_string</i>	Name string up to 32 characters";type="string";range="[1,32]

Default

N/A.

Usage Guidelines

Use this command to allocate a nickname for use by the local RBridge. The nickname is a 16-bit number that is unique within the TRILL network. Each nickname identifies a distribution tree rooted at the local RBridge and is used to identify an RBridge for the purpose of learning the unicast MAC address to RBridge mapping. All of the nicknames are advertised to the other RBridges in the TRILL network in the Nickname sub-TLV as part of the Group Address TLV. The optional nickname *nickname_string* is locally significant and allows the network administrator to reference the nickname by an easily remembered descriptive string. The *nickname_string* parameter has a maximum length of 32 octets and must start with a character. If the nickname's string name is not specified, the show output commands will indicate this by displaying the nickname value prefixed with the string "noname_" as the string name.

Each RBridge must have a minimum of one nickname. When TRILL is enabled for the first time, a default nickname is automatically created and the nickname value is randomly generated with the default priority of 0x40. The default nickname string name is set to Default_nickname. The default nickname cannot be deleted or modified and the parameters for the Default_nickname are not configurable. Up to four additional nicknames may be created. The default nickname is only used when



the TRILL Hello protocol is enabled. If the TRILL Hello protocol is not enabled, the `nickname_id` must be specified and must be unique within the TRILL network. The RBridge default nickname is not used when the TRILL Hello protocol is disabled.

Example

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



delete trill nickname

```
delete trill nickname [nickname_id | nickname_string | all]
```

Description

This command deletes a locally created nickname. Deleting a nickname removes the nickname from the TRILL network.

Syntax Description

trill	Transparent Interconnection of Lots of Links.
nickname	Nickname.
<i>nickname_id</i>	Identifier between 1 and 0xFFBF in hex.
<i>nickname_string</i>	Name string up to 32 characters.
all	All nicknames.

Default

N/A.

Usage Guidelines

Use this command to delete a locally created nickname. Deleting a nickname removes the nickname from the TRILL network. This may cause a temporary network traffic disruption. The distribution tree associated with the deleted nickname cannot be used by any other RBridge to forward multicast or flood traffic. If the nickname is associated with a computed in-use distribution tree, another nickname



(associated with a tree that is possibly rooted at another RBridge) must be chosen, if one is available. RBridges must flush MAC address FDB entries that are associated with the deleted nickname. The default nickname cannot be deleted. Specifying the **all** keyword deletes all user-created nicknames on the local RBridge, except for the default nickname.

Example

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



disable trill

disable trill

Description

This command disables the TRILL protocol on the switch. When disabled, the RBridge does not send TRILL Hellos, calculate forwarding trees, and exchange LSPs and does clear the Link State Database and terminates all TRILL packet forwarding.

Syntax Description

trill	Transparent Interconnection of Lots of Links.
--------------	---

Default

N/A.

Usage Guidelines

Use this command to disable the TRILL protocol on the switch. When disabled, the RBridge does not send TRILL Hellos, calculate forwarding trees, and exchange LSPs, but does clear the Link State Database and terminates all TRILL packet forwarding. The *All-RBridges* and *All-ISIS-RBridges* functional MAC addresses are removed from the local delivery filter table. Packets received with the RBridge functional MAC addresses are treated as generic broadcast packets. The disable command will not reset any TRILL configuration. Use the **unconfigure trill** command to clear all related TRILL configuration.



Example

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



enable trill

enable trill

Description

This command enables TRILL data plane forwarding on the switch. Once enabled, the switch is capable of operating as an RBridge.

Syntax Description

trill	Transparent Interconnection of Lots of Links.
--------------	---

Default

N/A.

Usage Guidelines

Use this command to enable TRILL data plane forwarding on the switch. Once enabled, the switch is capable of operating as an RBridge. By default, The default VLAN, VLAN ID of one, is automatically configured to support TRILL forwarding. The default access VLAN tag space defaults to one through 4094. When TRILL is enabled, the *All-RBridges* and *All-ISIS-RBridges* functional MAC addresses are installed and any packets sent to those addresses will be received and processed by the switch. TRILL data packets are only transmitted on TRILL configured VLANs. TRILL Hello packets are transmitted on all ports that are members of a TRILL configured VLANs and any VLAN that has a VLAN ID that matches a VLAN ID in the TRILL Access VLAN tag space.

Example



History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



show trill

```
show trill {detail}
```

Description

This command displays general configuration information related to the TRILL. This information includes enabled status, nicknames, priorities, configured number of Dtrees to compute, support, and use, timer configurations, System ID, Designated VLAN, Access VLAN IDs, MTU Size, MTU Probe status, and Tree Pruning status.

Syntax Description

show	Show information.
trill	Transparent Interconnection of Lots of Links.
detail	Show detailed TRILL information.

Default

N/A.

Usage Guidelines

Use this command to display general configuration information related to the TRILL. This information includes enabled status, nicknames, priorities, configured number of Dtrees to compute, support, and use, timer configurations, System ID, Designated VLAN, Access VLAN IDs, MTU Size, MTU Probe status, and Tree Pruning status.

Example

```
#show trill
TRILL           : Enabled      System ID       : 0004.966D.5D9D (Switch
MAC)
Hello Interval  : 10 sec      Inhibition timer : 30 sec
MTU Probe       : Disabled
Configured MTU  : 1470      Negotiated MTU   : 1470

TRILL Maintenance Mode: Disabled
```



```

Nickname:
0x0008 (Default_nickname)   Nickname-Priority: 64   Root-Priority:
32768
Desired Designated VLAN : Default
Network VLANs           : Default
Access Tags              :

#show trill detail
TRILL:                   Enabled
TRILL Maintenance Mode:  Disabled

Configured Nickname
* 0x0002 ()              Nickname-Priority: 19   Root-Priority: 32768
Generated Nickname
0x00cd (Default_nickname) Nickname-Priority: 64   Root-Priority:
32768
(*) Active

System ID:                0004.966D.6F4E (Switch MAC)
Desired Designated VLAN: Default
DTrees To Compute:       1
DTrees To Support:       1
DTrees To Use:           1
Bypass Pseudonode:       Enabled
Inhibit Time:            10 sec
RPF Check:                Enabled
maintenance Window:      Enabled

VLANs
Network:                  Default
Access VLAN IDs:          100

MTU
Configured MTU Size:      1470 octets
Negotiated MTU Size:      1470 octets
Probe:                    Enabled
Probe Fail Count:         3

Tree Pruning
VLANs:                    Enabled

TRILL Hello Timer
Hello Protocol:           Enabled
Multiplier:               3
Interval:                  10 sec

TRILL LSP Timer
Generation Interval:      30 sec
Refresh Interval:         90 sec
Lifetime:                  1200 sec
Transmit Interval:        33 sec
Retransmit Interval:      5 sec
Checksum:                  Enabled

TRILL SPF Timer
Minimum Backoff Delay:    500 ms
Maximum Backoff Delay:    50000 ms

```




```
TRILL CSNP Timer
Interval:          10 sec
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



show trill distribution-tree

```
show trill distribution-tree { [ pruning [vlan] | rpf] {dtree_nickname} }
```

Description

This command displays the egress RBridge forwarding database for the TRILL network. The tabular display output includes the egress RBridge nickname, System ID, hop count, next hop nickname, next hop System ID, and port number associated with the next hop RBridge.

Syntax Description

show	Show information.
trill	Transparent Interconnection of Lots of Links.
distribution-tree	TRILL Campus wide Distribution Tree.
pruning	Show pruning information.
vlan	Show pruning information related to VLAN destinations.
rpf	Show TRILL Reverse Path Forwarding table.
<i>dtree_nickname</i>	Nickname of dtree (in hex value);type="hex_t";range="[1,65471].

Default

N/A.

Usage Guidelines

Use this command to display the egress RBridge forwarding database for the TRILL network. The tabular display output includes the egress RBridge nickname, System ID, hop count, next hop nickname, next hop System ID, and port number associated with the next hop RBridge.



Example

```
show trill distribution-tree
```

```
D-Tree D-Tree      Hop  Adj.  Adjacency
Name   System ID    Count Name   System ID   Port
-----
0x0002 1111.1111.0001    1 0x0001 1111.1111.0000  1
                                     0x0003 0004.966D.5D85  2
-----
```

Hop Count: Number of D-Tree hops to the furthest node from local RBridge

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



show trill lsdb

```
show trill lsdb {lsp-id lsp-id | detail}
```

Description

This command displays the links state database associated with TRILL network. The display is essentially the same as the output displayed by the show isis lsdb command.

Syntax Description

show	Show information.
trill	Transparent Interconnection of Lots of Links.
lsdb	TRILL Campus wide Link State Database.
lsp-id	TRILL Link State PDU.
<i>lsp_id</i>	TRILL Link State PDU ID, In the format of xxxx.xxxx.xxxx.xx-xx";type="string";range="[1,21].
detail	Detailed information.

Default

N/A.



Usage Guidelines

Use this command to display the links state database associated with TRILL network. The display is essentially the same as the output displayed by the `show isis lsdb` command.

Example

```
#show trill lsdb
```

```
TRILL Link State Database
  LSP ID                Seq Num    Checksum    Hold Time    OL Flag
-----
  0004.966D.5D85.00-00   0x00000002  0x0f69      1164         0
* 0004.966D.6F4E.00-00   0x00000002  0x07c7      1165         0
  0204.966D.6528.00-00   0x00000002  0x47aa      1165         0
-----
(*) - Self
OL  - Overload
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



show trill neighbor

```
show trill neighbor {nickname nickname | system-id system | mac-address mac-address }
```

Description

This command displays network information related to TRILL RBridge neighbors. By default, the tabular display output includes the neighbor RBridge MAC Address, System ID, Nickname, connected Port, RBridge priority, and operational state.

Syntax Description

show	Show information.
trill	Transparent Interconnection of Lots of Links.
neighbor	TRILL Campus wide neighbors.
nickname	Show TRILL neighbor associated with a particular nickname.
<i>nickname</i>	TRILL neighbor nickname between 1 and 0xFFBF in hex"; type="hex_t";range="[1,65471.]
system-id	Show TRILL neighbor associated with a particular system id.



<i>system</i>	TRILL neighbor system id. In the format of xxxx.xxxx.xxxx";type="string";range="[1,14].
mac-address	Show TRILL neighbor associated with a particular MAC address.
<i>mac-address</i>	TRILL neighbor MAC address. In the format of xx:xx:xx:xx:xx:xx";type="mac_t.

Default

N/A.

Usage Guidelines

Use this command to display network information related to TRILL RBridge neighbors. By default, the tabular display output includes the neighbor RBridge MAC Address, System ID, Nickname, connected Port, RBridge priority, and operational state. The RBridge operational states are Detect, 2Way, Report, and DR. The display out can be filtered by optionally specifying the keyword nickname, system-id, or mac-address and the appropriate associated parameter. Specifying one of the optional filter keywords displays additional information related to the specified object in detail format.

Example

```
#show trill neighbor
Neighbor
Mac                System ID      Nickname Port  Hold DRB
Time Priority State
-----
00:04:96:6d:5d:85 0004.966D.5D85 0x00c5  1:3   21      64 REPORT
02:04:96:6d:65:28 0204.966D.6528 0x941e  1:5    8      64 REPORT/DR
-----
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



show trill ports

```
show trill ports port_list {counters {no-refresh | detail}}
```

Description

This command displays operational TRILL information associated with the switch ports. By default, the tabular display output includes the port number, port state, link state, DRB election priority, TRILL port



type, TRILL link type, link metric, and associated VLAN. Displayed port states include *Disabled*, *Enabled*, *Forwarding*, and *Inhibited*.

Syntax Description

show	Show information.
trill	Transparent Interconnection of Lots of Links.
ports	TRILL Ports.
<i>port_list</i>	Port list separated by a comma or (-).
counters	TRILL data packet counters.
no-refresh	Page by page display without auto-refresh.
detail	Detailed information.

Default

N/A.

Usage Guidelines

This command displays operational TRILL information associated with the switch ports. By default, the tabular display output includes the port number, port state, link state, DRB election priority, TRILL port type, TRILL link type, link metric, and associated VLAN. Displayed port states include Disabled, Enabled, Forwarding, and Inhibited. The link states include Active, Ready, Port not Present, and Loopback. The link types are Broadcast and P2P. The port types are Access, Trunk, and Hybrid. Optionally specifying the *port_number* displays detailed TRILL information pertaining to the specified port. Optionally specifying the keyword *counters* displays incoming and outgoing TRILL packet counts per port. Counters associated with RPF check failures, hop count exceeded, and TRILL option drops are maintained globally and are displayed using the `show trill counters` command.

Example

```
#show trill ports
```

```

      Port  Link  DRB Election Port      Link
Port State State Priority      Type      Type      Metric      Designated
-----
1:3  E F  A                64 Trunk  Broadcast  20000 * Default
1:5  E F  A                64 Trunk  Broadcast  20000 * Default
-----

```

```
Port State: D-Disabled, E-Enabled
```

```
           F-Forwarding, S-Suspended, d-Protocol Disabled
```

```
Link State: A-Active, R-Ready, NP-Port not present, L-Loopback
```

```
Metric:      (*)-Auto Metric
```

```
(!)-TRILL maintenance mode enabled
```



History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



show trill rbridges

show trill rbridges

Description

This command displays status information related to RBridges in the TRILL network.

Syntax Description

show	Show information.
trill	Transparent Interconnection of Lots of Links
rbridges	TRILL campus-wide RBridges

Default

N/A.

Usage Guidelines

Use this command to display status information related to RBridges in the TRILL network. The tabular display output includes the RBridge's nickname, system ID, DRB priority, root tree priority, and status. RBridges highlighted with an asterisk indicate RBridge information associated with the local switch.

Example

```
#show trill rbridges
Nickname System ID      Nickname Root      Status      Next Hop Port  Metric      Hop
Count
-----
# 0x0081 1111.1111.0000      64 32768  NotActive
  0x5ac5 1111.1111.0001      64 32768  Active      0x5ac5 3      1
1
% 0x8000 1111.1111.0000      193 32768  Active
-----
-
(#)-Automatically created default nickname for local RBridge
(%)-User created nickname for local RBridge
```



(*)-Static nexthop
 (!)-TRILL maintenance mode enabled

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



traceroute trill

Short reference description.

```
traceroute trill rbridge_nickname
```

Description

This command traces the path TRILL takes through the network.

Syntax Description

trill	Transparent Interconnection of Lots of Links.
<i>rbridge_nickname</i>	TRILL campus-wide RBridge nickname in hex value; type="hex_t";range="[1,65471]"

Default

N/A.

Usage Guidelines

Use this command to trace the path TRILL takes through the network.

Example

```
traceroute trill 0x3e
traceroute to 0x003e, 64 hops max

Hop      Nickname   Time                Nexthop
         Nickname
1        0x001d     60 msec            0x003e2
         0x003e     358 msec
traceroute to 0x003e completed
```



History

This command was first available in ExtremeXOS 15.4.

Platform Availability

TRILL is supported on the BD-X series and Summit X670 and X770 series of switches.



20 Commands for Virtual Routers

```
clear counters vr
configure vr add ports
configure vr add protocol
configure vr delete ports
configure vr delete protocol
configure vr description
configure vr rd
configure vr route-target
configure vr vpn-id
create virtual-router
delete virtual-router
disable snmp trap l3vpn
disable virtual-router
enable snmp trap l3vpn
enable virtual-router
show counters vr
show virtual-router
unconfigure vr description
unconfigure vr rd
unconfigure vr vpn-id
virtual-router
```

This chapter describes commands for:

- Creating and deleting a virtual router (VR) or virtual router forwarding instance (VRF).
- Configuring and managing VRs and VRFs.
- Displaying information about VRs and VRFs.



Note

In this chapter, "VR" refers to all types of VR (User VR, VRF, Default VR, etc.). If there is a distinction, it is clarified in the command.

For an introduction to VRs and VRFs, see the ExtremeXOS Concepts Guide.

clear counters vr

```
clear counters {vr} vpn-vrf-name
```

Description

Clears statistics information for a VPN Virtual Routing and Forwarding instance (VPN VRF).

Syntax Description

<i>vpn-vrf-name</i>	Specifies the name of a VPN VRF.
---------------------	----------------------------------

Default

N/A.

Usage Guidelines

This command can help to debug control path issues for a VPN VRF. Issuing a global XOS “clear counter” command will also clear VRF counters. This command clears the following counters:

- Route add operation count.
- Route delete operation count.
- Routes dropped count.

This command is supported only on VPN VRFs.

Example

The following command clears the counters for VPN VRF red:

```
Switch.19 # clear counters vr red
```

History

This command was first introduced in XOS Release 15.3.

Platform Availability

This command is available on BlackDiamond8000 c-, xl-, and xm-series modules, and Summit X460, X480, X670, and X770 switches.

configure vr add ports

```
configure vr vr-name add ports port_list
```

Description

Assigns a list of ports to the VR specified.



Syntax Description

<i>vr_name</i>	Specifies the name of the VR.
<i>port_list</i>	Specifies the ports to add to the VR.

Default

By default, all ports are assigned to the VR, VR-Default.

Usage Guidelines

When a new VR is created, by default, no ports are assigned, no VLAN interface is created, and no support for any routing protocols is added. Use this command to assign ports to a VR. Since all ports are initially assigned to VR-Default, you might need to delete the desired ports first from the VR where they reside, before you add them to the desired VR.

If you plan to assign VR ports to a VLAN, be aware that the ports that you add to a VLAN and the VLAN itself cannot be explicitly assigned to different VRs. When multiple VRs are defined, consider the following guidelines while adding ports to a VR:

- A VLAN can belong (either through explicit or implicit assignment) to only one VR.
- If a VLAN is not explicitly assigned to a VR, then the ports added to the VLAN must be explicitly assigned to a single VR.
- If a VLAN is explicitly assigned to a VR, then the ports added to the VLAN must be explicitly assigned to the same VR or to no VR.
- If a port is added to VLANs that are explicitly assigned to different VRs, the port must be explicitly assigned to no VR.

Example

The following command adds all the ports on slot 2 to the VR vr-acme:

```
configure vr vr-acme add ports 2:*
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond8000 c-, xl-, and xm-series modules, and Summit X460, X480, X670, and X770 switches.

configure vr add protocol

```
configure vr vr_name [add | delete] protocol [ospf | ospf3 | rip | ripng | bgp | isis | pim]
```



Description

Starts a Layer 3 protocol instance for a VR or VRF.

Syntax Description

<i>vr_name</i>	Specifies the name of a VR or a VRF.
protocol	Specifies a Layer 3 protocol that you can add or delete.
name	Specifies the name of a VR or a VRF. The following protocols are supported on VRs: RIP, RIPng, OSPF, OSPFv3,BGP,PIM. IS-IS, and MPLS. The following protocols are supported on VRFs: BGP.
add	Adds a routing protocol to VRF for PE – CE communication .
delete	Specifies the name of a VR or a VRF.

Default

By default, none of the dynamic protocols are added to a User VR or a VRF.

Usage Guidelines

When a new VR or VRF is created, by default, no ports are assigned, no VLAN interface is created, and no support for any routing protocols is added.

MPLS is the only protocol that you can add to or delete from the VR-Default VR. When MPLS is enabled on a switch, the default configuration adds MPLS to VR-Default. You cannot add or delete any other protocols from VR-Default, and you cannot add or delete any protocols from the other system VRs, VR-Mgmt and VR-Control.



Note

You must delete the MPLS protocol from VR-Default before you can add it to a user VR. MPLS can be active on only one VR within a switch.

When you add a protocol to a VRF, the parent VR starts that protocol, if it was not already running, and adds a protocol instance to support the VRF.

If a previously configured protocol instance is deleted, the CE routes imported from that protocol into the VRF RIB is removed.

Example

The following command starts RIP on the VR vr-acme:

```
configure vr vr-acme add protocol rip
```



The following command starts a BGP protocol instance for VRF vr-widget:

```
configure vr vr-widget add protocol bgp
```

History

This command was first available in ExtremeXOS 11.0.

MPLS protocol support was added in ExtremeXOS 12.4.

Support for the OSPFv3 and RIPng protocols on user VRs was added in ExtremeXOS 12.5.

Support for the BGP protocol on VRFs was added in ExtremeXOS 12.6.0-BGP.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond8000 c-, xl-, and xm-series modules, and Summit X460, X480, X670, and X770 switches.

configure vr delete ports

```
configure vr vr-name delete ports port_list
```

Description

Removes a list of ports from the VR specified.

Syntax Description

<i>vr-name</i>	Specifies the name of the VR.
<i>port_list</i>	Specifies the ports to remove from the VR.

Default

By default, all ports are assigned to the VR, VR-Default.

Usage Guidelines

When a new VR is created, by default, no ports are assigned, no VLAN interface is created, and no support for any routing protocols is added. Use this command to remove ports from a VR. Since all ports are initially assigned to VR-Default, you might need to delete the desired ports first from the VR where they reside, before you add them to the desired VR.



Example

The following command removes all the ports on slot 2 from the VR `vr-acme`:

```
configure vr vr-acme delete ports 2:*
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond8000 c-, xl-, and xm-series modules, and Summit X460, X480, X670, and X770 switches.

configure vr delete protocol

```
configure vr vr-name delete protocol protocol-name
```

Description

Stops and removes a Layer 3 protocol instance for a VR or VRF.

Syntax Description

<i>vr-name</i>	Specifies the name of the VR.
<i>protocol-name</i>	Specifies the Layer 3 protocol. The following protocols are supported on VRs: RIP, RIPng, OSPF, OSPFv3, BGP, PIM. IS-IS, and MPLS. The following protocols are supported on VRFs: BGP.

Default

N/A.

Usage Guidelines

MPLS is the only protocol that you can add to or delete from the VR-Default VR. When MPLS is enabled on a switch, the default configuration adds MPLS to VR-Default. You cannot add or delete any other protocols from VR-Default, and you cannot add or delete any protocols from the other system VRs, VR-Mgmt and VR-Control.



Note

You must delete the MPLS protocol from VR-Default before you can add it to a user VR. MPLS can be active on only on VR within a switch.



When you delete a protocol from a VRF, the protocol instance is deleted on the parent VR and the CE routes imported from that protocol into the VRF Routing Information Base (RIB) are removed. The parent VR continues to run the protocol until that protocol is removed from the VR.

Example

The following command shuts down and removes RIP from the VR `vr-acme`:

```
configure vr vr-acme delete protocol rip
```

The following command deletes the BGP protocol instance for VRF `vr-widget`:

```
configure vr vr-widget delete protocol bgp
```

History

This command was first available in ExtremeXOS 11.0.

MPLS protocol support was added in ExtremeXOS 12.4.

Support for the OSPFv3 and RIPng protocols on user VRs was added in ExtremeXOS 12.5.

Support for the BGP protocol on VRFs was added in ExtremeXOS 12.6.0-BGP.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond8000 c-, xl-, and xm-series modules, and Summit X460, X480, X670, and X770 switches.

configure vr description

```
configure vr vr_name description string
```

Description

Use this command to configure a description for the specified VR or VRF.

Syntax Description

<i>vr_vrf_name</i>	Specifies the name of a user VR or a VRF.
<i>string</i>	Specifies a text string to describe the VR. If the text string contains space characters, the entire string must be enclosed with double quote characters (" ").



Default

No description.

Usage Guidelines

This command allows you to add comments about a VRF/VR entity. Entering a NULL string on the CLI will unconfigure the description string for the VRF/VR. If the description string has spaces in it, then the string must be enclosed with in double quotes (" ").

This text message appears in the `show virtual-router` command display when the command specifies a VR name. For VPN VRFs, this message is returned for a mplsL3VPN MIB query of the MIB variable mplsL3VpnVrfDescription.

Example

The following command configures a description for the VRF named *corporate*:

```
configure vr corporate description "VRF for the corporate intranet"
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond8000 c-, xl-, and xm-series modules, and Summit X460, X480, X670, and X770 switches.

configure vr rd

```
configure vr vrf_name rd [2_byte_as_num:4_byte_number | ip_address:2_byte_number | 4_byte_as_num:2_byte_number]
```

Description

Use this command to configure a route-distinguisher (RD) for a VPN VRF.

Syntax Description

<i>vrf_name</i>	Specifies the name of a VPN VRF.
<i>rd</i>	Specifies a Route distinguisher for a VRF. It can be either ASN-related, where it is represented as <i>2-byte AS number:4-byte num</i> . It can be IP-address based where it is represented as <i>4-byte IP address:2-byte number</i> .
<i>2_byte_as_num</i>	Specifies a 2-byte Autonomous System (AS) number.



<i>4_byte_number</i>	Specifies a 4-byte number to further identify the RD. This number can be chosen by organization that configures the RD, and this number does not need to match any other network configuration parameters.
<i>ip_address</i>	Specifies an IP address to include as part of the RD.
<i>4_byte_as_num</i>	Specifies a 4-byte AS number.
<i>2_byte_number</i>	Specifies a 2-byte number to further identify the RD. This number can be chosen by the organization that configures the RD, and this number does not need to match any other network configuration parameters.

Default

N/A.

Usage Guidelines

The RD can be specified in the following formats:

- *2_byte_as_num:4_byte_number*

Here is an example: 9643:7000

- *ip_address:2_byte_number*

Here is an example: 10.203.134.5:324

- *4_byte_as_num:2_byte_number*

Note

Although route distinguisher is 8-bytes wide, this CLI accepts only 6 bytes value. The first two bytes ("type" field) is deduced from the values entered on the CLI, so it is redundant to set the type field on the CLI too. If *2-byte as-num:4-byte numis* entered on the CLI, the type field is automatically set to 0. If *ip_address:2_byte_number* is entered on the CLI, the type field is automatically set to 1. Type 2 (4-bytes AS number) Route Distinguisher is not supported.



Route distinguisher is a mandatory parameter for a VRF. Without this parameter, a VRF cannot be active and local VPN routes cannot be advertised across the SP's backbone to the remote VPN sites.

Use this command to configure or change the RD for a VPN VRF. If you use this command to change the RD, the Layer 3 VPN associated with that VPN VRF is reset by automatically disabling and re-enabling the VRF.

RD is added to the beginning of the VPN customer's IPv4 prefix to make globally unique VPNv4 prefixes. You must configure RD for a VRF to be functional. This command is not applicable for a heavy-weight traditional VR.



Example

The following examples configure RDs using the two of the supported formats:

```
configure vr corporate-extreme rd 10.203.134.5:324

configure vr corporate-guest rd 9643:7000
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on BlackDiamond8900 xl- and xm-series modules, and Summit X460, X480, X670, and X770 switches.

configure vr route-target

```
configure vr vrf_name route-target [import | export | both] [add | delete]
[route_target_extended_community]
```

Description

Use this command to add or delete entries in the import and export lists for route target extended communities for a specified VPN VRF.

Syntax Description

<i>vrf_name</i>	Specifies the name of a VPN VRF.
import	Specifies that the specified route target extended community is to be added to or deleted from the import list for the VRF.
export	Specifies that the specified route target extended community is to be added to or deleted from the export list for the VRF.
both	Specifies that the specified route target extended community is to be added to or deleted from both the import and export lists for the VRF.
add	Specifies that the specified route target extended community is to be added to the specified list for the VRF.
delete	Specifies that the specified route target extended community is to be deleted from the specified list for the VRF.
<i>route_target_extended_community</i>	Specifies the route target extended community. It can be represented in two formats. ASN-related or IP-Address-related.



Default

No default route targets. If you do not specify the **import** or **export** options at the CLI, by default **both** is assumed.

Usage Guidelines

This command creates lists of import and export route target extended communities for the specified VRF. Route Target attributes are used to control the VPNv4 route distribution by BGP. Learned routes (from the PE) that carry a specific route target extended community are imported into all VRFs configured with that extended community as an import target. Routes learned from a VRF site are labeled with export route target extended communities configured for that VRF. This is used to control the VRFs into which the route is imported.

A route target extended community can be specified in the following formats:

- *2_byte_as_num:4_byte_number*
- *ip_address:2_byte_number*
- *4_byte_as_num:2_byte_number*

To configure multiple route target extended communities in import or export lists, execute this command with **add** option multiple times, once for each extended community. Use the **delete** option to remove an extended community from an import or export list. You cannot use this command for a heavy-weight traditional VR.

Example

The following examples configure route target extended communities using the two supported formats:

```
configure vr corporate-extreme route-target both add 172.16.186.230:9823
configure vr corporate-guest route-target both add 9643:7002
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on BlackDiamond8900 xl- and xm-series modules, and Summit X460, X480, X670, and X770 switches.

configure vr vpn-id

```
configure vr vrf_name vpn-id 3_byte_oui:4_vpn_index
```

Description

This command configures a globally unique identifier for a VPN VRF.



Syntax Description

<i>vrf_name</i>	Specifies the name of a VPN VRF.
<i>3_byte_oui</i>	Specifies an organizationally unique identifier (OUI). The IEEE organization assigns this identifier to companies. The OUI is restricted to three bytes and must be entered in hexadecimal format.
<i>4_vpn_index</i>	Identifies the VPN within a company. This VPN index is restricted to 4 bytes and must be entered in hexadecimal format.

Default

VPN ID is not configured for VRFs.

Usage Guidelines

The VPN ID uniquely identifies a VPN. This command is only applicable for a VPN VRF. Each VRF configured in a PE router can have a VPN ID. Use the same VPN ID for the VRFs on other PE routers that belong to the same VPN. Ensure that the VPN ID is unique for each VPN in the Service Provider network.

The *oui* and *vpn_index* parameters must be entered on the CLI in hex format.

Example

The following example assigns VPN ID ac:9f3c8 to a VRF named *corporate-extreme*:

```
configure vr corporate-extreme vpn-id ac:9f3c8
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on BlackDiamond8900 xl- and xm-series modules and Summit X460, X480, X670, and X770 switches.

create virtual-router

```
create virtual-router name {type [vrf | vpn-vrf {vr parent_vr_name}]}
```

Description

Use the create virtual-router command to create a user VR or VRF.



Syntax Description

type	Specifies the type of virtual router you are creating.
vrf	Specifies that you are creating a new L3 or IP routing domain.
vpn-vrf	Specifies that you are creating a new L3 or IP routing domain that supports L3VPNs.
parent_vr_name	Specifies the parent VR that supports the VRF you are creating.

Default

If no **type** is specified, then the default is to create a user virtual router. A virtual router creates separate L3 Routing Domains.

If *parent_vr_name* parameter is not specified, the VRF will be created under the VR of the current CLI context. Default is VR-Default.

Usage Guidelines

All VRFs are created under default VR or a user created VR. VPN-VRFs can be created in any VR but for L3VPNs to work VPN-VRFs should be created under a parent VR where MPLS is configured. There is a single namespace maintained by XOS configuration manager and it contains VRs and VRFs. Hence name for a VR or a VRF must be unique in EXOS.

A VR or VRF name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 31 characters. The name must be unique among the object names on the switch, and the name is case insensitive. For information on VR and VRF name guidelines and a list of reserved names, see [Object Names](#) in the *ExtremeXOS Concepts Guide*.

When a new VR is created, by default, no ports are assigned, no VLAN interface is created, and no support for any routing protocols is added. A protocol process is started in the parent VR when a protocol instance is added to a VRF. If you do not specify a VR type, this command creates a user VR.

VRFs are supported as children of user VRs or VR-Default. If a *parent_vr_name* is specified when a VRF is created, the new VRF is created under that parent, provided that the parents supports VRFs. If no parent is specified, the VRF is assigned to the VR for the current VR context, or to VR-Default if the current VR context does not support VRFs.



Note

To support Layer 3 VPNs, a VPN VRF must be created under the VR that supports MPLS. The software supports MPLS on only one VR.

Example

The following command creates the VR vr-acme:

```
create virtual-router vr-acme
```



The following command creates the non-VPN VRF vrf1:

```
create virtual-router vrf1 type vrf
```

History

This command was first available in ExtremeXOS 11.0.

Support for non-VPNVRFs was added in ExtremeXOS 12.5.

Support for VPN VRFs was added in ExtremeXOS 12.6.0-BGP.

Support for L3 VPN VRFs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond8000 c-, xl-, and xm-series modules, and Summit X460, X480, X670, and X770 switches.

delete virtual-router

```
delete virtual-router vr-name
```

Description

This command deletes a VR or VRF.

Syntax Description

<i>vr-name</i>	Specifies the name of the VR or VRF.
----------------	--------------------------------------

Default

N/A.

Usage Guidelines

Only user VRs and VRFs can be deleted.

Before you delete a user VR, you must delete all VLANs and protocols assigned to the VR, and you must delete any child VRFs. All of the ports assigned to a deleted VR are made available to assign to other VRs.

Before you delete a VRF, you must delete all VLANs and stop all protocols that are assigned to that VRF. All of the ports assigned to a deleted VRF are deleted and made available to assign to other VRs and VRFs. Any routing protocol instance that is assigned to the VRF is deleted gracefully.



Example

The following command deletes the VR vr-acme:

```
delete virtual-router vr-acme
```

The following command deletes the VRF vrf1:

```
delete virtual-router vrf1
```

History

This command was first available in ExtremeXOS 11.0.

Support for non-VPN VRFs was added in ExtremeXOS 12.5.

Support for VPN VRFs was added in ExtremeXOS 12.6.0-BGP.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond8000 c-, xl-, and xm-series modules, and Summit X460, X480, X670, and X770 switches.

disable snmp trap l3vpn

```
disable snmp trap l3vpn {vr name}
```

Description

This command disables Layer 3 VPN MIB notification traps for the child VPN VRFs of the specified VR.

Syntax Description

<i>vr-name</i>	Specifies the name of the parent VR where this RFC 4382 scalar is applied. If <i>vr-name</i> is not provided, then this command is applied to the VR in the current context.
----------------	--

Default

Disabled.

Usage Guidelines

None.



Example

The following example disables SNMP traps for Layer 3 VPNs on the default VR:

```
disable snmp traps l3vpn vr vr-default
```

History

This command was first available in ExtremeXOS 12.6.0-BGP.

Platform Availability

This command is available on BlackDiamond8900 xl- and xm-series modules and Summit X460, X480, X670, and X770 switches.

disable virtual-router

```
disable virtual-router vrf-name
```

Description

Disables a VRF.



Note

This command is only applicable for VRFs.

Syntax Description

<i>vrf-name</i>	Specifies the name of the VRF.
-----------------	--------------------------------

Default

Enabled.

Usage Guidelines

When you disable a VRF, the software does the following:

- Disables Layer 3 protocols.
- Marks static routes as inactive and removes them from the hardware forwarding tables.
- Flushes the IP ARP and IPv6 neighbor-discovery caches.



Example

The following command disables VRF vrf1:

```
disable virtual-router vrf1
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond8000 c-, xl-, and xm-series modules, and Summit X460, X480, X670, and X770 switches.

enable snmp trap l3vpn

```
enable snmp trap l3vpn {vr_name}
```

Description

This command enables Layer 3 VPN MIB notification traps for the child VPN VRFs of the specified VR.

Syntax Description

<i>vr_name</i>	Specifies the name of the parent VR where this RFC 4382 scalar is applied. If <i>vr_name</i> is not provided, then this command is applied to the VR in the current context.
----------------	--

Default

Disabled.

Usage Guidelines

This command enables generation of the following Layer 3 VPN SNMP traps:

- mplsL3VpnVrfUp—Sent when the first IP VLAN becomes active and the administrative state is enabled.
- mplsL3VpnVrfDown—Sent when the last active IP VLAN becomes inactive, or the administrative state is disabled.



Example

The following example enables SNMP traps for Layer 3 VPNs on the default VR:

```
enable snmp traps l3vpn vr vr-default
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on BlackDiamond8900 xl- and xm-series modules and Summit X460, X480, X670, and X770 switches.

enable virtual-router

```
enable virtual-router vrf-name
```

Description

Enables a VRF.



Note

This command does not affect virtual routers.

Syntax Description

<i>vrf-name</i>	Specifies the name of the VR or VRF instance.
-----------------	---

Default

Enabled.

Usage Guidelines

This command is used to administratively enable or disable a VRF. The VRF specific commands are still accepted and retained by the switch. This configuration has an operational impact on the VRF.

When you enable a VRF, the software does the following:

- Enables Layer 3 protocols for the VRF.
- Marks static routes as active and adds them to the hardware forwarding tables.



Example

The following command enables VRF vrf1:

```
enable virtual-router vrf1
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond8000 c-, xl-, and xm-series modules, and Summit X460, X480, X670, and X770 switches.

show counters vr

```
show counters vr {vpn-vrf-name}
```

Description

Displays statistics information about VPN VRF operation.

Syntax Description

<i>vpn-vrf-name</i>	Specifies the name of a VPN VRF.
---------------------	----------------------------------

Default

N/A.

Usage Guidelines

This command displays counters that show:

- The total number of IP unicast and multicast routes.
- Route add operation count.
- Route delete operation count.
- Routes dropped count.

Note



The total route count displayed for this command can exceed the total route count displayed by the `show iproute` command because the `show iproute` command displays either unicast or multicast routes, but not both.

This command is supported only on VPN VRFs.



Example

The following command displays the counters for VPN VRF red:

```
Switch.19 # show counters vr red
Num of Current Routes: 4           Num of Routes Dropped: 0
Num of Route Add:      12          Num of Route Del:      6
Num of Current Routes: 10         Num of Routes Dropped: 0
Num of Route Add:      5           Num of Route Del:      2
```

History

This command was introduced in ExtremeXOS 15.3.

Platform Availability

This command is available on BlackDiamond8000 c-, xl-, and xm-series modules, and Summit X460, X480, X670, and X770 switches.

show virtual-router

```
show virtual-router {name}
```

Description

Displays information about VRs and VRFs.

Syntax Description

<i>name</i>	Specifies the name of a VR or VRF.
-------------	------------------------------------

Default

N/A.

Usage Guidelines

The output display differs for the following options:

- `show virtual-router`—displays information about all VRs and VRFs.
- `show virtual-router vr_name`—displays information about a user VR or VR-Default.
- `show virtual-router vrf_name`—displays information about the named VRF.



Example

The following command displays the VR and VRF configurations on the switch:

```
Switch.19 # show virtual-router
-----
Virtual          Number of   Number of   Flags
Router           Vlans      Ports
-----
VR-Control       0          0           -----S46
VR-Default       32         18          boprimORS46
xvr              1          0           b-----F46
VR-Mgmt          1          0           -----S46
-----
Flags : Virtual Router Type(
      S) System, (U) User, (F) Non-VPN VRF, (N) VPN VRF
      : Virtual Router Admin State
      (-) Disabled (4) IPv4 Enabled, (6) IPv6 Enabled
      : Routing protocols configured on the virtual router
      (b) BGP, (i) ISIS, (m) MPLS, (o) OSPF, (p) PIM, (r) RIP,
      (O) OSPFv3, (R) RIPng
System Totals      :
Total VRs          :      4   Max VRs          : 1066
Total User VRs    :      0   Max User VRs     :   63
Total Non-VPN VRFs :      1   Max VRFs        : 1000
Total VPN VRFs    :      0   Total System VRs :    3
Total Protocols   :      8   Max Protocols    :   64
Max IPv4 Vlans    :    4096 Max Ipv6 Vlans   : 1024
Total IPv4 Vlans  :     13   Total Ipv6 Vlans :    8
Active IPv4 Vlans :     12   Active Ipv6 Vlans :    8
Inactive IPv4 Vlans :      1   Inactive Ipv6 Vlans :    0
```

The following command displays information about VR-Default:

```
Switch.20 # show virtual-router "VR-Default"
Virtual Router      : VR-Default                Type : System
Description         : Default VR
IPv4 Admin State   : Enabled                IPv4 Forwarding : Enabled
IPv6 Admin State   : Enabled                IPv6 Forwarding : Enabled
Operational State  : Up
IPv4 Route Sharing : Enabled                IPv6 Route Sharing : Disabled
Protocols Configured :
-----
Protocol   Process      Configuration      Protocol
Name       Name         Module Name        Instances
-----
BGP        bgp          bgp                2
OSPF       ospf         ospf                1
PIM        pim          pim                1
RIP        rip          rip                1
ISIS       isis         isis                1
MPLS       mpls         mpls                1
OSPFv3    ospfv3       ospfv3              1
RIPng     ripng        ripng               1
-----
VRFs Configured   :
```



```

-----
Virtual Router          Flags
-----
xvr                     b-----F46
-----
Flags : Virtual Router Type(
      S) System, (U) User, (F) Non-VPN VRF, (N) VPN VRF
      : Virtual Router Admin State
      (-) Disabled (4) IPv4 Enabled, (6) IPv6 Enabled
      : Routing protocols configured on the virtual router
      (b) BGP, (i) ISIS, (m) MPLS, (o) OSPF, (p) PIM, (r) RIP,
      (O) OSPFv3, (R) RIPng
Port List                : 1:2-14, 1:16, 1:19-20, 2:1-2
VLANS                    : Default, e1, foo1,
                        foo10, foo11, foo12,
                        foo13, foo14, foo15,
                        foo16, foo17, foo18,
                        foo19, foo2, foo3,
                        foo4, foo5, foo6,
                        foo7, foo8, foo9,
                        fra_mil1, fra_mil2, fra_mil3,
                        fra_mil4, lo, loo,
                        v23, v77, vk-01,
                        vlan1, vlan2

Virtual Router Totals :
Total Non-VPN VRFs    :    1
Total VPN VRFs        :    0
Total Protocols       :    8   Max Protocols           :    8
Total Ports           :   18
Total Vlans           :   32
Total IPv4 Vlans      :   12   Total Ipv6 Vlans      :    8
Active IPv4 Vlans     :   11   Active Ipv6 Vlans     :    8
Inactive IPv4 Vlans   :    1   Inactive Ipv6 Vlans   :    0

```

The following command displays detailed information for a VRF:

```

t15.3 # sh virtual-router xvr
Virtual Router          : xvr                               Type : Non-VPN VRF
IPv4 Admin State       : Enabled                           IPv4 Forwarding      :
Enabled
IPv6 Admin State       : Enabled                           IPv6 Forwarding      :
Enabled
Operational State     : Up
IPv4 Route Sharing    : Enabled                           IPv6 Route Sharing   :
Disabled
Parent VR              : VR-Default
Protocols Configured  :
-----
Protocol   Process      Configuration      Protocol
Name       Name         Module Name        Instances
-----
BGP        bgp          bgp-3              2
-----
VLANS      : xlan
Virtual Router Totals :

```



```

Total Protocols      :    1      Max Protocols      :    8
Total Ports          :    0
Total Vlans          :    1
Total IPv4 Vlans    :    1      Total Ipv6 Vlans    :    0
Active IPv4 Vlans   :    1      Active Ipv6 Vlans   :    0
Inactive IPv4 Vlans :    0      Inactive Ipv6 Vlans :    0
    
```

The following command displays information for user VR region1:

```

t16.1 # sh virtual-router
-----
Virtual          Number of   Number of   Flags
Router           Vlans      Ports
-----
uservr-1         26         0           bo---m--U46
  xxx            1         0           b-----N46
VR-Control       0         0           -----S46
VR-Default       12        14          bopri-ORS46
VR-Mgmt          1         0           -----S46
-----
Flags : Virtual Router Type
      (S) System, (U) User, (F) Non-VPN VRF, (N) VPN VRF
      : Virtual Router Admin State
      (-) Disabled (4) IPv4 Enabled, (6) IPv6 Enabled
      : Routing protocols configured on the virtual router
      (b) BGP, (i) ISIS, (m) MPLS, (o) OSPF, (p) PIM, (r) RIP,
      (O) OSPFv3, (R) RIPng
System Totals      :
Total VRs          :    5      Max VRs          : 1066
Total User VRs    :    1      Max User VRs    :   63
Total Non-VPN VRFs :    0      Max VRFs        : 1000
Total VPN VRFs    :    1      Total System VRs :    3
Total Protocols   :   10      Max Protocols   :   64
Max IPv4 Vlans    :  4096      Max Ipv6 Vlans  :  1024
Total IPv4 Vlans  :   19      Total Ipv6 Vlans :    9
Active IPv4 Vlans :   18      Active Ipv6 Vlans :    9
Inactive IPv4 Vlans :    1      Inactive Ipv6 Vlans :    0
    
```

Show virtual-router detail for a user created VR

```

t16.3 # sh virtual-router "uservr-1"
Virtual Router      : uservr-1                      Type : User
IPv4 Admin State   : Enabled      IPv4 Forwarding   : Enabled
IPv6 Admin State   : Enabled      IPv6 Forwarding   : Enabled
Operational State  : Up
IPv4 Route Sharing : Enabled      IPv6 Route Sharing : Disabled
L3VPN SNMP Traps   : Disabled
Protocols Configured :
-----
Protocol   Process      Configuration   Protocol
Name       Name         Module Name     Instances
-----
BGP        bgp-5        bgp-5           2
OSPF       ospf-5       ospf-5          1
MPLS       mpls-5       mpls-5          1
-----
VRFs Configured   :
-----
    
```



```

Virtual Router                               Flags
-----
xxx                                           b-----N46
-----

Flags : Virtual Router Type
      : (S) System, (U) User, (F) Non-VPN VRF, (N) VPN VRF
      : Virtual Router Admin State
      : (-) Disabled (4) IPv4 Enabled, (6) IPv6 Enabled
      : Routing protocols configured on the virtual router
      : (b) BGP, (i) ISIS, (m) MPLS, (o) OSPF, (p) PIM, (r) RIP,
      : (O) OSPFv3, (R) RIPng

Route Exports into L3VPN (BGP):
-----
--
VPN VRF Policy                               Route Type   Flags       Priority
-----
--
xxx                                           Direct       EO          2048
  None
vpn2                                           Static       EO          2048
  None
-----
--
Flags: (E) Export Enabled, (L) Export Operationally Off due to Low Memory,
      (O) Export Operationally On

VLANS                                         : foo1, foo10, foo11,
                                           : foo12,foo13, foo14,
                                           : foo15,foo16, foo17,
                                           : foo18,foo19, foo2,
                                           : foo3, foo4,  foo5,
                                           : foo6, foo7,  foo8,
                                           : foo9, lan-uvr-1, lo,
                                           : lo-uvr-1, v77, v88,
                                           : vlan1, vlan2

Virtual Router Totals :
Total Non-VPN VRFs   : 0   Total VPN VRFs       : 1
Total Protocols      : 3   Max Protocols        : 8
Total Ports          : 0   Total Vlans          : 26
Total IPv4 Vlans     : 7   Total Ipv6 Vlans    : 1
Active IPv4 Vlans    : 7   Active Ipv6 Vlans   : 1
Inactive IPv4 Vlans  : 0   Inactive Ipv6 Vlans : 0

Show virtual router for a VPN VRF t16.2 # sh virtual-router "xxx"
Virtual Router       : xxx                               Type : VPN VRF
IPv4 Admin State     : Enabled   IPv4 Forwarding    : Enabled
IPv6 Admin State     : Enabled
                    : IPv6 Forwarding : Enabled
Operational State    : Up
IPv4 Route Sharing   : Enabled   IPv6 Route Sharing : Disabled
Parent VR            : uservr-1
VPN ID                :
VPN RD                : 1:1

```




```

Export RT          : 1:1
Import RT         : 1:1
Protocols Configured :
-----
Protocol   Process      Configuration   Protocol
Name       Name         Module Name     Instances
-----
BGP        bgp-5        bgp-3           2
-----
Route Exports into L3VPN (BGP):
-----
--
VPN VRF          Route Type      Flags           Priority
  Policy
-----
--
xxx              Direct          EO              2048
  None
-----
--
Flags: (E) Export Enabled, (L) Export Operationally Off due to Low Memory,
      (O) Export Operationally On
VLANS          : xlan
Virtual Router Totals :
Total Protocols   : 1   Max Protocols   : 8
Total Ports       : 0   Total Vlans     : 1
Total IPv4 Vlans  : 1   Total Ipv6 Vlans : 0
Active IPv4 Vlans : 1   Active Ipv6 Vlans : 0
Inactive IPv4 Vlans : 0   Inactive Ipv6 Vlans : 0

```

The current and configured values for **max-gateways** now apply to IPv6 gateway sets as well as IPv4, so these values are added to the output of `show ipconfig ipv6`.

```

# show virtual-router
Virtual Router          Number of      Number of
Flags                   Vlans
-----
VR-Boston              2             2   -----
U46
VR-Control             0             0   -----
S46
VR-Default             1             278
boprimORS46
VR-Mgmt               1             0   -----
S46
-----
--
Flags : Virtual Router Type
(S) System, (U) User, (F) Non-VPN VRF, (N) VPN VRF
: Virtual Router Admin State
(-) Disabled (4) IPv4 Enabled, (6) IPv6 Enabled
: Routing protocols configured on the virtual router
(b) BGP, (i) ISIS, (m) MPLS, (o) OSPF, (p) PIM, (r) RIP,
(O) OSPFv3, (R) RIPng

```



```

System Totals      :
Total VRs          :      4   Max      VRs          :    256
Total User VRs     :      1   Max      User VRs       :     63
Total Non-VPN VRFs :      0   Max      VRFs          :    190
Total VPN VRFs     :      0
Total System VRs   :      3
Total Protocols    :      8   Max
      Protocols    :      64
Max IPv4 Vlans     :    512   Max      IPv6 Vlans       :    512
Total IPv4 Vlans   :      0
Active IPv4 Vlans  :      0
Inactive IPv4 Vlans :      0
Max Shared GWs (Cur) :    32
Total IPv6 Vlans   :      0
Active IPv6 Vlans  :      0
Inactive IPv6 Vlans :      0
Max Shared GWs (Cfg) :    32

```

History

A command similar to this command was available in ExtremeXOS 10.1 (show vr).

This command was first available in ExtremeXOS 11.0.

Support for non-VPN VRFs was added in ExtremeXOS 12.5.

The show output for **max-gateways** was added in ExtremeXOS 15.3.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond8000 c-, xl-, and xm-series modules, and Summit X460, X480, X670, and X770 switches.

unconfigure vr description

unconfigure vr *name* **description**

Description

Removes a description for the specified VR or VRF.

Syntax Description

<i>name</i>	Specifies the name of a user VR or a VRF.
-------------	---

Default

No description.

Usage Guidelines

None.



Example

The following command removes a description for the VRF named corporate:

```
unconfigure vr corporate description
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond8000 c-, xl-, and xm-series modules, and Summit X460, X480, X670, and X770 switches.

unconfigure vr rd

```
unconfigure vr vrf_name rd
```

Description

This command removes the configuration for a VPN VRF RD.

Syntax Description

<i>vrf_name</i>	Specifies the name of a VPN VRF.
-----------------	----------------------------------

Default

N/A.

Usage Guidelines

None.

Example

The following examples unconfigure RDs using the two of the supported formats:

```
unconfigure vr corporate-extreme rd  
unconfigure vr corporate-guest rd
```

History

This command was first available in ExtremeXOS 12.5.



Platform Availability

This command is available on BlackDiamond8900 xl- and xm-series modules and Summit X460, X480, X670, and X770 switches.

unconfigure vr vpn-id

```
unconfigure vr vrf_name vpn-id
```

Description

This command removes the configuration for a globally unique identifier for a VPN VRF.

Syntax Description

<i>vrf_name</i>	Specifies the name of a VPN VRF.
-----------------	----------------------------------

Default

N/A.

Usage Guidelines

None.

Example

The following example removes VPN ID ac:9f3c8 from the VRF named corporate-extreme:

```
unconfigure vr corporate-extreme vpn-id
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on BlackDiamond8900 xl- and xm-series modules and Summit X460, X480, X670, and X770 switches.

virtual-router

```
virtual-router {vr-name}
```



Description

Changes the VR context.

Syntax Description

<i>vr-name</i>	Specifies the name of the VR.
----------------	-------------------------------

Default

N/A.

Usage Guidelines

Use this command to change the VR context for subsequent commands. When you issue the command, the prompt changes to reflect the VR domain. Configuration commands for Layer 3 routing protocols, creating VLANs, and deleting VLANs apply only to the current VR context.

Use this command with no name, or use the name VR-Default to return to the default configuration domain.

Under a VR configuration domain, any VR commands are applied only to that VR. The VR commands consist of all the BGP, OSPF, PIM and RIP commands, and the commands listed in the following table.

Table 27: VR Commands

[enable disable] ipforwarding
clear iparp ⁹
clear counters iparp ^a
configure iparp ^a
configure iparp [add delete] ^a
[enable disable] iparp ^a
show iparp ^a
configure iproute [add delete] ^a
show iproute ^a
show ipstats ^a
rtlookup
create [vlan vman] <vlan-name>
[enable disable] igmp
[enable disable] igmp snooping ^a
[enable disable] ipmcforwarding
show igmp

⁹ Other commands are available with these listed.



Table 27: VR Commands (continued)

show igmp snooping
show igmp group
show igmp snooping cache
[enable disable] mld
[enable disable] mld snooping
show mld
show mld snooping
show mld group

The VR context simplifies configuration because you do not have to specify the VR for each individual protocol configuration command. The current VR context is indicated in the command line interface (CLI) prompt.

For example, if you wish to configure OSPF for the user VR `vr-manufacturing`, you would change the VR context to that of `vr-manufacturing`. All the subsequent OSPF commands would apply to that VR, unless the context is changed again.

A VR is identified by a name (up to 32 characters long). The name must be unique among the VLAN and VR names on the switch. For backward compatibility, you cannot name a virtual router `VR-0`, `VR-1`, or `VR-2`. VR names are case insensitive.

When a new VR is created, by default, no ports are assigned, no VLAN interface is created, and no support for any routing protocols is added.

Example

The following command changes the VR context to `vr-acme`:

```
virtual-router vr-acme
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond8000 c-, xl-, and xm-series modules, and Summit X460, X480, X670, and X770 switches.



21 Policy Manager Commands

check policy
check policy attribute
edit policy
refresh policy
show policy

This chapter describes commands for:

- Creating and configuring policy files for IP access lists (ACLs).
- Creating and configuring policy files for routing policies.

Policies are a generalized category of features that impact forwarding and route forwarding decisions. Access policies are used primarily for security and quality of service (QoS) purposes.

IP access lists (also referred to as Access Lists or ACLs) consist of IP access rules and are used to perform packet filtering and forwarding decisions on traffic traversing the switch. Each packet on an interface is compared to the access list in sequential order and is either forwarded to a specified QoS profile or dropped. Additionally, for the BlackDiamond 8800 series and Summit family switches only, packets can be metered using ACLs. Using access lists has no impact on switch performance.

Access lists are typically applied to traffic that crosses Layer 3 router boundaries, but it is possible to use access lists within a Layer 2 VLAN. Extreme products are capable of performing this function with no additional configuration.

Routing policies are used to control the advertisement or recognition of routes from routing protocols, such as RIP, OSPF, IS-IS, or BGP. Routing policies can be used to 'hide' entire networks or to trust only specific sources for routes or ranges of routes. The capabilities of routing policies are specific to the type of routing protocol involved, but are sometimes more efficient and easier to implement than access lists.

Note



Although ExtremeXOS does not prohibit mixing ACL and routing type entries in a policy file, it is strongly recommended that you do not mix the entries, and you use separate policy files for ACL and routing policies.

check policy

```
check policy policy-name {access-list}
```

Description

Checks the syntax of the specified policy.

Syntax Description

<i>policy-name</i>	Specifies the policy to check.
access-list	Specifies that an access list specific check is performed.

Default

N/A.

Usage Guidelines

Use this command to check the policy syntax before applying it. If any errors are found, the line number and a description of the syntax error are displayed. A policy that contains syntax errors will not be applied.

This command can only determine if the syntax of the policy file is correct and can be loaded into the policy manager database. Since a policy can be used by multiple applications, a particular application may have additional constraints on allowable policies.

Example

The following example checks the syntax of the policy zone5:

```
check policy zone5
```

If no syntax errors are discovered, the following message is displayed:

```
Policy file check successful.
```

History

This command was available in ExtremeXOS 10.1.

The success message and the access-list keyword was added in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

check policy attribute

```
check policy attribute {attr}
```

Description

Displays the syntax of the specified policy attribute.



Syntax Description

<code>attr</code>	Specifies the attribute check.
-------------------	--------------------------------

Default

N/A.

Usage Guidelines

Use this command to display the syntax of policy attributes. The command displays any additional keywords to use with this attribute, and the types of values expected.

Policy attributes are used in the rule entries that make up a policy file.

For each attribute, this command displays which applications use the attribute, and whether the attribute is a match condition or a set (action, action modifier) condition.

The current applications are:

- ACL—access-lists.
- RT—routing profiles, route maps.
- CLF—CLEAR-Flow.

The syntax display does not show the text synonyms for numeric entries. For example, the `icmp-type` match condition allows you to specify either an integer or a text synonym for the condition. Specifying `icmp-type 8` or `icmp-type echo-request` are equivalent, but the syntax display shows only the numeric option.

Note



The syntax displayed is used by the policy manager to verify the syntax of policy files. The individual applications are responsible for implementing the individual attributes. Inclusion of a particular policy attribute in this command output does not imply that the attribute has been implemented by the application. See the documentation of the particular application for detailed lists of supported attributes.

Example

The following example displays the syntax of the policy attribute `icmp-type`:

```
check policy attribute icmp-type
```

The following is sample output for this command:

```
( match ) ( ACL )
icmp-type <uint32 val>
```



History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

edit policy

edit policy *filename*

Description

Edits a policy text file.

Syntax Description

<i>filename</i>	Specifies the filename of the policy text file.
-----------------	---

Default

N/A.

Usage Guidelines

This command edits policy text files that are on the switch. All policy files use “.pol” as the filename extension, so to edit the text file for the policy boundary use boundary.pol as the filename. If you specify the name of a file that does not exist, you will be informed and the file will be created.

This command spawns a VI-like editor to edit the named file. For information on using VI, if you are not familiar with it, do a web search for “VI editor basic information”, and you should find many resources. The following is only a short introduction to the editor.

Edit operates in one of two modes; command and input. When a file first opens, you are in the command mode. To write in the file, use the keyboard arrow keys to position your cursor within the file, then press one of the following keys to enter input mode:

- i - To insert text ahead of the initial cursor position.
- a - To append text after the initial cursor position.

To escape the input mode and return to the command mode, press the Escape key.

There are several commands that can be used from the command mode. The following are the most commonly used:

- dd - To delete the current line.
- yy - To copy the current line.
- p - To paste the line copied.
- :w - To write (save) the file.



- `:q` - To quit the file if no changes were made.
- `:q!` - To forcefully quit the file without saving changes.
- `:wq` - To write and quit the file.

Refresh Policy

After you have edited the text file for a policy that is currently active, you will need to refresh the policy if you want the changes to be reflected in the policy database. When you refresh the policy, the text file is read, the syntax is checked, the policy information is added to the policy manager database, and the policy then takes effect. Use the following command to refresh a policy:

```
refresh policy policy-name
```

If you just want to check to be sure the policy contains no syntax errors, use the following command:

```
check policy policy-name {access-list}
```

Example

The following command allows you to begin editing the text file for the policy boundary:

```
edit policy boundary.pol
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

refresh policy

```
refresh policy policy-name
```

Description

Refreshes the specified policy.

Syntax Description

<i>policy-name</i>	Specifies the policy to refresh.
--------------------	----------------------------------

Default

N/A.



Usage Guidelines

Use this command when a new policy file for a currently active policy has been downloaded to the switch, or when the policy file for an active policy has been edited. This command reprocesses the text file and updates the policy database.

Before 12.6.1 there was no support to refresh the policies that are associated to the local VPP. For network VPP, you can achieve policy refresh by changing the policy timestamp file. Beginning in release 11.4, the policy manager uses Smart Refresh to update the ACLs. When a change is detected, only the ACL changes needed to modify the ACLs are sent to the hardware, and the unchanged entries remain. This behavior avoids having to blackhole packets because the ACLs have been momentarily cleared. Smart Refresh works well for minor changes, however, if the changes are too great, the refresh reverts to the earlier behavior. To take advantage of Smart Refresh, disable access-list refresh blackholing by using the command:

```
disable access-list refresh blackhole
```

If you attempt to refresh a policy that cannot take advantage of Smart Refresh while blackholing is enabled, you will receive a message similar to the following:

```
Incremental refresh is not possible given the configuration of policy
<name>. Note, the current setting for Access-list Refresh Blackhole is
Enabled. Would you like to perform a full refresh? (Yes/No) [No]:
```

If blackholing is not enabled, you will receive a message similar to the following:

```
Incremental refresh is not possible given the configuration of policy
<name>. Note, the current setting for Access-list Refresh Blackhole is
Disabled. WARNING: If a full refresh is performed, it is possible packets
that should be denied may be forwarded through the switch during the time
the access list is being installed. Would you like to perform a full
refresh? (Yes/No) [No]:
```

If you attempt to refresh a policy that is not currently active, you will receive an error message.

For an ACL policy, the command is rejected if there is a configuration error or hardware resources are not available.

Example

The following example refreshes the policy zone5:

```
refresh policy zone5
```

History

This command was first available in ExtremeXOS 11.0.

Smart Refresh was added in ExtremeXOS 11.4.



Platform Availability

This command is available on all platforms.

show policy

```
show policy {policy-name | detail}
```

Description

Displays the specified policy.

Syntax Description

<i>policy-name</i>	Specifies the policy to display.
detail	Show the policy in detail.

Default

If no policy name is specified, all policies are shown

Usage Guidelines

Use this command to display which clients are using the specified policy. The detail option displays the rules that make up the policy.

Example

The following example displays all policies on the switch:

```
show policy
```

The following is sample output for the command:

```
Switch # sh policy
Policies at Policy Server:
PolicyName                ClientUsage    Client          BindCount
-----
p1                          1              acl             1
p2                          1              acl             1
vlanV1                      1              acl             1
-----
Total Policies : 3
```

History

This command was available in ExtremeXOS 10.1.



Platform Availability

This command is available on all platforms.



22 ACL Commands

clear access-list counter
clear access-list meter
configure access-list
configure access-list add
configure access-list delete
configure access-list network-zone
configure access-list rule-compression port-counters
configure access-list vlan-acl-precedence
configure access-list width
configure access-list zone
configure flow-redirect add nexthop
configure flow-redirect delete nexthop
configure flow-redirect health-check
configure flow-redirect nexthop
configure flow-redirect no-active
configure flow-redirect vr
create access-list
create access-list zone
create access-list network-zone
create flow-redirect
delete access-list
delete access-list network-zone
delete access-list zone
delete flow-redirect
disable access-list permit to-cpu
disable access-list refresh blackhole
enable access-list permit to-cpu
enable access-list refresh blackhole
refresh access-list network-zone
show access-list
show access-list configuration
show access-list counter
show access-list dynamic
show access-list dynamic counter
show access-list dynamic rule
show access-list interface
show access-list network-zone

```

show access-list usage acl-mask port
show access-list usage acl-range port
show access-list usage acl-rule port
show access-list usage acl-slice port
show access-list width
show flow-redirect
unconfigure access-list

```

This chapter describes commands for:

- Creating and configuring IP access lists (ACLs)

IP access lists (also referred to as Access Lists or ACLs) consist of IP access rules and are used to perform packet filtering and forwarding decisions on traffic traversing the switch. Each packet on an interface is compared to the access list in sequential order and is either forwarded to a specified QoS profile or dropped. Additionally, for the BlackDiamond X8 series switches, BlackDiamond 8800 series, E4G-200 and E4G-400 switches, SummitStack, and Summit family switches only, packets can be metered using ACLs. Using access lists has no impact on switch performance.

Access lists are typically applied to traffic that crosses Layer 3 router boundaries, but it is possible to use access lists within a Layer 2 VLAN. Extreme products are capable of performing this function with no additional configuration.

Note



Although ExtremeXOS does not prohibit mixing ACL and routing type entries in a policy file, we strongly recommend that you do not mix the entries, and that you use separate policy files for ACL and routing policies.

clear access-list counter

```

clear access-list {dynamic} counter {countername} {any | ports port_list | vlan
vlan_name} {ingress | egress}

```

Description

Clears the specified access list counters.

Syntax Description

dynamic	Specifies that the counter is from a dynamic ACL.
<i>countername</i>	Specifies the ACL counter to clear.
any	Specifies the wildcard ACL.
<i>port_list</i>	Specifies to clear the counters on these ports.
<i>vlan_name</i>	Specifies to clear the counters on the VLAN.



ingress	Clear the ACL counter for packets entering the switch on this interface.
egress	Clear the ACL counter for packets leaving the switch from this interface (BlackDiamond X8 series switches, BlackDiamond 8000 c-, xl-, and xm-series modules, E4G-200 and E4G-400 switches, and Summit X460, X480, and X670 series switches only).

Default

The default direction is ingress; the default ACL type is non-dynamic.

Usage Guidelines

Use this command to clear the ACL counters. If you do not specify an interface, or the any option, you will clear all the counters.

Example

The following example clears all the counters of the ACL on port 2:1:

```
clear access-list counter port 2:1
```

The following example clears the counter counter2 of the ACL on port 2:1

```
clear access-list counter counter2 port 2:1
```

History

This command was first available in ExtremeXOS 10.1.

The VLAN option was first available in ExtremeXOS 11.0.

The egress and dynamic options were first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

The egress option is available on BlackDiamond X8 series switches, BlackDiamond 8000 c- xl-, and xm-series modules, the E4G-200 and E4G-400 switches, and Summit X460, X480, X670, and X770 switches only.

clear access-list meter

```
clear access-list meter {meter_name} [any | ports port_list | vlan vlan_name]
```



Description

Clears the specified access list meters.

Syntax Description

<i>meter_name</i>	Specifies the ACL meter to clear.
<i>port_list</i>	Specifies to clear the counters on these ports.
<i>vlan_name</i>	Specifies to clear the counters on the VLAN.

Default

N/A.

Usage Guidelines

Use this command to clear the out-of-profile counters associated with the meter configuration.

Example

The following example clears all the out-of-profile counters for the meters of the ACL on port 2:1:

```
clear access-list meter port 2:1
```

The following example clears the out-of-profile counters for the meter meter2 of the ACL on port 2:1:

```
clear access-list meter meter2 port 2:1
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available only on the BlackDiamond X8 series switches, BlackDiamond 8000 series modules, SummitStack, the E4G-200 and E4G-400 switches, and the Summit family of switches.

configure access-list

```
configure access-list aclname [any | ports port_list | vlan vlan_name] {ingress | egress}
```

Description

Configures an access list to the specified interface.



Syntax Description

<i>aclname</i>	Specifies the ACL policy file name.
any	Specifies that this ACL is applied to all interfaces as the lowest precedence ACL.
<i>port_list</i>	Specifies the ingress or egress port list on which the ACL is applied.
<i>vlan_name</i>	Specifies the VLAN on which the ACL is applied.
ingress	Apply the ACL to packets entering the switch on this interface.
egress	Apply the ACL to packets leaving the switch from this interface. (BlackDiamond X8 series switches, BlackDiamond 8000 c-, xl-, xm-series modules, E4G-200 and E4G-400 switches, and Summit X460, X480, X670, and X770 switches only).

Default

The default direction is ingress.

Usage Guidelines

The access list applied in this command is contained in a text file created either externally to the switch or using the `edit policy` command. The file is transferred to the switch using TFTP before it is applied to the ports. The ACL name is the file name without its “.pol” extension. For example, the ACL `blocknetfour` would be in the file `blocknetfour.pol`. For more information on policy files, see the *ExtremeXOS Concepts Guide*.

Specifying the keyword `any` applies the ACL to all the ports, and is referred to as the wildcard ACL. This ACL is evaluated for ports without a specific ACL applied to it, and is also applied to packets that do not match the ACL applied to the interface.

Example

The following command configures the ACL policy `test` to port 1:2 at ingress:

```
configure access-list test ports 1:2
```

The following command configures the ACL `mydefault` as the wildcard ACL:

```
configure access-list mydefault any
```

The following command configures the ACL policy `border` as the wildcard egress ACL:

```
configure access-list border any egress
```



History

This command was first available in ExtremeXOS 10.1.

The VLAN option was first available in ExtremeXOS 11.0.

The egress option was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

The egress options are available on BlackDiamond X8 series switches, BlackDiamond 8000 c- xl-, and xm-series modules, E4G-200 and E4G-400 switches, and Summit X460, X480, X670, and X770 switches only.

configure access-list add

```
configure access-list add dynamic_rule [ [[first | last] {priority p_number}
{zone zone} ] | [[before | after] rule] | [ priority p_number {zone zone} ] ]
[ any | vlan vlan_name | ports port_list ] {ingress | egress}
```

Description

Configures a dynamic ACL rule to the specified interface and sets the priority and zone for the ACL.

Syntax Description

<i>dynamic_rule</i>	Specifies a dynamic ACL rule.
first	Specifies that the new dynamic rule is to be added as the first rule.
last	Specifies that the new dynamic rule is to be added as the last rule.
priority	Priority of rule within a zone.
<i>p_number</i>	Specifies the priority number of the rule within a zone. The range is from 0 (highest priority) to 7 (lowest priority).
<i>zone</i>	Specifies the ACL zone for the rule.
before rule	Specifies that the new dynamic rule is to be added before an existing dynamic rule.
after rule	Specifies that the new dynamic rule is to be added after an existing dynamic rule.
any	Specifies that this ACL is applied to all interfaces.
<i>vlan_name</i>	Specifies the VLAN on which this ACL is applied.
<i>port_list</i>	Specifies the ports on which this ACL is applied.



ingress	Apply the ACL to packets entering the switch on this interface.
egress	Apply the ACL to packets leaving the switch from this interface (BlackDiamond X8 series switches, BlackDiamond 8000 c- xl-, and xm-series modules, E4G-200 and E4G-400 switches, and Summit X460, X480, X670, and X770 series switches only).

Default

The default direction is ingress.

Usage Guidelines

The dynamic rule must first be created before it can be applied to an interface. Use the following command to create a dynamic rule:

```
create access-list dynamic-rule conditions actions {non-permanent}
```

When a dynamic ACL rule is applied to an interface, you will specify its precedence among any previously applied dynamic ACLs. All dynamic ACLs have a higher precedence than any ACLs applied through ACL policy files.

Specifying the keyword `any` applies the ACL to all the ports, and is referred to as the wildcard ACL. This ACL is evaluated for ports without a specific ACL applied to them, and is also applied to packets that do not match the ACL applied to the interface.

The `priority` keyword can be used to specify a sub-zone within an application's space. For example, to place ACLs into three sub-zones within the CLI application, you can use three priority numbers, such as 2, 4, and 7.

Configuring priority number 1 is the same as configuring first priority. Configuring priority number 8 is the same as configuring last priority.

Example

The following command applies the dynamic ACL `icmp-echo` as the first (highest precedence) dynamic ACL to port 1:2 at ingress:

```
configure access-list add icmp-echo first ports 1:2
```

The following command applies the dynamic ACL `udpacl` to port 1:2, with a higher precedence than rule `icmp-echo`:

```
configure access-list add udpacl before icmp-echo ports 1:2
```

History

This command was first available in ExtremeXOS 11.3.



Platform Availability

This command is available on all platforms.

The egress option is available on BlackDiamond X8 series switches, BlackDiamond 8000 c- xl-, and xm-series modules, E4G-200 and E4G-400 switches, and Summit X460, X480, X670, and X770 series switches only.

configure access-list delete

```
configure access-list delete ruleName [ any | vlan vlan_name | ports port_list | all ] { ingress | egress }
```

Description

Removes a dynamic ACL rule from the specified interface.

Syntax Description

<i>ruleName</i>	Specifies a dynamic ACL rule name.
any	Deletes this ACL as the wildcard ACL.
<i>vlan_name</i>	Specifies the VLAN on which this ACL is deleted.
<i>port_list</i>	Specifies the ports on which this ACL is deleted.
all	Deletes this ACL from all interfaces.
ingress	Deletes the ACL for packets entering the switch on this interface.
egress	Deletes the ACL for packets leaving the switch from this interface (BlackDiamond X8 series switches, BlackDiamond 8000 c-, xl-, and xm-series modules, E4G-200 and E4G-400 switches, and Summit X460, X480, X670, and X770 series switches only).

Default

The default direction is ingress.

Usage Guidelines

Specifying the keyword all removes the ACL from all interfaces it is used on.

Example

The following command removes the dynamic ACL icmp-echo from the port 1:2:

```
configure access-list delete icmp-echo ports 1:2
```



History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

The egress option is available on BlackDiamond X8 series switches, BlackDiamond 8000 c-, xl-, and xm-series modules, E4G-200 and E4G-400 switches, and Summit X460, X480, X670, and X770 series switches only.

configure access-list network-zone

```
configure access-list network-zone zone_name [add | delete] [mac-address
macaddress {macmask} | ipaddress [ipaddress {netmask} | ipNetmask |
ipv6_address_mask]]
```

Description

Adds or removes IP and MAC addresses to and from the network-zone.

Syntax Description

network-zone	Logical group of remote devices.
<i>zone_name</i>	Specifies the network-zone name.
add	Adds a logical group of entities to the network-zone.
delete	Deletes a logical group of entities to the network-zone.
mac-address	MAC address.
<i>macaddress</i>	Specifies the MAC address to be added/removed to/from the network-zone.
<i>macmask</i>	Specifies the MAC Mask. Example FF:FF:FF:00:00:00.
ipaddress	Specifies IPv4 address.
<i>ipaddress</i>	Specifies the IP address.
<i>netmask</i>	Specifies IP netmask.
<i>ipNetmask</i>	Specifies the IP address/Netmask.
<i>ipv6_address_mask</i>	Specifies IPv6 address/IPv6 prefix length.

Default

N/A.

Usage Guidelines

Use this command to to add or remove IP/MAC addresses to/from the network-zone.



Example

The following command adds an IPv6 IP address to network-zone “zone1”:

```
Switch# configure access-list network-zone zone1 add ipaddress
11.1.1.1/32
```

If you try to add the same IP/MAC with the same or narrow mask, the configuration is rejected, with the following error message.

```
Switch #configure access-list network-zone "zone1" add ipaddress 11.1.1.1/24
Error: Network Zone "zone1" - Zone already has the same entity value with
same or wider mask.
```

If you try to add more than eight attributes to a network-zone, the following error message is printed.

```
Switch #configure access-list network-zone "zone1" add ipaddress 11.1.1.1/24
Error: Network Zone "zone1" - Reached maximum number of attributes. Unable
to add more.
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on all platforms.

configure access-list rule-compression port-counters

```
configure access-list rule-compression port-counters [shared | dedicated]
```

Description

Switches between ACL configuration modes.

Syntax Description

shared	Sharing is “on” for counter rules.
dedicated	Sharing is “off” for counter rules.

Default

Dedicated.



Usage Guidelines

Use this command to switch between two ACL configuration modes. In the first mode, “port-counters shared”, similar port-based ACL rules with counters are allowed to share the same hardware entry. This uses less space but provides an inaccurate counter value. In the second mode, “port-counters dedicated”, similar port-based ACL rules with counters are not allowed to share the same hardware entry, thereby consuming more entries but providing a precise count.

Only ACLs that are entered after this command is entered are affected. The command does not affect any ACLs that are already configured.

To configure all ACLs in shared mode, configure access-list rule-compression port-counters shared must be entered before any ACLs are configured or have been saved in the configuration when a switch is booted.

This is a global setting for the switch; that is, the option does not support setting some ACL rules with shared counters and some with dedicated counters.

To view the results of the configuration use the `show access-list configuration` command.

Example

The following command configures ACL rules with counters to share the same hardware entry:

```
configure access-list rule-compression port-counters shared
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on the BlackDiamond X8 series switches, BlackDiamond 8800 series switches, E4G-200 and E4G-400 switches, and the Summit series switches.

configure access-list vlan-acl-precedence

```
configure access-list vlan-acl-precedence [dedicated | shared]
```

Description

Configures precedence mode for policy-file based ACLs that are applied on a VLAN.

Syntax Description

dedicated	Allocates exclusive precedence for VLAN-based ACLs.
shared	VLAN-based ACLs share the precedence with other ACLs.



Default

Dedicated.

Usage Guidelines

The following feature applies to only policy-file based ACLs that are applied on a VLAN. Use this command to switch between two VLAN-based ACL configuration modes. In the shared `vlan-acl-precedence` mode, VLAN-based ACL rules share the same precedence with other types of ACL rules. This is the default mode and provides the same behavior as in the previous software releases. In the dedicated `vlan-acl-precedence` mode, VLAN-based ACL rules have different precedence compared to other types of ACL rules. The dedicated mode yields improved installation performance for VLAN-based access-lists but may affect hardware rule utilization in some configurations.

After configuring, you are prompted to reboot the system for the changes to take effect.

Example

The following command allocates exclusive precedence for VLAN-based static ACL rules:

```
configure access-list vlan-acl-precedence dedicated
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.

configure access-list width

```
configure access-list width [double | single] [slot slotNo | all]
```

Description

Configures the TCAM width of a module or switch.

Syntax Description

double	Specifies a double wide ACL TCAM. Provides double wide ACL key with additional qualifiers.
single	Specifies a single wide ACL TCAM.
<i>slotNo</i>	Specifies the slot to configure.
all	Specifies all slots.



Default

Single.

Usage Guidelines

Use this feature to configure the width of the ACL TCAM key of a slot or switch to be either double wide or single wide.

The switch must be rebooted for the configuration change to take effect.

If you attempt to configure a double wide mode on a slot or switch that does not support it, an error message is displayed.

To display the configured mode, use the `show access-list width` command.

Example

The following command configures slot 1 to use double wide mode:

```
configure access-list width double slot 1
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available only on BlackDiamond X8 series switches, BlackDiamond 8000 c-, xl-, and xm-series modules E4G-200 and E4G-400 switches, and Summit X460, X480, X670, and X770 series switches whether or not included in a SummitStack.

configure access-list zone

```
configure access-list zone name zone-priority number
```

```
configure access-list zone name move-application appl_name to-zone name  
application-priority number
```

```
configure access-list zone name {add} application appl_name application_priority  
number
```

```
configure access-list zone name delete application appl_name
```

Description

Configures the priority of a zone; moves an application from one zone to another at a specified priority; adds an application to a zone with a specified priority, or changes the priority of an application within a zone; deletes an application from a zone.



Syntax Description

<i>name</i>	Specifies a zone name.
zone-priority <i>number</i>	Sets the priority of the zone.
move-application <i>appl_name</i>	Specifies the name of an application to be moved.
to-zone <i>name</i>	Specifies the zone to which the application is moved.
application-priority <i>number</i>	Sets the priority of the application within the zone. The range is from 0 (highest priority) to 7 (lowest priority).
add	Adds an application to a zone at a specified priority.
application <i>appl_name</i>	Specifies the application to be added to the zone.
application_priority number <i>number</i>	Sets the priority of a new or existing application within a zone. The range is from 0 (highest priority) to 7 (lowest priority).

Default

N/A.

Usage Guidelines

To configure the priority of a specific zone, use the syntax:

```
configure access-list zone name zone-priority number
```

To move an application from one zone to another, and set its priority in the new zone, use the syntax:

```
configure access-list zone name move-application appl-name to-zone name
application-priority number
```

To add an application to a zone and specify its priority or to change the priority of an application within a zone, use the syntax:

```
configure access-list zone name {add} application appl-name application_priority
number
```

To delete an application from a zone, use the syntax:

```
configure access-list zone name delete application appl-name
```

Example

The following command adds the CLI application to the zone myzone at a priority of 6:

```
configure access-list zone myzone add cli application-priority 6
```

History

This command was first available in ExtremeXOS 11.6.



Platform Availability

This command is available on all platforms.

configure flow-redirect add nexthop

```
configure flow-redirect flow_redirect_name add nexthop ipaddress priority number
```

Description

Adds a nexthop for the named flow redirection policy.

Syntax Description

<i>flow_redirect_name</i>	Specifies the name of the flow redirection policy.
<i>ipaddress</i>	Specifies the IPv4 or IPv6 address of a new nexthop.
<i>number</i>	Specifies the priority value for the nexthop.

Default

N/A.

Usage Guidelines

Use this command to add a new nexthop for the named flow redirection policy. You can specify an IPv4 address or an IPv6 unicast IP address (IPv6 multicast addresses are not supported). After you enter an IP address, the redirection policy only accepts addresses from the same family as the first address specified. For example, if the first IP address added is an IPv6 unicast address, you cannot add an IPv4 address to the policy.

The priority value can range from a low of “1” to a high of “254.” The nexthop with the highest priority among multiple ones is preferred as the working nexthop. When each added nexthop has the same priority, the first one configured is preferred.

Example

The following command adds a nexthop 10.1.1.1 for the flow redirection policy flow10 with a priority of 100:

```
configure flow-redirect flow10 add nexthop 10.1.1.1 priority 100.
```

History

This command was first available in ExtremeXOS 12.1.

Support for IPv6 flow-redirection policies was added in ExtremeXOS 12.7.



Platform Availability

This command is available for IPv4 and IPv6 flow-redirection policies on the platforms listed for the Policy Based Routing feature in [Feature License Requirements](#).

configure flow-redirect delete nexthop

```
configure flow-redirect flow_redirect_name delete nexthop {ipaddress | all }
```

Description

Deletes a single or all nexthops for the named flow redirection policy.

Syntax Description

<i>flow_redirect_name</i>	Specifies the name of the flow redirection policy.
<i>ipaddress</i>	Specifies the IPv4 or IPv6 address of the nexthop.
all	Specifies that all configured nexthops are to be deleted.

Default

N/A.

Usage Guidelines

Use this command to delete a nexthop for the named flow redirection policy. If the deleted nexthop is the working nexthop for the policy-based routing entry, another is selected from the remaining active next hops, based on priority.

Example

The following command deletes the nexthop 10.1.1.1 from the flow redirection policy flow10:

```
configure flow-redirect flow10 delete nexthop 10.1.1.1
```

The following command deletes all configured nexthop's from the flow redirection policy exflow:

```
configure flow-redirect exflow delete nexthop all
```

History

This command was first available in ExtremeXOS 12.1.

Support for IPv6 flow-redirection policies was added in ExtremeXOS 12.7.



Platform Availability

This command is available for IPv4 and IPv6 flow-redirection policies on the platforms listed for the Policy Based Routing feature in [Feature License Requirements](#).

configure flow-redirect health-check

```
configure flow-redirect flow_redirect_name health-check [ping | arp | neighbor-discovery]
```

Description

Configures health checking for a specific flow redirection policy.

Syntax Description

<i>flow_redirect_name</i>	Specifies the name of the flow redirection policy.
ping	Specifies ping health checking.
arp	Specifies ARP health checking for IPv4.
neighbor-discovery	Specifies Neighbor Discovery health checking for IPv6.

Default

Ping is the default.

Usage Guidelines

Use this command to configure health checking for a specific named flow redirection policy.

Example

The following command specifies arp health checking for the flow redirection policy flow10:

```
configure flow-redirect flow10 health-check arp
```

History

This command was first available in ExtremeXOS 12.1.

Support for IPv6 flow-redirection policies was added in ExtremeXOS 12.7.

Platform Availability

This command is available for IPv4 and IPv6 flow-redirection policies on the platforms listed for the Policy Based Routing feature in [Feature License Requirements](#).



configure flow-redirect nexthop

```
configure flow-redirect flow_redirect_name nexthop ip_address ping health-check  
interval seconds miss number
```

Description

Configures the ping interval and miss count for a nexthop in the flow redirection policy.

Syntax Description

<i>flow_redirect_name</i>	Specifies the name of the flow redirection policy.
<i>ip_address</i>	Specifies the IPv4 or IPv6 address of the nexthop.
<i>seconds</i>	Specifies the number of seconds between pings. The default is "2".
<i>number</i>	Specifies the number of misses allowed. The default is "2".

Default

N/A.

Usage Guidelines

Use this command to set a ping interval and miss count. When the ping response is not received within the interval $seconds * (number + 1)$, the nexthop is considered to be dead and a new candidate is selected from the remaining active nexthops.

Example

The following command configures a ping interval of 3 and miss count of 3 for the nexthop 10.1.1.1 in the flow redirection policy flow 3:

```
configure flow-redirect flow3 nexthop 10.1.1.1 ping interval 3 miss 3
```

History

This command was first available in ExtremeXOS 12.1.

Support for IPv6 flow-redirection policies was added in ExtremeXOS 12.7.

Platform Availability

This command is available for IPv4 and IPv6 flow-redirection policies on the platforms listed for the Policy Based Routing feature in [Feature License Requirements](#).



configure flow-redirect no-active

```
configure flow-redirect flow_redirect_name no-active [drop|forward]
```

Description

Configures packets to either follow the normal routing table or be dropped.

Syntax Description

<i>flow_redirect_name</i>	Specifies the name of the flow redirection policy.
drop	Specifies that the packets are to be dropped.
forward	Specifies that the packets are to follow the normal routing table.

Default

The default is forward.

Usage Guidelines

Use this command to set a drop or forward configuration for packets to be applied when all configured next hops become unreachable.

Example

The following command configures packets of the flow redirection policy flow3 to be dropped when all configured next hops become unreachable:

```
configure flow-redirect flow3 no-active drop
```

History

This command was first available in ExtremeXOS 12.1.

Support for IPv6 flow-redirection policies was added in ExtremeXOS 12.7.

Platform Availability

This command is available for IPv4 and IPv6 flow-redirection policies on the platforms listed for the Policy Based Routing feature in [Feature License Requirements](#).

configure flow-redirect vr

```
configure flow-redirect flow_redirect_name vr vr_name
```



Description

Configures a virtual router for a flow redirection policy.

Syntax Description

<i>flow_redirect_name</i>	Specifies the name of the flow redirection policy.
<i>vr_name</i>	Specifies the name of the virtual router.

Default

The default virtual router is VR-Default.

Usage Guidelines

Because ACLs do not recognize the virtual router concept, one policy-based routing can be used for multiple virtual routing entries when a VLAN-based virtual router is used for one port. This configuration of a VR into a flow-redirect makes a policy-based routing work for a specific VR.

Example

The following command configures virtual router mgmt for flow redirection policy flow3:

```
configure flow-redirect flow3 vr mgmt
```

History

This command was first available in ExtremeXOS 12.1.

Support for IPv6 flow-redirection policies was added in ExtremeXOS 12.7.

Platform Availability

This command is available for IPv4 and IPv6 flow-redirection policies on the platforms listed for the Policy Based Routing feature in [Feature License Requirements](#).

create access-list

```
create access-list dynamic_rule conditions actions {non_permanent}
```

Description

Creates a dynamic ACL.



Syntax Description

<i>dynamic_rule</i>	Specifies the dynamic ACL name. The name can be from 1-32 characters long.
<i>conditions</i>	Specifies the match conditions for the dynamic ACL.
<i>actions</i>	Specifies the actions for the dynamic ACLs.
non_permanent	Specifies that the ACL is not to be saved.

Default

By default, ACLs are permanent.

Usage Guidelines

This command creates a dynamic ACL rule. Use the `configure access-list add` command to apply the ACL to an interface.

The conditions parameter is a quoted string of match conditions, and the actions parameter is a quoted string of actions. Multiple match conditions or actions are separated by semi-colons. A complete listing of the match conditions and actions is in [ACLs](#) in the *ExtremeXOS Concepts Guide*.

Dynamic ACL rule names must be unique, but can be the same as used in a policy-file based ACL. Any dynamic rule counter names must be unique. For name creation guidelines and a list of reserved names, see [Object Names](#) in the *ExtremeXOS Concepts Guide*.

By default, ACL rules are saved when the save command is executed, and persist across system reboots. Configuring the optional keyword non-permanent means the ACL will not be saved.

Example

The following command creates a dynamic ACL that drops all ICMP echo-request packets on the interface:

```
create access-list icmp-echo "protocol icmp;icmp-type echo-request" "deny"
```

The created dynamic ACL will take effect after it has been configured on the interface. The previous example creates a dynamic ACL named icmp-echo that is equivalent to the following ACL policy file entry:

```
entry icmp-echo {
  if {
    protocol icmp;
    icmp-type echo-request;
  } then {
    deny;
  }
}
```



The following command creates a dynamic ACL that accepts all the UDP packets from the 10.203.134.0/24 subnet that are destined for the host 140.158.18.16, with source port 190 and a destination port in the range of 1200 to 1250:

```
create access-list udpacl "source-address 10.203.134.0/24;destination-address
140.158.18.16/32;protocol udp;source-port 190;destination-port 1200 -
1250;" "permit"
```

The previous example creates a dynamic ACL entry named udpacl that is equivalent to the following ACL policy file entry:

```
entry udpacl {
  if {
    source-address 10.203.134.0/24;
    destination-address 140.158.18.16/32;
    protocol udp;
    source-port 190;
    destination-port 1200 - 1250;
  } then {
    permit;
  }
}
```

History

This command was first available in ExtremeXOS 11.3.

The non-permanent option was added in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

create access-list zone

```
create access-list zone name zone-priority number
```

Description

Creates a dynamic ACL zone, and sets the priority of the zone.

Syntax Description

<i>name</i>	Specifies the dynamic ACL zone name. The name can be from 1-32 characters long.
zone-priority <i>number</i>	Specifies priority of the zone. The range is from 1 (highest priority) to 4294967295 (lowest priority).



Default

The denial of service, system, and security zones are configured by default, and cannot be deleted.

Usage Guidelines

This command creates a dynamic ACL zone. You can configure the priority of the zone in relation to the default zones or to other configured zones.

Example

The following command creates a new zone, called myzone, with a priority of 2:

```
create access-list myzone zone-priority 2
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

create access-list network-zone

```
create access-list network-zone zone_name
```

Description

Creates a network-zone with a specified name.

Syntax Description

access-list	Access list
network-zone	Network zone
<i>zone_name</i>	Network zone name

Default

N/A.

Usage Guidelines

Use this command to create a network-zone with a specified name. The network-zone can then be associated with the policy file using either the "source-zone" or "destination-zone" attribute.



Example

```
Switch# create access-list network-zone zone1
```

If the user tries to create a network-zone that was already created, the following error message will be displayed on the console, and the command will be rejected.

```
Switch#create access-list network-zone zone1  
Error: Network Zone "zone1" already exists.
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on all platforms.

create flow-redirect

```
create flow-redirect flow_redirect_name
```

Description

Creates a named flow redirection policy.

Syntax Description

<i>flow_redirect_name</i>	Specifies the name of the flow redirection policy.
---------------------------	--

Default

N/A.

Usage Guidelines

Use this command to create a named flow redirection policy to which nexthop information can be added.

For name creation guidelines and a list of reserved names, see [Object Names](#) in the *ExtremeXOS Concepts Guide*.



Example

The following command creates a flow redirection policy names flow3:

```
create flow-redirect flow3
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on BlackDiamond 8000 c-, e-, xl-, and xm-series modules, E4G-200 and E4G-400 switches, Summit X460, X480, X670, and X770 series switches, and SummitStack.

delete access-list

```
delete access-list dynamic_rule
```

Description

Deletes a dynamic ACL.

Syntax Description

<i>dynamic_rule</i>	Specifies the dynamic ACL name.
---------------------	---------------------------------

Default

N/A.

Usage Guidelines

This command deletes a dynamic ACL rule. Before you delete a dynamic ACL, it must be removed from any interfaces it is applied to. Use the `configure access-list delete` command to remove the ACL from an interface.

Example

The following command deletes the dynamic ACL icmp-echo:

```
delete access-list icmp-echo
```



History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

delete access-list network-zone

delete access-list network-zone *zone_name*

Description

This command is used to delete a network-zone and all configurations that belong to that zone.

Syntax Description

<i>zone_name</i>	Network-zone name
------------------	-------------------

Default

N/A.

Usage Guidelines

Use this command to delete a network-zone and all configurations belonging to that zone.

Example

```
Switch# delete access-list network-zone zone1
```

If the user tries to delete a network-zone that is bound with one or more policy files, the following error message will be displayed, and the command will be rejected.

```
Switch # delete access-list network-zone zone1
Error: Network Zone "zone1" - Unable to delete zone. Zone has one
or more policies.
```

History

This command was first available in ExtremeXOS 15.2.



Platform Availability

This command is available on all platforms.

delete access-list zone

```
delete access-list zone name
```

Description

Deletes an ACL zone.

Syntax Description

<i>name</i>	Specifies the zone name.
-------------	--------------------------

Default

N/A.

Usage Guidelines

This command deletes an ACL zone. You must remove all applications from a zone before you can delete the zone. To delete an application from a zone, use the command `configure access-list zone name delete application appl-name`.

You cannot delete the default zones.

Example

The following command deletes the zone `my_zone`:

```
delete access-list zone my_zone
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

delete flow-redirect

```
delete flow-redirect flow_redirect_name
```



Description

Deletes the named flow redirection policy.

Syntax Description

<i>flow_redirect_name</i>	Specifies the name of the flow redirection policy.
---------------------------	--

Default

N/A.

Usage Guidelines

Use this command to delete a named flow-redirection policy. Before it can be deleted, all nexthop information must be deleted, otherwise an error message is displayed.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules, E4G-200 and E4G-400 switches, Summit X480, X670, and X770 series switches, and SummitStack.

disable access-list permit to-cpu

disable access-list permit to-cpu

Description

Allows special packets to be blocked by low priority ACLs.

Syntax Description

This command has no arguments or variables.

Default

Enabled.



Usage Guidelines

This command allows ACLs to deny certain special packets from reaching the CPU, even if the packets match ACLs that would otherwise deny them. The special packets include STP and EAPS BPDUs, and ARP replies for the switch.

When this feature is disabled, these same packets will be denied if an ACL is applied that contains a matching entry that denies the packets. Contrary to expectations, the packets will still be denied if there is a higher precedence entry that permits the packets.

To enable this feature, use the following command:

```
enable access-list permit to-cpu
```

Example

The following command enables ACLs to deny STP BPDU packets from reaching the switch CPU:

```
disable access-list permit to-cpu
```

History

This command was first available in ExtremeXOS 11.3.2.

Platform Availability

This command is available only on the BlackDiamond X8 series switches, BlackDiamond 8800 series switches, E4G-200 and E4G-400 switches and SummitStack switches.

disable access-list refresh blackhole

```
disable access-list refresh blackhole
```

Description

Disables blackholing of packets during ACL refresh.

Syntax Description

This command has no arguments or variables.

Default

The feature is enabled.



Usage Guidelines

When access control lists (ACLs) are refreshed, this feature provides that any packets arriving during the refresh will be blackholed.

If you disable this feature, the ACLs will be refreshed as described in the `refresh policy` command.

To enable this feature, use the following command:

```
enable access-list refresh blackhole
```

Example

The following command disables dropping of packets during an ACL refresh:

```
disable access-list refresh blackhole
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

enable access-list permit to-cpu

```
enable access-list permit to-cpu
```

Description

Enables control packets to reach CPU, even if an ACL would deny them.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command allows control packets to reach the CPU, even if the packets match ACLs that would otherwise deny them. The control packets include STP and EAPS BPDUs, and ARP replies for the switch.



If this feature is disabled, these same packets will be denied if an ACL is applied that contains a matching entry that denies the packets. Contrary to expectations, when this feature is disabled, the packets will still be denied if there is a higher precedence entry that permits the packets.

To disable this feature, use the following command:

```
disable access-list permit to-cpu
```

Example

The following command enables STP BPDU packets to reach the switch CPU, despite any ACL:

```
enable access-list permit to-cpu
```

History

This command was first available in ExtremeXOS 11.3.2.

Platform Availability

This command is available only on the BlackDiamond X8 series switches, BlackDiamond 8800 series switches, E4G-200 and E4G-400 switches and SummitStack switches.

enable access-list refresh blackhole

Enables blackholing of packets during ACL refresh.

```
enable access-list refresh blackhole
```

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

When access control lists (ACLs) are refreshed, this command provides that any packets arriving during the refresh will be blackholed. As the ACL is being refreshed, packets may arrive while the ACL is in an indeterminate state, and packets may be permitted that otherwise are dropped. This feature protects the switch during an ACL refresh.



To disable this feature, use the following command:

```
disable access-list refresh blackhole
```

Example

The following command enables dropping of packets during an ACL refresh:

```
enable access-list refresh blackhole
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

refresh access-list network-zone

```
refresh access-list network-zone [zone_name | all]
```

Description

This command is used to refresh a specific network zone, or all the network zones.

Syntax Description

network-zone	Specifies the logical group of remote devices.
<i>zone_name</i>	Specifies the network_zone name.
all	Refresh all the network-zones.

Default

N/A.

Usage Guidelines

Use this command to refresh a specific network zone, or all the network zones.

When you issue the command to refresh a network-zone, or all network-zones, it can take a long time to clear the CLI because each individual policy must be converted before it is refreshed. The command succeeds, or fails, only after it receives a response for all policy refresh results from the hardware.



If the refresh fails for a specific zone, the following error message will be printed on the console.

```
Switch # refresh access-list network-zone zone1
ERROR: Refresh failed for network-zone "zone1".
```

Example

The following example refreshes all policies in “zone1”:

```
refresh access-list network-zone zone1
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on all platforms.

show access-list

```
show access-list {any | ports port_list | vlan vlan_name} {ingress | egress}
```

Description

Displays the ACLs configured on an interface.

Syntax Description

<i>aclname</i>	Specifies the ACL name. The name can be from 1-32 characters long.
any	Specifies the wildcard ACL.
<i>port_list</i>	Specifies which ports' ACLs to display.
<i>vlan_name</i>	Specifies which VLAN's ACL to display.
ingress	Display ingress ACLs.
egress	Display egress ACLs (BlackDiamond X8 series switches, BlackDiamond 8000 c- and xl-series modules, E4G-200 and E4G-400 switches, and Summit X480 series switches only).

Default

The default is to display all interfaces, ingress.



Usage Guidelines

The ACL with the port and VLAN displayed as an asterisk (*) is the wildcard ACL.

If you do not specify an interface, the policy names for all the interfaces are displayed, except that dynamic ACL rule names are not displayed. To display dynamic ACLs use the following commands:

```
show access-list dynamic
```

```
show access-list dynamic rule rule {detail}
```

If you specify an interface, all the policy entries, and dynamic policy entries are displayed.

Example

The following command displays all the interfaces configured with an ACL:

```
show access-list
```

The output from this command is similar to:

Vlan Name	Port	Policy Name	Dir	Rules	Dyn Rules
*	3:6	TCP_flag	ingress	3	2
*	3:8	qos_hongkong	ingress	3	0
*	2:1	tc_2.4	ingress	4	0
*	2:7	tcp	ingress	1	0
v1	*	tcp	ingress	1	0
*	*	firewall1	ingress	2	1

The following command displays the ingress access list entries configured on the VLAN v1006:

```
show access-list v1006 ingress
```

The output from this command is similar to the following:

```
# RuleNo 1
entry dacl13 {          #Dynamic Entry
if match all {
ethernet-destination-address 00:01:05:00:00:00 ;
} then {
count c13 ;
redirect 1.1.5.100 ;
} }
# RuleNo 2
entry dacl14 {          #Dynamic Entry
if match all {
ethernet-source-address 00:01:05:00:00:00 ;
} then {
count c14 ;
qosprofile qp7 ;
```




```
  } }  
  # RuleNo 3  
  entry dacl13 {  
  if match all {  
  ethernet-destination-address 00:01:05:00:00:00 ;  
  } then {  
  count c13 ;  
  redirect 1.1.5.100 ;  
  } }  
}
```

History

This command was first available in ExtremeXOS 10.1.

The *aclname* option was removed in ExtremeXOS 11.1.

The ingress, egress, any, ports, and vlan options were added in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

The egress option is available on BlackDiamond X8 series switches, BlackDiamond 8000 c- xl-, and xm-series modules, E4G-200 and E4G-400 switches, and Summit X480, X670, and X770 series switches.

show access-list configuration

show access-list configuration

Description

Displays the ACL configuration.

Syntax Description

There are no arguments or variables for this command.

Default

N/A.

Usage Guidelines

This command displays the state of the ACL configuration, set by the following commands:

```
enable access-list refresh blackhole  
enable access-list permit to-cpu  
configure access-list rule-compression port-counters
```



```
configure access-list vlan-acl-precedence
```

Example

The following command displays the state of the ACL configuration:

```
show access-list configuration
```

On BlackDiamond X8 series switches, E4G-200 and E4G-400 switches, and BlackDiamond 8800 and Summit series switches, the output from this command is similar to the following:

```
Access-list Refresh Blackhole: Enabled
Access-list Permit To-CPU: Enabled
Access-list configured vlan-acl precedence mode: Dedicated or Shared
Access-list operational vlan-acl-precedence mode: Dedicated or Shared
Access-list Rule-compression Port-counters: Dedicated or Shared
```

History

This command was first available in ExtremeXOS 11.0.

The Access-list Permit to CPU configuration was added in ExtremeXOS 11.3.2.

The Access-list Rule-compression Port Counters configuration was added in ExtremeXOS 12.3.

The Access-list Configured VLAN-ACL Precedence Mode configuration was added in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.

show access-list counter

```
show access-list counter {countername} {any | ports port_list | vlan vlan_name}
{ingress | egress}
```

Description

Displays the specified access list counters.

Syntax Description

<i>countername</i>	Specifies the ACL counter to display.
<i>port_list</i>	Specifies to display the counters on these ports.
<i>vlan_name</i>	Specifies to display the counters on the VLAN.



ingress	Specifies to display ingress counters.
egress	Specifies to display egress counters (BlackDiamond X8 series switches, BlackDiamond 8000 c-, xl-, and xm-series modules, E4G-200 andn E4G-400 switches, and Summit X460, X480, X670, and X770 series switches only).

Default

The default direction is ingress.

Usage Guidelines

Use this command to display the ACL counters.

Example

The following example displays all the counters for all ACLs:

```
show access-list counter
```

On a BlackDiamond 8000 c-, e-, xl-, or xm-series module, the output of this command is similar to the following:

```

Policy Name      Vlan Name      Port  Direction
Counter Name    Packet Count   Byte Count
=====
don1             *              2:1 ingress
source1111
source2222
0

```

History

This command was first available in ExtremeXOS 10.1.

The egress option was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

The egress option is available on BlackDiamond X8 series switches, BlackDiamond 8000 c- xl-, and xm-series modules, E4G-200 and E4G-400 switches, and Summit X460, X480, X670, and X770 series switches only.

show access-list dynamic

```
show access-list dynamic
```



Description

Displays the names of existing dynamic ACLs and a count of how many times each is used.

Syntax Description

There are no arguments or variables for this command.

Default

N/A.

Usage Guidelines

This command displays the names of existing dynamic ACLs, and how many times the ACL is used (bound to an interface).

To see the conditions and actions for a dynamic ACL, use the following command:

```
show access-list dynamic rule rule {detail}
```

Example

The following command displays names of all the dynamic ACLs:

```
show access-list dynamic
```

The following is sample output for this command:

```
Dynamic Rules:
Udpacl                Bound to 1 interfaces
icmp-echo             Bound to 1 interfaces
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

show access-list dynamic counter

```
show access-list dynamic counter {{countername} any | {countername} ports
port_list | {countername} vlan vlan_name} {ingress | egress}
```



Description

Displays the dynamic ACL counters.

Syntax Description

<i>countername</i>	Display the counter.
any	Specifies the wildcard ACL.
<i>port_list</i>	Specifies which ports' ACLs to display.
<i>vlan_name</i>	Specifies which VLAN's ACL to display.
ingress	Display ingress ACLs.
egress	Display egress ACLs (BlackDiamond X8 series switches, BlackDiamond 8000 c- or xl-series modules, E4G-200 and E4G-400 switches, and Summit X480 series switches only).

Default

The default is to display all interfaces, ingress.

Usage Guidelines

None.

Example

The following command displays all the dynamic ACL counters:

```
show access-list dynamic counter
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

The egress option is available on BlackDiamond X8 series switches, BlackDiamond 8000 c- xl-, and xm-series modules, E4G-200 and E4G-400 switches, and Summit X460, X480, X670, and X770 series switches only.

show access-list dynamic rule

```
show access-list dynamic rule [rule | rule_li ] detail
```



Description

Displays the syntax of a dynamic ACL.

Syntax Description

<i>rule</i>	Specifies the rule to display.
<i>rule_li</i>	Specifies the dynamic rule name for Lawful Intercept account only. You must have lawful intercept user privileges to specify this variable.
detail	Specifies to display where the ACL has been applied.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the syntax of the dynamic ACL `udpacl`:

```
show access-list dynamic rule udpacl
```

The output of the command is similar to the following:

```
entry udpacl {
  if match all {
    source-address 10.203.134.0/24 ;
    destination-address 140.158.18.16/32 ;
    protocol udp ;
    source-port 190 ;
    destination-port 1200 - 1250 ;
  } then {
    permit ;
  } }
}
```

The following command displays where the dynamic ACL `udpacl` has been applied:

```
show access-list dynamic rule udpacl
```

The output of the command is similar to the following:

```
Rule udpacl has been applied to the following interfaces.
Vlan Name   Port   Direction
=====
*           1     ingress
```



The lawful intercept user can display the names of the existing dynamic ACLs and a count of how many times each is used when the following command is issued:

```
* (pacman debug) X460-24p.1 > show access-list dynamic
Dynamic Rules: ((*)- Rule is non-permanent )
(*)hclag_arp_0_4_96_51_fe_b2  Bound to 0 interfaces for application
HealthCheckLAG
(*)idmgmt_def_blacklist      Bound to 0 interfaces for application
IdentityManager
(*)idmgmt_def_whitelist     Bound to 0 interfaces for application
IdentityManager
(*)mirror-data              Bound to 2 interfaces for application CLI
```

Use the following command to see the conditions and actions for a dynamic ACL:

```
* (pacman debug) X460-24p.2 > show access-list dynamic rule "mirror-data"
entry mirror-data {
if match all {
    source-address 10.66.9.8/24 ;
    protocol udp ;
} then {
    permit ;
    mirror law_mirror ;
} }
```

History

This command was first available in ExtremeXOS 11.3.

The **detail** keyword was added in ExtremeXOS 11.4.

The *rule_li* variable was added in ExtremeXOS 15.3.2.

Platform Availability

This command is available on all platforms.

show access-list interface

```
show access-list {rule rule {start} } [ any | port port | vlan vlan_name ] {zone
zone_name { appl-name appl_name {priority number } } } {ingress | egress} {detail}
```

Description

Displays the specified ACL zones, including their priority, applications, and the application priorities.

Syntax Description

any	Displays all zones on the specified interface.
port port	Displays all ACLs associated with the specified ports.



vlan <i>vlan_name</i>	Displays all ACLs associated with the specified VLAN.
<i>zone_name</i>	Specifies a zone to be displayed.
appl-name <i>appl_name</i>	Displays information by application within a zone.
priority <i>number</i>	Displays ACLs of the specified priority only, within an application area.
ingress	Displays ACLs applied to traffic in the ingress direction.
egress	Displays ACLs applied to traffic in the egress direction.
detail	Displays all ACLs applied to the specified interface.

Default

N/A.

Usage Guidelines

Use this command to display the ACL zones, applications, and priorities.

Specifying a zone will show all the ACLs installed in the particular zone. Specifying a priority within a zone will show all the ACLs installed at a particular priority within a zone.

Use the detail keyword to display all ACLs installed on a given interface.

Example

The following example displays the detailed view of the ACLs on port 1:1:

```
show access-list port 1:1 detail
```

The output of this command is similar to the following:

```
* BD-PC.1 # show access-list port 1:1 detail
RuleNo      Application      Zone          Sub Zone
=====
      1      CLI           myZone       1
entry mac1 {
if match all {
ethernet-source-address 00:0c:29:e5:94:c1 ;
destination-address 192.168.11.144/32 ;
} then {
count mac1 ;
} }
      2      CLI           myZone       5
entry mac51 {
if match all {
ethernet-source-address 00:0c:29:e5:94:51 ;
} then {
count mack51;
} }
      3      CLI           myZone       5
```




```

entry mac52 {
if match all {
ethernet-source-address 00:0c:29:e5:94:52 ;
} then {
count mac52 ;
} }

```

The following example displays the detailed view of the priority 5 ACLs in the zone myzone on port 1:1:

```

* BD-PC.2 # show access-list port 1:1 zone myZone priority 5 detail
RuleNo      Application      Zone      Sub Zone
=====
      2      CLI              myZone    5
  entry mac51 {
if match all {
ethernet-source-address 00:0c:29:e5:94:51 ;
} then {
count mack51;
} }
      3      CLI              myZone    5
  entry mac52 {
if match all {
ethernet-source-address 00:0c:29:e5:94:52 ;
} then {
count mac52 ;
} }

```

The following example displays the priority 5 ACLs in the zone myzone on port 1:1:

```

BD-PC.2 # show access-list port 1:1 zone myZone priority 5
#Dynamic Entries ((*)- Rule is non-permanent )
RuleNo      Name                      Application      Zone
Sub-Zone
1           mac51                      CLI              myZone          5
2           mac52                      CLI              myZone          5

```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

show access-list network-zone

```
show access-list network-zone {zone_name}
```



Description

Displays the network-zones configured, the number of attributes configured, and the number of policy files that have the specified zones in it.

Syntax Description

network-zone	Specifies the logical group of remote devices.
<i>zone_name</i>	Specifies the network-zone name.

Default

N/A.

Usage Guidelines

Use this command to display detailed information about a particular network-zone, the attributes configured in the zone, and the policies bound to the zone.

Example

The following example displays network-zone statistics for all configured zones:

```
Switch # sh access-list network-zone
=====
Network Zone                No. of    No. of Policies
Entities      Bound
=====
zone1                5          2
zone2                3          1
zone3                0          0
=====
Total Network Zones : 3
```

This example displays statistics for the specified zones, “zone1”, and “zone2”:

```
Switch #show access-list network-zone zone1
Network-zone      : zone1
Total Attributes : 3
Attributes       : 10.1.1.1 / 32
10.1.1.1 / 30
10.1.1.0 / 24
No. of Policies  : 1
Policies        : test
Switch # sh access-list network-zone zone2
Network-zone      : zone2
No. of Entities  : 3
Entities         : 00:00:00:00:00:22 / ff:ff:ff:ff:ff:ff
00:00:00:00:00:23 / ff:ff:ff:ff:00:00
```



```
00:00:00:00:00:24 / ff:ff:ff:ff:ff:00  
No. of Policies : 0
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on all platforms.

show access-list usage acl-mask port

```
show access-list usage acl-mask port port
```

Description

Displays the number of ACL masks consumed by the ACLs on a particular port.

Syntax Description

<i>port</i>	Displays the usage on the specified port.
-------------	---

Default

N/A.

Usage Guidelines

Use this command to display how many masks are currently consumed on a port.

Example

The following example displays the ACL mask usage on port 1:1:

```
Switch.8 # show access-list usage acl-mask port 1:1  
Used: 3 Available: 12
```

History

This command was first available in ExtremeXOS 11.4.



Platform Availability

This command is available on the BlackDiamond X8 series switches, BlackDiamond 8800 series switches, E4G-200 and E4G-400 switches, SummitStack, and Summit family switches only.

show access-list usage acl-range port

```
show access-list usage acl-range port port
```

Description

Displays the number of Layer 4 port ranges consumed by the ACLs on the slices that support a particular port.

Syntax Description

<i>port</i>	Specifies to display the usage for the slices that support this port.
-------------	---

Default

N/A.

Usage Guidelines

The BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, and xl-series modules, E4G-200 and E4G-400 switches, SummitStack, and Summit family switches can support a total of 16 Layer4 port ranges among the slices that support each group of 24 ports.

Use this command to display how many of these Layer4 ranges are currently consumed by the ACLs on the slices that support a particular port. The output of this command also displays which ports share the same slices as the specified port.

Example

The following example displays the Layer4 range usage on port 9:1:

```
Switch.3 # show access-list usage acl-range port 9:1
Ports 9:1-9:12, 9:25-9:36
L4 Port Ranges:  Used: 4  Available: 12
```

History

This command was first available in ExtremeXOS 11.5.



Platform Availability

This command is available only on BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules, E4G-200 and E4G-400 switches, SummitStack, and Summit family switches.

show access-list usage acl-rule port

```
show access-list usage acl-rule port port
```

Description

Displays the number of ACL rules consumed by the ACLs on a particular port or on the slices that support a particular port.

Syntax Description

<i>port</i>	Specifies to display the usage on this port.
-------------	--

Default

N/A.

Usage Guidelines

Use this command to display the rules used per slice, and also display the rule usage of the specified port.

The slice support for the BlackDiamond X8 series switches, BlackDiamond 8000 series modules, E4G-200 and E4G-400 switches, and Summit family switches that use this mechanism is as follows:

- BlackDiamond 8800 e-series modules—Each group of 24 ports has 8 slices with each slice having enough memory for 128 ingress rules and actions.
- Summit X460 series switches and E4G-200 and E4G-400 switches—
 - Each group of 24 ports has 4 slices with each slice having enough memory for 128 egress rules.
 - Each group of 24 ports has 16 slices with each slice having enough memory for 256 ingress rules.
- Summit X480 series switches—
 - Each group of 48 ports has 4 slices with each slice having enough memory for 256 egress rules.
 - Each group of 48 ports has 16 internal slices with each slice having enough memory for 512 ingress rules plus the external slice.
- Summit X670 switches and BlackDiamond X8 series switches—
 - Each group of 48 ports has 4 slices with each slice having enough memory for 256 egress rules.
 - Each group of 48 ports has 10 slices; the first 4 (0-3) slices hold 128 ingress rules each, and the last 6 (4-9) slices hold 256 ingress rules each, which adds up to 2048 ingress rules.
- BlackDiamond 8000 c- and xl-series modules—
 - 10G1Xc—

Its single port has 4 slices with each slice having enough memory for 128 egress rules.



-
- Its single port has 16 slices with each slice having enough memory for 256 ingress rules.
 - G8Xc—
 - Its 8 ports have 4 slices with each slice having enough memory for 128 egress rules.
 - Its 8 ports have 16 slices with each slice having enough memory for 256 ingress rules.
 - 10G4Xc/10G8Xc—
 - Each group of 2 ports has 4 slices with each slice having enough memory for 128 egress rules.
 - Each group of 2 ports has 16 slices with each slice having enough memory for 256 ingress rules.
 - 10G24X-c—
 - Each group of 12 ports has 4 slices with each slice having enough memory for 128 egress rules.
 - Each group of 12 ports has 12 slices with each of the first 8 slices having enough memory for 128 ingress rules and each of the last 4 slices having enough memory for 256 ingress rules, which adds up to 2048 ingress rules.
 - G96T-c—
 - Each group of 48 ports has 4 slices with each slice having enough memory for 256 egress rules.
 - Each group of 48 ports has 16 slices with each slice having enough memory for 512 ingress rules.
 - G48Tc/G48Xc/G24Xc—
 - Each group of 24 ports has 4 slices with each slice having enough memory for 128 egress rules.
 - Each group of 24 ports has 16 slices with each slice having enough memory for 256 ingress rules.
 - G48X-xl/G48T-xl—
 - Its 48 ports have 4 slices with each slice having enough memory for 256 egress rules.
 - Its 48 ports have 16 slices with each slice having enough memory for 512 ingress rules.
 - 10G8X-xl—
 - Each group of 4 ports has 4 slices with each slice having enough memory for 256 egress rules.
 - Each group of 4 ports has 16 slices with each slice having enough memory for 512 ingress rules.
 - 40G6X-xm—
 - Each group of 24 ports has 4 slices with each slice having enough memory for 256 egress rules.
 - Each group of 24 ports has 10 slices with each slice having enough memory for 256 ingress rules.

Note



Egress ACLs are supported on BlackDiamond X8 series switches, BlackDiamond 8000 c- xl-, and xm-series modules, E4G-200 and E4G-400 switches, and Summit X460, X480, and X670 series switches only.



Example

The following example displays the ACL rule usage on port 5:

```
Switch.3 # show access-list usage acl-rule port 5
Ports 1-12, 25-36
Total Rules:      Used: 34  Available: 990
```

The following example displays the ACL ingress and egress rule usage on port 5:1:

```
Switch.4 # show access-list usage acl-rule port 5:1
Ports 5:1-5:48
Total Ingress/Egress Rules:
Used: 11  Available: 8181
Used: 1  Available: 1023
```

History

This command was first available in ExtremeXOS 11.4.

This command was modified to support BlackDiamond 8000 e-series modules in ExtremeXOS 11.5.

Platform Availability

This command is available only on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, E4G-200 and E4G-400 switches, and Summit family switches, whether or not included in a SummitStack.

show access-list usage acl-slice port

```
show access-list usage acl-slice port port
```

Description

Displays the number of ACL slices and rules consumed by the ACLs on the slices that support a particular port.

Syntax Description

<i>port</i>	Specifies to display the usage for the slices that support this port.
-------------	---

Default

N/A.



Usage Guidelines

Use this command to display how many slices and how many rules per each slice are currently consumed by the ACLs on the slices that support a particular port. This command also displays which ports share the same slices as the specified port.

The slice support for the BlackDiamond X8 series switches, BlackDiamond 8000 series modules, E4G-200 and E4G-400 switches, and Summit family switches that use this mechanism is as follows:

- BlackDiamond 8800 e-series modules—Each group of 24 ports has 8 slices with each slice having enough memory for 128 ingress rules and actions.
- Summit X460 series switches and E4G-200 and E4G-400 switches—
 - Each group of 24 ports has 4 slices with each slice having enough memory for 128 egress rules.
 - Each group of 24 ports has 16 slices with each slice having enough memory for 256 ingress rules.
- Summit X480 series switches—
 - Each group of 48 ports has 4 slices with each slice having enough memory for 256 egress rules.
 - Each group of 48 ports has 16 internal slices with each slice having enough memory for 512 ingress rules plus the external slice.
- Summit X670 switches and BlackDiamond X8 series switches—
 - Each group of 48 ports has 4 slices with each slice having enough memory for 256 egress rules.
 - Each group of 48 ports has 10 slices; the first 4 (0-3) slices hold 128 ingress rules each, and the last 6 (4-9) slices hold 256 ingress rules each, which adds up to 2048 ingress rules.
- BlackDiamond 8000 c- and xl-series modules—
 - 10G1Xc—
 - Its single port has 4 slices with each slice having enough memory for 128 egress rules.
 - Its single port has 16 slices with each slice having enough memory for 256 ingress rules.
 - G8Xc—
 - Its 8 ports have 4 slices with each slice having enough memory for 128 egress rules.
 - Its 8 ports have 16 slices with each slice having enough memory for 256 ingress rules.
 - 10G4Xc/10G8Xc—
 - Each group of 2 ports has 4 slices with each slice having enough memory for 128 egress rules.
 - Each group of 2 ports has 16 slices with each slice having enough memory for 256 ingress rules.
 - 10G24X-c—
 - Each group of 12 ports has 4 slices with each slice having enough memory for 128 egress rules.
 - Each group of 12 ports has 12 slices with each of the first 8 slices having enough memory for 128 ingress rules and each of the last 4 slices having enough memory for 256 ingress rules, which adds up to 2048 ingress rules.
 - G96T-c—
 - Each group of 48 ports has 4 slices with each slice having enough memory for 256 egress rules.
 - Each group of 48 ports has 16 slices with each slice having enough memory for 512 ingress rules.
 - G48Tc/G48Xc/G24Xc—



Each group of 24 ports has 4 slices with each slice having enough memory for 128 egress rules.

Each group of 24 ports has 16 slices with each slice having enough memory for 256 ingress rules.

- G48X-xl/G48T-xl—

Its 48 ports have 4 slices with each slice having enough memory for 256 egress rules.

Its 48 ports have 16 slices with each slice having enough memory for 512 ingress rules.

- 10G8X-xl—

Each group of 4 ports has 4 slices with each slice having enough memory for 256 egress rules.

Each group of 4 ports has 16 slices with each slice having enough memory for 512 ingress rules.

- 40G6X-xm—

Each group of 24 ports has 4 slices with each slice having enough memory for 256 egress rules.

Each group of 24 ports has 10 slices with each slice having enough memory for 256 ingress rules.



Note

Egress ACLs are supported on BlackDiamond X8 series switches, BlackDiamond 8000 c, xl-, and xm-series modules, E4G-200 and E4G-400 switches, and Summit X460, X480, and X670 series switches only.

Beginning with ExtremeXOS 12.5, you can reserve or allocate a slice for a specific feature such that rules for the feature will not share a slice with other components. A text string has been added at the end of the output for each slice that indicates which feature, if any, is reserving the slice. See the example below.

Example

The following example displays the ACL slice usage on port 8:1:

```
Switch.3 # show access-list usage acl-slice port 8:1
Ports 8:1-8:12, 8:25-8:36
Slices:           Used: 3 Available: 5
Slice 5 Rules:    Used: 9 Available: 119
Slice 6 Rules:    Used: 1 Available: 127
Slice 7 Rules:    Used: 24 Available: 104
```

The following example displays the ACL ingress and egress slice usage on port 4:1:

```
Switch.4 # show access-list usage acl-slice port 4:1
Ports 4:1-4:48
Stage: INGRESS
Slices:           Used: 2 Available: 14
Slice 0 Rules:    Used: 0 Available: 512
Slice 1 Rules:    Used: 0 Available: 512
```



```

Slice 2 Rules:   Used: 0   Available: 512
Slice 3 Rules:   Used: 0   Available: 512
Slice 4 Rules:   Used: 0   Available: 512
Slice 5 Rules:   Used: 0   Available: 512
Slice 6 Rules:   Used: 0   Available: 512
Slice 7 Rules:   Used: 0   Available: 512
Slice 8 Rules:   Used: 0   Available: 512
Slice 9 Rules:   Used: 0   Available: 512
Slice 10 Rules:  Used: 0   Available: 512
Slice 11 Rules:  Used: 0   Available: 512
Slice 12 Rules:  Used: 0   Available: 512
Slice 13 Rules:  Used: 0   Available: 512 For: user/other
Slice 14 Rules:  Used: 1   Available: 511 Reserved for: <feature name>
Slice 15 Rules:  Used: 10  Available: 502 For: system
Stage: EGRESS
Slices:          Used: 1   Available: 3
Slice 0 Rules:   Used: 0   Available: 256
Slice 1 Rules:   Used: 0   Available: 256
Slice 2 Rules:   Used: 0   Available: 256
Slice 3 Rules:   Used: 1   Available: 255 Reserved for: <feature name>
Stage: LOOKUP
Slices:          Used: 1   Available: 3
Slice 0 Rules:   Used: 0   Available: 512
Slice 1 Rules:   Used: 0   Available: 512
Slice 2 Rules:   Used: 0   Available: 512
Slice 3 Rules:   Used: 49  Available: 463
Stage: EXTERNAL
Slices:          Used: 0   Available: 0

```

In this example, selected slices are allocated or reserved as follows:

- For: user/other—The slice is used by user ACLs and/or other switch features.
- Reserved for: *feature name*—The slice is reserved for the named feature, for instance VLAN statistics. Rules for this feature may not share a slice with other features or user ACLs.
- For: system—The slice contains only rules used for certain specific switch features. User ACLs may not share a slice with a system slice.

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available only on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, E4G-200 and E4G-400 switches, and Summit family switches, whether or not included in a SummitStack.

show access-list width

```
show access-list width [slot slotNo | all]
```



Description

Displays the wide ACL mode configured on the supported switch or slot.

Syntax Description

<i>slotNo</i>	Specifies the slot to display.
all	Specifies all slots.

Default

N/A.

Usage Guidelines

Use this feature to display the width of the ACL TCAM key configured on a module or switch as being double wide or single wide.

Example

The following command displays the wide key mode on all slots:

```
show access-list width slot all
```

Following is sample output for this command:

```
show access-list width {slot <slotNo>|all}
Slot  Type                               Width (Configured)
----  -
1     G48Ta                                   single
2     8900-G96T-c                             single          single
3     G8X                                     single
4
5     G48Xc                                   single
6     8900-10G8X-x1                           double          double
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on the BlackDiamond X8 series switches, BlackDiamond 8000 c-, xl-, and xm-series modules, E4G-200 and E4G-400 switches, and Summit X460, X480, and X670 switches only whether or not included in a SummitStack.



show flow-redirect

```
show flow-redirect {flow_redirect_name}
```

Description

Displays nexthop ipaddresses, up/down status, health-checking (ping/ARP/ND) and ACL bindings.

Syntax Description

<i>flow_redirect_name</i>	Specifies the name of the flow redirection policy.
---------------------------	--

Default

N/A.

Usage Guidelines

None.

Example

The following example displays information for all redirection policies:

```
Switch.1 # show flow-redirect
Name          Nexthop  Active          VR Name      Inactive  Health
Count        IP Address          Nexthops    Check
=====
pkh           2          2001:400::100  VR-Default  Forward  PING
ND: Neighbor Discovery
```

The next example displays an IPv6 redirection policy for a longer IPv6 address, which causes a two-line display for the related redirection policy:

```
Switch.13 # sh flow-redirect
Name          Nexthop  Active          VR Name      Inactive  Health
Count        IP address          Nexthops    Check
=====
pbr1          2          2004:1000:1000:1000::10
VR-Default  Forward  PING
ND: Neighbor Discovery
```

This example displays information for a specified IPv6 redirection policy:

```
* Switch.14 # show flow-redirect "pbr1"
Name          : pbr1          VR Name      : VR-Default
Inactive Nexthops: Forward    Health Check : PING
Nexthop Count  : 2
```



```

Active IP Address : 2004:1000:1000:1000::10
Index      State      Priority  IP Address      Status Interval Miss
=====
0          Disabled  200      2003::10       DOWN    2          2
1          Enabled   100      2004:1000:1000:1000::10
UP        2          2

```

History

This command was first available in ExtremeXOS 12.1.

Support for IPv6 flow-redirection policies was added in ExtremeXOS 12.7.

Platform Availability

This command is available for IPv4 and IPv6 flow-redirection policies on the platforms listed for the Policy Based Routing feature in [Feature License Requirements](#).

unconfigure access-list

```

unconfigure access-list policy_name {any | ports port_list | vlan vlan_name}
{ingress | egress}

```

Description

Removes a policy file ACL from the specified interface.

Syntax Description

<i>policy_name</i>	Specifies the ACL policy name. The name can be from 1-32 characters long.
<i>aclname</i>	Specifies the ACL name.
<i>port_list</i>	Specifies the ingress or egress port list on which the ACL is applied.
<i>vlan_name</i>	Specifies the VLAN on which the ACL is applied.
ingress	Remove the ACL for packets entering the switch on this interface.
egress	Remove the ACL for packets leaving the switch from this interface (BlackDiamond X8 series switches, BlackDiamond 8000 c- and xl-series modules, E4G-200 and E4G-400 switches, and Summit X460, X480, and X670 series switches only).

Default

The default direction is ingress.



Usage Guidelines

This command removes ACLs that are contained in ACL policy files. To remove dynamic ACLs, use the following command:

```
configure access-list delete ruleName [ any | vlanvlan_name | portsport_list | all ] {ingress | egress}
```

To remove all non-dynamic ACLs from all interfaces, do not specify any ports or VLANs.

Example

The following command removes the ACL from port 1:2:

```
unconfigure access-list ports 1:2
```

The following command removes the ACLs from ports 1:2-6:3 and 7:1:

```
unconfigure access-list ports 1:2-6:3,7:1
```

The following command removes the wildcard ACL:

```
unconfigure access-list any
```

The following command removes all ACLs from all the interfaces, including the wildcard ACL:

```
unconfigure access-list
```

History

This command was first available in ExtremeXOS 10.1.

The VLAN option was first available in ExtremeXOS 11.0.

The egress option was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

The egress options are available on BlackDiamond X8 series switches, BlackDiamond 8000 c-, xl-, xm-series modules, E4G-200 and E4G-400 switches, and Summit X460, X480, and X670 series switches only.



23 QoS Commands

```
clear counters wred
configure diffserv examination code-point qosprofile
configure diffserv replacement code-point
configure dot1p type
configure meter
configure port shared-packet-buffer
configure ports qosprofile
configure ports rate-limit egress
configure qosprofile
configure qosprofile qp8 weight
configure qosprofile wred
configure qosscheduler weighted-deficit-round-robin
create meter
create qosprofile
delete meter
delete qosprofile
disable diffserv examination ports
disable diffserv replacement ports
disable dot1p examination ports
disable dot1p replacement ports
enable diffserv examination ports
enable diffserv replacement ports
enable dot1p examination ports
enable dot1p replacement ports
show access-list meter
show diffserv examination
show diffserv replacement
show dot1p
show meter
show ports congestion
show ports qosmonitor
show ports qosmonitor {congestion}
show ports wred
show qosprofile
show wredprofile
unconfigure diffserv examination
unconfigure diffserv replacement
```

unconfigure qosprofile unconfigure qosprofile wred

This chapter describes commands for:

- Configuring Quality of Service (QoS) profiles.
- Creating traffic groupings and assigning the groups to QoS profiles.
- Configuring, enabling, and disabling explicit class-of-service traffic groupings (802.1p and DiffServ).
- Configuring traffic grouping priorities.
- Metering using ACLs—BlackDiamond 8800 series switches, SummitStack, and Summit family switches only.
- Verifying configuration and performance.
- Egress traffic rate limiting—BlackDiamond X8 series switches, BlackDiamond 8800 series switches, E4G-200 and E4G-400 switches, SummitStack, and Summit family switches only.

For an introduction to QoS features, see the *ExtremeXOS Concepts Guide*.

clear counters wred

clear counters wred

Description

Clears weighted random early detection (WRED) statistics for all ports.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following example clears the WRED statistics for all ports:

```
* Switch.20 # clear counters wred
```

History

This command was first available in ExtremeXOS 12.7.



Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8900 c-, xl-, and xm-series modules, E4G-200 and E4G-400 switches, and Summit X460, X480, and X670 switches.

configure diffserv examination code-point qosprofile

The syntax for BlackDiamond X8 series switches, BlackDiamond 8800, E4G-200 and E4G-400 switches, and Summit family switches is:

```
configure diffserv examination code-point code_point {qosprofile} qosprofile
```

Description

Configures the default ingress DiffServ code point (DSCP) to QoS profile mapping.

Syntax Description

code-point	Specifies a DiffServ code point (a 6-bit value in the IP-TOS byte in the IP header). Supported values are 0 to 63.
<i>qosprofile</i>	Specifies the QoS profile to which the DiffServ code point is mapped.

Default

See Table 28: Default DiffServ Code Point-to-QoS Profile Mapping on page 1451 below.

Usage Guidelines

You can specify up to 64 different code points for each port. Code point values are grouped and assigned to the default QoS profiles as shown in the following table.

Table 28: Default DiffServ Code Point-to-QoS Profile Mapping

Code Point	BlackDiamond X8 Series Switches, BlackDiamond 8800 Series Switches, E4G-200 and E4G-400 Switches, SummitStack, and Summit Family Switches QoSProfile
0-7	QP1
8-15	QP1
16-23	QP1
24-31	QP1
32-39	QP1
40-47	QP1
48-55	QP1
56-63	QP8



Example

The following command specifies that code point 25 be assigned to QP2:

```
configure diffserv examination code-point 25 qosprofile qp2
```

History

This command was first available in ExtremeXOS 11.0.

The ports keyword was first available in ExtremeXOS 12.2.2.

Platform Availability

This command is available on all platforms.

configure diffserv replacement code-point

The syntax for BlackDiamond X8 series switches, BlackDiamond 8800, E4G-200 and E4G-400 switches, SummitStack, and Summit family switches is:

```
configure diffserv replacement [{qosprofile} qosprofile | priority priority]  
code-point code_point
```

Description

Configures the egress Diffserv replacement mapping for either a QoS profile or an 802.1p priority value.

Syntax Description

<i>qosprofile</i>	Specifies a QoS profile.
value	Specifies an 802.1p priority value to map to a code point.
<i>code_point</i>	Specifies a 6-bit value to be used as the replacement DSCP in the IPv4 or IPv6 header.

Default

N/A.

Usage Guidelines



Note

We recommend that you use the qosprofile *qosprofile* value to configure this parameter.



Egress packets contain the DSCP assigned to the QoS profile, which can be selected by the 802.1p code point or by an ACL. The default 802.1p priority value to QoS profile to DSCP mapping is shown in the following table.

Table 29: Default QoS Profile-to-802.1p Priority Value-to-Code Point

802.1p Priority Value	BlackDiamond X8 Series Switches, BlackDiamond 8800 Series Switches, E4G-200 and E4G-400 Switches, SummitStack, and Summit Family Switches QoS Profile	DSCP
0	QP1	0
1	QP1	8
2	QP1	16
3	QP1	24
4	QP1	32
5	QP1	40
6	QP1	48
7	QP8	56

Example

The following command specifies that a code point value of 5 should be used to replace the DiffServ (TOS) bits in packets in QP2:

```
configure diffserv replacement qosprofile qp2 code-point 5
```

History

This command was first available in ExtremeXOS 11.0.

The ports keyword was first available in ExtremeXOS 12.2.2.

Platform Availability

This command is available on all platforms.

configure dot1p type

The syntax for BlackDiamond X8 series switches, BlackDiamond 8800, E4G-200 and E4G-400 switches, SummitStack, and Summit family switches is:

```
configure dot1p type dot1p_priority {qosprofile} qosprofile
```



Description

Configures an 802.1p priority to QoS profile mapping for the specified ports.

Syntax Description

<i>dot1p_priority</i>	Specifies the 802.1p priority value. The value is an integer between 0 and 7.
<i>qosprofile</i>	Specifies a specific QoS profile. The value range is QP1 to QP8.

Default

The default mapping of each 802.1p priority value to QoS profile is shown in the following table.

Table 30: Default 802.1p Priority Value-to-QoS Profile Mapping

802.1p Priority Value	BlackDiamond X8 Series Switches, BlackDiamond 8800 Series Switches, E4G-200 and E4G-400 Switches, SummitStack, and Summit Family Switches Default QoS Profile
0	QP1
1	QP1
2	QP1
3	QP1
4	QP1
5	QP1
6	QP1
7	QP8

Usage Guidelines

An 802.1p priority value seen on ingress can be mapped to a particular QoS profile and with specific bandwidth management and priority behavior.

BlackDiamond X8 Series Switches, BlackDiamond 8800 Series Switches, E4G-200 and E4G-400 Switches, SummitStack, and Summit Family Switches Only

You must create the QoS profile first, using the `create qosprofile [QP2 | QP3 | QP4 | QP5 | QP6 | QP7]` command, to map the 802.1p information to QoS profile 2 through 7.

SummitStack only.

You must create the QoS profile first, using the `create qosprofile [QP2 | QP3 | QP4 | QP5 | QP6 | QP7]`, to map the 802.1p information to QoS profile 2 through 6. You cannot create QP7 in a SummitStack.



Example

The following commands reassign (from the default) the QoS profiles associated with 802.1p priority values 1 and 2:

```
configure dot1p type 2 qosprofile qp2
configure dot1p type 1 qosprofile qp3
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure meter

```
configure meter metername {committed-rate cir [Gbps | Mbps | Kbps]} {max-burst-size burst-size [Kb | Mb]} {out-actions [drop | set-drop-precedence {dscp [none | dscp-value]}]}
```

Description

Configures an ACL meter to provide ingress traffic rate shaping on BlackDiamond X8 series switches, BlackDiamond 8800 series switches, E4G-200 and E4G-400 switches, SummitStack, and Summit family switches. On BlackDiamond c-, xl-, and xm-series modules, you can use this command to configure meters for ingress and egress rate limiting.

Syntax Description

<i>metername</i>	Specifies the ACL meter name.
committed-rate	Specifies the committed information rate in gigabits per second (Gbps), megabits per second (Mbps), or kilobits per second (Kbps).
max-burst-size	Specifies the maximum burst size or peak burst size in kilobits (Kb) or megabits (Mb) on BlackDiamond X8 series switches, BlackDiamond 8800 series switches, E4G-200 and E4G-400 switches, SummitStack, and Summit family switches.
out-actions	Specifies actions to take if traffic exceeds the profile.
drop	Specifies to drop out of profile traffic on BlackDiamond X8 series switches, BlackDiamond 8800 series switches, E4G-200 and E4G-400 switches, SummitStack, and Summit family switches.
set-drop-precedence	Specifies to mark packet for high drop precedence on BlackDiamond X8 series switches, BlackDiamond 8800 series switches, E4G-200 and E4G-400 switches, SummitStack, and Summit family switches.



dscp	Specifies to set DSCP on BlackDiamond X8 series switches, BlackDiamond 8800 series switches, E4G-200 and E4G-400 switches, SummitStack, and Summit family switches.
none	Specifies to leave the DSCP value unchanged on BlackDiamond X8 series switches, BlackDiamond 8800 series switches, E4G-200 and E4G-400 switches, SummitStack, and Summit family switches.

Default

By default, a newly committed meter has no maximum burst size, no committed rate, and a default action of drop.

Usage Guidelines

The meter configured with this command is associated with an ACL rule by specifying the meter name using the meter action modifier within the rule.

The committed-rate keyword specifies the traffic rate allowed for this meter, and the configured rate operates as described in [Table 31: Rate Configuration Notes](#) on page 1456 below. The rate you specify is rounded up to the next granularity increment value. For example, if you configure a 1 Mbps committed rate for a platform with a 64Kbps granularity increment, this value falls between the increment values of 960 Kbps and 1024 Kbps, so the effective committed rate is set to 1024 Kbps. Also, note that some platforms listed below require an adjustment to the expected rate to calculate the configured rate.

Table 31: Rate Configuration Notes

Platform	Granularity	Notes
BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules, E4G-200 and E4G-400 switches, and Summit family switches	64Kbps	Specify the traffic rate in Kbps, Mbps, or Gbps. The range is 64Kbps to 1 Gbps for GE ports and 1Mbps to 10 Gbps for 10GE ports. Add 20 bytes per frame to the expected rate to determine the configured rate.

The max-burst-size keyword specifies the maximum number of consecutive bits that are allowed to be in-profile at wire-speed. The max-burst-size parameter can be specified in Kb, Mb, or Gb. The specified max-burst-size is rounded down to the nearest supported size. The max-burst-size range on BlackDiamond X8 series switches, BlackDiamond 8000 series switches, E4G-200 and E4G-400 switches, and Summit switches is 32Kb to 128Mb.

The keyword out-actions specifies the action that is taken when a packet is out-of-profile. The supported actions include dropping the packet, marking the drop precedence for the packet, or setting the DSCP value in the packet. The keyword drop indicates that any out-of-profile packet is immediately dropped. The keyword set-drop-precedence marks out-of-profile packets with high drop precedence. If the optional keyword set-dscp is specified, the DSCP value, as specified by the parameter *dscp-value*, is written into the out-of-profile packet. Setting the DSCP value to none leaves the DSCP value in the packet unchanged.



On BlackDiamond X8 series switches, BlackDiamond 8900 xm-series modules and Summit X670 series switches, the meters behave as follows:

- QP1-4 support one unicast queue and one multicast queue for each QoS profile. The metering configuration for each of these QoS profiles applies to both the unicast and the multicast traffic for the profile.
- Configuration of maximum bandwidth metering on QP5-8 causes the configuration of the maximum meter on the supporting multicast queue to be set to the maximum bandwidth configured on QP5-8.

Example

The following command configures the ACL meter `maximum_bandwidth`, assigns it a rate of 10 Mbps, and sets the out of profile action to drop:

```
configure meter maximum_bandwidth committed-rate 10 Mbps out-action drop
```

History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure port shared-packet-buffer

```
configure port port_list shared-packet-buffer [percent | default]
```

Description

Configures the maximum amount of the shared packet buffer to be used by the specified ports.



Syntax Description

<i>port_list</i>	Specifies a list of ports or slots and ports.
<i>percent</i>	Specifies the maximum portion of the shared packet buffer to allot. The range is 0 to 100 percent.



Note

On some platforms or I/O modules, the hardware provides a limited number of settings. In these cases, ranges of percentage values achieve the same setting.



Note

You can view the configured percentage value using the `show ports port-list info detail` command.



Note

You can view the effect of this command using the `show ports port-list buffer` command.

Default

Platform	Percentage
BlackDiamond 8000 e-series modules BlackDiamond 8800 c-series modules	25%
BlackDiamond 8900 xl- and xm-series modules BlackDiamond 8900-G96T-c modules E4G-200 and E4G-400 switches BlackDiamond X8 series switches Summit X460, X480 series switches	20%
BlackDiamond 8900-10G24X-c modules	50%

Usage Guidelines

It is possible to overcommit the shared packet buffer using this command.

Example

The following command sets the shared packet buffer for port 1:1 to 50%:

```
configure port 1:1 shared-packet-buffer 50
```

History

This command was first available in ExtremeXOS 12.5.



Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 c-, e- xl-, and xm-series modules, E4G-200 and E4G-400 switches, Summit family switches, and SummitStack.

configure ports qosprofile

```
configure ports port_list {qosprofile} qosprofile
```

Description

Creates a port-based traffic group, which configures one or more ingress ports to use a particular egress QoS profile.

Syntax Description

<i>port_list</i>	Specifies a list of ports or slots and ports.
<i>qosprofile</i>	Specifies a QoS profile.

Default

All ingress ports have the default qosprofile of QP1.

Usage Guidelines

This command assigns traffic ingressing the specified port to a specified egress QoS profile. Extreme switches support eight egress QoS profiles (QP1 to QP8) for each port. SummitStack does not permit configuration of QP7.

Example

The following command configures port 5 on slot 5 of a modular switch to use QoS profile QP3:

```
configure ports 5:5 qosprofile QP3
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.



configure ports rate-limit egress

```
configure ports port_list rate-limit egress [no-limit | cir-rate [Kbps | Mbps | Gbps] {max-burst-size burst-size [Kb | Mb]}]
```

Description

Configures an egress traffic rate limit for a port or groups of ports.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
no-limit	Specifies traffic be transmitted without limit; use to reconfigure or unconfigure previous rate-limiting parameters.
<i>cir-rate</i>	Specifies the desired rate limit in Kbps, Mbps, or Gbps.
max-burst-size	Specifies the maximum burst size or peak burst size in kilobits (Kb) or megabits (Mb).

Default

No-limit.

Usage Guidelines

Port speed limits the egress traffic, as follows:

- 1 Gbps port—64 Kbps increments.
- 10 Gbps port—1 Mbps increments.

If the specified egress limit (*cir-rate*) is not a multiple of 64 Kbps for a 1 Gbps port or 1 Mbps for a 10Gbps port, the specified value is rounded down to the nearest appropriate multiple based on the port type.

Use the **no-limit** parameter to:

- Unconfigure egress rate limiting on the port(s).
- Reconfigure existing egress rate limiting on the port(s).

The *max-burst-size* parameter is the amount of traffic above the value in the *cir-rate* parameter that is allowed to burst from the port(s) for a short duration.

Example

The following command configures egress rate-limiting on slot 3 port 1 on a modular switch for 3 Mbps and a maximum burst size of 5 M bits:

```
configure port 3:1 rate-limit egress 3 Mbps max-burst-size 5 Mb
```



History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure qosprofile

BlackDiamond X8 series switches, BlackDiamond 8800 series switches, E4G-200 and E4G-400 switches, SummitStack, and Summit family switches:

```
configure qosprofile egress qosprofile [{minbw minbw_number} {maxbw maxbw_number} | {peak_rate peak_bps [K | M]}] [ports [port_list | all]]
```

```
configure qosprofile qosprofile [{minbw minbw_number} {maxbw maxbw_number} | {{committed_rate committed_bps [K | M]} {peak_rate peak_bps [K | M]} | [ports [port_list | all]]]
```

```
configure {qosprofile} qosprofile [[{maxbuffer buffer_percentage} {use-strict-priority}] | [maxbuffer buffer_percentage ports [port_list | all]]]
```

Description

Modifies the default egress QoS profile parameters.

Syntax Description

minbw	The minimum bandwidth (minbw) option specifies the committed information rate as a percentage of the maximum port speed. The range is 0 to 100%, and the default value is 0. When autonegotiation is off, the CIR is the specified percentage of the configured port speed. When autonegotiation is on, the CIR is the specified percentage of the maximum port speed.
maxbw	The maximum bandwidth (maxbw) option specifies the peak rate as a percentage of the maximum port speed. The range is 0 to 100%, and the default value is 100. When autonegotiation is off, the peak rate is the specified percentage of the configured port speed. When autonegotiation is on, the peak rate is the specified percentage of the maximum port speed (the switch does not detect the negotiated port speed).
peak_rate	Specifies a peak rate in Kbps (k) bits or Mbps (m).
committed_rate	Specifies a committed information rate in Kbps (k) bits or Mbps (m).
<i>port_list</i>	Specifies a list of slots and ports to which the parameters apply. Specify ports in the following formats: 3-5, 2:5, 2:6-2:8.



<i>buffer_percentage</i>	<p>When used without a port-list, specifies the percentage of the total buffer you are reserving for this QoS profile on all ports for which an override has not been configured. The range is 1 to 100; the default setting is 100.</p> <p>When used with a port-list, specifies a percentage override of the maxbuffer setting for the QoS profile specified. The range is 1-10000; the default is 100 (i.e., no override). Setting 100% is equivalent to unconfiguring the maxbuffer override.</p> <p>Available only on BlackDiamond 8800 series chassis, Summit series switches and SummitStack.</p>
qosprofile	Specifies a QoS profile name.
use-strict-priority	<p>When the global qosscheduler configuration (configure qosscheduler command) is set to weighted-round-robin, this option overrides the global configuration for the specified QoS profile, so that it operates in strict-priority-mode. This enables hybrid strict-priority and weighted-round-robin scheduling operation. This option is available only on BlackDiamond 8800 series switches, Summit family switches, and SummitStack.</p>
weight-value	<p>Specifies the weight value used for queue service weighting in the weighted-round-robin scheduler for this QoS profile. Range is 1-15 or 1-127 depending on hardware type. 0=strict-priority. Default is 1.</p> <p>This command enables the user to input a weight for queues in the weighted-round-robin scheduler or weighted-deficit-round-robin scheduler. The weight of both WRR and WDRR algorithms have been extended to 1-127 on the following hardware: X480, X460, E4G-400, X440, E4G-200, 8900-G96T-c, 8900-10G24X-c, 8900-MSM128, 8900-G48T-xl, 8900-G48X-xl, 8900-10G8X-xl, X670, 8900-40G6X-xm, BDX-MM1, BDXA-FM960, BDXA-FM480, BDXA-40G24X, BDXA-40G12X, G48Ta, G48Xa, 10G4Xa, 10G4Ca, G48Te2, G24Xc, G48Xc, G48Tc, 10G4Xc, 10G8Xc, MSM-48, S-G8Xc, S-10G1Xc, 8500-G24X-e, 8500-G48T-e, S-10G2Xc</p> <p>This option is available only on BlackDiamond 8800 series switches, Summit family switches, and SummitStack.</p>
all	Specifies this applies to all ports on the device.

Default

- QoS profiles—QP1 and QP8 on SummitStack and Summit family switches; QP1 through QP8 on all other switches.
- Minimum bandwidth—0%
- Maximum bandwidth—100%
- Maximum buffer—100%
- Maxbuffer override—100% (no override)
- Weight—1
- Priority—By default, each qosprofile is assigned a different priority level:
 - QP1 - 1, Low (the lowest priority)
 - QP2 - 2, LowHi
 - QP3 - 3, Normal
 - QP4 - 4, NormalHi
 - QP5 - 5, Medium
 - QP6 - 6, MediumHi



- QP7 - 7, High
- QP8 - 8, HighHi (highest priority)

Usage Guidelines



Note

You can view the effect of setting the buffer-percentage using the `show ports port-list buffer` command.



Note

You can view the configured buffer-percentage value using the `show qosprofile` or `show qosprofile ports port-list` commands, respectively.

The maximum bandwidth value can be configured as either:

- an absolute percentage of the total maximum link speed, regardless of the currently configured or negotiated speed.

OR

- an absolute peak rate in Mbps or Kbps.

On BlackDiamond X8 series switches, BlackDiamond 8800, E4G-200 and E4G-400 switches, SummitStack, and Summit family switches, QoS profiles QP1 and QP8 are preconfigured. If you want to use a QoS profile in the range of QP2 through QP7, you must first create the QoS profile. QoS profile QP7 is reserved on SummitStack for stack management and cannot be created or modified.

When specified without a port-list, the maxbuffer parameter can configure a reduction in the maximum amount of packet buffer space allotted to the specified QoS profile. If you reduce the allotment below the default value of 100%, the reduction releases packet buffer space to the shared packet buffer. Regardless of the setting for this parameter, the system does not drop any packets as long as reserved packet buffer memory for the port and QoS profile or shared packet memory for the port remains available.



Note

The configuration defined by the maxbuffer attribute in this command can be overridden on a per-port basis if the port is specified along with the maxbuffer parameter.

When specified with a port-list, the maxbuffer setting overrides the system-wide reduction of packet buffer reservation set with the `configure qosprofile maxbuffer` command for the specified QoS profile. If the packet buffer reservation is reduced to 75 percent for the entire QoS profile, the specified ports are allotted 75% of the allotment for the specified QoS profile. If for specified ports the maxbuffer is set to 200 percent, the packet buffer reservation will be set to 200 percent of the normal packet buffer reservation for those ports, thus overriding the maxbuffer percentage set for the QoS profile.



Note

The packet buffer configuration feature is provided for expert users who fully understand the impact of buffer configuration changes. Improper buffer configuration can stop traffic flow through QoS profiles and ports for which no direct configuration change was made.



A range of ports has its own packet buffer pool. The maxbuffer override capability allows you to overcommit the packet buffer pool for the port range. When a packet buffer pool is overcommitted by more than 20%, the following message appears in the system log:

```
Warning: Packet memory is overcommitted by <percentage> for ports in
range <port-range>
```

It is also possible to configure maxbuffer overrides such that the size of the shared portion of the buffer pool is reduced to zero. If some port and QoS profile in the port range for that buffer pool does not have sufficient reserved packet memory to accommodate larger packets, it will be impossible for that port and QoS profile to transmit any packets of the larger size. In this case, the following message appears in the system log

```
: Warning: At least one port and QoS profile in port range <port-range>
cannot transmit packets larger than <packet-size> because of packet
memory configuration.
```

The weight-value parameter does not apply when the switch is configured for strict priority scheduling, which is the default configuration. To configure the type of scheduling you want to use for the entire switch, use the `configure qoscheduler [strict-priority | weighted-round-robin | weighted-deficit-round-robin]` command.

The weight-value parameter configures the relative weighting for each QoS profile. Because each QoS profile has a default weight of 1, all QoS profiles have equal weighting. If you configure a QoS profile with a weight of 4, that specified QoS profile is serviced 4 times as frequently as the remaining QoS profiles, which still have a weight of 1. If you configure all QoS profiles with a weight of 16, each QoS profile is serviced equally but for a longer period.

When the switch is configured for weighted-round-robin mode, the `use-strict-priority` option overrides the switch configuration for the specified QoS profile on all ports. Among QoS profiles configured with the `use-strict-priority` option, QoS profile QP8 has the highest priority and QP1 has the lowest priority. All strict-priority QoS profiles are serviced first according to their priority level, and then all other QoS profiles are serviced based on their configured weight.



Note

If you specify `use-strict-priority`, lower-priority queues and weighted-round-robin queues are not serviced at all as long as higher-priority queues have any remaining packets.

Example

On a BlackDiamond X8 series switch, BlackDiamond 8800 series switch, E4G-200 or E4G-400 switch, or a Summit family switch, the following command overrides the maximum buffer setting configured on QoS profile qp1 for port1:1:

```
configure qosprofile qp1 maxbuffer 75 port 1:1
```

History

This command was first available in ExtremeXOS 10.1.



Committed and peak rates were added in ExtremeXOS 11.0. Also in ExtremeXOS 11.0, ports were made mandatory.

Support for Summit X440, X460, X480, and X670 series switches and BlackDiamond 8000 c-, xl-, and xm-series modules was added in the respective platform introduction releases.

The use-strict-priority option was added in ExtremeXOS 12.3.

The ability to configure a maxbuffer override was added in ExtremeXOS 12.5.

Support for the BlackDiamond X8 series switches and E4G-200 and E4G-400 switches was added in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms with specific parameter exceptions as noted in the Syntax Description above.

configure qosprofile qp8 weight

```
configure qosprofile qp8 weight weight_value
```

Description

This command enables the user to input a weight value for queue service weighting in the weighted-round-robin scheduler or weighted-deficit-round-robin scheduler for this QoS profile. The weight value of both WRR and WDRR algorithms have been extended to 1-127 on this supported hardware (refer to the Concepts Guide for supported hardware).

Syntax Description

<i>weight_value</i>	Range is 1-15 or 1-127 depending on hardware type.
---------------------	--

Default

Strict priority.

Usage Guidelines

Use this command to input a weight value for queue service weighting in the weighted-round-robin scheduler or weighted-deficit-round-robin scheduler for this QoS profile. The weight value of both WRR and WDRR algorithms have been extended to 1-127 on this supported hardware (refer to the *ExtremeXOS Concepts Guide* for supported hardware).

Example

Example output not yet available and will be provided in a future release.



History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8900 c-, xl-, and xm-series modules, E4G-200 and E4G-400 switches, and Summit X460, X480, X670, and X770 switches.


configure qosprofile wred

```
configure {qosprofile} {egress} qosprofile [wred [{color [tcp [green | red] | non-tcp [any|red]]] [{min-threshold min_thresh} {max-threshold } {max-drop-rate max_drop_rate}]}] | avg-weight avg_weight]] ports [port_list | all]
```

Description

Configures WRED on the specified QoS profile for the specified port.

Syntax Description

egress	This optional parameter specifies an egress QoS profile.
<i>qosprofile</i>	Specifies a QoS profile name. Valid names are QP1 to QP8.
	<div style="border: 1px solid #ccc; padding: 5px; margin: 5px 0;">  <p>Note On Summit X670 and BlackDiamond X8 switches, QoS profiles QP5 to QP8 do not support WRED.</p> </div>
color	Specifies the WRED color to be configured. Valid colors for BlackDiamond 8900 c- and xl-series modules, E4G-200 and E4G-400 switches, and Summit X460 and X480 switches are: TCP green, TCP red, and non-TCP any. Valid colors for BlackDiamond X8 series switches, BlackDiamond 8900-40G6X-xm modules and Summit X670 switches are: TCP green, TCP red, non-TCP any, and non-TCP red.
green	Specifies that the WRED configuration applies to TCP traffic that is marked green.
non-tcp any	Specifies that the WRED configuration applies to any non-TCP traffic.
red	Specifies that the WRED configuration applies to TCP traffic that is marked red.
<i>min_thresh</i>	Specifies the minimum threshold for the specified WRED color. The range is 1 to 100 percent.
max_threshold	Specifies the maximum threshold for the specified WRED color. The range is 1 to 100 percent.
<i>max_drop_rate</i>	Specifies the maximum drop rate for the specified WRED color. The range is 1 to 100 percent.
<i>port_list</i>	Specifies a list of slots and ports to which the parameters apply. Specify ports in the following formats: 3-5, 2:5, 2:6-2:8.



non-tcp red	Specifies that the WRED configuration applies to non-TCP traffic that is marked red.
avg_weight	Specifies the weight constant for calculating the average queue size for the specified QoS profile. The range is 1 to 15.
all	Specifies that this command applies to all ports on the device.

Default

- Minimum threshold—100%
- Maximum threshold—100%
- Maximum drop rate—100%
- Average weight—4

Usage Guidelines

The *max_drop_rate*, *min_threshold*, and *max_threshold* parameters apply to the specified color. The *avg_weight* parameter applies to all colors on the specified QoS profile. Increasing the *avg_weight* value reduces the probability that traffic is dropped. Conversely, decreasing the *avg_weight* value increases the probability that traffic is dropped.

Example

The following example configures WRED settings for port 2:1, QoS profile qp3, color green:

```
* Switch.24 # configure qosprofile qp3 wred color tcp green min-threshold 80
max-threshold 95 max-drop-rate 75 ports 2:1
```

The following example configures the average weight for port 2:1, QoS profile qp2:

```
* Switch.26 # configure qosprofile qp2 wred avg-weight 4 ports 2:1
```

The following example configures WRED settings for non-TCP traffic on port 4, QoS profile qp3:

```
* Switch.2 # configure qosprofile qp3 wred color non-tcp any min-threshold 10
ports 4
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8900 c-, xl-, and xm-series modules, E4G-200 and E4G-400 switches, and Summit X460, X480, X670, and X770 switches.



configure qosscheduler weighted-deficit-round-robin

```
configure qosscheduler [strict-priority | weighted-round-robin | weighted-deficit-round-robin]
```

Description

This command specifies the scheduling algorithm that the switch uses to service QoS profiles on BlackDiamond X8, BlackDiamond 8K, and Summit platforms.

Syntax Description

strict-priority	Specifies the switch services the higher-priority QoS profiles first.
weighted-round-robin	Specifies the switch services all QoS profiles based on the configured weighting for each QoS profile.
weighted-deficit-round-robin	Allows you to use a credit-based algorithm in order to sample the size of the packet while scheduling various queues.

Default

Strict-priority.

Usage Guidelines

The configured QoS scheduling algorithm applies to all switch ports, but you can override this configuration for a QoS profile using the following command:

```
configure qosprofile qosprofile use-strict-priority
```

In strict-priority mode, QoS profile QP8 has the highest priority and QP1 has the lowest priority.

Note



Queues are serviced using the configured scheduling algorithm until all of the minBws are satisfied, then all queues are serviced using the configured scheduling algorithm until all of the maxBws are satisfied.

Example

The following command configures the switch for weighted-round-robin servicing:

```
configure qosscheduler weighted-round-robin
```

The following command configures the switch for weighted-deficit-round-robin servicing:

```
configure qosscheduler weighted-deficit-round-robin
```



This command specifies the scheduling algorithm the switch uses to service QoS profiles on BD8K, Summit and BDX8 platforms. Weighted-deficit-round-robin mode of scheduling allows you to use a credit based algorithm in order to sample in the size of the packet while scheduling various queues.

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on

- X480
- X460
- E4G-400
- X440
- E4G-200
- 8900-G96T-c
- 8900-10G24X-c
- 8900-MSM128
- 8900-G48T-xl
- 8900-G48X-xl
- 8900-10G8X-xl
- X670
- X770
- 8900-40G6X-xm
- BDX-MM1
- BDXA-FM960
- BDXA-FM480
- BDXA-40G24X
- BDXA-40G12X
- G48Ta
- G48Xa
- 10G4Xa
- 10G4Ca
- G48Te2
- G24Xc
- G48Xc
- G48Tc
- 10G4Xc
- 10G8Xc
- MSM-48
- S-G8Xc
- S-10G1Xc
- 8500-G24X-e



- 8500-G48T-e
- S-10G2Xc
- BlackDiamond X8 series switches
- BlackDiamond 8800 series switches
- E4G-200 and E4G-400 switches
- SummitStack and the Summit family switches.

create meter

create meter *meter-name*

Description

On BlackDiamond X8 series switches, BlackDiamond 8800 series switches, E4G-200 and E4G-400 switches, SummitStack, and Summit family switches, this command creates a meter for ingress traffic rate limiting. On BlackDiamond c-, xl-, and xm-series modules, you can use this command to create meters for ingress and egress rate limiting.

Syntax Description

<i>meter-name</i>	Specifies the meter name.
-------------------	---------------------------

Default

N/A.

Usage Guidelines

Meter names must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but they cannot contain spaces. The maximum allowed length for a name is 32 characters. For meter name guidelines and a list of reserved names, see [Object Names](#) in the *ExtremeXOS Concepts Guide*.

Example

The following command creates the meter `maximum_bandwidth`:

```
create meter maximum_bandwidth
```

History

This command was available in ExtremeXOS 11.1.



Platform Availability

This command is available on BlackDiamond X8 series, E4G-200 and E4G-400 switches, SummitStack, and Summit family switches.

create qosprofile

```
create qosprofile [QP2 | QP3 | QP4 | QP5 | QP6 | QP7]
```

Description

Creates a QoS profile.

Syntax Description

QP2...QP7	Specifies the QoS profile you want to create.
------------------	---

Default

N/A.

Usage Guidelines

The BlackDiamond X8 series switches, BlackDiamond 8800 series switches, E4G-200 and E4G-400 switches, SummitStack, and Summit family switches allow dynamic creation and deletion of QoS profiles QP2 to QP7. Creating a QoS profile dynamically does not cause loss of traffic.

QoS profiles QP1 and QP8 are part of the default configuration and cannot be deleted. You must create a QoS profile in the range of QP2 to QP7 before you can configure it or assign it to traffic groups.

QoS profile QP7 cannot be created in a SummitStack; this queue is reserved for control traffic.

Note



The sFlow application uses QP2 to sample traffic on SummitStack and Summit family switches; any traffic grouping using QP2 can encounter unexpected results when sFlow is enabled on these specific devices.

Example

The following command creates QoS profile QP3:

```
create qosprofile qp3
```

History

This command was first available in ExtremeXOS 11.1.



Platform Availability

This command is available on all platforms.

delete meter

```
delete meter meter-name
```

Description

Deletes a meter.

Syntax Description

<i>meter-name</i>	Specifies the meter name.
-------------------	---------------------------

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the meter `maximum_bandwidth`:

```
delete meter maximum_bandwidth
```

History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8800 series switches, E4G-200 and E4G-400 switches, SummitStack, and Summit family switches.

delete qosprofile

```
delete qosprofile [QP2 | QP3 | QP4 | QP5 | QP6 | QP7]
```



Description

Deletes a user-created QoS profile.

Syntax Description

QP2...QP7	Specifies the user-created QoS profile you want to delete.
------------------	--

Default

N/A.

Usage Guidelines

You cannot delete the default QoS profiles of QP1 and QP8. On a SummitStack, you also cannot delete QoS profile QP7. If you attempt to delete QoS profile QP7, the system returns an error.

All configuration information associated with the specified QoS profile is removed.

Example

The following command deletes the user-created QoS profile QP3:

```
delete qosprofile qp3
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

disable diffserv examination ports

```
disable diffserv examination ports [port_list | all]
```

Description

Disables the examination of the DiffServ field in an IP packet.

Syntax Description

<i>port_list</i>	Specifies a list of ports or slots and ports to which the parameters apply.
all	Specifies that DiffServ examination should be disabled for all ports.



Default

Disabled.

Usage Guidelines

The diffserv examination feature is disabled by default.

Example

The following command disables DiffServ examination on the specified ports:

```
disable diffserv examination ports 5:3,5:5,6:6
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable diffserv replacement ports

```
disable diffserv replacement ports [port_list | all]
```

Description

Disables the replacement of DiffServ code points in packets transmitted by the switch.

Syntax Description

<i>port_list</i>	Specifies a list of ports or slots and ports on which Diffserv replacement will be disabled.
all	Specifies that DiffServ replacement should be disabled for all ports.

Default

N/A.

Usage Guidelines

The DiffServ replacement feature is disabled by default.



Example

The following command disables DiffServ replacement on selected ports:

```
disable diffserv replacement ports 1:2,5:5,6:6
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

disable dot1p examination ports

```
disable dot1p examination ports [port_list | all]
```

Description

Prevents examination of the 802.1p priority field as part of the QoS configuration.

Syntax Description

<i>port_list</i>	Specifies a list of ports or slots and ports.
all	Specifies that dot1p replacement should be disabled for all ports.

Default

Enabled.

Usage Guidelines

The 802.1p examination feature is enabled by default. To free ACL resources, disable this feature whenever another QoS traffic grouping is configured. (See [ACLs](#) for information on available ACL resources.)



Note

If you disable this feature when no other QoS traffic grouping is in effect, 802.1p priority enforcement of 802.1q tagged packets continues.

SummitStack Only.

Dot1p examination cannot be disabled for priority values 5 and 6. However, the precedence of the examination is lowered so that all other traffic grouping precedences are higher. The mappings you configure with the `configure dot1p type` command remain in effect.



Example

The following command disables 802.1p value examination on ports 1 to 5:

```
disable dot1p examination ports 1-5
```

History

This command was available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

disable dot1p replacement ports

```
disable dot1p replacement ports [port_list | all]
```

Description

Disables the ability to overwrite 802.1p priority values for a given set of ports.

Syntax Description

<i>port_list</i>	Specifies a list of ports or slots and ports to which the parameters apply.
all	Specifies that 802.1p replacement should be disabled for all ports.

Default

N/A.

Usage Guidelines

The dot1p replacement feature is disabled by default.

Beginning with ExtremeXOS version 11.4 on the 1 Gigabit Ethernet ports, 802.1p replacement always happens when you configure the DiffServ traffic grouping.

Example

The following command disables 802.1p value replacement on all ports:

```
disable dot1p replacement ports all
```



History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

enable diffserv examination ports

```
enable diffserv examination ports [port_list | all]
```

Description

Enables the DiffServ field of an IP packet to be examined in order to select a QoS profile.

Syntax Description

<i>port_list</i>	Specifies a list of ports or slots and ports to which the parameters apply.
all	Specifies that DiffServ examination is enabled for all ports.

Default

Disabled.

Usage Guidelines

The Diffserv examination feature is disabled by default.

If you are using DiffServ for QoS parameters, we recommend that you also configure 802.1p or port-based QoS parameters to ensure that high-priority traffic is not dropped prior to reaching the MSM/MM on modular switches.

Example

The following command enables DiffServ examination on selected ports:

```
enable diffserv examination ports 1:1,5:5,6:2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.



enable diffserv replacement ports

```
enable diffserv replacement ports [port_list | all]
```

Description

Enables the DiffServ code point to be overwritten in IP packets transmitted by the switch.

Syntax Description

<i>port_list</i>	Specifies a list of ingress ports or slots and ports on which to enable DiffServ replacement.
all	Specifies that DiffServ replacement should be enabled for all ports.

Default

N/A.

Usage Guidelines

The Diffserv replacement feature functions for IPv4 and IPv6 traffic and is disabled by default.

Note



The port in this command is the ingress port.

This command affects only that traffic in traffic groupings based on explicit packet class of service information and physical/logical configuration.

Example

The following command enables DiffServ replacement on specified ports:

```
enable diffserv replacement ports 5:3,5:5,6:2
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

enable dot1p examination ports

```
enable dot1p examination ports [port_list | all]
```



Description

Enables egress QoS profile selection based on the 802.1p bits in the incoming frame.

Syntax Description

<code>port_list</code>	Specifies a list of ports on which to enable the dot1p examination feature.
<code>all</code>	Specifies that dot1p examination should be enabled for all ports.

Default

Enabled.

Usage Guidelines

To increase available ACLs, you can disable the 802.1p examination feature if you are not running QoS or are running QoS using DiffServ. See the *ExtremeXOS Concepts Guide* for information on ACL limitations on these platforms.

Use this command to re-enable the 802.1p examination feature.

Example

The following command enables dot1p examination on ports 1 to 5:

```
enable dot1p examination ports 1-5
```

History

This command was available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

enable dot1p replacement ports

```
enable dot1p replacement ports [port_list | all]
```

Description

Allows the 802.1p priority field to be overwritten on egress according to the QoS profile to 802.1p priority mapping for a given set of ports.



Syntax Description

<i>port_list</i>	Specifies a list of ports or slots and ports.
all	Specifies that dot1p replacement should be enabled for all ports.

Default

N/A.

Usage Guidelines

The dot1p replacement feature is disabled by default.

By default, 802.1p priority information is not replaced or manipulated, and the information observed on ingress is preserved when transmitting the packet.



Note

The port in this command is the ingress port.

If 802.1p replacement is enabled, the 802.1p priority information that is transmitted is determined by the hardware queue that is used when transmitting the packet.



Note

This command affects only that traffic in traffic groupings based on explicit packet class of service information and physical/logical configuration.

Beginning with ExtremeXOS version 11.4 on the 1 Gigabit Ethernet ports, 802.1p replacement always happens when you configure the DiffServ traffic grouping.



Note

Enabling dot1p replacement on all ports may take some time to complete.

Example

The following command enables dot1p replacement on all ports:

```
enable dot1p replacement ports all
```

History

This command was available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.



show access-list meter

```
show access-list meter {meter-name} [any | ports port_list | vlan vlan_name]
```

Description

Displays the specified access list meter statistics and configurations.

Syntax Description

<i>meter-name</i>	Specifies the ACL meter to display.
<i>port_list</i>	Specifies to display the meters on these ports.
<i>vlan_name</i>	Specifies to display the meters on the VLAN.

Default

N/A.

Usage Guidelines

Use this command to display the ACL meters.

Example

The following example displays access list meter information for port 7:1:

```
Switch.8 # show access-list meter mtr1 port 7:1
Policy Name      Vlan Name      Port
Committed      Committed Burst Peak Rate  Peak Burst
Out-of-Profile
Meter           Rate (Kbps) Size (Kb)      (Kbps)      Size(kb)      Packet
Count
=====
=
irl1            *              7:1
a.              mtr1           10            20           10           20           0
```

History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.



show diffserv examination

The syntax for BlackDiamond X8 series switches, BlackDiamond 8800, SummitStack, and Summit family switches is:

show diffserv examination

Description

Displays the DiffServ-to-QoS profile mapping.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

Once you alter the default mappings, the “->” in the display (shown below) becomes “* >”.

Example

Because the BlackDiamond X8 series switches, BlackDiamond 8800 series switches, SummitStack, and Summit family switches have 2 default QoS profiles, you see different displays depending on the platform.

The following is sample output from a BlackDiamond 8810 switch:

```
show diffserv examination
CodePoint->QOSProfile mapping:
00->QP1 01->QP1 02->QP1 03->QP1 04->QP1 05->QP1 06->QP1 07->QP1
08->QP1 09->QP1 10->QP1 11->QP1 12->QP1 13->QP1 14->QP1 15->QP1
16->QP1 17->QP1 18->QP1 19->QP1 20->QP1 21->QP1 22->QP1 23->QP1
24->QP1 25->QP1 26->QP1 27->QP1 28->QP1 29->QP1 30->QP1 31->QP1
32->QP1 33->QP1 34->QP1 35->QP1 36->QP1 37->QP1 38->QP1 39->QP1
40->QP1 41->QP1 42->QP1 43->QP1 44->QP1 45->QP1 46->QP1 47->QP1
48->QP1 49->QP1 50->QP1 51->QP1 52->QP1 53->QP1 54->QP1 55->QP1
56->QP8 57->QP8 58->QP8 59->QP8 60->QP8 61->QP8 62->QP8 63->QP8
```

History

This command was first available in ExtremeXOS 10.1.

The ports keyword was first available in ExtremeXOS 12.2.2.



Platform Availability

This command is available on all platforms.

show diffserv replacement

These values are placed in egress packets when DiffServ replacement is enabled.

The syntax for BlackDiamond X8 series switches, BlackDiamond 8800, SummitStack, and Summit family switches is:

```
show diffserv replacement
```

Description

Displays the DiffServ replacement code-point values assigned to each QoS profile.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

Once you alter the default mappings, the “->” in the display (shown below) becomes “* >”.

Example

The following is sample output from a BlackDiamond 8810 switch:

```
show diffserv replacement
QoSProfile->CodePoint mapping:
QP1->00
QP8->56
```

History

This command was first available in ExtremeXOS 10.1.

The ports keyword was first available in ExtremeXOS 12.2.2.

Platform Availability

This command is available on all platforms.



show dot1p

The syntax for BlackDiamond X8 series switches, BlackDiamond 8800, SummitStack, and Summit family switches is:

show dot1p

Description

Displays the 802.1p-to-QoS profile mappings.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

N/A.

Example

The BlackDiamond X8 series switches, BlackDiamond 8800 series switches, SummitStack, and Summit family switches have 2 default QoS profiles.

Following is sample output from the show dot1p command on the BlackDiamond 8810 switch:

```
show dot1p
802.1p Priority Value      QOS Profile
0                QP1
1                QP1
2                QP1
3                QP1
4                QP1
5                QP1
6                QP1
7                QP8
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.



show meter

show meter *meter_name*

Description

Displays the configured meters.

Syntax Description

<i>meter_name</i>	Specifies the meter name.
-------------------	---------------------------

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays meters on the switch:

```
show meter
```

The following is sample output from this command:

```
-----
Name      Committed Rate(Kbps)  Peak Rate(Kbps)
-----
peggy          1000000--
```

Note



When you are using a BlackDiamond 8800 series switches, SummitStack, or Summit family switches, you configure a peak rate for QoS meters using the `configure meter metername {committed-ratecir [Gbps | Mbps | Kbps]} {max-burst-sizeburst-size [Kb | Mb]} {out-actions [drop | set-drop-precedence {dscp [none | dscp-value]}}` command.

History

This command was first available in ExtremeXOS 11.4.



Platform Availability

This command is available on all platforms.

show ports congestion

```
show ports port_list congestion {no-refresh}
```

Description

Displays the port egress congestion statistics (dropped packets) for the specified ports on the front panel.

Syntax Description

<i>port_list</i>	Specifies one or more slots and ports.
no-refresh	Specifies a static snapshot of data instead of the default dynamic display.

Default

Displays the port congestion statistics for all ports in real-time.

Usage Guidelines

The bottom line in the real-time display shows keys that you can press to change the display. For example, you can clear the counters or page up or down through the list of ports.

Note



If you are displaying congestion statistics in real time and another CLI session resets the counters for a port you are monitoring, the counters displayed in your session for that port are also reset.

If you specify the `no-refresh` parameter, the system displays a snapshot of the data at the time you issue the command.

Note



Packets can be dropped at multiple locations along the path through the hardware. The per-port congestion counters count all dropped packets for all ports.



If you do not specify a port number or range in the command, dropped packet counts are displayed for all ports.

Note



To display the congestion statistics for the QoS profiles on a port, use the `show ports port_list qosmonitor {congestion} {no-refresh}` command.

On BlackDiamond 8900 xm-series modules and Summit X670 series switches, QP1-4 support one unicast and one multicast queue for each QoS profile. The congestion counters for QP1-4 tally the unicast and multicast traffic for these QoS profiles. Congestion counters for QP5-8 tally only the unicast traffic for these QoS profiles.

Example

The following example shows the packets dropped due to congestion for all ports in real time:

```
BD-8810.1 # show ports congestion
Port Congestion Monitor                                     Tue May 27 13:02:37
2008
Port          Link          Packet
State        Drop
=====
==
1:1           R             0
1:2           R             0
1:3           A             96
1:4           R             0
2:1           R             0
2:2           A            28513
2:3           R             0
2:4           R             0
2:5           R             0
2:6           R             0
2:7           R             0
2:8           R             0
3:1           R             0
3:2           R             0
3:3           R             0
3:4           R             0
=====
==
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
0->clear counters U->page up D->page down ESC->exit
```

The following example shows a snapshot display of the packets dropped due to congestion for all ports:

```
BD-8810.1 # show ports congestion no-refresh
Port          Link          Packet
State        Drop
=====
==
1:1           R             0
```



```

1:2      R      0
1:3      A      96
1:4      R      0
2:1      R      0
2:2      A      28513
2:3      R      0
2:4      R      0
2:5      R      0
2:6      R      0
2:7      R      0
2:8      R      0
3:1      R      0
3:2      R      0
3:3      R      0
3:4      R      0
5:1      R      0
=====
==
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback

```

History

This command was first available in ExtremeXOS 12.2.2.

Platform Availability

This command is available on all platforms.

show ports qosmonitor

```
show ports port_list qosmonitor {ingress | egress} {bytes | packets} {no-refresh}
```

Note



This description describes command operation on BlackDiamond X8 and BlackDiamond 8800 series switches. For a description of a similar command for operation on BlackDiamond X8 series switches, BlackDiamond 8800, SummitStack, and Summit family switches, see the command description for `show ports qosmonitor {congestion}`.

Description

Displays egress traffic counts or ingress traffic counts for each QoS profile on the specified ports.

Syntax Description

<i>port_list</i>	Specifies one or more slots and ports.
egress	Specifies the display of egress traffic counts. Default.
bytes	Specifies to display ingress or egress traffic counts in bytes.



packets	Specifies to display ingress or egress traffic counts in packets.
no-refresh	Specifies a static snapshot of data instead of the default dynamic display.

Default

Displays egress packet counts in real-time.

Usage Guidelines

The bottom line in the real-time display shows keys that you can press to change the display. For example, the spacebar toggles the display between QoS traffic counts in either packets or bytes.

If you specify the no-refresh parameter, the system displays a snapshot of the data at the time you issue the command.

If you do not specify a port number or range of ports when displaying ingress or egress traffic counts, traffic counts are displayed for all ports.

Example

The following example shows the egress packet counts for the specified ports:

```
# show ports 1:1-1:2 qosmonitor
Qos Monitor Req Summary                               Thu Mar  2 10:58:23
2006
Port      QP1      QP2      QP3      QP4      QP5      QP6      QP7      QP8
Pkt      Pkt      Pkt      Pkt      Pkt      Pkt      Pkt      Pkt
Xmts     Xmts     Xmts     Xmts     Xmts     Xmts     Xmts     Xmts
=====
==
1:1          0        0        0        0        0        0        0        0
1:2          0        0        0        0        0        0        0        0
=====
==
> indicates Port Display Name truncated past 8 characters
Spacebar->Toggle screen 0->Clear counters  U->Page up  D->Page down  ESC->exit
```

The following example shows the ingress packet counts for the specified ports:

```
# show ports 1:1-1:2 qosmonitor ingress
Qos Monitor Req Summary                               Thu Mar  2 10:59:28
2006
Port      IQP1     IQP2     IQP3     IQP4     IQP5     IQP6     IQP7     IQP8
Pkt      Pkt      Pkt      Pkt      Pkt      Pkt      Pkt      Pkt
Xmts     Xmts     Xmts     Xmts     Xmts     Xmts     Xmts     Xmts
=====
==
1:1          0        0        0        0        0        0        0        0
1:2          0        0        0        0        0        0        0        0
=====
==
```



> indicates Port Display Name truncated past 8 characters
 Spacebar->Toggle screen 0->Clear counters U->Page up D->Page down ESC->exit

History

This command was first available in ExtremeXOS 10.1.

The ingress information was added in ExtremeXOS 11.0.

Also, you must specify the ports in ExtremeXOS 11.0.

The egress and no-refresh keywords were added in ExtremeXOS 11.3.

The bytes and packets keywords, as well as the toggling functionality, were added in ExtremeXOS 11.4.

Platform Availability

This command is available on BlackDiamond X8 and 8800.

show ports qosmonitor {congestion}

```
show ports port_list qosmonitor {congestion} {no-refresh}
```



Note

This description describes command operation on the BlackDiamond X8 series switches, BlackDiamond 8800, SummitStack, and Summit family switches.

Description

Displays egress packet counts or dropped-traffic counts for each QoS profile on the specified ports.

Syntax Description

<i>port_list</i>	Specifies one or more slots and ports.
congestion	Specifies the display of packets dropped at ingress due to port congestion.
no-refresh	Specifies a static snapshot of data instead of the default dynamic display.

Default

Displays egress packet counts in real-time.



Usage Guidelines

The bottom line in the real-time display shows keys that you can press to change the display. For example, the spacebar toggles the display between egress packet counts and ingress dropped-packet counts.



Note

This command does not work properly if another CLI session is displaying congestion statistics in real time.

If you specify the no-refresh parameter, the system displays a snapshot of the data at the time you issue the command.



Note

Packets can be dropped at multiple locations along the path through the hardware. Due to hardware limitations, the dropped-packet counters for QoS profiles cannot count dropped packets from all possible locations. Because of these limitations, the sum of all dropped packets for all QoS profiles can be less than the per port count displayed with the command:

```
show ports port_list congestion {no-refresh}.
```

On BlackDiamond X8 series switches, and BlackDiamond 8500 and 8800 c-, and e-series modules, you can display packet counts for one port per slot or module at a time. You can simultaneously display packet counts for multiple ports, but they must be from different slots or modules. The dropped packet display is limited to the 8 most-significant digits. This limitation does not apply to BlackDiamond 8900 series modules.

When you display the packet counts for a port, this action configures the hardware to monitor that port. If the slot or module hardware was previously configured to monitor a different port, the counters are reset for the new port. If the selected port is the last port displayed on the module, the counters are not reset.



Note

On BlackDiamond X8 series switches, BlackDiamond 8900 xm-series modules and Summit X670 series switches, QP1-4 support one unicast and one multicast queue for each QoS profile. The QoS monitor counters for QP1-4 tally the unicast and multicast traffic for these QoS profiles. QoS monitor counters for QP5-8 tally only the unicast traffic for these QoS profiles.

Example

The following example shows the egress packet counts for the specified ports:

```
# show ports 2:1, 3:6 qosmonitor
Qos Monitor Req Summary                               Thu Mar  2 10:58:23
2006
Port          QP1      QP2      QP3      QP4      QP5      QP6      QP7      QP8
Pkt          Pkt      Pkt      Pkt      Pkt      Pkt      Pkt      Pkt
Xmts         Xmts     Xmts     Xmts     Xmts     Xmts     Xmts     Xmts
=====
==
```



```

2:1      0      0      0      0      0      0      0      0
3:6      0      0      0      0      0      0      0      0
=====
==
> indicates Port Display Name truncated past 8 characters
Spacebar->Toggle screen 0->Clear counters U->Page up D->Page down ESC->exit

```

The next example shows the dropped packet counts for the specified ports:

```

# show ports 2:1, 3:6 qosmonitor congestion
QoS Monitor Req Summary                               Thu Jun 12 01:17:14 2008
Port   QP1     QP2     QP3     QP4     QP5     QP6     QP7     QP8
Pkt    Pkt     Pkt     Pkt     Pkt     Pkt     Pkt     Pkt
Cong   Cong    Cong    Cong    Cong    Cong    Cong    Cong
=====
==
2:1    0       0       0       0       0       0       0       0
3:6    8745    0       129     0       0       0       0       0
=====
==
> indicates Port Display Name truncated past 8 characters
Spacebar->Toggle screen 0->Clear counters U->Page up D->Page down ESC->exit

```

History

This command was first available in ExtremeXOS 10.1.

You must specify the ports in ExtremeXOS 11.0.

The no-refresh keyword was added in ExtremeXOS 11.3.

The congestion keyword was added in ExtremeXOS 12.2, and the toggling functionality was modified to switch between egress packets and dropped packets.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, SummitStack, and Summit family switches.

show ports wred

```
show ports port_list wred {no-refresh}
```

Description

Displays WRED statistics for the specified ports or all ports.



Syntax Description

<i>port_list</i>	Specifies a list of slots and ports to display. Specify ports in the following formats: 3-5, 2:5, 2:6-2:8.
no-refresh	Specifies a static snapshot of data instead of the default dynamic display.

Default

N/A.

Usage Guidelines

If no port or port list is specified, this command displays the WRED statistics for all ports. If WRED is not configured on a port, the statistics for that port display as 0.

The drop counters in the display represent packets that were dropped based on the WRED congestion avoidance algorithm. The Green Pkt Drop column counts in-profile TCP and non-TCP packets that have been dropped. The Red Pkt Drop column counts out-of-profile TCP and non-TCP packets that have been dropped.

Note



The values in the Yellow Pkt Drop column are always 0 in this release because the yellow traffic color is not supported at this time. For Summit X460 and X480 switches and E4G-200 and E4G-400 switches, the values in the Green Pkt Drop column are always 0 because the hardware does not provide a drop counter for the green traffic color.

Example

The following example displays the WRED statistics for port list 2:1-9:

```
* Switch.243 # show ports 2:1-9 wred no-refresh
Port WRED Stats Monitor
=====
Port      Link      Green      Yellow      Red
State     Pkt Drop  Pkt Drop  Pkt Drop
=====
2:1       A         0          0           0
2:2       R         0          0           0
2:3       R         0          0           0
2:4       R         0          0           0
2:5       R         0          0           0
2:6       R         0          0           0
2:7       R         0          0           0
2:8       R         0          0           0
2:9       R         0          0           0
=====
Link State: A-Active, R-Ready, NP-Port Not Present L-Loopback
```



History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8900 c-, xl-, and xm-series modules, E4G-200 and E4G-400 switches, and Summit X460, X480, and X670 switches.




show qosprofile

```
show qosprofile [ all | port_list ]
```

Description

Displays QoS information on the switch.

Syntax Description

<i>ports</i>	Specifies to display information for specified ports.  Note This parameter is available only on BlackDiamond c-, e-, xl-, and xm-series modules, and Summit X460, X480, X670, and X770 series switches.
all	Specifies all ports.  Note This parameter is available only on BlackDiamond c-, e-, xl-, and xm-series modules, and Summit X460, X480, X670, and X770 series switches.
<i>port_list</i>	Specifies a list of slots and ports.  Note This parameter is available only on BlackDiamond c-, e-, xl-, and xm-series modules, and Summit X460, X480, X670, and X770 series switches.

Default

Displays egress QoS information for all ports.

Usage Guidelines

The displayed QoS profile information differs depending on the platform you are running on. The following section shows examples for different platforms.



Example

The display varies depending on your platform.

All Summit series, BD8K, BDX, BlackDiamond X8 Series Switches, BlackDiamond 8800 Modules, E4G only, whether or not included in a SummitStack

The following shows the information that appears when you omit the optional port parameter:

```
BD-8810Rack3.3 # show qosprofile
QP1   Weight = 1      Max Buffer Percent = 100
QP2   Weight = 1      Max Buffer Percent = 100
QP8   Weight = 1      Max Buffer Percent = 100
```

The following example shows how the display appears when the switch is configured for weighted-round-robin mode and some QoS profiles are configured for strict priority mode:

```
BD-8810.7 # show qosprofile
QP1   Weight = 1      Max Buffer Percent = 100
QP2   Weight = 1      Max Buffer Percent = 100
QP3   Weight = 1      Max Buffer Percent = 100
QP5   Strict-Priority Max Buffer Percent = 100
QP8   Strict-Priority Max Buffer Percent = 100
```

All Summit series, BD8K, BDX, BlackDiamond X8 Series Switches, BlackDiamond 8800 Modules, E4G only, whether or not included in a SummitStack

When you add the optional port parameter, the switch displays the following sample output:

```
Switch.6 # show qosprofile ports 1:1-2
Port: 1:1
QP1   MinBw =    20% MaxBw =    50% MaxBuf =   100%
QP8   MinBw =     0% MaxBw =   100% MaxBuf =  1000%
Port: 1:2
QP1   MinBw =     0% MaxBw =   100% MaxBuf =   100%
QP8   MinBw =     0% MaxBw =   100% MaxBuf =   100%
```



Note

This last sample output is not available on the XGS2 ports.

History

This command was first available in ExtremeXOS 10.1.

The ingress information was added in ExtremeXOS 11.0.



Platform Availability

This command is available on all platforms.

The ports parameter is available on all Summit series, BD8K, BDX, BlackDiamond X8 Series Switches, BlackDiamond 8800 Modules, and E4G whether or not included in a SummitStack.

show wredprofile

```
show wredprofile {ports [port_list | all]}
```

Description

Displays WRED configuration data for the specified ports or all ports.

Syntax Description

<i>port_list</i>	Specifies a list of slots and ports to display. Specify ports in the following formats: 3-5, 2:5, 2:6-2:8.
all	Specifies that this command applies to all ports on the device.

Default

N/A.

Usage Guidelines

If no port or port list is specified, this command displays the default WRED configuration values.

Example

The following example displays the WRED settings for port 2:

```
* Switch.9 # show wredprofile ports 2
Port: 2
WRED configuration parameters
=====
QoS      Packet      Min      Max      Max      Avg.
Profile  Type        Color   Thresh  Thresh  Drop-Rate  Weight
=====
QP1     TCP         Green   100%   100%   100%      4
QP1     TCP         Red     100%   100%   100%      4
QP1     non-TCP    Any     100%   100%   100%      4
QP1     non-TCP    Red     100%   100%   100%      4
QP3     TCP         Green   10%    20%    100%      4
QP3     TCP         Red     100%   100%   100%      4
QP3     non-TCP    Any     100%   100%   100%      4
QP3     non-TCP    Red     100%   100%   100%      4
QP8     TCP         Green   100%   100%   100%      4
QP8     TCP         Red     100%   100%   100%      4
```



QP8	non-TCP	Any	100%	100%	100%	4
QP8	non-TCP	Red	100%	100%	100%	4

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8900 c-, xl-, and xm-series modules, E4G-200 and E4G-400 switches, and Summit X460, X480, X670, and X770 switches.

unconfigure diffserv examination

The syntax for BlackDiamond 8800, SummitStack, and Summit family switches is:

```
unconfigure diffserv examination
```

Description

Disables DiffServ traffic groups.

Syntax Description

N/A.

Default

Disabled.

Usage Guidelines

Use this command to disable DiffServ code point examination.

Example

The following command disables DiffServ code point examination:

```
unconfigure diffserv examination
```

History

This command was first available in ExtremeXOS 11.0.



Platform Availability

This command is available on all platforms.

unconfigure diffserv replacement

The syntax for BlackDiamond X8 series switches, BlackDiamond 8800, SummitStack, and Summit family switches is:

```
unconfigure diffserv replacement
```

Description

Resets all DiffServ replacement mappings to the default values.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

Use this command to reset all DiffServ replacement mappings to default values.

Example

The following command resets the DiffServ replacement mappings to their default values:

```
unconfigure diffserv examination
```

History

This command was first available in ExtremeXOS 11.0.

The ports keyword was first available in ExtremeXOS 12.2.2.

Platform Availability

This command is available on all platforms.

unconfigure qosprofile



```
unconfigure qosprofile {ports [port_list|all]}
```



Description

Returns the rate-shaping parameters for all QoS profiles on the specified ports to the default values.

Syntax Description

<i>port_list</i>	<p>Specifies the ports on which to unconfigure QoS profiles.</p> <hr/> <p>Note</p>  <p>This parameter is available only on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit X460 and X480 series switches, whether or not included in a SummitStack.</p>
all	<p>Specifies that this command applies to all ports on the device.</p> <hr/> <p>Note</p>  <p>This parameter is available only on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit X460 and X480 series switches, whether or not included in a SummitStack.</p>

Default

The default values for egress bandwidth on all supported platforms are:

- Minimum bandwidth—0%
- Maximum bandwidth—100%

The default values for egress priority and ingress QoS profiles differ by platform as described in the following sections.

The platform-specific default values for the two default egress QoS profiles (QP1 and QP8) on the BlackDiamond 8800 series switches, SummitStack, and Summit family switches are:

- Maximum buffer—100% (as set by the `configure qosprofile` command)
- Maximum buffer override—100% (as set by the `configure qosprofile` command)
- Weight—1
- WRED—See the `configure qosprofile wred` command description.

Usage Guidelines

None.

Example

The following command resets the QoS profiles for all ports to default settings:

```
unconfigure qosprofile
```



History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

The egress and ports parameters are available only on the BlackDiamond X8 series switches, BlackDiamond c-, e-, xl-, and xm-series modules, and Summit X460, X480, X670, and X770 series switches, whether or not included in a SummitStack.

unconfigure qosprofile wred

```
unconfigure qosprofile wred {ports [port_list | all]}
```

Description

Removes the WRED configuration for all QoS profiles on the specified port or all ports.

Syntax Description

<i>port_list</i>	Specifies a list of slots and ports from which the WRED configuration is removed. Specify ports in the following formats: 3-5, 2:5, 2:6-2:8.
all	Specifies that this command applies to all ports on the device.

Default

N/A.

Usage Guidelines

None.

Example

The following example removes the WRED configuration for port 3:

```
* Switch.24 # unconfigure qosprofile wred port 3
```

History

This command was first available in ExtremeXOS 12.7.



Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8900 c-, xl-, and xm-series modules, E4G-200 and E4G-400 switches, and Summit X460, X480, X670, and X770 switches.



24 Network Login Commands

```
clear netlogin state
configure netlogin add mac-list
configure netlogin add proxy-port
configure netlogin agingtime
configure netlogin allowed-refresh-failures
configure netlogin authentication database-order
configure netlogin authentication failure vlan
configure netlogin authentication service-unavailable vlan
configure netlogin banner
configure netlogin base-url
configure netlogin delete mac-list
configure netlogin delete proxy-port
configure netlogin dot1x eapol-transmit-version
configure netlogin dot1x guest-vlan
configure netlogin dot1x timers
configure netlogin dynamic-vlan
configure netlogin dynamic-vlan uplink-ports
configure netlogin local-user
configure netlogin local-user security-profile
configure netlogin mac timers reauth-period
configure netlogin move-fail-action
configure netlogin port allow egress-traffic
configure netlogin ports mode
configure netlogin ports no-restart
configure netlogin ports restart
configure netlogin redirect-page
configure netlogin session-refresh
configure netlogin vlan
configure vlan netlogin-lease-timer
create netlogin local-user
delete netlogin local-user
disable netlogin
disable netlogin authentication failure vlan ports
disable netlogin authentication service-unavailable vlan ports
disable netlogin dot1x guest-vlan ports
disable netlogin logout-privilege
disable netlogin ports
```

```
disable netlogin reauthenticate-on-refresh
disable netlogin redirect-page
disable netlogin session-refresh
enable netlogin
enable netlogin authentication failure vlan ports
enable netlogin authentication service-unavailable vlan ports
enable netlogin dot1x guest-vlan ports
enable netlogin logout-privilege
enable netlogin ports
enable netlogin reauthentication-on-refresh
enable netlogin redirect-page
enable netlogin session-refresh
show banner netlogin
show netlogin
show netlogin authentication failure vlan
show netlogin authentication service-unavailable vlan
show netlogin banner
show netlogin guest-vlan
show netlogin local-users
show netlogin mac-list
unconfigure netlogin allowed-refresh-failures
unconfigure netlogin authentication database-order
unconfigure netlogin authentication failure vlan
unconfigure netlogin authentication service-unavailable vlan
unconfigure netlogin banner
unconfigure netlogin dot1x guest-vlan
unconfigure netlogin local-user security-profile
unconfigure netlogin session-refresh
unconfigure netlogin vlan
```

This chapter describes commands for configuring network login.

Network login is a feature designed to control the admission of user packets into a network by giving network access only to users that have been properly authenticated. Network login is controlled by an administrator on a per port, per VLAN basis and uses an integration of DHCP, user authentication over the web interface, user authentication by MAC address, or 802.1x client software, and a RADIUS server to provide a user database or specific configuration details.

Network login has two modes of operation:

- Campus mode, used when a port in a VLAN will move to another VLAN when authentication has been completed successfully. This mode is for the roaming user who will not always be using the same port for authentication. Campus mode requires a DHCP server and a RADIUS server configured for Extreme Network Login.
- ISP mode, used when the port and VLAN used will remain constant. All network settings are configured for that VLAN.



A DHCP server is included to support network login functionality.

clear netlogin state

```
clear netlogin state {port port_list}
```

Description

Clears and initializes the network login sessions on a VLAN port.

Syntax Description

<i>port_list</i>	Specifies the ports to clear.
------------------	-------------------------------

Default

None.

Usage Guidelines

Clear the states of every MAC learned on this VLAN port and put the port back to unauthenticated state. The port will be moved to its original VLAN if configured in campus mode.

Example

The following command clears the Network Login state of port 2:9:

```
clear netlogin state port 2:9
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure netlogin add mac-list

```
configure netlogin add mac-list [mac {mask} | default] {encrypted} {password}  
{ports port_list}
```

Description

Adds an entry to the MAC address list for MAC-based network login.



Syntax Description

<i>mac</i>	Specifies the MAC address to add.
<i>mask</i>	Specifies the number of bits to use for the mask.
default	Specifies the default entry.
encrypted	Used to display encrypted form of password in configuration files. Do not use.
<i>password</i>	Specifies the password to send for authentication.
ports	Specifies the port or port list to use for authentication.

Default

If no password is specified, the MAC address will be used.

Usage Guidelines

Use this command to add an entry to the MAC address list used for MAC-based network login.

If no match is found in the table of MAC entries, and a default entry exists, the default will be used to authenticate the client. All entries in the list are automatically sorted in longest prefix order.

Associating a MAC Address to a Port

You can configure the switch to accept and authenticate a client with a specific MAC address. Only MAC addresses that have a match for the specific ports are sent for authentication. For example, if you associate a MAC address with one or more ports, only authentication requests for that MAC addresses received on the port(s) are sent to the RADIUS server. The port(s) block all other authentication requests that do not have a matching entry. This is also known as secure MAC.

To associate a MAC address with one or more ports, specify the ports option when using the `configure netlogin add mac-list [mac {mask} | default] {encrypted} {password} {portsport_list}` command.

You must enable MAC-based network login on the switch and the specified ports before using this command. If MAC-based network login is not enabled on the specified port(s), the switch displays a warning message similar to the following:

```
WARNING: Not all specified ports have MAC-Based NetLogin enabled.
```

If this occurs, make sure to enable MAC-based network login.

Example

The following command adds the MAC address 10:20:30:40:50:60 with the password foo to the list:

```
configure netlogin add mac-list 10:20:30:40:50:60 password foo
```



The following command associates MAC address 10:20:30:40:50:70 with ports 2:2 and 2:3. This means authentication requests from MAC address 10:20:30:40:50:70 are only accepted on ports 2:2 and 2:3:

```
configure netlogin add mac-list mac 10:20:30:40:50:70 ports 2:2-2:3
```

History

This command was first available in ExtremeXOS 11.1.

The ports option was added in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

configure netlogin add proxy-port

```
configure netlogin add proxy-port tcp_port {http | https}
```

Description

Configure the ports that will be hijacked and redirected for HTTP or HTTPS traffic.

Syntax Description

<i>tcp_port</i>	Specifies the port to be hijacked.
-----------------	------------------------------------

Default

HTTP traffic.

Usage Guidelines

This command allows you to configure the ports that will be hijacked and redirected for HTTP or HTTPS traffic. For each hijacked proxy port, you must specify whether the port is to be used for HTTP or HTTPS traffic.

No more than 5 such ports are supported in addition to ports 80 and ports 443. Attempts to add more than 5 ports generate an error.

History

This command was first available in ExtremeXOS 12.1.



Platform Availability

This command is available on all platforms.

configure netlogin agingtime

configure netlogin agingtime *minutes*

Description

Lets you configure network login aging.

Syntax Description

<i>minutes</i>	Specifies the aging time in minutes.
----------------	--------------------------------------

Default

The default value is 5.

Usage Guidelines

Use this command to configure the aging time for network login. The aging time is the time after which learned clients that failed authentication or did not attempt to authenticate are removed from the system. This prevents the switch from keeping all clients ever seen on a network-login-enabled port.

The range can be from 0 to 3000, where 0 indicates no age out.

Example

The following command specifies an aging time of 15 minutes:

```
configure netlogin agingtime 15
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure netlogin allowed-refresh-failures

configure netlogin allowed-refresh-failures *num_failures*



Description

Sets the number refresh failures.

Syntax Description

<i>num_failures</i>	Specifies the number of refresh failures. The range is from 0 to 5.
---------------------	---

Default

The default is 0.

Usage Guidelines

This command allows you to set the number of refresh failures allowed. You can set the number of failures to be from between 0 to 5. The default value is 0.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

configure netlogin authentication database-order

```
configure netlogin [mac | web-based] authentication database-order [[radius] | [local] | [radius local] | [local radius]]
```

Description

Configures the order of database authentication protocols to use.

Syntax Description

mac	Specifies MAC-based authentication.
web-based	Specifies Web-based authentication.
radius	Specifies an authentication order from only the RADIUS database.
local	Specifies an authentication order from only the local database.
radius local	Specifies an authentication order of RADIUS database first, followed by local.
local radius	Specifies an authentication order of local database first, followed by RADIUS.



Default

By default, the authentication order is RADIUS, local-user database.

Usage Guidelines

Use this command in situations where, when both a network login RADIUS server and a local-user database are configured, you want to have control over which database to use first. If one authentication fails, the other database is tried; if that authentication is successful, the switch authenticates the network login user.

Example

The following command sets the database authentication order to local-user database, RADIUS:

```
configure netlogin mac authentication database-order local radius
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

configure netlogin authentication failure vlan

```
configure netlogin authentication failure vlan vlan_name {ports port_list}
```

Description

Configures authentication failure VLAN on network login enabled ports.

Syntax Description

<i>vlan_name</i>	Specifies the name of the authentication failure VLAN.
<i>port_list</i>	Specifies one or more ports or slots and ports. If the ports keyword is not used, the command applies to all ports.

Default

By default, authentication failure VLAN is configured on all network login enabled ports if no port is specifically configured.



Usage Guidelines

Use this command to configure authentication failure VLAN on network login enabled ports. When a supplicant fails authentication, it is moved to the authentication failure VLAN and is given limited access until it passes the authentication either through RADIUS or local. Depending on the authentication database order for that particular network login method (MAC, web or dot1x), the other database is used to authenticate the client. If the final result is an authentication failure and if the authentication failure VLAN is configured and enabled on that port, the client is moved to that location.

There four different authentication orders which can be configured per authentication method currently. They are:

- RADIUS.
- local.
- RADIUS, local.
- local, RADIUS.

In each case, you must consider the end result in deciding whether to authenticate the client in authentication failure VLAN or authentication service unavailable VLAN (if configured).

For example, when netlogin mac authentication database order is local, radius, if the authentication of a MAC client fails through a local database, RADIUS is used for authentication. If RADIUS also fails authentication, the client is moved to authentication failure VLAN. The same is true for all authentication database orders (radius,local; local,radius; radius; local).

If authentication through local fails but passes through RADIUS, the client is moved to the appropriate destination VLAN.

If the local authentication fails and the RADIUS server is not available, the client is not moved to authentication failure VLAN.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

configure netlogin authentication service-unavailable vlan

```
configure netlogin authentication service-unavailable vlan vlan_name {ports  
port_list}
```

Description

Configures authentication service unavailable VLAN on network login enabled ports.



Syntax Description

<i>vlan_name</i>	Specifies the name of the service-unavailable VLAN.
<i>port_list</i>	Specifies one or more ports or slots and ports. If the ports keyword is not used, the command applies to all ports.

Default

Defaults to all network login enabled ports.

Usage Guidelines

This command configures authentication service unavailable VLAN on the specified network login enabled ports. Authentication service unavailable VLAN is configured on all the network login enabled ports, if no port is specifically mentioned. When an authentication service is not available to authenticate the network login clients, they are moved to the authentication service-unavailable VLAN and are given limited access until the authentication service is available either through RADIUS or local. Depending on the authentication database order for that particular network login method (MAC, web or dot1x), the other database is used to authenticate the client. If the final result is an authentication failure and if the authentication failure VLAN is configured and enabled on that port, the client is moved to that location.



Note

The local database can be configured for MAC and Web authentication method only, not for dot1x.

There are four different authentication orders which can be configured per authentication method currently. They are:

- RADIUS.
- Local.
- RADIUS, local.
- Local, RADIUS.

In each case, you must consider the end result in deciding whether to authenticate the client in authentication failure VLAN or authentication service unavailable VLAN (if configured).

For example, when netlogin mac authentication database order is local, radius, if the authentication of a MAC client fails through a local database, RADIUS is used for authentication. If RADIUS also fails authentication, the client is moved to authentication failure VLAN. The same is true for all authentication database orders (radius,local; local,radius; radius; local).

If authentication through local fails but passes through RADIUS, the client is moved to appropriate destination VLAN.

If the local authentication fails and the RADIUS server is not available, the client is not moved to authentication failure VLAN.

Authentication service is considered to be unavailable in the following cases:

- For local authentication if the user entry is not present in the local database.



- For RADIUS in the following cases:
 - the RADIUS server is not running.
 - the RADIUS server is not configured on the switch.
 - the RADIUS server is configured but not enabled on the switch.

Note



If web is enabled on a port where dot1x or MAC are also enabled, the authentication failure/service-unavailable VLAN configuration is not applicable to those clients where dot1x or MAC clients which fail authentication or where authentication service is not available.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

configure netlogin banner

configure netlogin banner *banner*

Description

Configures the network login page banner.

Syntax Description

<i>banner</i>	Specifies the HTML code for the banner.
---------------	---

Default

The default banner is the Extreme Networks logo.

Usage Guidelines

The banner is a quoted, HTML string, that will be displayed on the network login page. The string is limited to 1024 characters.

This command applies only to the web-based authentication mode of network login.



Example

The following command configures the network login page banner:

```
configure netlogin banner "<html><head>Please Login</head></html>"
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure netlogin base-url

```
configure netlogin base-url url
```

Description

Configures the base URL for network login.

Syntax Description

<i>url</i>	Specifies the base URL for network login.
------------	---

Default

The base URL default value is "network-access.com."

Usage Guidelines

When you login using a web browser, you are redirected to the specified base URL, which is the DNS name for the switch.

You must configure a DNS name of the type "www.xx...xx.xxx" or "xx...xx.xxx".

This command applies only to the web-based authentication mode of network login.

Example

The following command configures the network login base URL as access.net:

```
configure netlogin base-url access.net
```



History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure netlogin delete mac-list

```
configure netlogin delete mac-list [mac {mask} | default]
```

Description

Deletes an entry from the MAC address list for MAC-based network login.

Syntax Description

<i>mac</i>	Specifies the MAC address to delete.
<i>mask</i>	Specifies the number of bits to use for the mask.
default	Specifies the default entry.

Default

N/A.

Usage Guidelines

Use this command to delete an entry from the MAC address list used for MAC-based network login.

Example

The following command deletes the MAC address 10:20:30:40:50:60 from the list:

```
configure netlogin delete mac-list 10:20:30:40:50:60
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.



configure netlogin delete proxy-port

```
configure netlogin delete proxy-port tcp_port
```

Description

Configure the ports that are to be hijacked and redirected for HTTP or HTTPS traffic.

Syntax Description

<i>tcp_port</i>	Specifies the port to be hijacked.
-----------------	------------------------------------

Default

N/A.

Usage Guidelines

This command allows you to unconfigure the ports that will be hijacked and redirected for HTTP or HTTPS traffic.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

configure netlogin dot1x eapol-transmit-version

```
configure netlogin dot1x eapol-transmit-version eapol-version
```

Description

Configures the default EAPOL version sent in transmitted packets for network login.

Syntax Description

<i>eapol-version</i>	Specifies the EAPOL version. Choices are "v1" or "v2".
----------------------	--

Default

The default is "v1".



Usage Guidelines

Although the ExtremeXOS software supports EAPOL version 2, some clients do not yet accept the version 2 EAPOL packets. The packet format for the two versions is the same.

Example

The following command changes the EAPOL version to 2:

```
configure netlogin dot1x eapol-transmit-version v2
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure netlogin dot1x guest-vlan

```
configure netlogin dot1x guest-vlan vlan_name {ports port_list}
```

Description

Configures a guest VLAN for 802.1x authentication network login.

Syntax Description

<i>vlan_name</i>	Specifies the name of the guest VLAN.
<i>port_list</i>	Specifies one or more ports or slots and ports. If the ports keyword is not used, the command applies to all ports.

Default

N/A.



Usage Guidelines

This command configures the guest VLAN for 802.1x on the current virtual router (VR).

Note



Beginning with ExtremeXOS 11.6, you can configure guest VLANs on a per port basis, which allows you to configure more than one guest VLAN per VR. In ExtremeXOS 11.5 and earlier, you can only configure guest VLANs on a per VLAN basis, which allows you to configure only one guest VLAN per VR.

If you do not specify any ports, the guest VLAN is configured for all ports.

Each port can have a different guest VLAN.

A guest VLAN provides limited or restricted network access if a supplicant connected to a port does not respond to the 802.1x authentication requests from the switch. A port always moves untagged into the guest VLAN.

Keep in mind the following when configuring guest VLANs:

- You must create a VLAN and configure it as a guest VLAN before enabling the guest VLAN feature.
- Configure guest VLANs only on network login ports with 802.1x enabled.
- Movement to guest VLANs is not supported on network login ports with MAC-based or web-based authentication.
- 802.1x must be the only authentication method enabled on the port for movement to guest VLAN.
- No supplicant on the port has 802.1x capability.
- You configure only one guest VLAN per virtual router interface.

Note



The supplicant does not move to a guest VLAN if it fails authentication after an 802.1x exchange; the supplicant moves to the guest VLAN only if it does not respond to an 802.1x authentication request.

Modifying the Supplicant Timer

By default, the switch attempts to authenticate the supplicant every 30 seconds for a maximum of three tries. If the supplicant does not respond to the authentication requests, the client moves to the guest VLAN. The number of authentication attempts is not a user-configured parameter.

To modify the supplicant response timer, use the following command and specify the `supp-resp-timeout` parameter:

```
configure netlogin dot1x timers [{server-timeout server_timeout} {quiet-
periodquiet_period} {reauth-period reauth_period {reauth-maxmax_num_reauths}}
{supp-resp-timeoutsupp_resp_timeout}]
```

If a supplicant on a port in the guest VLAN becomes 802.1x-capable, the switch starts processing the 802.1x responses from the supplicant. If the supplicant is successfully authenticated, the port moves from the guest VLAN to the destination VLAN specified by the RADIUS server.



Enabling Guest VLANs

To enable the guest VLAN, use the following command:

```
enable netlogin dot1x guest-vlan ports [all |ports]
```

Example

The following command creates a guest VLAN for 802.1x named guest for all ports:

```
configure netlogin dot1x guest-vlan guest
```

The following command creates a guest VLAN named guest for ports 2 and 3:

```
configure netlogin dot1x guest-vlan guest ports 2,3
```

History

This command was first available in ExtremeXOS 11.2.

The ports option was added in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure netlogin dot1x timers

```
configure netlogin dot1x timers [{server-timeout server_timeout} {quiet-period quiet_period} {reauth-period reauth_period {reauth-max max_num_reauths}} {supp-  
resp-timeout supp_resp_timeout}]
```

Description

Configures the 802.1x timers for network login.

Syntax Description

server-timeout	Specifies the timeout period for a response from the RADIUS server. The range is 1 to 120 seconds.
quiet-period	Specifies the time for which the switch will not attempt to communicate with the supplicant after authentication has failed. The range is 0 to 65535 seconds.
reauth-period	Specifies time after which the switch will attempt to re-authenticate an authenticated supplicant. The range is 0, 30 to 7200 seconds.



reauth-max	Specifies the maximum reauthentication counter value. The range is 1 to 10.
supp-resp-timeout	Specifies the time for which the switch will wait for a response from the supplicant. The range is 1 to 120 seconds.

Default

The defaults are as follows:

- server-timeout—30 seconds.
- quiet-period—60 seconds.
- reauth-period—3600 seconds.
- reauth-max—3.
- supp-resp-timeout—30 seconds.

Usage Guidelines

To disable re-authentication, specify 0 for the reauth-period parameter. (If reauth-period is set to 0, reauth-max value doesn't apply.)

If you attempt to configure a timer value that is out of range (not supported), the switch displays an error message. The following is a list of sample error messages:

- server-timeout—`ERROR: RADIUS server response timeout out of range (1..120 sec)`
- quiet-period—`%% Invalid number detected at '^' marker. %% Input number must be in the range [0, 65535].`
- reauth-period—`ERROR: Re-authentication period out of range (0, 30..7200 sec)`
- reauth-counter—`ERROR: Re-authentication counter value out of range (1..10)`
- supp-resp-timeout—`ERROR: Supplicant response timeout out of range (1..120 sec)`

To display the 802.1x timer settings, use the `show netlogin` and `show netlogin dot1x` commands.

Example

The following command changes the 802.1x server-timeout to 10 seconds:

```
configure netlogin dot1x timers server-timeout 10
```

History

This command was first available in ExtremeXOS 11.1.

The reauth-max keyword was added in ExtremeXOS 12.1.



Platform Availability

This command is available on all platforms.

configure netlogin dynamic-vlan

```
configure netlogin dynamic-vlan [disable | enable]
```

Description

Configures the switch to automatically and dynamically create a VLAN after receiving authentication requests from one or more supplicants (clients).

Syntax Description

disable	Specifies that the switch does not automatically create dynamic VLANs. This is the default behavior.
enable	Specifies that the switch automatically create dynamic VLANs.

Default

The default is disabled.

Usage Guidelines

Use this command to configure the switch to dynamically create a VLAN. If configured for dynamic VLAN creation, the switch automatically creates a supplicant VLAN that contains both the supplicant's physical port and one or more uplink ports.

A dynamically created VLAN is only a Layer 2 bridging mechanism; this VLAN does not work with routing protocols to forward traffic. After the switch unauthenticates all of the supplicants from the dynamically created VLAN, the switch deletes that VLAN.

Note



Dynamically created VLANs do not support the session refresh feature of web-based network login because dynamically created VLANs do not have an IP address. Also, dynamic VLANs are not supported on ports when STP and network login are both configured on the ports.

By dynamically creating and deleting VLANs, you minimize the number of active VLANs configured on your edge switches. In addition, the RADIUS server forwards VSA information to dynamically create the VLAN thereby simplifying switch management. A key difference between dynamically created VLANs and other VLANs is that the switch does not save dynamically created VLANs. Even if you use the save command, the switch does not save a dynamically created VLAN.



Supported Vendor Specific Attributes

To prevent conflicts with existing VLANs on the switch, the RADIUS server uses Vendor Specific Attributes (VSAs) to forward VLAN information, including VLAN ID, to the switch. The following list specifies the supported VSAs for configuring dynamic network login VLANs:

- Extreme: Netlogin-VLAN-ID (VSA 209).
- IETF: Tunnel-Private-Group-ID (VSA 81).
- Extreme: Netlogin-Extended-VLAN (VSA 211).



Note

If the ASCII string only contains numbers, it is interpreted as the VLAN ID. Dynamic VLANs only support numerical VLAN IDs; VLAN names are not supported.

The switch automatically generates the VLAN name in the following format: `SYS_NLD_TAG` where `TAG` specifies the VLAN ID. For example, a dynamic network login VLAN with an ID of 10 has the name `SYS_NLD_0010`.

Specifying the Uplink Ports

To specify one or more ports as tagged uplink ports that are added to the dynamically created VLAN, use the following command:

```
configure netlogin dynamic-vlan uplink-ports [port_list | none]
```

The uplink ports send traffic to and from the supplicants from the core of the network.

By default the setting is none. For more information about this command, see the usage guidelines for `configure netlogin dynamic-vlan uplink-ports`.

Viewing Status Information

To display summary information about all of the VLANs on the switch, including any dynamic VLANs currently operating on the switch, use the following command:

```
show vlan
```

If the switch dynamically creates a VLAN, the VLAN name begins with `SYS_NLD_` and the output contains a `d` flag for the dynamically created VLAN.

To display the status of dynamic VLAN configuration on the switch, use the following command:

```
show netlogin
```

The switch displays the current state of dynamic VLAN creation (enabled or disabled) and the uplink port(s) associated with the dynamic VLAN.



Example

The following command automatically adds ports 1:1-1:2 to the dynamically created VLAN as uplink ports:

```
configure netlogin dynamic-vlan uplink-ports 1:1-1:2
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure netlogin dynamic-vlan uplink-ports

```
configure netlogin dynamic-vlan uplink-ports [port_list | none]
```

Description

Specifies which port(s) are added as tagged, uplink ports to the dynamically created VLANs for network login.

Syntax Description

<i>port_list</i>	Specifies one or more ports to add to the dynamically created VLAN for network login.
none	Specifies that no ports are added. This is the default setting.

Default

The default setting is none.

Usage Guidelines

Use this command to specify which port(s) are used as uplink ports and added to the dynamically created VLAN for network login. The uplink ports send traffic to and from the supplicants from the core of the network.

Uplink ports should not be configured for network login (network login is disabled on uplink ports). If you specify an uplink port with network login enabled, the configuration fails and the switch displays an error message similar to the following:

```
ERROR: The following ports have NetLogin enabled: 1, 2
```

If this occurs, select a port with network login disabled.



Enabling Dynamic Network Login VLANs

To configure the switch to dynamically create a VLAN upon receiving an authentication response, use the following command:

```
configure netlogin dynamic-vlan [disable | enable]
```

By default, the setting is disabled. For more detailed information about this command, see the usage guidelines `configure netlogin dynamic-vlan uplink-ports`.

Viewing Status Information

To display summary information about all of the VLANs on the switch, including any dynamic VLANs currently operating on the switch, use the following command:

```
show vlan
```

If the switch dynamically creates a VLAN, the VLAN name begins with SYS_NLD_ and the output contains a d flag for the dynamically created VLAN.

To display the status of dynamic VLAN configuration on the switch, use the following command:

```
show netlogin
```

The switch displays the current state of dynamic VLAN creation (enabled or disabled) and the uplink port(s) associated with the dynamic VLAN.

Example

The following command configures the switch to add ports 1:1-1:2 to the dynamically created network login VLAN:

```
configure netlogin dynamic-vlan uplink-ports 1:1-1:2
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure netlogin local-user

```
configure netlogin local-user user-name {vlan-vsa [[{tagged | untagged}  
[vlan_name | vlan_tag]] | none}]
```

Description

Configures an existing local network login account.



Syntax Description

<i>user-name</i>	Specifies the name of an existing local network login account.
tagged	Specifies that the client be added as tagged.
untagged	Specifies that the client be added as untagged.
<i>vlan_name</i>	Specifies the name of the destination VLAN.
<i>vlan_tag</i>	Specifies the VLAN ID, tag, of the destination VLAN.
none	Specifies that the VSA 211 wildcard (*) is applied, only if you do not specify tagged or untagged.

Default

N/A.

Usage Guidelines

Use this command to modify the attributes of an existing local network login account. You can update the following attributes associated with a local network login account:

- Password of the local network login account.
- Destination VLAN attributes including: adding clients tagged or untagged, the name of the VLAN, and the VLAN ID.



Note

Passwords are case-sensitive and must have a minimum of 1 character and a maximum of 32 characters.

You must create a local network login account before using this command. To create a local network login user name and password, use the following command:

```
create netlogin local-user user-name {encrypted} {password} {vlan-vsa [{tagged | untagged} [vlan_name] | vlan_tag]} {security-profilesecurity_profile}
```

If the switch displays a message similar to the following:

```
* Switch # configure netlogin local-user purplenet
^
%% Invalid input detected at '^' marker.
```

You might be attempting to modify a local network login account that is not present on the switch, or you might have incorrectly entered the account name. To confirm the names of the local network login accounts on your switch, use the following command:

```
show netlogin local-users
```

Additional Requirements

This command applies only to the web-based and MAC-based modes of network login. 802.1x network login does not support local database authentication.



You must have administrator privileges to use this command. If you do not have administrator privileges, the switch displays a message similar to the following:

```
This user does not have permissions for this command.
```

Passwords are case-sensitive. Passwords must have a minimum of 0 characters and a maximum of 32 characters. If you attempt to create a password with more than 32 characters, the switch displays the following message after you re-enter the password:

```
Password cannot exceed 32 characters
```

Example

This section contains the following examples:

- Updating the password.
- Modifying destination VLAN attributes.

Updating the Password

The following command updates the password of an existing local network login account:

```
configure netlogin local-user megtest
```

After you enter the local network login user name, press [Enter]. The switch prompts you to enter a password; however, the switch does not display the password. At the prompt enter the new password:

```
password:
```

After you enter the new password, press [Enter]. The switch then prompts you to re-enter the password:

```
Reenter password:
```

Updating VLAN Attributes

You can add a destination VLAN, change the destination VLAN, or remove the destination from an existing local network login account. This example changes the destination VLAN for the specified local network login account:

```
configure netlogin local-user megtest vlan-vsa green
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.



configure netlogin local-user security-profile

```
configure netlogin local-user user-name security-profile security_profile
```

Description

Changes a previously associated security profile.

Syntax Description

<i>user-name</i>	Specifies the name of an existing local network login account.
<i>security_profile</i>	Specifies a security profile string during account creation.

Default

N/A.

Usage Guidelines

Use this command to change any previously associated security profiles on the switch.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

configure netlogin mac timers reauth-period

```
configure netlogin mac timers reauth-period reauth_period
```

Description

Configures the reauthentication period for network login MAC-based authentication.

Syntax Description

reauth_period	Specifies time after which the switch will attempt to re-authenticate an authenticated supplicant. The range is 0, 30 to 7200 seconds.
----------------------	--

Default

The default is 0 (disabled).



Usage Guidelines

This command allows you to configure the reauth-period for network login MAC-based authentication. The session-timeout configuration on the RADIUS server overrides the reauth-period if it has been configured.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

configure netlogin move-fail-action

```
configure netlogin move-fail-action [ authenticate | deny ]
```

Description

Configures the action network login takes if a VLAN move fails. This can occur if two clients attempt to move to an untagged VLAN on the same port.

Syntax Description

authenticate	Specifies that the client is authenticated.
deny	Specifies that the client is not authenticated. This is the default setting.

Default

The default setting is deny.

Usage Guidelines

Use this command to specify how network login behaves if a VLAN move fails. Network login can either authenticate the client on the current VLAN or deny the client.

The following describes the parameters of this command if two clients want to move to a different untagged VLAN on the same port:

- **authenticate**—Network login authenticates the first client that requests a move and moves that client to the requested VLAN. Network login authenticates the second client but does not move that client to the requested VLAN. The second client moves to the first client's authenticated VLAN.
- **deny**—Network login authenticates the first client that requests a move and moves that client. Network login does not authenticate the second client.

To view the current move-fail-action setting on the switch, use the `show netlogin` command.



Example

The following command configures network login to authenticate the client on the current VLAN:

```
configure netlogin move-fail-action authenticate
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure netlogin port allow egress-traffic

```
configure netlogin ports [port_list | all] allow egress-traffic [none | unicast | broadcast | all_cast]
```

Description

Configures the egress traffic in an unauthenticated state.

Syntax Description

all	Specifies all network login ports.
<i>port_list</i>	Specifies one or more network login ports.
none	Specifies that no traffic is sent out if no authenticated clients exist on the VLAN.
unicast	Specifies that the unicast flooding traffic for the VLANs on the network login enabled port be sent.
broadcast	Specifies that the broadcast traffic for the VLANs on the network login enabled port be sent.
all_cast	Specifies that the broadcast and unicast flooding traffic for the VLANs on the network login enabled port be sent.

Default

The default is none.

Usage Guidelines

This command allows you to configure the egress traffic in an unauthenticated state on a per-port basis.



History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

configure netlogin ports mode

```
configure netlogin ports [all | port_list] mode [mac-based-vlans | port-based-vlans]
```

Description

Configures the network login port's mode of operation.

Syntax Description

all	Specifies all netlogin ports.
<i>port_list</i>	Specifies one or more network login ports.
mac-based-vlans	Allows more than one untagged VLAN.
port-based-vlans	Allows only one untagged VLAN. This is the default behavior.

Default

The default setting is port-based-vlans.

Usage Guidelines

Use this command to configure network login MAC-based VLANs on a network login port.

If you modify the mode of operation to mac-based-vlans and later disable all network login protocols on that port, the mode of operation automatically returns to port-based-vlans.

When you change the network login port's mode of operation, the switch deletes all currently known supplicants from the port and restores all VLANs associated with that port to their original state. In addition, by selecting mac-based-vlans, you are unable to manually add or delete untagged VLANs from this port. Network login now controls these VLANs.

With network login MAC-based operation, every authenticated client has an additional FDB flag that indicates a translation MAC address. If the supplicant's requested VLAN does not exist on the port, the switch adds the requested VLAN.

Important Rules and Restrictions

This section summarizes the rules and restrictions for configuring network login MAC-based VLANs:



- If you attempt to configure the port's mode of operation before enabling network login, the switch displays an error message similar to the following:

```
ERROR: The following ports do not have NetLogin enabled; 1
```

To enable network login on the switch, use the following command to enable network login and to specify an authentication method (for example, 802.1x—identified as dot1x in the CLI):

```
enable netlogin dot1x
```

To enable network login on the ports, use the following command to enable network login and to specify an authentication method (for example, 802.1x—identified as dot1x in the CLI):

```
enable netlogin ports 1:1 dot1x
```

- On ExtremeXOS versions prior to 12.0 on switches other than the Summit family, 10 Gigabit Ethernet ports such as those on the 10G4X I/O module and the uplink ports on Summit family switches do not support network login MAC-based VLANs.

If you attempt to configure network login MAC-based VLANs on 10 Gigabit Ethernet ports, the switch displays an error message similar to the following:

```
ERROR: The following ports do not support the MAC-Based VLAN mode; 1, 2, 10
```

In ExtremeXOS version 12.0, BlackDiamond 8800 10G4Xa modules support MAC-based VLANs on 10 Gigabit Ethernet ports. The BlackDiamond 8800 10G4X 10 Gigabit Ethernet ports do not support MAC-based VLANs.

- You can have a maximum of 1,024 MAC addresses per I/O module or per Summit family switch.

Displaying FDB Information

To view network login-related FDB entries, use the following command:

```
show fdb netlogin [all | mac-based-vlans]
```

The following is sample output from the show fdb netlogin mac-based-vlans command:

Mac	Vlan	Age	Use	Flags	Port List
00:04:96:10:51:80	VLONE(0021)	0086	0000	n m	v 1:11
00:04:96:10:51:81	VL TWO(0051)	0100	0000	n m	v 1:11
00:04:96:10:51:91	VL TWO(0051)	0100	0000	n m	v 1:11

Flags : d - Dynamic, s - Static, p - Permanent, n - NetLogin, m - MAC, i - IP, x - IPX, l - lockdown MAC, M - Mirror, B - Egress Blackhole, b - Ingress Blackhole, v - NetLogin MAC-Based VLAN.

The flags associated with network login include:

- v—Indicates the FDB entry was added because the port is part of a MAC-based virtual port/VLAN combination.
- n—Indicates the FDB entry was added by network login.



Displaying Port and VLAN Information

To view information about the VLANs that are temporarily added in MAC-based mode for network login, use the following command:

```
show ports port_list information detail
```

The following is sample output from this command:

```
Port: 1
Virtual-router: VR-Default
Type: UTP
Random Early drop: Disabled
Admin state: Enabled with auto-speed sensing auto-duplex
Link State: Active, 100Mbps, full-duplex
Link Counter: Up 1 time(s)
VLAN cfg:
Name: Default, Internal Tag = 1(MAC-Based), MAC-limit = No-limit
...<truncated output>
Egress 802.1p Replacement: Disabled
NetLogin: Enabled
NetLogin authentication mode: Mac based
NetLogin port mode: MAC based VLANs
Smart redundancy: Enabled
Software redundant port: Disabled
auto-polarity: Enabled
```

The added output displays information about the mode of operation for the network login port.

- VLAN cfg—The term MAC-based appears next to the tag number.
- Netlogin port mode—This output was added to display the port mode of operation. Mac based appears as the network login port mode of operation.

To view information about the ports that are temporarily added in MAC-based mode for network login, due to discovered MAC addresses, use the following command:

```
show vlan detail
```

The following is sample output from this command:

```
VLAN Interface with name Default created by user
Tagging: 802.1Q Tag 1
Priority: 802.1P Priority 0
Virtual router: VR-Default
STPD: s0(Disabled,Auto-bind)
Protocol: Match all unfiltered protocols
Loopback: Disable
NetLogin: Disabled
Rate Shape: Disabled
QosProfile: None configured
Ports: 26. (Number of active ports=2)
Untag: *1um, *2, 3, 4, 5, 6, 7,
8, 9, 10, 11, 12, 13, 14,
15, 16, 17, 18, 19, 20, 21,
22, 23, 24, 25, 26
```



Flags: (*) Active, (!) Disabled, (g) Load Sharing port
 (b) Port blocked on the vlan, (a) Authenticated NetLogin Port
 (u) Unauthenticated NetLogin port, (m) Mac-Based port

The flags associated with network login include:

- a—Indicates an authenticated network login port.
- u—Indicates an unauthenticated network login port.
- m—Indicates that the network login port operates in MAC-based mode.

Example

The following command configures the network login ports mode of operation:

```
configure netlogin ports 1:1-1:10 mode mac-based-vlans
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

configure netlogin ports no-restart

```
configure netlogin ports [all | port_list] no-restart
```

Description

Disables the network login port restart feature.

Syntax Description

all	Specifies all network login ports.
<i>port_list</i>	Specifies one or more network login ports.

Default

The default setting is no-restart; the network login port restart feature is disabled.

Usage Guidelines

Use this command to disable the network login port restart feature on a network login port.



Configure network login port restart on ports with directly attached supplicants. If you use a hub to connect multiple supplicants, only the last unauthenticated supplicant causes the port to restart.

Displaying the Port Restart Configuration

To display the network login settings on the port, including the configuration for port restart, use the following command:

```
show netlogin port port_list
```

Output from this command includes the enable/disable state for network login port restart.

Example

The following command disables network login port restart on port 1:1:

```
configure netlogin ports 1:1 no-restart
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure netlogin ports restart

```
configure netlogin ports [all | port_list] restart
```

Description

Enables the network login port restart feature.

Syntax Description

all	Specifies all network login ports.
<i>port_list</i>	Specifies one or more network login ports.

Default

The default setting is no-restart; the network login port restart feature is disabled.



Usage Guidelines

Use this command to enable the network login port restart feature on a network login port. This allows network login to restart specific network login-enabled ports when the last authenticated supplicant releases, regardless of the configured protocols on the port.

Configure network login port restart on ports with directly attached supplicants. If you use a hub to connect multiple supplicants, only the last unauthenticated supplicant causes the port to restart.

Displaying the Port Restart Configuration

To display the network login settings on the port, including the configuration for port restart, use the following command:

```
show netlogin port port_list
```

Output from this command includes the enable/disable state for network login port restart.

Example

The following command enables network login port restart on port 1:1:

```
configure netlogin ports 1:1 restart
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure netlogin redirect-page

```
configure netlogin redirect-page url
```

Description

Configures the redirect URL for Network Login.

Syntax Description

<i>url</i>	Specifies the redirect URL for Network Login.
------------	---

Default

The redirect URL default value is “http://www.extremenetworks.com”; the default port value is 80.



Usage Guidelines

In ISP mode, you can configure network login to be redirected to a base page after successful login using this command. If a RADIUS server is used for authentication, then base page redirection configured on the RADIUS server takes priority over this configuration.

You must configure a complete URL starting with `http://` or `https://`

You can also configure a specific port location at a specific target URL location. For example, you can configure a target port 8080 at `extremenetworks.com` with the following command:

```
configure netlogin redirect-page "www.extremenetworks.com:8080"
```

To support `https`, you must first download and install the separate Extreme Networks SSH software module (`ssh.xmod`). This additional module allows you to configure both SSH2 and SSL on the switch. For more information about SSH2, see [Security](#) in the ExtremeXOS Concepts Guide. For information about installing the SSH module, see [Software Upgrade and Boot Options](#) in the *ExtremeXOS Concepts Guide*.

This command applies only to the web-based authentication mode of Network Login.

Example

The following command configures the redirect URL as `http://www.extremenetworks.com/support`:

```
configure netlogin redirect-page http://www.extremenetworks.com/support
```

History

This command was first available in ExtremeXOS 11.1.

Support for HTTPS was introduced in ExtremeXOS 11.2.

Target port support was introduced in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

configure netlogin session-refresh

```
configure netlogin session-refresh {refresh_seconds}
```

Description

Configures network login session refresh.



Syntax Description

<code>refresh_seconds</code>	Specifies the session refresh time for network login in seconds.
------------------------------	--

Default

Enabled, with a value of 180 seconds for session refresh.

Usage Guidelines

Network login sessions can refresh themselves after a configured timeout. After the user has been logged in successfully, a logout window opens which can be used to close the connection by clicking on the Logout link. Any abnormal closing of this window is detected on the switch and the user is logged out after a time interval as configured for session refresh. The session refresh is enabled and set to 360 seconds by default. The value can range from 1 to 3600 seconds. When you configure the network login session refresh for the logout window, ensure that the FDB aging timer is greater than the network login session refresh timer.

This command applies only to the web-based authentication mode of network login.

Example

The following command enables network login session refresh and sets the refresh time to 100 seconds:

```
configure netlogin session-refresh 100
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

configure netlogin vlan

```
configure netlogin vlan vlan_name
```

Description

Configures the VLAN for Network Login.

Syntax Description

<code>vlan</code>	Specifies the VLAN for Network Login.
-------------------	---------------------------------------



Default

N/A.

Usage Guidelines

This command will configure the VLAN used for unauthenticated clients. One VLAN needs to be configured per VR. To change the VLAN, network login needs to be disabled. Network login can only be enabled when a VLAN is assigned (and no ports are configured for it).

By default no VLAN is assigned for network login.

Example

The following command configures the VLAN login as the network login VLAN:

```
configure netlogin vlan login
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure vlan netlogin-lease-timer

```
configure vlan vlan_name netlogin-lease-timer seconds
```

Description

Configures the timer value returned as part of the DHCP response for clients attached to networklogin-enabled ports.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN to which this timer value applies.
<i>seconds</i>	Specifies the timer value, in seconds.

Default

10 seconds.



Usage Guidelines

The timer value is specified in seconds.

This command applies only to the web-based authentication mode of network login.

Example

The following command sets the timer value to 15 seconds for VLAN corp:

```
configure vlan corp netlogin-lease-timer 15
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

create netlogin local-user

```
create netlogin local-user user-name {encrypted} {password} {vlan-vsa [[{tagged | untagged} [vlan_name] | vlan_tag]]} {security-profile security_profile}
```

Description

Creates a local network login user name and password.

Syntax Description

<i>user-name</i>	Specifies a new local network login user name. User names must have a minimum of 1 character and a maximum of 32 characters.
encrypted	The encrypted option is used by the switch to encrypt the password. Do not use this option through the command line interface (CLI).
<i>password</i>	Specifies a local network login user password. Passwords must have a minimum of 0 characters and a maximum of 32 characters.
tagged	Specifies that the client be added as tagged.
untagged	Specifies that the client be added as untagged.
<i>vlan_name</i>	Specifies the name of the destination VLAN.
<i>vlan_tag</i>	Specifies the VLAN ID, tag, of the destination VLAN.
<i>security_profile</i>	Specifies a security profile string during account creation.



Default

N/A.

Usage Guidelines

Use this command to create a local network login account and to configure the switch to use its local database for network login authentication. This method of authentication is useful in the following situations:

- If both the primary and secondary (if configured) RADIUS servers timeout or are unable to respond to authentication requests.
- If no RADIUS servers are configured.
- If the RADIUS server used for network login authentication is disabled.

If any of the above conditions are met, the switch checks for a local user account and attempts to authenticate against that local account.

Extreme Networks recommends creating a maximum of 64 local accounts. If you need more than 64 local accounts, we recommend using RADIUS for authentication. For more information about RADIUS authentication, see the *ExtremeXOS Concepts Guide*.

You can also specify the destination VLAN to enter upon a successful authentication.



Note

If you do not specify a password or the keyword encrypted, you are prompted for one.

Additional Requirements

This command applies only to the web-based and MAC-based modes of network login. 802.1x network login does not support local database authentication.

You must have administrator privileges to use this command. If you do not have administrator privileges, the switch displays a message similar to the following:

```
This user does not have permissions for this command.
```

User names are not case-sensitive. Passwords are case-sensitive. User names must have a minimum of 1 character and a maximum of 32 characters. Passwords must have a minimum of 0 characters and a maximum of 32 characters. If you use RADIUS for authentication, we recommend that you use the same user name and password for both local authentication and RADIUS authentication.

If you attempt to create a user name with more than 32 characters, the switch displays the following messages:

```
%% Invalid name detected at '^' marker. %% Name cannot exceed 32
characters.
```

If you attempt to create a password with more than 32 characters, the switch displays the following message after you re-enter the password:

```
Password cannot exceed 32 characters
```



Modifying an Existing Account

To modify an existing local network login account, use the following command:

```
configure netlogin local-user user-name {vlan-vsa [{tagged | untagged}  
[vlan_name | vlan_tagw]] | none}
```

Displaying Local Network Login Accounts

To display a list of local network login accounts on the switch, including VLAN information, use the following command:

```
show netlogin local-users
```

Example

The following command creates a local network login user name and password:

```
create netlogin local-user megtest
```

After you enter the local network login user name, press [Enter]. The switch prompts you to enter a password (the switch does not display the password):

```
password:
```

After you enter the password, press [Enter]. The switch then prompts you to re-enter the password:

```
Reenter password:
```

The following command creates a local network login user name, password, and associates a destination VLAN with this account:

```
create netlogin local-user accounting vlan-vsa blue
```

As previously described, the switch prompts you to enter and confirm the password.

History

This command was first available in ExtremeXOS 11.2.

The vlan-vsa parameter and associated options were added in ExtremeXOS 11.3.

The security-profile parameter was added in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.



delete netlogin local-user

```
delete netlogin local-user user-name
```

Description

Deletes a specified local network login user name and its associated password.

Syntax Description

<i>user-name</i>	Specifies a local network login user name.
------------------	--

Default

N/A.

Usage Guidelines

Use the `show netlogin local-users` command to determine which local network login user name you want to delete from the system. The `show netlogin local-users` output displays the user name and password in a tabular format.

This command applies only to web-based and MAC-based modes of network login. 802.1x network login does not support local database authentication.

You must have administrator privileges to use this command.

Example

The following command deletes the local network login megtest along with its associated password:

```
delete netlogin local-user megtest
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

disable netlogin

```
disable netlogin [ {dot1x} {mac} {web-based} ]
```



Description

Disables network login modes.

Syntax Description

dot1x	Specifies 802.1x authentication.
mac	Specifies MAC-based authentication.
web-based	Specifies web-based authentication.

Default

All types of authentication are disabled.

Usage Guidelines

Any combination of authentication types can be disabled on the same switch. To enable an authentication mode, use the following command:

```
enable netlogin [{dot1x} {mac} {web-based}]
```

Example

The following command disables MAC-based network login:

```
disable netlogin mac
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

disable netlogin authentication failure vlan ports

```
disable netlogin authentication failure vlan ports [ports | all]
```

Description

Disables the configured authentication failure VLAN on the specified ports.



Syntax Description

all	Specifies all ports included in the authentication failure VLAN.
<i>ports</i>	Specifies one or more ports or slots and ports on which the authentication failure VLAN is enabled.

Default

All ports.

Usage Guidelines

Use this command to disable the configured authentication failure VLAN on either the specified ports, or all ports.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

disable netlogin authentication service-unavailable vlan ports

```
disable netlogin authentication service-unavailable vlan ports [ports | all]
```

Description

Disable the configured authentication service-unavailable VLAN on the specified ports.

Syntax Description

ports	Specifies one or more ports or slots and ports on which the authentication service-unavailable VLAN is enabled.
all	Specifies all ports included in the authentication service-unavailable VLAN.

Default

All ports.

Usage Guidelines

Use this command to disable the configured authentication service-unavailable VLAN on the specified ports, or on all ports.



History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

disable netlogin dot1x guest-vlan ports

```
disable netlogin dot1x guest-vlan ports [all | ports]
```

Description

Disables the guest VLAN on the specified 802.1x network login ports.

Syntax Description

all	Specifies all ports included in the guest VLAN.
<i>ports</i>	Specifies one or more ports included in the guest VLAN.

Default

Disabled.

Usage Guidelines

Use this command to disable the guest VLAN feature.

Enabling Guest VLANs

To enable the guest VLAN, use the following command:

```
enable netlogin dot1x guest-vlan ports [all | ports]
```

Example

The following command disables the guest VLAN on all ports:

```
disable netlogin dot1x guest-vlan ports all
```

History

This command was first available in ExtremeXOS 11.2.



Platform Availability

This command is available on all platforms.

disable netlogin logout-privilege

```
disable network login logout-privilege
```

Description

Disables network login logout window pop-up.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command controls the logout window pop-up on the web-based network client. This command applies only to the web-based authentication mode of network login. When disabled, the logout window pop-up will no longer appear. However, if session refresh is enabled, the login session will be terminated after the session refresh timeout.

Example

The following command disables network login logout-privilege:

```
disable netlogin logout-privilege
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

disable netlogin ports

```
disable netlogin ports ports [ {dot1x} {mac} {web-based} ]
```



Description

Disables network login on a specified port for a particular method.

Syntax Description

<i>ports</i>	Specifies the ports for which network login should be disabled.
dot1x	Specifies 802.1x authentication.
mac	Specifies MAC-based authentication.
web-based	Specifies web-based authentication.

Default

Network login is disabled by default.

Usage Guidelines

Network login must be disabled on a port before you can delete a VLAN that contains that port.

This command applies to the MAC-based, web-based, and 802.1x mode of network login. To control which authentication mode is used by network login, use the following commands:

```
enable netlogin [{dot1x} {mac} {web-based}] disable netlogin [{dot1x}
{mac} {web-based}]
```

Example

The following command disables dot1x and web-based network login on port 2:9:

```
disable netlogin ports 2:9 dot1x web-based
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

disable netlogin reauthenticate-on-refresh

```
disable netlogin reauthenticate-on-refresh
```

Description

Disables network login reauthentication on refresh.



Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

The web-based Netlogin client's session is periodically refreshed by sending an HTTP request which acts as a keep-alive without actually re-authenticating the user's credentials with the back-end RADIUS server or local database. If reauthenticate-on-refresh is enabled, re-authentication occurs with the session refresh.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

disable netlogin redirect-page

disable netlogin redirect-page

Description

Disables the network login redirect page function.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command disables the network login redirect page so that the client is sent to the originally requested page.

History

This command was first available in ExtremeXOS 12.1.



Platform Availability

This command is available on all platforms.

disable netlogin session-refresh

disable netlogin session-refresh

Description

Disables network login session refresh.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Network login sessions can refresh themselves after a configured timeout. After the user has been logged in successfully, a logout window opens which can be used to close the connection by clicking on the LogOut link. Any abnormal closing of this window is detected on the switch and the user is logged out after a time interval as configured for session refresh. The session refresh is enabled and set to three minutes by default.

This command applies only to the web-based authentication mode of network login.

Example

The following command disables network login session refresh:

```
disable netlogin session-refresh
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.



enable netlogin

```
enable netlogin [{dot1x} {mac} {web-based}]
```

Description

Enables network login authentication modes.

Syntax Description

dot1x	Specifies 802.1x authentication.
mac	Specifies MAC-based authentication.
web-based	Specifies web-based authentication.

Default

All types of authentication are disabled.

Usage Guidelines

Any combination of types of authentication can be enabled on the same switch. At least one of the authentication types must be specified on the command line.

To disable an authentication mode, use the following command:

```
disable netlogin [{dot1x} {mac} {web-based}]
```

Example

The following command enables web-based network login:

```
enable netlogin web-based
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

enable netlogin authentication failure vlan ports

```
enable netlogin authentication failure vlan ports [ports | all]
```



Description

Enables the configured authentication failure VLAN on the specified ports.

Syntax Description

all	Specifies all ports included in the authentication failure VLAN.
<i>ports</i>	Specifies one or more ports or slots and ports on which the authentication failure VLAN is enabled.

Default

All ports.

Usage Guidelines

Use this command to enable the configured authentication failure VLAN on either the specified ports, or all ports.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

enable netlogin authentication service-unavailable vlan ports

```
enable netlogin authentication service-unavailable vlan ports [ports | all]
```

Description

Enables the configured authentication service-unavailable VLAN on the specified ports.

Syntax Description

ports	Specifies one or more ports or slots and ports on which the service-unavailable VLAN is enabled.
all	Specifies all ports included in the service-unavailable VLAN.

Default

All ports.



Usage Guidelines

Use this command to enable the configured authentication service-unavailable VLAN on the specified ports, or on all ports.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

enable netlogin dot1x guest-vlan ports

```
enable netlogin dot1x guest-vlan ports [all | ports]
```

Description

Enables the guest VLAN on the specified 802.1x network login ports.

Syntax Description

all	Specifies all ports included in the guest VLAN.
<i>ports</i>	Specifies one or more ports or slots and ports on which the guest VLAN is enabled.

Default

Disabled.

Usage Guidelines

A guest VLAN provides limited or restricted network access if a supplicant connected to a port does not respond to the 802.1x authentication requests from the switch. A port always moves untagged into the guest VLAN.

Modifying the Supplicant Timer

By default, the switch attempts to authenticate the supplicant every 30 seconds for a maximum of three tries. If the supplicant does not respond to the authentication requests, the client moves to the guest VLAN. The number of authentication attempts is a user-configured parameter with allowed values in the range of 1 to 10.

To modify the supplicant response timer, use the following command and specify the `supp-resp-timeout` parameter:



```
configure netlogin dot1x timers [{server-timeout server_timeout} {quiet-  
periodquiet_period} {reauth-periodreauth_period} {reauth-maxmax_num_reauths}]  
{supp-resp-timeoutsupp_resp_timeout}]
```

Creating the Guest VLAN

Before you can enable the guest VLAN on the specified ports, you must create the guest VLAN. To create the guest VLAN, use the following command:

```
configure netlogin dot1x guest-vlan vlan_name {portsport_list}
```

Example

The following command enables the guest VLAN on all ports:

```
enable netlogin dot1x guest-vlan ports all
```

The following command enables the guest VLAN on ports 2 and 3:

```
enable netlogin dot1x guest-vlan ports 2,3
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable netlogin logout-privilege

```
enable netlogin logout-privilege
```

Description

Enables network login logout pop-up window.

Syntax Description

This command has no arguments or variables.

Default

Enabled.



Usage Guidelines

This command controls the logout window pop-up on the web-based network client. This command applies only to the web-based authentication mode of network login.

Example

The following command enables network login logout-privilege:

```
enable netlogin logout-privilege
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

enable netlogin ports

```
enable netlogin ports ports [{dot1x} {mac} {web-based}]
```

Description

Enables network login on a specified port for a particular authentication method.

Syntax Description

<i>ports</i>	Specifies the ports for which network login should be enabled.
dot1x	Specifies 802.1x authentication.
mac	Specifies MAC-based authentication.
web-based	Specifies web-based authentication.

Default

All methods are disabled on all ports.

Usage Guidelines

For campus mode network login with web-based clients, the following conditions must be met:

- A DHCP server must be available, and a DHCP range must be configured for the port or ports in the VLAN on which you want to enable Network Login.
- The switch must be configured as a RADIUS client, and the RADIUS server must be configured to enable the network login capability.



For ISP mode login, no special conditions are required. A RADIUS server must be used for authentication.

Network login is used on a per port basis. A port that is tagged can belong to more than one VLAN. In this case, network login can be enabled on one port for each VLAN.

Windows authentication is not supported via network login.

Example

The following command configures network login on port 2:9 using web-based authentication:

```
enable netlogin ports 2:9 web-based
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

enable netlogin reauthentication-on-refresh

```
enable netlogin reauthentication-on-refresh
```

Description

Enables network login reauthentication on refresh.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

The web-based Netlogin client's session is periodically refreshed by sending a HTTP request which acts as a keep-alive without actually re-authenticating the user's credentials with the back-end RADIUS server or local database. If reauthenticate-on-refresh is enabled, re-authentication occurs with the session refresh.



History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

enable netlogin redirect-page

enable netlogin redirect-page

Description

Enables the network login redirect page function.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command enables the network login redirect page so that the client is sent to the redirect page rather than the original page.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

enable netlogin session-refresh

enable netlogin session-refresh *{refresh_minutes}*

Description

Enables network login session refresh.



Syntax Description

<i>refresh_minutes</i>	Specifies the session refresh time for network login in minutes.
------------------------	--

Default

Enabled, with a value of three minutes for session refresh.

Usage Guidelines

Network login sessions can refresh themselves after a configured timeout. After the user has been logged in successfully, a logout window opens which can be used to close the connection by clicking on the Logout link. Any abnormal closing of this window is detected on the switch and the user is logged out after a time interval as configured for session refresh. The session refresh is enabled and set to three minutes by default. The value can range from 1 to 255 minutes. When you configure the network login session refresh for the logout window, ensure that the FDB aging timer is greater than the network login session refresh timer.

This command applies only to the web-based authentication mode of network login.

To reset the session refresh value to the default behavior, use this command without the minutes parameter.

Example

The following command enables network login session refresh and sets the refresh time to ten minutes:

```
enable netlogin session-refresh 10
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

show banner netlogin

```
show banner netlogin
```

Description

Displays the user-configured banner string for network login.



Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to view the banner that is displayed on the network login page.

Example

The following command displays the network login banner:

```
show banner netlogin
```

If a custom banner web page exists, show banner netlogin generates the following output:

```
***** Testing NETLOGIN BANNER at <system name>***** NOTE:
Banner is not in use. Overridden since custom login page
"netlogin_login_page.html" is present.
```

If a custom banner web page does not exist, show banner netlogin generates the following output:

```
***** Testing NETLOGIN BANNER at <system name>*****
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

show netlogin

```
show netlogin {port port_list vlan vlan_name} {dot1x {detail}} {mac} {web-based}
```

Description

Shows status information for network login.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>vlan_name</i>	Specifies the name of a VLAN.



dot1x	Specifies 802.1x information.
detail	Shows detailed information.
mac	Specifies MAC-based information.
web-based	Specifies web-based information.

Default

N/A.

Usage Guidelines

Depending on your configuration, software version, and the parameters you choose to display, the information reported by this command may include some or all of the following:

- Whether network login is enabled or disabled.
- The base-URL.
- The default redirect page.
- The logout privileges setting.
- The network login session-refresh setting and time.
- The MAC and IP address of supplicants.
- The type of authentication, 802.1x, MAC-based, or HTTP (web-based).
- The guest VLAN configurations, if applicable.
- The dynamic VLAN state and uplink ports, if configured.
- Whether network login port restart is enabled or disabled.
- Which order of authentication protocols is currently being used.

If you do not specify the authentication method, the switch displays information for all network login authentication methods.

Example

The following command shows the summary network login information:

```
show netlogin
```

The following is sample output from this command:

```
NetLogin Authentication Mode : web-based ENABLED; 802.1x ENABLED; mac-based
ENABLED
NetLogin VLAN                : "nvlan"
NetLogin move-fail-action    : Authenticate
NetLogin Client Aging Time   : 5 minutes
Dynamic VLAN Creation        : Enabled
Dynamic VLAN Uplink Ports    : 12
-----
Web-based Mode Global Configuration
-----
```



```

Base-URL                : network-access.com
Default-Redirect-Page   : http://www.yahoo.com
Logout-privilege       : YES
Netlogin Session-Refresh : ENABLED; 3 minutes
Authentication Database : Radius, Local-User database
-----

```

```

-----
802.1x Mode Global Configuration
-----

```

```

Quiet Period           : 60
Supplicant Response Timeout : 30
Re-authentication period : 200
RADIUS server timeout  : 30
EAPOL MPDU version to transmit : v1
Authentication Database : Radius
-----

```

```

-----
MAC Mode Global Configuration
-----

```

MAC Address/Mask	Password (encrypted)	Port(s)
00:00:86:3F:1C:35/48	yaqu	any
00:01:20:00:00:00/24	yaqu	any
00:04:0D:28:45:CA/48	=4253C5;500@	any
00:10:14:00:00:00/24	yaqu	any
00:10:A4:A9:11:3B/48	yaqu	any
00:10:A4:00:00:00/24	yaqu	any
Default	yaqu	any

```

Authentication Database : Radius, Local-User database
-----

```

```

Port: 5, Vlan: nvlan, State: Enabled, Authentication: mac-based, Guest
Vlan <Not Configured>: Disabled

```

MAC	IP address	Authenticated	Type	ReAuth-Timer	User

```

Port: 9, Vlan: nvlan, State: Enabled, Authentication: web-based, Guest
Vlan <Not Configured>: Disabled

```

MAC	IP address	Authenticated	Type	ReAuth-Timer	User

```

Port: 10, Vlan: nvlan, State: Enabled, Authentication: 802.1x, mac-based,
Guest Vlan <Not Configured>: Disabled

```

MAC	IP address	Authenticated	Type	ReAuth-Timer	User

```

Port: 17, Vlan: engr, State: Enabled, Authentication: mac-based, Guest
Vlan <Not Configured>: Disabled

```

MAC	IP address	Authenticated	Type	ReAuth-Timer	User

```

Port: 17, Vlan: mktg, State: Enabled, Authentication: mac-based, Guest
Vlan <Not Configured>: Disabled

```

MAC	IP address	Authenticated	Type	ReAuth-Timer	User

```

Port: 19, Vlan: corp, State: Enabled, Authentication: 802.1x, Guest Vlan
<Not Configured>: Disabled

```

MAC	IP address	Authenticated	Type	ReAuth-Timer	User
00:04:0d:50:e1:3a	0.0.0.0	No		0	
00040D50E13A					
00:10:dc:98:54:00	10.201.31.113	Yes, Radius	802.1x	24	md5isp7

```

-----

```



```

Port: 19, Vlan: nvlan, State: Enabled, Authentication: 802.1x, Guest Vlan
<Not Configured>: Disabled
MAC          IP address      Authenticated  Type      ReAuth-Timer  User
00:04:0d:50:e1:3a  0.0.0.0         No             802.1x    0
-----
Port: 19, Vlan: voice-ip, State: Enabled, Authentication: 802.1x, Guest
Vlan <Not Configured>: Disabled
MAC          IP address      Authenticated  Type      ReAuth-Timer  User
00:04:0d:50:e1:3a  0.0.0.0         Yes, Radius    802.1x    75
00040D50E13A
-----

```

The following command shows more detailed information, including the configured authentication methods:

```
show netlogin port 3:2 vlan "Default"
```

The following is sample output from this command:

```

Port: 2:1      Vlan: Default
Authentication: Web-Based, 802.1x
Port State:    Unauthenticated
Guest VLAN:    Not Enabled
DHCP:          Not Enabled
MAC           IP address     Auth  Type      ReAuth-Timer  User
00:0C:F1:E8:4E:13  0.0.0.0       No   802.1x    0              Unknown
00:01:30:F3:EA:A0  10.0.0.1      Yes  802.1x    0              testUser

```

The following command shows information about a specific port configured for network login:

```
show netlogin port 1:1
```

The following is sample output from this command:

```

Port          : 1:1
Port Restart  : Enabled
Vlan          : Default
Authentication: mac-based
Port State    : Enabled
Guest Vlan    : Disabled
MAC           IP address     Auth  Type      ReAuth-Timer  User
-----

```

The following command shows the details of the 802.1x mode:

```
show netlogin dot1x detail
```



The following is sample output from this command:

```

NetLogin Authentication Mode : web-based DISABLED; 802.1x ENABLED; mac-
based DISABLED
NetLogin VLAN                : "nl"
NetLogin move-fail-action    : Deny
-----
802.1x Mode Global Configuration
-----
Quiet Period                  : 30
Supplicant Response Timeout  : 30
Re-authentication period     : 3600
RADIUS server timeout        : 30
EAPOL MPDU version to transmit : v1
Guest VLAN                   : destVlan
-----
Port: 1:1, Vlan: Default, State: Enabled, Authentication: 802.1x, Guest
Vlan: destVlan
MAC
00:00:86:53:c3:14   : IP=0.0.0.0           Auth=Yes User= testUser
: AuthPAE state=AUTHENTICATED BackAuth state=IDLE
: ReAuth time left=3595      ReAuth count=1
: Quiet time left=37
00:02:03:04:04:05   : IP=0.0.0.0           Auth=No  User=
: AuthPAE state=CONNECTING  BackAuth state=IDLE
: ReAuth time left=0        ReAuth count=2
: Quiet time left=37
-----

```

For 802.1x, if re-authentication is disabled, the re-authentication period appears as follows:

```

Re-authentication period      : 0 (Re-authentication disabled)

```

The following command:

```

show netlogin port 5:4 dot1x

```

Generates the following sample output:

```

Port                : 5:4
Port Restart        : Disabled
Vlan                : corp
Authentication      : 802.1x
Port State          : Enabled
Guest Vlan          : Enabled
MACIP addressAuthenticatedTypeReAuth-TimerUser
00:10:dc:92:53:2d10.201.31.119Yes,Radius802.1x14md5isp4
-----

```



The following command:

```
sh netlogin port 5:4 dot1x detail
```

Generates the following sample output:

```
Port: 5:4
Port Restart: Disabled
Vlan: corp
Authentication: 802.1x
Port State: Enabled
Guest Vlan: Enabled
MAC
00:10:dc:92:53:2d : IP=10.201.31.119 Auth=Yes User=md5isp4
: AuthPAE state=AUTHENTICATED BackAuth state=IDLE
: ReAuth time left=8 ReAuth count=0
: Quiet time left=0
-----
```

History

This command was first available in ExtremeXOS 11.1.

Information about the guest VLAN was added in ExtremeXOS 11.2.

Information about the configured port MAC list was added in ExtremeXOS 11.3.

Information about dynamic VLANs and network login port restart was added in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

show netlogin authentication failure vlan

```
show netlogin authentication failure vlan {vlan_name}
```

Description

Displays the authentication failure VLAN related configuration details.

Syntax Description

<i>vlan_name</i>	Specifies the name of a failure VLAN.
------------------	---------------------------------------

Default

N/A.



Usage Guidelines

Use this command to display configuration details for the authentication failure VLAN.

Example

The following command displays:

```
show netlogin authentication failure vlan
```

The following is sample output from this command:

```
-----
Authentication Service unavailable
Vlan
port                Status
-----
corp
1:2                 Disabled
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

show netlogin authentication service-unavailable vlan

```
show netlogin authentication service-unavailable vlan {vlan_name}
```

Description

Displays the authentication service-unavailable VLAN related configuration details.

Syntax Description

<i>vlan_name</i>	Specifies the name of the authentication service-unavailable VLAN.
------------------	--

Default

N/A.



Usage Guidelines

Use this command to display configuration details for the service-unavailable VLAN.

Example

The following command displays:

```
show netlogin authentication service-unavailable vlan
```

The following is sample output from this command:

```
-----
---
server-unavailable Vlanport                Status
-----
---
xyz 2:1                                Disabled
abc3:1                                Enabled
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

show netlogin banner

show netlogin banner

Description

Displays the user-configured banner string for network login.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to view the banner that is displayed on the network login page.



Example

The following command displays the network login banner:

```
show netlogin banner
```

If a custom banner web page exists, show banner netlogin generates the following output:

```
***** Testing NETLOGIN BANNER at <system name>*****
NOTE: Banner is not in use. Overridden since custom login page
"netlogin_login_page.html" is present.
```

If a custom banner web page does not exist, nothing is displayed.

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

show netlogin guest-vlan

```
show netlogin guest-vlan {vlan_name}
```

Description

Displays the configuration for the guest VLAN feature.

Syntax Description

vlan_name	Specifies the name of a guest VLAN.
------------------	-------------------------------------

Default

N/A.

Usage Guidelines

Use this command to display the guest VLANs configured on the switch.

If you specify the vlan_name, the switch displays information for only that guest VLAN.

The output displays the following information in a tabular format:

- Port—Specifies the 802.1x enabled port configured for the guest VLAN.



- Guest-vlan—Displays the enabled/disabled state of the guest VLAN feature.
- Vlan—Specifies the name of the guest VLAN.

Example

The following command displays the local network login list:

```
show netlogin guest-vlan
```

The following is sample output from this command:

Port	Guest-vlan	Vlan
5:1	Disabled	gv11
5:2	Enabled	gv12
5:3	Disabled	gv13
5:4	Enabled	gv14

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

show netlogin local-users

```
show netlogin local-users
```

Description

Displays the local network login users configured on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to display the list of local network login users and associated VLANs.



If you associated a VLAN with a local network login user, the output displays the name of the VLAN. If you have not associated a VLAN with a local network login user, the output displays not configured.

The Extended-VLAN VSA column displays the name of the VLAN and the following information:

- *not configured*—Specifies that you have not associated a VLAN with a local network login user.
- *—Specifies the movement based on the incoming port's traffic. For example, the VLAN behaves like VSA 203 if identified with a VLAN name or VSA 209 if identified with a VLAN ID.
- T—Specifies a tagged client.
- U—Specifies an untagged client.

In addition, this output is useful to determine which local network login user you want to modify or delete from the system.

Example

The following command displays the local network login list:

```
show netlogin local-users
```

The following is sample output from this command:

Netlogin Local User Name	Password (encrypted)	Extended-VLAN VSA
000000000012	Iqyydz\$MP7AG.VAmwOoqiKX2u13H1	U hallo
00008653C314	BoO28L\$oRVvKv8.wmxcorhhXxQY40	* default
megtest	w7iMbp\$lBL34/dLx4G4M8aAdiCvI	<not configured>
testUser	/Jhouw\$iHE15steebwhOibgj6pZq.	T testVlan

History

This command was first available in ExtremeXOS 11.2.

The output was modified to include VLAN information in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

show netlogin mac-list

```
show netlogin mac-list
```

Description

Displays the MAC address list for MAC-based network login.



Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to display the MAC address list used for MAC-based network login.

MAC-based authentication is VR aware, so there is one MAC list per VR.

Example

The following command displays the MAC address list:

```
show netlogin mac-list
```

The following is sample output from this command:

MAC Address/Mask	Password (encrypted)	Port(s)
00:00:00:00:00:10/48	<not configured>	1:1-1:5
00:00:00:00:00:11/48	<not configured>	1:6-1:10
00:00:00:00:00:12/48	<not configured>	any
00:01:30:70:0C:00/48	yaquany	
00:01:30:32:7D:00/48	ravdqsrany	
00:04:96:00:00:00/24	<not configured>	any

History

This command was first available in ExtremeXOS 11.1.

Information about the configured port MAC list was added in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

unconfigure netlogin allowed-refresh-failures

```
unconfigure netlogin allowed-refresh-failures
```

Description

Restores the number refresh failures to the default value.



Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command allows you to restore the number of refresh failures allowed to the default value of 0.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

unconfigure netlogin authentication database-order

```
unconfigure netlogin [mac | web-based] authentication database-order
```

Description

Restores the default order of database authentication protocols to use.

Syntax Description

mac	Specifies the MAC address to add.
mask	Specifies the number of bits to use for the mask.
default	Specifies the default entry.
encrypted	Used to display encrypted form of password in configuration files. Do not use.
password	Specifies the password to send for authentication.
ports	Specifies the port or port list to use for authentication.

Default

By default, the authentication order is RADIUS, local-user database.

Usage Guidelines

Use this command to restore the default configuration order for the database authentication protocols. For details see [Feature License Requirements](#).



Example

The following command sets the database authentication order to RADIUS, local user database for MAC-based authentication:

```
unconfigure netlogin mac authentication database-order
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

unconfigure netlogin authentication failure vlan

```
unconfigure netlogin authentication failure vlan vlan_name {ports port_list}
```

<i>vlan_name</i>	Specifies the name of the authentication failure VLAN.
<i>port_list</i>	Specifies one or more ports or slots and ports. If the ports keyword is not used, the command applies to all ports.

Description

Disables authentication failure VLAN on network login enabled ports.

Syntax Description

Default

N/A.

Usage Guidelines

Use this command to disable authentication failure VLAN on network login enabled ports. When a supplicant fails authentication, it is moved to the authentication failure vlan and is given limited access until it passes the authentication.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.



unconfigure netlogin authentication service-unavailable vlan

```
unconfigure netlogin authentication service-unavailable vlan vlan_name {ports
port_list}
```

Description

Unconfigures authentication service unavailable VLAN on network login enabled ports.

Syntax Description

<i>vlan_name</i>	Specifies the name of the authentication service-unavailable VLAN.
<i>port_list</i>	Specifies one or more ports or slots and ports. If the ports keyword is not used, the command applies to all ports.

Default

Defaults to all network login enabled ports.

Usage Guidelines

This command unconfigures authentication service unavailable VLAN on the specified network login enabled ports. Authentication service unavailable VLAN is unconfigured on all the network login enabled ports, if no port is specifically mentioned.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

unconfigure netlogin banner

```
unconfigure netlogin banner
```

Description

Unconfigures the network login page banner.

Syntax Description

This command has no arguments or variables.



Default

The default banner is the Extreme Networks logo.

Usage Guidelines

Use this command to unconfigure a netlogin banner.

After the command is issued, the configured banner specified is no longer displayed.

Example

The following command unconfigures the network login page banner:

```
unconfigure netlogin banner
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

unconfigure netlogin dot1x guest-vlan

```
unconfigure netlogin dot1x guest-vlan {ports port_list | vlan_name}
```

Description

Unconfigures the guest VLAN feature for 802.1x authentication.

Syntax Description

<i>port_list</i>	Specifies one or more ports included in the guest VLAN.
<i>vlan_name</i>	Specifies all ports included in the guest VLAN.

Default

N/A.

Usage Guidelines

Use this command to unconfigure the guest VLAN for 802.1x authentication.



If you do not specify one or more ports or the VLAN name, this command unconfigures all of the 802.1x ports configured for the guest VLAN feature.

If you specify one or more ports, this command unconfigures the specified 802.1x ports for the guest VLAN feature.

If you specify the VLAN name, this command unconfigures all of the 802.1x ports configured for the specified guest VLAN.

Example

The following command unconfigures the guest VLAN feature for 802.1x authentication:

```
unconfigure netlogin dot1x guest-vlan
```

History

This command was first available in ExtremeXOS 11.2.

The ports option was added in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

unconfigure netlogin local-user security-profile

```
unconfigure netlogin local-user user-name security-profile
```

Description

Clears a previously associated security profile.

Syntax Description

<i>user-name</i>	Specifies the name of an existing local network login account.
------------------	--

Default

N/A.

Usage Guidelines

Use this command to clear any previously associated security profiles on the switch.



History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

unconfigure netlogin session-refresh

unconfigure netlogin session-refresh

Description

Restores the session refresh value to the default.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command allows you to restore the session refresh to the default value of 180 seconds.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

unconfigure netlogin vlan

unconfigure netlogin vlan

Description

Unconfigures the VLAN for network login.



Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command unconfigures the VLAN used for unauthenticated clients. One VLAN needs to be configured per VR. To change the VLAN, network login needs to be disabled.

Example

The following command unconfigures the network login VLAN:

```
unconfigure netlogin vlan
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.



25 Commands for Identity Management

```
clear counters identity-management
configure identity-management access-list
configure identity-management blacklist
configure identity-management database memory-size
configure identity-management detection
configure identity-management greylist
configure identity-management kerberos snooping aging time
configure identity-management kerberos snooping force-aging time
configure identity-management kerberos snooping forwarding
configure identity-management kerberos snooping server
configure identity-management list-precedence
configure identity-management ports
configure identity-management role
configure identity-management role add child-role
configure identity-management role add dynamic-rule
configure identity-management role add policy
configure identity-management role-based-vlan
configure identity-management role delete child-role
configure identity-management role delete dynamic-rule
configure identity-management role delete policy
configure identity-management role match-criteria inheritance
configure identity-management role priority
configure identity-management stale-entry aging-time
configure identity-management whitelist
configure ldap domain
configure ldap domain add server
configure ldap domain base-dn
configure ldap domain bind-user
configure ldap domain delete server
configure ldap domain netlogin
create identity-management role
create ldap domain
configure ldap hierarchical-search-oid
delete identity-management role
delete ldap domain
```

```
disable identity-management
disable snmp traps identity-management
enable identity-management
enable snmp traps identity-management
refresh identity-management role
show identity-management
show identity-management blacklist
show identity-management entries
show identity-management greylist
show identity-management list-precedence
show identity-management role
show identity-management statistics
show identity-management whitelist
show ldap domain
show ldap statistics
unconfigure identity-management
unconfigure identity-management list-precedence
unconfigure ldap domains
```

This chapter describes commands for:

- Role-based user management
- Enabling and disabling identity management
- Configuring and unconfiguring identity management
- Managing kerberos snooping with identity management
- Displaying identity management information
- Clearing the identity management counters

clear counters identity-management

```
clear counters identity-management
```

Description

Clears the identity management feature counters.

Syntax Description

This command has no arguments or variables.

Default

N/A.



Usage Guidelines

This command clears the following identity management statistics counters:

- High memory usage level reached count
- Critical memory usage level reached count
- Max memory usage level reached count
- Normal memory usage level trap sent
- High memory usage level trap sent
- Critical memory usage level trap sent
- Max memory usage level trap sent
- Event notification sent

You can view these counters with the `show identity-management statistics` command.



Note

The clear counters command also clears these counters. The following counters relate to active entries and are not cleared: Total number of users logged in, Total number of login instances, and Total memory used.

Example

The following command clears the identity management feature counters:

```
Switch.4 # clear counters identity-management
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

configure identity-management access-list

```
configure identity-management access-list source-address [mac | ip]
```

Description

Configures the access-list source-address type.

Syntax Description

mac	Specifies MAC addresses.
ip	Specifies IP addresses.



Default

MAC addresses.

Usage Guidelines

The identity management feature can install ACLs for identities based on the source MAC or source IP address. By default the MAC address of the identity is used to install the ACLs. Every network entity has a MAC address, but not all network devices have an IP address, so we recommend that you use the default mac selection to install ACLs for network entities based on the source MAC address.

You must disable the identity management feature with the `disable identity-management` command before you use this command.

Example

The following command configures the identity management feature to use MAC-based ACLs:

```
* Switch.4 # configure identity-management access-list source-address mac
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on all platforms.

configure identity-management blacklist

```
configure identity-management blacklist add [mac mac_address {macmask} | ip
ip_address {netmask} | ipNetmask] | user user_name] configure identity-management
blacklist delete [all | mac mac_address {macmask} | ip ip_address {netmask} |
ipNetmask] | user user_name]
```

Description

Adds or deletes an entry in the identity manager blacklist.

Syntax Description

add	Adds the specified identity to the blacklist.
delete	Deletes the specified identity from the blacklist.
all	Specifies that all identities are to be deleted from the blacklist. This option is available only when the delete attribute is specified.
<i>mac_address</i>	Specifies an identity by MAC address.



<i>macmask</i>	Specifies a MAC address mask. For example: FF:FF:FF:00:00:00.
<i>ip_address</i>	Specifies an identity by IP address.
<i>netmask</i>	Specifies a mask for the specified IP address.
<i>ipNetmask</i>	Specifies an IP network mask.
<i>user_name</i>	specifies an identity by user name.

Default

N/A.

Usage Guidelines

The software supports up to 512 entries in the blacklist. When you add an identity to the blacklist, the switch searches the whitelist for the same identity. If the identity is already in the whitelist, the switch displays an error.

It is possible to configure an identity in both lists by specifying different attributes in each list. For example, you can add an identity username to the blacklist and add the MAC address for that user's laptop in the whitelist. Because the blacklist has priority over the whitelist, the username is denied access to the switch from all locations.

If you add a new blacklist entry that is qualified by a MAC or IP address, the identity manager does the following:

- Reviews the identities already known to the switch. If the new blacklist entry is an identity known on the switch, all existing ACLs (based on user roles or whitelist configuration) for the identity are removed.
- When a blacklisted MAC-based identity is detected or already known, a Deny All ACL is programmed for the identity MAC address for the port on which the identity is detected.
- When a blacklisted IP-based identity is detected or already known, a Deny All ACL is programmed for the identity IP address for the port on which the identity is detected.
- The ACL for blacklisted MAC and IP addresses precedes any ACLs based on user names (including Kerberos snooping) that may have been previously configured on the port. This ensures that a Kerberos exchange cannot complete when initiated for blacklisted identities.

If you add a new blacklist entry that is qualified by a username (with or without a domain name), the identity manager does the following:

- Reviews the identities already known to the switch. If the new blacklist entry is an identity known on the switch, a Deny All ACL is programmed for the identity MAC address on all ports to which the identity is connected.
- When a new blacklisted username-based identity accesses the switch, a Deny All ACL is programmed for the identity MAC address on the port on which the identity was detected.



- The ACL for a blacklisted username follows any ACLs based on Kerberos snooping. This ensures that a Kerberos exchange for another user can complete when initiated from the same MAC address.



Note

Identity manager programs ingress ACLs. Blacklisted devices can receive traffic from the network, but they cannot send traffic into the network.

Deny All ACLs for blacklisted entries exist as long as the identity remains in the identity manager database.

If you delete an identity from the blacklist, identity manager checks to see if the identity is in the local database. If the identity is known to the switch, the switch does the following:

- Removes the Deny All ACL from the port to which the identity connected.
- Initiates the role determination procedure for the switch port to which the known identity connected. This ensures that the appropriate role is applied to the identity that is no longer blacklisted.



Note

The role determination process can trigger an LDAP refresh to collect identity attributes for role determination.

Example

The following command adds a MAC address to the blacklist:

```
* Switch.4 # configure identity-management blacklist add mac 00:01:05:00:03:18
```

The following command deletes a user name from the blacklist:

```
* Switch.5 # configure identity-management blacklist delete user
bill_jacob@b.com
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on all platforms.

configure identity-management database memory-size

```
configure identity-management database memory-size Kbytes
```



Description

Configures the maximum amount of memory that is allocated to the identity management database.

Syntax Description

<i>Kbytes</i>	Specifies the maximum amount of memory to be used for maintaining identity information. The range is 64 to 49152 KB.
---------------	--

Default

512 KB.

Usage Guidelines

If the current memory usage is higher than the memory size specified in the `configure identity-management database memory-size` command, the command is not successful and a warning message appears. The message indicates that the current memory usage level is higher than the configured level and that the memory can be freed only when existing identities log out or disconnect.

Example

The following command allocates 4096 kilobytes to the identity management database:

```
* Switch.4 # configure identity-management database memory-size 4096
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

configure identity-management detection

```
configure identity-management detection [on | off] [fdb | iparp | ipsecurity | kerberos | lldp | netlogin | all] ports [port_list | all]
```

Description

This command provides the administrator a way to enable/disable the detection of the identities that are triggered through any of the following protocols:

- FDB
- IPARP
- IPSecurity DHCP Snooping



- LLDP
- Netlogin
- Kerberos

Syntax Description

detection	Detection of the identities.
on	Detection of identities on.
off	Detection of identities off.
fdb	FDB identities.
iparp	IPARP identities.
ipsecurity	Identities detected through DHCP snooping entries.
kerberos	Kerberos identities.
lldp	LLDP identities.
all	All identities.

Default

On.

Usage Guidelines

The identity manager detects the identities using the following protocols:

- FDB
- IPARP
- IPSECURITY DHCP Snooping
- LLDP
- Netlogin
- Kerberos

By default, Identity Management detects identities through all the above mentioned protocols.

This feature provides the administrator a way to enable/disable the detection of the identities that are triggered through any of the above said protocols. The administrator can control the identity detection through any of the protocol trigger at the port level. This configuration can be applied to identity management enabled ports only. EXOS displays an error if this configuration is applied for the identity management disabled ports.

Note



All types of Netlogin identity will not be detected if the netlogin detection is disabled.
Enabling Kerberos identity detection will not create identities for the previously authenticated Kerberos clients.



Example

```
* Slot-1 Stack.1 # configure identity-management detection off fdb ports 1:3-6
* Slot-1 Stack.2 # configure identity-management detection off ipsecurity
ports 1:3-6
* Slot-1 Stack.3 # configure identity-management detection off kerberos ports
1:1, 2:5-8
* Slot-1 Stack.4 # configure identity-management detection off netlogin ports
1:1-24, 2:1-24
The effect of these commands can be seen by issuing the show identity-
management command
* Slot-1 Stack.5 # show identity-management
Identity Management : Enabled
Stale entry age out (effective) : 180 Seconds (180 Seconds)
Max memory size : 512 Kbytes
Enabled ports : 1:1-24, 2:1-24
FDB Detection Disabled ports : 1:3-6
IPARP Detection Disabled ports : None
IPSecurity Detection Disabled ports : 2:1
Kerberos Detection Disabled ports : 1:1, 2:5-8
LLDP Detection Disabled ports : None
Netlogin Detection Disabled ports : 1:1-24, 2:1-24
SNMP trap notification : Enabled
Access list source address type : IP
Kerberos aging time (DD:HH:MM) : 00:08:00
Kerberos force aging time (DD:HH:MM) : None
Valid Kerberos servers : none configured(all valid)
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on all platforms.

configure identity-management greylist

```
configure identity-management greylist add user username identity-management
greylist delete [all | user username]
```

Description

This command enables a network administrator to choose usernames whose identity is not required to be maintained. These user names are added to greylist. Identity Management module does not create an identity when greylist users log in.

Syntax Description

<i>username</i>	Specifies an identity by user name.
-----------------	-------------------------------------



Default

N/A.

Usage Guidelines

The software supports up to 512 entries in greylist. Administrator can configure username as part of greylist. When such configuration takes place, identity manager takes following action.

- Checks if the same entry is present in blacklist/whitelist. If yes, command is rejected with appropriate error message.
- Checks if this entry is ineffective because of existing entries in blacklist/whitelist. During this check, precedence of greylist is also taken into account.
 - E.g: New entry being configured into greylist is: Richard@corp. Assume blacklist has higher precedence and it has an entry "Richard". In this case, new entry is ineffective and the configuration is rejected giving the details.
- If no conflict is found, greylist is updated.
- IDM checks if any existing identity matches the new entry in greylist. If match is found, location/identity will be deleted and unknown identity is created with the same MAC.

If greylist user is the only user logged into the device, unknown identity is created and user is kept in unauthenticated role. However if actual user is present along with greylist user, no additional policy is applied for greylist user. Greylist user will get access permissions same as that of actual user logged in.

When user deletes an entry from greylist, identity manager will:

1. Delete the entry and updates the list.
2. User identity is constructed based on NetLogin details, if deleted username is found in NetLogin authenticated user database.

Example

The following command adds an username to the greylist:

```
configure identity-management greylist add user Richard@corp
```

The following command deletes an username from the greylist:

```
configure identity-management greylist del user Richard@corp
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms.



configure identity-management kerberos snooping aging time

```
configure identity-management kerberos snooping aging time minutes
```

Description

Specifies the aging time for Kerberos snooping entries.

Syntax Description

<i>minutes</i>	Specifies the aging time in minutes. The range is 1 to 65535 minutes.
----------------	---

Default

N/A.

Usage Guidelines

Kerberos does not provide any service for un-authentication or logout. Kerberos does provide a ticket lifetime, but that value is encrypted and cannot be detected during snooping.

To enable the aging and removal of snooped Kerberos entries, this timer defines a maximum age for the snooped entry. When a MAC address with a corresponding Kerberos entry in Identity Manager is aged out, the Kerberos snooping timer starts. If the MAC address becomes active before the Kerberos snooping timer expires, the timer is reset and the Kerberos entry remains active. If the MAC address is inactive when the Kerberos snooping timer expires, the Kerberos entry is removed.

Example

The following command configures the aging time for 600 minutes:

```
* Switch.4 # configure identity-management kerberos snooping aging time 600
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

configure identity-management kerberos snooping force-aging time

```
configure identity-management kerberos snooping force-aging time [none | minutes]
```



Description

Configures the switch to remove all Kerberos snooping entries after the specified time expires.

Syntax Description

<i>minutes</i>	Specifies the aging time in minutes. The range is 1 to 65535 minutes.
none	Disables the Kerberos force-aging feature.

Default

N/A.

Usage Guidelines

If Kerberos force aging is enabled, we recommend that the Kerberos snooping force aging time be set to the same value as the Kerberos ticket lifetime.

Example

The following command removes all Kerberos snooping entries after 600 minutes:

```
* Switch.4 # configure identity-management kerberos snooping force-aging time
600
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on all platforms.

configure identity-management kerberos snooping forwarding

```
configure identity-management kerberos snooping forwarding [fast-path | slow-
path]
```

Description

When identity management is enabled on a port, Kerberos packets are software-forwarded. With this command, you can report if shared folder access via identity management-enabled ports is slow if there exists other CPU-bound traffic.



Syntax Description

forwarding	Configure how customer Kerberos authentication packets are forwarded by this system.
fast-path	Forward customer snooped Kerberos packets in hardware (default).
slow-path	Forward customer snooped Kerberos packets in software. This option is recommended only for systems with low CPU-bound traffic.

Default

Fast-path.

Usage Guidelines

Use this command to report if shared folder access via identity management-enabled ports is slow if there exists other CPU-bound traffic.

Example

The following show command displays the modified Kerberos information:

```
X460-48p.14 # sh identity-management
Identity Management           : Enabled
Stale entry age out (effective) : 180 Seconds (180 Seconds)
Max memory size              : 512 Kbytes
Enabled ports                 : 1
SNMP trap notification       : Enabled
Access list source address type : MAC
Kerberos aging time (DD:HH:MM) : None
Kerberos force aging time (DD:HH:MM) : None
Kerberos snooping forwarding : Fast path
Kerberos snooping forwarding : Slow path
Valid Kerberos servers       : none configured(all valid)
LDAP Configuration:
-----
LDAP Server      : No LDAP Servers configured
Base-DN         : None
Bind credential  : anonymous

LDAP Configuration for Netlogin:
dot1x           : Enabled
mac             : Enabled
web-based       : Enabled
```

History

This command was first available in ExtremeXOS 15.1.3.



Platform Availability

This command is available on all platforms.

configure identity-management kerberos snooping server

```
configure identity-management kerberos snooping add server ip_address
```

```
configure identity-management kerberos snooping delete server [ip_address | all]
```

Description

Adds or deletes a Kerberos server to the Kerberos server list.

Syntax Description

<i>ip_address</i>	Specifies a Kerberos server IP address to add or delete.
all	Specifies that all Kerberos server list entries are to be deleted.

Default

No servers are in the Kerberos server list.

Usage Guidelines

When no servers are configured in the Kerberos server list, the Kerberos snooping feature processes responses from all Kerberos servers, which can expose the system to simulated logins. To avoid this exposure, you can configure a list of up to 20 valid Kerberos servers. When the Kerberos server list contains one or more entries, the switch only processes responses from the Kerberos servers in the list.

Example

The following command adds the Kerberos server at IP address 10.10.10.1 to the Kerberos server list:

```
* Switch.4 # configure identity-management kerberos snooping add server
10.10.10.1
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.



configure identity-management list-precedence

```
configure identity-management list-precedence listname1 listname2 listname3
```

Description

This command allows you to configure the precedence of list types. You must specify the list-names in the desired order of precedence. Listname1 will take precedence of all lists (i.e., highest precedence). Listname2 will take precedence over Listname3. When the user/device logs in, entries present in Listname1 will be searched at first to find matching role. Entries present in Listname2 will be searched after Listname1 and entries in Listname3 will be searched at last.

Syntax Description

<i>listname1</i>	Specifies the list type which has precedence over all list types.
<i>listname2</i>	Specifies the list type which has next precedence, after listname1.
<i>listname3</i>	Specifies the list type which has least precedence of all.

Default

greylist, blacklist, whitelist

Usage Guidelines

By default, greylist entries have higher precedence over blacklist and whitelist entries.

This means that IDM consults with greylist first upon detection of user, and then decides if identity needs to be created. If there is a greylist entry matching the incoming username, user identity is not created. If there is no matching greylist entry, IDM proceeds with role identification for the user. However, greylist precedence is configurable. Following are three possibilities for greylist precedence configuration.

1. greylist, blacklist, whitelist
2. blacklist, greylist, whitelist
3. blacklist, whitelist, greylist

It is important to notice that blackist always has higher precedence over whitelist for EXOS 15.1.2. In order to change the list precedence, Identity Management should be disabled first. Disabling IDM is required since there may be many users/devices already mapped to some roles and policies/ACLs applied. Considering the processing load of unmapping the roles and removing policies, changing precedence isn't allowed when IDM is enabled. When precedence configuration is changed, each entry present in the list with lower precedence (new precedence) is checked with each entry present in all the lists with higher precedence.



Example

The following example instructs that blacklist has precedence over all lists. Greylist has precedence over whitelist. Whitelist has least precedence.

```
configure identity-management list-precedence blacklist greylist whitelist
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms.

configure identity-management ports

```
configure identity-management {add | delete} ports [port_list | all]
```

Description

Adds or deletes identity management for the specified ports.

Syntax Description

add	Enables identity management on the specified port list.
delete	Disables identity management on the specified port list.
<i>port_list</i>	Specifies the ports to which this command applies.
all	Specifies that this command applies to all ports.

Default

No ports are in the identity management enabled port list.

Usage Guidelines

If neither the add nor the delete keyword is entered, identity management is enabled on the specified port list, and the new port list overrides any previous port list.

If identity management is enabled on a port and a user or device is connected to it, information about the user or device is present in the identity management database. If this port is removed from the identity-management enabled port list, the user or device information remains in the data base until the



user logs out or the device disconnects. However, once a port is deleted from enabled port list, no new information is added to the identity management database for that port.



Note

Kerberos identities are not detected when both server and client ports are added to identity management.

Example

The following command enables identity management on ports 2:3 and 2:5:

```
configure identity-management add ports 2:3,2:5
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.



configure identity-management role

```
configure identity-management role role_name { tag [ tag | none ] } {vr vr_name | none } }
```

Description

This command defines VLAN/VR membership to an identity management role.

Syntax Description

<i>role_name</i>	Name of the role.
tag	VLAN tag for dynamic VLAN creation for this role.
<i>tag</i>	VLAN tag between 1 and 4094.
vr	Virtual router name for dynamic VLAN creation for this role.
<i>vr_name</i>	Virtual router name.
none	None.

Default

N/A.



Usage Guidelines

Use this command to configure VLAN tag and the VR in which the dynamic VLAN has to be created for a role. By default the dynamic VLAN is created in VR-Default if the VR is not configured. The identity is placed in the base VLAN if no VLAN tag is configured for this role. The configured VLAN tag and VR can be set to none to unconfigure the same. VR-Mgmt is not allowed to configure. The VLAN tag and VR is applicable only to the user created roles.

Example

The following example configures role "r1" and tag 100:

```
x460-48t.10 # configure identity-management role "r1" tag 100 vr "VR-Default"
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

configure identity-management role add child-role

```
configure identity-management role role_name add child-role child_role
```

Description

Adds a child role to the specified role.

Syntax Description

<i>role_name</i>	Specifies the name of an existing role.
child-role	Specifies a name for the new child role (up to 32 characters).

Default

N/A.

Usage Guidelines

The child role name can include up to 32 characters. Role names must begin with an alphabetical letter, and only alphanumeric, underscore (_), and hyphen (-) characters are allowed in the remainder of the name. Role names cannot match reserved keywords. For more information on role name requirements and a list of reserved keywords, see *Object Names* of the *ExtremeXOS Concepts Guide*.

The following guidelines apply to child roles:



- A child role inherits all the policies applied to its parent and any higher levels above the parent.
- The software supports 5 levels of hierarchy.
- Each role can have a maximum of 8 child roles.
- Each child role can have only 1 parent role.

Example

The following example configures a child role named East for the existing role named India-Engr:

```
* Switch.66 # configure identity-management role "India-Engr" add child-role
East
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

configure identity-management role add dynamic-rule

```
configure identity-management role role_name [add dynamic-rule rule_name { first
| last | { [before | after] ref_rule_name}}]
```

Description

Adds a dynamic ACL rule for the specified role and specifies the order.

Syntax Description

<i>role_name</i>	Specifies the name of an existing role.
<i>rule_name</i>	Specifies the name of a dynamic ACL rule to add to the specified role.

Default

The order of the dynamic rule is last if the order is not explicitly specified.

Usage Guidelines

The maximum number of policies or ACL rules that can be applied to a particular role is restricted to 8. This count does not include the policies and rules inherited from a parent role. Since the maximum hierarchy depth is 5, the maximum number of policies and rules supported for a role at the maximum hierarchy depth is 40 (8 x 5).



When a dynamic ACL rule is added to a role, it is immediately installed for all identities mapped to that role and roles below it in the role hierarchy.

Example

The following example configures the role named India-Engr to use the ACL rule named india-Engr-rule:

```
* Switch.55 # configure identity-management role "India-Engr" add dynamic-
rule india-Engr-rule
```

History

This command was first available in ExtremeXOS 12.5.

This command was modified in EXOS 15.2.1 to specify order.

Platform Availability

This command is available on all platforms.

configure identity-management role add policy

```
configure identity-management role role_name add policy policy-name {first | last}
{[before | after] ref_policy_name}
```

Description

Adds a policy for the specified role and specifies the order.

Syntax Description

<i>role_name</i>	Specifies the name of an existing role.
<i>policy-name</i>	Specifies the name of a policy to add to the specified role.

Default

The order of the policy is last if the order is not explicitly specified.

Usage Guidelines

The maximum number of policies or ACL rules that can be applied to a particular role is restricted to 8. This count does not include the policies and rules inherited from a parent role. Since the maximum hierarchy depth is 5, the maximum number of policies and rules supported for a role at the maximum hierarchy depth is 40 (8 x 5).



When a policy is added to a role, it is immediately installed for all identities mapped to that role and all roles below it in the role hierarchy.

Example

The following example configures the role named India-Engr to use the policy named india-Engr-policy:

```
* Switch.44 # configure identity-management role "India-Engr" add policy
india-Engr-policy
```

History

This command was first available in ExtremeXOS 12.5.

This command was modified in EXOS 15.2.1 to specify order.

Platform Availability

This command is available on all platforms.



configure identity-management role-based-vlan

```
configure identity-management role-based-vlan [add | delete] ports [port_list |
all]
```

Description

This command defines

Syntax Description

role-based vlan	Associates the identity to a specific VLAN based on the identity's role.
add	Adds ports to the Identity Management role-based vlan enabled portlist.
delete	Deletes ports from the Identity Management role-based vlan enabled portlist.
ports	Configures Identity Management role-based VLAN on ports.
<i>port_list</i>	Configures Identity Management role-based VLAN on specified port list.
all	Configures Identity Management role-based VLAN on all ports.

Default

N/A.



Usage Guidelines

Use this command to configure the role-based VLAN feature for Identity Management enabled ports. This command requires the ports to be part of a base VLAN. Enabling role-based VLAN on Identity Management enabled ports allows the identity to be placed in the correct VLAN mapped to the role as configured by the administrator.



Note

You cannot enable the Identity Manager role-based VLAN feature on Netlogin enabled ports.

Example

The following example configures Identity Management on ports 1-3, and 5.

```
x460-48t.10 # configure identity-management role-based-vlan add ports
1-3,5
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

configure identity-management role delete child-role

```
configure identity-management role role_name delete child-role [child_role | all]
```

Description

Deletes one or all child roles from the specified role.

Syntax Description

<i>role_name</i>	Specifies the name of an existing role.
child-role	Specifies a name for a child role to delete.
all	Specifies that all child roles are to be deleted.

Default

N/A.

Usage Guidelines

None.



Example

The following example deletes the child role named East from the existing role named India-Engr:

```
* Switch.66 # configure identity-management role "India-Engr" delete child-
role East
```

The following command deletes all child roles from the existing role named India-Engr:

```
* Switch.66 # configure identity-management role "India-Engr" delete child-
role all
```

History

This command was first available in ExtremeXOS 12.5.

The all option was added in ExtremeXOS 12.6.

Platform Availability

This command is available on all platforms.

configure identity-management role delete dynamic-rule

```
configure identity-management role role_name delete dynamic-rule [rule_name |
all]
```

Description

Deletes one or all dynamic ACL rules for the specified role.

Syntax Description

<i>role_name</i>	Specifies the name of an existing role.
<i>rule_name</i>	Specifies the name of a dynamic ACL rule to delete from the specified role.
all	Specifies that all dynamic ACL rules are to be deleted.

Default

N/A.

Usage Guidelines

None.



Example

The following example deletes all dynamic rules from the role named India-Engr:

```
* Switch.55 # configure identity-management role "India-Engr" delete dynamic-rule all
```

History

This command was first available in ExtremeXOS 12.5.

The all option was added in ExtremeXOS 12.6.

Platform Availability

This command is available on all platforms.

configure identity-management role delete policy

```
configure identity-management role role_name delete policy [policy-name | all]
```

Description

Deletes one or all policies for the specified role.

Syntax Description

<i>role_name</i>	Specifies the name of an existing role.
<i>policy-name</i>	Specifies the name of a policy to delete from the specified role.
all	Specifies that all policies are to be deleted from the specified role.

Default

N/A.

Usage Guidelines

None.

Example

The following example deletes the policy named india-Engr-policy from the role named India-Engr:

```
* Switch.44 # configure identity-management role "India-Engr" delete policy india-Engr-policy
```



History

This command was first available in ExtremeXOS 12.5.

The all option was added in ExtremeXOS 12.6.

Platform Availability

This command is available on all platforms.

configure identity-management role match-criteria inheritance

```
configure identity-management role match-criteria inheritance
configure identity-management role match-criteria inheritance [on | off]
```

Description

This command enables or disables the match-criteria inheritance support. Check the current status by performing the "show identity-management" command.

Syntax Description

role	User role.
match-criteria	Match criteria for the role.
inheritance	Inheriting match criteria from parent role to child role.
on off	Specifies whether match criteria inheritance is on or off.

Default

Off.

Usage Guidelines

From EXOS 15.2, child roles can inherit the match criteria of the parent role. This helps the user since the match criteria need not be duplicated in all levels of hierarchy.

When match-criteria inheritance is on, for a user to be classified under a child role, he has to satisfy the match criteria of the child role and also all parent roles in the hierarchy.

Match criteria inheritance helps users in avoiding the need to duplicate match-criteria entries in the hierarchy.



Example

For example, there are roles called Employee, USEmployee and USSales in an organization hierarchy of a company XYZCorp.com. Till EXOS 15.1 (or with match-criteria inheritance off), the user has to create three roles like this:

```
* Switch.1 # create identity-management role Employee match-criteria "company
== XYZCorp.com;"
* Switch.2 # create identity-management role USEmployee match-criteria
"company == XYZCorp.com; AND country == USA;"
* Switch.3 # create identity-management role USSales match-criteria "company
== XYZCorp.com; AND country == USA; AND department = Sales"
* Switch.4 # configure identity-management role "Employee" add child-role
"USEmployee"
* Switch.5 # configure identity-management role "USEmployee" add child-role
"USSales"
```

Now this can be simplified into the following since child role inherits parent role's match criteria:

```
* Switch.1 # configure identity-management role match-criteria inheritance on
* Switch.2 # create identity-management role Employee match-criteria "company
== XYZCorp.com;"
* Switch.3 # create identity-management role USEmployee match-criteria
"country == USA;"
* Switch.4 # create identity-management role USSales match-criteria
"department = Sales"
* Switch.5 # configure identity-management role "Employee" add child-role
"USEmployee"
* Switch.6 # configure identity-management role "USEmployee" add child-role
"USSales"
```

History

This command was first available in ExtremeXOS 15.2

Platform Availability

This command is available on all platforms.

configure identity-management role priority

```
configure identity-management role role_name priority pri_value
```

Description

Configures a priority value for the specified role.



Syntax Description

<i>role_name</i>	Specifies the name of an existing role that you want to configure.
<i>pri_value</i>	Specifies the role priority; the lower the priority number, the higher the priority. The range of values is 1 to 255. Value 1 represents the highest priority, and value 255 represents the lowest priority.

Default

Priority=255.

Usage Guidelines

The role priority determines which role a user is mapped to when the user's attributes match the match-criteria of more than 1 role. If the user's attributes match multiple roles, the highest priority (lowest priority value) role applies. If the priority is the same for all matching roles, the role for which the priority was most recently set or modified is used.

Example

The following example configures the role named India-Engr to use the highest priority:

```
* Switch.33 # configure identity-management role "India-Engr" priority 1
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

configure identity-management stale-entry aging-time

```
configure identity-management stale-entry aging-time seconds
```

Description

Configures the stale-entry aging time for event entries in the identity management database.

Syntax Description

<i>seconds</i>	Specifies the period (in seconds) at which event entries are deleted. The range is 60 to 1800 seconds.
----------------	--



Default

180 seconds.

Usage Guidelines

The identity management database contains active entries, which correspond to active users and devices, and event entries, which record identity management events such as user logout or device disconnect. The active entries are automatically removed when a user logs out or a device disconnects. The event entries are automatically removed after a period defined by the stale-entry aging time.

Note



To capture active and event entries before they are deleted, you can use external management software such as Ridgeline™, which can access the switch using XML APIs. We recommend that the external client(s) that poll the identity management database be configured for polling cycles that are between one-third and two-thirds of the stale-aging time. This ensures that a new database entry or event does not age out before the next polling cycle.

The stale-entry aging time defines when event entries become stale. To preserve memory, the software periodically uses a cleanup process to remove the stale entries. You can configure the stale-entry aging time. The cleanup interval is defined by the software.

When memory usage is high, the software reduces both the stale-entry aging time and the cleanup interval to keep memory available for new entries. The following table shows how the database is managed as memory usage increases.

Table 32: Identity Management Database Usage Levels

Database Memory Usage Level	Database Memory Usage Level (Percent)	Effective Stale-Entry Aging Time	Description
Normal	Up to 80%	Configured stale-entry aging time	New identities and associated information (VLAN and IP addresses) are added to or updated in the database. Events are also added to the database. Events are deleted from the database after the configured stale-entry aging time.
High	Above 80% to 90%	The lower value of the following: 90 seconds or 50% of the configured stale-entry aging time	Identities and events are added to the database as for the normal usage level, but the effective stale-entry aging time is reduced to delete events sooner and free memory.



Table 32: Identity Management Database Usage Levels (continued)

Database Memory Usage Level	Database Memory Usage Level (Percent)	Effective Stale-Entry Aging Time	Description
Critical	Above 90%	15 seconds	The effective stale-entry aging time is further reduced to delete events sooner and free memory. No new identities are added to the database at this usage level, but updates (such as the addition or deletion of a VLAN or IP address) continue. At this level, the database might be missing active entries.
Maximum	Above 98%	15 seconds	At this level, the software does not process additions or updates to the database. The software only processes deletions. At this level, the database might be missing active entries.

Whenever the database usage level changes, an EMS message is logged, and if enabled, an SNMP trap is sent. If the switch changes the stale-entry aging time, the SNMP trap contains the new stale-entry aging time.

Note

If the database level regularly reaches the high usage level, or if it reaches the critical or maximum levels, it is time to investigate the cause of the issue. The solution might be to increase the database memory size.

External clients should be capable of adjusting the polling cycles. Because the aging cycle is shorter when memory is low, it is best if external clients can adjust their polling cycles in response to SNMP traps that announce a change in the stale-entry aging time.

Example

The following command configures the stale-entry aging time for 90 seconds:

```
* Switch.4 # configure identity-management stale-entry aging-time 90
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.



configure identity-management whitelist

```
configure identity-management whitelist add [mac mac_address {macmask} | ip
ip_address {netmask} | ipNetmask] | user user_name]configure identity-management
whitelist delete [all | mac mac_address {macmask} | ip ip_address {netmask} |
ipNetmask] | user user_name]
```

Description

Adds or deletes an identity in the identity manager whitelist.

Syntax Description

add	Adds the specified identity to the whitelist.
delete	Deletes the specified identity from the whitelist.
all	Specifies that all identities are to be deleted from the whitelist. This option is available only when the delete attribute is specified.
<i>mac_address</i>	Specifies an identity by MAC address.
<i>macmask</i>	Specifies a MAC address mask. For example: FF:FF:FF:00:00:00.
<i>ip_address</i>	Specifies an identity by IP address.
<i>netmask</i>	Specifies a mask for the specified IP address.
<i>ipNetmask</i>	Specifies an IP network mask.
<i>user_name</i>	Specifies an identity by user name.

Default

N/A.

Usage Guidelines

The software supports up to 512 entries in the whitelist. When you add an identity to the whitelist, the switch searches the blacklist for the same identity. If the identity is already in the blacklist, the switch displays an error.

It is possible to configure an identity in both lists by specifying different attributes in each list. For example, you can add an identity username to the whitelist and add the MAC address for that user's laptop in the blacklist. Because the blacklist has priority over the whitelist, identity access is denied from the user's laptop, but the user can access the switch from other locations.

If you add a new whitelist entry that is qualified by a MAC or IP address, the identity manager does the following:

- Reviews the identities already known to the switch. If the new whitelist entry is blacklisted (by specifying a different identity attribute), no action is taken.
- If the identity is not blacklisted and is known on the switch, all existing ACLs for the identity are removed.



- When a whitelisted MAC-based identity is detected or already known, an Allow All ACL is programmed for the identity MAC address for the port on which the identity is detected.
- When a whitelisted IP-based identity is detected or already known, an Allow All ACL is programmed for the identity IP address for the port on which the identity is detected.

If you add a new whitelist entry that is qualified by a username (with or without a domain name), the identity manager does the following:

- Reviews the identities already known to the switch. If the new whitelist entry is an identity known on the switch, an Allow All ACL is programmed for the identity MAC address on all ports to which the identity is connected.
- When a new whitelisted username-based identity accesses the switch, an Allow All ACL is programmed for the identity MAC address on the port on which the identity is detected.
- The ACL for a whitelisted username follows any ACLs based on Kerberos snooping.

Allow All ACLs for whitelisted entries exist as long as the identity remains in the identity manager database.

If you delete an identity from the whitelist, identity manager checks to see if the identity is in the local database. If the identity is known to the switch, the switch does the following:

- Removes the Allow All ACL from the port to which the identity connected.
- Initiates the role determination procedure for the switch port to which the known identity connected. This ensures that the appropriate role is applied to the identity that is no longer whitelisted.

**Note**

The role determination process can trigger an LDAP refresh to collect identity attributes for role determination.

Example

The following command adds an IP address to the whitelist:

```
* Switch.4 # configure identity-management whitelist add ip 10.0.0.1
```

The following command deletes a user name from the whitelist:

```
* Switch.5 # configure identity-management whitelist delete user john
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on all platforms.



configure ldap domain

```
configure ldap domain domain_name [default | non-default]
```

Description

This command is used to configure a previously added LDAP domain as default or non-default. If a domain is configured as default, older default domain, if any, will no longer be default since once only one domain can be default at a time.

Syntax Description

<i>domain_name</i>	Name of domain to be configured.
--------------------	----------------------------------

Default

N/A.

Usage Guidelines

Use this command to configure an LDAP domain as default or non-default.

Example

This command marks the LDAP domain sales.XYZCorp.com as the default domain.

```
configure ldap domain sales.XYZCorp.com default
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on all platforms.

configure ldap domain add server


```
configure {identity-management} ldap {domain domain_name} add server [host_ipaddr  
| host_name] {server_port} {client-ip client_ipaddr} {vr vr_name} {encrypted sasl  
digest-md5}
```



Description

This command adds an LDAP server under an LDAP domain and configures the parameters for contacting the server.

Syntax Description

<i>domain_name</i>	Specifies the LDAP domain under which this server should be added.
<i>host_ipaddr</i>	Specifies a IP address for an LDAP server to add.
<i>host_name</i>	Specifies a DNS hostname for an LDAP server to add.
<i>server_port</i>	Specifies a port number for the LDAP service. The default port number is 389.
<i>client_ipaddr</i>	Specifies the LDAP client IP address, which should be set to the IP address of the interface that will connect to the LDAP server.
<i>vr_name</i>	Specifies the VR name for the interface that will connect to the LDAP server. The default VR for LDAP client connections is VR-Mgmt.
encrypted sasl digest-md5	<p>Specifies that the LDAP client uses Digest RSA Data Security, Inc. MD5 Message-Digest Algorithm encryption over SASL (Simple Authentication and Security Layer) to communicate with the LDAP server. Note that this mechanism encrypts only the password credentials, and the LDAP information exchange uses plain text.</p> <hr/> <p>Note</p>  <p>To support Digest RSA Data Security, Inc. MD5 Message-Digest Algorithm over SASL, the LDAP client (bind user) password must be stored using 'reverse encryption,' and the <i>host_name</i> should be configured as the fully-qualified host name for the LDAP server.</p> <hr/>

Default

client-ipaddr is optional. If *client-ipaddr* is not specified, the LDAP client looks up the interface through which the LDAP server can be reached.

If *vr_name* is not specified, the LDAP client assumes it to be VR-Mgmt.

If "encrypted sasl digest-md5" is not specified, the LDAP client talks to the LDAP server using plain text.

Usage Guidelines

You can configure up to 8 LDAP servers under one LDAP domain. The LDAP servers are contacted in the order of configuration. If the first server does not respond before the timeout period expires, the second server is contacted. This process continues until an LDAP server responds, and then the



responding server marked as 'active'. Subsequent LDAP requests for that LDAP domain are sent to the 'active' server.



Note

If the switch cannot resolve the host name using a DNS server, the switch rejects the command and generates an error message.

As of 15.2, the "identity-management" keyword is now optional in this command.

Example

The following command configures LDAP client access to LDAP server LDAP1 using encrypted authentication:

```
* Switch.6 # configure identity-management ldap add server LDAP1 client-ip
10.10.2.1
encrypted sasl digest-md5
```

The following command adds the LDAP server LDAPServer1.sales.XYZCorp.com under the domain sales.XYZCorp.com and configures the LDAP client to contact it over VR-Default. It also configures the LDAP client to communicate with the server using digest-md5 encryption over SASL.

```
configure ldap domain sales.XYZCorp.com add server
LDAPServer1.sales.XYZCorp.com vr VR-Default encrypted sasl digest-md5
```

The following command adds the LDAP server 192.168.1.1 under the domain sales.XYZCorp.com and also configures the LDAP client to contact it through the interface 10.10.10.1 over VR-Mgmt.

```
configure ldap domain sales.XYZCorp.com add server 192.168.1.1 client-ip
10.10.10.1
```

History

This command was first available in ExtremeXOS 12.5.

This command was modified in ExtremeXOS 15.2 to make the identity management keyword optional.

Platform Availability

This command is available on all platforms.

configure ldap domain base-dn

```
configure {identity-management} ldap {domain [domain_name | all]} base-dn [base_dn
| none | default]
```



Description

Configures the LDAP base-dn to be used while searching an user under an LDAP domain.

Syntax Description

<i>domain_name</i>	Specifies the LDAP domain for which this base-dn is to be configured.
<i>base_dn</i>	Specifies the LDAP base domain under which the users are to be searched.
none	Specifies the LDAP root domain as the location under which the users are to be searched.
default	Restores the base_dn to it default value i.e., same as the domain name.

Default

By default base-dn is assumed to be the same as the domain name unless configured otherwise.

If a domain is not specified, the base-dn is configured for the default domain.

Usage Guidelines

LDAP base-dn is the LDAP directory root under which the users are to be searched. By default base-dn is assumed to be the same as the domain name.

For users upgrading from EXOS 15.1 and older versions, a domain is created with the same name as the base-dn in the older configuration. This domain is marked as the default domain. This can be changed later if required.

Example

The following commands configure the base-dn for the domain sales.XYZCorp.com.

The base-dn configured as XYZCorp.com means that XYZCorp.com is the base location to search for user information.

```
* Switch.11 # configure ldap domain sales.XYZCorp.com base-dn XYZCorp.com
```

The base-dn configured as none means that the directory root is the base location to search for user information.

```
* Switch.12 # configure ldap domain sales.XYZCorp.com base-dn none
```

History

This command was first available in ExtremeXOS 12.5.

This command was modified in ExtremeXOS 15.2 to add the **{domain [domain_name | all]}** option.



Platform Availability

This command is available on all platforms.


configure ldap domain bind-user

```
configure {identity-management} ldap {domain [domain_name|all] } bind-user
[user_name {encrypted} password | anonymous ]
```

Description

Configures the LDAP client credentials required for the switch to access an LDAP server.

Syntax Description

<i>domain_name</i>	Specifies the LDAP domain for which this bind-user is to be configured.
<i>user_name</i>	Specifies the user name for LDAP server access.
encrypted	Indicates that the specified password is encrypted.
<i>password</i>	Specifies the user password for LDAP server access.
	<div style="border: 1px solid black; padding: 5px;"> <p>Note</p>  <p>To support Digest RSA Data Security, Inc. MD5 Message-Digest Algorithm over SASL, the password must be stored using 'reverse encryption.'</p> </div>
anonymous	Specifies user anonymous for LDAP server access.

Default

If no domain is specified, the bind-user is configured for the default domain.

Usage Guidelines

The bind-user is an LDAP user who has read access to user information in the LDAP directory.

On many newer directory servers "anonymous" access is disabled. You may also find that though the LDAP bind succeeds, the anonymous user might be denied read access to user information.

Example

The following command configures the LDAP bind user as jsmith with password Extreme for the domain sales.XYZCorp.com:

```
* Switch.14 # configure ldap domain sales.XYZCorp.com bind-user jsmith
password Extreme
```



History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

configure ldap domain delete server

```
configure { identity-management } ldap {domain [domain_name|all]} delete server
[host_ipaddr | host_name] {server_port} {vr vr_name}
```

Description

This command is used to delete one or all LDAP servers from one or all LDAP domains.

Syntax Description

<i>domain_name</i>	Specifies the LDAP domain from which this server is to be deleted.
all	Specifies that all configured LDAP servers are to be deleted.
<i>host_ipaddr</i>	Specifies the IP address of the LDAP server to delete.
<i>host_name</i>	Specifies a DNS hostname of the LDAP server to delete.
<i>server_port</i>	Specifies a port number for the LDAP service to delete. The default port number is 389.

Default

If a domain is not specified, the server(s) under default domain is deleted.

Usage Guidelines

None.

Example

The following command deletes the LDAP server LDAPServer1.sales.XYZCorp.com from the domain sales.XYZCorp.com:

```
* Switch.8 # configure ldap domain sales.XYZCorp.com delete server
LDAPServer1.sales.XYZCorp.com
```

The following command deletes all LDAP servers from all LDAP domains:

```
* Switch.8 # configure ldap domain all delete server all
```



History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

configure ldap domain netlogin

```
configure {identity-management} ldap { domain [ domain_name | all ] } [enable | disable] netlogin [dot1x | mac | web-based]
```

Description

Enables or disables LDAP queries for the specified type of network login users.

Syntax Description

<i>domain_name</i>	Specifies the LDAP domain for which this configuration is to be applied.
dot1x	Enables or disables LDAP queries for dot1x network login.
mac	Enables or disables LDAP queries for MAC network login.
web-based	Enables or disables LDAP queries for web-based network login.

Default

LDAP queries are enabled for all types of network login.

Usage Guidelines

It may be necessary to disable LDAP queries for specific type of netlogin user, for example, netlogin mac users, whose username is the same as mac address. The LDAP directory might not contain useful information about these type of users and unnecessary LDAP queries can be avoided.



Note

LDAP queries are not sent for locally authenticated network login users.

Example

The following command enables LDAP queries for MAC network login:

```
* Switch.99 # configure identity-management ldap enable netlogin mac
```



The following command disables LDAP queries for dot1x network login:

```
* Switch.99 # configure identity-management ldap disable netlogin dot1x
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

create identity-management role

```
create identity-management role role_name match-criteria match_criteria {priority pri_value}
```

Description

Creates and configures an identity management role.

Syntax Description

<i>role_name</i>	Specifies a name for the new role (up to 32 characters).
<i>match_criteria</i>	Specifies an expression that identifies the users to be assigned to the new role.
<i>pri_value</i>	Specifies the role priority; the lower the priority number, the higher the priority. The range of values is 1 to 255. Value 1 represents the highest priority, and value 255 represents the lowest priority.

Default

Priority=255.

Usage Guidelines

The identity management feature supports a maximum of 64 roles.

The role name can include up to 32 characters. Role names must begin with an alphabetical letter, and only alphanumeric, underscore (_), and hyphen (-) characters are allowed in the remainder of the name. Role names cannot match reserved keywords, or the default role names reserved by identity manager. For more information on role name requirements and a list of reserved keywords, see [Object Names](#) of the *ExtremeXOS Concepts Guide*. The role names reserved by identity manager are:

- authenticated.
- blacklist.



- unauthenticated.
- whitelist.

The match-criteria is an expression or group of expressions consisting of identity attributes, operators and attribute values. The maximum number of attribute value pairs in a role match criteria is 16. The variables in the match criteria can be matched to attributes retrieved for the identity from an LDAP server, or they can be matched to attributes learned locally by identity manager.

Table 33: LDAP Match Criteria Attributes on page 1615 lists match criteria attributes that can be retrieved from an LDAP server.

Table 34: Locally Learned Match Criteria Attributes on page 1615 lists locally learned attributes that can be used for match criteria.

Table 35: Match Criteria Operators on page 1616 lists the match criteria operators.

Table 33: LDAP Match Criteria Attributes

LDAP Attribute Name	Value Type
l or location	String
company	String
co or country	String
department	String
employeeID	String
st or state	String
title	String
mail or email	String
memberOf	String

Table 34: Locally Learned Match Criteria Attributes

Attribute Description	Attribute Name	Value Type	Example
LLDP device name	device-model	String	<code>device-name == Avaya4300</code>
LLDP device capabilities	device-capability	String: OtherRepeaterBridgeWLAN access portRouterPhoneDOCSIS cable deviceStation only	<code>device-capability == Telephone</code>
LLDP device manufacturer name	device-manufacturer-name	String	<code>device-manufacturer-name == Avaya</code>



Table 34: Locally Learned Match Criteria Attributes (continued)

Attribute Description	Attribute Name	Value Type	Example
LLPD system description	device-description	String	<code>device-description==Dell EqualLogic Storage Array</code>
MAC address	mac	MAC	<code>mac == 00:01:e6:00:00:00/ ff:ff:ff:00:00:00</code>
MAC OUI	mac-oui	MAC	<code>mac-oui == 00:04:96</code>
IP address	ip-address	IP	<code>ip-address == 10.1.1.0/20</code>
User name	username	String	<code>userName == adam</code>
Port list	ports	Portlist	<code>ports == 1,5-8</code>

Table 35: Match Criteria Operators

Operator	Description
<code>==</code>	Equal. Creates a match when the value returned for the specified attribute matches the value specified in the role.
<code>!=</code>	Not equal. Creates a match when the value returned for the specified attribute does not match the value specified in the role.
<code>AND</code>	And. Creates a match when the two expressions joined by this operator are both true.
<code>contains</code>	Contains. Creates a match when the specified attribute contains the text specified in the role definition.
<code>;</code>	Semicolon. This delimiter separates expressions within the match criteria.

The role priority determines which role a user is mapped to when the user's attributes match the match-criteria of more than 1 role. If the user's attributes match multiple roles, the highest priority (lowest numerical value) role applies. If the priority is the same for all matching roles, the role for which the priority was most recently set or modified is used.

Example

The following examples create roles for the conditions described in the comments that precede the commands:

```
# Creates a role named "India-Engr" that matches employees from the
```



```

Engineering
# department who work in India
* Switch.22 # create identity-management role "India-Engr" match-criteria
"country==India; AND department==Engineering;"
# Creates a role named "US-Engr" that matches employees whose title is
Engineer and
# who work in United States
* Switch.23 # create identity-management role US-Engr match-criteria "title
contains Engineer; AND country == US;" priority 100
# Creates a role named "Avaya4300Device" for Avaya phones of type 4300 that
are
# manufactured by Avaya
* Switch.24 # create identity-management role "Avaya4300Device" match-
criteria "device-capability == Phone; AND device-name == Avaya4300; AND
device-manufacturer-name == Avaya;"
# Creates a role for all Extreme Networks switches with MAC-OUI "00:04:96"
* Switch.25 # create identity-management role "ExtremeSwitch" match-criteria
"mac-oui == 00:04:96;"
# Creates a role for all identities with IP address 1.2.3.1 - 1.2.3.255
* Switch.26 # create identity-management role "EngineeringDomain" match-
criteria "ip-Address == 1.2.3.0/255.255.255.0;"
# Creates a role for all phone devices with MAC_OUI of "00:01:e6"
* Switch.27 # create identity-management role "Printer" match-criteria "mac
== 00:01:e6:00:00:00/ff:ff:ff:00:00:00; device-capability == Phone;"
# Creates a role for the user name "adam" when he logs in from IP address
1.2.3.1 -
# 1.2.3.255.
* Switch.28 # create identity-management role "NotAccessibleUser" match-
criteria "userName == adam; AND "ip-Address == 1.2.3.0/24;"
# Creates a role named "secureAccess" for users who log in on ports 1, 5, 6,
7, and 8
# with IP addresses in the range of 10.1.1.1 to 10.1.1.255
create identity-management role "SecureAccess" match-criteria "ports ==
1,5-8; AND ip-address == 10.1.1.0/20;"
# Creates a role named "Prod-Engineers" for all the engineers who are under
LDAP group 'Production'.
Create identity-management role "Prod-Engineers" match-criteria
"title==Engineer; AND memberOf==Production;"

```

History

This command was first available in ExtremeXOS 12.5.

Support for matching locally learned attributes was added in ExtremeXOS 12.7.

Platform Availability

This command is available on all platforms.

create ldap domain

```
create ldap domain domain_name {default}
```



Description

This command is used to add an LDAP domain. The new domain can be added as the default. Older default domains, if any, will no longer be the default since once only one domain can be default at a time.

Syntax Description

<i>domain_name</i>	Name of new LDAP domain to be added
--------------------	-------------------------------------

Default

N/A.

Usage Guidelines

Use this command to add an LDAP domain.

You can see the LDAP domains added by using the show ldap domain command.

Supporting multiple domains gives EXOS the capability to send LDAP queries to gather information about users belonging to different domains but connected to the same switch.

You can add upto 8 LDAP domains.

Example

The following command creates an LDAP domain with the name "sales.XYZCorp.com and marks it as the default domain:

```
create ldap domain sales.XYZCorp.com default
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on all platforms.

configure ldap hierarchical-search-oid

```
configure ldap {domain [domain_name | all]} hierarchical-search-oid [ldap-matching-rule-in-chain | oid | none]
```



Description

Configures an OID to perform a hierarchical search if the LDAP server requires it.

Syntax Description

<i>domain_name</i>	Domain name on which to configure ldap.
all	All domains.
ldap-matching-rule-in-chain	Configures the OID 1.2.840.113556.1.4.1941.
<i>oid</i>	Object identifier.
none	Specifies that LDAP query should not include any OID for hierarchical search.

Default

N/A.

Usage Guidelines

Use this command to configure an OID to perform a hierarchical search if the LDAP requires it. The OID supplied with this command will be used to form the LDAP query. If a server does not require extended control OID, the none option can be selected.

Example

```
configure ldap domain abc.com hierarchical-search-oid
ldap_matching_rule_in_chain
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

delete identity-management role

```
delete identity-management role {role_name | all}
```

Description

Deletes one or all roles.



Syntax Description

<i>role_name</i>	Specifies a name of an existing role to delete.
all	Specifies that all roles are to be deleted.

Default

N/A.

Usage Guidelines

Any policy applied to users of a deleted role gets reverted. The users are placed under one of the other roles based on their attributes. Parent and child relationships to other roles are also deleted. For example, all child roles under the deleted role become orphans and hence they and their descendants no longer inherit the policies of the deleted role.

Example

The following example deletes the role named India-Engr:

```
* Switch.99 # delete identity-management role "India-Engr"
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

delete ldap domain

```
delete ldap domain [domain_name | all]
```

Description

This command is used to delete one or all LDAP domains.

When an LDAP domain is deleted, all LDAP servers added under that domain are also deleted. Also all LDAP configurations done for that domain are deleted.

Syntax Description

<i>domain_name</i>	Name of the LDAP domain that will be deleted.
--------------------	---



Default

N/A.

Usage Guidelines

Use this command to delete one or all LDAP domains.

When an LDAP domain is deleted, all LDAP servers added under that domain are also deleted. All LDAP configurations for that domain are also deleted.

Example

This command deletes the LDAP domain sales.XYZCorp.com

```
delete ldap domain sales.XYZCorp.com
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on all platforms.

disable identity-management

disable identity-management

Description

Disables the identity management feature, which tracks users and devices that connect to the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.



Usage Guidelines

Only admin-level users can execute this command.

Note



If the identity management feature is running and then disabled, all identity management database entries are removed and cannot be retrieved. If identity management is enabled later, the identity management feature starts collecting information about currently connected users and devices.

Example

The following command disables the identity management feature:

```
disable identity-management
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

disable snmp traps identity-management

```
disable snmp traps identity-management
```

Description

Disables the identity management feature to send SNMP traps for low memory conditions.

Syntax Description

This command has no arguments or variables.

Default

No traps are sent.

Usage Guidelines

None.



Example

The following command disables the identity management SNMP trap feature:

```
disable snmp traps identity-management
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

enable identity-management

enable identity-management

Description

Enables the identity management feature, which tracks users and devices that connect to the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Only admin-level users can execute this command.

After identity management is enabled, the software creates two dynamic ACL rules named `idm_black_list` and `idm_white_list`. These rules are removed if identity management is disabled.



Note

FDB entries are flushed on identity management enabled ports when this command is executed.



Example

The following command enables the identity management feature:

```
enable identity-management
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

enable snmp traps identity-management

```
enable snmp traps identity-management
```

Description

Enables the identity management feature to send SNMP traps for low memory conditions.

Syntax Description

This command has no arguments or variables.

Default

No traps are sent.

Usage Guidelines

The low memory conditions are described in the description for the `configure identity-management stale-entry aging-time seconds` command.

Example

The following command enables the identity management SNMP trap feature:

```
enable snmp traps identity-management
```

History

This command was first available in ExtremeXOS 12.4.



Platform Availability

This command is available on all platforms.

refresh identity-management role

```
refresh identity-management role user [user_name {domain domain_name} | all {role role_name}]
```

Description

Refreshes the role evaluation for the specified user, for all users, or for all users currently under the specified role.

Syntax Description

<i>user_name</i>	Specifies a user name for which role evaluation will be refreshed.
<i>domain_name</i>	Specifies a domain name for the specified user.
all	Specifies a refresh for all users associated with the specified role.
<i>role_name</i>	Specifies a role name for which all users will be refreshed.

Default

N/A.

Usage Guidelines

It may be necessary to refresh the role of a user due to a new role which might be better suited for the user or due to a change in LDAP attributes of the user which in turn might result in the user being classified under a different role. This command can be used in all such cases.

Example

The following example refreshes the role for user Tony:

```
* Switch.22 # refresh identity-management role user "Tony"
```

The following example refreshes the role for all users who are currently classified under the Marketing role:

```
* Switch.22 # refresh identity-management role all "Marketing"
```

History

This command was first available in ExtremeXOS 12.7.



Platform Availability

This command is available on all platforms.

show identity-management

show identity-management

Description

Displays the identity management feature configuration.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the identity management feature configuration:

```
# show identity-management
Identity Management           : Enabled
Stale entry age out (effective) : 180 Seconds (180 Seconds)
Max memory size              : 512 Kbytes
Enabled ports                 : 1-26
FDB Detection Disabled ports  : None
IPARP Detection Disabled ports : None
IPSecurity Detection Disabled ports : None
Kerberos Detection Disabled ports : None
LLDP Detection Disabled ports  : None
Netlogin Detection Disabled ports : None
SNMP trap notification       : Enabled
Match Criteria Inheritance    : On
Access list source address type : MAC
Kerberos aging time (DD:HH:MM) : None
Kerberos force aging time (DD:HH:MM) : None
Valid Kerberos servers       : none configured(all valid)
```

History

This command was first available in ExtremeXOS 12.4.



Kerberos Force Aging Time information was added in ExtremeXOS 12.6.

Platform Availability

This command is available on all platforms.

show identity-management blacklist

show identity-management blacklist

Description

Displays the identities in the identity manager blacklist.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the identities in the blacklist:

```
* Switch.93 # show identity-management blacklist
-----
Type          BlackList Entry
-----
MAC           01:02:03:04:05:06/ff:ff:ff:00:00:00
IP            1.2.3.4/255.255.255.0
User          john@mydomain.com
-----
> indicates entry value truncated past 35 characters
Number of BlackList Entries      : 3
```

History

This command was first available in ExtremeXOS 12.7.



Platform Availability

This command is available on all platforms.

show identity-management entries

```
show identity-management entries {user id_name} {domain domain} {ports port_list}
{mac mac_address} {vlan vlan_name} {ipaddress ip_address} {detail}
```

Description

Displays the entries in the identity management database.

Syntax Description

<i>id_name</i>	Limits the display to entries that contain the specified user ID.
<i>domain</i>	Limits the display to entries for the specified domain.
<i>port_list</i>	Limits the display to entries for the specified ports.
<i>mac_address</i>	Limits the display to entries that contain the specified MAC address.
<i>vlan_name</i>	Limits the display to entries that contain the specified VLAN name.
<i>ip_address</i>	Limits the display to entries that contain the specified IP address.
detail	Expands the display to include more information about identity management entries.

Default

N/A.

Usage Guidelines

Only admin-level users can execute this command.

The displayed ID Name is the actual user name when Network Login or Kerberos Snooping is enabled. For unknown users, the software creates a user name using the format: User_XXXXXXXXXXXXXXXX. The number in the user name is a 16-bit hash number that is generated using the user's port, MAC address, and IP address numbers.

The displayed Domain Name is displayed only if the client is discovered through Kerberos snooping or Dot1x and the domain name is supplied in the form of *domain\user*). The NetBIOS hostname is only displayed if this information was present in the Kerberos packets.

When the role is shown as multiple, the identity is connected through multiple ports/locations and different roles apply to each device.



Example

The following command displays all entries in the identity management database:

```
* Switch.4 # show identity-management entries
ID Name/          Flags  Port      MAC/          VLAN          Role
Domain Name      IP
-----
--
Unknown_00:00:00:> ----  1:3      00:00:00:00:00:22  v1(1)
unauthenticat>
-- NA --
00005A4B0000      -m--  1:4      00:00:5a:4b:d1:98  test126(1)    Phone
126.0.0.2(1)
00005A4B0000      -m--  1:4      00:00:5a:4b:d1:9c  test128(1)    Phone
128.0.0.2(1)
00005A4B0000      -m--  1:4      00:00:5a:4b:d1:9e  test129(1)    Phone
129.0.0.2(1)
.
.
000105000000      -m--  1:4      00:01:05:00:03:18  test150(1)    Phone
-- NA --
OTHER(00:04:96:1e> 1---  4:11     00:04:96:1e:32:80  -- NA --
unauthenticat>
-- NA --
joe                --k-  1        00:00:22:33:55:66  v1(1)
authenticated extreme                2.1.3.4(1)
bill              --k-  2        00:00:22:33:44:55  v1(2)          multiple
corp.extremenetworks.com  1.2.3.4(1)
Unknown_00:00:00:> ----  1        00:00:00:00:22:33  v1(1)
unauthenticat>
-- NA --
.
.
OTHER(02:04:96:51> 1---  4:3      02:04:96:51:77:c7  -- NA --
unauthenticat>
-- NA --
-----
--
Flags:                k - Kerberos Snooping, l - LLDP Device,
m - NetLogin MAC-Based, w - NetLogin Web-Based,
x - NetLogin 802.1x
Legend: >            - VLAN / ID Name / Domain / Role Name truncated to column width
(#)                - Total # of associated VLANs/IPs
-- NA --- No IP or VLAN associated
Total number of entries: 60
```

The following command shows the detail format:

```
* Switch.4 # show identity-management entries detail
- ID: "00005A4B0000", 1 Port binding(s)
Role: "Phone"
Port: 1:4, 24 MAC binding(s)
MAC: 00:00:5a:4b:d1:98, Flags: -m--, Discovered: Fri Sep 24 18:30:17 2010
```



```

1 VLAN binding(s)
VLAN: "test126", 1 IP binding(s)
IPv4: 126.0.0.2
Security Profile: ----, Security Violations: ----;
MAC: 00:00:5a:4b:d1:9c, Flags: -m--, Discovered: Fri Sep 24 18:30:17 2010
1 VLAN binding(s)
VLAN: "test128", 1 IP binding(s)
IPv4: 128.0.0.2
Security Profile: ----, Security Violations: ----;
MAC: 00:00:5a:4b:d1:9e, Flags: -m--, Discovered: Fri Sep 24 18:30:17 2010
1 VLAN binding(s)
VLAN: "test129", 1 IP binding(s)
IPv4: 129.0.0.2
Security Profile: ----, Security Violations: ----;
.
.
.
MAC: 00:00:5a:4b:d1:c8, Flags: -m--, Discovered: Fri Sep 24 18:30:17 2010
1 VLAN binding(s)
VLAN: "test150", 1 IP binding(s)
IPv4: 150.0.0.2
Security Profile: ----, Security Violations: ----;
- ID: "000071710000", 1 Port binding(s)
Role: "Phone"
Port: 1:5, 1 MAC binding(s)
MAC: 00:00:71:71:00:01, Flags: -m--, Discovered: Fri Sep 24 19:42:29 2010
1 VLAN binding(s)
VLAN: "palani", 0 IP binding(s)
- ID: "000105000000", 1 Port binding(s)
Role: "Phone"
Port: 1:4, 25 MAC binding(s)
MAC: 00:01:05:00:03:00, Flags: -m--, Discovered: Fri Sep 24 18:30:17 2010
1 VLAN binding(s)
VLAN: "test126", 0 IP binding(s)
MAC: 00:01:05:00:03:01, Flags: -m--, Discovered: Fri Sep 24 18:30:17 2010
1 VLAN binding(s)
VLAN: "test127", 0 IP binding(s)
MAC: 00:01:05:00:03:02, Flags: -m--, Discovered: Fri Sep 24 18:30:17 2010
1 VLAN binding(s)
VLAN: "test128", 0 IP binding(s)
.
.
.
MAC: 00:01:05:00:03:18, Flags: -m--, Discovered: Fri Sep 24 18:30:18 2010
1 VLAN binding(s)
VLAN: "test150", 0 IP binding(s)
- ID: "OTHER(00:04:96:1e:32:80)", 8 Port binding(s)
Role: "unauthenticated"
Port: 4:11, 1 MAC binding(s)
MAC: 00:04:96:1e:32:80, Flags: l---, Discovered: Fri Sep 24 18:30:17 2010
0 VLAN binding(s)
Port: 4:12, 1 MAC binding(s)
MAC: 00:04:96:1e:32:80, Flags: l---, Discovered: Fri Sep 24 18:30:17 2010
0 VLAN binding(s)
Port: 4:13, 1 MAC binding(s)
MAC: 00:04:96:1e:32:80, Flags: l---, Discovered: Fri Sep 24 18:30:17 2010
0 VLAN binding(s)
.

```



```

.
.
Port: 4:18, 1 MAC binding(s)
MAC: 00:04:96:1e:32:80, Flags: l---, Discovered: Fri Sep 24 18:30:17 2010
0 VLAN binding(s)
- ID: "OTHER(02:04:96:51:77:c7)", 2 Port binding(s)
Role: "unauthenticated"
Port: 1:1, 1 MAC binding(s)
MAC: 02:04:96:51:77:c7, Flags: l---, Discovered: Fri Sep 24 18:30:17 2010
0 VLAN binding(s)
Port: 4:3, 1 MAC binding(s)
MAC: 02:04:96:51:77:c7, Flags: l---, Discovered: Fri Sep 24 18:30:17 2010
0 VLAN binding(s)

```

```

-----
--
Flags:          k - Kerberos Snooping, l - LLDP Device,
m - NetLogin MAC-Based, w - NetLogin Web-Based,
x - NetLogin 802.1x
Security Profile:  a - ARP Validation, d - DoS Protection,
g - Gratuitous ARP Protection, r - DHCP Snooping
Security Violations: A - ARP Validation Violation, D - DoS Violation
G - Gratuitous ARP Violation, R - Rogue DHCP Server Detected

```

The following command example shows how domain names, NetBIOS hostnames, and multiple roles appear when in use:

```

Switch4 # show identity-management entries detail
- ID: "john", 1 Port binding(s)
Role: "IT-Engineer"
Domain: "XYZCorp.com", NetBios hostname: "JOHN-DESKTOP"
Port: 17 (Bld-1-Port-1), 1 MAC binding(s)
MAC: 00:00:5a:4b:d1:98, Flags: --k-, Discovered: Tue Nov 16 12:22:46 2010
Force Aging TTL: 00:00:02    Inactive Aging TTL: 00:00:03
1 VLAN binding(s)
VLAN: "corp", 1 IP binding(s)
IPv4: 126.0.0.2
Security Profile: -d--, Security Violations: ----;
- ID: "ramesh", 2 Port binding(s)
Role: "multiple"
Domain: "corp.extremenetworks.com"
Port: 1, 1 MAC binding(s)
MAC: 00:00:00:00:00:13, Flags: --k-, Discovered: Sat Apr 2 02:23:41 2011
Force Aging TTL: 00:00:02    Inactive Aging TTL: N/A
1 VLAN binding(s)
VLAN: "v1", 1 IP binding(s)
IPv4: 10.120.89.9
Role: "Engineer"
Security Profile: adgsr---, Security Violations: A-----,
Port: 2, 1 MAC binding(s)
MAC: 00:00:00:00:00:30, Flags: --k-, Discovered: Sat Apr 2 02:24:30 2011
Force Aging TTL: 00:00:02    Inactive Aging TTL: N/A
1 VLAN binding(s)
VLAN: "v2", 1 IP binding(s)
IPv4: 10.2.3.45
Role: "iphoneEngineer"
Security Profile: ----, Security Violations: ----;
-----

```



```
--
Flags:          k - Kerberos Snooping, l - LLDP Device,
m - NetLogin MAC-Based, w - NetLogin Web-Based,
x - NetLogin 802.1x
Security Profile:  a - ARP Validation, d - DoS Protection,
g - Gratuitous ARP Protection, r - DHCP Snooping
Security Violations: A - ARP Validation Violation, D - DoS Violation
G - Gratuitous ARP Violation, R - Rogue DHCP Server Detected
```

The following command example shows that you can specify multiple options, such as the user name and ports:

```
show identity-management entries user eelco ports 2:2
```

History

This command was first available in ExtremeXOS 12.4.

Kerberos Force Aging TTL and Inactive Aging TTL information was added in ExtremeXOS 12.6.

Support for multiple roles for a single identity was added in ExtremeXOS 12.7.

Platform Availability

This command is available on all platforms.

show identity-management greylist

```
show identity-management greylist
```

Description

Displays the identities in the identity manager greylist.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to display the identities in the identity manager greylist.



Example

```
* Switch.94 # show identity-management greylist
-----
Type          GreyList Entry
-----
User          june@mydomain.com
User          Richard@corp.acme.com
-----
> indicates entry value truncated past 35 characters
Number of GreyList Entries      : 2
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms.

show identity-management list-precedence

show identity-management list-precedence

Description

This command displays the order of list-precedence. The default list-precedence is: greylist blacklist whitelist.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to display the order of list-precedence.

Example

```
* Switch.97 # show identity-management list-precedence
List Precedence:
1.    Greylist
```



2. Blacklist
3. Whitelist

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms.

show identity-management role

```
show identity-management role {role_name} {detail}
```

Description

Displays summary or detailed configuration information for one or all roles.

Syntax Description

<i>role_name</i>	Specifies a name of an existing role to display.
all	Specifies that all roles are to be displayed.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays all roles that are configured on the switch:

```
* Switch.95 # show identity-management role
-----
Role Name           Priority   Child Roles   # Identities
*authenticated      255                0
*unauthenticated    255                0
extr-empl           255      extr-engr     2
extr-engr           255                0
*whitelist          0                0
*blacklist          0                3
-----
Flags : * - Default Roles
```




```
-----
Total number of role(s) configured : 6
```

The following command displays detailed information for all roles that are configured on the switch:

```
* Switch.96 # show identity-management role detail
Role name : extr-empl
Child Roles : engr
Match Criteria : "company==Extreme;"
Policies : extrPol
Identities : john_smith@d.com; MAC: 00:16:23:51:77:99; Port:8
bob_craig@e.com; MAC: 00:18:23:51:77:99; Port:9
Role name : engr
Child Roles : india-engr
Match Criteria : "department==Engineering;"
Policies : engrPol, extrPol
Identities : joe_hardy@b.com; MAC: 00:12:23:51:77:99; Port:10
Role name : india-engr
Child Roles : -
Match Criteria : "country=India; AND department=Engineering;"
Policies : indEngrPol, engrPol, extrPol
Identities : bill_jacob@b.com; MAC: 00:12:33:51:77:99; Port:11
Role name : marketing
Child Roles : -
Match Criteria : "department=Marketing;"
Policies : markrPol, extrPol
Identities : will_smith@a.com; MAC: 00:11:33:51:77:99; Port:14
Role Name: whitelist (Default Role)
Child Roles : ---
Priority : 0
Match Criteria : "Not Applicable"
Policies : --
Identities # : 0
Identities : --
Role Name: blacklist(Default Role)
Child Roles : ---
Priority : 0
Match Criteria : "Not Applicable"
Policies : --
Identities # : 3
Identities : Unknown_00:11:22:33:44:55; MAC: 00:11:22:33:44:55; Port:1
johndoe@extremenetworks.com; MAC: 00:01:02:03:04:05; Port:2
janedoe@extremenetworks.com; MAC: 00:02:04:06:08:10; Port:3
```

The next two examples display detailed information for a single role:

```
* Switch.97 # show identity-management role extr-empl detail
Role name : extr-empl
Child Roles : engr
Match Criteria : "company=Extreme;"
Policies : extrPol
Identities : johnsmith@extreme.com; MAC: 00:11:33:55:77:99; Port:4
bobcraig@extreme.com; MAC: 00:01:03:05:07:09; Port:5
* Switch.98 # show identity-management role NotAccessibleUser detail
Role name : NotAccessibleUser
```



```

Child Roles : engr
Match Criteria : "UserName = adam; AND IP-Address == 1.2.3.0/24; AND port ==
1;"
Policies : extrPol
Identities : adam; MAC: 00:00:11:22:33:44; Port: 1

```

History

This command was first available in ExtremeXOS 12.5.

MAC addresses were added to the displays for the detail option in ExtremeXOS 12.7.

Platform Availability

This command is available on all platforms.

show identity-management statistics

```
show identity-management statistics
```

Description

Displays operation statistics for the identity management feature.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

A user can login from multiple machines or ports, so the total number of login instances can be more than the number of unique users logged in.

Example

The following command displays the identity management feature statistics:

```

Switch.4 # show identity-management statistics
Total number of users logged in      :    2
Total number of login instances      :    2
Total memory used                    :    1 Kbytes
Total memory used by events          :    0 Kbytes
Total memory available               :   511 Kbytes
High memory usage level reached count :    0

```



```

Critical memory usage level reached count:    0
Max memory usage level reached count      :    0
Current memory usage level                 : Normal
Normal memory usage level trap sent        :    0
High memory usage level trap sent         :    0
Critical memory usage level trap sent     :    0
Max memory usage level trap sent          :    0
Event notification sent                    :    0
Total number of roles configured          :    3

```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

show identity-management whitelist

```
show identity-management whitelist
```

Description

Displays the identities in the identity manager whitelist.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the identities in the whitelist:

```

* Switch.94 # show identity-management whitelist
-----
Type          WhiteList Entry
-----
MAC           04:32:13:44:25:06/ff:ff:ff:00:00:00
IP            11.12.13.14/255.255.255.0

```



```
User      jane@mydomain.com
-----
> indicates entry value truncated past 35 characters
Number of WhiteList Entries      : 3
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on all platforms.

show ldap domain

```
show ldap domain {domain_name | all}
```

Description

This command displays the LDAP servers and other LDAP configuration details of one or all LDAP domains.

Syntax Description

<i>domain_name</i>	Displays the details of the specified domain.
all	Displays the details for all domains.

Default

N/A.

Usage Guidelines

Use this command to display the LDAP servers and other LDAP configuration details of one or all LDAP domains. The summary version (`show ldap domain`) displays the list of LDAP domains configured.

Example

```
# show ldap domain
-----
LDAP Domains
-----
XYZCorp.com (Default)
engg.XYZCorp.com
mktg.XYZCorp.com
```



```
sales.XYZCorp.com
-----
```

If no default domain is configured, this note appears at the bottom:

```
Note: No default domain configured
# show ldap domain all
-----
Domain(default) : XYZCorp.com
-----
Base-DN          : XYZCorp.com
Bind credential  : jsmith
LDAP Hierarchical Search OID : LDAP_MATCHING_RULE_IN_CHAIN
(1.2.840.113556.1.4.1941)
LDAP Configuration for Netlogin:
dot1x           : Enabled
mac             : Enabled
web-based       : Enabled
LDAP Server 1   : 192.168.2.101
Server Port     : 389
Client IP       : Any
Client VR       : VR-Mgmt
Security Mechanism : Plain Text
Status          : Active
LDAP Server 2   : 192.168.2.102
Server Port     : 389
Client IP       : Any
Client VR       : VR-Mgmt
Security Mechanism : Plain Text
Status          : Not Active
-----
Domain          : engg.XYZCorp.com
-----
Base-DN          : engg.XYZCorp.com
Bind credential  : pkumar
LDAP Hierarchical Search OID : 1.2.840.113345.1.4.1789
LDAP Configuration for Netlogin:
dot1x           : Enabled
mac             : Enabled
web-based       : Enabled
LDAP Server 1   : engsrv1.engg.XYZCorp.com(192.168.3.101)
Server Port     : 389
Client IP       : 192.168.10.31
Client VR       : VR-Mgmt
Security Mechanism : Plain Text
Status          : Active
LDAP Server 2   : 192.168.3.102
Server Port     : 389
Client IP       : 192.168.10.31
Client VR       : VR-Mgmt
Security Mechanism : Plain Text
Status          : Not Active
-----
Domain          : it.XYZCorp.com
-----
Base-DN          : it.XYZCorp.com
Bind credential  : asingh
LDAP Hierarchical Search OID : None
```



```

LDAP Configuration for Netlogin:
dotlx      : Enabled
mac        : Enabled
web-based  : Enabled
LDAP Server 1      : 192.168.4.101
Server Port      : 389
Client IP       : 192.168.10.31
Client VR       : VR-Mgmt
Security Mechanism : Plain Text
Status          : Not Active
LDAP Server 2      : 192.168.4.102
Server Port      : 389
Client IP       : 192.168.10.31
Client VR       : VR-Mgmt
Security Mechanism : Plain Text
Status          : Active
-----
Domain       : mktg.XYZCorp.com
-----
Base-DN      : mktg.XYZCorp.com
Bind credential : gprasad
LDAP Hierarchical Search OID : LDAP_MATCHING_RULE_IN_CHAIN
(1.2.840.113556.1.4.1941)
LDAP Configuration for Netlogin:
dotlx      : Enabled
mac        : Enabled
web-based  : Enabled
LDAP Server 1      : mktgsrv1.mktg.XYZCorp.com(192.168.5.101)
Server Port      : 389
Client IP       : Any
Client VR       : VR-Mgmt
Security Mechanism : Plain Text
Status          : Active
LDAP Server 2      : 192.168.5.102
Server Port      : 389
Client IP       : Any
Client VR       : VR-Mgmt
Security Mechanism : Plain Text
Status          : Not Active
-----
Domain       : sales.XYZCorp.com
-----
Base-DN      : sales.XYZCorp.com
Bind credential : masiq
LDAP Hierarchical Search OID : LDAP_MATCHING_RULE_IN_CHAIN
(1.2.840.113556.1.4.1941)
LDAP Configuration for Netlogin:
dotlx      : Enabled
mac        : Enabled
web-based  : Enabled
LDAP Server      : No LDAP Servers configured
# show ldap domain "engg.XYZCorp.com"
-----
Domain       : engg.XYZCorp.com
-----
Base-DN      : engg.XYZCorp.com
Bind credential : pkumar
LDAP Hierarchical Search OID : LDAP_MATCHING_RULE_IN_CHAIN

```



```
(1.2.840.113556.1.4.1941)
LDAP Configuration for Netlogin:
dot1x          : Enabled
mac           : Enabled
web-based     : Enabled
LDAP Server 1 : engsrv1.engg.XYZCorp.com(192.168.3.101)
Server Port   : 389
Client IP     : 192.168.10.31
Client VR     : VR-Mgmt
Security Mechanism : Plain Text
Status       : Active
LDAP Server 2 : 192.168.3.102
Server Port   : 389
Client IP     : 192.168.10.31
Client VR     : VR-Mgmt
Security Mechanism : Plain Text
Status       : Not Active
```

If the server was specified as a host name and the IP address was not resolved, this is shown:

```
LDAP Server1  : server1.domain.com(IP address unresolved)
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on all platforms.

show ldap statistics

show ldap statistics

Description

This command displays LDAP packet statistics per LDAP domain.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to show all LDAP related statistics per LDAP domain.



Example

```
Switch.21 # show ldap statistics
-----
Domain          : XYZCorp.com (default)
-----
LDAP Server 1   : 192.168.2.101
Server Port    : 389
Client VR      : VR-Mgmt
Status         : Active
Requests       : 12
Responses      : 12
Errors         : 0
LDAP Server 2   : 192.168.2.102
Server Port    : 389
Client VR      : VR-Mgmt
Status         : Not Active
Requests       : 0
Responses      : 0
Errors         : 0
-----
Domain          : engg.XYZCorp.com
-----
LDAP Server 1   : engsrv1.engg.XYZCorp.com(192.168.3.101)
Server Port    : 389
Client VR      : VR-Mgmt
Status         : Active
Requests       : 22
Responses      : 20
Errors         : 2
LDAP Server 2   : 192.168.3.102
Server Port    : 389
Client VR      : VR-Mgmt
Status         : Not Active
Requests       : 0
Responses      : 0
Errors         : 0
-----
Domain          : it.XYZCorp.com
-----
LDAP Server 1   : 192.168.4.101
Server Port    : 389
Client VR      : VR-Mgmt
Status         : Not Active
Requests       : 1
Responses      : 0
Errors         : 1
LDAP Server 2   : 192.168.4.102
Server Port    : 389
Client VR      : VR-Mgmt
Status         : Active
Requests       : 6
Responses      : 6
Errors         : 0
-----
Domain          : mktg.XYZCorp.com
-----
```




```

LDAP Server 1 : 192.168.5.101
Server Port  : 389
Client VR    : VR-Mgmt
Status       : Not Active
Requests     : 8
Responses    : 7
Errors       : 1
LDAP Server 2 : 192.168.5.102
Server Port  : 389
Client VR    : VR-Mgmt
Status       : Active
Requests     : 12
Responses    : 12
Errors       : 0
-----
Domain       : sales.XYZCorp.com
-----
LDAP Server  : No LDAP Servers configured

```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on all platforms.

unconfigure identity-management

```

unconfigure identity-management {[[database memory-size] | [stale-entry aging-
time] | [ports] | [kerberos snooping {aging time}]]}

```

Description

Sets the specified identity management configuration parameter to the default values.

Syntax Description

database memory-size	Sets the identity management database size to the default value.
stale-entry aging-time	Sets the stale-entry aging-time to the default value.
ports	Removes all ports from the identity management port list.
kerberos snooping aging time	Sets the kerberos snooping aging time to the default value (none).

Default

N/A.



Usage Guidelines

If no configuration parameters are specified, all configuration parameters are set to the default values.

Example

The following command sets all identity management configuration parameters to the default values:

```
* Switch.4 # unconfigure identity-management
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

unconfigure identity-management list-precedence

unconfigure identity-management list-precedence

Description

This command allows you to restore the order of list-precedence to factory defaults. The default list-precedence is: greylist blacklist whitelist.

Syntax Description

This command has no arguments or variables.

Default

greylist, blacklist, whitelist.

Usage Guidelines

Use this command to restore the order of list-precedence to factory defaults. The default list-precedence is: greylist blacklist whitelist.

Example

Example output not yet available and will be provided in a future release.



History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms.

unconfigure ldap domains

unconfigure ldap domains

Description

This command deletes all LDAP domains, and thereby all LDAP servers and other LDAP configurations done for those domains.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to delete all LDAP related configuration from the switch.

Example

The following command deletes all LDAP configurations, LDAP servers and LDAP domains.

```
Switch.25 # unconfigure ldap domains
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on all platforms.



26 Security Commands

SSH

User Authentication

Denial of Service

```
clear ip-security anomaly-protection notify cache
clear ip-security arp validation violations
clear ip-security dhcp-snooping entries
clear ip-security source-ip-lockdown entries ports
clear vlan dhcp-address-allocation
configure dos-protect acl-expire
configure dos-protect interval
configure dos-protect trusted ports
configure dos-protect type I3-protect alert-threshold
configure dos-protect type I3-protect notify-threshold
configure ip-security anomaly-protection icmp ipv4-max-size
configure ip-security anomaly-protection icmp ipv6-max-size
configure ip-security anomaly-protection notify cache
configure ip-security anomaly-protection notify rate limit
configure ip-security anomaly-protection notify rate window
configure ip-security anomaly-protection notify trigger off
configure ip-security anomaly-protection notify trigger on
configure ip-security anomaly-protection tcp
configure ip-security dhcp-snooping information check
configure ip-security dhcp-snooping information circuit-id port-information port
configure ip-security dhcp-snooping information circuit-id vlan-information
configure ip-security dhcp-snooping information option
configure ip-security dhcp-snooping information policy
configure ip-security dhcp-bindings add
configure ip-security dhcp-bindings delete
configure ip-security dhcp-bindings storage
configure ip-security dhcp-bindings storage filename
configure ip-security dhcp-bindings storage location
configure mac-lockdown-timeout ports aging-time
configure ports rate-limit flood
configure ports vlan
configure radius server client-ip
configure radius shared-secret
configure radius timeout
```

```
configure radius-accounting server client-ip
configure radius-accounting shared-secret
configure radius-accounting timeout
configure ssh2 key
configure sshd2 user-key add user
configure sshd2 user-key delete user
configure ssl certificate pregenerated
configure ssl certificate privkeylen
configure ssl privkey pregenerated
configure tacacs server client-ip
configure tacacs shared-secret
configure tacacs timeout
configure tacacs-accounting server
configure tacacs-accounting shared-secret
configure tacacs-accounting timeout
configure trusted-ports trust-for dhcp-server
configure trusted-servers add server
configure trusted-servers delete server
configure vlan dhcp-address-range
configure vlan dhcp-lease-timer
configure vlan dhcp-options
create sshd2 key-file
create sshd2 user-key
delete sshd2 user-key
disable dhcp ports vlan
disable dos-protect
disable iparp gratuitous protect vlan
disable ip-security anomaly-protection
disable ip-security anomaly-protection ip
disable ip-security anomaly-protection l4port
disable ip-security anomaly-protection tcp flags
disable ip-security anomaly-protection tcp fragment
disable ip-security anomaly-protection icmp
disable ip-security anomaly-protection notify
disable ip-security arp gratuitous-protection
disable ip-security arp learning learn-from-arp
disable ip-security arp learning learn-from-dhcp
disable ip-security arp validation
disable ip-security dhcp-bindings restoration
disable ip-security dhcp-snooping
disable ip-security source-ip-lockdown ports
disable mac-lockdown-timeout ports
```



```
disable radius
disable radius-accounting
disable ssh2
disable tacacs
disable tacacs-accounting
disable tacacs-authorization
disable web http
disable web https
download ssl certificate
download ssl privkey
enable dhcp ports vlan
enable dos-protect
enable dos-protect simulated
enable iparp gratuitous protect
enable ip-option loose-source-route
enable ip-security anomaly-protection
enable ip-security anomaly-protection icmp
enable ip-security anomaly-protection ip
enable ip-security anomaly-protection l4port
enable ip-security anomaly-protection notify
enable ip-security anomaly-protection tcp flags
enable ip-security anomaly-protection tcp fragment
enable ip-security arp gratuitous-protection
enable ip-security arp learning learn-from-arp
enable ip-security arp learning learn-from-dhcp
enable ip-security arp validation violation-action
enable ip-security dhcp-bindings restoration
enable ip-security dhcp-snooping
enable ip-security source-ip-lockdown ports
enable mac-lockdown-timeout ports
enable radius
enable radius-accounting
enable ssh2
enable tacacs
enable tacacs-accounting
enable tacacs-authorization
enable web http
enable web https
scp2
show dhcp-server
show dos-protect
show ip-security anomaly-protection notify cache ports
```



```
show ip-security arp gratuitous-protection
show ip-security arp learning
show ip-security arp validation
show ip-security arp validation violations
show ip-security dhcp-snooping entries
show ip-security dhcp-snooping information-option
show ip-security dhcp-snooping information circuit-id port-information
show ip-security dhcp-snooping information-option circuit-id vlan-information
show ip-security dhcp-snooping
show ip-security dhcp-snooping violations
show ip-security source-ip-lockdown
show mac-lockdown-timeout fdb ports
show mac-lockdown-timeout ports
show ports rate-limit flood
show radius
show radius-accounting
show ssh2 private-key
show sshd2 user-key
show ssl
show tacacs
show tacacs-accounting
show vlan dhcp-address-allocation
show vlan dhcp-config
show vlan security
ssh2
unconfigure ip-security dhcp-snooping information check
unconfigure ip-security dhcp-snooping information circuit-id port-information ports
unconfigure ip-security dhcp-snooping information circuit-id vlan-information
unconfigure ip-security dhcp-snooping information option
unconfigure ip-security dhcp-snooping information policy
unconfigure radius
unconfigure radius-accounting
unconfigure tacacs
unconfigure tacacs-accounting
unconfigure trusted-ports trust-for dhcp-server
unconfigure vlan dhcp
unconfigure vlan dhcp-address-range
unconfigure vlan dhcp-options
upload dhcp-bindings
```

This chapter describes commands for:

- Managing the switch using SSH2.
- Configuring switch user authentication through a RADIUS client.



- Configuring switch user authentication through TACACS+.
- Protecting the switch from Denial of Service attacks.

SSH

Secure Shell 2 (SSH2) is a feature of ExtremeXOS that allows you to encrypt session data between a network administrator using SSH2 client software and the switch. Configuration and policy files may also be transferred to the switch using the Secure Copy Program 2 (SCP2).

User Authentication

Remote Authentication Dial In User Service (RADIUS, RFC 2138) is a mechanism for authenticating and centrally administrating access to network nodes. The ExtremeXOS RADIUS client implementation allows authentication for SSH2, Telnet or console access to the switch.

Extreme switches are also capable of sending RADIUS accounting information. You can configure RADIUS accounting servers to be the same as the authentication servers, but this is not required.

Terminal Access Controller Access Control System Plus (TACACS+) is a mechanism for providing authentication, authorization, and accounting on a centralized server, similar in function to the RADIUS client. The ExtremeXOS version of TACACS+ is used to authenticate prospective users who are attempting to administer the switch. TACACS+ is used to communicate between the switch and an authentication database.



Note

You cannot use RADIUS and TACACS+ at the same time.

Denial of Service

You can configure ExtremeXOS to protect your Extreme switches in the event of a denial of service attack. During a typical denial of service attack, the CPU on the switch gets flooded with packets from multiple attackers, potentially causing the switch to fail. To protect against this type of attack, you can configure the software so that when the number of packets received is more than the configured threshold limit of packets per second, a hardware ACL is enabled.

clear ip-security anomaly-protection notify cache

```
clear ip-security anomaly-protection notify cache {slot [slot | all ]}
```

Description

Clear the local protocol anomaly event cache.



Syntax Description

<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

N/A.

Usage Guidelines

This command clears the local protocol anomaly event cache.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on the Summit X440, X460, X480, X670, and X770 platforms, whether or not included in a SummitStack, E4G-200 and E4G-400 switches, and the BlackDiamond X8 series switches and BlackDiamond 8000 c-, e-, xl-, and xm-series modules.

clear ip-security arp validation violations

clear ip-security arp validation violations

Description

Clears the violation counters.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command clears the ARP validation violation counters.

History

This command was first available in ExtremeXOS 12.1.



Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

clear ip-security dhcp-snooping entries

```
clear ip-security dhcp-snooping entries { vlan } vlan_name
```

Description

Clears the DHCP binding entries present on a VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN of the DHCP server.
------------------	--

Default

N/A.

Usage Guidelines

Use this command to clear the DHCP binding entries present on a VLAN. When an entry is deleted, all its associated entries (such as source IP lockdown, secured ARP, and so on) and their associated ACLs, if any, are also deleted.

Example

The following command clears the DHCP binding entry temporary from the VLAN:

```
clear ip-security dhcp-snooping entries temporary
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

clear ip-security source-ip-lockdown entries ports

```
clear ip-security source-ip-lockdown entries ports [ ports | all ]
```



Description

Clears locked-down source IP addresses on a per-port basis.

Syntax Description

<i>ports</i>	Specifies the port or ports to be cleared.
all	Specifies that all ports are to be cleared.

Default

N/A.

Usage Guidelines

Use this command to clear locked-down source IP addresses on a per port basis. This command deletes the entries on the indicated ports and clears the associated ACLs.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

clear vlan dhcp-address-allocation

```
clear vlan vlan_name dhcp-address-allocation [[all {offered | assigned | declined | expired}] | ipaddress]
```

Description

Removes addresses from the DHCP allocation table.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN of the DHCP server.
all	Specifies all IP addresses, or all IP addresses in a particular state.
offered	Specifies IP addresses offered to clients.
assigned	Specifies IP addresses offered to and accepted by clients.
declined	Specifies IP addresses declined by clients.



expired	Specifies IP addresses whose lease has expired and not renewed by the DHCP server.
<i>ipaddress</i>	Specifies a particular IP address.

Default

N/A.

Usage Guidelines

You can delete either a single entry, using the IP address, or all entries. If you use the all option, you can additionally delete entries in a specific state.

Example

The following command removes all the declined IP addresses by hosts on the VLAN temporary:

```
clear vlan temporary dhcp-address-allocation all declined
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure dos-protect acl-expire

```
configure dos-protect acl-expire seconds
```

Description

Configures the denial of service protection ACL expiration time.

Syntax Description

<i>seconds</i>	Specifies how long the ACL is in place.
----------------	---

Default

The default is 5 seconds.



Usage Guidelines

This command configures how long the DoS protection ACL remains in place.

Example

This example sets the ACL expiration time to 15 seconds:

```
configure dos-protect acl-expire 15
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure dos-protect interval

```
configure dos-protect interval seconds
```

Description

Configures the denial of service protection interval.

Syntax Description

<i>seconds</i>	Specifies how often the DoS protection counter is monitored.
----------------	--

Default

The default is one second.

Usage Guidelines

This command configures how often the DoS protection counter is monitored.

Example

This example sets the interval to 5 seconds:

```
configure dos-protect interval 5
```



History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure dos-protect trusted ports

```
configure dos-protect trusted-ports [ports [ports | all] | add-ports [ports-to-add | all] | delete-ports [ports-to-delete | all]]
```

Description

Configures the list of trusted ports.

Syntax Description

<i>ports</i>	Specifies the trusted ports list.
<i>ports-to-add</i>	Specifies the ports to add to the trusted ports list.
all	Specifies all the ports.
<i>ports-to-delete</i>	Specifies the ports to delete from the trusted ports list.

Default

N/A.

Usage Guidelines

Traffic from trusted ports will be ignored when DoS protect counts the packets to the CPU. If we know that a machine connected to a certain port on the switch is a safe "trusted" machine, and we know that we will not get a DoS attack from that machine, the port to which this machine is connected can be configured as a trusted port, even though a large amount of traffic is going through this port.

Example

This example sets the trusted port list to 3:1-3:7:

```
configure dos-protect trusted-ports ports 3:1-3:7
```

This example adds the trusted port 3:8 to the current list (use this command with a network administrator machine not connected to the internet that is attached to port 3:8):

```
configure dos-protect trusted-ports add-ports 3:8
```



History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure dos-protect type l3-protect alert-threshold

```
configure dos-protect type l3-protect alert-threshold packets
```

Description

Configures the denial of service protection alert threshold.

Syntax Description

<i>packets</i>	Specifies how many packets in an interval will cause an alert.
----------------	--

Default

The default is 4000 packets.

Usage Guidelines

This command configures how many packets received in an interval will cause a DoS protection alert. When an alert occurs, the packets are analyzed, and a temporary ACL is applied to the switch.

Example

This example sets the alert threshold to 8000 packets:

```
configure dos-protect type l3-protect alert-threshold 8000
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.



configure dos-protect type l3-protect notify-threshold

```
configure dos-protect type l3-protect notify-threshold packets
```

Description

Configures the denial of service protection notification threshold.

Syntax Description

<i>packets</i>	Specifies how many packets in an interval will cause a notification.
----------------	--

Default

The default is 3500 packets.

Usage Guidelines

This command configures how many packets received in an interval will cause a DoS protection notification.

Example

This example sets the notification threshold to 7500 packets:

```
configure dos-protect type l3-protect notify-threshold 7500
```

History

This command was first available in ExtremeXOS 11.1

Platform Availability

This command is available on all platforms.

configure ip-security anomaly-protection icmp ipv4-max-size

```
configure ip-security anomaly-protection icmp ipv4-max-size size {slot [slot |  
all ]}
```

Description

Configures the maximum IPv4 ICMP allowed size.



Syntax Description

<i>size</i>	Specifies the size of the IPv4 ICMP in bytes.
<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default size is 512 bytes.

Usage Guidelines

This command configures the IPv4 ICMP allowed size. The absolute maximum is 1023 bytes.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X440, X460, X480, X670, and X770 platforms, whether or not included in a SummitStack, and on the BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules.

configure ip-security anomaly-protection icmp ipv6-max-size

```
configure ip-security anomaly-protection icmp ipv6-max-size size {slot [ slot | all ]}
```

Description

Configures the maximum ipv6 ICMP allowed size.

Syntax Description

<i>size</i>	Specifies the size of the IPv6 ICMP in bytes.
<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default size is 512 bytes.



Usage Guidelines

This command configures the IPv6 ICMP allowed size. The absolute maximum is 16K bytes.

You can use this command to configure the maximum IPv6 ICMP packet size for detecting IPv6 ICMP anomalies. If the next header in the IPv6 ICMP packet is not 0x3A:ICMP, this anomaly is not detected. For example, an IPv6 ICMP packet with packet header 0x2c: Fragment Header is not detected.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X440, X460, X480, X670, and X770 platforms, whether or not included in a SummitStack, and on the BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules.

configure ip-security anomaly-protection notify cache

```
configure ip-security anomaly-protection notify cache size {slot [slot | all ]}
```

Description

Configures the size of local notification cache.

Syntax Description

<i>size</i>	Specifies the size of the local notification cache.
<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is 1000 events.

Usage Guidelines

This command configures the size of local notification cache. Cached events are stored in local memory. The range is between 1 and 1000 events per second. If the cache is full, newer events replace older events.

History

This command was first available in ExtremeXOS 12.0.



Platform Availability

This command is only available on the Summit X440, X460, X480, X670, and X770 platforms, whether or not included in a SummitStack, and on the BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules.

configure ip-security anomaly-protection notify rate limit

```
configure ip-security anomaly-protection notify rate limit value {slot [slot | all ]}
```

Description

Configures the rate limiting for protocol anomaly notification.

Syntax Description

<i>value</i>	Specifies the period of the rate limit.
<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is 10 events per second.

Usage Guidelines

This is a paired command with `configure ip-security anomaly-protection notify rate window` that configures the rate limiting for protocol anomaly notification. When the anomaly notification is enabled, in order to avoid overloading CPU, the system generates only the number of limited notifications in a period of window seconds. The range is from 1 to 100 events.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X440, X460, X480, X670, and X770 platforms, whether or not included in a SummitStack, and on the BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules.

configure ip-security anomaly-protection notify rate window

```
configure ip-security anomaly-protection notify rate window value {slot [slot | all ]}
```



Description

Configures the rate limiting for protocol anomaly notification.

Syntax Description

<i>value</i>	Specifies the period of the rate limit.
<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is 1 second.

Usage Guidelines

This is a paired command with `configure ip-security anomaly-protection notify rate limit` that configures the rate limiting for protocol anomaly notification. When the anomaly notification is enabled, in order to avoid overloading CPU, the system generates only the number of limited notifications in a period of window seconds. The range is between 1 and 300 seconds.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X440, X460, X480, X670, and X770 platforms, whether or not included in a SummitStack, and on the BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules.

configure ip-security anomaly-protection notify trigger off

```
configure ip-security anomaly-protection notify trigger off value {slot [slot | all ]}
```

Description

Configures an anomaly rate-based notification feature.

Syntax Description

<i>value</i>	Specifies the number of events for the trigger.
<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.



Default

The default is 1.

Usage Guidelines

This is a paired command with `configure ip-security anomaly-protection notify trigger on` that configures an anomaly rate-based notification feature. The anomaly notification is automatically triggered if the rate of anomaly events is greater than the configured ON value, and the notification is disabled if the rate falls below the value set in the `configure ip-security anomaly-protection notify trigger off` command.

The command takes effects after the anomaly notification is enabled.



Note

The value set in ON must be greater than or equal to the value set in OFF.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X440, X460, X480, X670, and X770 platforms, whether or not included in a SummitStack, and the BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules.

configure ip-security anomaly-protection notify trigger on

```
configure ip-security anomaly-protection notify trigger on value {slot [slot | all ]}
```

Description

Configures an anomaly rate-based notification feature.

Syntax Description

<i>value</i>	Specifies the number of events for the trigger.
<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is 1.



Usage Guidelines

This is a paired command with `configure ip-security anomaly-protection notify trigger off` that configures an anomaly rate-based notification feature. The anomaly notification is automatically triggered if the rate of anomaly events is greater than the configured ON value, and the notification is disabled if the rate falls below the value set in the `configure ip-security anomaly-protection notify trigger off` command.

The command takes effects after the anomaly notification is enabled.



Note

The value set in ON must be greater than or equal to the value set in OFF.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X440, X460, X480, X670, and X770 platforms, whether or not included in a SummitStack, and the BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules.

configure ip-security anomaly-protection tcp

```
configure ip-security anomaly-protection tcp min-header-size size {slot [ slot | all ]}
```

Description

Configures the minimum TCP header allowed.

Syntax Description

<i>size</i>	Specifies the size of the header in bytes.
<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default value is 20 bytes.

Usage Guidelines

This command configures the minimum TCP header allowed. It takes effect for both IPv4 and IPv6 TCP packets.



The range of the minimum TCP header may be between 8 and 255 bytes.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X440, X460, X480, X670, and X770 platforms, whether or not included in a SummitStack, and the BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules.

configure ip-security dhcp-snooping information check

```
configure ip-security dhcp-snooping information check
```

Description

Enables the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82) checking in the server-originated packets.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command enables the checking of the server-originated packets for the presence of option 82. In some instances, a DHCP server may not properly handle a DHCP request packet containing a relay agent option. Use this command to prevent DHCP reply packets with invalid or missing relay agent options from being forwarded to the client. With checking enabled, the following checks and actions are performed:

- When the option 82 is present in the packet, the MAC address specified in the remote-ID sub-option is the switch system MAC address. If the check fails, the packet is dropped.
- When option 82 is not present in the packet, the DHCP packet is forwarded with no modification.

To disable this check, use the following command:

```
unconfigure ip-security dhcp-snooping information check
```



Example

The following command enables DHCP relay agent option checking:

```
configure ip-security dhcp-snooping information check
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

configure ip-security dhcp-snooping information circuit-id port-information port

```
configure ip-security dhcp-snooping information circuit-id port-information
port_info port port
```

Description

Configures the port information portion of the circuit ID.

Syntax Description

<i>port_info</i>	Specifies the circuit ID port information in the format of <i>VLAN Info - Port Info</i> ; maximum length is 32 bytes.
<i>port</i>	Specifies the port for which DHCP Snooping should be enabled.

Default

The default value is the ASCII representation of the ingress port's SNMP ifIndex.

Usage Guidelines

This command allows you to configure the port information portion of the circuit ID whose format is *vlan_info - port_info* for each port. The parameter *port_info* is a string of up to 32 bytes in length. When a specific value is not configured for port information, the *port_info* defaults to the ASCII representation of the ingress ports's SNMP ifIndex.

History

This command was first available in ExtremeXOS 12.3.



Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

configure ip-security dhcp-snooping information circuit-id vlan-information

```
configure ip-security dhcp-snooping information circuit-id vlan-information
vlan_info {vlan} [vlan_name | all]
```

Description

Configures the VLAN info portion of the circuit ID of a VLAN.

Syntax Description

<i>vlan_info</i>	Specifies the circuit ID VLAN information for each VLAN in the format of <i>VLAN Info-Port Info</i> ; maximum length is 32 bytes.
<i>vlan_name</i>	Specifies the VLAN for which DHCP should be enabled.
all	Specifies all VLANs.

Default

The default value is the ASCII representation of the ingress VLAN's ID.

Usage Guidelines

This command allows you to configure the VLAN information portion of the circuit ID of a VLAN. The VLAN info is a string of characters of up to 32 bytes in length, and is entered in the format of *VLAN InfoPort Info*. When a specific value is not configured for a VLAN, *vlan_info* defaults to the ASCII representation of the ingress VLAN's ID.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

configure ip-security dhcp-snooping information option

```
configure ip-security dhcp-snooping information option
```



Description

Enables the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82).

Syntax Description

This command has no arguments or variables.

Default

The default is unconfigured.

Usage Guidelines

This command enables the DHCP relay agent option (option 82), which is inserted into client-originated DHCP packets before they are forwarded to the server.

To disable the DHCP relay agent option (option 82), use the following command:

```
unconfigure ip-security dhcp-snooping information option
```

Example

The following command enable the DHCP relay agent option:

```
configure ip-security dhcp-snooping information information option
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

configure ip-security dhcp-snooping information policy

```
configure ip-security dhcp-snooping information policy [drop | keep | replace]
```

Description

Configures the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82) policy.



Syntax Description

drop	Specifies to drop the packet.
keep	Specifies to keep the existing option 82 information in place.
replace	Specifies to replace the existing data with the switch's own data.

Default

The default value is replace.

Usage Guidelines

Use this command to set a policy for the relay agent. Packets can be dropped, the option 82 information can be replaced (the default), or the packet can be forwarded with the information unchanged.

Example

The following command configures the DHCP relay agent option 82 policy to keep:

```
configure ip-security dhcp-snooping information information policy keep
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

configure ip-security dhcp-bindings add

```
configure ip-security dhcp-binding add ip ip_address mac mac_address {vlan}
vlan_name server-port server_port client-port client_port lease-time seconds
```

Description

Creates a DHCP binding.

Syntax Description

<i>ip_address</i>	Specifies the IP address for the DHCP binding.
<i>mac_address</i>	Specifies the MAC address for the DHCP binding.



<i>vlan_name</i>	Specifies the name of the VLAN for the DHCP binding.
<i>server_port</i>	Specifies the server port for the DHCP binding.
<i>client_port</i>	Specifies the client port for the DHCP binding.
<i>seconds</i>	Specifies the number of seconds for the lease.

Default

N/A.

Usage Guidelines

This commands allows you to add a DHCP binding in order to re-create the bindings after reboot and to allow IP Security features to work with clients having static IP addresses.



Note

Setting the lease-time to 0 causes the DHCP binding to be static; in other words, it is not aged-out if no DHCP renew occurs. This is for use with clients using static IP addresses.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

configure ip-security dhcp-bindings delete

```
configure ip-security dhcp-binding delete ip ip_address {vlan} vlan_name
```

Description

Deletes a DHCP binding.

Syntax Description

<i>ip_address</i>	Specifies the IP address for the DHCP binding.
<i>vlan_name</i>	Specifies the name of the VLAN for the DHCP binding.

Default

N/A.



Usage Guidelines

This command allows you to delete a DHCP binding created with the command `configure ip-security dhcp-binding add ip ip_address mac mac_address {vlan}vlan_name server-portserver_port client-portclient_port lease-timeseconds`.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

configure ip-security dhcp-bindings storage

```
configure ip-security dhcp-bindings storage [write-interval minutes | write-threshold num_changed_entries]
```

Description

Configures DHCP bindings file storage upload variables.

Syntax Description

<i>minutes</i>	Specifies the number of minutes for the write interval.
<i>num_changed_entries</i>	Specifies the limit for the write threshold.

Default

The default write threshold is 50 entries; the default write interval is 30 minutes.

Usage Guidelines

This command allows you to configure the upload variables for the DHCP bindings file that you created with the command `configure ip-security dhcp-bindings storage filename name` and specified the location of with the command `configure ip-security dhcp-bindings storage location server [primary | secondary] ip_address [hostname]{vrvr-name} tftp`.

For redundancy, the DHCP bindings file is uploaded to both the primary and the secondary server. The failure of one upload (for example, due to a TFTP server timeout) does not affect the upload of any other.

When the maximum file size limit is reached, no additional DHCP bindings can be uploaded until one of the older bindings is removed.



The point at which DHCP bindings can be uploaded can be configured to work in one of the following ways:

- **Periodic upload:** Upload every N minutes, provided that DHCP bindings have changed since the last upload.
- **Upload based on number of yet-to-be uploaded entries:** Allows you to configure the maximum number of changed entries that are allowed to accumulate before being uploaded.

The write interval is configurable from 5 minutes to 1 day, with a default value of 30 minutes. The default value of the write threshold is 50 entries, with a minimum of 25 and maximum of 200.

Additions and deletions are considered changes, but updates are not, which means that DHCP renewals of existing leases are not counted.

By default, the write interval is in effect, but not the write-threshold. You may change whichever of these you wish by explicitly configuring the value.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

configure ip-security dhcp-bindings storage filename

```
configure ip-security dhcp-bindings storage filename name
```

Description

Creates a storage file for DHCP binding information.

Syntax Description

<i>name</i>	Specifies the name of the DHCP binding storage file.
-------------	--

Default

N/A.

Usage Guidelines

This commands allows you to configure the filename with which the DHCP bindings storage file is created on the external server when it is uploaded to the external server. The text file resides on an external server. You can configure the server with the command `configure ip-security dhcp-`



```
bindings storage location server [primary | secondary] ip_address | hostname]
{vrvr-name} tftp.
```

The bindings file must have a .xsf extension. If the input filename doesn't already have a .xsf extension, one is added automatically.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

configure ip-security dhcp-bindings storage location

```
configure ip-security dhcp-bindings storage location server [primary | secondary]
ip_address | hostname]{vr vr-name} tftp
```

Description

Specifies the server location for the DHCP bindings storage file. The uploads can be made to any tftp server regardless of the virtual router that it is present in.

Syntax Description

<i>ip_address</i>	Specifies the IP address location for the bindings storage file.
<i>hostname</i>	Specifies the hostname of the server.
<i>vr-name</i>	Specifies the virtual router name.

Default

N/A.

Usage Guidelines

This commands allows you to specify where you want to store the DHCP storage file that you created with the command `configure ip-security dhcp-bindings storage filename name`.



Example

The following command configures storage to the primary server 10.1.1.14:

```
configure ip-security dhcp-bindings storage location server primary 10.1.1.14
vr "VR-Default" tftp
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

configure mac-lockdown-timeout ports aging-time

```
configure mac-lockdown-timeout ports [all | port_list] aging-time seconds
```

Description

Configures the MAC address lock down timeout value in seconds for the specified port or group of ports or for all ports on the switch.

Syntax Description

all	Specifies all ports.
<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>seconds</i>	Configures the length of the time out value in seconds. The default is 15 seconds; the range is 15 to 2,000,000 seconds.

Default

The default is 15 seconds.

Usage Guidelines

This timer overrides the FDB aging time.

This command only sets the duration of the MAC address lock down timer. To enable the lock down timeout feature, use the following command:

```
enable mac-lockdown-timeout ports [all | port_list]
```



Example

The following command configures the MAC address lock down timer duration for 300 seconds for ports 2:3, 2:4, and 2:6:

```
configure mac-lockdown-timeout ports 2:3, 2:4, 2:6 aging-time 300
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure ports rate-limit flood

```
configure ports port_list rate-limit flood [broadcast | multicast | unknown-destmac] [no-limit | pps]
```

Description

Limits the amount of ingress flooded traffic; minimizes network impact of broadcast loops.

Syntax Description

<i>port_list</i>	Specifies the port number. On a stand-alone switch, this value is just the port number, and on a modular switch, this value is the slot and port number.
broadcast	Specifies all broadcast packets.
multicast	Specifies all flooded multicast packets (known IP multicast caches are still forwarded at line rate).
unknown-destmac	Specifies all packets with unknown MAC DAs.
no-limit	Specifies unlimited rate.
<i>pps</i>	Packets per second allowed; range is from 0 to 262,144.

Default

No limit.

Usage Guidelines

Use this command to limit the amount of ingress flooding traffic and to minimize the network impact of broadcast loops.

To display results, use the `show ports rate-limit flood` command.



Example

The following command rate limits broadcast packets on port 3 on a stand-alone switch to 500 pps:

```
configure ports 3 rate-limit flood broadcast 500
```

History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available on the BlackDiamond X8, BlackDiamond 8800 series switches, SummitStack, and the Summit family of switches.

configure ports vlan

```
configure ports port_list {tagged tag} vlan vlan_name [limit-learning number
{action [blackhole | stop-learning]} | lock-learning | unlimited-learning |
unlock-learning]
```

Description

Configures virtual ports for limited or locked MAC address learning.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
tagged tag	Specifies the port-specific VLAN tag. When there are multiple ports specified in the <i>port_list</i> , the same tag is used for all of them.
<i>vlan_name</i>	Specifies the name of the VLAN.
limit-learning number	Specifies a limit on the number of MAC addresses that can be dynamically learned on the specified ports.
blackhole	Specifies that blackhole entries are allowed.
stop-learning	Specifies that the learning be halted to protect the switch from exhausting FDB resources by not creating blackhole entries.
lock-learning	Specifies that the current FDB entries for the specified ports should be made permanent static, and no additional learning should be allowed.
unlimited-learning	Specifies that there should not be a limit on MAC addresses that can be learned.
unlock-learning	Specifies that the port should be unlocked (allow unlimited, dynamic learning).



Default

Unlimited, unlocked learning.

Usage Guidelines

If you have enabled ESRP, see the *ExtremeXOS Concepts Guide* for information about using this feature with ESRP.

Limited learning

The limited learning feature allows you to limit the number of dynamically-learned MAC addresses per VLAN. When the learned limit is reached, all new source MAC addresses are blackholed at both the ingress and egress points. This prevents these MAC addresses from learning and responding to Internet control message protocol (ICMP) and address resolution protocol (ARP) packets.

If the limit you configure is greater than the current number of learned entries, all the current learned entries are purged.

Dynamically learned entries still get aged, and can be cleared. If entries are cleared or aged out after the learning limit has been reached, new entries will then be able to be learned until the limit is reached again.

Permanent static and permanent dynamic entries can still be added and deleted using the `create fdbentry` and `delete fdbentry` commands. These override any dynamically learned entries.

For ports that have a learning limit in place, the following traffic still flows to the port:

- Packets destined for permanent MACs and other non-blackholed MACs.
- Broadcast traffic.
- EDP traffic.

Traffic from the permanent MAC and any other non-blackholed MACs will still flow from the virtual port.

If you configure a MAC address limit on VLANs that participate in an Extreme Standby Router Protocol (ESRP) domain, you should add an additional back-to-back link (that has no MAC address limit on these ports) between the ESRP-enabled switches. Doing so prevents ESRP protocol data units (PDUs) from being dropped due to MAC address limit settings.

Stop learning

When stop-learning is enabled with learning-limit configured, the switch is protected from exhausting FDB resources by not creating blackhole entries. Any additional learning and forwarding is prevented, but packet forwarding from FDB entries is not impacted.

Port lockdown

The port lockdown feature allows you to prevent any additional learning on the virtual port, keeping existing learned entries intact. This is equivalent to making the dynamically-learned entries permanent static, and setting the learning limit to zero. All new source MAC addresses are blackholed.



Locked entries do not get aged, but can be deleted like any other permanent FDB entries. The maximum number of permanent lockdown entries is 1024. Any FDB entries above will be flushed and blackholed during lockdown.

For ports that have lockdown in effect, the following traffic still flows to the port:

- Packets destined for the permanent MAC and other non-blackholed MACs.
- Broadcast traffic.
- EDP traffic.

Traffic from the permanent MAC will still flow from the virtual port.

Once the port is locked down, all the entries become permanent and will be saved across reboot.

When you remove the lockdown using the unlock-learning option, the learning-limit is reset to unlimited, and all associated entries in the FDB are flushed.

To display the locked entries on the switch, use the following command:

```
show fdb
```

Locked MAC address entries have the "I" flag.

To verify the MAC security configuration for the specified VLAN or ports, use the following commands:

```
show vlan vlan name security show ports port_list info detail
```

Example

The following command limits the number of MAC addresses that can be learned on ports 1, 2, 3, and 6 in a VLAN named accounting, to 128 addresses:

```
configure ports 1, 2, 3, 6 vlan accounting learning-limit 128
```

The following command locks ports 4 and 5 of VLAN accounting, converting any FDB entries to static entries, and prevents any additional address learning on these ports:

```
configure ports 4,5 vlan accounting lock-learning
```

The following command removes the learning limit from the specified ports:

```
configure ports 1, 2, vlan accounting unlimited-learning
```

The following command unlocks the FDB entries for the specified ports:

```
configure ports 4,5 vlan accounting unlock-learning
```



The following example illustrates use of the **tagged** keyword:

```
configure ports 1 tag 10 vlan accounting learning-limit 128
configure ports 1 vlan accounting learning-limit 128
configure ports 4 tag 10 vlan accounting lock-learning
configure ports 4 vlan accounting lock-learning
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure radius server client-ip

```
configure radius {mgmt-access | netlogin} [primary | secondary] server [ipaddress
| hostname] {udp_port} client-ip [ipaddress] {vr vr_name}
```

Description

Configures the primary and secondary RADIUS authentication server.

Syntax Description

mgmt-access	Specifies the RADIUS authentication server for switch management.
netlogin	Specifies the RADIUS authentication server for network login.
primary	Configures the primary RADIUS authentication server.
secondary	Configures the secondary RADIUS authentication server.
<i>ipaddress</i>	The IP address of the server being configured.
<i>hostname</i>	The host name of the server being configured.
<i>udp_port</i>	The UDP port to use to contact the RADIUS authentication server.
<i>ipaddress</i>	The IP address used by the switch to identify itself when communicating with the RADIUS authentication server.
<i>vr_name</i>	Specifies the virtual router on which the client IP is located. NOTE: User-created VRs are supported only on the platforms listed for this feature in Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i> .

Default

The following lists the default behavior of this command:

- The UDP port setting is 1812.
- The virtual router used is VR-Mgmt, the management virtual router.



- Switch management and network login use the same primary and secondary RADIUS servers for authentication.

Usage Guidelines

Use this command to specify RADIUS server information.

Use of the *hostname* parameter requires that DNS be enabled.

The RADIUS server defined by this command is used for user name authentication and CLI command authentication.

Beginning with ExtremeXOS 11.2, you can specify one pair of RADIUS authentication servers for switch management and another pair for network login. To specify RADIUS authentication servers for switch management (Telnet, SSH, and console sessions), use the *mgmt-access* keyword. To specify RADIUS authentication servers for network login, use the *netlogin* keyword. If you do not specify a keyword, switch management and network login use the same pair of RADIUS authentication servers.

If you are running ExtremeXOS 11.1 or earlier and upgrade to ExtremeXOS 11.2, you do not lose your existing RADIUS server configuration. Both switch management and network login use the RADIUS authentication server specified in the older configuration.

Example

The following command configures the primary RADIUS server on host *radius1* using the default UDP port (1812) for use by the RADIUS client on switch 10.10.20.30 using a virtual router interface of *VR-Default*:

```
configure radius primary server radius1 client-ip 10.10.20.30 vr vr-Default
```

The following command configures the primary RADIUS server for network login authentication on host *netlog1* using the default UDP port for use by the RADIUS client on switch 10.10.20.31 using, by default, the management virtual router interface:

```
configure radius netlogin primary server netlog1 client-ip 10.10.20.31
```

History

This command was first available in ExtremeXOS 10.1.

The *mgmt-access* and *netlogin* keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.



configure radius shared-secret

```
configure radius {mgmt-access | netlogin} [primary | secondary] shared-secret
{encrypted} string
```

Description

Configures the authentication string used to communicate with the RADIUS authentication server.

Syntax Description

mgmt-access	Specifies the switch management RADIUS authentication server.
netlogin	Specifies the network login RADIUS authentication server.
primary	Configures the authentication string for the primary RADIUS server.
secondary	Configures the authentication string for the secondary RADIUS server.
encrypted	Indicates that the string is already encrypted.
<i>string</i>	The string to be used for authentication.

Default

Unconfigured.

Usage Guidelines

The secret must be the same between the client switch and the RADIUS server.

The RADIUS server must first be configured for use with the switch as a RADIUS client.

The mgmt-access keyword specifies the RADIUS server used for switch management authentication.

The netlogin keyword specifies the RADIUS server used for network login authentication.

If you do not specify the mgmt-access or netlogin keywords, the secret applies to both the primary or secondary switch management and netlogin RADIUS servers.

The encrypted keyword is primarily for the output of the show configuration command, so the shared secret is not revealed in the command output. Do not use it to set the shared secret.

Example

The following command configures the shared secret as “purplegreen” on the primary RADIUS server for both switch management and network login:

```
configure radius primary shared-secret purplegreen
```



The following command configures the shared secret as “redblue” on the primary switch management RADIUS server:

```
configure radius mgmt-access primary shared-secret redblue
```

History

This command was first available in ExtremeXOS 10.1.

The encrypted keyword was added in ExtremeXOS 11.0.

The mgmt-access and netlogin keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure radius timeout

```
configure radius {mgmt-access | netlogin} timeout seconds
```

Description

Configures the timeout interval for RADIUS authentication requests.

Syntax Description

mgmt-access	Specifies the switch management RADIUS authentication server.
netlogin	Specifies the network login RADIUS authentication server.
<i>seconds</i>	Specifies the number of seconds for authentication requests. Range is 3 to 120 seconds.

Default

The default is 3 seconds.

Usage Guidelines

This command configures the timeout interval for RADIUS authentication requests. When the timeout has expired, another authentication attempt will be made. After three failed attempts to authenticate, the alternate server will be used. After six failed attempts, local user authentication will be used.

The mgmt-access keyword specifies the RADIUS server used for switch management authentication.

The netlogin keyword specifies the RADIUS server used for network login authentication.



If you do not specify the `mgmt-access` or `netlogin` keywords, the timeout interval applies to both switch management and netlogin RADIUS servers.

Example

The following command configures the timeout interval for RADIUS authentication to 10 seconds. After 30 seconds (three attempts), the alternate RADIUS server will be used. After 60 seconds (six attempts) local user authentication is used.

```
configure radius timeout 10
```

History

This command was first available in ExtremeXOS 10.1.

The `mgmt-access` and `netlogin` keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure radius-accounting server client-ip

```
configure radius-accounting {mgmt-access | netlogin} [primary | secondary] server
[ipaddress | hostname] {tcp_port} client-ip [ipaddress] {vr vr_name}
```

Description

Configures the RADIUS accounting server.

Syntax Description

mgmt-access	Specifies the RADIUS accounting server for switch management.
netlogin	Specifies the RADIUS accounting server for network login.
primary	Configure the primary RADIUS accounting server.
secondary	Configure the secondary RADIUS accounting server.
<i>ipaddress</i>	The IP address of the accounting server being configured.
<i>hostname</i>	The host name of the accounting server being configured.
<i>tcp_port</i>	The UDP port to use to contact the RADIUS accounting server.
<i>ipaddress</i>	The IP address used by the switch to identify itself when communicating with the RADIUS accounting server.
<i>vr_name</i>	Specifies the virtual router on which the client IP is located. NOTE: User-created VRs are supported only on the platforms listed for this feature in Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i>



Default

The following lists the default behavior of this command:

- The UDP port setting is 1813.
- The virtual router used is VR-Mgmt, the management virtual router.
- Switch management and network login use the same RADIUS accounting server.

Usage Guidelines

Use this command to specify the radius accounting server.

The accounting server and the RADIUS authentication server can be the same.

Use of the *hostname* parameter requires that DNS be enabled.

Beginning with ExtremeXOS 11.2, you can specify one pair of RADIUS accounting servers for switch management and another pair for network login. To specify RADIUS accounting servers for switch management (Telnet, SSH, and console sessions), use the *mgmt-access* keyword. To specify RADIUS accounting servers for network login, use the *netlogin* keyword. If you do not specify a keyword, switch management and network login use the same pair of RADIUS accounting servers.

If you are running ExtremeXOS 11.1 or earlier and upgrade to ExtremeXOS 11.2, you do not lose your existing RADIUS accounting server configuration. Both switch management and network login use the RADIUS accounting server specified in the older configuration.

Example

The following command configures RADIUS accounting on host *radius1* using the default UDP port (1813) for use by the RADIUS client on switch 10.10.20.30 using a virtual router interface of VR-Default for both management and network login:

```
configure radius-accounting primary server radius1 client-ip 10.10.20.30 vr
vr-Default
```

The following command configures RADIUS accounting for network login on host *netlog1* using the default UDP port for use by the RADIUS client on switch 10.10.20.31 using the default virtual router interface:

```
configure radius-accounting netlogin primary server netlog1 client-ip
10.10.20.31
```

History

This command was first available in ExtremeXOS 10.1.

The *mgmt-access* and *netlogin* keywords were added in ExtremeXOS 11.2.



Platform Availability

This command is available on all platforms.

configure radius-accounting shared-secret

```
configure radius-accounting {mgmt-access | netlogin} [primary | secondary]
shared-secret {encrypted} string
```

Description

Configures the authentication string used to communicate with the RADIUS accounting server.

Syntax Description

mgmt-access	Specifies the switch management RADIUS accounting server.
netlogin	Specifies the network login RADIUS accounting server.
primary	Configures the authentication string for the primary RADIUS accounting server.
secondary	Configures the authentication string for the secondary RADIUS accounting server.
encrypted	Indicates that the string is already encrypted.
<i>string</i>	The string to be used for authentication.

Default

Unconfigured.

Usage Guidelines

The secret must be the same between the client switch and the RADIUS accounting server.

The `mgmt-access` keyword specifies the RADIUS accounting server used for switch management.

The `netlogin` keyword specifies the RADIUS accounting server used for network login.

If you do not specify the `mgmt-access` or `netlogin` keywords, the secret applies to both the primary or secondary switch management and `netlogin` RADIUS accounting servers.

The `encrypted` keyword is primarily for the output of the `show configuration` command, so the shared secret is not revealed in the command output. Do not use it to set the shared secret.



Example

The following command configures the shared secret as “purpleaccount” on the primary RADIUS accounting server for both management and network login:

```
configure radius primary shared-secret purpleaccount
```

The following command configures the shared secret as “greenaccount” on the primary management RADIUS accounting server:

```
configure radius mgmt-access primary shared-secret greenaccount
```

History

This command was first available in ExtremeXOS 10.1.

The encrypted keyword was added in ExtremeXOS 11.0.

The mgmt-access and netlogin keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure radius-accounting timeout

```
configure radius-accounting {mgmt-access | netlogin} timeout seconds
```

Description

Configures the timeout interval for RADIUS-Accounting authentication requests.

Syntax Description

mgmt-access	Specifies the switch management RADIUS accounting server.
netlogin	Specifies the network login RADIUS accounting server.
<i>seconds</i>	Specifies the number of seconds for accounting requests. Range is 3 to 120 seconds.

Default

The default is 3 seconds.



Usage Guidelines

This command configures the timeout interval for RADIUS-Accounting authentication requests. When the timeout has expired, another authentication attempt will be made. After three failed attempts to authenticate, the alternate server will be used.

The `mgmt-access` keyword specifies the RADIUS accounting server used for switch management.

The `netlogin` keyword specifies the RADIUS accounting server used for network login.

If you do not specify the `mgmt-access` or `netlogin` keywords, the timeout interval applies to both switch management and netlogin RADIUS accounting servers.

Example

This example configures the timeout interval for RADIUS-Accounting authentication to 10 seconds. After 30 seconds (three attempts), the alternate RADIUS server will be used:

```
configure radius-accounting timeout 10
```

History

This command was first available in ExtremeXOS 10.1.

The `mgmt-access` and `netlogin` keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure ssh2 key

```
configure ssh2 key {pregenerated}
```

Description

Generates the Secure Shell 2 (SSH2) host key.

Syntax Description

pregenerated	Indicates that the SSH2 authentication key has already been generated. The user will be prompted to enter the existing key.
---------------------	---

Default

The switch generates a key for each SSH2 session.



Usage Guidelines

Secure Shell 2 (SSH2) is a feature of ExtremeXOS that allows you to encrypt session data between a network administrator using SSH2 client software and the switch or to send encrypted data from the switch to an SSH2 client on a remote system. Configuration, policy, image, and public key files may also be transferred to the switch using the Secure Copy Program 2 (SCP2).

SSH2 functionality is not present in the base ExtremeXOS software image, but is available as an additional, installable module. Before you can access any SSH2 commands, you must install the module. Without the module, the SSH2 commands do not appear on the command line. To install the module, see the instructions in [Software Upgrade and Boot Options](#).

After you have installed the SSH2 module, you must generate a host key and enable SSH2. To generate an SSH2 host key, use the `configure ssh2 key` command. To enable SSH2, use the `enable ssh2` command.

An authentication key must be generated before the switch can accept incoming SSH2 sessions. This can be done automatically by the switch, or you can enter a previously generated key.

If you elect to have the key generated, the key generation process can take up to ten minutes, and cannot be canceled after it has started. Once the key has been generated, you should save your configuration to preserve the key.

To use a key that has been previously created, use the pregenerated keyword. Use the `show ssh2 private-key` command to list and copy the previously generated key. Then use the `configure ssh2 key {pregenerated}` command where “pregenerated” represents the key that you paste.

Note



Keys generated by ExtremeXOS cannot be used on switches running ExtremeWare images, and keys generated by ExtremeWare cannot be used on switches running ExtremeXOS images.

The key generation process generates the SSH2 private host key. The SSH2 public host key is derived from the private host key, and is automatically transmitted to the SSH2 client at the beginning of an SSH2 session.

To view the status of SSH2 on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for SSH2 sessions, whether a valid key is present, and the TCP port and virtual router that is being used.

Example

The following command generates an authentication key for the SSH2 session:

```
configure ssh2 key
```

The command responds with the following messages:

```
WARNING: Generating new server host key This will take approximately 10
minutes and cannot be canceled. Continue? (y/n)
```



If you respond yes, the command begins the process.

To configure an SSH2 session using a previously generated key, use the following command:

```
configure ssh2 key pregenerated <pre-generated key>
```

Enter the previously-generated key (you can copy and paste it from the saved configuration file; a part of the key pattern is similar to 2d:2d:2d:2d:20:42:45:47:).

History

This command was first available in the ExtremeXOS 11.0 SSH module.

Platform Availability

This command is available on all platforms.

configure sshd2 user-key add user

```
configure sshd2 user-key key_name add user user_name
```

Description

Associates a user to a key.

Syntax Description

<i>key_name</i>	Specifies the name of the public key.
<i>user_name</i>	Specifies the name of the user.

Default

N/A.

Usage Guidelines

This command associates (or binds) a user to a key.

Example

The following example binds the key id_dsa_2048 to user admin:

```
configure sshd2 user-key id_dsa_2048 add user admin
```



History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

configure sshd2 user-key delete user

```
configure sshd2 user-key key_name delete user user_name
```

Description

Disassociates a user to a key.

Syntax Description

<i>key_name</i>	Specifies the name of the public key.
<i>user_name</i>	Specifies the name of the user.

Default

N/A.

Usage Guidelines

This command disassociates (or unbinds) a user to a key.

Example

The following example unbinds the key `id_dsa_2048` from user `admin`:

```
configure sshd2 user-key id_dsa_2048 delete user admin
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.



configure ssl certificate pregenerated

`configure ssl certificate pregenerated`

Description

Obtains the pre-generated certificate from the user.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

You must upload or generate a certificate for SSL server use. With this command, you copy and paste the certificate into the command line followed by a blank line to end the command. The following security algorithms are supported:

- RSA for public key cryptography (generation of certificate and public-private key pair, certificate signing). RSA key size between 1024 and 4096 bits.
- Symmetric ciphers (for data encryption): RC4, DES, and 3DES.
- Message Authentication Code (MAC) algorithms: RSA Data Security, Inc. MD5 Message-Digest Algorithm and SHA.

This command is also used when downloading or uploading the configuration. Do not modify the certificate stored in the uploaded configuration file because the certificate is signed using the issuer's private key.

The certificate and private key file should be in PEM format and generated using RSA as the cryptography algorithm.

Similar to SSH2, before you can use any SSL commands, you must first download and install the separate Extreme Networks SSH software module (ssh.xmod). This additional module allows you to configure both SSH2 and SSL on the switch. SSL is packaged with the SSH module; therefore, if you do not install the module, you are unable to configure SSL. If you try to execute SSL commands without installing the module first, the switch notifies you to download and install the module. To install the module, see the instructions in *Software Upgrade and Boot Options* of the *ExtremeXOS Concepts Guide*.

The Converged Network Analyzer (CNA) Agent requires SSL to encrypt communication between the CNA Agent and the CNA Server. For more information about the CNA Agent, see [CNA Agent](#).



Example

The following command obtains the pre-generated certificate from the user:

```
configure ssl certificate pregenerated
```

Next, you open the certificate and then copy and paste the certificate into the console/Telnet session, followed by a blank line to end the command.

History

This command was first available in the ExtremeXOS 11.2 and supported with the SSH module.

Platform Availability

This command is available on all platforms.

configure ssl certificate privkeylen

```
configure ssl certificate privkeylen length country code organization org_name  
common-name name
```

Description

Creates a self signed certificate and private key that can be saved in the EEPROM.

Syntax Description

<i>length</i>	Specifies the private key length in bytes. Valid values are between 1024 and 4096.
<i>code</i>	Specifies the country code in 2-character form.
<i>org_name</i>	Specifies the organization name. The organization name can be up to 64 characters long.
<i>name</i>	Specifies the common name. The common name can be up to 64 characters long.

Default

N/A.

Usage Guidelines

This command creates a self signed certificate and private key that can be saved in the EEPROM. The certificate generated is in the PEM format.

Any existing certificate and private key is overwritten.

The size of the certificate depends on the RSA key length (privkeylen) and the length of the other parameters (country, organization name, and so forth) supplied by the user. If the RSA key length is



1024, then the certificate is approximately 1 kb. For an RSA key length of 4096, the certificate length is approximately 2 kb, and the private key length is approximately 3 kb.

Similar to SSH2, before you can use any SSL commands, you must first download and install the separate Extreme Networks SSH software module (ssh.xmod). This additional module allows you to configure both SSH2 and SSL on the switch. SSL is packaged with the SSH module; therefore, if you do not install the module, you are unable to configure SSL. If you try to execute SSL commands without installing the module first, the switch notifies you to download and install the module. To install the module, see the instructions in [Software Upgrade and Boot Options](#) of the *ExtremeXOS Concepts Guide*.

The CNA Agent requires SSL to encrypt communication between the CNA Agent and the CNA Server. For more information about the CNA Agent, see [CNA Agent Commands](#).

Example

The following command creates an SSL certificate in the USA for a website called bigcats:

```
configure ssl certificate privkeylen 2048 country US organization IEEE common-  
name bigcats
```

History

This command was first available in the ExtremeXOS 11.2 and supported with the SSH module.

Platform Availability

This command is available on all platforms.

configure ssl privkey pregenerated

```
configure ssl privkey pregenerated
```

Description

Obtains the pre-generated private key from the user.

Syntax Description

This command has no arguments or variables.

Default

N/A.



Usage Guidelines

This command is also used when downloading or uploading the configuration. The private key is stored in the EEPROM, and the certificate is stored in the configuration file.

With this command, you copy and paste the private key into the command line followed by a blank line to end the command. The following security algorithms are supported:

- RSA for public key cryptography (generation of certificate and public-private key pair, certificate signing). RSA key size between 1024 and 4096 bits.
- Symmetric ciphers (for data encryption): RC4, DES, and 3DES.
- Message Authentication Code (MAC) algorithms: RSA Data Security, Inc. MD5 Message-Digest Algorithm and SHA.

The certificate and private key file should be in PEM format and generated using RSA as the cryptography algorithm.

Similar to SSH2, before you can use any SSL commands, you must first download and install the separate Extreme Networks SSH software module (ssh.xmod). This additional module allows you to configure both SSH2 and SSL on the switch. SSL is packaged with the SSH module; therefore, if you do not install the module, you are unable to configure SSL. If you try to execute SSL commands without installing the module first, the switch notifies you to download and install the module. To install the module, see the instructions in [Software Upgrade and Boot Options](#) of the *ExtremeXOS Concepts Guide*.

The CNA Agent requires SSL to encrypt communication between the CNA Agent and the CNA Server. For more information about the CNA Agent, see [CNA Agent Commands](#).

Example

The following command obtains the pre-generated private key from the user:

```
configure ssl privkey pregenerated
```

Next, you open the certificate and then copy and paste the certificate into the console/Telnet session, followed by a RETURN to end the command.

History

This command was first available in the ExtremeXOS 11.2 and supported with the SSH module.

Platform Availability

This command is available on all platforms.

configure tacacs server client-ip

```
configure tacacs [primary | secondary] server [ipaddress | hostname] {tcp_port}  
client-ip ipaddress {vr vr_name}
```



Description

Configures the server information for a TACACS+ authentication server.

Syntax Description

primary	Configures the primary TACACS+ server.
secondary	Configures the secondary TACACS+ server.
<i>ipaddress</i>	The IP address of the TACACS+ server being configured.
<i>hostname</i>	The host name of the TACACS+ server being configured.
<i>tcp_port</i>	The TCP port to use to contact the TACACS+ server.
<i>ipaddress</i>	The IP address used by the switch to identify itself when communicating with the TACACS+ server.
<i>vr_name</i>	Specifies the virtual router on which the client IP is located. NOTE: User-created VRs are supported only on the platforms listed for this feature in Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i> .

Default

TACACS+ uses TCP port 49. The default virtual router is VR-Mgmt, the management virtual router.

Usage Guidelines

Use this command to configure the server information for a TACACS+ server.

To remove a server, use the following command:

```
unconfigure tacacs server [primary | secondary]
```

Use of the *hostname* parameter requires that DNS be enabled.

Example

The following command configures server tacacs1 as the primary TACACS+ server for client switch 10.10.20.35 using a virtual router interface of VR-Default:

```
configure tacacs primary server tacacs1 client-ip 10.10.20.35 vr vr-Default
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.



configure tacacs shared-secret

```
configure tacacs [primary | secondary] shared-secret {encrypted} string
```

Description

Configures the shared secret string used to communicate with the TACACS+ authentication server.

Syntax Description

primary	Configures the authentication string for the primary TACACS+ server.
secondary	Configures the authentication string for the secondary TACACS+ server.
encrypted	Indicates that the string is already encrypted.
<i>string</i>	The string to be used for authentication.

Default

N/A.

Usage Guidelines

The secret must be the same between the client switch and the TACACS+ server.

The encrypted keyword is primarily for the output of the show configuration command, so the shared secret is not revealed in the command output. Do not use it to set the shared secret.

Example

The following command configures the shared secret as “purplegreen” on the primary TACACS+ server:

```
configure tacacs-accounting primary shared-secret purplegreen
```

History

This command was first available in ExtremeXOS 10.1.

The encrypted keyword was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.



configure tacacs timeout

```
configure tacacs timeout seconds
```

Description

Configures the timeout interval for TACAS+ authentication requests.

Syntax Description

<i>seconds</i>	Specifies the number of seconds for authentication requests. Range is 3 to 120 seconds.
----------------	---

Default

The default is 3 seconds.

Usage Guidelines

Use this command to configure the timeout interval for TACACS+ authentication requests.

To detect and recover from a TACACS+ server failure when the timeout has expired, the switch makes one authentication attempt before trying the next designated TACACS+ server or reverting to the local database for authentication. In the event that the switch still has IP connectivity to the TACACS+ server, but a TCP session cannot be established, (such as a failed TACACS+ daemon on the server), failover happens immediately regardless of the configured timeout value.

For example, if the timeout value is set for 3 seconds (the default value), it will take 3 seconds to fail over from the primary TACACS+ server to the secondary TACACS+ server. If both the primary and the secondary servers fail or are unavailable, it takes approximately 6 seconds to revert to the local database for authentication.

Example

The following command configures the timeout interval for TACACS+ authentication to 10 seconds:

```
configure tacacs timeout 10
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.



configure tacacs-accounting server

```
configure tacacs-accounting [primary | secondary] server [ipaddress | hostname]
{udp_port} client-ip ipaddress {vr vr_name}
```

Description

Configures the TACACS+ accounting server.

Syntax Description

primary	Configures the primary TACACS+ accounting server.
secondary	Configures the secondary TACACS+ accounting server.
<i>ipaddress</i>	The IP address of the TACACS+ accounting server being configured.
<i>hostname</i>	The host name of the TACACS+ accounting server being configured.
tcp_port	The TCP port to use to contact the TACACS+ server.
<i>ipaddress</i>	The IP address used by the switch to identify itself when communicating with the TACACS+ accounting server.
<i>vr_name</i>	Specifies the virtual router on which the client IP is located. NOTE: User-created VRs are supported only on the platforms listed for this feature in Feature License Requirements the <i>ExtremeXOS Concepts Guide</i> .

Default

Unconfigured. The default virtual router is VR-Mgmt, the management virtual router.

Usage Guidelines

You can use the same TACACS+ server for accounting and authentication.

To remove a server, use the following command:

```
unconfigure tacacs server [primary | secondary]
```

Example

The following command configures server tacacs1 as the primary TACACS+ accounting server for client switch 10.10.20.35 using a virtual router interface of VR-Default:

```
configure tacacs-accounting primary server tacacs1 client-ip 10.10.20.35 vr
vr-Default
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure tacacs-accounting shared-secret

```
configure tacacs-accounting [primary | secondary] shared-secret {encrypted}
string
```

Description

Configures the shared secret string used to communicate with the TACACS+ accounting server.

Syntax Description

primary	Configures the authentication string for the primary TACACS+ accounting server.
secondary	Configures the authentication string for the secondary TACACS+ accounting server.
<i>string</i>	The string to be used for authentication.

Default

N/A.

Usage Guidelines

Secret needs to be the same as on the TACACS+ server.

The encrypted keyword is primarily for the output of the show configuration command, so the shared secret is not revealed in the command output. Do not use it to set the shared secret.

Example

The following command configures the shared secret as “tacacsaccount” on the primary TACACS+ accounting server:

```
configure tacacs-accounting primary shared-secret tacacsaccount
```

History

This command was first available in ExtremeXOS 10.1.



The encrypted keyword was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure tacacs-accounting timeout

```
configure tacacs-accounting timeout seconds
```

Description

Configures the timeout interval for TACACS+ accounting authentication requests.

Syntax Description

<i>seconds</i>	Specifies the number of seconds for accounting requests. Range is 3 to 120 seconds.
----------------	---

Default

The default is 3 seconds.

Usage Guidelines

This command configures the timeout interval for TACACS+ accounting authentication requests.

To detect and recover from a TACACS+ accounting server failure when the timeout has expired, the switch makes one authentication attempt before trying the next designated TACACS+ accounting server or reverting to the local database for authentication. In the event that the switch still has IP connectivity to the TACACS+ accounting server, but a TCP session cannot be established, (such as a failed TACACS+ daemon on the accounting server), failover happens immediately regardless of the configured timeout value.

For example, if the timeout value is set for 3 seconds (the default value), it takes 3 seconds to fail over from the primary TACACS+ accounting server to the secondary TACACS+ accounting server. If both the primary and the secondary servers fail or are unavailable, it takes approximately 6 seconds to revert to the local database for authentication.

Example

The following command configures the timeout interval for TACACS+ accounting authentication to 10 seconds:

```
configure tacacs-accounting timeout 10
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure trusted-ports trust-for dhcp-server

```
configure trusted-ports [ports|all] trust-for dhcp-server
```

Description

Configures one or more trusted DHCP ports.

Syntax Description

<i>ports</i>	Specifies one or more ports to be configured as trusted ports.
all	Specifies all ports to be configured as trusted ports.

Default

N/A.

Usage Guidelines

To configure trusted DHCP ports, you must first enable DHCP snooping on the switch. To enable DHCP snooping, use the following command:

```
enable ip-security dhcp-snooping {vlan} vlan_name ports [all |ports] violation-  
action [drop-packet {[block-mac | block-port] [durationduration_in_seconds |  
permanently] | none}] [snmp-trap]
```

Trusted ports do not block traffic; rather, the switch forwards any DHCP server packets that appear on trusted ports. Depending on your DHCP snooping configuration, the switch drops packets and can disable the port temporarily, disable the port permanently, blackhole the MAC address temporarily, blackhole the MAC address permanently, and so on.

If you configure one or more trusted ports, the switch assumes that all DHCP server packets on the trusted port are valid.

Displaying DHCP Trusted Server Information

To display the DHCP snooping configuration settings, including DHCP trusted ports if configured, use the following command: `show ip-security dhcp-snooping {vlan} vlan_name`

To display any violations that occur, including those on DHCP trusted ports if configured, use the following command: `show ip-security dhcp-snooping violations {vlan} vlan_name`



Example

The following command configures ports 2:2 and 2:3 as trusted ports:

```
configure trusted-ports 2:2-2:3 trust-for dhcp-server
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

configure trusted-servers add server

```
configure trusted-servers {vlan} vlan_name add server ip_address trust-for dhcp-server
```

Description

Configures and enables a trusted DHCP server on the switch.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN name.
<i>ip_address</i>	Specifies the IP address of the trusted DHCP server.

Default

N/A.

Usage Guidelines

If you configured trusted DHCP server, the switch forwards only DHCP packets from the trusted servers. The switch drops DHCP packets from other DHCP snooping-enabled ports.

You can configure a maximum of eight trusted DHCP servers on the switch.

If you configure a port as a trusted port, the switch assumes that all DHCP server packets on that port are valid.

Displaying DHCP Trusted Server Information

To display the DHCP snooping configuration settings, including DHCP trusted servers if configured, use the following command:



```
show ip-security dhcp-snooping {vlan} vlan_name
```

To display any violations that occur, including those on the DHCP trusted servers if configured, use the following command:

```
show ip-security dhcp-snooping violations {vlan} vlan_name
```

Example

The following command configures a trusted DHCP server on the switch:

```
configure trusted-servers vlan purple add server 10.10.10.10 trust-for dhcp-server
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

configure trusted-servers delete server

```
configure trusted-servers vlan vlan_name delete server ip_address trust-for dhcp-server
```

Description

Deletes a trusted DHCP server from the switch.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN name.
<i>ip_address</i>	Specifies the IP address of the trusted DHCP server.

Default

N/A.

Usage Guidelines

Use this command to delete a trusted DHCP server from the switch.



Displaying DHCP Trusted Server Information

To display the DHCP snooping configuration settings, including DHCP trusted servers if configured, use the following command:

```
show ip-security dhcp-snooping {vlan} vlan_name
```

To display any violations that occur, including those on the DHCP trusted servers if configured, use the following command:

```
show ip-security dhcp-snooping violations {vlan} vlan_name
```

Example

The following command deletes a trusted DHCP server from the switch:

```
configure trusted-servers vlan purple delete server 10.10.10.10 trust-for
dhcp-server
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

configure vlan dhcp-address-range

```
configure vlan vlan_name dhcp-address-range ipaddress1 - ipaddress2
```

Description

Configures a set of DHCP addresses for a VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN on whose ports DHCP will be enabled.
<i>ipaddress1</i>	Specifies the first IP address in the DHCP address range to be assigned to this VLAN.
<i>ipaddress2</i>	Specifies the last IP address in the DHCP address range to be assigned to this VLAN.

Default

N/A.



Usage Guidelines

The following error conditions are checked: `ipaddress2 >= ipaddress1`, the range must be in the VLAN's network, the range does not contain the VLAN's IP address, and the VLAN has an IP address assigned.

Example

The following command allocates the IP addresses between 192.168.0.20 and 192.168.0.100 for use by the VLAN temporary:

```
configure temporary dhcp-address-range 192.168.0.20 - 192.168.0.100
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure vlan dhcp-lease-timer

```
configure vlan vlan_name dhcp-lease-timer lease-timer
```

Description

Configures the timer value in seconds returned as part of the DHCP response.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN on whose ports netlogin should be disabled.
<i>lease-timer</i>	Specifies the timer value, in seconds.

Default

N/A.

Usage Guidelines

The timer value is specified in seconds. The timer value range is 0 - 4294967295, where 0 indicates the default (not configured) value of 7200 second.



Example

The following command configures the DHCP lease timer value for VLAN corp:

```
configure vlan corp dhcp-lease-timer <lease-timer>
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure vlan dhcp-options

```
configure {vlan} vlan_name dhcp-options [code option_number [16-bit value1
{value2 {value3 {value4}}} | 32-bit value1 {value2 {value3 {value4}}} | flag [on
| off] | hex string_value | ipaddress ipaddress1 {ipaddress2 {ipaddress3
{ipaddress4}}} | string string_value] | default-gateway | dns-server {primary |
secondary} | wins-server] ipaddress
```

Description

Configures the DHCP options returned as part of the DHCP response by a switch configured as a DHCP server.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN on which to configure DHCP.
code	Specifies the generic DHCP option code.
<i>option_number</i>	Specifies the DHCP Option number.
16-bit	Specifies that one to four 16-bit unsigned integer values associated with selected DHCP option.
32-bit	Specifies that one to four 32-bit unsigned integer values associated with selected DHCP option.
flag	Specifies that 1 byte value associated with selected DHCP option number.
hex	Specifies that hexadecimal string associated with selected DHCP option number.
string	Specifies that a string is associated with selected DHCP option number.
<i>string_value</i>	The string value associated with specified option.
default-gateway	Specifies the router option.
dns-server	Specifies the Domain Name Server (DNS) option.



primary	Specifies the primary DNS option.
secondary	Specifies the secondary DNS option.
wins-server	Specifies the NetBIOS name server (NBNS) option.
<i>ipaddress</i>	The IP address associated with the specified option.

Default

N/A.

Usage Guidelines

This command configures the DHCP options that can be returned to the DHCP client. For the default-gateway option you are only allowed to configure an IP address that is in the VLAN's network range. For the other options, any IP address is allowed.

The options below represent the following BOOTP options specified by RFC2132:

- default-gateway—Router option, number 3.
- dns-server—Domain Name Server option, number 6.
- wins-server—NetBIOS over TCP/IP Name Server option, number 44.

Example

The following command configures the DHCP server to return the IP address 10.10.20.8 as the router option:

```
configure vlan <name> dhcp-options default-gateway 10.10.20.8
```

History

This command was first available in ExtremeXOS 11.0.

The primary and secondary DNS options were added in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

create sshd2 key-file

```
create sshd2 key-file {host-key | user-key} key_name
```

Description

Creates a file for the user-key or host-key.



Syntax Description

host-key	Specifies the name of the host-key.
user-key	Specifies the name of the user-key.
<i>key_name</i>	Specifies the name of the public key.

Default

N/A.

Usage Guidelines

This command is used to write the user or the host public key in a file. The key files will be created with a .ssh file extension; this enables the administrator to copy the public key files to another server.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

create sshd2 user-key

```
create sshd2 user-key key_name key {subject subject} {comment comment}
```

Description

Creates a user key.

Syntax Description

<i>key_name</i>	Specifies the name of the public key.
<i>key</i>	Specifies the key. Note: The key cannot have any spaces in it.
<i>subject</i>	Specifies the subject.
<i>comment</i>	Specifies the comment (an optional field).

Default

N/A.



Usage Guidelines

This command is used to enter, or cut and paste, your public key. You can also enter the public key into the switch by using the SCP or SFTP client that is connected to the switch.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

delete sshd2 user-key

```
delete sshd2 user-key key_name
```

Description

Deletes a user key.

Syntax Description

<i>key_name</i>	Specifies the name of the public key to be deleted.
-----------------	---

Default

N/A.

Usage Guidelines

This command is used to delete a user key. The key is deleted regardless of whether or not it is bound to a user.



Note

If a user is bound to the key, they are first unbound or unassociated, and then the key is deleted.

Example

The following example shows the SSH user key `id_dsa_2048` being deleted:

```
delete sshd2 user-key id_dsa_2048
```



History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

disable dhcp ports vlan

```
disable dhcp ports port_list vlan vlan_name
```

Description

Disables DHCP on a specified port in a VLAN.

Syntax Description

<i>port_list</i>	Specifies the ports for which DHCP should be disabled.
<i>vlan_name</i>	Specifies the VLAN on whose ports DHCP should be disabled.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables DHCP for port 6:9 in VLAN corp:

```
disable dhcp ports 6:9 vlan corp
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.



disable dos-protect

disable dos-protect

Description

Disables denial of service protection.

Syntax Description

There are no arguments or variables for this command.

Default

Default is disabled.

Usage Guidelines

None.

Example

The following command disables denial of service protection:

```
disable dos-protect
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

disable iparp gratuitous protect vlan

disable iparp gratuitous protect vlan *vlan-name*

Description

Disables gratuitous ARP protection on the specified VLAN.

Syntax Description

<i>vlan-name</i>	Specifies the VLAN.
------------------	---------------------



Default

Disabled.

Usage Guidelines

Hosts can launch man-in-the-middle attacks by sending out gratuitous ARP requests for the router's IP address. This results in hosts sending their router traffic to the attacker, and the attacker forwarding that data to the router. This allows passwords, keys, and other information to be intercepted.

To protect against this type of attack, the router will send out its own gratuitous ARP request to override the attacker whenever a gratuitous ARP broadcast with the router's IP address as the source is received on the network.

This command disables gratuitous ARP protection.

Example

The following command disables gratuitous ARP protection for VLAN corp:

```
disable iparp gratuitous protect vlan corp
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

disable ip-security anomaly-protection

```
disable ip-security anomaly-protection {slot [ slot | all ]}
```

Description

Disables all anomaly checking options.

Syntax Description

<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.



Default

The default is disabled.

Usage Guidelines

This command disables all anomaly checking options, including IP address, UDP/TCP port, TCP flag and fragment, and ICMP anomaly checking.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X440, X460, X480, X670, and X770 platforms, whether or not included in a SummitStack, and the BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules.

disable ip-security anomaly-protection ip

```
disable ip-security anomaly-protection ip { slot [ slot | all ] }
```

Description

Disables source and destination IP address checking.

Syntax Description

<i>slot</i>	Specifies the slot.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command disables source and destination IP addresses checking. This checking takes effect for both IPv4 and IPv6 packets. When enabled, the switch drops IPv4/IPv6 packets if its source IP address are the same as the destination IP address. In most cases, the condition of source IP address being the same as the destination IP address indicates a Layer 3 protocol error. (These kind of errors are found in LAND attacks.)

History

This command was first available in ExtremeXOS 12.0.



Platform Availability

This command is only available on the Summit X440, X460, X480, X670, and X770 platforms, whether or not included in a SummitStack, and the BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules.

disable ip-security anomaly-protection l4port

```
disable ip-security anomaly-protection l4port [tcp | udp | both] {slot [slot | all ]}
```

Description

Disables TCP and UDP ports checking.

Syntax Description

tcp	Specifies that the TCP port be disabled for checking.
udp	Specifies that the UDP port be disabled for checking.
both	Specifies both the TCP and UDP ports be disabled for checking.
<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command disables TCP and UDP ports checking. This checking takes effect for both IPv4 and IPv6 TCP and UDP packets. When enabled, the switch drops TCP and UDP packets if its source port is the same as its destination port. In most cases, when the condition of source port is the same as that of the destination port, it indicates a Layer4 protocol error. (This type of error can be found in a BALT attack.)

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X440, X460, X480, X670, and X770 platforms, whether or not included in a SummitStack, and the BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules.



disable ip-security anomaly-protection tcp flags

```
disable ip-security anomaly-protection tcp flags {slot [ slot | all ]}
```

Description

Disables TCP flag checking.

Syntax Description

<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command disables TCP flag checking. This checking takes effect for both IPv4 and IPv6 TCP packets. When enabled, the switch drops TCP packets if one of following condition is true:

- TCP SYN flag==1 and the source port<1024
- TCP control flag==0 and the sequence number==0
- TCP FIN, URG, and PSH bits are set, and the sequence number==0
- TCP SYN and FIN both are set.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X440, X460, X480, and X670 platforms, whether or not included in a SummitStack, and the BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules.

disable ip-security anomaly-protection tcp fragment

```
disable ip-security anomaly-protection tcp fragment {slot [ slot | all ]}
```

Description

Disables TCP fragment checking.



Syntax Description

<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command disables TCP fragment checking. This checking takes effect for IPv4/IPv6. When it is enabled, the switch drops TCP packets if one of following condition is true:

- For the first IPv4 TCP fragment (its IP offset field==0), if its TCP header is less than the minimum IPv4 TCP header allowed size.
- If its IP offset field==1 (for IPv4 only).

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X440, X460, X480, and X670 platforms, whether or not included in a SummitStack, and the BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules.

disable ip-security anomaly-protection icmp

```
disable ip-security anomaly-protection icmp {slot [ slot | all ]}
```

Description

Disables ICMP size and fragment checking.

Syntax Description

<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.



Usage Guidelines

This command disables ICMP size and fragment checking. This checking takes effect for both IPv4 and IPv6 TCP packets. When enabled, the switch drops ICMP packets if one of following condition is true:

- Fragmented ICMP packets for IPv4 packets.
- IPv4 ICMP pings packets with payload size greater than the maximum IPv4 ICMP-allowed size. (The maximum allowed size is configurable.)
- IPv6 ICMP ping packets with payload size > the maximum IPv6 ICMP-allowed size. (The maximum allowed size is configurable.)

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X440, X460, X480, and X670 platforms, whether or not included in a SummitStack, and the BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules.

disable ip-security anomaly-protection notify

```
disable ip-security anomaly-protection notify [log | snmp | cache] {slot [ slot | all ]}
```

Description

Disables protocol anomaly notification.

Syntax Description

log	Specifies the switch to send the notification to a log file.
snmp	Specifies the switch to send an SNMP trap when an event occurs.
cache	Specifies the switch to send the notification to cache.
<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command disables anomaly notification. When enabled, any packet failed to pass enabled protocol checking is sent to XOS Host CPU and notifies the user. There are three different types of notifications:



- `log`: log anomaly events in the switch log system; you can view and manage this log with the `show log` and `configure log` commands.
- `snmp`: the anomaly events generate SNMP traps.
- `cache`: logs the most recent and unique anomaly events in memory; rebooting the switch will cause all the logged events to be lost (the number of cached events is configured by command).

When disabled, the switch drops all violating packets silently.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X440, X460, X480, and X670 platforms, whether or not included in a SummitStack, and the BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules.

disable ip-security arp gratuitous-protection

```
disable ip-security arp gratuitous-protection {vlan} [all | vlan_name]
```

Description

Disables gratuitous ARP protection on one or all VLANs on the switch.

Syntax Description

all	Specifies all VLANs configured on the switch.
<i>vlan_name</i>	Specifies the VLAN.

Default

By default, gratuitous ARP protection is disabled.

Usage Guidelines

Beginning with ExtremeXOS 11.6, this command replaces the `disable iparp gratuitous protect vlan` command.

Hosts can launch man-in-the-middle attacks by sending out gratuitous ARP requests for the router's IP address. This results in hosts sending their router traffic to the attacker, and the attacker forwarding that data to the router. This allows passwords, keys, and other information to be intercepted.

To protect against this type of attack, the router will send out its own gratuitous ARP request to override the attacker whenever a gratuitous ARP broadcast with the router's IP address as the source is received on the network.



This command disables gratuitous ARP protection.

Example

The following command disables gratuitous ARP protection for VLAN corp:

```
disable ip-security arp gratuitous-protection vlan corp
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

disable ip-security arp learning learn-from-arp

```
disable ip-security arp learning learn-from-arp {vlan} vlan_name ports [all | ports]
```

Description

Disables ARP learning on the specified VLAN and member ports.

Syntax Description

<i>vlan_name</i>	Specifies the name of the VLAN to which this rule applies.
all	Specifies all ingress ports.
<i>ports</i>	Specifies one or more ingress ports.

Default

By default, ARP learning is enabled.

Usage Guidelines

You can disable ARP learning so that the only entries in the ARP table are either manually added or those created by DHCP secured ARP; the switch does not add entries by tracking ARP requests and replies. By disabling ARP learning and adding a permanent entry or configuring DHCP secured ARP, you can centrally manage and allocate client IP addresses and prevent duplicate IP addresses from interrupting network operation.

To manually add a permanent entry to the ARP table, use the following command:



```
configure iparp add ip_addr {vrvr_name} mac
```

To configure DHCP secure ARP as a method to add entries to the ARP table, use the following command:

```
enable ip-security arp learning learn-from-dhcp vlan <vlan_name> ports [all | ports] {poll-interval interval_in_seconds} {retries number_of_retries}
```

Displaying ARP Information

To display how the switch builds an ARP table and learns MAC addresses for devices on a specific VLAN and associated member ports, use the following command:

```
show ip-security arp learning {vlan} vlan_name
```

To view the ARP table, including permanent and DHCP secured ARP entries, use the following command:

```
show iparp {ip_address | mac | vlan vlan_name | permanent} {vrvr_name}
```



Note

DHCP secured ARP entries are stored as static entries in the ARP table.

Example

The following command disables ARP learning on port 1:1 of the VLAN learn:

```
disable ip-security arp learning learn-from-arp vlan learn ports 1:1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

disable ip-security arp learning learn-from-dhcp

```
disable ip-security arp learning learn-from-dhcp {vlan} vlan_name ports [all | ports]
```

Description

Disables DHCP secured ARP learning for the specified VLAN and member ports.



Syntax Description

<i>vlan_name</i>	Specifies the name of the VLAN to which this rule applies.
all	Specifies all ingress ports.
<i>ports</i>	Specifies one or more ingress ports.

Default

By default, DHCP secured ARP learning is disabled.

Usage Guidelines

Use this command to disable DHCP secured ARP learning.

Displaying ARP Information

To display how the switch builds an ARP table and learns MAC addresses for devices on a specific VLAN and associated member ports, use the following command:

```
show ip-security arp learning {vlan} vlan_name
```

To view the ARP table, including permanent and DHCP secured ARP entries, use the following command:

```
show iparp {ip_address | mac | vlanvlan_name | permanent} {vrvr_name}
```

Example

The following command disables DHCP secured ARP learning on port 1:1 of the VLAN learn:

```
disable ip-security arp learning learn-from-dhcp vlan learn ports 1:1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

disable ip-security arp validation

```
disable ip-security arp validation {vlan} vlan_name [all | ports]
```



Description

Disables ARP validation for the specified VLAN and member ports.

Syntax Description

<i>vlan_name</i>	Specifies the name of the VLAN to which this rule applies.
all	Specifies all ports.
<i>ports</i>	Specifies one or more ports.

Default

By default, ARP validation is disabled.

Usage Guidelines

Use this command to disable ARP validation.

Displaying ARP Validation Information

To display information about ARP validation, use the following command:

```
show ip-security arp validation {vlan} vlan_name
```

Example

The following command disables ARP validation on port 1:1 of the VLAN valid:

```
disable ip-security arp validation vlan valid ports 1:1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

disable ip-security dhcp-bindings restoration

```
disable ip-security dhcp-bindings restoration
```

Description

Disables the download and upload of DHCP bindings.



Syntax

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

The command allows you to disable the download and upload of the DHCP bindings, essentially disabling the DHCP binding functionality. The default is disabled.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

disable ip-security dhcp-snooping

```
disable ip-security dhcp-snooping {vlan} vlan_name ports [all | ports]
```

Description

Disables DHCP snooping on the switch.

Syntax Description

<i>vlan_name</i>	Specifies the name of the DHCP-snooping VLAN.
all	Specifies all ports to stop receiving DHCP packets.
<i>ports</i>	Specifies one or more ports to stop receiving DHCP packets.

Default

By default, DHCP snooping is disabled.

Usage Guidelines

Use this command to disable DHCP snooping on the switch.



Example

The following command disables DHCP snooping on the switch:

```
disable ip-security dhcp-snooping vlan snoop ports 1:1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

disable ip-security source-ip-lockdown ports

```
disable ip-security source-ip-lockdown ports [all | ports]
```

Description

Disables the source IP lockdown feature on one or more ports.

Syntax Description

all	Specifies all ports for which source IP lockdown should be disabled.
<i>ports</i>	Specifies one or more ports for which source IP lockdown should be disabled.

Default

By default, source IP lockdown is disabled on the switch.

Usage Guidelines

To display the source IP lockdown configuration on the switch, use the following command:

```
show ip-security source-ip-lockdown
```

Example

The following command disables source IP lockdown on ports 1:1 and 1:4:

```
disable ip-security source-ip-lockdown ports 1:1, 1:4
```



History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

disable mac-lockdown-timeout ports

```
disable mac-lockdown-timeout ports [all | port_list]
```

Description

Disables the MAC address lock down timeout feature for the specified port or group of ports or for all ports on the switch.

Syntax Description

all	Specifies all ports.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

By default, the MAC address lock down feature is disabled.

Usage Guidelines

If you disable the MAC lock down timer on a port, existing MAC address entries for the port will time out based on the FDB aging period.

Example

The following command disables the MAC address lock down timer set for ports 2:3 and 2:4:

```
disable mac-lockdown-timeout ports 2:3, 2:4
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.



disable radius

```
disable radius {mgmt-access | netlogin}
```

Description

Disables the RADIUS client.

Syntax Description

mgmt-access	Specifies the switch management RADIUS authentication server.
netlogin	Specifies the network login RADIUS authentication server.

Default

RADIUS authentication is disabled for both switch management and network login by default.

Usage Guidelines

Use the `mgmt-access` keyword to disable RADIUS authentication for switch management functions.

Use the `netlogin` keyword to disable RADIUS authentication for network login.

If you do not specify a keyword, RADIUS authentication is disabled on the switch for both management and network login.

Example

The following command disables RADIUS authentication on the switch for both management and network login:

```
disable radius
```

The following command disables RADIUS authentication on the switch for network login:

```
disable radius netlogin
```

History

This command was first available in ExtremeXOS 10.1.

The `mgmt-access` and `netlogin` keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.



disable radius-accounting

```
disable radius-accounting {mgmt-access | netlogin}
```

Description

Disables RADIUS accounting.

Syntax Description

mgmt-access	Specifies the switch management RADIUS accounting server.
netlogin	Specifies the network login RADIUS accounting server.

Default

RADIUS accounting is disabled for both switch management and network login by default.

Usage Guidelines

Use the mgmt-access keyword to disable RADIUS accounting for switch management functions.

Use the netlogin keyword to disable RADIUS accounting for network login.

If you do not specify a keyword, RADIUS accounting is disabled on the switch for both management and network login.

Example

The following command disables RADIUS accounting on the switch for both management and network login:

```
disable radius-accounting
```

The following command disables RADIUS accounting on the switch for network login:

```
disable radius-accounting netlogin
```

History

This command was first available in ExtremeXOS 10.1.

The mgmt-access and netlogin keywords were added in ExtremeXOS 11.2.



Platform Availability

This command is available on all platforms.

disable ssh2

disable ssh2

Description

Disables the SSH2 server for incoming SSH2 sessions to switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

SSH2 options (non-default port setting) are not saved when SSH2 is disabled.

To view the status of SSH2 on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for SSH2.

Example

The following command disables the SSH2 server:

```
disable ssh2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable tacacs

disable tacacs



Description

Disables TACACS+ authentication.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables TACACS+ authentication for the switch:

```
disable tacacs
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable tacacs-accounting

```
disable tacacs-accounting
```

Description

Disables TACACS+ accounting.

Syntax Description

This command has no arguments or variables.

Default

N/A.



Usage Guidelines

None.

Example

The following command disables TACACS+ accounting:

```
disable tacacs-accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable tacacs-authorization

disable tacacs-authorization

Description

Disables TACACS+ authorization.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This disables CLI command authorization but leaves user authentication enabled.

Example

The following command disables TACACS+ CLI command authorization:

```
disable tacacs-authorization
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable web http

disable web http

Description

Disables the hypertext transfer protocol (HTTP) access to the switch on the default port (80).

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Use this command to disallow users from connecting with HTTP. Disabling HTTP access forces user to use a secured HTTPS connection if web HTTPS is enabled.

Use the following command to enable web HTTPS:

```
enable web https
```

Example

The following command disables HTTP on the default port:

```
disable web http
```

History

This command was first available in the ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.



disable web https

disable web https

Description

Disables the secure socket layer (SSL) access to the switch on the default port (443).

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Use this command to disable SSL before changing the certificate or private key.

Similar to SSH2, before you can use any SSL commands, you must first download and install the separate Extreme Networks SSH software module (ssh.xmod). This additional module allows you to configure both SSH2 and SSL on the switch. SSL is packaged with the SSH module; therefore, if you do not install the module, you are unable to configure SSL. If you try to execute SSL commands without installing the module first, the switch notifies you to download and install the module. To install the module, see the instructions in *Software Upgrade and Boot Options* of the *ExtremeXOS Concepts Guide*.

Example

The following command disables SSL on the default port:

```
disable web https
```

History

This command was first available in the ExtremeXOS 11.2 and supported with the SSH module.

Platform Availability

This command is available on all platforms.

download ssl certificate

download ssl ipaddress certificate cert_file



Description

Permits downloading of a certificate key from files stored in a TFTP server.

Syntax Description

<i>ipaddress</i>	Specifies the IP address of the TFTP server.
<i>cert_file</i>	Specifies the name of the certificate key.

Default

N/A.

Usage Guidelines

If the download operation is successful, any existing certificate is overwritten. After a successful download, the software attempts to match the public key in the certificate against the private key stored. If the private and public keys do not match, the switch displays a warning message similar to the following: Warning: The Private Key does not match with the Public Key in the certificate. This warning acts as a reminder to also download the private key.



Note

You can only download a certificate key in the VR-Mgmt virtual router.

Downloaded certificates and keys are not saved across switch reboots unless you save your current switch configuration. Once you issue the save command, the downloaded certificate is stored in the configuration file and the private key is stored in the EEPROM.

Similar to SSH2, before you can use any SSL commands, you must first download and install the separate Extreme Networks SSH software module (ssh.xmod). This additional module allows you to configure both SSH2 and SSL on the switch. SSL is packaged with the SSH module; therefore, if you do not install the module, you are unable to configure SSL. If you try to execute SSL commands without installing the module first, the switch notifies you to download and install the module. To install the module, see the instructions in [Software Upgrade and Boot Options](#) of the *ExtremeXOS Concepts Guide*.

You can purchase and obtain SSL certificates from Internet security vendors.

Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for remote IP addresses.

When specifying a remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z).
- Numerals (0-9).
- Period (.).
- Colon (:).



When configuring an IP address for your network server, remember the requirements listed above.

Remote Filename Character Restrictions

This section provides information about the characters supported by the switch for remote filenames.

When specifying a remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z).
- Numerals (0-9).
- Period (.).
- Dash (-).
- Underscore (_).
- Slash (/).

When naming a remote file, remember the requirements listed above.

Example

The following command downloads a certificate from a TFTP server with the IP address of 123.45.6.78:

```
download ssl 123.45.6.78 certificate g0ethner1
```

History

This command was first available in the ExtremeXOS 11.2 and supported with the SSH module.

Platform Availability

This command is available on all platforms.

download ssl privkey

```
download ssl ipaddress privkey key_file
```

Description

Permits downloading of a private key from files stored in a TFTP server.

Syntax Description

<i>ipaddress</i>	Specifies the IP address of the TFTP server.
<i>key_file</i>	Specifies the name of the private key file.



Default

N/A.

Usage Guidelines

If the operation is successful, the existing private key is overwritten.

After a successful download, a check is performed to find out whether the private key downloaded matches the public key stored in the certificate. If the private and public keys do not match, the switch displays a warning similar to the following: Warning: The Private Key does not match with the Public Key in the certificate. This warning acts as a reminder to also download the corresponding certificate.

The certificate and private key file should be in PEM format and generated using RSA as the cryptography algorithm.

Downloaded certificates and keys are not saved across switch reboots unless you save your current switch configuration. Once you issue the save command, the downloaded certificate is stored in the configuration file and the private key is stored in the EEPROM.

Similar to SSH2, before you can use any SSL commands, you must first download and install the separate Extreme Networks SSH software module (ssh.xmod). This additional module allows you to configure both SSH2 and SSL on the switch. SSL is packaged with the SSH module; therefore, if you do not install the module, you are unable to configure SSL. If you try to execute SSL commands without installing the module first, the switch notifies you to download and install the module. To install the module, see the instructions in *Software Upgrade and Boot Options of the ExtremeXOS Concepts Guide*.

Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for remote IP addresses.

When specifying a remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z).
- Numerals (0-9).
- Period (.).
- Colon (:).

When configuring an IP address for your network server, remember the requirements listed above.

Remote Filename Character Restrictions

This section provides information about the characters supported by the switch for remote filenames.

When specifying a remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z).
- Numerals (0-9).
- Period (.).
- Dash (-).



- Underscore (_).
- Slash (/).

When naming a remote file, remember the requirements listed above.

Example

The following command downloads a private key from a TFTP server with the IP address of 123.45.6.78:

```
download ssl 123.45.6.78 privkey t00Ts1e
```

History

This command was first available in the ExtremeXOS 11.2 and supported with the SSH module.

Platform Availability

This command is available on all platforms.

enable dhcp ports vlan

```
enable dhcp ports port_list vlan vlan_name
```

Description

Enables DHCP on a specified port in a VLAN.

Syntax Description

<i>port_list</i>	Specifies the ports for which DHCP should be enabled.
<i>vlan_name</i>	Specifies the VLAN on whose ports DHCP should be enabled.

Default

Disabled.

Usage Guidelines

None.



Example

The following command enables DHCP for port 5:9 in VLAN corp:

```
enable dhcp ports 5:9 vlan corp
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

enable dos-protect

enable dos-protect

Description

Enables denial of service protection.

Syntax Description

This command has no arguments or variables.

Default

The default is disabled.

Usage Guidelines

None.

Example

The following command enables denial of service protection.

```
enable dos-protect
```

History

This command was first available in ExtremeXOS 11.1.



Platform Availability

This command is available on all platforms.

enable dos-protect simulated

```
enable dos-protect simulated
```

Description

Enables simulated denial of service protection.

Syntax Description

This command has no arguments or variables.

Default

The default is disabled.

Usage Guidelines

If simulated denial of service is enabled, no ACLs are created. This mode is useful to gather information about normal traffic levels on the switch. This will assist in configuring denial of service protection so that legitimate traffic is not blocked.

Example

The following command enables simulated denial of service protection:

```
enable dos-protect simulated
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on all platforms.

enable iparp gratuitous protect

```
enable iparp gratuitous protect vlan vlan-name
```



Description

Syntax Description

<code>vlan-name</code>	Specifies the VLAN.
------------------------	---------------------

Default

By default, gratuitous ARP is disabled.

Usage Guidelines

Beginning with ExtremeXOS 11.6, the command replaces this command for configuring gratuitous ARP.

Hosts can launch man-in-the-middle attacks by sending out gratuitous ARP requests for the router's IP address. This results in hosts sending their router traffic to the attacker, and the attacker forwarding that data to the router. This allows passwords, keys, and other information to be intercepted.

To protect against this type of attack, the router will send out its own gratuitous ARP request to override the attacker whenever a gratuitous ARP broadcast with the router's IP address as the source is received on the network.

Example

The following command enables gratuitous ARP protection for VLAN corp:

```
enable iparp gratuitous protect vlan corp
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

enable ip-option loose-source-route

```
enable ip-option loose-source-route
```

Description

Enables processing of the loose source route IP option in the IPv4 packet header.



Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This enables the switch to forward IP packets that have the loose source route IP option (0x83) enabled.

Source routing is used when a sending host specifies the router interfaces that the packet must traverse on it's way to it's destination.

With loose source routing enabled, the packet is forwarded if the routing table has a reverse path to the source IP address of the packet.

Example

The following command enables processing of the loose source route IP option:

```
enable ip-option loose-source-route
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable ip-security anomaly-protection

```
enable ip-security anomaly-protection {slot [ slot | all ]}
```

Description

Enables all anomaly checking options.

Syntax Description

<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.



Default

The default is disabled.

Usage Guidelines

This commands enables all anomaly checking options, including IP address, UDP/TCP port, TCP flag and fragment, and ICMP anomaly checking.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X440, X460, X480, and X670 platforms, whether or not included in a SummitStack, and the BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules.

enable ip-security anomaly-protection icmp

```
enable ip-security anomaly-protection icmp {slot [ slot | all ]}
```

Description

Enables ICMP size and fragment checking.

Syntax Description

<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command enables ICMP size and fragment checking. This checking takes effect for both IPv4 and IPv6 TCP packets. When enabled, the switch drops ICMP packets if one of following condition is true:

- Fragmented ICMP packets.
- IPv4 ICMP ping packets with payload size greater than the maximum IPv4 ICMP-allowed size. (The maximum allowed size is configurable.)
- IPv6 ICMP ping packets with payload size > the maximum IPv6 ICMP-allowed size. (The maximum allowed size is configurable.)



History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X440, X460, X480, and X670 platforms, whether or not included in a SummitStack, and the BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules.

enable ip-security anomaly-protection ip

```
enable ip-security anomaly-protection ip { slot [ slot | all ] }
```

Description

Enables source and destination IP address checking.

Syntax Description

<i>slot</i>	Specifies the slot.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command enables source and destination IP addresses checking. This checking takes effect for both IPv4 and IPv6 packets. When enabled, the switch drops IPv4/IPv6 packets if its source IP address are the same as the destination IP address. In most cases, the condition of source IP address being the same as the destination IP address indicates a Layer 3 protocol error. (These kind of errors are found in LAND attacks.)

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X440, X460, X480, and X670 platforms, whether or not included in a SummitStack, and the BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules.



enable ip-security anomaly-protection l4port

```
enable ip-security anomaly-protection l4port [tcp | udp | both] {slot [ slot |
all ]}
```

Description

Enables TCP and UDP ports checking.

Syntax Description

tcp	Specifies that the TCP port be enabled for checking.
udp	Specifies that the UDP port be enabled for checking.
both	Specifies both the TCP and UDP ports be enabled for checking.
<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command enabled TCP and UDP ports checking. This checking takes effect for both IPv4 and IPv6 TCP and UDP packets. When enabled, the switch drops TCP and UDP packets if its source port is the same as its destination port. In most cases, when the condition of source port is the same as that of the destination port, it indicates a Layer4 protocol error. (This type of error can be found in a BALT attack.)

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X440, X460, X480, and X670 platforms, whether or not included in a SummitStack, and the BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules.

enable ip-security anomaly-protection notify

```
enable ip-security anomaly-protection notify [log | snmp | cache] {slot [ slot |
all ]}
```



Description

Enables protocol anomaly notification.

Syntax Description

log	Specifies the switch to send the notification to a log file.
snmp	Specifies the switch to send an SNMP trap when an event occurs.
cache	Specifies the switch to send the notification to cache.
<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command enables anomaly notification. When enabled, any packet failed to pass enabled protocol checking is sent to XOS Host CPU and notifies the user. There are three different types of notifications:

- **log**: The anomaly events are logged into EMS log.
- **snmp**: The anomaly events generate SNMP traps.
- **cache**: The most recent and unique anomaly events are stored in memory for review and investigation.

When disabled, the switch drops all violating packets silently.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X440, X460, X480, and X670 platforms, whether or not included in a SummitStack, and the BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules.

enable ip-security anomaly-protection tcp flags

```
enable ip-security anomaly-protection tcp flags {slot [ slot | all ]}
```

Description

Enables TCP flag checking.



Syntax Description

<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command Enables TCP flag checking. This checking takes effect for both IPv4 and IPv6 TCP packets. When enabled, the switch drops TCP packets if one of following condition is true:

- TCP SYN flag==1 and the source port<1024
- TCP control flag==0 and the sequence number==0
- TCP FIN, URG, and PSH bits are set, and the sequence number==0
- TCP SYN and FIN both are set.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X440, X460, X480, and X670 platforms, whether or not included in a SummitStack, and the BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules.

enable ip-security anomaly-protection tcp fragment

```
enable ip-security anomaly-protection tcp fragment {slot [ slot | all ]}
```

Description

Enables TCP fragment checking.

Syntax Description

<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.



Usage Guidelines

This command enables TCP fragment checking. This checking takes effect for IPv4/IPv6. When it is enabled, the switch drops TCP packets if one of following condition is true:

- For the first IPv4 TCP fragment (its IP offset field==0), if its TCP header is less than the minimum IPv4 TCP header allowed size.
- For the first IPv6 TCP fragment (its IP offset field==0), if its TCP header is less than the minimum IPv6 TCP header allowed size.
- If its IP offset field==1 (for IPv4 only).

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X440, X460, X480, and X670 platforms, whether or not included in a SummitStack, and the BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules.

enable ip-security arp gratuitous-protection

```
enable ip-security arp gratuitous-protection {vlan} [all | vlan_name]
```

Description

Enables gratuitous ARP protection on one or all VLANs on the switch.

Syntax Description

all	Specifies all VLANs configured on the switch.
<i>vlan_name</i>	Specifies the VLAN.

Default

By default, gratuitous ARP protection is disabled.

Usage Guidelines

Beginning with ExtremeXOS 11.6, this command replaces the `enable iparp gratuitous protect` command.

Hosts can launch man-in-the-middle attacks by sending out gratuitous ARP requests for the router's IP address. This results in hosts sending their router traffic to the attacker, and the attacker forwarding that data to the router. This allows passwords, keys, and other information to be intercepted.



To protect against this type of attack, the router will send out its own gratuitous ARP request to override the attacker whenever a gratuitous ARP broadcast with the router's IP address as the source is received on the network.

Beginning with ExtremeXOS 11.6, if you enable both DHCP secured ARP and gratuitous ARP protection, the switch protects its own IP address and those of the hosts that appear as secure entries in the ARP table.

To protect the IP addresses of the hosts that appear as secure entries in the ARP table, use the following commands to enable DHCP snooping, DHCP secured ARP, and gratuitous ARP on the switch:

- `enable ip-security dhcp-snooping {vlan} vlan_name ports [all | ports] violation-action [drop-packet {[block-mac | block-port] [durationduration_in_seconds | permanently] | none}] {snmp-trap}`
- `enable ipsecurity arp learning learn-from-arp`
- `enable ip-security arp gratuitous-protection {vlan} [all | vlan_name]`

Displaying Gratuitous ARP Information

To display information about gratuitous ARP, use the following command:

```
show ip-security arp gratuitous-protection
```

Example

The following command enables gratuitous ARP protection for VLAN corp:

```
enable ip-security arp gratuitous-protectection vlan corp
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

enable ip-security arp learning learn-from-arp

```
enable ip-security arp learning learn-from-arp {vlan} vlan_name ports [all | ports]
```

Description

Enables ARP learning for the specified VLAN and member ports.



Syntax Description

<i>vlan_name</i>	Specifies the name of the VLAN to which this rule applies.
all	Specifies all ingress ports.
<i>ports</i>	Specifies one or more ingress ports.

Default

By default, ARP learning is enabled.

Usage Guidelines

ARP is part of the TCP/IP suite used to associate a device's physical address (MAC address) with its logical address (IP address). The switch broadcasts an ARP request that contains the IP address, and the device with that IP address sends back its MAC address so that traffic can be transmitted across the network. The switch maintains an ARP table (also known as an ARP cache) that displays each MAC address and its corresponding IP address.

By default, the switch builds its ARP table by tracking ARP requests and replies, which is known as ARP learning.

Displaying ARP Information

To display how the switch builds an ARP table and learns MAC addresses for devices on a specific VLAN and associated member ports, use the following command:

```
show ip-security arp learning {vlan} vlan_name
```

To view the ARP table, including permanent and DHCP secured ARP entries, use the following command:

```
show iparp {ip_addre | mac | vlanvlan_name | permanent} {vrvr_name}
```

Example

The following command enables ARP learning on port 1:1 of the VLAN learn:

```
enable ip-security arp learning learn-from-arp vlan learn ports 1:1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.



enable ip-security arp learning learn-from-dhcp

```
enable ip-security arp learning learn-from-dhcp {vlan} vlan_name ports [all | ports]
```

Description

Enables DHCP secured ARP learning for the specified VLAN and member ports.

Syntax Description

<i>vlan_name</i>	Specifies the name of the VLAN to which this rule applies.
all	Specifies all ingress ports.
<i>ports</i>	Specifies one or more ingress ports.

Default

By default, DHCP secured ARP learning is disabled.

Usage Guidelines

Use this command to configure the switch to add the MAC address and its corresponding IP address to the ARP table as a secure ARP entry. The switch does not update secure ARP entries, regardless of the ARP requests and replies seen by the switch. DHCP secured ARP is linked to the “DHCP snooping” feature. The same DHCP bindings database created when you enabled DHCP snooping is also used by DHCP secured ARP to create secure ARP entries. The switch only removes secure ARP entries when the corresponding DHCP entry is removed from the trusted DHCP bindings database.



Note

If you enable DHCP secured ARP on the switch, ARP learning continues, which allows insecure entries to be added to the ARP table.

The default ARP timeout (configure `iparp timeout`) and ARP refresh (enable `iparp refresh`) settings do not apply to DHCP secured ARP entries. The switch removes DHCP secured ARP entries upon any DHCP release packet received from the DHCP client.

Displaying ARP Information

To display how the switch builds an ARP table and learns MAC addresses for devices on a specific VLAN and associated member ports, use the following command:

```
show ip-security arp learning {vlan} vlan_name
```

To view the ARP table, including permanent and DHCP secured ARP entries, use the following command:

```
show iparp {ip_address | mac | vlanvlan_name | permanent} {vrvr_name}
```



Example

The following command enables DHCP secured ARP learning on port 1:1 of the VLAN learn and uses the default polling and retry intervals:

```
enable ip-security arp learning learn-from-dhcp vlan learn ports 1:1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

enable ip-security arp validation violation-action

```
enable ip-security arp validation {destination-mac} {source-mac} {ip} {vlan}
vlan_name [all | ports] violation-action [drop-packet {[block-port] [duration
duration_in_seconds | permanently]]] {snmp-trap}
```

Description

Enables ARP validation for the specified VLAN and member ports.

Syntax Description

destination-mac	Specifies that the switch checks the ARP payload for the MAC destination address in the Ethernet header and the receiver's host address in the ARP response.
source-mac	Specifies that the switch checks ARP requests and responses for the MAC source address in the Ethernet header and the sender's host address in the ARP payload.
ip	Specifies the switch checks the IP address in the ARP payload and compares it to the DHCP bindings database. If the IP address does exist in the DHCP bindings table, the switch verifies that the MAC address is the same as the sender hardware address in the ARP request. If not, the packet is dropped.
<i>vlan_name</i>	Specifies the name of the VLAN to which this rule applies.
all	Specifies all ports to participate in ARP validation.
<i>ports</i>	Specifies one or more ports to participate in ARP validation.
drop-packet	Specifies that the switch drops the invalid ARP packet.
block-port	Indicates that the switch blocks invalid ARP requests on the specified port.



<i>duration_in_seconds</i>	Specifies the switch to temporarily disable the specified port upon receiving an invalid ARP request. The range is seconds.
permanently	Specifies the switch to permanently disable the port upon receiving an invalid ARP request.
snmp-trap	Specifies the switch to send an SNMP trap when an event occurs.

Default

By default, ARP validation is disabled.

Usage Guidelines

The violation action setting determines what action(s) the switch takes when an invalid ARP is received.

Depending on your configuration, the switch uses the following methods to check the validity of incoming ARP packets:

- **Drop packet**—The switch confirms that the MAC address and its corresponding IP address are in the DHCP binding database built by DHCP snooping. This is the default behavior when you enable ARP validation. If the MAC address and its corresponding IP address are in the DHCP bindings database, the entry is valid. If the MAC address and its corresponding IP address are not in the DHCP bindings database, the entry is invalid, and the switch drops the ARP packet.
- **IP address**—The switch checks the IP address in the ARP payload. If the switch receives an IP address in the ARP payload that is in the DHCP binding database, the entry is valid. If the switch receives an IP address that is not in the DHCP binding database, for example 255.255.255.255 or an IP multicast address, the entry is invalid or unexpected.
- **Source MAC address**—The switch checks ARP requests and responses for the source MAC address in the Ethernet header and the sender's host address in the ARP payload. If the source MAC address and sender's host address are the same, the entry is valid. If the source MAC source and the sender's host address are different, the entry is invalid.
- **Destination MAC address**—The switch checks the ARP payload for the destination MAC address in the Ethernet header and the receiver's host address. If the destination MAC address and the target's host address are the same, the entry is valid. If the destination MAC address and the target's host address are different, the entry is invalid.

Any violation that occurs causes the switch to generate an Event Management System (EMS) log message. You can configure to suppress the log messages by configuring EMS log filters. For more information about EMS, see the EMS commands in [Commands for Status Monitoring and Statistics](#).

Displaying ARP Validation Information

To display information about ARP validation, use the following command:

```
show ip-security arp validation {vlan} vlan_name
```



Example

The following command enables ARP validation on port 1:1 of the VLAN valid:

```
enable ip-security arp validation vlan valid ports 1:1 drop-packet
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

enable ip-security dhcp-bindings restoration

```
enable ip-security dhcp-bindings restoration
```

Description

Enables download and upload of DHCP bindings.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

The command allows you to enable the download and upload of the DHCP bindings, essentially enabling the DHCP binding functionality. The default is disabled.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.



enable ip-security dhcp-snooping

```
enable ip-security dhcp-snooping {vlan} vlan_name ports [all | ports] violation-
action [drop-packet {[block-mac | block-port] [duration duration_in_seconds |
permanently] | none}]] {snmp-trap}
```

Description

Enables DHCP snooping for the specified VLAN and ports.

Syntax Description

<i>vlan_name</i>	Specifies the name of the DHCP-snooping VLAN. Create and configure the VLAN before enabling DHCP snooping.
all	Specifies all ports to receive DHCP packets.
<i>ports</i>	Specifies one or more ports to receive DHCP packets.
drop-packet	Indicates that the switch drop the rogue DHCP packet received on the specified port.
block-mac	Indicates that the switch blocks rogue DHCP packets from the specified MAC address on the specified port. The MAC address is added to the DHCP bindings database.
block-port	Indicates that the switch blocks rogue DHCP packets on the specified port. The port is added to the DHCP bindings database.
<i>duration_in_seconds</i>	Specifies that the switch temporarily disable the specified port upon receiving a rogue DHCP packet. The range is seconds.
permanently	Specifies that the switch to permanently disable the specified port upon receiving a rogue DHCP packet.
none	Specifies that the switch takes no action when receiving a rogue DHCP packet; the switch does not drop the packet.
snmp-trap	Specifies the switch to send an SNMP trap when an event occurs.

Default

By default, DHCP snooping is disabled.

Usage Guidelines

Use this command to enable DHCP snooping on the switch.



Note

Snooping IP fragmented DHCP packets is not supported.

The violation action setting determines what action(s) the switch takes when a rouge DHCP server packet is seen on an untrusted port or the IP address of the originating server is not among those of the



configured trusted DHCP servers. The DHCP server packets are DHCP OFFER, ACK and NAK. The following list describes the violation actions:

- **block-mac**—The switch automatically generates an ACL to block the MAC address on that port. The switch does not blackhole that MAC address in the FDB. The switch can either temporarily or permanently block the MAC address.
- **block-port**—The switch blocks all incoming rogue DHCP packets on that port. The switch disables the port either temporarily or permanently to block the traffic on that port.
- **none**—The switch takes no action to drop the rogue DHCP packet or block the port, and so on. In this case, DHCP snooping continues to build and manage the DHCP bindings database and DHCP forwarding will continue in hardware as before.

Any violation that occurs causes the switch to generate an Event Management System (EMS) log message. You can configure to suppress the log messages by configuring EMS log filters. For more information about EMS, see the EMS commands in [Commands for Status Monitoring and Statistics](#).

Displaying DHCP Snooping Information

To display the DHCP snooping configuration settings, use the following command:

```
show ip-security dhcp-snooping {vlan} vlan_name
```

To display the DHCP bindings database, use the following command:

```
show ip-security dhcp-snooping entries {vlan} vlan_name
```

To display any violations that occur, use the following command:

```
show ip-security dhcp-snooping violations {vlan} vlan_name
```

Example

The following command enables DHCP snooping on the switch and has the switch block DHCP packets from port 1:1:

```
enable ip-security dhcp-snooping vlan snoop ports 1:1 violation-action drop-  
packet block-port
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

enable ip-security source-ip-lockdown ports

```
enable ip-security source-ip-lockdown ports [all | ports]
```



Description

Enables the source IP lockdown feature on one or more ports.

Syntax Description

all	Specifies all ports for which source IP lockdown should be enabled.
ports	Specifies one or more ports for which source IP lockdown should be enabled.

Default

By default, source IP lockdown is disabled on the switch.

Usage Guidelines

Source IP lockdown prevents IP address spoofing by automatically placing source IP address filters on specified ports. If configured, source IP lockdown allows only traffic from a valid DHCP-assigned address obtained by a DHCP snooping-enabled port or an authenticated static IP address to enter the network.

To configure source IP lockdown, you must enable DHCP snooping on the ports connected to the DHCP server and DHCP client before you enable source IP lockdown. You must enable source IP lockdown on the ports connected to the DHCP client, not on the ports connected to the DHCP server. The same DHCP bindings database created when you enable DHCP snooping is also used by the source IP lockdown feature to create ACLs that permit traffic from DHCP clients. All other traffic is dropped. In addition, the DHCP snooping violation action setting determines what action(s) the switch takes when a rouge DHCP server packet is seen on an untrusted port.

To enable DHCP snooping, use the following command:

```
enable ip-security dhcp-snooping {vlan} vlan_name ports [all | ports] violation-
action [drop-packet {[block-mac | block-port] [duration duration_in_seconds |
permanently] | none}] {snmp-trap}
```

Displaying Source IP Lockdown Information

To display the source IP lockdown configuration on the switch, use the following command:

```
show ip-security source-ip-lockdown
```

Example

The following command enables source IP lockdown on ports 1:1 and 1:4:

```
enable ip-security source-ip-lockdown ports 1:1, 1:4
```

History

This command was first available in ExtremeXOS 11.6.



Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

enable mac-lockdown-timeout ports

```
enable mac-lockdown-timeout ports [all | port_list]
```

Description

Enables the MAC address lock down timeout feature for the specified port or group of ports or for all ports on the switch.

Syntax Description

all	Specifies all ports.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

By default, the MAC address lock down timeout feature is disabled.

Usage Guidelines

You cannot enable the MAC lock down timer on a port that also has the lock learning feature enabled.

Example

The following command enables the MAC address lock down timeout feature for ports 2:3, 2:4, and 2:6:

```
enable mac-lockdown-timeout ports 2:3, 2:4, 2:6
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

enable radius

```
enable radius {mgmt-access | netlogin}
```



Description

Enables the RADIUS client on the switch.

Syntax Description

mgmt-access	Specifies the switch management RADIUS authentication server.
netlogin	Specifies the network login RADIUS authentication server.

Default

RADIUS authentication is disabled for both switch management and network login by default.

Usage Guidelines

Before you enable RADIUS on the switch, you must configure the servers used for authentication and configure the authentication string (shared secret) used to communicate with the RADIUS authentication server.

To configure the RADIUS authentication servers, use the following command:

```
configure radius {mgmt-access | netlogin} [primary | secondary] server [ipaddress
| hostname] {udp_port} client-ip [ipaddress] {vrvr_name}
```

To configure the shared secret, use the following command:

```
configure radius {mgmt-access | netlogin} [primary | secondary] shared-secret
{encrypted} string
```

If you do not specify a keyword, RADIUS authentication is enabled on the switch for both management and network login. When enabled, all web, Telnet, and SSH logins are sent to the RADIUS servers for authentication. When used with a RADIUS server that supports ExtremeXOS CLI authorization, each CLI command is sent to the RADIUS server for authorization before it is executed.

Use the `mgmt-access` keyword to enable RADIUS authentication for switch management functions.

Use the `netlogin` keyword to enable RADIUS authentication for network login.

Example

The following command enables RADIUS authentication on the switch for both management and network login:

```
enable radius
```

The following command enables RADIUS authentication on the switch for network login:

```
enable radius netlogin
```



History

This command was first available in ExtremeXOS 10.1.

The `mgmt-access` and `netlogin` keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable radius-accounting

```
enable radius-accounting {mgmt-access | netlogin}
```

Description

Enables RADIUS accounting.

Syntax Description

mgmt-access	Specifies the switch management RADIUS accounting server.
netlogin	Specifies the network login RADIUS accounting server.

Default

RADIUS accounting is disabled for both switch management and network login by default.

Usage Guidelines

The RADIUS client must also be enabled.

Before you enable RADIUS accounting on the switch, you must configure the servers used for accounting and configure the authentication string (shared secret) used to communicate with the RADIUS accounting server.

To configure the RADIUS accounting servers, use the following command:

```
configure radius-accounting {mgmt-access | netlogin} [primary | secondary] server
[ipaddress | hostname] {tcp_port} client-ip [ipaddress] {vr vr_name}
```

To configure the shared secret, use the following command:

```
configure radius-accounting {mgmt-access | netlogin} [primary | secondary] shared-secret
{encrypted} string
```

If you do not specify a keyword, RADIUS accounting is enabled on the switch for both management and network login.

Use the `mgmt-access` keyword to enable RADIUS accounting for switch management functions.



Use the `netlogin` keyword to enable RADIUS accounting for network login.

Example

The following command enables RADIUS accounting on the switch for both management and network login:

```
enable radius-accounting
```

The following command enables RADIUS accounting for network login:

```
enable radius-accounting netlogin
```

History

This command was first available in ExtremeXOS 10.1.

The `mgmt-access` and `netlogin` keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable ssh2

```
enable ssh2 {access-profile [access_profile | none]} {port tcp_port_number} {vr
vr_name | all | default}
```

Description

Enables SSH2 server to accept incoming sessions from SSH2 clients.

Syntax Description

<code>access_profile</code>	Specifies an ACL policy.
<code>none</code>	Cancels a previously configured ACL policy.
<code>port</code>	Specifies a TCP port number. The default is port 22.
<code>vr_name</code>	Specifies a virtual router name. NOTE: User-created VRs are supported only on the platforms listed for this feature in Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i> .
<code>all</code>	Specifies that SSH is enabled on all virtual routers.
<code>default</code>	Specifies that SSH is enabled on the default virtual router.



Default

The SSH2 feature is disabled by default.

Usage Guidelines

SSH2 enables the encryption of session data. You must be logged in as an administrator to enable SSH2.

SSH2 functionality is not present in the base ExtremeXOS software image, but is in an additional, installable module. Before you can access any SSH2 commands, you must install the module. Without the module, the commands do not appear on the command line. To install the module, see the instructions in [Software Upgrade and Boot Options](#)

After you have installed the SSH2 module, you must generate a host key and enable SSH2. To generate an SSH2 host key, use the `configure ssh2 key` command. To enable SSH2, use the `enable ssh2` command.

Use the port option to specify a TCP port number other than the default port of 22. You can only specify ports 22 and 1024 through 65535.

Using ACLs to Control SSH Access

You can specify a list of predefined clients that are allowed SSH2 access to the switch. To do this, you configure an ACL policy to permit or deny a specific list of IP addresses and subnet masks for the SSH port. You must create an ACL policy file before you can use the access-profile option. If the ACL policy file does not exist on the switch, the switch returns an error message indicating that the file does not exist.

Use the none option to cancel a previously configured ACL.

In the ACL policy file for SSH2, the source-address field is the only supported match condition. Any other match conditions are ignored.

Policy files can also be configured using the following command:

```
configure ssh2 access-profile [ access_profile | [[addrule ] [first | [[before | after]previous_rule]]] | delete rule | none ]
```

Creating an ACL Policy File

To create an ACL policy file, use the `edit policy` command. For more information about creating and implementing ACL policy files, see [Policy Manager](#) and [ACLs](#) in the *ExtremeXOS Concepts Guide*.

If you attempt to implement a policy that does not exist on the switch, an error message similar to the following appears:

```
Error: Policy /config/MyAccessProfile_2.pol does not exist on file system
```

If this occurs, make sure the policy you want to implement exists on the switch. To confirm the policies on the switch, use the `ls` command. If the policy does not exist, create the ACL policy file.



Viewing SSH Information

To view the status of SSH2 sessions on the switch, use the **show management** command. This command displays information about the switch including the enable/disable state for SSH2 sessions and whether a valid key is present.

Example

The following command enables the SSH2 feature:

```
enable ssh2
```

The next example assumes you have already created an ACL to apply to SSH.

The following command applies the ACL MyAccessProfile_2 to SSH:

```
enable ssh2 access-profile MyAccessProfile_2
```

History

This command was first available in the ExtremeXOS 11.0 SSH module.

The access-profile and none options were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable tacacs

enable tacacs

Description

Enables TACACS+ authentication.

Syntax Description

This command has no arguments or variables.

Default

Disabled.



Usage Guidelines

After they have been enabled, all web and Telnet logins are sent to one of the two TACACS+ servers for login name authentication.

Example

The following command enables TACACS+ user authentication:

```
enable tacacs
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable tacacs-accounting

```
enable tacacs-accounting
```

Description

Enables TACACS+ accounting.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

If accounting is used, the TACACS+ client must also be enabled.

Example

The following command enables TACACS+ accounting for the switch:

```
enable tacacs-accounting
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable tacacs-authorization

enable tacacs-authorization

Description

Enables CLI command authorization.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

When enabled, each command is transmitted to the remote TACACS+ server for authorization before the command is executed. TACACS+ authentication must also be enabled to use TACACS+ authorization. Use the following command to enable authentication:

```
enable tacacs
```

Example

The following command enables TACACS+ command authorization for the switch:

```
enable tacacs-authorization
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.



enable web http

enable web http

Description

Enables hypertext transfer protocol (HTTP) access to the switch on the default HTTP port (80).

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

If HTTP access has been disabled, use this command to enable HTTP access to the switch.

Example

The following command enables HTTP on the default port:

```
enable web http
```

History

This command was first available in the ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

enable web https

enable web https

Description

Enables secure socket layer (SSL) access to the switch on the default port (443).

Syntax Description

This command has no arguments or variables.



Default

Disabled.

Usage Guidelines

Use this command to allow users to connect using a more secure HTTPS connection.

Similar to SSH2, before you can use any SSL commands, you must first download and install the separate Extreme Networks SSH software module (ssh.xmod). This additional module allows you to configure both SSH2 and SSL on the switch. SSL is packaged with the SSH module; therefore, if you do not install the module, you are unable to configure SSL. If you try to execute SSL commands without installing the module first, the switch notifies you to download and install the module. To install the module, see the instructions in [Software Upgrade and Boot Options](#) of the *ExtremeXOS Concepts Guide*.

To use secure HTTP access (HTTPS) for web-based login connections, you must specify HTTPS as the protocol when configuring the redirect URL. For more information about configuring the redirect URL, see the `configure netlogin redirect-page` command.

Prior to ExtremeXOS 11.2, the SSH module did not include SSL. To use SSL, you must download and install the current SSH module.

If you are currently running SSH with ExtremeXOS 11.0 or 11.1, and you want to use SSL for secure HTTPS web-based login, you must upgrade your core software image to ExtremeXOS 11.2 or later, install the SSH module that works in concert with that core software image, and reboot the switch.

Example

The following command enables SSL on the default port:

```
enable web https
```

History

This command was first available in the ExtremeXOS 11.2 and supported with the SSH module.

Platform Availability

This command is available on all platforms.

scp2

```
scp2 {vr vr_name} {cipher [3des | blowfish]} {port portnum} user [hostname |  
ipaddress]:remote_file local_file
```

or




```
scp2 {vr vr_name} {cipher [3des | blowfish]} {port portnum} local_file user
[hostname | ipaddress]:remote_file
```

Description

The first command initiates an SCP2 client session to a remote SCP2 server and copies a configuration or policy file from the remote system to the switch.

The second command initiates an SCP2 client session to a remote SCP2 server and copies a configuration or policy file from the switch to a remote system.

Syntax Description

<i>vr_name</i>	Specifies the virtual router. The default virtual router is VR-Mgmt.
	 Note User-created VRs are supported only on the platforms listed for this feature in Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i> .
3des	Specifies that the 3des cipher should be used for encryption. This is the default.
blowfish	Specifies that the blowfish cipher should be used for encryption.
<i>portnum</i>	Specifies the TCP port number to be used for communicating with the SSH2 client. The default is port 22.
<i>user</i>	Specifies a login name for the remote host.
<i>hostname</i>	Specifies the name of the remote host.
<i>ipaddress</i>	Specifies the IP address of the remote host.
<i>remote_file</i>	Specifies the name of the remote file (configuration file, policy file, image file, public key file) to be transferred.
<i>local_file</i>	Specifies the name of the local file (configuration file, policy file, image file, public key file) to be transferred.

Default

The default settings for SSH2 parameters are as follows:

- cipher—3des encryption.
- port—22.
- compression—off.
- vr_name—VR-Mgmt.

Usage Guidelines

You must be running the SSH2 module (ssh.xmod), which is under Export Control, in order to use the SCP2 command.

SSH2 does not need to be enabled on the switch in order to use this command.



This command logs into the remote host as *user* and accesses the file *remote_file*. You will be prompted for a password from the remote host, if required.

Host Name, User Name, and Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name, user name, or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z).
- Numerals (0-9).
- Period (.).
- Dash (-) Permitted for host and user names.
- Underscore (_) Permitted for host and user names.
- Colon (:).
- At symbol (@). Permitted only for user names.
- Slash (/). Permitted only for user names.

When naming the host, creating a user name, or configuring the IP address, remember the requirements listed above.

Remote Filename Character Restrictions

This section provides information about the characters supported by the switch for remote filenames.

When specifying a remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z).
- Numerals (0-9).
- Period (.).
- Dash (-).
- Underscore (_).
- Slash (/).

When naming a remote file, remember the requirements listed above.

Example

The following command copies the configuration file test.cfg on host system1 to the switch:

```
scp2 admin@system1:test.cfg localtest.cfg
```

The following command copies the configuration file engineering.cfg from the switch to host system1:

```
scp2 engineering.cfg admin@system1:engineering.cfg
```



The following command copies the file Anna5.xsf from the default virtual router to 150.132.82.140:

```
scp2 vr vr-default Anna5.xsf root@150.132.82.140:Anna5.xsf
Upload /config/Anna5.xsf to
Connecting to 150.132.82.140...
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show dhcp-server

```
show dhcp-server {vlan vlan_name}
```

Description

Displays the DHCP server's configuration and address allocation on a specified VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN of the DHCP server of interest.
------------------	--

Default

N/A.

Usage Guidelines

If no VLAN is specified, the configuration and address allocation for the servers on all the VLANs is displayed.

Example

The following command displays the configuration and address allocation for the DHCP server for the VLAN test:

```
show dhcp-server vlan test
```

The following is sample output from this command:

```
# show dhcp-server
VLAN "Default":
```



```

DHCP Address Range      : Not configured
Netlogin Lease Timer   : Not configured (Default = 10 seconds)
DHCP Lease Timer       : Not configured (Default = 7200 seconds)
DHCP Option Code  12   : hex "11:22:33:44:45"
DHCP Option Code  69   : ipaddress 10.0.0.1 10.0.0.2
Ports DHCP Enabled    : No ports enabled

```

History

This command was first available in ExtremeXOS 11.0.

The output is modified to show primary and secondary DNS servers in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

show dos-protect

```
show dos-protect {detail}
```

Description

Displays DoS protection configuration and state.

Syntax Description

detail	Specifies to display statistics in addition to configuration and state.
---------------	---

Default

N/A.

Usage Guidelines

Use this command to display the DoS protection settings. Using the detail option will also display the following cumulative statistics:

- trusted.
- notify.
- alerts.

Example

The following command displays the DoS protection settings for the switch:

```
show dos-protect
```



The following is sample output from this command:

```
dos-protect is disabled
dos-protect settings:
interval:      1 (measurement interval secs)
acl expire time: 5 (secs)
trusted ports:
no trusted ports configured
type L3-Protect:
notify threshold: 3500 (level to log a message)
alert threshold: 4000 (level to generate an ACL)
```

The following command displays detailed DoS protection settings for the switch:

```
show dos-protect detail
```

The following is sample output from this command:

```
dos-protect is enabled
dos-protect settings:
interval:      1 (measurement interval secs)
acl expire time: 5 (secs)
trusted ports:
1:2
type L3-Protect:
notify threshold: 3500 (level to log a message)
alert threshold: 4000 (level to generate an ACL)
dos-protect statistics:
trusted:      1301
notify:       0
alerts:       0
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

show ip-security anomaly-protection notify cache ports

```
show ip-security anomaly-protection notify cache ports port_list
```

Description

Displays most anomaly notification caches.



Syntax Description

<code>port_list</code>	Specifies one or more ports or slots and ports.
------------------------	---

Default

N/A.

Usage Guidelines

This command displays most anomaly notification caches.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is only available on the Summit X440, X460, X480, and X670 platforms, whether or not included in a SummitStack, and the BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules.

show ip-security arp gratuitous-protection

show ip-security arp gratuitous-protection

Description

If configured for gratuitous ARP, displays the gratuitous ARP protection configuration on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The switch displays the name of each VLAN configured for gratuitous ARP.

If you do not have gratuitous ARP configured, the switch does not display any VLAN information.



Example

The following command displays the gratuitous ARP configuration on the switch:

```
show ip-security arp gratuitous-protection
```

The following is sample output from this command:

```
Gratuitous ARP Protection enabled on following VLANs:
Default, test
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

show ip-security arp learning

```
show ip-security arp learning {vlan} vlan_name
```

Description

Displays how the switch builds an ARP table and learns MAC addresses for devices on a specific VLAN and associated member ports.

Syntax Description

<i>vlan_name</i>	Specifies the name of the VLAN.
------------------	---------------------------------

Default

N/A.

Usage Guidelines

The switch displays the following ARP learning information:

- Port—The member port of the VLAN.
- Learn-from—The method the port uses to build the ARP table. The methods are:
 - ARP—ARP learning is enabled. The switch uses a series of requests and replies to build the ARP table.
 - DHCP—DHCP secured ARP is enabled. The switch uses DHCP snooping to build the ARP table.



- None—Both DHCP secured ARP and ARP learning are disabled.

Example

The following command displays how the switch builds its ARP table for the VLAN learn:

```
show ip-security arp learning vlan learn
```

The following is sample output from this command:

Port	Learn-from
2:1	ARP
2:2	DHCP, poll 300 sec, retries 3
2:3	ARP
2:4	None
2:5	ARP
2:6	ARP
2:7	ARP
2:8	ARP

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

show ip-security arp validation

```
show ip-security arp validation {vlan} vlan_name
```

Description

Displays ARP validation information for the specified VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the name of the VLAN.
------------------	---------------------------------

Default

N/A.



Usage Guidelines

The switch displays the following ARP validation information:

- Port—Indicates the port that received the ARP entry.
- Validation—Indicates how the entry is validated.
- Violation-action—Determines what action(s) the switch takes when an invalid ARP is received.

Example

The following command displays ARP validation on for the VLAN valid:

```
show ip-security arp validation vlan valid
```

The following is sample output from this command:

```
-----
Port      Validation      Violation-action
-----
7         DHCP            drop-packet, block-port for 120 seconds, snmp-
trap
23        DHCP            drop-packet, block-port for 120 seconds, snmp-
trap
-----
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

show ip-security arp validation violations

```
show ip-security arp validation violations {vlan} vlan_name ports [ports | all]
```

Description

Displays the violation count on an ARP validation.

Syntax Description

<i>vlan_name</i>	Specifies the name of the VLAN.
<i>ports</i>	Specifies the name of the port.
all	Specifies all ports.



Default

N/A.

Usage Guidelines

The switch displays the following ARP validation information:

- Port—Indicates the port that received the ARP entry.
- Validation—Indicates how the entry is validated.
- Violation count—Indicates the number of violations for each port.

Example

The following command displays ARP validation violation counts on all ports:

```
show ip-security arp validation violations ragu ports all
```

The following is sample output from this command:

```
-----
Port      Validation Violation Count
-----
1:1 ip,DHCP 1233
1:3 ip,DHCP 3425
1:4 ip,DHCP 5654
1:5 ip,DHCP 0
1:6 ip,DHCP 3645
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

show ip-security dhcp-snooping entries

```
show ip-security dhcp-snooping entries {vlan} vlan_name
```

Description

Displays the DHCP bindings database on the switch.



Syntax Description

<code>vlan_name</code>	Specifies the name of the DHCP-snooping VLAN.
------------------------	---

Default

N/A.

Usage Guidelines

The switch displays the following DHCP bindings database information:

- VLAN—The name of the DHCP-snooping VLAN.
- IP Addr—The IP address of the untrusted interface or client.
- MAC Addr—The MAC address of the untrusted interface or client.
- Port—The port number where the untrusted interface or client attempted to access the network.

Example

The following command displays the DHCP bindings database on the switch:

```
show ip-security dhcp-snooping entries vlan dhcpVlan
```

The following is sample output from this command:

```
-----
Vlan: dhcpVlan
-----
Server      Client
IP Addr    MAC Addr          Port      Port
-----
172.16.100.9  00:90:27:c6:b7:65  1:1      1:2
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

show ip-security dhcp-snooping information-option

```
show ip-security dhcp-snooping information-option
```



Description

Displays the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82) settings.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays DHCP relay agent option (option 82) settings. For example, the following command:

```
show ip-security dhcp-snooping information-option
```

Generates the following output:

```
Information option insertion: Enabled Information option checking :  
Disabled Information option policy : Drop
```

The following command:

```
show ip-security dhcp-snooping information-option
```

Generates the following output:

```
Information option insertion: Enabled Information option checking :  
Enabled Information option policy : Keep
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

show ip-security dhcp-snooping information circuit-id port-information

```
show ip-security dhcp-snooping information circuit-id port-information ports  
[port_list | all ]
```

Description

Displays the port information portion of the circuit ID for the indicated port(s).



Syntax Description

<i>port_list</i>	Specifies one or more ports.
all	Specifies all ports.

Default

N/A.

Usage Guidelines

This command displays the port information portion of the circuit ID for the indicated ports.

Example

The following command:

```
# show ip-security dhcp-snooping information circuit-id port-information
ports 1-7
```

Displays the following output:

```
Port          Circuit-ID Port information string
-----
1             portinfostring1
2             portinfostring2
3             portinfostring3
4             portinfostring4
5             portinfostring5
Port          Circuit-ID Port information string
-----
6             1006
7             1007
```

Note: The full Circuit ID string has the form '<Vlan Info>--<Port Info>'

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.



show ip-security dhcp-snooping information-option circuit-id vlan-information

```
show ip-security dhcp-snooping information-option circuit-id vlan-information
[{vlan} vlan_name]
```

Description

Displays the VLAN information portion of the circuit ID for the indicated VLAN.

Syntax Description

<code>vlan_name</code>	Specifies a vlan name.
------------------------	------------------------

Default

N/A.

Usage Guidelines

This command displays the VLAN information portion of the circuit ID for the indicated VLAN. When a VLAN is not specified, the circuit ID information for all the VLANs is displayed.

Example

The following command:

```
show ip-security dhcp-snooping information-option circuit-id vlan-information
vlan Mktg
```

Displays the following output:

```
Vlan                               Circuit-ID vlan information string
----                               -
Mktg                               DSLAM1
Note: The full Circuit ID string has the form <Vlan Info>-<Port ifIndex>.
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.



show ip-security dhcp-snooping

```
show ip-security dhcp-snooping {vlan} vlan_name
```

Description

Displays the DHCP snooping configurations on the switch.

Syntax Description

<i>vlan_name</i>	Specifies the name of the DHCP-snooping VLAN.
------------------	---

Default

N/A.

Usage Guidelines

The switch displays the following DHCP snooping information:

- DHCP snooping enabled on ports—The ports that have DHCP snooping enabled.
- Trusted ports—The ports configured as trusted ports.
- Trusted DHCP servers—The servers configured as trusted DHCP servers.
- Port—The specific port that has DHCP snooping enabled.
- Violation-action—The action the switch takes upon detecting a rogue DHCP packet on the port.

Example

The following command displays the DHCP snooping settings for the switch:

```
show ip-security dhcp-snooping vlan "Default"
```

The following is sample output from this command:

```
DHCP Snooping enabled on ports: 7, 9, 11
Trusted Ports: None
Trusted DHCP Servers: None
Bindings Restoration      : Enabled
Bindings Filename        : dhcpsonia.xsf
Bindings File Location    :
Primary Server   : 10.1.1.14, VR-Default, TFTP
Secondary Server: None
Bindings Write Interval : 5 minutes
Bindings last uploaded at:
-----
Port          Violation-action
-----
7             none
```



9	none
11	none

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

show ip-security dhcp-snooping violations

```
show ip-security dhcp-snooping violations {vlan} vlan_name
```

Description

Displays the MAC addressed from which the rouge DHCP packet was received by the switch.

Syntax Description

<i> vlan_name </i>	Specifies the name of the DHCP-snooping VLAN.
--------------------	---

Default

N/A.

Usage Guidelines

The switch displays the following DHCP snooping information:

- Port—The specific port that received the rouge DHCP packet.
- Violating MAC—The MAC address from which the rouge DHCP was received by the switch.

Example

The following command displays the DHCP snooping violations for the VLAN green:

```
show ip-security dhcp-snooping violations green
```

The following is sample output from this command:

```
Violations seen on following ports
-----
Port                Violating MAC
```



```
-----  
2:3          00-0c-11-a0-3e-12
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

show ip-security source-ip-lockdown

show ip-security source-ip-lockdown

Description

Displays the source IP lockdown configuration on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The switch displays the following source IP lockdown information:

- Port—Indicates the port that has DHCP snooping enabled and is configured for source IP lockdown.
- Locked IP Address—Indicates a valid DHCP-assigned address obtained by a DHCP snooping-enabled port or an authenticated static IP address.

Example

The following command displays the source IP configuration on the switch:

```
show ip-security source-ip-lockdown
```

The following is sample output from this command:

```
Ports          Locked IP Address  
23 10.0.0.101
```



History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

show mac-lockdown-timeout fdb ports

```
show mac-lockdown-timeout fdb ports [all | port_list]
```

Description

Displays the MAC entries that are learned on the specified port or group of ports or for all ports on the switch along with the aging time of each port.

Syntax Description

all	Specifies all ports
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

If a port is down, the command displays all of the MAC entries that are maintained locally in the software.

The MAC entries learned on the specified port are displayed only if the MAC lock down timeout feature is enabled on the port. If you specify a port on which this feature is disabled, the MAC entries learned on that port are not displayed.

The switch displays the following information:

- Mac—The MAC address that defines the entry.
- Vlan—The VLAN name and ID for the entry.
- Age—The age of the entry, in seconds.
- Flags—Flags that define the type of entry:
 - B—Egress Blackhole.
 - b—Ingress Blackhole.
 - F—Entry in the hardware FDB.
 - L—Entry in the software.
- Port—The port on which the MAC address has been learned.



Example

The following command displays information about the MAC address lock down timeout settings for ports 2:3 and 2:4:

```
show mac-lockdown-timeout fdb ports 2:3, 2:4
```

The following is sample output from this command:

```

Mac                Vlan      Age  Flags  Port
-----
00:00:01:02:03:04v1(4094)0010F2:3
00:00:01:00:00:02v1(4094)0030FB b2:3
00:00:0A:02:03:04v2(4093)0050L2:4
00:00:0B:02:03:04v2(4093)0090F2:4
Flags : (F) Entry as in h/w FDB, (L) Entry in s/w and not in h/w
        (B) Egress Blackhole, (b) Ingress Blackhole
Total: 4 Entries in FDB: 3Entries in s/w: 1

```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

show mac-lockdown-timeout ports

```
show mac-lockdown-timeout ports [all | port_list]
```

Description

Displays information about the MAC address lock down timeout feature for the specified port or group of ports or for all ports on the switch.

Syntax Description

all	Specifies all ports
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

N/A.



Usage Guidelines

The switch displays the following MAC address timeout information:

- Port—Indicates the port number that you specified in the command.
- MAC Lockdown Timeout—Specifies the enabled/disabled state of the MAC address lock down timeout feature.
- Timeout (in seconds)—Specifies the timeout value for the specified ports. By default, the timeout value is 15 seconds. Even if MAC address lock down is disabled, the default timeout value is displayed.

Example

The following command displays information about the MAC address lock down timeout settings for ports 2:3, 2:4, and 2:6:

```
show mac-lockdown-timeout ports 2:3, 2:4, 2:6
```

The following is sample output from this command:

```
Ports   MAC Lockdown Timeout   Timeout (in seconds)
=====
2:3 Enabled300
2:4 Enabled 300
2:6Disabled                15
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

show ports rate-limit flood

```
show ports {port_list} rate-limit flood {no-refresh}
```

Description

Displays rate-limit discard statistics.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
no-refresh	Specifies a static snapshot of data.



Default

N/A.

Usage Guidelines

This command displays the per port ingress rate-limit flood traffic counter as well as information about received packets that have not been discarded due to rate-limiting.

It is used to show the results of the `configure ports port_list rate-limit flood [broadcast | multicast | unknown-destmac] [no-limit | pps]` command.

Note



As part of the system health check, the system polls the Rate-limit Flood Counters every 5 minutes and looks for non-zero counters on a port. A HAL.RateLimit.Info log message is logged when this is first detected on a port to alert the user that something in the network has triggered the rate limiting to occur. The message is not be logged again unless the counters are cleared.

Example

The following command displays information for port 1:1 without a screen refresh on a BlackDiamond 8800 switch:

```
show port 1:1 rate-limit flood no-refresh
```

Following is sample output from this command:

```
BD-8810.1 # show port 1:1 rate-limit flood no-refresh
Port Rate-Limit Discard Monitor Tue May 27 13:02:37 2008
Port   Link   Rx Pkt   Rx Byte Rx Pkt Rx Pkt   Flood Rate
State  Count  Count   Bcast  Mcast           Exceeded
=====
1:1    R      5225    65230  2112    0        2112
=====
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
```

Following is sample out put from this command:

```
* switch # show ports rate-limit flood
Port Rate-Limit Discard Monitor           Wed Oct  8 13:15:00
2008
Port   Link   Rx Pkt   Rx Byte   Rx Pkt   Rx Pkt   Flood
Rate
State  Count  Count   Bcast     Mcast     Exceeded
=====
==
```




```

1      A      0      0      0
0      0
2      A      0      0      0
0      0
3      R      0      0      0
0      0
4      R      0      0      0
0      0
5      R      0      0      0
0      0
6      R      0      0      0
0      0
7      R      0      0      0
0      0
8      R      0      0      0
0      0
9      R      0      0      0
0      0
10     R      0      0      0
0      0
11     R      0      0      0
0      0
12     R      0      0      0
0      0
13     R      0      0      0
0      0
14     R      0      0      0
0      0
15     R      0      0      0
0      0
16     R      0      0      0
0      0

```

```
=====
```

```
==
```

```

> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
0->Clear Counters U->page up D->page down ESC->exit

```

The following command displays information without a screen refresh on a Summit switch:

```
show ports rate-limit flood no-refresh
```

Following is sample out put from this command:

```

* switch # show ports rate-limit flood no-refresh
Port Rate-Limit Discard Monitor
Port      Link      Rx Pkt      Rx Byte      Rx Pkt      Rx Pkt      Flood
Rate
State     Count      Count      Bcast        Mcast        Exceeded
=====
==
1      A      0      0      0      0
0      0
2      A      0      0      0      0
0      0

```



```

3      R      0      0      0
0      0
4      R      0      0      0
0      0
5      R      0      0      0
0      0
6      R      0      0      0
0      0
7      R      0      0      0
0      0
8      R      0      0      0
0      0
9      R      0      0      0
0      0
10     R      0      0      0
0      0
11     R      0      0      0
0      0
12     R      0      0      0
0      0
13     R      0      0      0
0      0
14     R      0      0      0
0      0
15     R      0      0      0
0      0
16     R      0      0      0
0      0
17     R      0      0      0
0      0
18     R      0      0      0
0      0
19     A      0      0      0
0      0
20     A      0      0      0
0      0
21     R      0      0      0
0      0
22     R      0      0      0
0      0
23     R      0      0      0
0      0
24     R      0      0      0
0      0
25     R      0      0      0
0      0
26     R      0      0      0
0      0
27     R      0      0      0
0      0
28     R      0      0      0
0      0

```

```
=====
```

```
==
```

```

> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present L-Loopback

```



History

This command was first available in ExtremeXOS 12.2.

Platform Availability

This command is available on the Summit family switches, whether or not included in a SummitStack, BlackDiamond X8 series switches, and BlackDiamond 8000 c-, e-, xl-, and xm-series modules.

show radius

```
show radius {mgmt-access | netlogin}
```

Description

Displays the current RADIUS client configuration and statistics.

Syntax Description

mgmt-access	Specifies configuration and statistics for the switch management RADIUS authentication server.
netlogin	Specifies configuration and statistics for the network login RADIUS authentication server.

Default

N/A.

Usage Guidelines

If you do not specify a keyword, configuration details related to both management and network login are displayed. The output from this command displays the status of RADIUS and RADIUS accounting (enabled or disabled) and the primary and secondary servers for RADIUS and RADIUS accounting.

Use the `mgmt-access` keyword to display only RADIUS configuration details related to management access.

Use the `netlogin` keyword to only RADIUS configuration details related to network login.

Example

The following command displays the current RADIUS client configuration and statistics for both management and network login:

```
show radius
```



The following is sample output from this command:

```

Switch Management Radius: enabled
Switch Management Radius server connect time out: 3 seconds
Switch Management Radius Accounting: disabled
Switch Management Radius Accounting server connect time out: 3 seconds
Netlogin Radius: enabled
Netlogin Radius server connect time out: 3 seconds
Netlogin Radius Accounting: disabled
Netlogin Radius Accounting server connect time out: 3 seconds
Primary Switch Management Radius server:
Server name      :
IP address       : 10.100.1.100
Server IP Port: 1812
Client address: 10.116.3.101 (VR-Mgmt)
Shared secret   : g~`#uovpkkpvi~`
Access Requests : 0                Access Accepts      : 0
Access Rejects  : 0                Access Challenges   : 0
Access Retransmits: 0            Client timeouts     : 0
Bad authenticators: 0            Unknown types       : 0
Round Trip Time : 0
Secondary Switch Management Radius server:
Server name      :
IP address       : 10.100.1.101
Server IP Port: 1812
Client address: 10.116.3.101 (VR-Mgmt)
Shared secret   : g~`#uovpkkpvi~`
Access Requests : 0                Access Accepts      : 0
Access Rejects  : 0                Access Challenges   : 0
Access Retransmits: 0            Client timeouts     : 0
Bad authenticators: 0            Unknown types       : 0
Round Trip Time : 0
Primary Netlogin Radius server:
Server name      :
IP address       : 10.100.1.200
Server IP Port: 1812
Client address: 10.116.3.101 (VR-Mgmt)
Shared secret   : g~`#uovpkkpvi~`
Access Requests : 0                Access Accepts      : 0
Access Rejects  : 0                Access Challenges   : 0
Access Retransmits: 0            Client timeouts     : 0
Bad authenticators: 0            Unknown types       : 0
Round Trip Time : 0
Secondary Netlogin Radius server:
Server name      :
IP address       : 10.100.1.201
Server IP Port: 1812
Client address: 10.116.3.101 (VR-Mgmt)
Shared secret   : g~`#uovpkkpvi~`
Access Requests : 0                Access Accepts      : 0
Access Rejects  : 0                Access Challenges   : 0
Access Retransmits: 0            Client timeouts     : 0
Bad authenticators: 0            Unknown types       : 0
Round Trip Time : 0

```



History

This command was first available in ExtremeXOS 10.1.

The `mgmt-access` and `netlogin` keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show radius-accounting

```
show radius-accounting {mgmt-access | netlogin}
```

Description

Displays the current RADIUS accounting client configuration and statistics.

Syntax Description

mgmt-access	Specifies configuration and statistics for the switch management RADIUS accounting server.
netlogin	Specifies configuration and statistics for the network login RADIUS accounting server.

Default

N/A.

Usage Guidelines

If you do not specify a keyword, configuration details related to both management and network login are displayed. The output from this command displays information about the status and configuration of RADIUS accounting.

Use the `mgmt-access` keyword to display only RADIUS accounting configuration details related to management access.

Use the `netlogin` keyword to display only RADIUS accounting configuration details related to network login.

Example

The following command displays RADIUS accounting client configuration and statistics for both management and network login:

```
show radius-accounting
```



The following is sample output from this command:

```

Switch Management Radius Accounting: disabled
Switch Management Radius Accounting server connect time out: 3 seconds
Netlogin Radius Accounting: disabled
Netlogin Radius Accounting server connect time out: 3 seconds
Primary Switch Management Accounting server:
Server name      :
IP address       : 10.100.1.100
Server IP Port   : 1813
Client address   : 10.116.3.101 (VR-Mgmt)
Shared secret    : g~`#uovpkkpvi~`
Acct Requests    : 0                      Acct Responses    : 0
Acct Retransmits : 0                      Timeouts          : 0
Secondary Switch Management Accounting server:
Server name      :
IP address       : 10.100.1.101
Server IP Port   : 1813
Client address   : 10.116.3.101 (VR-Mgmt)
Shared secret    : g~`#uovpkkpvi~`
Acct Requests    : 0                      Acct Responses    : 0
Acct Retransmits : 0                      Timeouts          : 0
Primary Netlogin Accounting server:
Server name      :
IP address       : 10.100.1.200
Server IP Port   : 1813
Client address   : 10.116.3.101 (VR-Mgmt)
Shared secret    : g~`#uovpkkpvi~`
Acct Requests    : 0                      Acct Responses    : 0
Acct Retransmits : 0                      Timeouts          : 0
Secondary Netlogin Accounting server:
Server name      :
IP address       : 10.100.1.201
Server IP Port   : 1813
Client address   : 10.116.3.101 (VR-Mgmt)
Shared secret    : g~`#uovpkkpvi~`
Acct Requests    : 0                      Acct Responses    : 0
Acct Retransmits : 0                      Timeouts          : 0

```

History

This command was first available in ExtremeXOS 10.1.

The mgmt-access and netlogin keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show ssh2 private-key

```
show ssh2 private-key
```



Description

Displays the ssh2 server's private key.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the ssh server's private key which can be used to configure the key later or on another switch by using the `configure ssh2 key {pregenerated}` command. The key is saved in the switch's EEPROM.

To erase the key from the EEPROM, use the unconfigure switch command.

History

This command was first available in ExtremeXOS 12.1.

This command was added to ExtremeXOS 11.6 SR, and ExtremeXOS 12.0 SR.

Platform Availability

This command is available on all platforms.

show sshd2 user-key

```
show sshd2 user-key {key_name {users}}
```

Description

Displays the user names bound to a key.

Syntax Description

<i>key_name</i>	Specifies the name of the public key.
users	Specifies the name of the users.

Default

N/A.



Usage Guidelines

This command displays the names of the users that are bound to a public key.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

show ssl

show ssl {detail}

Description

Displays the secure socket layer (SSL) configuration.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the following information:

- HTTPS port configured. This is the port on which the clients will connect.
- Length of the RSA key (the number of bits used to generate the private key).
- Basic information about the stored certificate.

Similar to SSH2, before you can use any SSL commands, you must first download and install the separate Extreme Networks SSH software module (ssh.xmod). This additional module allows you to configure both SSH2 and SSL on the switch. SSL is packaged with the SSH module; therefore, if you do not install the module, you are unable to configure SSL. If you try to execute SSL commands without installing the module first, the switch notifies you to download and install the module. To install the module, see the instructions in [Software Upgrade and Boot Options](#) of the *ExtremeXOS Concepts Guide*.

If you attempt to use this command before installing the SSH module, the switch displays a message similar to the following:



SSL Module: Not Installed.



Note

The switch utilizes the SSH module for SSL functionality. You do not install an SSL module, only the SSH module.

Example

The following command displays the SSL configuration:

```
show ssl
```

The following is sample output from this command:

```
HTTPS Port Number: 443
Private Key matches with the Public Key in certificate. (or Private key does
not match with the Public Key in the certificate)
RSA Key Length: 1024
Certificate:
Data:
Version: 1 (0x0)
Serial Number: 6 (0x6)
Signature Algorithm: md5WithRSAEncryption
Issuer: C=AU, O=CryptSoft Pty Ltd, CN=Test CA (1024 bit)
Validity
Not Before: Oct 16 22:31:03 2000 GMT
Not After : Jan 14 22:31:03 2003 GMT
Subject: C=AU, O=CryptSoft Pty Ltd, CN=Server test cert (512 bit)
```

History

This command was first available in the ExtremeXOS 11.2 and supported with the SSH module.

Platform Availability

This command is available on all platforms.

show tacacs

show tacacs

Description

Displays the current TACACS+ configuration and statistics.



Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The output of this command displays the following information:

- TACACS+—The current state of TACACS+, enabled or disabled.
- TACACS+ Authorization—The current state of TACACS+ authorization, enabled or disabled.
- TACACS+ Accounting—The current state of TACACS+ accounting, enabled or disabled.
- TACACS+ Server Connect Timeout—The amount of time configured to detect and recover from a TACACS+ server failure.
- Primary TACACS+ Server—Describes information about the primary TACACS+ server, including:
 - The name of the primary TACACS+ server.
 - The IP address of the primary TACACS+ server.
 - The TCP port to use to contact the primary TACACS+ server.
 - The IP address and VR used by the switch.
 - The shared secret configured for the primary TACACS+ server.
- Secondary TACACS+ Server—Contains the same type of output as the primary TACACS+ server for the secondary TACACS+ server, if configured.
- TACACS+ Acct Server Connect Timeout—The amount of time configured to detect and recover from a TACACS+ accounting server failure.
- TACACS+ Accounting Server parameters, if configured. Contains the same type of output as the TACACS+ server for the TACACS+ accounting server(s), if configured.

Example

The following command displays TACACS+ client configuration and statistics:

```
show tacacs
```

The following is sample output from this command:

```
TACACS+: enabled
TACACS+ Authorization: enabled
TACACS+ Accounting : enabled
TACACS+ Server Connect Timeout sec: 3
Primary TACACS+ Server:
Server name      :
IP address       : 10.201.31.238
Server IP Port: 49
Client address: 10.201.31.65 (VR-Default)
Shared secret   : qijxou
```



```
Secondary TACACS+ Server:
Server name      :
IP address       : 10.201.31.235
Server IP Port   : 49
Client address   : 10.201.31.65 (VR-Default)
Shared secret    : qijxou
TACACS+ Acct Server Connect Timeout sec: 3
Primary TACACS+ Accounting Server:
Server name      :
IP address       : 10.201.31.238
Server IP Port   : 49
Client address   : 10.201.31.65 (VR-Default)
Shared secret    : qijxou
Secondary TACACS+ Accounting Server:
Server name      :
IP address       : 10.201.31.235
Server IP Port   : 49
Client address   : 10.201.31.65 (VR-Default)
Shared secret    : qijxou
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show tacacs-accounting

show tacacs-accounting

Description

Displays the current TACACS+ accounting client configuration and statistics.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The output of this command displays the following information:



- TACACS+ Accounting—The current state of TACACS+ accounting, enabled or disabled.
- TACACS+ Accounting Server Connect Timeout—The amount of time configured to detect and recover from a TACACS+ server failure.
- Primary TACACS+ Accounting Server—Describes information about the primary TACACS+ accounting server, including:
 - The name of the primary TACACS+ accounting server.
 - The IP address of the primary TACACS+ accounting server.
 - The TCP port to use to contact the primary TACACS+ accounting server.
 - The IP address and VR used by the switch.
 - The shared secret configured for the primary TACACS+ accounting server.
- Secondary TACACS+ Accounting Server—Contains the same type of output as the primary TACACS+ accounting server for the secondary TACACS+ accounting server, if configured.

Example

The following command displays TACACS+ accounting client configuration and statistics:

```
show tacacs-accounting
The following is sample output of this command:
TACACS+ Accounting : enabled
TACACS+ Acct Server Connect Timeout sec: 3
Primary TACACS+ Accounting Server:
Server name      :
IP address       : 10.201.31.238
Server IP Port   : 49
Client address   : 10.201.31.85 (VR-Default)
Shared secret    : qijxou
Secondary TACACS+ Accounting Server:Not configured
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show vlan dhcp-address-allocation

```
show {vlan} vlan_name dhcp-address-allocation
```

Description

Displays the DHCP server's address allocation on a specified VLAN.



Syntax Description

<code>vlan_name</code>	Specifies the VLAN of the DHCP server of interest.
------------------------	--

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the configuration of the DHCP for the VLAN corp:

```
show vlan corp dhcp-address-allocation
```

The following is sample output from this command:

```
=====
IP                MAC                State      Lease Renewal Time
=====
10.0.0.2         00:02:03:04:05:00  Offered   0000:00:10
10.0.0.3         00:08:03:04:05:00  Assigned  0000:59:09
10.0.0.4         ee:1c:00:04:05:00  Assigned  0000:59:09
=====
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show vlan dhcp-config

```
show {vlan} vlan_name dhcp-config
```

Description

Displays the DHCP server's configuration for the specified VLAN.

Syntax Description

<code>vlan_name</code>	Specifies the VLAN of the DHCP server of interest.
------------------------	--



Default

N/A.

Usage Guidelines

None.

Example

The following command displays the configuration of the DHCP server for the VLAN test:

```
show vlan test dhcp-config
```

The following is sample output from this command:

```

      DHCP Address Range           : 10.10.10.100-
>10.10.10.200
      Netlogin Lease Timer         : Not configured
(Default = 10 seconds)
      DHCP Lease Timer             : Not configured
(Default = 7200 seconds)
      Primary DNS Server           : 1.1.1.1
      Secondary DNS Server         : 2.2.2.2
      Ports DHCP Enabled           : 23

```

History

This command was first available in ExtremeXOS 11.0.

The output is modified to show primary and secondary DNS servers in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

show vlan security

```
show {vlan} vlan_name security
```

Description

Displays the MAC limit-learning and lock-learning information for the specified VLAN.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
------------------	------------------------



Default

N/A.

Usage Guidelines

The switch displays the following information:

- **Port**—Indicates the port on which the MAC address has been learned.
- **Limit**—Indicates that there is either a limited or unlimited amount of learned entries.
- **State**—Indicates that the current FDB entries for the port are permanent, no additional entries are learned, or that the port allows unlimited, dynamic learning.
- **Learned**—Specifies the number of learned entries.
- **Blackholed**—Specifies the number of blackholed entries.
- **Locked**—Specifies the number of locked entries.

Example

The following command displays the security setting of the DHCP server for the VLAN corp:

```
show vlan blue security
```

The following is sample output from this command:

Port	Limit	State	Learned	Blackholed	Locked
24	Unlimited	Unlocked	0	0	0

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

ssh2

```
ssh2 {cipher [3des | blowfish]} {port portnum} {compression [on | off]} {user
username} {username} [host | ipaddress] {remote command} {vr vr_name}
```


Description

Initiates an SSH2 client session to a remote SSH2 server.



Syntax Description

3des	Specifies that the 3des cipher should be used for encryption. This is the default.
blowfish	Specifies that the blowfish cipher should be used for encryption.
<i>portnum</i>	Specifies the TCP port number to be used for communicating with the SSH2 client. The default is port 22.
on	Specifies that the data is to be compressed.
off	Specifies that compression is not to be used. This is the default.
<i>username</i>	Specifies a login name for the remote host, as an alternate to the username@host parameter. Can be omitted if it is the same as the username on the switch.
<i>host</i>	Specifies the name of the remote host.
<i>ipaddress</i>	Specifies the IP address of the remote host.
<i>remote command</i>	Specifies a command to be passed to the remote system for execution. The switch does not support remote commands. The option is only valid if the remote system is a system, such as a UNIX workstation, that accepts remote commands.
<i>vr_name</i>	Specifies the virtual router. The default virtual router is VR-Mgmt.



Note
User-created VRs are supported only on the platforms listed for this feature in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

Default

The default settings for SSH2 parameters are as follows:

- cipher—3des encryption.
- port—22.
- compression—off.
- vr_name—VR-Mgmt.

Usage Guidelines

You must be running the SSH2 module (ssh.xmod), which is under Export Control, in order to use the SSH2 client command.

SSH2 does not need to be enabled on the switch in order to use this command.

Typically, this command is used to establish a secure session to a remote switch. You are prompted for your password. Once you have logged in successfully, all ExtremeXOS command you enter are executed on the remote switch. When you terminate the remote session, commands will then resume being executed on the original switch.

Host Name, User Name, and Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for host names and remote IP addresses.



When specifying a host name, user name, or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z).
- Numerals (0-9).
- Period (.).
- Dash (-) Permitted for host and user names.
- Underscore (_) Permitted for host and user names.
- Colon (:) Permitted for host names and remote IP addresses.
- At symbol (@) Permitted only for user names.

When naming the host, creating a user name, or configuring the IP address, remember the requirements listed above.

Remote Filename Character Restrictions

This section provides information about the characters supported by the switch for remote filenames.

When specifying a remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z).
- Numerals (0-9).
- Period (.).
- Dash (-).
- Underscore (_).
- Slash (/).

When naming a remote file, remember the requirements listed above.

Example

The following command establishes an SSH2 session on switch engineering1:

```
ssh2 admin@engineering1
```

The following command establishes an SSH2 session with the switch named BlackDiamond8810 over TCP port 2050 with compression enabled:

```
ssh2 compression on port 2050 admin@BlackDiamond8810
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms with the SSH2 module installed.



unconfigure ip-security dhcp-snooping information check

unconfigure ip-security dhcp-snooping information check

Description

Disables the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82) checking in the server-originated packets.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command disables the checking of the server-originated packets for the presence of option 82 so the packets will be forwarded normally.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

unconfigure ip-security dhcp-snooping information circuit-id port-information ports

unconfigure ip-security dhcp-snooping information circuit-id port-information ports [*port_list* | **all**]

Description

Unconfigures the port information portion of the circuit ID.

Syntax Description

<i>port_list</i>	Specifies the port(s) for which port information of the circuit-ID is unconfigured.
all	Specifies all ports.



Default

The default is all.

Usage Guidelines

This command unconfigures the port information portion of the circuit ID string for the indicated ports thereby restoring it to the default (ifIndex value).

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

unconfigure ip-security dhcp-snooping information circuit-id vlan-information

```
unconfigure ip-security dhcp-snooping information circuit-id vlan-information
{vlan} [vlan_name | all]
```

Description

Unconfigures the VLAN info portion of the circuit ID of a VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN for which VLAN information of the circuit-ID is unconfigured.
all	Specifies all VLANs.

Default

The default is all.

Usage Guidelines

This command unconfigures the VLAN information portion of the circuit ID of a VLAN, restoring it to the default.

History

This command was first available in ExtremeXOS 12.1.



Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

unconfigure ip-security dhcp-snooping information option

unconfigure ip-security dhcp-snooping information option

Description

Disables the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82).

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command disables the DHCP relay agent option (option 82), which is inserted into client-originated DHCP packets before they are forwarded to the server.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

unconfigure ip-security dhcp-snooping information policy

unconfigure ip-security dhcp-snooping information policy

Description

Unconfigures the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82) policy.

Syntax Description

This command has no arguments or variables.



Default

N/A.

Usage Guidelines

This command unconfigures the DHCP relay agent option information policy to the default value of replace.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

unconfigure radius

```
unconfigure radius {mgmt-access | netlogin} {server [primary | secondary]}
```

Description

Unconfigures the RADIUS client configuration.

Syntax Description

mgmt-access	Specifies the switch management RADIUS authentication server.
netlogin	Specifies the network login RADIUS authentication server.
primary	Unconfigures the primary RADIUS server.
secondary	Unconfigures the secondary RADIUS server.

Default

Unconfigures both primary and secondary servers for management and network login.

Usage Guidelines

If you do not specify any keywords, this command unconfigures both the primary and secondary servers for management and network login.

The following list describes the available keywords:

- **mgmt-access**—Use this keyword to unconfigure only the server(s) for management functions.
- **netlogin**—Use this keyword to unconfigure only the server(s) for network login.



- **primary**—Use this keyword to specify only the primary RADIUS sever.
- **secondary**—Use this keyword to specify only the secondary RADIUS server.

Example

The following command unconfigures the secondary RADIUS server settings for both management and network login:

```
unconfigure radius server secondary
```

The following command unconfigures the secondary RADIUS server settings for only network login:

```
unconfigure radius netlogin server secondary
```

The following command unconfigures all RADIUS server settings for only management functions:

```
unconfigure radius mgmt-access
```

History

This command was first available in ExtremeXOS 10.1.

The **mgmt-access** and **netlogin** keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

unconfigure radius-accounting

```
unconfigure radius-accounting {mgmt-access | netlogin} {server [primary | secondary]}
```

Description

Unconfigures the RADIUS accounting server configuration.

Syntax Description

mgmt-access	Specifies the switch management RADIUS accounting server.
netlogin	Specifies the network login RADIUS accounting server.
primary	Unconfigures the primary RADIUS accounting server.
secondary	Unconfigures the secondary RADIUS accounting server.



Default

Unconfigures both the primary and secondary accounting servers for management and network login.

Usage Guidelines

If you do not specify any keywords, this command unconfigures both the primary and secondary accounting servers for management and network login.

The following list describes the available keywords:

- `mgmt-access`—Use this keyword to unconfigure only the accounting server(s) for management functions.
- `netlogin`—Use this keyword to unconfigure only the accounting server(s) for network login.
- `primary`—Use this keyword to specify only the primary RADIUS accounting sever.
- `secondary`—Use this keyword to specify only the secondary RADIUS accounting server.

Example

The following command unconfigures the secondary RADIUS accounting server settings for both management and network login:

```
unconfigure radius-accounting server secondary
```

The following command unconfigures the secondary RADIUS accounting server settings for only network login:

```
unconfigure radius-accounting netlogin server secondary
```

The following command unconfigures all RADIUS accounting server settings for only management functions:

```
unconfigure radius-accounting mgmt-access
```

History

This command was first available in ExtremeXOS 10.1.

The `mgmt-access` and `netlogin` keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

unconfigure tacacs

```
unconfigure tacacs {server [primary | secondary]}
```



Description

Unconfigures the TACACS+ server configuration.

Syntax Description

primary	Unconfigures the primary TACACS+ server.
secondary	Unconfigures the secondary TACACS+ server.

Default

Unconfigures both the primary and secondary TACACS+ servers.

Usage Guidelines

None.

Example

The following command unconfigures all TACACS+ servers settings:

```
unconfigure tacacs
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

unconfigure tacacs-accounting

```
unconfigure tacacs-accounting {server [primary | secondary]}
```

Description

Unconfigures the TACACS+ accounting server configuration.

Syntax Description

primary	Unconfigures the primary TACACS+ accounting server.
secondary	Unconfigures the secondary TACACS+ accounting server.



Default

Unconfigures both the primary and secondary TACACS+ accounting servers.

Usage Guidelines

None.

Example

The following command unconfigures all TACACS+ accounting servers settings:

```
unconfigure tacacs-accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

unconfigure trusted-ports trust-for dhcp-server

```
unconfigure trusted-ports [ports | all] trust-for dhcp-server
```

Description

Unconfigures, disables one or more DHCP trusted ports.

Syntax Description

<i>ports</i>	Specifies one or more trusted ports.
all	Specifies all trusted ports.

Default

N/A.

Usage Guidelines

Use this command to disable one or more DHCP trusted ports.



Displaying DHCP Trusted Server Information

To display the DHCP snooping configuration settings, including DHCP trusted ports if configured, use the following command:

```
show ip-security dhcp-snooping {vlan} vlan_name
```

To display any violations that occur, including those on DHCP trusted ports if configured, use the following command:

```
show ip-security dhcp-snooping violations {vlan} vlan_name
```

Example

The following command unconfigures ports 2:2 and 2:3 as trusted ports:

```
unconfigure trusted-ports 2:2-2:3 trust-for dhcp-server
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 series modules, and Summit Family switches.

unconfigure vlan dhcp

```
unconfigure vlan vlan_name dhcp
```

Description

Unconfigure all the DHCP configuration information for the specified VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN on which to unconfigure DHCP.
------------------	--

Default

N/A.

Usage Guidelines

None.



Example

The following command unconfigures the DHCP server for the VLAN temporary:

```
unconfigure temporary dhcp
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

unconfigure vlan dhcp-address-range

```
unconfigure vlan vlan_name dhcp-address-range
```

Description

Unconfigure the DHCP address range information for the specified VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN on which to unconfigure DHCP.
------------------	--

Default

N/A.

Usage Guidelines

None.

Example

The following command unconfigures the DHCP address range for the VLAN temporary:

```
unconfigure temporary dhcp-address-range
```

History

This command was first available in ExtremeXOS 11.0.



Platform Availability

This command is available on all platforms.

unconfigure vlan dhcp-options

```
unconfigure {vlan} vlan_name dhcp-options {[ default-gateway | dns-server
{primary | secondary} | wins-server]}
```

Description

Unconfigure the DHCP option information for the specified VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN on which to unconfigure DHCP.
default-gateway	Specifies the router option.
dns-server	Specifies the Domain Name Server (DNS) option.
primary	Specifies the primary DNS option.
secondary	Specifies the secondary DNS option.
wins-server	Specifies the NetBIOS name server (NBNS) option.

Default

N/A.

Usage Guidelines

None.

Example

The following command unconfigures the DHCP options for the VLAN temporary:

```
unconfigure temporary dhcp-options
```

History

This command was first available in ExtremeXOS 11.0.

The primary and secondary DNS options were added in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.



upload dhcp-bindings

upload dhcp-bindings

Description

Upload the DHCP bindings immediately on demand.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command enables the functionality to allow you to upload DHCP bindings on demand.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.



27 CLEAR-Flow Commands

```
disable clear-flow
enable clear-flow
show clear-flow
show clear-flow acl-modified
show clear-flow rule
show clear-flow rule-all
show clear-flow rule-triggered
```

This chapter describes commands for:

- Enabling and disabling CLEAR-Flow.
- Displaying CLEAR-Flow rules.
- Displaying triggered CLEAR-Flow rules.

CLEAR-Flow is a broad framework for implementing security, monitoring, and anomaly detection in ExtremeXOS software. Instead of simply looking at the source and destination of traffic, CLEAR-Flow allows you to specify certain types of traffic that require more attention. Once certain criteria for this traffic are met, the switch can either take an immediate, pre-determined action, or send a copy of the traffic off-switch for analysis.

CLEAR-Flow is an extension to Access Control Lists (ACLs). You create ACL policy rules to count packets of interest. CLEAR-Flow rules are added to the policy to monitor these ACL counter statistics. The CLEAR-Flow agent monitors the counters for the situations of interest to you and your network. You can monitor the cumulative value of a counter, the change to a counter over a sampling interval, the ratio of two counters, or even the ratio of the changes of two counters over an interval. For example, you can monitor the ratio between TCP SYN and TCP packets. An abnormally large ratio may indicate a SYN attack.

If the rule conditions are met, the CLEAR-Flow actions configured in the rule are executed. The switch can respond by installing an ACL that will block or rate limit the traffic, executing a set of CLI commands, or sending a report using a SNMP trap or EMS log message.

Note



CLEAR-Flow is available on platforms with an Edge, Advanced Edge, or Core license. These include BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm- series modules, E4G-200 and E4G-400 switches, and Summit X440, X460, X480, and X670 series switches. For more license information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*. CLEAR-Flow is supported only on ingress. Any limitations on a given platform for a regular ACL also hold true for CLEAR-Flow.

disable clear-flow

disable clear-flow

Description

Disable the CLEAR-Flow agent.

Syntax Description

This command has no arguments or variables.

Default

CLEAR-Flow is disabled by default.

Usage Guidelines

When the CLEAR-Flow agent is disabled, sampling stops and the and all rules are left in the current state. It will not reset actions that were taken while CLEAR-Flow was enabled.

Example

The following example disables CLEAR-Flow on the switch:

```
disable clear-flow
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available only on BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules, E4G-200 and E4G-400 switches, and Summit X440, X460, X480, and X670 series switches.

enable clear-flow

enable clear-flow

Description

Enable the CLEAR-Flow agent.



Syntax Description

This command has no arguments or variables.

Default

CLEAR-Flow is disabled by default.

Usage Guidelines

When the CLEAR-Flow agent is enabled, sampling begins and actions are taken based on the CLEAR-Flow rules that are configured on the switch.

Example

The following example enables CLEAR-Flow on the switch:

```
enable clear-flow
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available only on BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules, E4G-200 and E4G-400 switches, and Summit X440, X460, X480, and X670 series switches.

show clear-flow

```
show clear-flow
```

Description

Displays the status of the CLEAR-Flow agent, any CLEAR-Flow policies on each interface, and the number of CLEAR-Flow rules.

Syntax Description

This command has no arguments or variables.

Default

N/A.



Usage Guidelines

None.

Example

The following display shows output for the command `show clear-flow`:

```
clear-flow: Enabled
VLAN      Port    Policy Name          No. of CF Rules
=====
*         2:1    CFexample            6
*         2:26   CFexample            6
*         2:40   CFexample            6
Default   *       CFexample            6
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available only on BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules, E4G-200 and E4G-400 switches, and Summit X440, X460, X480, and X670 series switches.

show clear-flow acl-modified

```
show clear-flow acl-modified
```

Description

Displays the ACLs modified by CLEAR-Flow actions.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the ACLs that have been modified by CLEAR-Flow rules that have been triggered.



Example

The following display shows output for the command `show clear-flow acl-modified`:

```

Policy Name      Vlan Name      Port Rule Name      Default ACL      CF Added
Actions         Actions
=====
==
clearFlow      *              2:26 acl-rule-4        D                QP1
clearFlow      *              2:26 acl-rule-3        D                D
clearFlow      *              2:26 acl-rule-2        D                M
clearFlow      *              2:26 acl-rule-1        D                P
clearFlow      Default        *      acl-rule-4        D                QP1
clearFlow      Default        *      acl-rule-3        D                D
clearFlow      Default        *      acl-rule-2        D                M
clearFlow      Default        *      acl-rule-1        D                P
=====
==
Total Entries: 8
Notation:
P - Permit, D- Deny, M - mirror enabled, m - mirror disabled

```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available only on BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules, E4G-200 and E4G-400 switches, and Summit X440, X460, X480, and X670 series switches.

show clear-flow rule

```
show clear-flow [port port | vlan vlanname | any] {rule rulename} {detail}
```

Description

Displays the CLEAR-Flow rules, values, and configuration.

Syntax Description

<i>port</i>	Specifies the port.
<i>vlanname</i>	Specifies the VLAN.
any	Specifies the wildcard interface.
<i>rulename</i>	Specifies the entry name of a CLEAR-Flow rule.
detail	Display detailed information.



Default

N/A.

Usage Guidelines

If you issue the command without the rule keyword, all of the CLEAR-Flow rules for the policy on the port, VLAN, and the wildcard are displayed. If you specify a rule name, only that rule will be displayed. The detail keyword displays detailed information about the rule.

Example

The following display shows output for the command `show clear-flow port 2:6`:

```

Rule Name      Type Period  Last          Rel  Threshold  TCNT
NumAction
Value          Oper          If  Else
=====
=
rule-count     CN 30        16892762     >   100        7   3   3
rule-delta     DT 30        7762385      >  1000       1   4   3
rule-delta-2   DT 5         0            >  1000       0   4   3
rule-delta-ratio DR 30        0            >   20        0   2   0
rule-ratio     RT 30        0            >   10        0   3   3
rule-ratio-2   RT 5         0            >   10        0   3   3
=====
=
Total Entries: 6
Notation:
Threshold Type: CN - Count, DT - Delta, RT - Ratio, DR - DeltaRatio
TCNT - Number of times expression is continuously evaluated to be true

```

The following display shows output for the command `show clear-flow port 2:6 rule rule-delta detail`:

```

Rule Name: rule-delta      Sample Period: 30      Hysteresis: 20
=====
==
DELTA(counter1) = 0 sampled at 24 seconds ago
Expression evaluation is currently FALSE
if (DELTA(counter1) > 1000) then {
PERMIT:   Allow ACL rule acl-rule-3
SYSLOG:   [INFO] [Delta $ruleValue counter $counter1 offset $counterOffset1
delTime $deltaTime delay $delayTime]
CLI:      [disable port $port]
QOS:      Set rule acl-rule-4 qos value to QP6
} else {
DENY:     Block ACL rule acl-rule-3
QOS:      Set rule acl-rule-4 qos value to QP1
CLI:      [enable port $port]
}

```



History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available only on BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules, E4G-200 and E4G-400 switches, and Summit X440, X460, X480, and X670 series switches.

show clear-flow rule-all

show clear-flow rule-all

Description

Displays all the CLEAR-Flow rules on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following display shows output for the command show clear-flow rule-all:

```

Policy Name      Vlan Name      Port Rule Name  Last Value OP Threshold  TCNT
Sec
=====
==
clearFlow      *              2:1 rule-count 1              > 100          0    11
clearFlow      *              2:1 rule-delta 1              > 1000        0    11
clearFlow      *              2:1 rule-delta 0              > 1000        0    4
clearFlow      *              2:1 rule-delta 0              > 20          0    11
clearFlow      *              2:1 rule-ratio 0              > 10          0    11
clearFlow      *              2:1 rule-ratio 0              > 10          0    4
clearFlow      *              2:26 rule-count 9030635       > 100          1    10
clearFlow      *              2:26 rule-delta 9030635       > 1000        1    10
clearFlow      *              2:26 rule-delta 0              > 1000        0    4
clearFlow      *              2:26 rule-delta 0              > 20          0    10
clearFlow      *              2:26 rule-ratio 0              > 10          0    10

```



```

clearFlow      *           2:26 rule-ratio 0           > 10           0           4
clearFlow      Default    *     rule-count 36666439 > 100         1           10
clearFlow      Default    *     rule-delta 36666439 > 1000        1           10
clearFlow      Default    *     rule-delta 0         > 1000        0           4
clearFlow      Default    *     rule-delta 0         > 20          0           10
clearFlow      Default    *     rule-ratio 0         > 10          0           10
clearFlow      Default    *     rule-ratio 0         > 10          0           4
=====
==
Total Entries: 18
Notation:
TCNT - Number of times expression is continuously evaluated to be true
Sec - Number of seconds elapsed from last sampled data

```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available only on BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules, E4G-200 and E4G-400 switches, and Summit X440, X460, X480, and X670 series switches.

show clear-flow rule-triggered

```
show clear-flow rule-triggered
```

Description

Displays the triggered CLEAR-Flow rules.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the rules that have been triggered; in other words, the rule threshold has been reached.



Example

The following display shows output for the command `show clear-flow rule-triggered`:

```

Policy Name      Vlan Name      Port Rule Name  Last Value OP Threshold  TCNT
Sec
=====
==
clearFlow      *              2:26 rule-count  9130377   >  100         2    25
clearFlow      *              2:26 rule-delta  99742     >  1000        2    25
clearFlow      Default        *    rule-count  37069465  >  100         2    25
clearFlow      Default        *    rule-delta  403026    >  1000        2    25
=====
==
Total Entries:  4
Notation:
TCNT - Number of times expression is continuously evaluated to be true
Sec - Number of seconds elapsed from last sampled data

```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available only on BlackDiamond X8 series switches, BlackDiamond 8000 c-, e-, xl-, and xm-series modules, E4G-200 and E4G-400 switches, and Summit X440, X460, X480, and X670 series switches.



28 EAPS Commands

```
clear eaps counters
configure eaps add control vlan
configure eaps add protected vlan
configure eaps cfm
configure eaps config-warnings off
configure eaps config-warnings on
configure eaps delete control vlan
configure eaps delete protected vlan
configure eaps failtime
configure eaps failtime expiry-action
configure eaps fast-convergence
configure eaps hello-pdu-egress
configure eaps hellotime
configure eaps mode
configure eaps multicast add-ring-ports
configure eaps multicast send-igmp-query
configure eaps multicast temporary-flooding
configure eaps multicast temporary-flooding duration
configure eaps name
configure eaps port
configure eaps priority
configure eaps shared-port common-path-timers
configure eaps shared-port link-id
configure eaps shared-port mode
configure eaps shared-port segment-timers expiry-action
configure eaps shared-port segment-timers health-interval
configure eaps shared-port segment-timers timeout
configure forwarding L2-protocol fast-convergence
configure ip-arp fast-convergence
create eaps
create eaps shared-port
delete eaps
delete eaps shared-port
disable eaps
enable eaps
show eaps
show eaps cfm groups
```

```
show eaps counters
show eaps counters shared-port
show eaps shared-port
show eaps shared-port neighbor-info
show vlan eaps
unconfigure eaps shared-port link-id
unconfigure eaps shared-port mode
unconfigure eaps port
```

This chapter describes commands for completing the following Ethernet Automatic Protection Switching (EAPS) tasks:

- Configuring EAPS
- Displaying EAPS information

For an introduction to EAPS, see the *ExtremeXOS Concepts Guide*.

clear eaps counters

clear eaps counters

Description

Clears, resets the counters gathered by EAPS for all of the EAPS domains and any EAPS shared ports configured on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to clear, reset the EAPS counters.

The counters continue to increment until you clear the information. By clearing the counters, you can see fresh statistics for the time period you are monitoring.

To display information about the EAPS counters, use the following commands:

- `show eaps counters` —This command displays summary EAPS counter information.
- `show eaps counters shared-port` —If configured for EAPS shared ports, this command displays summary EAPS shared port counter information.



Example

The following command clears, resets all of the counters for EAPS:

```
clear eaps counters
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure eaps add control vlan

```
configure eaps name add control {vlan} vlan_name
```

Description

Adds the specified control VLAN to the specified EAPS domain.

Syntax Description

<i>name</i>	Specifies the name of an EAPS domain.
<i>vlan_name</i>	Specifies the name of the control VLAN.

Default

N/A.

Usage Guidelines

You must configure one control VLAN for each EAPS domain. The control VLAN is used only to send and receive EAPS messages.

The control VLAN must be configured as follows:

- The VLAN must NOT be assigned an IP address, to avoid loops in the network.
- Only ring ports can be added as members of the control VLAN.
- The ring ports of the control VLAN must be tagged.

A control VLAN cannot belong to more than one EAPS domain. When the EAPS domain is active, you cannot delete or modify the configuration of the control VLAN.



By default, EAPS protocol data units (PDUs) are automatically assigned to QoS profile QP8. This ensures that the control VLAN messages reach their intended destinations. You do not need to configure a QoS profile for the control VLAN.

The VLAN must already exist before you can add it as a control VLAN. If you attempt to add a VLAN that does not exist, the switch displays a message similar to the following:

```
* Switch.8 # configure eaps megtest add control foo^%% Invalid input
detected at '^' marker.
```

To create the VLAN, use the `create vlan` command.

Example

The following command adds the control VLAN keys to the EAPS domain `eaps_1`.

```
configure eapseaps_1add control vlan keys
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure eaps add protected vlan

```
configure eaps name add protected {vlan} vlan_name
```

Description

Adds the specified protected VLAN to the specified EAPS domain.

Syntax Description

<i>name</i>	Specifies the name of an EAPS domain.
<i>vlan_name</i>	Specifies the name of the protected VLAN.

Default

N/A.



Usage Guidelines

You must configure one or more protected VLANs for each EAPS domain. The protected VLANs are the data-carrying VLANs.

A protected VLAN can be added to one or more EAPS domains.

When you configure a protected VLAN, the ring ports of the protected VLAN must be tagged (except in the case of the default VLAN). As long as the ring is complete, the master node blocks the protected VLANs on its secondary port.

The VLAN must already exist before you can add it as a protected VLAN. If you attempt to add a VLAN that does not exist, the switch displays a message similar to the following:

```
* Switch.5 # configure eaps megtest add protected foo^%% Invalid input
detected at '^' marker.
```

To create the VLAN, use the [create vlan](#) command.

Example

The following command adds the protected VLAN orchid to the EAPS domain eaps_1:

```
configure eapseaps_1add protected vlan orchid
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure eaps cfm

```
configure eaps cfm [add | delete] group group_name
```

Description

Notifies the CFM that EAPs is interested in notifications for the specified MEP and RMEP pair.

Syntax Description

cfm	Connectivity Fault Management.
add	Add a MEP group.
delete	Delete a MEP group.
group <i>group_name</i>	MEP group to bind.



Default

N/A.

Usage Guidelines

This command notifies CFM that EAPs is interested in notifications for this MEP and RMEP pair. This MEP should already be bound to a physical port, so when notification is received, EAPS associates that notification with a ring-port failure.

Example

The following command deletes the control VLAN keys from the EAPS domain eaps_1:

```
configure eaps cfm add
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on all EXOS platforms; however, not all platforms support hardware-based CFM. Platforms with no hardware-based CFM support are limited to software-based CFM transmit intervals of 100ms., or higher. Hardware-based intervals can go as low as 3.3ms.

Currently, only the x460 and E4G platforms support hardware-based CFM.

configure eaps config-warnings off

```
configure eaps config-warnings off
```

Description

Disables the loop protection warning messages displayed when configuring specific EAPS parameters.

Syntax Description

This command has no arguments or variables.

Default

By default, loop protection warnings are enabled and displayed when configuring specific EAPS parameters.



Usage Guidelines

This is a global EAPS command. You configure the warning message display on a per switch basis, not per EAPS domain.

When configuring the following EAPS parameters, the switch displays loop protection warning messages:

- Adding EAPS primary or secondary ring ports to a VLAN
- Deleting a protected VLAN
- Disabling the global EAPS setting on the switch
- Disabling an EAPS domain
- Configuring an EAPS domain as a transit node
- Unconfiguring EAPS primary or secondary ring ports from an EAPS domain

We recommend that you keep the loop protection warning messages enabled. If you have considerable knowledge and experience with EAPS, you might find the EAPS loop protection warning messages unnecessary. For example, if you use a script to configure your EAPS settings, disabling the warning messages allows you to configure EAPS without replying to each interactive yes/no question.

To confirm the setting on the switch, use the following command:

```
show eaps {eapsDomain} {detail}
```

Example

The following command disables the loop protection warning messages:

```
configure eaps config-warnings off
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure eaps config-warnings on

```
configure eaps config-warnings on
```

Description

Enables the loop protection warning messages displayed when configuring specific EAPS parameters.



Syntax Description

This command has no arguments or variables.

Default

By default, loop protection warnings are enabled and displayed when configuring specific EAPS parameters.

Usage Guidelines

This is a global EAPS command. You configure the warning message display on a per switch basis, not per EAPS domain.

When configuring the following EAPS parameters, the switch displays loop protection warning messages:

- Adding EAPS primary or secondary ring ports to a VLAN
- Deleting a protected VLAN
- Disabling the global EAPS setting on the switch
- Disabling an EAPS domain
- Configuring an EAPS domain as a transit node
- Unconfiguring EAPS primary or secondary ring ports from an EAPS domain

We recommend that you keep the loop protection warning messages enabled.

Example

The following command enables the loop protection warning messages:

```
configure eaps config-warnings on
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure eaps delete control vlan

```
configure eaps name delete control {vlan} vlan_name
```

Description

Deletes the specified control VLAN from the specified EAPS domain.



Syntax Description

<i>name</i>	Specifies the name of an EAPS domain.
<i>vlan_name</i>	Specifies the name of the control VLAN.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the control VLAN keys from the EAPS domain eaps_1:

```
configure eapseaps_1 delete control vlan keys
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure eaps delete protected vlan

```
configure eaps name delete protected {vlan} vlan_name
```

Description

Deletes the specified protected VLAN from the specified EAPS domain.

Syntax Description

<i>name</i>	Specifies the name of an EAPS domain.
<i>vlan_name</i>	Specifies the name of the protected VLAN.

Default

N/A.



Usage Guidelines

To prevent loops in the network, you must delete the ring ports (the primary and the secondary ports) from the protected VLAN **before** deleting the protected VLAN from the EAPS domain. Failure to do so can cause a loop in the network.

The switch displays by default a warning message and prompts you to delete the VLAN from the EAPS domain. When prompted, do one of the following:

- Enter `y` to delete the VLAN from the specified EAPS domain.
- Enter `n` or press [Return] to cancel this action.

If you have considerable knowledge and experience with EAPS, you might find the EAPS loop protection warning messages unnecessary. For more information, see the `configure eaps config-warnings off` command.

Useful show Commands

Use the following show commands to display information about your EAPS domain, including protected VLANs and primary and secondary ports:

- `show vlan` —This command displays summary information for all of the VLANs on the device. If the VLAN is a protected VLAN, the P flag appears in the flag column. To see more detailed information about the protected VLAN, use the following command: `show vlanvlan_name` .
- `show eaps` —This command displays summary EAPS domain information, including the name of the domain and the primary and secondary ports. To see more detailed information, including the name of the protected VLAN and the primary and secondary ports, use the `show eapseapsDomain` command.
- `show vlan eaps` —This command displays whether the VLAN is a control or partner VLAN for an EAPS domain. This command also displays if the VLAN is not a member of any EAPS domain.

Example

The following command deletes the protected VLAN orchid from the EAPS domain eaps_1:

```
configure eapseaps_1delete protected vlan orchid
```

The switch displays the following warning message and prompts you to confirm this action:

```
WARNING: Make sure EAPS ring-ports are deleted from the VLAN first.
Otherwise deleting the VLAN from the EAPS domain could cause a loop in
the network! Are you sure you want to remove the VLAN before deleting
EAPS ring-ports.? (y/n)
```

Enter `y` to delete the VLAN from the specified EAPS domain. Enter `n` to cancel this action.

History

This command was first available in ExtremeXOS 11.0.

The interactive messages were added in ExtremeXOS 11.4.



Platform Availability

This command is available on all platforms.

configure eaps failtime

```
configure eaps name failtime seconds milliseconds
```

Description

Configures the period after which the master node declares a failure if no hello PDUs are received.

Syntax Description

<i>name</i>	Specifies the name of an EAPS domain.
<i>seconds</i>	Specifies the number of seconds the master node waits before the failtimer expires. Default is 3 seconds, and the range is 0 to 300 seconds.
<i>milliseconds</i>	Specifies the number of milliseconds to wait before the failtimer expires. The range is 300 to 999 milliseconds.

Default

The default is 3 seconds.

Usage Guidelines

Use the failtime keyword and its associated seconds parameter to specify the amount of time the master node waits before the failtimer expires. The failtime period (seconds plus milliseconds) must be set greater than the configured value for hellotime. The default value is three seconds.

Increasing the failtime value reduces the likelihood of false failure detections caused by network congestion.



Note

You configure the action taken when the failtimer expires by using the `configure eaps failtime expiry-action` command.

In ExtremeXOS 11.0, the failtimer range was 2 to 60 seconds.

Example

The following command configures the failtimer value for the EAPS domain eaps_1 to 15 seconds:

```
configure eapseaps_1failtime15 0
```



The following command configures the failtimer value for the EAPS domain `eaps_2` to 300 milliseconds:

```
configure eapseaps_2failtime0 300
```

History

This command was first available in ExtremeXOS 11.0.

The range for the failtimer was changed to 2 to 300 seconds in ExtremeXOS 11.1. The default value for the failtimer remains unchanged.

The milliseconds parameter was added in ExtremeXOS 12.4.2.

Platform Availability

This command is available on all platforms.

configure eaps failtime expiry-action

```
configure eaps name failtime expiry-action [open-secondary-port | send-alert]
```

Description

Configures the action taken when the failtimer expires.

Syntax Description

<i>name</i>	Specifies the name of an EAPS domain.
open-secondary-port	Specifies to open the secondary port when the failtimer expires.
send-alert	Specifies that a critical message is sent to the syslog when the failtimer expires.

Default

Default is `send-alert`.

Usage Guidelines

By default the action is to send an alert if the failtimer expires. Instead of going into a Failed state, the master node remains in a Complete or Init state, maintains the secondary port blocking, and writes a critical error message to syslog warning the user that there is a fault in the ring. An SNMP trap is also sent.



If the EAPS ring contains non-EAPS devices, you must use the open-secondary-port parameter.

Note



Use caution when setting the failtimer expiry action to open-secondary port. Using this configuration, if the master node loses three consecutive hello PDUs, the failtimer expires—but there might not be a break in the ring. Opening the secondary port in this situation creates a loop.

Example

The following command configures the failtimer expiry action for EAPS domain eaps_1:

```
configure eapseaps_1 failtimeexpiry-action open-secondary-port
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure eaps fast-convergence

```
configure eaps fast-convergence[off | on]
```

Description

Enables EAPS to converge more quickly.

Syntax Description

off	Turns fast-convergence off. Default is off.
on	Turns fast-convergence on.

Default

Default is off.

Usage Guidelines

This command acts on the switch, not per domain.



In certain environments to keep packet loss to a minimum when the ring is broken, configure EAPS with fast-convergence turned on. If fast convergence is turned on, you can view the configuration with the `show eaps` command.



Note

If fast-convergence is turned on, the link filters on all EAPS ring ports are turned off. This can result problems if the port's hardware encountered a problem and started "flapping" between link-up/link-down states.

Example

The following command configures fast convergence for all of the EAPS domains on the switch:

```
configure eaps fast-convergence on
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure eaps hello-pdu-egress

```
configure eaps name hello-pdu-egress [primary-port | secondary-port]
```

Description

Configures the port through which a master node sends EAPS hello PDUs.

Syntax Description

<i>name</i>	Specifies the name of an EAPS domain.
-------------	---------------------------------------

Default

Default is the primary port.



Usage Guidelines

This command is provided for special network topologies that use spatial reuse and require that all EAPS hello PDUs travel in the same direction on the ring.



Note

We recommend the default (primary-port) configuration for this command.

Example

The following command configures the master switch to send EAPS hello packets from the secondary port:

```
configure eaps "domain12" hello-pdu-egress secondary-port
```

History

This command was first available in ExtremeXOS 12.4.2.

Platform Availability

This command is available on all platforms.

configure eaps hellotime

```
configure eaps name hellotime seconds milliseconds
```

Description

Configures the period at which the master node sends EAPS hello PDUs to verify ring connectivity.

Syntax Description

<i>name</i>	Specifies the name of an EAPS domain.
<i>seconds</i>	Specifies the number of seconds to wait between transmission of hello PDUs on the control VLAN. The range is 0 to 15 seconds.
<i>milliseconds</i>	Specifies the number of milliseconds to wait between transmission of hello PDUs on the control VLAN. The range is 0 to 999 milliseconds.

Default

Default is 1 second.



Usage Guidelines

Use the `hellotime` keyword and its associated parameters to specify the amount of time the master node waits between transmissions of hello PDUs on the control VLAN. Increasing the `hellotime` value results in a reduced load on the processor and less traffic on the EAPS ring.



Note

The hello PDU timer value must be smaller than the fail timer value to prevent false failure detection. If you change the hello PDU timer, verify that the fail timer value remains larger.

This command applies only to the master node. If you configure the hello PDU timer for a transit node, the timer value is ignored. If you later reconfigure that transit node as the master node, the master node uses the configured hello PDU timer value.

In ExtremeXOS 11.0, the range is 1 to 15 seconds. If you are running ExtremeXOS 11.0 with the hello timer value greater than 15 seconds and you upgrade to ExtremeXOS 11.1 or later, you must modify the hello timer to be within the 1 to 15 seconds range.

Example

The following command configures the `hellotime` value for the EAPS domain `eaps_1` to 300 milliseconds:

```
configure eapseaps_1hellotime0 300
```

History

This command was first available in ExtremeXOS 11.0.

The range for the hello timer was changed to 1 to 15 seconds in ExtremeXOS 11.1. The default value for the hello timer remains unchanged.

Support for a specific number of milliseconds was added in ExtremeXOS 12.4.2.

Platform Availability

This command is available on all platforms.

configure eaps mode

```
configure eaps name mode [master | transit]
```

Description

Configures the switch as either the EAPS master node or as an EAPS transit node for the specified domain.



Syntax Description

<i>name</i>	Specifies the name of an EAPS domain.
master	Specifies that this switch should be the master node for the named EAPS domain.
transit	Specifies that this switch should be the transit node for the named EAPS domain.

Default

N/A.

Usage Guidelines

One node (or switch) on the ring must be configured as the master node for the specified domain; all other nodes (or switches) on the ring are configured as transit nodes for the same domain.

If you configure a switch to be a transit node for an EAPS domain, the switch displays by default messages to:

- Remind you to configure a master node in the EAPS domain.
- Notify you that changing a master node to a transit node might cause a loop in the network. If you have not assigned a new master node before changing the current master node to a transit node, you might cause a loop in the network.

When prompted, do one of the following:

- Enter `y` to identify the switch as a transit node.
- Enter `n` or press [Return] to cancel this action.

If you have considerable knowledge and experience with EAPS, you might find the EAPS loop protection warning messages unnecessary. For more information, see the `configure eaps config-warnings off` command.

Example

The following command identifies this switch as the master node for the domain named `eaps_1`:

```
configure eapseaps_1mode master
```

The following command identifies this switch as a transit node for the domain named `eaps_1`:

```
configure eapseaps_1mode transit
```

The switch displays the following warning message and prompts you to confirm this action:

```
WARNING: Make sure this specific EAPS domain has a Master node in the
ring. If you change this node from EAPS master to EAPS transit, you could
cause a loop in the network. Are you sure you want to change mode to
transit? (y/n)
```



Enter `y` to identify the switch as a transit node. Enter `n` to cancel this action.

History

This command was first available in ExtremeXOS 11.0.

The interactive messages were added in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure eaps multicast add-ring-ports

```
configure eaps multicast add-ring-ports [on | off]
```

Description

Configures the switch to add previously blocked ring ports to existing multicast groups when an EAPS topology change occurs.

Syntax Description

on	Enables the multicast add-ring-ports feature.
off	Disables the multicast add-ring-ports feature.

Default

Off.

Usage Guidelines

When this feature is set to on and an EAPS topology change occurs, multicast traffic is fastpath forwarded using the switch hardware during the topology transition. The on setting improves multicast forwarding performance during the transition.



Note

EAPS multicast flooding must be enabled before this feature will operate. For information on enabling EAPS multicast flooding, see the `configure eaps multicast temporary-flooding` command description.

When this feature is set to off and an EAPS topology change occurs, multicast traffic is slowpath forwarded using the CPU during the topology transition. The off setting reduces multicast forwarding performance during the transition.

For other methods of supporting multicast traffic during an EAPS topology change, see the descriptions for the following commands:



- `configure eaps multicast send-igmp-query`
- `configure eaps multicast temporary-flooding`

Example

The following command enables the add-ring-ports feature:

```
configure eaps multicast add-ring-ports on
```

History

This command was first available in ExtremeXOS 12.1.2.

Platform Availability

This command is available on all platforms.

configure eaps multicast send-igmp-query

```
configure eaps multicast send-igmp-query [on | off]
```

Description

Configures the switch to send IGMP query messages to all protected VLANs when an EAPS topology change occurs.

Syntax Description

on	Enables the multicast send-igmp-query feature.
off	Disables the multicast send-igmp-query feature.

Default

On.

Usage Guidelines

When this feature is set to on and an EAPS topology change occurs, the switch sends IGMP query messages to all protected VLANs. If the protected VLANs in the node detecting (and generating) the topology change do not have IP address, a query is generated with the source IP address set to the querier address in that VLAN.

In a EAPS ring with many protected VLANs, the many responses can impact switch performance. This is the default behavior and was the only method for supporting multicast traffic during EAPS topology changes prior to release 12.1.2.



When this feature is set to off and an EAPS topology change occurs, the switch does not automatically send IGMP queries to all protected VLANs during the topology transition. The off setting improves switch performance during the transition, but you should use one of the following commands to see that multicast traffic is supported during and after the topology change:

- `configure eaps multicast add-ring-ports`
- `configure eaps multicast temporary-flooding`

Example

The following command disables the send-igmp-query feature:

```
configure eaps multicast send-igmp-query off
```

History

This command was first available in ExtremeXOS 12.1.2.

Platform Availability

This command is available on all platforms.

configure eaps multicast temporary-flooding

```
configure eaps multicast temporary-flooding [on | off]
```

Description

Configures the switch to temporarily enable multicast flooding when an EAPS topology change occurs.

Syntax Description

on	Enables the multicast temporary-flooding feature.
off	Disables the multicast temporary-flooding feature.

Default

Off.

Usage Guidelines

When this feature is set to on and an EAPS topology change occurs, the switch temporarily enables multicast flooding to all protected VLANs for the duration specified by the following command:

```
configure eaps multicast temporary-flooding duration
```



If you change the configuration to off, topology changes that occur after this command do not result in temporary flooding. For example, if you change the configuration to off while flooding is in progress for a protected VLAN or set of protected VLANs (due to an EAPS topology change), the flooding continues for the configured duration period. New topology changes on the protected VLANs do not cause flooding.

When this feature is set to off and an EAPS topology change occurs, the switch does not enable flooding to all protected VLANs during the topology transition. The default switch response for multicast traffic during an EAPS topology change is that defined by the following command:

```
configure eaps multicast send-igmp-query
```

You can also use the following command to configure the switch response for multicast traffic during an EAPS topology change:

```
configure eaps multicast add-ring-ports
```

Example

The following command enables the temporary-flooding feature:

```
configure eaps multicast temporary-flooding on
```

History

This command was first available in ExtremeXOS 12.1.2.

Platform Availability

This command is available on all platforms.

configure eaps multicast temporary-flooding duration

```
configure eaps multicast temporary-flooding duration seconds
```

Description

Configures the duration for which the switch temporarily enables multicast flooding when an EAPS topology change occurs.

Syntax Description

<i>seconds</i>	Specifies the period (in seconds) for which the switch enables multicast flooding.
----------------	--

Default

15 seconds.



Usage Guidelines

The flooding duration configuration applies only when the temporary-flooding feature is enabled with the following command:

```
configure eaps multicast temporary-flooding
```

Example

The following command configures the temporary-flooding feature duration for 30 seconds:

```
configure eaps multicast temporary-flooding duration 30
```

History

This command was first available in ExtremeXOS 12.1.2.

Platform Availability

This command is available on all platforms.

configure eaps name

```
configure eaps old_name name new_name
```

Description

Renames an existing EAPS domain.

Syntax Description

<i>old_name</i>	Specifies the current name of an EAPS domain.
<i>new_name</i>	Specifies a new name for the EAPS domain.

Default

N/A.

Usage Guidelines

If you use the same name across categories (for example, STPD and EAPS names), we recommend that you specify the identifying keyword as well as the actual name. If you do not use the keyword, the system might return an error message.



Example

The following command renames EAPS domain eaps-1 to eaps-5:

```
configure eaps eaps-1 name eaps-5
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure eaps port

```
configure eaps name [primary | secondary] port ports
```

Description

Configures a node port as the primary or secondary port for the specified EAPS domain.

Syntax Description

<i>name</i>	Specifies the name of an EAPS domain.
primary	Specifies that the port is to be configured as the primary port.
secondary	Specifies that the port is to be configured as the secondary port.
<i>ports</i>	Specifies one port or slot and port.

Default

N/A.

Usage Guidelines

Each node on the ring connects through two ring ports. One port must be configured as the primary port; the other must be configured as the secondary port.

The primary and secondary ports have significance only on a master node. The health-check messages are sent out the primary port of the master node, and the master node blocks the protected VLANs on the secondary port.

The master node's secondary EAPS port cannot be configured on ports that are already configured as follows:

- Shared-port



- MLAG ISC port

There is no distinction between the primary and secondary ports on a transit node.

Beginning with ExtremeXOS 11.1, if you have a primary or secondary port that is a member of a load-shared group, you do not need to disable your EAPS domain and remove that ring port when modifying the load-shared group. For more information about configuring load sharing on your switch, see [Configuring Slots and Ports on a Switch](#) in the ExtremeXOS Concepts Guide.

For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

Messages Displayed when Adding EAPS Ring Ports to a VLAN

If you attempt to add EAPS ring ports to a VLAN that is not protected by EAPS, the switch prompts you by default to confirm this action. For example, if you use the `configure vlan vlan_name add ports port_list` command, and the ports that you are attempting to add to the VLAN are currently used by EAPS as either primary or secondary ring ports, the switch displays the following message:

```
Make sure <vlan_name> is protected by EAPS. Adding EAPS ring ports to a
VLAN could cause a loop in the network. Do you really want to add these
ports (y/n)
```

Enter y to add the ports to the VLAN. Enter n or press [Return] to cancel this action.

If you see this message, either configure the VLAN as an EAPS protected VLAN by using the `configure eaps add protected vlan` command or add ports that the EAPS domain does not use as primary or secondary ring ports.

If you have considerable knowledge and experience with EAPS, you might find the EAPS loop protection warning messages unnecessary. For more information, see the `configure eaps config-warnings off`.

Example

The following command adds port 1 of the module installed in slot 8 to the EAPS domain `eaps_1` as the primary port:

```
configure eapseaps_1primary port8:1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.



configure eaps priority

```
configure eaps name priority {high | normal}
```

Description

Configures an EAPS domain priority.

Syntax Description

<i>name</i>	Specifies the name of an EAPS domain.
-------------	---------------------------------------

Default

Normal.

Usage Guidelines

Extreme Networks recommends that no more than 200 protected VLANs be configured as high priority domains. Priority protection works best when the majority of protected VLANs are configured for normal priority and a relatively small percentage of the protected VLANs are configured as high priority domains.

When EAPS domains on two separate physical rings share a common link (shared-port configuration) and have one or more protected VLANs in common, the domains must be configured with the same domain priority.

When EAPS domain priority is configured on separate physical rings that are connected to the same switch, the priorities on each ring are serviced independently. For example, if there is a break on both Ring A and Ring B, the high priority domains on each ring are serviced before the lower priority domains. However, the switch does not attempt to process the high priority domains on Ring B before servicing the normal priority domains on Ring A.

For a high priority domain to get priority over normal priority domains, all switches in the EAPS domain must support high priority domains. If high priority domains are configured on a switch that is in a ring with one or more switches that do not support high priority domains (software releases before ExtremeXOS Release 12.5), the high priority domain operates as a normal priority domain.

Example

The following command configures the eaps_1 domain as a high priority domain:

```
configure eapseaps_1 priority high
```

History

This command was first available in ExtremeXOS 12.5.



Platform Availability

This command is available on all platforms.

configure eaps shared-port common-path-timers

```
configure eaps shared-port port common-path-timers {[health-interval | timeout]
seconds}
```

Description

Configures the common path health interval or timeout value.

Syntax Description

<i>port</i>	Specifies the port number of the common link port.
health-interval	Specifies the interval for health check messages on the common link.
timeout	Specifies the timeout value for the common link.
<i>seconds</i>	Specifies the amount of health interval, in seconds.

Default

N/A.

Usage Guidelines

This command allows you to configure the length of the common path health interval, in seconds, for a given port. The range is from 1 to 10 seconds.

Example

The following command configures a common-link health interval of 5 seconds on port 1:1.

```
configure eaps shared-port 1:1 common-path-timers health-interval 5
```

The following command configures a segment timeout of 10 seconds on port 1:1.

```
configure eaps shared-port 1:1 common-path-timers timeout 10
```

History

This command was first available in ExtremeXOS 12.1.



Platform Availability

This command is available on all platforms with the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure eaps shared-port link-id

```
configure eaps shared-port ports link-id id
```

Description

Configures the link ID of the shared port.

Syntax Description

<i>ports</i>	Specifies the port number of the common link port.
<i>id</i>	Specifies the link ID of the port. The link ID range is 1 to 65535.

Default

N/A.

Usage Guidelines

Each common link in the EAPS network must have a unique link ID. The controller and partner shared ports belonging to the same common link must have matching link IDs. No other instance in the network should have that link ID.

If you have multiple adjacent common links, we recommend that you configure the link IDs in ascending order of adjacency. For example, if you have an EAPS configuration with three adjacent common links, moving from left to right of the topology, configure the link IDs from the lowest to the highest value.

Example

The following command configures the EAPS shared port 1:1 to have a link ID of 1.

```
configure eaps shared-port 1:1 link-id 1
```

History

This command was first available in ExtremeXOS 11.1.



Platform Availability

This command is available on all platforms with the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure eaps shared-port mode

```
configure eaps shared-port ports mode controller | partner
```

Description

Configures the mode of the shared port.

Syntax Description

<i>ports</i>	Specifies the port number of the shared port.
<i>controller</i>	Specifies the controller mode. The controller is the end of the common link responsible for blocking ports when the common link fails thereby preventing the superloop.
<i>partner</i>	Specifies partner mode. The partner is responsible only for sending and receiving health-check messages.

Default

N/A.

Usage Guidelines

The shared port on one end of the common link must be configured to be the controller. This is the end responsible for blocking ports when the common link fails thereby preventing the superloop.

The shared port on the other end of the common link must be configured to be the partner. This end does not participate in any form of blocking. It is responsible only for sending and receiving health-check messages.

Example

The following command configures the shared port 1:1 to be the controller.

```
configure eaps shared-port 1:1 mode controller
```

History

This command was first available in ExtremeXOS 11.1.



Platform Availability

This command is available on all platforms with the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure eaps shared-port segment-timers expiry-action

```
configure eaps shared-port port segment-timers expiry-action [segment-down | send-alert]
```

Description

Configures the action taken when the segment timeout timer expires.

Syntax Description

<i>port</i>	Specifies the port number of the common link port.
segment-down	Marks the segment as DOWN if the segment timer expires. No link-status-query is sent to verify that links are down.
send-alert	If the segment timer expires, the switch keeps segments up, but sends a warning message to the log. The segment fail flag is set, an SNMP trap is sent, and a link-status-query is sent to verify if any links are down.

Default

Default is send-alert.

Usage Guidelines

By default, the action is to send an alert if the segment timeout timer expires. Instead of the segment going into a failed state and being marked as down, the segment remains in a segment up state with the failed flag set. The switch writes a critical error message to the syslog warning the user that there is a fault in the segment. An SNMP trap is also sent.

Note



Use caution when setting the segment-timeout expiry action to segment-down. Using this configuration, if the controller or partner node loses three consecutive hello PDUs, the failtimer expires—but there might not be a break in the segment. Opening a blocked port in this situation creates a loop.

The following describes some general recommendations for using this command:

- When you configure your Extreme Networks switches as the partner and controller, respectively, make sure that their segment timer configurations are identical.

For example, if you have a partner switch with the segment-timeout expiry action set to send-alert, make sure the controller switch has its segment-timeout expiry action set to send-alert.



However, if you have a partner switch with the segment-timeout expiry action set to send-alert, and the controller switch does not have a segment timer configuration, you must configure the partner switch's segment-timeout expiry action to segment-down.

- If you have a network containing non-Extreme Networks switches or non-EAPS devices, set the segment-timeout expiry action to segment-down.

The following events can cause a ring segment failure:

- There is a hardware failure.
- The controller or partner received a Link Down message from the partner or controller, respectively.
- The segment timer expires and the expiry action was set to segment-down. This means that either the controller or partner did not receive health check messages during the defined segment timeout period.

To view shared-port information, including shared-port segment status, use the following command:

```
show eaps shared-port {port}{detail}
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms with the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure eaps shared-port segment-timers health-interval

```
configure eaps shared-port port segment-timers health-interval seconds
```

Description

Configures the shared-port health interval timeout.

Syntax Description

<i>port</i>	Specifies the port number of the common link port.
<i>seconds</i>	Specifies the amount of health interval, in seconds.

Default

N/A.



Usage Guidelines

This command allows you to configure the length of the shared-port health interval timeout, in seconds, for a given port.

Example

The following command configures a shared-port health interval timeout of 10 seconds on port 1:1.

```
configure eaps shared-port 1:1 segment-timers health-interval 10
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms with the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure eaps shared-port segment-timers timeout

```
configure eaps shared-port port segment-timers timeout seconds
```

Description

Configures the shared-port timeout.

Syntax Description

<i>port</i>	Specifies the port number of the common link port.
<i>seconds</i>	Specifies the amount of health interval, in seconds.

Default

N/A.

Usage Guidelines

This command allows you to configure the length of the shared-port timeout, in seconds, for a given port.



Example

The following command configures a shared-port timeout of 10 seconds on port 1:1.

```
configure eaps shared-port 1:1 segment-timers timeout 10
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms with the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure forwarding L2-protocol fast-convergence

configure forwarding L2-protocol fast-convergence on | off

Description

Configures the switch to flooding the unicast traffic during L2 protocol convergence.

Syntax Description

on	Used to avoid flooding the unicast traffic during L2 protocol convergence.
off	Used to Temporarily flooding unicast traffic during L2 protocol convergence. (default)

Default

On.

Usage Guidelines

Use this command to influence the L2-protocol convergence when topology changes in the network to minimize the congestion.

Example

The following command will influence the L2-Protocol control traffic:

```
configure forwarding L2-protocol fast-convergence off
```



History

This command was first available in ExtremeXOS 15.1.3.

Platform Availability

This command available on all Summit, BD8K, BD-X8 platforms.

configure ip-arp fast-convergence

configure ip-arp fast-convergence [on | off]

Description

This command improves IP convergence for IP traffic.

Syntax Description

on	Fast-convergence on.
off	Fast-convergence off (default).

Default

Off.

Usage Guidelines

Use this command for quick recovery when running IP traffic over an EAPS ring.

Example

The following example shows output from the configure ip-arp fast-convergence on command:

```
E4G200-1.2 # show iparp
VR          Destination      Mac                Age  Static  VLAN
VID  Port
VR-Default  10.109.1.2      00:04:96:52:2b:16  0    NO     box1-box2
950   3
VR-Default  10.109.1.6      00:04:96:52:2a:f2  0    NO     box1-box3
951   1
Dynamic Entries  :          2          Static
Entries          :          0
Pending Entries  :          0
In Request       :          1          In
Response         :          1
Out Request      :          1          Out
Response         :          1
Failed Requests  :          0
Proxy Answered  :          0
```



```

Rx Error      :          0          Dup IP
Addr         :          0.0.0.0
Rejected Count :          Rejected IP      :
Rejected Port :          Rejected I/F     :
Max ARP entries :          8192          Max ARP pending entries :
256
ARP address check: Enabled          ARP refresh      :
Enabled
Timeout      :          20 minutes          ARP Sender-Mac Learning  :
Disabled
Locktime     :          1000 milliseconds
Retransmit Time :          1000 milliseconds
Reachable Time :          900000 milliseconds (Auto)
Fast Convergence :          Off
E4G200-1.3 #
E4G200-1.4 # show iparp
VR           Destination      Mac           Age  Static  VLAN
VID  Port
VR-Default  10.109.1.2         00:04:96:52:2b:16  1    NO    box1-box2
950  3
VR-Default  10.109.1.6         00:04:96:52:2a:f2  1    NO    box1-box3
951  1
Dynamic Entries :          2          Static
Entries         :          0
Pending Entries :          0
In Request      :          1          In
Response        :          1
Out Request     :          1          Out
Response        :          1
Failed Requests :          0
Proxy Answered  :          0
Rx Error       :          0          Dup IP
Addr          :          0.0.0.0
Rejected Count :          Rejected IP      :
Rejected Port :          Rejected I/F     :
Max ARP entries :          8192          Max ARP pending entries :
256
ARP address check: Enabled          ARP refresh      :
Enabled
Timeout      :          20 minutes          ARP Sender-Mac Learning  :
Disabled
Locktime     :          1000 milliseconds
Retransmit Time :          1000 milliseconds
Reachable Time :          900000 milliseconds (Auto)
Fast Convergence :          On
E4G200-1.5 #

```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on all platforms.



create eaps

create eaps *name*

Description

Creates an EAPS domain with the specified name.

Syntax Description

<i>name</i>	Specifies the name of an EAPS domain to be created. Can be up to 32 characters in length.
-------------	---

Default

N/A.

Usage Guidelines

An EAPS domain name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see [Object Names](#) in the ExtremeXOS Concepts Guide.

Example

The following command creates EAPS domain eaps_1:

```
create eaps eaps_1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

create eaps shared-port

create eaps shared-port *ports*

Description

Creates an EAPS shared port on the switch.



Syntax Description

<code>ports</code>	Specifies the port number of the common link port.
--------------------	--

Default

N/A.

Usage Guidelines

To configure a common link, you must create a shared port on each switch on either end of the common link.

Example

The following command creates a shared port on the EAPS domain.

```
create eaps shared-port 1:2
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms with the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

delete eaps

```
delete eaps name
```

Description

Deletes the EAPS domain with the specified name.

Syntax Description

<code>name</code>	Specifies the name of an EAPS domain to be deleted.
-------------------	---

Default

N/A.



Usage Guidelines

None.

Example

The following command deletes EAPS domain eaps_1:

```
delete eaps eaps_1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

delete eaps shared-port

```
delete eaps shared-port ports
```

Description

Deletes an EAPS shared port on a switch.

Syntax Description

<i>ports</i>	Specifies the port number of the Common Link port.
--------------	--

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes shared port 1:1.

```
delete eaps shared-port 1:1
```



History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms with the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable eaps

```
disable eaps {name}
```

Description

Disables the EAPS function for a named domain or for an entire switch.

Syntax Description

<i>name</i>	Specifies the name of an EAPS domain.
-------------	---------------------------------------

Default

Disabled for the entire switch.

Usage Guidelines

To prevent loops in the network, the switch displays by default a warning message and prompts you to disable EAPS for a specific domain or the entire switch. When prompted, do one of the following:

- Enter *y* to disable EAPS for a specific domain or the entire switch.
- Enter *n* or press [Return] to cancel this action.

If you have considerable knowledge and experience with EAPS, you might find the EAPS loop protection warning messages unnecessary. For more information, see the `configure eaps config-warnings off`.

Example

The following command disables the EAPS function for entire switch:

```
disable eaps
```

The switch displays the following warning message and prompts you to confirm this action:



WARNING: Disabling EAPS on the switch could cause a loop in the network!
Are you sure you want to disable EAPS? (y/n) Enter y to disable EAPS on the switch. Enter n to cancel this action.

The following command disables the EAPS function for the domain eaps-1:

```
disable eaps eaps-1
```

The switch displays the following warning message and prompts you to confirm this action:

WARNING: Disabling specific EAPS domain could cause a loop in the network!

Are you sure you want to disable this specific EAPS domain? (y/n)

Enter y to disable the EAPS function for the specified domain. Enter n to cancel this action.

History

This command was first available in ExtremeXOS 11.0.

The interactive messages were added in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

enable eaps

```
enable eaps {name}
```

Description

Enables the EAPS function for a named domain or for an entire switch.

Syntax Description

<i>name</i>	Specifies the name of an EAPS domain.
-------------	---------------------------------------

Default

Disabled.

Default command enables EAPS for the entire switch.



Usage Guidelines



Note

If you use the same name across categories (for example, STPD and EAPS names), you must specify the identifying keyword as well as the actual name.

To configure and enable an EAPS, complete the following steps:

- 1 Create EAPS domain and assign the name.
- 2 Configure the control VLAN.
- 3 Configure the protected VLAN(s).
- 4 Add the control VLAN to EAPS domain.
- 5 Add the protected VLAN(s) to EAPS domain.
- 6 Configure EAPS mode, master or transit.
- 7 Configure EAPS port, secondary and primary.
- 8 If desired, configure timeout and action for failtimer expiration*.
- 9 If desired, configure the hello time for the health-check packets*.
- 10 Enable EAPS for the entire switch.
- 11 If desired, enable Fast Convergence*.
- 12 Enable EAPS for the specified domain.

Although you can enable EAPS prior to configuring these steps, the EAPS domain(s) does not run until you configure these parameters.

* These steps can be configured at any time, even after the EAPS domains are running.

You must enable EAPS globally and specifically for each named EAPS domain.

Example

The following command enables the EAPS function for entire switch:

```
enable eaps
```

The following command enables the EAPS function for the domain eaps-1:

```
enable eaps eaps-1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.



show eaps

```
show eaps {eapsDomain} {detail}
```

Description

Displays EAPS status information.

Syntax Description

<i>eapsDomain</i>	Specifies the name of an EAPS domain.
detail	Specifies all available detail for each domain.

Default

N/A.

Usage Guidelines

If you enter the show eaps command without a keyword, the command displays less than with the detail keyword.

Use the optional eapsDomain parameter to display status information for a specific EAPS domain.

Some state values are different on a transit node than on a master node.

When you enter the show eaps command without a domain name, the switch displays the following fields:

EAPS Enabled:	Current state of EAPS on this switch: Yes—EAPS is enabled on the switch.No—EAPS is not enabled.
EAPS Fast Convergence:	Displays only when Fast Convergence is on.
EAPS Display Config Warnings:	Displays the setting for loop protection messages: On—Loop protection messages are displayed (this is the default behavior).Off —Loop protection messages are not displayed.
EAPS Multicast Add Ring Ports:	Displays the configuration of the multicast add-ring-ports feature as configured with the <code>configure eaps multicast add-ring-ports</code> command.
EAPS Multicast Send IGMP Query:	Displays the configuration of the multicast send-igmp-query feature as configured with the <code>configure eaps multicast send-igmp-query</code> command.
EAPS Multicast Temporary Flooding:	Displays the configuration of the multicast temporary-flooding feature as configured with the <code>configure eaps multicast temporary-flooding</code> command.
EAPS Multicast Temporary Flooding Duration:	Displays the duration configuration for the multicast temporary-flooding feature as configured with the <code>configure eaps multicast temporary-flooding duration</code> command.



Number of EAPS instances:	Number of EAPS domains created. The maximum number of EAPS domains per switch is 128.
Domain:	Entries in this column identify the name of an EAPS domain.
State:	<p>On a transit node, the command displays one of the following states: Idle—The EAPS domain has been enabled, but the configuration is not complete. Links-Up—This EAPS domain is running, and both its ports are up and in the FORWARDING state. Links-Down—This EAPS domain is running, but one or both of its ports are down. Preforwarding—This EAPS domain is running, and both of its ports are up, but one of them is in a temporary BLOCKED state.</p> <p>On a master node, the command displays one of the following states: Idle—The EAPS domain has been enabled, but the configuration is not complete. Init—The EAPS domain has started but has not yet determined the status of the ring. The secondary port is in a BLOCKED state. Complete—The ring is in the COMPLETE state for this EAPS domain. Failed—There is a break in the ring for this EAPS domain. Pre-Init—The EAPS domain has started operation for Init state and has sent a request to lower hardware layers to block the secondary port. It is in transient state waiting for acknowledgement from hardware layer indicating the operation is completed. Pre-Complete—The EAPS domain has started operation for Complete state and has sent a request to lower hardware layers to block the secondary port. It is in transient state waiting for acknowledgement from the hardware layer indicating the operation is completed. [Failtimer Expired]—When the failtimer expires and it's action is set to send-alert, this flag is set. This flag indicates there is a misconfiguration or hardware problem in the EAPS ring. The EAPS master node continues to remain in COMPLETE or INIT state with it's secondary port blocking.</p>
Mo:	The configured EAPS mode for this switch: transit (T) or master (M).
Primary/Secondary port:	The port numbers assigned as the EAPS primary and secondary ports. On the master node, the port distinction indicates which port is blocked to avoid a loop.
Prio	The EAPS domain priority, which is H for high priority or N for normal priority.

When you enter the show eaps command with a domain name or the detail keyword, the switch displays the following fields:

Name:	Identifies the EAPS domain displayed.
Priority	The EAPS domain priority, which is either High or Normal.



State:	<p>On a transit node, the command displays one of the following states: Idle—The EAPS domain has been enabled, but the configuration is not complete. Links-Up—This EAPS domain is running, and both its ports are up and in the FORWARDING state. Links-Down—This EAPS domain is running, but one or both of its ports are down. Preforwarding—This EAPS domain is running, and both of its ports are up, but one of them is in a temporary BLOCKED state.</p> <p>On a master node, the command displays one of the following states: Idle—The EAPS domain has been enabled, but the configuration is not complete. Init—The EAPS domain has started but has not yet determined the status of the ring. The secondary port is in a BLOCKED state. Complete—The ring is in the COMPLETE state for this EAPS domain. Failed—There is a break in the ring for this EAPS domain. Pre-Init—The EAPS domain has started operation for Init state and has sent a request to lower hardware layers to block the secondary port. It is in transient state waiting for acknowledgement from hardware layer indicating the operation is completed. Pre-Complete—The EAPS domain has started operation for Complete state and has sent a request to lower hardware layers to block the secondary port. It is in transient state waiting for acknowledgement from the hardware layer indicating the operation is completed. [Failtimer Expired]—When the failtimer expires and it's action is set to send-alert, this flag is set. This flag indicates there is a misconfiguration or hardware problem in the EAPS ring. The EAPS master node continues to remain in COMPLETE or INIT state with it's secondary port blocking.</p>
[Running: ...]	Yes—This EAPS domain is running. No—This EAPS domain is not running.
Enabled:	Indicates whether EAPS is enabled on this domain. Y—EAPS is enabled on this domain. N—EAPS is not enabled.
Mode:	The configured EAPS mode for this switch: transit (T) or master (M).
Primary/Secondary port:	The port numbers assigned as the EAPS primary and secondary ports. On the master node, the port distinction indicates which port is blocked to avoid a loop.
Port status:	Unknown—This EAPS domain is not running, so the port status has not yet been determined. Up—The port is up and is forwarding data. Down—The port is down. Blocked—The port is up, but data is blocked from being forwarded.
Tagstatus:	Tagged status of the control VLAN: Tagged—The control VLAN has this port assigned to it, and the port is tagged in the VLAN. Untagged—The control VLAN has this port assigned to it, but the port is untagged in the control VLAN. Undetermined—Either a VLAN has not been added as the control VLAN to this EAPS domain or this port has not been added to the control VLAN.
Hello timer interval:	The configured value of the timer in seconds and milliseconds, specifying the time that the master node waits between transmissions of health check packets.
Fail timer interval:	The configured value of the timer in seconds, specifying the time that the master node waits before the failtimer expires.
Failtimer expiry action:	Displays the action taken when the failtimer expires: Send-alert—Sends a critical message to the syslog when the failtimer expires. Open-secondary-port—Opens the secondary port when the failtimer expires. Displays only for master nodes.



Preforwarding Timer interval: ¹⁰	The configured value of the timer. This value is set internally by the EAPS software. The set value is 15 seconds. Note: If two links in an EAPS domain go down at the same time and one link comes back up, it takes 15 seconds for the reconnected link to start receiving traffic again. Displays only for transit nodes.
Last valid EAPS update:	Indicates the last time a hello packet was received.
EAPS Domain Controller Vlan:	Lists the assigned name and ID of the control VLAN.
EAPS Domain Protected Vlan(s):	Lists the assigned names and VLAN IDs of all the protected VLANs configured on this EAPS domain.
Number of Protected Vlans:	The count of protected VLANs configured on this EAPS domain.

Example

The following command displays information for all EAPS domains:

```
Switch.5 # show eaps
EAPS Enabled: Yes
EAPS Fast-Convergence: Off
EAPS Display Config Warnings: On
EAPS Multicast Add Ring Ports: Off
EAPS Multicast Send IGMP Query: On
EAPS Multicast Temporary Flooding: Off
EAPS Multicast Temporary Flooding Duration: 15 sec
Number of EAPS instances: 2
# EAPS domain configuration :
-----
--
Domain          State           Mo En Pri  Sec  Control-Vlan VID  Count
Prio
-----
--
d1              Idle           T  N  1   2   cv1             (101 ) 0   H
d2              Links-Up       T  Y  3:8 3:16 c2             (1001) 100 H
-----
--
```

The following command displays information for EAPS domain d1:

```
Switch.7 # show eaps d1
Name: d1          Priority: High
State: Idle       Running: No
Enabled: No       Mode: Transit
Primary port: 1   Port status: Unknown  Tag status: Undetermined
Secondary port: 2 Port status: Unknown  Tag status: Undetermined
Hello timer interval: 1 sec 0 millisc
Fail timer interval: 3 sec 0 millisc
Fail Timer expiry action: Send alert
Last valid EAPS update: From Master Id 00:01:30:f9:9c:b0, at Wed Jun 9
09:09:35 2004
```

¹⁰ These fields apply only to transit nodes; they are not displayed for a master node.



```

EAPS Domain has following Controller Vlan:
Vlan Name          VID
c1                  1000
EAPS Domain has following Protected Vlan(s):
Vlan Name          VID
p_1                1
p_2                2
p_3                3
p_4                4
p_5                5
p_6                6
p_7                7
p_8                8
p_9                9
p_10               10
p_11               11
p_12               12
p_13               13
p_14               14
p_15               15
p_16               16
p_17               17
p_18               18
p_19               19
p_20               20
p_21               21
p_22               22
p_23               23
p_24               24
p_25               25
p_26               26
p_27               27
p_28               28
p_29               29
p_30               30

```

The following command displays information on EAPS domain domain12, which is configured to send hello packets on the secondary port:

```

Switch.9 # show eaps "domain12"
Name: domain12      Priority: High
State: Complete    Running: Yes
Enabled: Yes        Mode: Master
Primary port:      17          Port status: Up Tag status: Tagged
Secondary port:    27          Port status: Blocked Tag status: Tagged
Hello Egress Port: Secondary
Hello timer interval: 0 sec 100 millisec
Fail timer interval: 0 sec 300 millisec
Fail Timer expiry action: Send alert
Last update: From Master Id 00:04:96:34:e3:43, at Tue May 11 15:39:29 2010
EAPS Domain has following Controller Vlan:
Vlan Name          VID
vlanc12            1002
EAPS Domain has following Protected Vlan(s):
Vlan Name          VID
pvlan11            204
pvlan12            205

```



```
pvlan13                206
Number of Protected Vlans: 3
```



Note

You might see a slightly different display, depending on whether you display the master node or the transit node.

The display from the `show eaps detail` command shows all the information shown in the `show eaps eapsDomain` command, but displays information for all configured EAPS domains.

For the CFM support in EAPS, the existing `show eaps` output places a “!” next to a CFM monitored ring port if the CFM indicates the MEP group for that port is down.

```
X480-48t.1 # sh eaps
EAPS Enabled: Yes
EAPS Fast-Convergence: Off
EAPS Display Config Warnings: Off
EAPS Multicast Add Ring Ports: Off
EAPS Multicast Send IGMP Query: On
EAPS Multicast Temporary Flooding: Off
EAPS Multicast Temporary Flooding Duration: 15 sec
Number of EAPS instances: 1
# EAPS domain configuration :
-----
----
Domain          State          Mo En Pri   Sec   Control-Vlan VID   Count
Prio
-----
----
d2              Failed          M  Y  !41   31    v2              (101 )
1              N
-----
----
Flags : (!) CFM Down
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show eaps cfm groups

```
show eaps cfm groups
```

Description

Displays summary EAPS CFM groups information.



Syntax Description

There are no keywords or variables for this command.

Default

N/A.

Usage Guidelines

The following command displays EAPS CFM group information:

```
X480-48t.2 # sh eaps cfm groups
-----
--
MEP Group Name                Status Port    MEP ID
-----
--
eapsCfmGrp1                   Up           41         11
eapsCfmGrp2                   Up           31         12
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on all platforms.

show eaps counters

```
show eaps counters [eapsDomain | global]
```

Description

Displays summary EAPS counter information.

Syntax Description

<i>eapsDomain</i>	Specifies the name of an EAPS domain. The switch displays counter information for only that domain.
global	Displays EAPS counter information when the events counted are not applicable to any specific EAPS domain.

Default

N/A.



Usage Guidelines

If you specify the name of an EAPS domain, the switch displays counter information related to only that domain. If you specify the global keyword, the switch displays EAPS counter information when the events counted are not applicable to any specific EAPS domain. The output displayed is for all configured EAPS domains, not just one specific EAPS domain.

Viewing and maintaining statistics on a regular basis allows you to see how well your network is performing. If you keep simple daily records, you will see trends emerging and notice problems arising before they cause major network faults.

Clearing the Counters

The counters continue to increment until you clear the information. By clearing the counters, you can see fresh statistics for the time period you are monitoring. To clear, reset the EAPS counters, use one of the following commands:

- `clear counters`
- `clear eaps counters`

Understanding the Output

The following table describes the significant fields and values in the display output of the `show eaps counters eapsDomain` command:

Field	Description
Rx-Health	Indicates the EAPS domain received EAPS Health PDUs.
Rx-RingUp-FlushFdb	Indicates the EAPS ring is up, and the EAPS domain received EAPS RingUp-FlushFdb PDUs to flush the FDB.
Rx-RingDown-FlushFdb	Indicates the EAPS ring is down, and the EAPS domain received EAPS RingDown-FlushFdb PDUs to flush the FDB.
Rx-Link-Down	Indicates the EAPS domain received EAPS Link-Down PDUs and took down the link.
Rx-Flush-Fdb	Indicates the EAPS domain received EAPS Flush-Fdb PDUs and flushed the FDB.
Rx-Suspend-Prefwd-Timer	Indicates the EAPS domain received EAPS Suspend-Preforward-Timer PDUs. NOTE: Switches running ExtremeWare send this PDU during an MSM/MM failover. Switches running ExtremeXOS 10.1 or later do not send or receive this PDU.
Rx-Query-Link-Status	Indicates the EAPS domain received EAPS Query-Link-Status PDUs.
Rx-Link-Up	Indicates the EAPS domain received EAPS Link-Up PDUs and brought the link back up.
Rx-Unknown	Indicates the EAPS domain dropped unknown EAPS PDUs.
Rx-Another-Master	Indicates the EAPS domain dropped EAPS PDUs because there is another Master switch in the same EAPS domain.
Rx-Unconfigured-Port	Indicates the EAPS domain dropped EAPS PDUs because the ingress port is not configured to be a ring port for the EAPS domain and the corresponding control VLAN.



Field	Description
Rx-Health-Pdu-Pri-Port	Indicates the EAPS domain dropped EAPS Health PDUs because the primary port received them instead of the secondary port. NOTE: The secondary port of the Master switch must receive EAPS Health PDUs, not the primary port.
Tx-Health	Indicates the EAPS domain sent EAPS Health PDUs.
Tx-RingUp-FlushFdb	Indicates the EAPS ring is up, and the EAPS domain sent EAPS RingUp-FlushFdb PDUs to flush the FDB.
Tx-RingDown-FlushFdb	Indicates the EAPS ring is down, and the EAPS domain sent EAPS RingDown-FlushFdb PDUs to flush the FDB.
Tx-Link-Down	Indicates the EAPS domain sent EAPS Link-Down PDUs because the link went down.
Tx-Flush-Fdb	Indicates the EAPS domain sent EAPS Flush-Fdb PDUs because the FDB needs to be flushed.
Tx-Suspend-Prefwd-Timer	Indicates the EAPS domain sent EAPS Suspend-Preforward-Timer PDUs. NOTE: Switches running ExtremeWare send this PDU during an MSM/MM failover. Switches running ExtremeXOS 10.1 or later do not send or receive this PDU. This counter should remain at 0.
Tx-Query-Link-Status	Indicates the EAPS domain sent EAPS Query-Link-Status PDUs.
Tx-Link-Up	Indicates the EAPS domain sent EAPS Link-Up PDUs and the link is up.
Tx-Unknown	Indicates the number of unknown EAPS PDUs sent by the EAPS domain. NOTE: Unknown EAPS PDUs can be a new type of PDU that the switch does not track in the sending routine.
Tx-Transmit-Err	Indicates the number of EAPS PDUs the EAPS domain was unable to send because of an error.
Fw-Link-Down	Indicates the number of EAPS Link-Down PDUs received by the EAPS domain and forwarded in slow path.
Fw-Flush-Fdb	Indicates the number of EAPS Flush-Fdb PDUs received by the EAPS domain and forwarded in slow path.
FW-Query-Link-Status	Indicates the number of EAPS Query-Link-Status PDUs received by the EAPS domain and forwarded in slow path.
Fw-Unknown	Indicates the number of unknown EAPS PDUs forwarded in slow path. NOTE: Unknown EAPS PDUs can be a new type of PDU that the switch does not track in the forwarding routine.
Fw-Transmit-Er	Indicates the number of EAPS PDUs the EAPS domain was unable to forward in slow path because of an error.

Note



Rx and Fw counters—If a PDU is received, processed, and consumed, only the Rx counter increments. If a PDU is forwarded in slow path, both the Rx counter and Fw counter increment.

The following table describes the significant fields and values in the display output of the `show eaps counters` global command:



Field	Description
Rx-Failed	Indicates an error occurred when receiving packets from the Layer 2 forwarding engine.
Rx-Invalid-Vlan-Intf	Indicates that the VLAN interface for the incoming VLAN cannot be found.
Rx-Undersize-Pkt	Indicates the length of the packet is less than the length of the header.
Rx-Invalid-8021Q-Tag	Indicates the VlanTypeLength field in the Ethernet header does not match the default Ethernet value for the 802.1Q tag.
Rx-Invalid-SNAP-Type	Indicates an invalid Subnetwork Access Protocol (SNAP) value in the Ethernet header.
Rx-Invalid-OUI	Indicates the Organizational Unique Identifier (OUI) value in the Ethernet header does not match 00:EO:2B.
Rx-EEP-Unsupported-Version	Indicates an unsupported Extreme Encapsulation Protocol (EEP) version. The EEP version should be 1.
Rx-EEP-Invalid-Length	Indicates the length of the EEP header is greater than the length of the packet.
Rx-EEP-Checksum-Invalid	Indicates the EEP checksum is invalid.
Rx-Domain-Invalid	Indicates the control VLAN's incoming PDU is not associated with an EAPS domain.
Rx-Lif-Invalid	Indicates that EAPS is unable to determine the logical interface (LIF) for the ingress port.
Rx-Lif-Down	Indicates the LIF for the ingress port is in the Down state.
Tx-Failed	Indicates an error occurred when sending packets to the Layer 2 forwarding engine.

Example

The following command displays the counters for a specific EAPS domain named eaps1:

```
show eaps counters eaps1
```

The following is sample output from this command:

```
Counters for EAPS domain: eaps1
Rx Stats
Rx-Health : 0
Rx-Ringup-Flushfdb : 0
Rx-Ringdown-Flushfdb : 0
Rx-Link-Down : 0
Rx-Flush-Fdb : 0
Rx-Suspend-Prefwd-Timer : 0
Rx-Query-Link-Status : 0
Rx-Link-Up : 0
Rx Dropped
Rx-Unknown : 0
Rx-Another-Master : 0
```




```

Rx-Unconfigured-Port      : 0
Rx-Health-Pdu-Pri-Port   : 0
Tx Stats
Tx-Health                  : 5011
Tx-Ringup-Flushfdb        : 0
Tx-Ringdown-Flushfdb      : 0
Tx-Link-Down              : 0
Tx-Flush-Fdb              : 0
Tx-Suspend-Prefwd-Timer   : 0
Tx-Query-Link-Status      : 3342
Tx-Link-Up                : 0
Tx Dropped
Tx-Unknown                 : 0
Tx-Transmit-Err           : 0
Fw Stats
Fw-Link-Down              : 0
Fw-Flush-Fdb              : 0
Fw-Query-Link-Status      : 0
Fw Dropped
Fw-Unknown                 : 0
Fw-Transmit-Err           : 0

```

The following command displays the global EAPS counters:

```
show eaps counters global
```

The following is sample output from this command:

```

Global counters for EAPS:
Rx-Failed : 0
Rx-Invalid-Vlan-Intf : 0
Rx-Undersize-Pkt : 0
Rx-Invalid-SNAP-Type : 0
Rx-Invalid-OUI : 0
Rx-EEP-Unsupported-Version : 0
Rx-EEP-Invalid-Length : 0
Rx-EEP-Checksum-Invalid : 0
Rx-Domain-Invalid : 0
Rx-Failed : 0
Rx-Lif-Invalid : 0
Rx-Lif-Down : 0
Tx-Failed : 0

```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.



show eaps counters shared-port

```
show eaps counters shared-port [global | port {segment-port segport
{eapsDomain}}]
```

Description

Displays summary EAPS shared port counter information.

Syntax Description

global	Displays general counter information for all configured EAPS shared port instances. The output displayed is calculated for all configured EAPS shared ports; not just one specific shared port instance.
<i>port</i>	Identifies the port number of the specified common link port.
<i>segport</i>	Identifies the segment port. The segment port is the other ring port of an EAPS domain that is not the shared-port.
<i>eapsDomain</i>	Specifies the name of the EAPS domain. If no EAPS domain is specified, all counters for all EAPS domains on the specified segment port are displayed.

Default

N/A.

Usage Guidelines

If the switch is configured for EAPS shared ports, use this command to display an array of counters associated with the EAPS shared port functionality.

If you specify the global keyword, the switch displays general counter information for all configured EAPS shared port instances. The output displayed is calculated for all configured EAPS shared ports; not just one specific shared port instance.

If you specify a particular EAPS shared port, the switch displays counter information related to only that shared port.

If you specify a particular EAPS segment port, the switch displays counter information related to only that segment port for the specified EAPS domain.

Viewing and maintaining statistics on a regular basis allows you to see how well your network is performing. If you keep simple daily records, you will see trends emerging and notice problems arising before they cause major network faults.

Clearing the Counters

The counters continue to increment until you clear the information. By clearing the counters, you can see fresh statistics for the time period you are monitoring. To clear, reset the EAPS counters, including the shared port counters, use one of the following commands:



- `clear counters`
- `clear eaps counters`

Understanding the Output

The following table describes the significant fields and values in the display output of the `show eaps counters shared-port` global command:

Field	Description
Rx-Invalid-Instance	Displays the number of dropped EAPS shared-port PDUs because there is not a valid EAPS shared port instance for the incoming port.
Rx-Unknown	Displays the number of unknown EAPS PDUs dropped by the shared port instances.
Fw-Invalid-Instance	Displays the number of EAPS shared-port PDUs that could not be forwarded in slow path because the shared port instances could not find a valid EAPS shared port instance for the outgoing port.

The following table describes the significant fields and values in the display output of the `show eaps counters shared-port portsegment-port segport eapsDomain` command:

Field	Description
Rx-Seg-Health	Indicates the shared port instance received EAPS shared ports Segment-Health-Check PDUs.
Rx-Path-Detect	Indicates the shared port instance received EAPS shared ports Path-Detect PDUs.
Rx-Flush-Notify	Indicates the shared port instance received EAPS shared ports Flush-Notify PDUs and flushed the FDB. If this PDU reaches a port of the shared ports pair that initiated the PDU, the shared port instance might terminate the PDU. Otherwise, the shared port instance forwards the PDU.
Rx-Unknown	Displays the number of unknown EAPS PDUs dropped by the shared port instance.
Rx-Seg-Health-Dropped	Displays the number of EAPS shared ports Segment-Health-Check PDUs dropped by the shared port instance. This counter increments if the Segment-Health-Check PDU returns to the sending switch. If that occurs, the switch drops the Segment-Health-Check PDU.
Rx-Path-Detect-Dropped	Displays the number of EAPS shared ports Path-Detect PDUs dropped by the shared port instance. This counter increments in the following situations: If the packet's Fwd-id matches the EAPS shared port's Link-Id, the port is not in the blocking state, and the incoming port is a segment port. If the packet's Link-Id matches the EAPS shared port's Link-Id, the port is not in the blocking state, and the incoming port is a segment port.
Rx-Flush-Notify-Dropped	Displays the number of EAPS shared ports Flush-Notify-Dropped PDUs dropped by the shared port instance. This counter increments in the following situations: If the Flush-Notify-Dropped PDU returns to the sending switch. If the packet's Fwd-Id matches the EAPS shared port's Link-Id and the port is not in the blocking state.



Field	Description
Rx-Dropped-Invalid-Port	Displays the number of EAPS shared ports PDUs dropped by the shared port instance because it does not exist.
Tx-Seg-Health	Indicates the shared port instance sent EAPS shared ports Segment-Health-Check PDUs.
Tx-Path-Detect	Indicates the shared port instance sent EAPS shared ports Path-Detect PDUs. NOTE: This counter appears under Common Link Port Stats and should always be 0.
Tx-Flush-Notify	Indicates the shared port instance sent EAPS shared ports Flush-Notify PDUs to flush the FDB. NOTE: This counter appears under Common Link Port Stats and should always be 0.
Tx-Flush-Fdb	Indicates the shared port instance sent EAPS Flush-Fdb PDUs because the FDB needs to be flushed. NOTE: This counter appears under Common Link Port Stats and should always be 0.
Tx-Unknown	Indicates the number of unknown EAPS PDUs sent by the shared port instance. NOTE: Unknown EAPS PDUs can be a new type of PDU that the switch does not track in the sending routine.
Tx-Transmit-Err	Indicates the number of EAPS PDUs the shared port instance was unable to send because of an error.
Fw-Seg-Health	Indicates the number of EAPS shared ports Segment-Health-Check PDUs received by the shared port instance and forwarded in slow path.
Fw-Path-Detect	Indicates the number of EAPS shared ports Path-Detect PDUs received by the shared port instance and forwarded in slow path.
Fw-Flush-Notify	Indicates the number of EAPS Flush-Notify PDUs received by the shared port instance and forwarded in slow path to flush the FDB.
Fw-Flush-Fdb	Indicates the number of EAPS Flush-Fdb PDUs received by the shared port instance and forwarded in slow path.
Fw-Unknown	Indicates the number of unknown EAPS PDUs forwarded in slow path. NOTE: Unknown EAPS PDUs can be a new type of PDU that the switch does not track in the forwarding routine.
Fw-Transmit-Err	Indicates the number of EAPS PDUs the shared port instance was unable to forward in slow path because of an error.

Example

The following command displays global, high-level counter information for EAPS shared port:

```
show eaps counters shared-port global
```

The following is sample output from this command:

```
Global counters for EAPS Shared-Ports:
Rx Dropped
Rx-Invalid-Instance : 0
```



```

Rx-Unknown          : 0
Fw Dropped
Fw-Invalid-Instance : 0

```

The following example assumes that port 17 is configured as an EAPS shared port. The following command displays counter information the specified EAPS shared port:

```
show eaps counters shared-port 17
```

The following is sample output from this command:

```

Counters for EAPS Shared-Port 17:
Common Link Port Stats
Rx Stats
Rx-Seg-Health          : 0
Rx-Path-Detect        : 0
Rx-Flush-Notify       : 0
Rx Dropped
Rx-Seg-Health-Dropped : 0
Rx-Path-Detect-Dropped : 0
Rx-Flush-Notify-Dropped : 0
Rx-Dropped-Invalid-Port : 0
Tx Stats
Tx-Seg-Health          : 0
Tx-Path-Detect        : 0
Tx-Flush-Notify       : 0
Tx-Flush-Fdb          : 0
Tx Dropped
Tx-Unknown            : 0
Tx-Transmit-Err       : 0
Fw Stats
Fw-Seg-Health          : 0
Fw-Path-Detect        : 0
Fw-Flush-Notify       : 0
Fw Dropped
Fw-Unknown            : 0
Fw-Transmit-Err       : 0

```

The following example assumes that port 1:2 is configured as an EAPS shared port and port 1:1 is a segment port. The following command displays counter information the specified EAPS shared port, segment port, and EAPS domain:

```
show eaps counters shared-port 1:2 segment-port 1:1 eaps1
```

The following is sample output from this command:

```

Counters for EAPS Shared-Port 1:2, Segment Port: 1:1, EAPS Domain: eaps1
Rx Stats
Rx-Seg-Health          : 0
Rx-Path-Detect        : 0
Rx-Flush-Notify       : 0
Rx-Seg-Health-Dropped : 0

```



```

Rx-Path-Detect-Dropped : 0
Rx-Flush-Notify-Dropped : 0
Rx-Dropped-Invalid-Port : 0
Tx Stats
Tx-Seg-Health : 2275
Tx-Path-Detect : 0
Tx-Flush-Notify : 0
Tx-Flush-Fdb : 0
Tx-Transmit-Err : 0
Tx-Unknown : 0
Fw Stats
Fw-Seg-Health : 0
Fw-Path-Detect : 0
Fw-Flush-Notify : 0
Fw-Transmit-Err : 0
Fw-Unknown : 0

```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms with the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show eaps shared-port

```
show eaps shared-port {port} {detail}
```

Description

Displays shared-port information for one or more EAPS domains.

Syntax Description

<i>port</i>	Specifies a shared-port.
detail	Specifies to display the status of all segments and VLANs.

Default

N/A.

Usage Guidelines

If you enter the `show eaps shared-port` command without the `detail` keyword, the command displays a summary of status information for all configured EAPS shared ports.



If you specify an EAPS shared-port, the command displays information about that specific port and the related segment ports. The segment ports are sorted in ascending order based on their port number. You can use this order and your knowledge of the EAPS topology to determine which segment port becomes the active-open port if the common link fails. For more information, see [Common Link Fault Detection and Response](#) in the ExtremeXOS Concepts Guide.

You can use the detail keyword to display more detailed status information about the segments and VLANs associated with each shared port.

The following table describes the significant fields and values in the display output of the `show eaps shared-port {port {detail}}` commands:

Field	Description
Shared Port	Displays the port number of the shared port.
Mode	Indicates whether the switch on either end of the common link is a controller or partner. The mode is configured by the user.
Link ID	The link ID is the unique common link identifier configured by the user.
Up	Displays one of the following: Yes—Indicates that the link ID and the mode are configured.No—Indicates that the link ID or the mode is not configured.
State	Displays one of the following states: Idle—Shared-port instance is not running.Ready—The EAPS shared-port instance is running, the neighbor can be reached, and the common link is up.Blocking—The EAPS shared-port instance is running, the neighbor cannot be reached, or the common link is down.Preforwarding—The EAPS shared-port instance is in a blocking state, and the common link came up. To prevent a superloop, a temporary blocking state is created before going into Ready state.
Domain Count	Indicates the number of EAPS domains sharing the common link.
VLAN Count	Indicates the total number of VLANs that are protected under the EAPS domains sharing this common link.
Nbr	Yes—Indicates that the EAPS instance on the other end of the common link is configured with matching link ID and opposite modes. For example, if one end of the common link is configured as a controller, the other end must be configured as a partner.Err—Indicates that the EAPS instance on the other end of the common link is configured with a matching link ID, but the modes are configured the same. For example, both modes are configured as controller, or both modes are configured as partner.No—The neighbor on the other end of the common link cannot be reached. Indicates one or more of the following: - The switch on the other end of the common link is not running. - The shared port has not been created. - The link IDs on each side of the common link do not match. - The common link, and any other segment, between the controller and partner are not fully connected.
RB ID	The ID of the root blocker. If the value is none, there are not two or more common-link failures.
RB State	None—This EAPS shared-port is not the root blocker.Active—This EAPS shared-port is the root blocker and is currently active.Inactive—This EAPS shared-port is the root blocker but is currently inactive.
Active Open (available with the detail keyword)	None—Indicates that there is no Active-Open port on the VLAN.Port #—Indicates the port that is Active-Open and is in a forwarding state.



Field	Description
Segment Timer expiry action	Segment down—Specifies that if the controller or partner switch detects a down segment, that segment stays down and a query is not sent through the ring. The switch marks the segment status as Down. Send alert—Specifies that if the controller or partner switch detects a down segment, that switch keeps the segment up and sends a warning message to the log (default). The switch sends a trap alert and sets the failed flag [F].
Segment Port (available with the detail keyword or by specifying a shared port)	Identifies the segment port of an EAPS ring that shares the common link.
Status (available with the detail keyword or by specifying a shared port)	Up—Connectivity is established between the segment and the EAPS shared-port on the common link neighbor. Down—There is a break in the path between the segment and the EAPS shared-port on the common link neighbor. Blocking-Up—The path is Up, but due to the root blocker being in the Active state, this port is blocked to prevent a loop. Blocking-Down—The root blocker is in the Active state; however, the path is Down. Because the path is Down, there is no need to block the root blocker port to prevent a loop. [F]—The segment timer has expired but has not received an explicit link-down notification. The segment port remains in the Up state, with the timer expired flag set to True.
EAPS Domain (available with the detail keyword or by specifying a shared port)	The EAPS domain assigned to the segment port.
Vlan-port count (available with the detail keyword or by specifying a shared port)	The total number of VLANs being protected on this segment port.
Adjacent Blocking Id (available with the detail keyword or by specifying a shared port)	None—The neighbor on this port is not reporting a Controller in the Blocking state. <i>Link-Id</i> —The neighbor on this port is a controller in the Blocking state with a link ID of <i>Link-Id</i> .
Segment RB Id (available with the detail keyword or by specifying a shared port)	None—The neighbor on this port is not aware of a root blocker in the network. <i>RB-Id</i> —The neighbor on this port has determined that there is a root blocker in the network with a link ID of <i>RB-Id</i> .
Vlan (available with the detail keyword or by specifying a shared port)	Displays a list of VLANs protected by the segment port.
Virtual-port Status (available with the detail keyword or by specifying a shared port)	This information appears for the Controller, when it is in either the Blocking or Preforwarding state. Active-Open—This VLAN or port is in the Forwarding state and has connectivity to the neighboring EAPS shared port via this port. Open—This VLAN or port is in the Forwarding state but does not have connectivity to the neighboring EAPS shared port via this port. Blocked—This VLAN or port is in the Blocking state to prevent a loop in the network. Down—This port's link is down. Active—At this moment, this VLAN or port is not being handled by EAPS shared port. Rather, this VLAN or port is being handled by the regular EAPS protocol.
Bvlan	When a common link connects an access VLAN (CVLAN or SVLAN) to a core VLAN (BVLAN), this field displays the BVLAN name. For more information, see Common Link Fault Detection and Response in the <i>ExtremeXOS Concepts Guide</i> .



Example

The following command displays shared-port information for all EAPS shared ports on a switch:

```
show eaps shared-port
EAPS shared-port count: 1
-----
--
Link          Domain Vlan      RB      RB
Shared-port  Mode          Id  Up State  count  count Nbr State  Id
-----
--
10:1         Controller  1   Y  Ready    2     1   Yes None  None
Segment Timer expiry action: Send alert
-----
--
```

The following command displays detailed information for all EAPS shared ports:

```
show eaps shared-port detail
EAPS shared-port count: 1
-----
--
Link          Domain Vlan      RB      RB
Shared-port  Mode          Id  Up State  count  count Nbr State  Id
-----
--
4:1          Controller  10  Y  Blocking  2     1   Yes Active  10
Segment Timer expiry action: Send alert
Segment Port: 5:7, Status: Blocking-Up
EAPS Domain:          d1
Vlan-port count:     1
Adjacent Blocking Id: None
Segment RB Id:       None
Vlan                Virtual-port Status
p_1                  Blocked
Segment Port: 2:11, Status: Down
EAPS Domain:          d2
Vlan-port count:     1
Adjacent Blocking Id: 20
Segment RB Id:       None
Vlan                Virtual-port Status
p_1                  Open
Vlan: p_1, Vlan-port count: 2, Active Open: None
Segment Port        Virtual-port Status
5:7                  Blocked
2:11                 Open
```

The following command displays detailed information for an EAPS shared port that is in the Blocking state:

```
* Switch.2 # show eaps shared-port 1:24
-----
--
```



```

Link                               Domain Vlan      RB      RB
Shared-port Mode                   Id  Up State      count  count Nbr State   Id
-----
--
1:24      Controller  10  Y  Blocking      3      5      Yes None   None
Segment Health Check interval:      1 sec
Segment Timeout:                      3 sec
Segment Fail Timer expiry action:    Send alert
Common Path Health Check interval:   1 sec
Common Path Timeout:                 3 sec
Segment Port: 3:35  Status: Up
EAPS Domain:          d3
Vlan-port count:      3
Adjacent Blocking Id:  None
Segment RB Id:        None
Segment Port: 3:36  Status: Up
EAPS Domain:          d2
Vlan-port count:      3
Adjacent Blocking Id:  None
Segment RB Id:        None
Segment Port: 3:38  Status: Up
EAPS Domain:          d1
Vlan-port count:      5
Adjacent Blocking Id:  None
Segment RB Id:        None
Vlan: data1,          Vlan-port count: 3,   Active Open: 3:38 Bvlan: metro1
Vlan: data2,          Vlan-port count: 3,   Active Open: 3:38 Bvlan: metro1
Vlan: data3,          Vlan-port count: 3,   Active Open: 3:38 Bvlan: metro2
Vlan: metro1,         Vlan-port count: 1,   Active Open: 3:38
Vlan: metro2,         Vlan-port count: 1,   Active Open: 3:38
-----
--

```



Note

The BVLAN information in the previous example appears only when a BVLAN configuration is present.

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms with the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show eaps shared-port neighbor-info

```
show eaps shared-port {port} neighbor-info {detail}
```



Description

Displays shared-port information from neighboring shared links for one or more EAPS domains.

Syntax Description

<code>port</code>	Specifies a shared-port.
<code>detail</code>	Specifies to display the status of all segments and VLANs.

Default

N/A.

Usage Guidelines

If you enter the command without the detail keyword, the command displays a summary of status information for all configured EAPS shared ports from neighboring shared links. If you specify an EAPS shared-port, the command displays information about that specific port. Otherwise, the command displays information about all of the shared-ports configured on the switch.

You can use the detail keyword to display more detailed status information about the segments and VLANs associated with each shared port. For full details of the significant fields and values in the display output of the command, see the relevant tables in the `show eaps shared port {port} {detail}` command description.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

show vlan eaps

```
show {vlan} vlan_name eaps
```

Description

Displays the EAPS configuration (control, partner, or not added to an EAPS domain) of a specific VLAN.

Syntax Description

<code>vlan_name</code>	Specifies a VLAN name.
------------------------	------------------------



Default

N/A.

Usage Guidelines

Use this command to see if the specified VLAN is associated with an EAPS domain.

The output of this command displays whether the VLAN is a control or partner VLAN for an EAPS domain. This command also displays if the VLAN is not a member of any EAPS domain.

If a VLAN is a partner VLAN for more than one EAPS domain, all of the EAPS domains that the VLAN is a partner of appears in the output.

Example

The following command displays the EAPS configuration for the control VLAN orange in EAPS domain eaps1:

```
show vlan orange eaps
```

The following is sample output from this command:

```
Vlan is Control in following EAPS domain:  
eaps1
```

The following command displays the EAPS configuration for the protected VLAN purple in EAPS domain eaps1:

```
show vlan purple eaps
```

The following is sample output from this command:

```
Vlan is Protected in following EAPS domain(s):  
eaps1
```

The following command displays information about the VLAN default not participating in EAPS:

```
show vlan default eaps
```

The following is sample output from this command:

```
Vlan has not been added to any EAPS domain
```



History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

unconfigure eaps shared-port link-id

unconfigure eaps shared-port *ports* link-id

Description

Unconfigures an EAPS link ID on a shared port on the switch.

Syntax Description

<i>ports</i>	Specifies the port number of the Common Link port.
--------------	--

Default

N/A.

Usage Guidelines

None.

Example

The following command unconfigures the link ID on shared port 1:1.

```
unconfigure eaps shared-port 1:1 link-id
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms with the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



unconfigure eaps shared-port mode

```
unconfigure eaps shared-port ports mode
```

Description

Unconfigures the EAPS shared port mode.

Syntax Description

<i>ports</i>	Specifies the port number of the Common Link port.
--------------	--

Default

N/A.

Usage Guidelines

None.

Example

The following command unconfigures the shared port mode on port 1:1:

```
unconfigure eaps shared-port 1:1 mode
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms with the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure eaps port

```
unconfigure eaps eapsDomain [primary | secondary] port
```

Description

Sets the specified port's internal configuration state to INVALID.



Syntax Description

<i>eapsDomain</i>	Specifies the name of an EAPS domain.
primary	Specifies that the primary port should be unconfigured.
secondary	Specifies that the secondary port should be unconfigured.

Default

N/A.

Usage Guidelines

Unconfiguring an EAPS port sets its internal configuration state to INVALID, which causes the port to appear in the Idle state with a port status of Unknown when you use the `show eaps detail` command to display the status information about the port.

To prevent loops in the network, the switch displays by default a warning message and prompts you to unconfigure the specified EAPS primary or secondary ring port. When prompted, do one of the following:

- Enter `y` to unconfigure the specified port.
- Enter `n` or press [Return] to cancel this action.

If you have considerable knowledge and experience with EAPS, you might find the EAPS loop protection warning messages unnecessary. For more information, see the `configure eaps config-warnings off`.

Example

The following command unconfigures this node's EAPS primary ring port on the domain `eaps_1`:

```
unconfigureeapseaps_1primary port
```

The switch displays the following warning message and prompts you to confirm this action:

```
WARNING: Unconfiguring the Primary port from the EAPS domain could cause
a loop in the network! Are you sure you want to unconfigure the Primary
EAPS Port? (y/n)
```

Enter `y` to continue and unconfigure the EAPS primary ring port. Enter `n` to cancel this action.

The switch displays a similar warning message if you unconfigure the secondary EAPS port.

History

This command was first available in ExtremeXOS 11.0.

The interactive messages were added in ExtremeXOS 11.4.



Platform Availability

This command is available on all platforms.



29 ERPS Commands

```
clear counters erps
configure erps dynamic-state clear
configure erps add control vlan
configure erps add protected vlan
configure erps cfm md-level
configure erps cfm port ccm-interval
configure erps cfm port group
configure erps cfm port mepid
configure erps delete control vlan
configure erps delete protected vlan
configure erps name
configure erps neighbor port
configure erps notify-topology-change
configure erps protection-port
configure erps revert
configure erps ring-ports east | west
configure erps subring-mode
configure erps timer guard
configure erps timer hold-off
configure erps timer periodic
configure erps timer wait-to-block
configure erps timer wait-to-restore
configure erps topology-change
create erps ring
debug erps
debug erps show
delete erps
disable erps
disable erps block-vc-recovery
disable erps ring-name
disable erps topology-change
enable erps
enable erps block-vc-recovery
enable erps ring-name
enable erps topology-change
run erps force-switch | manual-switch
show erps
```

```

show erps ring-name
show erps statistics
unconfigure erps cfm
unconfigure erps neighbor-port
unconfigure erps notify-topology-change
unconfigure erps protection-port
unconfigure erps ring-ports west
configure erps cfm protection group

```

This chapter describes commands for completing the following Ethernet Ring Protection Switching (ERPS) tasks:

- Configuring ERPS
- Displaying ERPS information

For an introduction to ERPS (also known as ITU-T standard G.8032), see the *ExtremeXOS Concepts Guide*.

clear counters erps

```
clear counters erps ring-name
```

Description

Clear statistics on the specified ERPS ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
------------------	--

Default

N/A.

Usage Guidelines

Use this command to clear statistics on the specified ERPS ring.

Example

The following command clears statistics on the ERPS ring named “ring1”:

```
clear counters erps ring1
```



History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps dynamic-state clear

```
configure erps ring-name dynamic-state [force-switch | manual-switch | clear]
port slot:port
```

Description

Clear force and manual switch triggers to the ERPS ring/sub-ring.

Syntax Description

dynamic-state	Configure force/manual/clear switch on the active ERPS ring.
force-switch	Force switch operation.
manual-switch	Manual switch operation.
clear	Clear.

Default

N/A.

Usage Guidelines

Use this command to clear force and manual switch triggers to the ERPS ring/sub-ring.

Example

The following command clears force and manual switch triggers of an ERPS ring named "ring1":

```
configure erps ring1 dynamic-state clear
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.



configure erps add control vlan

```
configure erps ring-name add control {vlan} vlan_name
```

Description

Add a control VLAN on the ERPS ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
control	VLAN that carries ERPS control traffic.
<i>vlan_name</i>	Alphanumeric string identifying the VLAN to be used for control traffic.

Default

N/A.

Usage Guidelines

Use this command to add a control VLAN on the ERPS ring. This is the VLAN that carries ERPS control traffic.



Note

Other VLAN types such as VMAN, SVLAN, CVLAN and BVLAN will not be used for control traffic. A control VLAN cannot be deleted from a ring that has CFM configured.

Example

The following command adds a control VLAN named “vlan10” to an ERPS ring named “ring1”:

```
configure erps ring1 add control vlan vlan10
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps add protected vlan

```
configure erps ring-name add protected {vlan} vlan_name
```



Description

Add a protected VLAN on the ERPS ring. This is a data VLAN that ERPS will protect.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
<i>vlan_name</i>	Alphanumeric string identifying the data VLAN to be added that ERPS will protect. This can be a VLAN, SVLAN, BVLAN or VMAN.

Default

N/A.

Usage Guidelines

Use this command to add a protected data VLAN on the ERPS ring. This VLAN will be protected by ERPS, and it can be a VLAN, SVLAN, BVLAN or VMAN.



Note

The SVLAN-BVLAN combination cannot both be added to the same ring or sub-ring.

Example

The following command adds a protected VLAN named “vlan10” to an ERPS ring named “ring1”:

```
configure erps ring1 add protected vlan vlan10
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps cfm md-level

```
configure erps ring-name cfm md-level level
```

Description

Specify the connectivity fault management (CFM) maintenance domain level for an ERPS ring.



Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
<i>level</i>	Maintenance domain level specified for the ERPS ring.

Default

N/A.

Usage Guidelines

Use this command to specify the CFM maintenance domain level for an ERPS ring.

Example

The following command sets the CFM maintenance domain level to 6 for an ERPS ring named “ring1”:

```
configure erps ring1 cfm md-level 6
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps cfm port ccm-interval

```
configure erps ring-name cfm port [east | west] ccm-interval [100 | 1000 | 10000 | 60000 | 600000]
```

Description

Specify the time interval for transmitting CFM connectivity check messages (CCM) on a port of an ERPS ring.

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
east	East port.
west	West port.
100	100 milliseconds.
1000	1000 milliseconds.
10000	10000 milliseconds.



60000	60000 milliseconds.
600000	600000 milliseconds.

Default

N/A.

Usage Guidelines

Use this command to specify the time interval at which CCMs are transmitted for a port of an ERPS ring.

Example

The following command sets the CCM time interval to 1000 for the east port of an ERPS ring named “ring1”:

```
configure erps ring1 cfm port east ccm-interval 1000
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps cfm port group

```
configure erps ring_name cfm port [east | west] [add | delete] group group_name
```

Description

Associates or disassociates fault monitoring entities on the ERPS ring ports.

Syntax Description

<i>ring_name</i>	Alphanumeric string that identifies the ERPS ring.
east	East port.
west	West port.
add	Associates a CFM Down-MEP entity.
delete	Disassociates a CFM Down-MEP entity.



group	Specifies a CFM Down-MEP group.
<i>group_name</i>	Specifies the name of the Down MEP group.

Default

N/A.

Usage Guidelines

Use this command to associate or disassociate fault monitoring entities on the ERPS ring ports.

Example

The following command associates fault monitoring on the group "group1":

```
configure erps ring1 cfm port east add group1
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms running ExtremeXOS.

configure erps cfm port mepid

```
configure erps ring-name cfm port [east | west] mepid mepid remote-mepid rmepid
```

Description

Specify the maintenance end point identifier for the connectivity fault management (CFM) on a port of an ERPS ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
east	East port.
west	West port.
mepid	Maintenance End Point identifier for the ring ports.
<i>rmepid</i>	Remote Maintenance End Point identifier for the ring ports.



Default

N/A.

Usage Guidelines

Use this command to specify the maintenance end point identifier for CFM on a port of an ERPS ring.

Example

The following command specifies the maintenance end point identifier for the east port of an ERPS ring named "ring1":

```
configure erps ring1 cfm port east mepid 1 remote-mepid 3
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps delete control vlan

```
configure erps ring-name delete control {vlan} vlan_name
```

Description

Delete a control VLAN on the ERPS ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
<i>vlan_name</i>	Alphanumeric string identifying the VLAN used for control traffic.

Default

N/A.



Usage Guidelines

Use this command to delete a control VLAN from the ERPS ring. This is the VLAN that carries ERPS control traffic.



Note

Other VLAN types such as VMAN, SVLAN, CVLAN and BVLAN will not be used for control traffic.

A control VLAN cannot be deleted from a ring that has CFM configured.

Example

The following command deletes a control VLAN named “vlan10” from an ERPS ring named “ring1”:

```
configure erps ring1 delete control vlan vlan10
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps delete protected vlan

```
configure erps ring-name delete protected {vlan} vlan_name
```

Description

Delete a protected data VLAN from the ERPS ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
<i>vlan_name</i>	Alphanumeric string identifying the data VLAN to be deleted from the ERPS ring.

Default

N/A.

Usage Guidelines

Use this command to delete a protected VLAN from the ERPS ring.



Example

The following command deletes a protected VLAN named “vlan10” from an ERPS ring named “ring1”:

```
configure erps ring1 delete protected vlan vlan10
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps name

```
configure erps old-ring-name name new-ring-name
```

Description

Rename the ERPS ring/sub-ring.

Syntax Description

<i>old-ring-name</i>	Alphanumeric string that identifies the ERPS ring.
<i>new-ring-name</i>	New alphanumeric string identifying the ERPS ring.

Default

N/A.

Usage Guidelines

Use this command to rename the ERPS ring or sub-ring.

Example

The following command an ERPS ring from “ring1” to “ring2”:

```
configure erps ring1 name ring2
```

History

This command was first available in ExtremeXOS 15.1.



Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps neighbor port

```
configure erps ring-name neighbor-port port
```

Description

Add RPL (ring protection link) neighbor configuration for the ERPS ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
<i>port</i>	The slot:port number for RPL neighbor.

Default

N/A.

Usage Guidelines

Use this command to add RPL neighbor configuration for the ERPS ring.



Note

This command implicitly makes the node on which it is configured the RPL neighbor.

Example

The following command adds RPL neighbor on port 5 to an ERPS ring named “ring1”:

```
configure erps ring1 neighbor-port 5
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps notify-topology-change

```
configure {erps} ring-name notify-topology-change {eaps} domain_name
```



Description

Add an ERPS sub-ring to the EAPS domain.

Syntax Description

<i>ring-name</i>	Alphanumeric string identifying the ERPS sub-ring.
<i>domain_name</i>	Alphanumeric string identifying the EAPS domain.

Default

N/A.

Usage Guidelines

Use this command to add an ERPS sub-ring to the EAPS domain.

Example

Example output not yet available and will be provided in a future release.

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps protection-port

```
configure erps ring-name protection-port port
```

Description

Add ring protection link (RPL) owner configuration for the ERPS ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
<i>port</i>	The slot:port number for the ring protection link (RPL) owner.

Default

N/A.



Usage Guidelines

Use this command to add ring protection link (RPL) owner configuration for the ERPS ring.



Note

This command implicitly makes the node on which it is configured the RPL owner.

Example

The following command adds RPL owner configuration on port 5 to an ERPS ring named “ring1”:

```
configure erps ring1 protection-port 5
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps revert

```
configure {erps} ring-name revert [ enable | disable ]
```

Description

Add or delete ERPS revert operation along with the “wait-to-restore” time interval.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
enable	Enable revert mode to ERPS ring.
disable	Disable revert mode from ERPS ring.

Default

The default is the revertive mode (enable).

Usage Guidelines

Use this command to enable/disable a G.8032 ring to revert to the original ring protection link (RPL) block state.



Example

The following command disables revert mode from an ERPS ring named “ring1”:

```
configure erps ring1 revert disable
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps ring-ports east | west

```
configure erps ring-name ring-ports [east | west] port
```

Description

Add ring ports on the ERPS ring. The ring ports connect the switch to the ERPS ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
east	Add the ring port to the east port of the switch.
west	Add the ring port to the west port of the switch.
<i>port</i>	The slot:port number for the ring port.

Default

N/A.

Usage Guidelines

Use this command to add ring ports on the ERPS ring. The ring ports can be added to the east or west port of the switch. The ring ports connect the switch to the ERPS ring.

Example

The following command adds port 5 as a ring port on the east port of the switch for an ERPS ring named “ring1”:

```
configure erps ring1 add ring-ports east 5
```



History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps subring-mode

```
configure erps ring_name subring-mode [no-virtualChannel | virtualChannel]
```

Description

Configures sub-ring mode.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
no-virtualChannel	No Virtual Channel required to complete it's control path.
virtualChannel	Virtual Channel required to complete it's control path.

Default

N/A.

Usage Guidelines

Use this command to add or delete ERPS sub-rings.

Example

The following example configures a virtual channel for the control path:

```
configure erps ring1 subring-mode virtualChannel
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms that are running ExtremeXOS.



configure erps timer guard

```
configure {erps} ring-name timer guard [ default | milliseconds ]
```

Description

Configure a guard timer to control when the node should act on received R-APS (ring automatic protection switching) messages.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
default	The default value, 500 milliseconds.
<i>milliseconds</i>	The interval for the guard timer in milliseconds, with a range of 10 to 2000.

Default

The default is 500 milliseconds.

Usage Guidelines

Use this command to configure a guard timer to control when the node should act on received R-APS messages.

Example

The following command sets the guard timer to 1000 milliseconds for an ERPS ring named "ring1":

```
configure erps ring1 timer guard 1000
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps timer hold-off

```
configure {erps} ring-name timer hold-off [ default | milliseconds ]
```

Description

Configure a hold-off timer to control when a signal fault is relayed.



Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
default	The default value, 0 milliseconds.
<i>milliseconds</i>	The interval for the hold-off time in milliseconds, with a range of 0 to 10000.

Default

The default is 0 milliseconds.

Usage Guidelines

Use this command to configure a hold-off timer to control when a signal fault is relayed.

Example

The following command sets the hold-off timer to 1000 milliseconds for an ERPS ring named "ring1":

```
configure erps ring1 timer hold-off 1000
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps timer periodic

```
configure {erps} ring-name timer periodic [ default | milliseconds ]
```

Description

Configure a periodic timer to control the interval between signal failures.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
default	The default value, 5000 milliseconds.
<i>milliseconds</i>	The interval for the periodic time in milliseconds, with a range of 2000 to 7000.



Default

The default is 5000 milliseconds.

Usage Guidelines

Use this command to configure a periodic timer to control the interval between signal failure.

Example

The following command sets the periodic timer to 6000 milliseconds for an ERPS ring named “ring1”:

```
configure erps ring1 timer periodic 6000
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps timer wait-to-block

```
configure {erps} ring-name timer wait-to-block [ default | milliseconds]
```

Description

Configure a wait-to-block timer for revertive operations on RPL owner initiated reversion.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
default	The default value, 5000 milliseconds.
<i>milliseconds</i>	The time interval to wait before restoring, with a range of 5000 to 7000 milliseconds.

Default

The default is 5000 milliseconds.

Usage Guidelines

Use this command to configure a wait-to-block timer for revertive operations on RPL owner-initiated reversion.



Example

The following command sets the wait-to-block timer to 6000 milliseconds for an ERPS ring named “ring1”:

```
configure erps ring1 timer wait-to-block 6000
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps timer wait-to-restore

```
configure {erps} ring-name timer wait-to-restore [ default | milliseconds ]
```

Description

Configure a time interval to wait before restoring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
default	The default value, 1000 milliseconds.
<i>milliseconds</i>	The time interval to wait before restoring, with a range of 0 to 720000 milliseconds.

Default

The default is 1000 milliseconds.

Usage Guidelines

Use this command to configure a time interval to wait before restoring.

Example

The following command sets the wait-to-restore timer to 3000 milliseconds for an ERPS ring named “ring1”:

```
configure erps ring1 timer wait-to-restore 3000
```



History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps topology-change

```
configure erps ring-name [add | delete] topology-change ring-list
```

Description

Identify the rings to which topology change events need to be propagated.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
add	Add rings/sub-rings to topology change propagation list.
delete	Delete rings/sub-rings from topology change propagation list.
<i>ring-list</i>	List of ERPS rings/sub-rings to which topology change needs to be propagated.

Default

N/A.

Usage Guidelines

Use this command to add or delete ERPS rings/sub-rings from the topology change propagation list.

Example

Example output not yet available and will be provided in a future release.

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.



create erps ring

create erps *ring-name*

Description

Creates an ERPS ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
------------------	--

Default

N/A.

Usage Guidelines

Use this command to create an ERPS ring.

Example

The following command creates an ERPS ring named “ring1”:

```
create erps ring1
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

debug erps

debug erps [*options*]

Description

Debugs an ERPS ring.



Syntax Description

options	Different options to enable looking at debug information.
----------------	---

Default

N/A.

Usage Guidelines

Use this command to debug an ERPS ring.

Example

The following command debugs an ERPS ring:

```
debug erps [options]
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

debug erps show

```
debug erps show ring-name
```

Description

Debugs ERPS ring by checking "show" output.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
------------------	--

Default

N/A.

Usage Guidelines

To debug this feature, check the output of "show erps" and "show erps ring" to see if the node state is as expected. In steady state, the node should be in "Idle" or "Protected" state.



Check the output of "show erps ring statistics" to see if any error/dropped counters are incrementing. If they are check the state of the ring ports and trace these links to the neighbor node to see the state of the links. The output of "show log" after turning on the filters for ERPS should provide more information on what is happening on the switch.

Example

Example output not yet available and will be provided in a future release.

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

delete erps

```
delete erps ring-name
```

Description

Deletes an ERPS ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
------------------	--

Default

N/A.

Usage Guidelines

Use this command to delete an ERPS ring.

Example

The following command deletes an ERPS ring named "ring1":

```
delete erps ring1
```



History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

disable erps

disable erps

Description

Disable ERPS (Ethernet Ring Protection Switching/ITU-T G.8032 standard).

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

Use this command to disable ERPS.

Example

The following command disables ERPS:

```
disable erps
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

disable erps block-vc-recovery

disable erps *ring-name* **block-vc-recovery**



Description

Disables the ability on ERPS rings to block virtual channel recovery to avoid temporary loops .

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
block-vc-recovery	Block on Virtual channel recovery.

Default

N/A.

Usage Guidelines

Use this command to disable the ability on ERPS rings to block on virtual channel recovery to avoid temporary loops. This is done on interconnected nodes for sub-ring configurations.

Example

The following example disables a virtual channel recovery block on “ring1”:

```
diabale erps ring1 block-vc-recovery
```

History

This command was first available in ExtremeXOS 15.13.

Platform Availability

This command is available on all platforms that are running ExtremeXOS.

disable erps ring-name

```
disable erps ring-name
```

Description

Disable an existing ERPS ring/sub-ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
------------------	--



Default

N/A.

Usage Guidelines

Use this command to disable an existing ERPS ring/sub-ring.

Example

The following example disables an existing ERPS ring identified as “ring1”:

```
disable erps ring1
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

disable erps topology-change

```
disable erps ring-name topology-change
```

Description

Disable the ability of ERPS to set the topology-change bit to send out Flush events.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS sub-ring.
topology-change	Topology change propagation control.

Default

N/A.

Usage Guidelines

Use this command to disable the ability of ERPS to set the topology-change bit to send out Flush events.



Example

The following example disables the ability to set the topology-change bit for an existing ERPS sub-ring identified as “ring1”:

```
disable erps ring1 topology-change
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

enable erps

enable erps

Description

Enable ERPS (Ethernet Ring Protection Switching/ITU-T G.8032 standard).

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

Use this command to enable ERPS.

Example

```
enable erps
```

History

This command was first available in ExtremeXOS 15.1.



Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

enable erps block-vc-recovery

```
enable erps ring-name block-vc-recovery
```

Description

Enable ability on ERPS rings to block virtual channel recovery to avoid temporary loops .

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
block-vc-recovery	Block on Virtual channel recovery.

Default

N/A.

Usage Guidelines

Use this command to enable ability on ERPS rings to block on virtual channel recovery to avoid temporary loops. This is done on interconnected nodes for sub-ring configurations.

Example

The following example enables a virtual channel recovery block on “ring1”:

```
enable erps ring1 block-vc-recovery
```

History

This command was first available in ExtremeXOS 15.13.

Platform Availability

This command is available on all platforms that are running ExtremeXOS.

enable erps ring-name

```
enable erps ring-name
```



Description

Enable an existing ERPS ring/sub-ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
------------------	--

Default

N/A.

Usage Guidelines

Use this command to enable an existing ERPS ring/sub-ring.

Example

The following example enables an existing ERPS ring identified as “ring1”:

```
enable erps ring1
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

enable erps topology-change

```
enable erps ring-name topology-change
```

Description

Enable the ability of ERPS to set the topology-change bit to send out Flush events.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS sub-ring.
topology-change	Topology change propagation control.



Default

N/A.

Usage Guidelines

Use this command to enable the ability of ERPS to set the topology-change bit to send out Flush events.

Example

The following example enables the ability to set the topology-change bit for an existing ERPS sub-ring identified as “ring1”:

```
enable erps ring1 topology-change
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

run erps force-switch | manual-switch

```
run erps ring-name [force-switch | manual-switch] {port} port
```

Description

Set up force and manual switch triggers to the ERPS ring/sub-ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
force-switch	Force switch operation.
manual-switch	Manual switch operation.
<i>port</i>	The slot:port number for the ring port.

Default

N/A.



Usage Guidelines

Use this command to set up force and manual switch triggers to the ERPS ring/sub-ring.

Example

The following command sets up force switch operation on port 6 of an ERPS ring named “ring1”:

```
run erps ring1 force-switch port 6
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

show erps

show erps

Description

Display global information for ERPS.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

Use this command to display global information for ERPS.

Example

```
# show erps

ERPS Enabled: Yes
ERPS Display Config Warnings: On
ERPS Multicast Add Ring Ports: Off
ERPS Multicast Send IGMP Query: On
```




```
ERPS Multicast Temporary Flooding: Off
ERPS Multicast Temporary Flooding Duration: 15 sec
Number of ERPS instances: 1
# ERPS ring configuration :
```

```
-----
--
Ring          State          Type    East    West    Control-Vlan  VID
-----
R1            Protection    R r     21      +20     cv1           (1000)
-----
```

```
--
where State: Init/Idle/Protection/Manual-Switch/Force-Switch/Pending
      Type: (I) Interconnected node, (N) RPL Neighbor,
            R) RPL Owner, (X) Ring node
      Flags: (n) Non-revertive, (r) Revertive,
            (+) RPL Protection Port, (^) RPL Neighbor Port
            (f) Force Switch Port, (m) Manual Switch Port
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

show erps ring-name

```
show erps ring-name
```

Description

Display specific details for an ERPS ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
------------------	--

Default

N/A.

Usage Guidelines

Use this command to display specific details for an ERPS ring.



Example

The following example displays details for an ERPS ring named "R1":

```
# show erps "R1"

Name: R1
Operational State: Protection enabled          Node Type: RPL Owner, Revertive
Configured State : Enabled

East Ring Port : 21  MepId: 1  Remote MepId: 3      Status: Blocked
West Ring Port : +20 MepId: 2  Remote MepId: 4      Status: Blocked

Periodic timer interval: 5000 millisec (Enabled)
Hold-off timer interval: 0 millisec (Enabled)
Guard timer interval   : 500 millisec (Enabled)
WTB timer interval     : 5500 millisec (Enabled)
WTR timer interval     : 1000 millisec (Enabled)

Ring MD Level          : 1
CCM Interval East     : 1000 millisec
CCM Interval West     : 1000 millisec
Notify Topology Change : -----
Subring Mode          : Virtual Channel

ERPS Control Vlan: cv1          VID:1000
Topology Change Propagation List: None
Topology Change Propagation : Disabled
ERPS Ring's Sub-Ring(s): None
ERPS Ring has following Protected Vlan(s):
  Vlan Name          VID
  pvl                1001
Number of Protected Vlans: 1
(+) RPL Protection Port, (^) RPL Neighbor Port
(f) Force Switch Port, (m) Manual Switch Port
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

show erps statistics

```
show erps ring-name statistics
```

Description

Display control packet and event statistics for an ERPS ring.



Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
------------------	--

Default

N/A.

Usage Guidelines

Use this command to display control packet and event statistics for an ERPS ring.

Example

The following example displays statistics for an ERPS ring named "R1":

```
# show erps "R1" statistics
port      Sent      Received  Dropped  Blocked  Un-blocked  SF  SF-clear
          R-APS   R-APS    R-APS    events  events
-----
2:1       2309    3400      4         5         0         0       0
1:20      100     45        0         0        10       2000    100
-----
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

unconfigure erps cfm

```
unconfigure {erps} ring-name cfm
```

Description

Unconfigure the CFM maintenance association for the ERPS ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
------------------	--



Default

N/A.

Usage Guidelines

Use this command to unconfigure connectivity fault management (CFM) for the ERPS ring.

Example

The following command unconfigures connectivity fault management on an ERPS ring named “ring1”:

```
unconfigure erps ring1 cfm
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

unconfigure erps neighbor-port

```
unconfigure erps ring-name neighbor-port
```

Description

Delete the ring protection link (RPL) neighbor configuration for the ERPS ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
------------------	--

Default

N/A.

Usage Guidelines

See Description.



Example

The following command deletes RPL neighbor configuration for the ERPS ring named “ring1”:

```
unconfigure erps ring1 neighbor-port
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

unconfigure erps notify-topology-change

```
unconfigure {erps} ring-name notify-topology-change {eaps} domain_name
```

Description

Delete an ERPS sub-ring from the EAPS domain.

Syntax Description

<i>ring-name</i>	Alphanumeric string identifying the ERPS sub-ring.
<i>domain_name</i>	Alphanumeric string identifying the EAPS domain.

Default

N/A.

Usage Guidelines

Use this command to delete an ERPS sub-ring from the EAPS domain.

Example

Example output not yet available and will be provided in a future release.

History

This command was first available in ExtremeXOS 15.1.



Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

unconfigure erps protection-port

```
unconfigure erps ring-name protection-port
```

Description

Delete ring protection link (RPL) owner configuration for the ERPS ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
------------------	--

Default

N/A.

Usage Guidelines

Use this command to delete ring protection link (RPL) owner configuration for the ERPS ring.

Example

The following command deletes RPL owner configuration on an ERPS ring named “ring1”:

```
unconfigure erps ring1 protection-port
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

unconfigure erps ring-ports west

```
unconfigure erps ring-name ring-ports west
```

Description

Delete ring ports on the ERPS ring.



Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
west	Delete the ring port on the west port of the switch.

Default

N/A.

Usage Guidelines

Use this command to delete ring ports on the ERPS ring. Ring ports are the ports of the switch that connect it to the ERPS ring. This command deletes the ring port on the west port of the switch.



Note

On unconfiguring the west port, the node is treated as an interconnected node.

Example

The following command deletes the ring ports on the west port of the switch for an ERPS ring named “ring1”:

```
unconfigure erps ring1 ring-ports west
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps cfm protection group

```
configure erps ring_name cfm protection [add delete] group cfm_group
```

Description

Associates or disassociates a CFM UP MEP group for subring protection across the main ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
east	East port.



west	West port.
add	Associates a CFM Up-MEP entity.
delete	Disassociates a CFM Up-MEP entity.
group	Specifies a CFM Up-MEP group.
<i>group_name</i>	Specifies the name of the Up MEP group.

Default

N/A.

Usage Guidelines

Use this command to associate or disassociate a CFM UP MEP group for subring protection across the main ring.

Example

The following command associates a CFM UP MEP group for subring protection on the group "group1":

```
configure erps ring1 cfm protection add group1
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms running ExtremeXOS.



30 STP Commands

STP

RSTP

MSTP

Spanning Tree Domains

STP Rules and Restrictions

clear counters stp

configure mstp format

configure mstp region

configure mstp revision

configure stpd add vlan

configure stpd default-encapsulation

configure stpd delete vlan

configure stpd description

configure stpd flush-method

configure stpd forwarddelay

configure stpd hellotime

configure stpd maxage

configure stpd max-hop-count

configure stpd mode

configure stpd ports active-role disable

configure stpd ports active-role enable

configure stpd ports bpdu-restrict

configure stpd ports cost

configure stpd ports edge-safeguard disable

configure stpd ports edge-safeguard enable

configure stpd ports link-type

configure stpd ports mode

configure stpd ports port-priority

configure stpd ports priority

configure stpd ports restricted-role disable

configure stpd ports restricted-role enable

configure stpd priority

configure stpd tag

configure vlan add ports stpd

create stpd

delete stpd

disable stpd

```
disable stpd auto-bind
disable stpd ports
disable stpd rapid-root-failover
enable stpd
enable stpd auto-bind
enable stpd ports
enable stpd rapid-root-failover
show stpd
show stpd ports
show vlan stpd
unconfigure mstp region
unconfigure stpd
unconfigure stpd ports link-type
```

This chapter describes commands for:

- Creating, configuring, enabling, and disabling Spanning Tree Protocol (STP) on the switch.
- Enabling and disabling Rapid Spanning Tree Protocol (RSTP) on the switch.
- Enabling and disabling Multiple Spanning Tree Protocol (MSTP) on the switch.
- Displaying and resetting STP settings on the switch.

STP

STP is a bridge-based mechanism for providing fault tolerance on networks. STP is a part of the 802.1D bridge specification defined by the IEEE Computer Society. To explain STP in terms used by the 802.1D specification, the switch is referred to as a bridge.

STP allows you to implement parallel paths for network traffic and ensure that redundant paths are:

- Disabled when the main paths are operational.
- Enabled if the main path fails.



Note

STP and Extreme Standby Router Protocol (ESRP) cannot be configured on the same Virtual LAN (VLAN) simultaneously.

RSTP

The Rapid Spanning Tree Protocol (RSTP) IEEE 802.1w provides an enhanced spanning tree algorithm that improves the convergence speed of bridged networks.

RSTP takes advantage of point-to-point links in the network and actively confirms that a port can safely transition to the forwarding state without relying on any timer configurations. If a network topology change or failure occurs, RSTP rapidly recovers network connectivity by confirming the change locally before propagating that change to other devices across the network. For broadcast links, there is no difference in convergence time between STP and RSTP.



RSTP supersedes legacy STP protocols, supports the existing STP parameters and configurations, and allows for seamless interoperability with legacy STP.

MSTP

MSTP logically divides a Layer 2 network into regions.

Each region has a unique identifier and contains multiple spanning tree instances (MSTIs). An MSTI is a spanning tree domain that operates within and is bounded by a region. MSTIs control the topology inside the regions. The Common and Internal Spanning Tree (CIST) is a single spanning tree domain that interconnects MSTP regions. The CIST is responsible for creating a loop-free topology by exchanging and propagating BPDUs across regions to form a Common Spanning Tree (CST).

MSTP uses RSTP as its converging algorithm and is interoperable with the legacy STP protocols: STP (802.1D) and RSTP (802.1w).

Spanning Tree Domains

The switch can be partitioned into multiple virtual bridges. Each virtual bridge can run an independent spanning tree instance. Each spanning tree instance is called a Spanning Tree Domain (STPD). Each STPD has its own root bridge and active path. After an STPD is created, one or more VLANs can be assigned to it.

A port can belong to multiple STPDs. In addition, a VLAN can span multiple STPDs.

The key points to remember when configuring VLANs and STP are:

- Each VLAN forms an independent broadcast domain.
- STP blocks paths to create a loop-free environment.
- Within any given STPD, all VLANs belonging to it use the same spanning tree.

Member VLANs

When you add a VLAN to an STPD, that VLAN becomes a member of the STPD. The two types of member VLANs in an STPD are:

- Carrier.
- Protected.

Carrier VLAN

A carrier VLAN defines the scope of the STPD, which includes the physical and logical ports that belong to the STPD and if configured, the 802.1Q tag used to transport Extreme Multiple Instance Spanning Tree Protocol (EMISTP) or Per VLAN Spanning Tree (PVST+) encapsulated Bridge Protocol Data Units



(BPDUs). Only one carrier VLAN can exist in a given STPD, although some of its ports can be outside the control of any STPD at the same time.

**Note**

If you use EMISTP or PVST+, the STPD ID must be identical to the VLAN ID of the carrier VLAN in that STPD.

If you have an 802.1D configuration, we recommend that you configure the StpdID to be identical to the VLAN ID of the carrier VLAN in that STPD.

If you configure MSTP, you do not need carrier VLANs for MSTP operation. With MSTP, you configure a CIST that controls the connectivity of interconnecting MSTP regions and sends BPDUs across the regions to communicate the status of MSTP regions. All VLANs participating in the MSTP region have the same privileges.

Protected VLAN

Protected VLANs are all other VLANs that are members of the STPD. These VLANs “piggyback” on the carrier VLAN. Protected VLANs do not transmit or receive STP BPDUs, but they are affected by STP state changes and inherit the state of the carrier VLAN. Protected VLANs can participate in multiple STPD, but any particular port in the VLAN can belong to only one STPD. Also known as non-carrier VLANs.

If you configure MSTP, all member VLANs in an MSTP region are protected VLANs. These VLANs do not transmit or receive STP BPDUs, but they are affected by STP state changes communicated by the CIST to the MSTP regions. MSTIs cannot share the same protected VLAN; however, any port in a protected VLAN can belong to multiple MSTIs.

STPD Modes

An STPD has three modes of operation:

- 802.1D mode

Use this mode for backward compatibility with previous STP versions and for compatibility with third-party switches using IEEE standard 802.1D. When configured in this mode, all rapid configuration mechanisms are disabled.

- 802.1w mode

Use this mode for compatibility with Rapid Spanning Tree (RSTP). When configured in this mode, all rapid configuration mechanisms are enabled. The benefit of this mode is available on point-to-point and edge ports only.

You enable or disable RSTP on a per STPD basis only. You do not enable RSTP on a per port basis.

- MSTP mode

Use this mode for compatibility with Multiple Spanning Tree (MSTP, 802.1s). MSTP is an extension of RSTP and offers the benefit of better scaling with fast convergence. When configured in this mode, all rapid configuration mechanisms are enabled. The benefit of MSTP is available only on point-to-point links and when you configure the peer in MSTP or 802.1w mode. If you do not select point-to-point links and the peer is not configured in 802.1w mode, the STPD fails back to 802.1D mode.



You can create only one MSTP region on the switch, and all switches that participate in the region must have the same regional configurations. You enable or disable an MSTP on a per STPD basis only. You do not enable MSTP on a per port basis.

By default, the:

- STPD operates in 802.1D mode.
- Default device configuration contains a single STPD called s0.
- Default VLAN is a member of STPD s0 with autobind enabled.

All STP parameters default to the IEEE 802.1D values, as appropriate.

Encapsulation Modes

You can configure ports within an STPD to accept and transmit specific BPDU encapsulations. This STP port encapsulation is separate from the STP mode of operation. For example, you can configure a port to accept the PVST+ BPDU encapsulation while running in 802.1D mode.

An STP port has three possible encapsulation modes:

- 802.1D mode

This mode is used for backward compatibility with previous STP versions and for compatibility with third-party switches using IEEE standard 802.1D. BPDUs are sent untagged in 802.1D mode. Because of this, any given physical interface can have only one STPD running in 802.1D mode.

This encapsulation mode supports the following STPD modes of operation: 802.1D, 802.1w, and MSTP.

- Extreme Multiple Instance Spanning Tree Protocol (EMISTP) mode

EMISTP mode is proprietary to Extreme Networks and is an extension of STP that allows a physical port to belong to multiple STPDs by assigning the port to multiple VLANs. EMISTP adds significant flexibility to STP network design. BPDUs are sent with an 802.1Q tag having an STPD instance Identifier (STPD ID) in the VLAN ID field.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

- Per VLAN Spanning Tree (PVST+) mode

This mode implements PVST+ in compatibility with third-party switches running this version of STP. The STPDs running in this mode have a one-to-one relationship with VLANs, and send and process packets in PVST+ format.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

These encapsulation modes are for STP ports, not for physical ports. When a physical port belongs to multiple STPDs, it is associated with multiple STP ports. It is possible for the physical port to run in different modes for different domains to which it belongs.

MSTP STPDs use only 802.1D BPDU encapsulation mode. The switch prevents you from configuring EMISTP or PVST+ encapsulation mode for MSTP STPDs.



STP Rules and Restrictions

This section summarizes the rules and restrictions for configuring STP as follows:

- The carrier VLAN must span all ports of the STPD. (This is not applicable to MSTP.)
- The STPD ID must be the VLAN ID of the carrier VLAN; the carrier VLAN cannot be partitioned. (This is not applicable to MSTP.)
- A default VLAN cannot be partitioned. If a VLAN traverses multiple STPDs, the VLAN must be tagged.
- An STPD can carry, at most, one VLAN running in PVST+ mode, and its STPD ID must be identical with that VLAN ID. In addition, the PVST+ VLAN cannot be partitioned.
- The default VLAN of a PVST+ port must be identical with the native VLAN on the PVST+ device connected to that port.
- If an STPD contains both PVST+ and non-PVST+ ports, that STPD must be enabled. If that STPD is disabled, the BPDUs are flooded in the format of the incoming STP port, which may be incompatible with those of the connected devices.
- The 802.1D ports must be untagged; and the EMISTP/PVST+ ports must be tagged in the carrier VLAN.
- An STPD with multiple VLANs must contain only VLANs that belong to the same virtual router instance.
- STP and network login operate on the same port as follows:
 - STP (802.1D), RSTP (802.1W), and MSTP (802.1S) support both network login and STP on the same port.
 - At least one VLAN on the intended port should be configured both for STP and network login.
 - When STP blocks a port, network login does not process authentication requests and BPDUs are the only traffic in and out of the port. All user data forwarding stops.
 - When STP places a port in forwarding state, network login operates and BPDUs and user data flow in and out of the port. The forwarding state is the only STP state that allows network login and user data forwarding.
 - When RSTP is used with network login campus mode, autobind must be enabled on all VLANs that support RSTP and network login campus mode.
 - When RSTP is used with network login campus mode on a port, dynamic VLANs cannot be supported.
- STP cannot be configured on the following ports:
 - A mirroring target port.
 - A software-controlled redundant port.
- MSTP and 802.1D STPDs cannot share a physical port.
- Only one MSTP region can be configured on a switch.
- In an MSTP environment, A VLAN can belong to either a CIST or one of the MSTIs.
- A VLAN can belong to only one MSTP domain.
- MSTP is not interoperable with PVST+.
- The CIST can operate without any member VLANs.



clear counters stp

```
clear counters stp {[all | diagnostics | domains | ports]}
```

Description

Clears, resets all STP statistics and counters.

Syntax Description

all	Specifies all STP domain, port, and diagnostics counters.
diagnostics	Specifies STP diagnostics counters.
domains	Specifies STP domain counters.
ports	Specifies STP port counters.

Default

N/A.

Usage Guidelines

If you do not enter a parameter, the result is the same as specifying the all parameter: the counters for all domains, ports, and diagnostics are reset.

Enter one of the following parameters to reset the STP counters on the switch:

- all—Specifies the counters for all STPDs and ports, and clears all STP counters.
- diagnostics—Clears the internal diagnostic counters.
- domains—Clears the domain level counters.
- ports—Clears the counters for all ports and leaves the domain level counters.

Viewing and maintaining statistics on a regular basis allows you to see how well your network is performing. If you keep simple daily records, you will see trends emerging and notice problems arising before they cause major network faults. By clearing the counters, you can see fresh statistics for the time period that you are monitoring.

Example

The following command clears all of the STP domain, port, and diagnostic counters:

```
clear counters stp
```

History

This command was first available in ExtremeXOS 10.1.



Platform Availability

This command is available on all platforms.

configure mstp format

```
configure mstp format format_identifier
```

Description

Configures the number used to identify the MSTP BPDUs sent in the MSTP region.

Syntax Description

<i>format_identifier</i>	Specifies a number that MSTP uses to identify all BPDUs sent in the MSTP region. The default is 0. The range is 0 to 255.
--------------------------	---

Default

The default value used to identify the MSTP BPDU is 0.

Usage Guidelines

For a switch to be part of an MSTP region, you must configure each switch in the region with the same MSTP configuration attributes, also known as MSTP region identifiers. These identifiers consist of the following:

- Region Name—The name of the MSTP region.
- Format Selector—The number used to identify the format of MSTP BPDUs. The default is 0.
- Revision Level—This identifier is reserved for future use; however, the switch uses and displays a default of 3.

You can configure only one MSTP region on the switch at any given time.

The switches contained in a region transmit and receive BPDUs that contain information relevant to only that MSTP region. By having devices look at the region identifiers, MSTP discovers the logical boundary of a region.

If you have an active MSTP region, Extreme Networks recommends that you disable all active STPDs in the region before modifying the value used to identify MSTP BPDUs on all participating switches.

Example

The following command configures the number 2 to identify the MSTP BPDUs sent within an MSTP region:

```
configure mstp format 2
```



History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure mstp region

```
configure mstp region regionName
```

Description

Configures the name of an MSTP region on the switch.

Syntax Description

<i>regionName</i>	Specifies a user-defined name for the MSTP region. May be up to 32 characters.
-------------------	--

Default

By default, the switch uses the MAC address of the switch to generate an MSTP region.

Before you configure the MSTP region, it also has the following additional defaults:

- MSTP format Identifier—0.
- MSTP Revision Level—3.

Usage Guidelines

The maximum length for a name is 32 characters. Names can contain alphanumeric characters and underscores (`_`) but cannot be any reserved keywords, for example, `mstp`. Names must start with an alphabetical character, for example, `a`, `Z`.

By default, the switch uses the unique MAC address of the switch to generate an MSTP region. Since each MAC address is unique, every switch is in its own region by default.

For multiple switches to be part of an MSTP region, you must configure each switch in the region with the same MSTP configuration attributes, also known as MSTP region identifiers. These identifiers consist of the following:

- Region Name—The name of the MSTP region.
- Format Selector—The number used to identify the format of MSTP BPDUs. The default is 0.
- Revision Level—This identifier is reserved for future use; however, the switch uses and displays a default of 3.

You can configure only one MSTP region on the switch at any given time.



The switches inside a region exchange BPDUs that contain information for MSTIs. The switches connected outside of the region exchange CIST information. By having devices look at the region identifiers, MSTP discovers the logical boundary of a region.

If you have an active MSTP region, we recommend that you disable all active STPDs in the region before renaming the region on all of the participating switches.

Viewing MSTP Information

To view the MSTP configuration on the switch, use the `show stpd` command. Output from this command contains global MSTP settings, including the name of the MSTP region, the number or tag that identifies all of the BPDUs sent in the MSTP region, and the reserved MSTP revision level. If configured, the output also displays the name of the Common and Internal Spanning Tree (CIST), and the number of Multiple Spanning Tree Instances (MSTIs).

Example

The following command creates an MSTP region named purple:

```
configure mstp region purple
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure mstp revision

```
configure mstp revision revision
```

Description

Configures the revision number of the MSTP region.

Syntax Description

<i>revision</i>	This parameter is reserved for future use.
-----------------	--

Default

The default value of the revision level is 3.



Usage Guidelines

Although this command is displayed in the CLI, it is reserved for future use. Please do not use this command.

If you accidentally configure this command, remember that each switch in the region must have the same MSTP configuration attributes, also known as MSTP region identifiers. These identifiers consist of the following:

- Region Name—The name of the MSTP region.
- Format Selector—The number used to identify the format of MSTP BPDUs. The default is 0.
- Revision Level—This identifier is reserved for future use; however, the switch uses and displays a default of 3.

Example

The following command returns the MSTP revision number to 3, the default revision number:

```
configure mstp revision 3
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure stpd add vlan

```
configure stpd stpd_name add vlan vlan_name ports [all | port_list] {[dot1d | emistp | pvst-plus]}
```

Description

Adds all ports or a list of ports within a VLAN to a specified STPD.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>vlan_name</i>	Specifies a VLAN name.
all	Specifies all of the ports in the VLAN to be included in the STPD.
<i>port_list</i>	Specifies the port or ports to be included in the STPD.
dot1d	Specifies the STP encapsulation mode of operation to be 802.1D.



emistp	Specifies the STP encapsulation mode of operation to be EMISTP.
pvst-plus	Specifies the STP encapsulation mode of operation to be PVST+.

Default

Ports in the default STPD (s0) are in dot1d mode.

Ports in user-created STPDs are in `emistp` mode.

Usage Guidelines

To create an STP domain, use the `create stpd` command. To create a VLAN, use the `create vlan` command.

In an EMISTP or PVST+ environment, this command adds a list of ports within a VLAN to a specified STPD provided the carrier VLAN already exists on the same set of ports. You can also specify the encapsulation mode for those ports.

In an MSTP environment, you do not need a carrier VLAN. A CIST controls the connectivity of interconnecting MSTP regions and sends BPDUs across the regions to communicate region status. You must use the dot1d encapsulation mode in an MSTP environment.

You cannot configure STP on the following ports:

- Mirroring target ports.
- Software-controlled redundant ports.

If you see an error similar to the following:

```
Error: Cannot add VLAN default port 3:5 to STP domain
```

You might be attempting to add:

- A carrier VLAN port to a different STP domain than the carrier VLAN belongs.
- A VLAN/port for which the carrier VLAN does not yet belong.



Note

This restriction is enforced only in an active STP domain and when you enable STP to make sure you have a legal STP configuration.

Care must be taken to ensure that ports in overlapping domains do not interfere with the orderly working of each domain's protocol.

By default, when the switch boots for the first time, it automatically creates a VLAN named default with a tag value of 1 and STPD s0. The switch associates VLAN default to STPD s0. All ports that belong to this VLAN and STPD are in 802.1D encapsulation mode with autobind enabled. If you disable autobind on the VLAN default, that configuration is saved across a reboot.



Naming Conventions

If your STPD has the same name as another component, for example a VLAN, we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keywords `stpd` and `vlan` are optional.

STP Encapsulations Modes

You can specify the following STP encapsulation modes:

- `dot1d`—This mode is reserved for backward compatibility with previous STP versions. BPDUs are sent untagged in 802.1D mode. Because of this, any given physical interface can have only one STPD running in 802.1D mode.

This encapsulation mode supports the following STPD modes of operation: 802.1D, 802.1w, and MSTP.

- `emistp`—This mode sends BPDUs with an 802.1Q tag having an STPD ID in the VLAN ID field.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

- `pvst-plus`—This mode implements PVST+ in compatibility with third-party switches running this version of STP. The STPDs running in this mode have a one-to-one relationship with VLANs, and send and process packets in PVST+ format.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

These encapsulation modes are for STP ports, not for physical ports. When a physical port belongs to multiple STPDs, it is associated with multiple STP ports. It is possible for the physical port to run in different modes for different domains for which it belongs.

MSTP STPDs use 802.1D BPDU encapsulation mode by default. To ensure correct operation of your MSTP STPDs, do not configure EMISTP or PVST+ encapsulation mode for MSTP STPDs.

STPD Identifier

An `StpdID` is used to identify each STP domain. You assign the `StpdID` when configuring the domain. An STPD ID must be identical to the VLAN ID of the carrier VLAN in that STPD and that VLAN cannot belong to another STPD.

MSTP uses two different methods to identify the STPDs that are part of the MSTP network. An instance ID of 0 identifies the Common and Internal Spanning Tree (CIST). The switch assigns this ID automatically when you configure the CIST STPD. A multiple spanning tree instance identifier identifies each STP domain that is part of an MSTP region. You assign the MSTI ID when configuring the STPD that participates in the MSTP region. In an MSTP region, MSTI IDs only have local significance. You can reuse MSTI IDs across MSTP regions.

Automatically Inheriting Ports--MSTP Only

In an MSTP environment, whether you manually or automatically bind a port to an MSTI in an MSTP region, the switch automatically binds that port to the CIST. The CIST handles BPDU processing for itself and all of the MSTIs; therefore, the CIST must inherit ports from the MSTIs in order to transmit and receive BPDUs.



Example

Create a VLAN named marketing and an STPD named STPD1 as follows:

```
create vlan marketing
create stpd stpd1
```

The following command adds the VLAN named marketing to the STPD STPD1, and includes all the ports of the VLAN in STPD1:

```
configure stpd stpd1 add vlan marketing ports all
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure stpd default-encapsulation

```
configure stpd stpd_name default-encapsulation [dot1d | emistp | pvst-plus]
```

Description

Configures the default encapsulation mode for all ports added to the specified STPD.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
dot1d	Specifies the STP encapsulation mode of operation to be 802.1d.
emistp	Specifies the STP encapsulation mode of operation to be EMISTP.
pvst-plus	Specifies the STP encapsulation mode of operation to be PVST+.

Default

Ports in the default STPD (s0) are dot1d mode.

Ports in user-created STPDs are in **emistp** mode.

Usage Guidelines

Care must be taken to ensure that ports in overlapping domains do not interfere with the orderly working of each domain's protocol.



By default, when the switch boots for the first time, it automatically creates a VLAN named default with a tag value of 1 and STPD s0. The switch associates VLAN default to STPD s0. All ports that belong to this VLAN and STPD are in 802.1d encapsulation mode with autobind enabled. If you disable autobind on the VLAN default, that configuration is saved across a reboot.

MSTP STPDs use 802.1D BPDU encapsulation mode by default. To ensure correct operation of your MSTP STPDs, do not configure EMISTP or PVST+ encapsulation mode for MSTP STPDs.

Naming Conventions

If your STPD has the same name as another component, for example a VLAN, we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional. For name creation guidelines and a list of reserved names, see [Object Names](#) in the *ExtremeXOS Concepts Guide*.

STP Encapsulation Modes

You can specify the following STP encapsulation modes:

- `dot1d`—This mode is reserved for backward compatibility with previous STP versions. BPDUs are sent untagged in 802.1D mode. Because of this, any given physical interface can have only one STPD running in 802.1D mode.

This encapsulation mode supports the following STPD modes of operation: 802.1D, 802.1w, and MSTP.

- `emistp`—This mode sends BPDUs with an 802.1Q tag having an STPD ID in the VLAN ID field.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

- `pvst-plus`—This mode implements PVST+ in compatibility with third-party switches running this version of STP. The STPDs running in this mode have a one-to-one relationship with VLANs and send and process packets in PVST+ format.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

Note



These encapsulation modes are for STP ports, not for physical ports. When a physical port belongs to multiple STPDs, it is associated with multiple STP ports. It is possible for the physical port to run in different modes for different domains for which it belongs.

STPD Identifier

An `StpdID` is used to identify each STP domain. You assign the `StpdID` when configuring the domain. An STPD ID must be identical to the VLAN ID of the carrier VLAN in that STP domain, and that VLAN cannot belong to another STPD.

MSTP uses two different methods to identify the STPDs that are part of the MSTP network. An instance ID of 0 identifies the Common and Internal Spanning Tree (CIST). The switch assigns this ID automatically when you configure the CIST STPD. A multiple spanning tree instance identifier identifies each STP domain that is part of an MSTP region. You assign the MSTI ID when configuring the STPD



that participates in the MSTP region. In an MSTP region, MSTI IDs only have local significance. You can reuse MSTI IDs across MSTP regions.

Example

The following command specifies that all ports subsequently added to the STPD STPD1 be in PVST+ encapsulation mode unless otherwise specified or manually changed:

```
configure stpd stpd1 default-encapsulation pvst-plus
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure stpd delete vlan

```
configure stpd stpd_name delete vlan vlan_name ports [all | port_list]
```

Description

Deletes one or more ports in the specified VLAN from an STPD.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>vlan_name</i>	Specifies a VLAN name.
all	Specifies that all of the ports in the VLAN are to be removed from the STPD.
<i>port_list</i>	Specifies the port or ports to be removed from the STPD.

Default

N/A.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keywords `stpd` and `vlan` are optional.

In EMISTP and PVST+ environments, if the specified VLAN is the carrier VLAN, all protected VLANs on the same set of ports are also removed from the STPD.



You also use this command to remove autobind ports from a VLAN. ExtremeXOS records the deleted ports so that the ports are not automatically added to the STPD after a system restart.

When a port is deleted on the MSTI, it is automatically deleted on the CIST as well.

Example

The following command removes all ports of a VLAN named Marketing from the STPD STPD1:

```
configure stpd stpd1 delete vlan marketing ports all
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure stpd description

```
configure {stpd} stpd_name description [stpd-description | none]
```

Description

Adds or overwrites the STP domain description field.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>stpd-description</i>	Specifies an STPD description.
none	Clears the STPD string.

Default

The STP domain description string is empty.

Usage Guidelines

Use this command to add or overwrite the STP domain description field.

The maximum STP domain description length is 180 characters.

The *stpd-description* must be in quotes if the string contains any spaces.



To display the description, use the `show stpd stpd_name` command. When no STP domain description is configured, Description is not displayed in the output.

To clear the STP domain description string, either specify the keyword `none` in this command or use the `unconfigure stpd {stpd_name}` command.

Example

The following command adds the description “this is s0 domain” to the STPD named s0:

```
configure stpd s0 description "this is s0 domain"
```

History

This command was first available in ExtremeXOS 12.4.4.

Platform Availability

This command is available on all platforms.

configure stpd flush-method

```
configure stpd flush-method [vlan-and-port | port-only]
```

Description

Configures the method used by STP to flush the FDB during a topology change.

Syntax Description

vlan-and-port	Specifies a VLAN and port combination flush method.
port-only	Specifies a port flush method.

Default

The default flush method is `vlan-and-port`.

Usage Guidelines

For scaled up configurations where there are more than 1000 VLANs and more than 70 ports participating in STP, the number of messages exchanged between STP/FDB/HAL modules can consume a lot of system memory during an STP topology change using the default configuration for flush method. In such situations, setting the flush method to “port-only” can help reduce the system memory consumption.



Example

The following command sets the flush method to port-only:

```
configure stpd flush-method port-only
```

History

This command was available in ExtremeXOS 12.4.5.

Platform Availability

This command is available on all platforms.

configure stpd forwarddelay

```
configure stpd stpd_name forwarddelay seconds
```

Description

Specifies the time (in seconds) that the ports in this STPD spend in the listening and learning states when the switch is the root bridge.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>seconds</i>	Specifies the forward delay time in seconds. The default is 15 seconds, and the range is 4 to 30 seconds.

Default

The default forward delay time is 15 seconds.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

The range for the *seconds* parameter is 4 through 30 seconds.



Example

The following command sets the forward delay from STPD1 to 20 seconds:

```
configure stpd stpd1 forwarddelay 20
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure stpd hellotime

```
configure stpd stpd_name hellotime seconds
```

Description

Specifies the time delay (in seconds) between the transmission of BPDUs from this STPD when it is the root bridge.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>seconds</i>	Specifies the hello time in seconds. The default is 2 seconds, and the range is 1 to 10 seconds.

Default

The default hello time is 2 seconds.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

In an MSTP environment, configure the hello timer only on the CIST, not on the MSTIs.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

The range for the *seconds* parameter is 1 through 10 seconds.



Example

The following command sets the time delay from STPD1 to 10 seconds:

```
configure stpd stpd1 hellotime 10
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure stpd maxage

```
configure stpd stpd_name maxage seconds
```

Description

Specifies the maximum age of a BPDU in the specified STPD.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>seconds</i>	Specifies the maxage time in seconds. The default is 20 seconds, and the range is 6 to 40 seconds.

Default

The default maximum age of a BPDU is 20 seconds.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

In an MSTP environment, configure the maximum age of a BPDU only on the CIST, not on the MSTIs.

The range for the *seconds* parameter is 6 through 40 seconds.

Note that the time must be greater than, or equal to $2 * (\text{Hello Time} + 1)$ and less than, or equal to $2 * (\text{Forward Delay} - 1)$.



Example

The following command sets the maximum age of STPD1 to 30 seconds:

```
configure stpd stpd1 maxage 30
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure stpd max-hop-count

```
configure stpd stpd_name max-hop-count hopcount
```

Description

Specifies the maximum hop count of a BPDU until the BPDU is discarded in the specified MSTP STP domain.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>hopcount</i>	Specifies the number of hops required to age out information and notify changes in the topology. The default is 20 hops, and the range is 6 to 40 hops.

Default

The default hop count of a BPDU is 20 hops.

Usage Guidelines

This command is applicable only in an MSTP environment.

If your STPD has the same name as another component, for example a VLAN, Extreme Networks recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

The range for the *hopcount* parameter is 6 through 40 hops.



In an MSTP environment, the hop count has the same purpose as the maxage timer for 802.1D and 802.1w environments.

The main responsibility of the CIST is to exchange or propagate BPDUs across regions. The switch assigns the CIST an instance ID of 0, which allows the CIST to send BPDUs for itself in addition to all of the MSTIs within an MSTP region. Inside a region, the BPDUs contain CIST records and piggybacked M-records. The CIST records contain information about the CIST, and the M-records contain information about the MSTIs. Boundary ports only exchange CIST record BPDUs.

On boundary ports, only CIST record BPDUs are exchanged. In addition, if the other end is an 802.1D or 802.1w bridge, the maxage timer is used for interoperability between the protocols.

Example

The following command sets the hop of the MSTP STPD, STPD2, to 30 hops:

```
configure stpd stpd2 max-hop-count 30
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure stpd mode

```
configure stpd stpd_name mode [dot1d | dot1w | mstp [cist | msti instance]]
```

Description

Configures the operational mode for the specified STP domain.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
dot1d	Specifies the STPD mode of operation to be 802.1D.
dot1w	Specifies the STPD mode of operation to be 802.1w, and rapid configuration is enabled.
mstp	Specifies the STPD mode of operation to be 802.1s, and rapid configuration is enabled.
cist	Configures the specified STPD as the common instance spanning tree for the MSTP region.



msti	Configures the specified STPD as a multiple spanning tree instance for the MSTP region.
<i>instance</i>	Specifies the Id of the multiple spanning tree instance. The range is 1 to 4,094.

Default

The STPD operates in 802.1D mode.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

If you configure the STP domain in 802.1D mode, the rapid reconfiguration mechanism is disabled.

If you configure the STP domain in 802.1w mode, the rapid reconfiguration mechanism is enabled. You enable or disable RSTP on a per STPD basis only. You do not enable RSTP on a per port basis.

If you configure the STP domain in MSTP mode, the rapid reconfiguration mechanism is enabled. You enable or disable MSTP on a per STPD basis only. You do not enable MSTP on a per port basis. MSTP STPDs use 802.1D BPDU encapsulation mode by default. To ensure correct operation of your MSTP STPDs, do not configure EMISTP or PVST+ encapsulation mode for MSTP STPDs.

You must first configure a Common and Internal Spanning Tree (CIST) before configuring any multiple spanning tree instances (MSTIs) in the region. You cannot delete or disable a CIST if any of the MSTIs are active in the system.

Example

The following command configures STPD s1 to enable the rapid reconfiguration mechanism and operate in 802.1w mode:

```
configure stpd s1 mode dot1w
```

The following command configures STPD s2 to operate as an MSTI in an MSTP domain:

```
configure stpd s2 mode mstp msti 3
```

History

This command was first available in ExtremeXOS 10.1.

The `mstp` parameter was added in ExtremeXOS 11.4.



Platform Availability

This command is available on all platforms.

configure stpd ports active-role disable

```
configure stpd stpd_name ports active-role disable port
```

Description

Allows a port to be selected as an alternate or backup port.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>port</i>	Specifies a port.

Default

The default is disabled.

Usage Guidelines

Use this command to revert to the default that allows a specified port to be elected to any STP port role.

Example

The following command disables an active role on STDP s1, port 6:3:

```
configure stpd s1 ports active-role disable 6:3
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

configure stpd ports active-role enable

```
configure stpd stpd_name ports active-role enable port
```



Description

Prevents a port from becoming an alternate or backup port.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>port</i>	Specifies a port.

Default

The default is disabled.

Usage Guidelines

Use this command to keep a port in an active role. It prevents a specified port from being elected to an alternate or backup role which puts the port in a blocking state.

The following describes the port role and state when RSTP stabilizes.

STP Port Role	Port State
Alternate (inactive)	Blocking
Backup (inactive)	Blocking
Root (active)	Forwarding
Designated (active)	Forwarding

This feature can be enabled on only one STP port in the STP domain.

The restricted port role cannot be combined with this feature.

An active port role (root or designated) cannot be enabled with an edge port.

To disable this command, use the `configure stpd ports active-role disable` command.

To view the status of the active role, use the `show stpd ports` command.

Example

The following command enables an active role on STDP s1, port 6:3:

```
configure stpd s1 ports active-role enable 6:3
```

History

This command was first available in ExtremeXOS 12.5.



Platform Availability

This command is available on all platforms.

configure stpd ports bpdu-restrict

```
configure {stpd} stpd_name ports bpdu-restrict [enable | disable] port_list
{recovery-timeout {seconds}}
```

Description

Configures BPDU Restrict.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
bpdu-restrict	Disables port as soon as a BPDU is received.
recovery-timeout	Time after which the port will be re-enabled.
<i>seconds</i>	Specifies the time in seconds. The range is 60 to 600. The default is 300.

Default

The default is disabled.

Usage Guidelines

Before using this command, the port(s) should be configured for edge-safeguard.

Example

The following command enables bpdu-restrict on port 2 of STPD s1:

```
configure stpd s1 ports bpdu-restrict enable 2
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.



configure stpd ports cost

```
configure stpd stpd_name ports cost [auto | cost] port_list
```

Description

Specifies the path cost of the port in the specified STPD.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
auto	Specifies the switch to remove any user-defined port cost value(s) and use the appropriate default port cost value(s).
<i>cost</i>	Specifies a numerical port cost value. The range is 1 through 200,000,000.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

The switch automatically assigns a default path cost based on the speed of the port, as follows:

- 10 Mbps port—the default cost is 2,000,000.
- 100 Mbps port—the default cost is 200,000.
- 1000 Mbps port—the default cost is 20,000.
- 10000 Mbps ports—the default cost is 2,000.

The default port cost for trunked ports is dynamically calculated based on the available bandwidth.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

The 802.1D-2004 standard modified the default port path cost value to allow for higher link speeds. If you have a network with both 802.1D-2004 and 802.1D-1998 compliant bridges, a higher link speed can create a situation whereby an 802.1D-1998 compliant bridge could become the most favorable transit path and possibly cause the traffic to span more bridges. To prevent this situation, configure the port path cost to make links with the same speed use the same path host value. For example, if you have 100 Mbps links on all bridges, configure the port path cost for the 802.1D-2004 compliant bridges to 19 instead of using the default 200,000.



Note

You cannot configure the port path cost on 802.1D-1998 compliant bridges to 200,000 because the path cost range setting is 1 to 65,535.



The range for the cost parameter is 1 through 200,000,000. If you configure the port cost, a setting of 1 indicates the highest priority.

If you configured a port cost value and specify the auto option, the switch removes the user-defined port cost value and returns to the default, automatically assigned, port cost value.

The auto port cost of a trunk port is calculated based on number member ports in the trunk port. Link up and down of the member port does not affect the trunk port cost, thus it does not trigger topology change. Only adding or removing a member port to/from the trunk port causes auto trunk port cost to change. Also, by so configuring a static trunk port cost, the value is frozen regardless of the number of member ports in the trunk port.

ExtremeXOS 11.5 and Earlier

If you have switches running ExtremeXOS 11.5 and earlier, the default costs are different than switches running ExtremeXOS 11.6 and later.

The range for the cost parameter is 1 through 65,535.

The switch automatically assigns a default path cost based on the speed of the port, as follows:

- 10 Mbps port—the default cost is 100.
- 100 Mbps port—the default cost is 19.
- 1000 Mbps port—the default cost is 4.
- 10000 Mbps ports—the default cost is 2.

Example

The following command configures a cost of 100 to slot 2, ports 1 through 5 in STPD s0:

```
configure stpd s0 ports cost 100 2:1-2:5
```

History

This command was first available in ExtremeXOS 10.1.

The auto option was added in ExtremeXOS 11.0.

The default costs were updated based on support for the 802.1D-2004 standard in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure stpd ports edge-safeguard disable

```
configure {stpd} stpd_name ports edge-safeguard disable port_list {bpd-restrict}
{recovery-timeout {seconds}}
```



Description

Disables the edge safeguard loop prevention on the specified RSTP or MSTP edge port.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>port_list</i>	Specifies one or more edge ports.
bpdu-restrict	Disables port as soon as a BPDU is received.
recovery-timeout	Time after which the port will be re-enabled.
<i>seconds</i>	Specifies the time in seconds. The range is 60 to 600. The default is 300.

Default

By default, this feature is disabled.

Usage Guidelines

This command applies only to ports that have already been configured as edge ports.

Loop prevention and detection on an edge port configured for RSTP or MSTP is called edge safeguard. An edge port configured with edge safeguard immediately enters the forwarding state and transmits BPDUs.

If you disable this feature, the edge port enters the forwarding state but no longer transmits BPDUs unless a BPDU is received by that edge port. This is the default behavior.

Recovery time starts as soon as the port becomes disabled. If no recovery-timeout is specified, the port is permanently disabled.

BPDU restrict can be disabled using the `configure stpd stpd_name ports bpdu-restrict disableport_list` command.

If edge safeguard is disabled, BPDU restrict is also disabled.

To view the status of the edge safeguard feature use the `show {stpd} stpd_name ports {[detail |port_list {detail}]}` command. You can also use the `show stpd {stpd_name | detail}` command to display the STPD configuration on the switch, including the enable/disable state for edge safeguard.



Note

In MSTP, configuring edge safeguard at CIST will be inherited in all MSTI.

To enable or re-enable edge safeguard, use one of the following commands:

- `configure {stpd} stpd_name ports edge-safeguard enableport_list {bpdu-restrict} {recovery-timeout {seconds}}`



- `configure stpd stpd_name ports link-type [[auto | broadcast | point-to-point]port_list | edgeport_list {edge-safeguard [enable | disable] {bpdu-restrict} {recovery-timeoutseconds}}]`

Example

The following command disables edge safeguard on RSTP edge port 4 in STPD s1 on a stand-alone switch:

```
configure stpd s1 ports edge-safeguard disable 4
```

The following command disables edge safeguard on the RSTP edge port on slot 2, port 3 in STPD s1 on a modular switch:

```
configure stpd s1 ports edge-safeguard disable 2:3
```

History

This command was first available in ExtremeXOS 11.4.

The BPDU Restrict function was added in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

configure stpd ports edge-safeguard enable

```
configure {stpd} stpd_name ports edge-safeguard enable port_list {bpdu-restrict}  
{recovery-timeout {seconds}}
```

Description

Enables the edge safeguard loop prevention on the specified RSTP or MSTP edge port.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>port_list</i>	Specifies one or more edge ports.
bpdu-restrict	Disables port as soon as a BPDU is received.
recovery-timeout	Time after which the port will be re-enabled.
<i>seconds</i>	Specifies the time in seconds. The range is 60 to 600. The default is 300.



Default

By default, this feature is disabled.

Usage Guidelines

This command applies only to ports that have already been configured as edge ports.

Loop prevention and detection on an edge port configured for RSTP or MSTP is called edge safeguard. You configure edge safeguard on RSTP or MSTP edge ports to prevent accidental or deliberate misconfigurations (loops) resulting from connecting two edge ports together or by connecting a hub or other non-STP switch to an edge port. Edge safeguard also limits the impact of broadcast storms that might occur on edge ports.

An edge port configured with edge safeguard immediately enters the forwarding state and transmits BPDUs. This advanced loop prevention mechanism improves network resiliency but does not interfere with the rapid convergence of edge ports.

Recovery time starts as soon as the port becomes disabled. If no recovery-timeout is specified, the port is permanently disabled.

BPDUs can be disabled using the `configure {stp} stpd_name ports bpdu-restrict [enable | disable] port_list {recovery-timeout {seconds}}` command and selecting `disable`.

If edge safeguard is disabled, BPDUs are also disabled.

To view the status of the edge safeguard feature use the `show {stp} stpd_name ports {[detail | port_list {detail}]}` command. You can also use the `show stpd {stp_name | detail}` command to display the STPD configuration on the switch, including the enable/disable state for edge safeguard.



Note

In MSTP, configuring edge safeguard at CIST will be inherited in all MSTI.

To disable edge safeguard, use one of the following commands:

- `configure {stp} stpd_name ports edge-safeguard disable port_list {bpdu-restrict} {recovery-timeout {seconds}}`
- `configure stpd stpd_name ports link-type [[auto | broadcast | point-to-point] port_list | edgeport_list {edge-safeguard [enable | disable] {bpdu-restrict} {recovery-timeout seconds}}`

Example

The following command enables edge safeguard on RSTP edge port 4 in STPD s1 on a stand-alone switch:

```
configure stpd s1 ports edge-safeguard enable 4
```



The following command enables edge safeguard on the RSTP edge port on slot 2, port 3 in STPD s1 on a modular switch:

```
configure stpd s1 ports edge-safeguard enable 2:3
```

History

This command was first available in ExtremeXOS 11.4.

The BPDU Restrict function was added in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

configure stpd ports link-type

```
configure stpd stpd_name ports link-type [[auto | broadcast | point-to-point]
port_list | edge port_list {edge-safeguard [enable | disable] {bpdu-restrict}
{recovery-timeout seconds}}]
```

Description

Configures the ports in the specified STPD as auto, broadcast, edge, or point-to-point link types.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
auto	Specifies the switch to automatically determine the port link type. An auto link behaves like a point-to-point link if the link is in full-duplex mode or if link aggregation is enabled on the port. Used for 802.1w configurations.
broadcast	Specifies a port attached to a LAN segment with more than two bridges. Used for 802.1D configurations. A port with broadcast link type cannot participate in rapid reconfiguration using RSTP or MSTP. By default, all STP.1D ports are broadcast links.
point-to-point	Specifies a port attached to a LAN segment with only two bridges. A port with point-to-point link type can participate in rapid reconfiguration. Used for 802.1w and MSTP configurations. By default, all 802.1w and MSTP ports are point-to-point link types.
<i>port_list</i>	Specifies one or more ports or slots and ports.
edge	Specifies a port that does not have a bridge attached. An edge port is placed and held in the STP forwarding state unless a BPDU is received by the port. Used for 802.1w and MSTP configurations.
edge-safeguard	Specifies that the edge port be configured with edge safeguard, a loop prevention and detection mechanism. Used for 802.1w and MSTP configurations.



enable	Specifies that edge safeguard be enabled on the edge port(s).
disable	Specifies that edge safeguard be disabled on the edge port(s).
bpdu-restrict	Disables port as soon as a BPDU is received.
recovery-timeout	Time after which the port will be re-enabled.
<i>seconds</i>	Specifies the time in seconds. The range is 60 to 600. The default is 300.

Default

STP.1D ports are broadcast link types 802.1w and MSTP ports are point-to-point link types.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword stpd is optional.

The default, broadcast links, supports legacy STP (802.1D) configurations. If the switch operates in 802.1D mode, any configured port link type will behave the same as the broadcast link type.

RSTP rapidly moves the designated ports of a point-to-point link type into the forwarding state. This behavior is supported by RSTP and MSTP only.

In an MSTP environment, configure the same link types for the CIST and all MSTIs.

Auto Link Type

An auto link behaves like a point-to-point link if the link is in full duplex mode or if link aggregation is enabled on the port; otherwise, an auto link behaves like a broadcast link. If a non-STP switch exists between several switches operating in 802.1w mode with auto links, the non-STP switch may negotiate full-duplex even though the broadcast domain extends over several STP devices.

Edge Link Type

RSTP does not send any BPDUs from an edge port nor does it generate topology change events when an edge port changes its state.

If you configure a port to be an edge port, the port immediately enters the forwarding state. Edge ports remain in the forwarding state unless the port receives a BPDU. In that case, edge ports enter the blocking state. The edge port remains in the blocking state until it stops receiving BPDUs and the message age timer expires.

Edge Safeguard

Loop prevention and detection on an edge port configured for RSTP or MSTP is called edge safeguard. You configure edge safeguard on RSTP or MSTP edge ports to prevent accidental or deliberate misconfigurations (loops) resulting from connecting two edge ports together or by connecting a hub or



other non-STP switch to an edge port. Edge safeguard also limits the impact of broadcast storms that might occur on edge ports.

An edge port configured with edge safeguard immediately enters the forwarding state and transmits BPDUs. This advanced loop prevention mechanism improves network resiliency but does not interfere with the rapid convergence of edge ports.

Recovery time starts as soon as the port becomes disabled. If no recovery-timeout is specified, the port is permanently disabled.

BPDUs restrict can be disabled using the `configure stpd stpd_name ports bpdu-restrict disableport_list` command.

If edge safeguard is disabled, BPDUs restrict is also disabled.

To configure a port as an edge port and enable edge safeguard on that port, use the `configure stpd stpd_name ports link-type edgeport_list edge-safeguard` command and specify `enable`.

To disable edge safeguard on the edge port, use the `configure stpd stpd_name ports link-type edgeport_list edge-safeguard` command and specify `disable`.

Two other commands are also available to enable and disable edge safeguard:

```
configure stpd ports edge-safeguard enable
configure stpd ports edge-safeguard disable
```

In MSTP, configuring edge safeguard at CIST will be inherited in all MSTI.

Example

The following command configures slot 2, ports 1 through 4 to be point-to-point links in STPD s1:

```
configure stpd s1 ports link-type point-to-point 2:1-2:4
```

The following command enables edge safeguard on the RSTP edge port on slot 2, port 3 in STPD s1 configured for RSTP:

```
configure stpd s1 ports link-type edge 2:3 edge-safeguard enable
```

History

This command was first available in ExtremeXOS 10.1.

The BPDUs Restrict function was added in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.



configure stpd ports mode

```
configure stpd stpd_name ports mode [dot1d | emistp | pvst-plus] port_list
```

Description

Configures the encapsulation mode for the specified port list.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
dot1d	Specifies the STP encapsulation mode of operation to be 802.1d.
emistp	Specifies the STP encapsulation mode of operation to be EMISTP.
pvst-plus	Specifies the STP encapsulation mode of operation to be PVST+.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Ports in the default STPD (s0) are dot1d mode.

Ports in user-created STPDs are in **emistp** mode.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword **stpd** is optional.

MSTP STPDs use 802.1D BPDU encapsulation mode by default. To ensure correct operation of your MSTP STPDs, do not configure EMISTP or PVST+ encapsulation mode for MSTP STPDs.

You can specify the following STP encapsulation modes:

- **dot1d**—This mode is reserved for backward compatibility with previous STP versions. BPDUs are sent untagged in 802.1D mode. Because of this, any given physical interface can have only one STPD running in 802.1D mode.

This encapsulation mode supports the following STPD modes of operation: 802.1D, 802.1w, and MSTP.

- **emistp**—This mode sends BPDUs with an 802.1Q tag having an STPD ID in the VLAN ID field.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

- **pvst-plus**—This mode implements PVST+ in compatibility with third-party switches running this version of STP. The STPDs running in this mode have a one-to-one relationship with VLANs, and send and process packets in PVST+ format.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.



Example

The following command configures STPD s1 with PVST+ packet formatting for slot 2, port 1:

```
configure stpd s1 ports mode pvst-plus 2:1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure stpd ports port-priority

```
configure stpd stpd_name ports port-priority priority port_list
```

Description

Specifies the port priority of the port in the specified STPD.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>priority</i>	Specifies a numerical port priority value. The range is 0 through 240 and is subject to the multiple of 16 restriction.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

The default is 128.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

By changing the priority of the port, you can make it more or less likely to become the root port or a designated port.



To preserve backward compatibility and to use ExtremeXOS 11.5 or earlier configurations, the existing `configure stpd ports priority` command is available in ExtremeXOS 11.6. If you have an ExtremeXOS 11.5 or earlier configuration, the switch interprets the port priority based on the 802.1D-1998 standard. If the switch reads a value that is not supported in ExtremeXOS 11.6, the switch rejects the entry. For example, if the switch reads the `configure stpd ports priority 16` command from an ExtremeXOS 11.5 or earlier configuration, (which is equivalent to the command `configure stpd ports priority 8` entered through CLI), the switch saves the value in the new ExtremeXOS 11.6 configuration as `configure stpd ports port-priority 128`.

A setting of 0 indicates the highest priority.

The range for the priority parameter is 0 through 240 and is subject to the multiple of 16 restriction.

Example

The following command assigns a priority of 32 to slot 2, ports 1 through 5 in STPD s0:

```
configure stpd s0 ports port-priority 32 2:1-2:5
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure stpd ports priority

```
configure stpd stpd_name ports priority priority port_list
```

Description

Specifies the port priority of the port in the specified STPD.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>priority</i>	Specifies a numerical port priority value. The range is 0 through 31 for STP and 0 through 15 for MSTP and RSTP.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

The default is 128.



Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

By changing the priority of the port, you can make it more or less likely to become the root port or a designated port.

To preserve backward compatibility and to use ExtremeXOS 11.5 or earlier configurations, the existing `configure stpd ports priority` command is available in ExtremeXOS 11.6. If you have an ExtremeXOS 11.5 or earlier configuration, the switch interprets the port priority based on the 802.1D-1998 standard. If the switch reads a value that is not supported in ExtremeXOS 11.6, the switch rejects the entry.

A setting of 0 indicates the highest priority.

The range for the priority parameter is 0 through 31 for STP and 0 through 15 for MSTP and RSTP.

ExtremeXOS 11.6 introduces support for a new ports priority command: `configure stpd ports port-priority`. When you save the port priority value in an ExtremeXOS 11.6 configuration, the switch saves it as the new command `configure stpd ports port-priority` with the corresponding change in priority values. The priority range of this command is 0 through 240 and is subject to the multiple of 16 restriction. For more information see `configure stpd ports port-priority`.

ExtremeXOS 11.5 and Earlier

If you have switches running ExtremeXOS 11.5 and earlier, the default value for the priority range are different than switches running ExtremeXOS 11.6.

The range for the priority parameter is 0 through 31.

The default is 16.

Example

The following command assigns a priority of 1 to slot 2, ports 1 through 5 in STPD s0:

```
configure stpd s0 ports priority 1 2:1-2:5
```

History

This command was first available in ExtremeXOS 10.1.

The priority range and behavior was updated based on support for the 802.1D-2004 standard in ExtremeXOS 11.6.



Platform Availability

This command is available on all platforms.

configure stpd ports restricted-role disable

```
configure stpd stpd_name ports restricted-role disable port_list
```

Description

Disables restricted role on the specified port inside the core network.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

The restricted role is disabled by default. If set, it can cause a lack of spanning tree connectivity. A network administrator enables the restricted role to prevent bridges external to a core region of the network from influencing the spanning tree active topology, possibly because those bridges are not under the full control of the administrator.



Note

Disabling Restricted Role at CIST is inherited by all MSTI.

Example

The following command disables restricted role for s1 on port 6:3:

```
configure stpd s1 ports restricted-role disable 6:3
```

History

This command was first available in ExtremeXOS 12.1.

This command was added to RSTP in ExtremeXOS 11.6 and 12.0.3.

Platform Availability

This command is available on all platforms.



configure stpd ports restricted-role enable

```
configure stpd stpd_name ports restricted-role enable port_list
```

Description

Enables restricted role on the specified port inside the core network.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

Enabling restricted role causes the port not to be selected as a root port even if it has the best spanning tree priority vector. Such a port is selected as an alternate port after the root port has been selected.

The restricted role is disabled by default. If set, it can cause a lack of spanning tree connectivity. A network administrator enables the restricted role to prevent bridges external to a core region of the network from influencing the spanning tree active topology, possibly because those bridges are not under the full control of the administrator.



Note

Restricted role should not be enabled with edge mode.

Enabling Restricted Role at CIST is inherited by all MSTI.

Example

The following command enables restricted role on port 6:3:

```
configure stpd s1 ports restricted-role enable 6:3
```

History

This command was first available in ExtremeXOS 12.1.

This command was added to RSTP in ExtremeXOS 11.6 and 12.0.3.



Platform Availability

This command is available on all platforms.

configure stpd priority

```
configure stpd stpd_name priority priority
```

Description

Specifies the bridge priority of the STPD.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>priority</i>	Specifies the bridge priority of the STPD. The range is 0 through 61,440 and is subject to the multiple of 4,096 restriction.

Default

The default priority is 32,768.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

By changing the bridge priority of the STPD, you can make it more or less likely to become the root bridge.

The range for the *priority* parameter is 0 through 61,440 and is subject to the multiple of 4,096 restriction. A setting of 0 indicates the highest priority.

If you have an ExtremeXOS 11.5 or earlier configuration that contains an STP or RSTP bridge priority that is not a multiple of 4,096, the switch rejects the entry and the bridge priority returns to the default value. The MSTP implementation already uses multiples of 4,096 to determine the bridge priority.

For example, to lower the numerical value of the priority (which gives the priority a higher precedence), you subtract 4,096 from the default priority: $32,768 - 4,096 = 28,672$. If you modify the priority by a value other than 4,096, the switch rejects the entry.

ExtremeXOS 11.5 and Earlier

If you have switches running ExtremeXOS 11.5 and earlier, the priority range is different than switches running ExtremeXOS 11.6 and later.



The range for the priority parameter is 0 through 65,535. A setting of 0 indicates the highest priority.

Example

The following command sets the bridge priority of STPD1 to 16,384:

```
configure stpd stpd1 priority 16384
```

History

This command was first available in ExtremeXOS 10.1.

The priority range and behavior was updated based on support for the 802.1D-2004 standard in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

configure stpd tag

```
configure stpd stpd_name tag stpd_tag
```

Description

Assigns an StpdID to an STPD.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>stpd_tag</i>	Specifies the VLAN ID of the carrier VLAN that is owned by the STPD.

Default

N/A.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.



An STPD ID is used to identify each STP domain. You assign the StpdID when configuring the domain. An STPD ID must be identical to the VLAN ID of the carrier VLAN in that STP domain, and that VLAN cannot belong to another STPD. Unless all ports are running in 802.1D mode, an STPD with ports running in either EMISTP mode or PVST+ mode must be configured with an STPD ID.

You must create and configure the VLAN, along with the tag, before you can configure the STPD tag. To create a VLAN, use the `create vlan` command. To configure the VLAN, use the `configure vlan` commands.

MSTP Only

MSTP uses two different methods to identify the STPDs that are part of the MSTP network. An instance ID of 0 identifies the CIST. The switch assigns this ID automatically when you configure the CIST STPD. To configure the CIST STPD, use the `configure stpd stpd_name mode [dot1d | dot1w | mstp [cist | mstiinstance]]` command.

An MSTI identifier (MSTI ID) identifies each STP domain that is part of an MSTP region. You assign the MSTI ID when configuring the STPD that participates in the MSTP region. Each STPD that participates in a particular MSTP region must have the same MSTI ID. To configure the MSTI ID, use the `configure stpd stpd_name mode [dot1d | dot1w | mstp [cist | mstiinstance]]` command.

Example

The following command assigns an StpdID to the purple_st STPD:

```
configure stpd purple_st tag 200
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure vlan add ports stpd

```
configure vlan vlan_name add ports [all | port_list] {tagged {tag} | untagged}
stpd stpd_name {[dot1d | emistp | pvst-plus]}
```

Description

Adds one or more ports in a VLAN to a specified STPD.



Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
all	Specifies all of the ports to be included in the STPD.
<i>port_list</i>	Specifies the port or ports to be included in the STPD.
tagged	Specifies the ports should be configured as tagged.
<i>tag</i>	Specifies the port-specific VLAN tag. When there are multiple ports specified in the <i>port_list</i> , the same tag is used for all of them. When unspecified port <i>tag</i> is equal to the VLAN tag.
untagged	Specifies the ports should be configured as untagged.
<i>stpd_name</i>	Specifies an STPD name on the switch.
dot1d	Specifies the STP encapsulation mode of operation to be 802.1d.
emistp	Specifies the STP encapsulation mode of operation to be EMISTP.
pvst-plus	Specifies the STP encapsulation mode of operation to be PVST+.

Default

Ports in the default STPD (s0) are in dot1d mode.

Ports in user-created STPDs are in emistp mode.

Usage Guidelines

To create a VLAN, use the `create vlan` command. To create an STP domain, use the `create stpd` command.

In an EMISTP or PVST+ environment, this command adds a list of ports to a VLAN and a specified STPD at the same time provided the carrier VLAN already exists on the same set of ports. You can also specify the encapsulation mode for those ports.

In an MSTP environment, you do not need a carrier VLAN. A CIST controls the connectivity of interconnecting MSTP regions and sends BPDUs across the regions to communicate region status. You must use the dot1d encapsulation mode in an MSTP environment.

You cannot configure STP on the following ports:

- Mirroring target ports.
- Software-controlled redundant ports.

If you see an error similar to the following:

```
Error: Cannot add VLAN default port 3:5 to STP domain
```

You might be attempting to add:

- A carrier VLAN port to a different STP domain than the carrier VLAN belongs.



- A VLAN/port for which the carrier VLAN does not yet belong.



Note

This restriction is only enforced in an active STP domain and when you enable STP to ensure you have a legal STP configuration.

Naming Conventions

If your VLAN has the same name as another component, for example an STPD, we recommend that you specify the identifying keyword as well as the name. If your VLAN has a name unique only to that VLAN, the keywords `vlan` and `stpd` are optional.

STP Encapsulation Modes

You can specify the following STP encapsulation modes:

- `dot1d`—This mode is reserved for backward compatibility with previous STP versions. BPDUs are sent untagged in 802.1D mode. Because of this, any given physical interface can have only one STPD running in 802.1D mode.

This encapsulation mode supports the following STPD modes of operation: 802.1D, 802.1w, and MSTP.

- `emistp`—This mode sends BPDUs with an 802.1Q tag having an STPD ID in the VLAN ID field.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

- `pvst-plus`—This mode implements PVST+ in compatibility with third-party switches running this version of STP. The STPDs running in this mode have a one-to-one relationship with VLANs, and send and process packets in PVST+ format.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

These encapsulation modes are for STP ports, not for physical ports. When a physical port belongs to multiple STPDs, it is associated with multiple STP ports. It is possible for the physical port to run in different modes for different domains for which it belongs.

MSTP STPDs use only 802.1D BDU encapsulation mode. The switch prevents you from configuring EMISTP or PVST+ encapsulation mode for MSTP STPDs.

Specify the port `tag` when you need to put multiple vlans into a broadcast domain.

Automatically Inheriting Ports--MSTP Only

In an MSTP environment, whether you manually or automatically bind a port to an MSTI in an MSTP region, the switch automatically binds that port to the CIST. The CIST handles BDU processing for itself and all of the MSTIs; therefore, the CIST must inherit ports from the MSTIs in order to transmit and receive BPDUs.



Example

The following command adds slot 1, port 2 and slot 2, port 3, members of a VLAN named Marketing, to the STPD named STPD1, and specifies that they be in EMISTP mode:

```
configure vlan marketing add ports 1:2, 2:3 tagged stpd stpd1 emistp
```

The following examples illustrate the *tag* variable in ExtremeXOS 15.4:

The following example configures vlan with tag 100 and port tag of 10 and 11 on two different ports:

```
create vlan exchange tag 100
config vlan exchange add ports 3 tagged 10
config vlan exchange add ports 4 tagged 11
```

The following example configures a VLAN with tag 100, and port tag of 10 and 11 on the same ports:

```
create vlan exchange tag 100
config vlan exchange add ports 3 tagged 10
config vlan exchange add ports 3 tagged 11
```

The following example configures VLAN with tag 100, and port tag of 10 on two ports and 11 on a different port:

```
create vlan exchange tag 100
config vlan exchange add ports 2:3,2:4 tagged 10
config vlan exchange add ports 2:5 tagged 11
```

History

This command was first available in ExtremeXOS 10.1.

The **nobroadcast** keyword was removed in ExtremeXOS 11.4.

The *tag* variable was added in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

create stpd

```
create stpd stpd_name {description stpd-description}
```

Description

Creates a user-defined STPD.



Syntax Description

<i>stp_name</i>	Specifies a user-defined STPD name to be created. May be up to 32 characters in length.
<i>stp-description</i>	Specifies an STP domain description string.

Default

The default device configuration contains a single STPD called s0.

When an STPD is created, the STPD has the following default parameters:

- State—disabled.
- StpdID—none.
- Assigned VLANs—none.
- Bridge priority—32,768.
- Maximum BPDU age—20 seconds.
- Hello time—2 seconds.
- Forward delay—15 seconds.
- Operational mode—802.1D.
- Rapid Root Failover—disabled.
- Default Binding Mode (encapsulation mode)—Ports in the default STPD (s0) are in 802.1d mode. Ports in user-created STPDs are in emistp mode.
- Maximum hop count (when configured for MSTP)—20 hops.
- STP domain description string—empty.

Usage Guidelines

The maximum length for a name is 32 characters. Names can contain alphanumeric characters and underscores (`_`) but cannot be any reserved keywords, for example, `stp` or `stp`. Names must start with an alphabetical character, for example, a, Z. For name creation guidelines and a list of reserved names, see [Object Names](#) in the *ExtremeXOS Concepts Guide*.

Each STPD name must be unique and cannot duplicate any other named STPDs on the switch. If you are uncertain about the STPD names on the switch, use the `show stpd` command to view the STPD names.

You can, however, re-use names across multiple categories of switch configuration. For example, you can use the name `Test` for an STPD and a VLAN. If you use the same name, we recommend that you specify the appropriate keyword when configuring the STPD. If you do not specify the appropriate keyword, the switch displays a message similar to the following:

```
%% Ambiguous command: "configure Test"
```

To view the names of the STPDs on the switch, enter `configure` and press [Tab]. Scroll to the end of the output to view the names.

The maximum length for an STPD description is 180 characters. The description must be in quotes if the string contains any spaces. To display the description, use the `show stpd stpd_name` command.



Each STPD has its own Root Bridge and active path. After the STPD is created, one or more VLANs can be assigned to it.

Example

The following example creates an STPD named purple_st:

```
create stpd purple_st
```

History

This command was first available in ExtremeXOS 10.1.

The STPD description option was added in ExtremeXOS 12.4.4.

Platform Availability

This command is available on all platforms.

delete stpd

```
delete stpd stpd_name
```

Description

Removes a user-defined STPD from the switch.

Syntax Description

<i>stpd_name</i>	Specifies a user-defined STPD name on the switch.
------------------	---

Default

N/A.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, we recommend that you specify the identifying keyword as well as the name. If you do not specify the stpd keyword, an error message similar to the following is displayed:

```
%% Ambiguous command: "delete Test"
```

In this example, to delete the STPD Test, enter delete stpd Test.

If you created an STPD with a name unique only to that STPD, the keyword stpd is optional.



The default STPD, s0, cannot be deleted.

In an MSTP environment, you cannot delete or disable a CIST if any of the MSTIs are active in the system.

Example

The following command deletes an STPD named purple_st:

```
delete stpd purple_st
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable stpd

```
disable stpd {stpd_name}
```

Description

Disables the STP protocol on a particular STPD or for all STPDs.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
------------------	---------------------------------------

Default

Disabled.

Usage Guidelines

After you have created the STPD with a unique name, the keyword stpd is optional.

If you want to disable the STP protocol for all STPDs, do not specify an STPD name.

In an MSTP environment, you cannot delete or disable a CIST if any of the MSTIs are active in the system.



Example

The following command disables an STPD named purple_st:

```
disable stpd purple_st
```

The following command disables the STP protocol for all STPDs on the switch:

```
disable stpd
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable stpd auto-bind

Disables the ability to automatically add ports to an STPD when they are added to a member VLAN.

```
disable stpd stpd_name auto-bind vlan vlan_name
```

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>vlan_name</i>	Specifies the name of a member VLAN with autobind enabled.

Default

The autobind feature is disabled on user-created STPDs. The autobind feature is enabled on the default VLAN that participates in the default STPD S0.

Usage Guidelines



Note

Ports already in the STPD remain in that domain (as if they were added manually).

If you create an STPD and a VLAN with unique names, the keywords `stpd` and `vlan` are optional.

Ports added to the STPD automatically when autobind is enabled are not removed when autobind is disabled. The ports are present after a switch reboot.

To view STP configuration status of the ports in a VLAN, use the following command:

```
show {vlan} vlan_name stpd
```



Example

The following example disables autobind on an STPD named s8:

```
disable stpd s8 auto-bind v5
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable stpd ports

Disables STP on one or more ports for a given STPD.

```
disable stpd stpd_name ports [all | port_list]
```

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
all	Specifies all ports for a given STPD.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Enabled.

Usage Guidelines

If you create the STPD with a unique name, the keyword `stpd` is optional.

Disabling STP on one or more ports puts those ports in the forwarding state; all BPDUs received on those ports are disregarded and dropped.

Use the `all` keyword to specify that all ports of a given STPD are disabled.

Use the `port_list` parameter to specify a list of ports of a given STPD are disabled.

If you do not use the default STPD, you must create one or more STPDs and configure and enable the STPD before you can use the `disable stpd ports` command.



Example

The following command disables slot 2, port 4 on an STPD named Backbone_st:

```
disable stpd backbone_st ports 2:4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable stpd rapid-root-failover

```
disable stpd stpd_name rapid-root-failover
```

Description

Disables rapid root failover for STP recovery times.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
------------------	---------------------------------------

Default

Disabled.

Usage Guidelines

This command is applicable for STPDs operating in 802.1D.

After you have created the STPD with a unique name, the keyword stpd is optional.

To view the status of rapid root failover on the switch, use the `show stpd` command. The `show stpd` command displays information about the STPD configuration on the switch including the enable/disable state for rapid root failover.

Example

The following command disables rapid root fail over on STPD Backbone_st:

```
disable stpd backbone_st rapid-root-failover
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable stpd

```
enable stpd {stpd_name}
```

Description

Enables the STP protocol for one or all STPDs.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
------------------	---------------------------------------

Default

Disabled.

Usage Guidelines

If you want to enable the STP protocol for all STPDs, do not specify an STPD name.

Example

The following command enables an STPD named Backbone_st:

```
enable stpd backbone_st
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable stpd auto-bind

```
enable stpd stpd_name auto-bind vlan vlan_name
```



Description

Automatically adds ports to an STPD when ports are added to a member VLAN.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>vlan_name</i>	Specifies the name of the VLAN to have autobind enabled.

Default

The autobind feature is disabled on user-created STPDs. The autobind feature is enabled on the default VLAN that participates in the default STPD S0.

If you enable autobind and add ports to a member VLAN, those ports are automatically added to the STPD.

Usage Guidelines

If you create an STPD and a VLAN with unique names, the keywords `stpd` and `vlan` are optional.

You cannot configure the autobind feature on a network login VLAN.

In an EMISTP or PVST+ environment, when you issue this command, any port or list of ports that you add to the carrier VLAN are automatically added to the STPD with autobind enabled. In addition, any port or list of ports that you remove from a carrier VLAN are automatically removed from the STPD. This allows the STPD to increase or decrease its span as you add ports to or remove ports from a carrier VLAN.

For MSTP, when you issue this command, any port or list of ports that gets automatically added to an MSTI are automatically inherited by the CIST. In addition, any port or list of ports that you remove from an MSTI protected VLAN are automatically removed from the CIST. For more information see the section. For more information, see [Automatically Inheriting Ports—MSTP Only](#).

Carrier VLAN

A carrier VLAN defines the scope of the STPD, which includes the physical and logical ports that belong to the STPD and the 802.1Q tag used to transport STP BPDUs in the encapsulation mode is EMISTP or PVST+. Only one carrier VLAN can exist in a given STPD, although some of its ports can be outside the control of any STPD at the same time.



Note

The STPD ID must be identical to the VLAN ID of the carrier VLAN in that STPD.

If you configure MSTP, you do not need a carrier VLAN. With MSTP, you configure a CIST that controls the connectivity of interconnecting MSTP regions and sends BPDUs across the regions to communicate the status of MSTP regions. All VLANs participating in the MSTP region have the same privileges.



Protected VLAN

Protected VLANs are all other VLANs that are members of the STPD. These VLANs “piggyback” on the carrier VLAN. Protected VLANs do not transmit or receive STP BPDUs, but they are affected by STP state changes and inherit the state of the carrier VLAN. Protected VLANs can participate in multiple STPDs, but any particular port in the VLAN can belong to only one STPD.

Enabling autobind on a protected VLAN does not expand the boundary of the STPD. However, the VLAN and port combinations are added to or removed from the STPD subject to the boundaries of the carrier VLAN.

If you configure MSTP, all member VLANs in an MSTP region are protected VLANs. These VLANs do not transmit or receive STP BPDUs, but they are affected by STP state changes communicated by the CIST to the MSTP regions. MSTIs cannot share the same protected VLAN; however, any port in a protected VLAN can belong to multiple MSTIs.

Automatically Inheriting Ports--MSTP Only

In an MSTP environment, whether you manually or automatically bind a port to an MSTI in an MSTP region, the switch automatically binds that port to the CIST. The CIST handles BPDU processing for itself and all of the MSTIs; therefore, the CIST must inherit ports from the MSTIs in order to transmit and receive BPDUs.

Displaying STP Information

To view STP configuration status of the ports on a VLAN, use the following command:

```
show {vlan} vlan_name stpd
```

Example

The examples in this section assume that you have already removed the ports from the Default VLAN.

To automatically add ports to an STPD running 802.1D, EMISTP, or PVST+ and to expand the boundary of the STPD, you must complete the following tasks:

- Create the carrier VLAN.
- Assign a VLAN ID to the carrier VLAN.
- Add ports to the carrier VLAN.
- Create an STPD (or use the default, S0).
- Enable autobind on the STPDs carrier VLAN.
- Configure the STPD tag (the STPD ID must be identical to the VLAN ID of the carrier VLAN in the STP domain).
- Enable STP.

The following example enables autobind on an STPD named s8 after creating a carrier VLAN named v5:

```
create vlan v5
configure vlan v5 tag 100
configure vlan v5 add ports 1:1-1:20 tagged
create stpd s8
```




```
enable stpd s8 auto-bind v5
configure stpd s8 tag 100
enable stpd s8
```

To automatically add ports to the CIST STPD and to expand the boundary of the STPD, you must complete the following tasks:

- Create a VLAN or use the Default VLAN. (In this example, the Default VLAN is used.)
- Create the MSTP region.
- Create the STPD to be used as the CIST, and configure the mode of operation for the STPD.
- Specify the priority for the CIST.
- Enable the CIST.

The following example enables autobind on the VLAN Default for the CIST STPD named s1:

```
configure mstp region 1
create stpd s1
configure stpd s1 mode mstp cist
configure stpd s1 priority 32768
enable stpd s1
```

The following example enables autobind on the VLAN math for the MSTI STPD named s2:

```
create vlan math
configure vlan math tag 2
configure vlan math add ports 2-3
configure mstp region 1
create stpd s2
configure stpd s2 mode mstp msti 1
configure stpd s2 priority 32768
enable stpd s2 auto-bind vlan math
configure stpd s2 ports link-type point-to-point 5-6
enable stpd s2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable stpd ports

```
enable stpd stpd_name ports [all | port_list]
```

Description

Enables the STP protocol on one or more ports.



Syntax Description

<i>stpd_name</i>	Specifies an STPD on the switch.
all	Specifies all ports for a given STPD.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Enabled.

Usage Guidelines

If you create an STPD with a unique name, the keyword `stpd` is optional.

If STP is enabled for a port, BPDUs are generated and processed on that port if STP is enabled for the associated STPD.

You must configure one or more STPDs before you can use the `enable stpd ports` command. To create an STPD, use the `create stpd stpd_name {description stpd-description}` command. If you have considerable knowledge and experience with STP, you can configure the STPD using the `configure stpd` commands. However, the default STP parameters are adequate for most networks.

Example

The following command enables slot 2, port 4 on an STPD named `Backbone_st`:

```
enable stpd backbone_st ports 2:4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable stpd rapid-root-failover

```
enable stpd stpd_name rapid-root-failover
```

Description

Enables rapid root failover for faster STP recovery times.



Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
------------------	---------------------------------------

Default

Disabled.

Usage Guidelines

This command is applicable for STPDs operating in 802.1D.

If you create an STPD with a unique name, the keyword `stpd` is optional.

To view the status of rapid root failover on the switch, use the `show stpd` command. The `show stpd` command displays information about the STPD configuration on the switch including the enable/disable state for rapid root failover.

Example

The following command enables rapid root fail over on STPD Backbone_st:

```
enable stpd backbone_st rapid-root-failover
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show stpd

```
show stpd {stpd_name | detail}
```

Description

Displays STPD settings on the switch.

Syntax Description

<i>stpd_name</i>	Specifies an STPD on the switch.
detail	Specifies that STPD settings should be shown for each STPD.



Default

N/A.

Usage Guidelines

If you specify the command without any options, the following STPD information appears:

- Name—The name of the STPD.
- Tag—The StpdID of the domain, if configured.
- Flags—The following flags communicate information about the current state of the STPD:
 - (C) Topology Change—A network topology change has occurred in the network.
 - (D) Disable—The STPD is disabled.
 - (E) Enable—The STPD is enabled.
 - (R) Rapid Root Failover—The STPD has been configured for rapid root failover.
 - (T) Topology Change Detected—The STPD has detected a change in the network topology.
 - (M) MSTP CIST—The STPD has been configured for MSTP, and the STPD is the common and internal spanning tree.
 - (I) MSTP MSTI—The STPD has been configured for MSTP, and the STPD is a multiple instance spanning tree.
- Ports—The number of ports that are part of the STPD.
- Bridge ID—The MAC addresses of the switch.
- Designated Root—The MAC address of the switch that is the designated root bridge.
- Rt Port—The root port.
- Rt Cost—The path cost to the root port.
- Total Number of STPDs—The total number of STPDs configured on the switch.
- STP Flush Method—The method used to flush the FDB during a topology change.

If you have an MSTP region and associated spanning trees configured on the switch, the command also displays the following global MSTP information:

- MSTP Region—The name of the MSTP region configured on the switch.
- Format Identifier—The number used by BPDUs to communicate within an MSTP region.
- Revision Level—This number is reserved for future use.
- Common and Internal Spanning Tree (CIST)—The name of the CIST that controls the connectivity of interconnecting MSTP regions.
- Total number of MST Instances (MSTI)—The number of MSTIs running in the MSTP region.

If you use the `show stpd` command and specify the name of an STPD, in addition to the data previously described, the command displays more detailed information about the STPD. If you specify the `detail` option, the switch displays the same type of information for all of the STPDs configured on the switch.

The additional output includes the following:

- STPD mode of operation.
- Autobind mode.
- Active VLANs.
- Timer information.
- Topology change information.



If you have MSTP configured, the command also displays the following information:

- Bridge role.
- CIST root.
- CIST regional root.
- MSTI instances.
- Master port (Displayed only on MSTI STPDs).

If your STPD has the same name as another component, for example a VLAN, we recommend that you specify the identifying keyword as well as the name. If you do not specify the stpd keyword, an error message similar to the following is displayed:

```
%% Ambiguous command: "show Test"
```

In this example, to view the settings of the STPD Test, enter show stpd Test.

If your STPD has a name unique only to that STPD, the keyword stpd is optional.

Example

The following command displays the STPD settings on a switch that has MSTP configured:

```
show stpd
```

The following is sample output from this command:

```
MSTP Global Configuration:
MSTP Region Name           : 00304841ed97
MSTP format Identifier     : 0
MSTP Revision Level       : 3
Common and Internal Spanning Tree (CIST) : ----
Total Number of MST Instances (MSTI) : 0
Name Tag  Flags  Ports Bridge ID      Designated Root  Rt Port Rt Cost
s0        0000 D----- 0 8000001030f99dc0 0000000000000000 -----
0
    Total number of STPDs:
1                                     STP Flush Method:
Port only
Flags: (C) Topology Change, (D) Disable, (E) Enable, (R) Rapid Root
Failover      (T) Topology Change Detected, (M) MSTP CIST , (I) MSTP
MSTI
```

The following command displays STPD settings on an STPD named Backbone_st:

```
show stpd backbone_st
```

The following is sample output from this command:

```
Stpd: backbone_st Stp: ENABLED      Number of Ports: 51
Description: this is backbone_st domain
```



```

Rapid Root Failover: Disabled
Operational Mode: 802.1W Default Binding Mode: 802.1D
802.1Q Tag: (none)
Ports: 1:1,1:2,2:1,2:2,3:1,3:2,4:1,4:2,5:1,5:2,
5:3,5:4,5:5,5:6,5:7,5:8,5:9,5:10,5:11,5:12,
5:13,5:14,5:15,5:16,5:17,5:18,5:19,5:20,5:21,5:22,
5:23,5:24,5:25,5:26,5:27,5:28,5:29,5:30,5:31,5:32,
5:33,5:34,5:35,5:36,5:37,5:38,5:39,5:40,5:41,5:42,
5:43
Participating Vlans: Default
Auto-bind Vlans: Default
Bridge Priority: 5000
BridgeID: 13:88:00:01:30:f4:06:80
Designated root: 0a:be:00:01:30:28:b7:00
RootPathCost: 19 Root Port: 28
MaxAge: 20s HelloTime: 2s ForwardDelay: 15s
CfgBrMaxAge: 20s CfgBrHelloTime: 2s CfgBrForwardDelay: 15s
Topology Change Time: 35s Hold time: 1s
Topology Change Detected: FALSE Topology Change: FALSE
Number of Topology Changes: 7
Time Since Last Topology Change: 4967s

```

The following is sample output for an STPD configured as the CIST (the output is similar for an STPD configured as an MSTI):

```

Stpd: s0 Stp: DISABLED Number of Ports: 0
Description: this is s0 domain
Rapid Root Failover: Disabled Operational Mode:
MSTP Default Binding Mode: 802.1d
MSTP Instance :CIST CIST : s0
802.1Q Tag: (none) Ports:
(none)
Participating Vlan Count: 1
Auto-bind Vlans Count: 1
Bridge Priority: 32768 BridgeID:
80:00:00:10:30:f9:9d:c0Bridge
Role : CIST Regional Root
CIST Root 80:00:00:10:30:f9:9d:c0CIST
Regional Root: 80:00:00:10:30:f9:9d:c0
Designated root: 00:00:00:00:00:00:00:00
RootPathCost: 0 External RootPathCost: 0 Root Port:
----- MaxAge:0sHelloTime:
0sForwardDelay:0s CfgBrMaxAge:20sCfgBrHelloTime:
2sCfgBrForwardDelay: 15s MaxHopCount: 20 CfgBrMaxHopCount : 20
Topology Change Time: 35s Hold time:
1s Topology Change Detected: FALSE Topology Change:
FALSE Number of Topology Changes:
0 Time Since Last Topology
Change: 0s
Participating Vlans :
(none) Auto-bind Vlans :
Default

```



History

This command was first available in ExtremeXOS 10.1.

Information about MSTP was added in ExtremeXOS 11.4.

Description was added in ExtremeXOS 12.4.4.

Platform Availability

This command is available on all platforms.

show stpd ports

```
show {stpd} stpd_name ports {[detail | port_list {detail}]}
```

Description

Displays the STP state of a port.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name.
<i>port_list</i>	Specifies one or more ports or slots and ports.
detail	Specifies more detailed information about one or more ports of the STPD.

Default

N/A.

Usage Guidelines

This command displays the following:

- STPD port configuration.
- STPD port encapsulation mode.
- STPD path cost.
- STPD priority.
- STPD state (root bridge, and so on).
- Port role (root designated, alternate and so on).
- STPD port state (forwarding, blocking, and so on).
- Configured port link type.
- Operational port link type.
- Edge port settings (inconsistent behavior, edge safeguard setting).
- Restricted role (enabled, disabled).



- MSTP port role (internal or boundary).
- Active port role.

To display more detailed information for one or more ports in the specified STPD, including participating VLANs, specify the detail option.

If you have MSTP configured and specify the detail option, this command displays additional information:

- MSTP internal path cost.
- MSTP timers.

If your STPD has the same name as another component, for example a VLAN, Extreme Networks recommends that you specify the identifying keyword as well as the name. If you do not specify the stpd keyword, an error message similar to the following is displayed:

```
%% Ambiguous command: "show Test ports"
```

In this example, to view all of the port settings of STPD Test, enter `show stpd Test ports`.

If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

Example

The following command displays the state of ports 1, 2, and 4 on an STPD named `s1`:

```
show stpd s1 ports
```

The following is sample output from this command:

```
Port   Mode   State      Cost  Flags      Priority Port ID Designated Bridge
1      EMISTP DISABLED 200000 e?pp-w---t 128      8001
00:00:00:00:00:00:00:00
2      EMISTP DISABLED 200000 e?pp-w---- 128      8002
00:00:00:00:00:00:00:00
4      EMISTP DISABLED 200000 e?pp-w---- 128      8004
00:00:00:00:00:00:00:00
Total Ports: 3
----- Flags: -----
1:          e=Enable, d=Disable
2: (Port role) R=Root, D=Designated, A=Alternate, B=Backup, M=Master
3: (Config type) b=broadcast, p=point-to-point, e=edge, a=auto
4: (Oper. type) b=broadcast, p=point-to-point, e=edge
5:          p=proposing, a=agree
6: (partner mode) d = 802.1d, w = 802.1w, m = mstp
7:          i = edgeport inconsistency
8:          S = edgeport safe guard active
s = edgeport safe guard configured but inactive
8:          G = edgeport safe guard bpdu restrict active in 802.1w and
mstp
g = edgeport safe guard bpdu restrict active in 802.1d
9:          B = Boundary, I = Internal
10:         r = Restricted Role, t = active Role
```



The following command displays the detailed information for the ports in STPD s1:

```
show stpd s1 ports 1 detail
```

The following is sample output from this command:

```
Stpd: s1          Port: 1 PortId: 8001    Stp: ENABLED    Path Cost: 20000
Port Mode: EMISTP
Port State:  DISABLED          Topology Change Ack: FALSE
Port Priority: 128
Designated Root:  00:00:00:00:00:00:00:00    Designated Cost: 0
Designated Bridge: 00:00:00:00:00:00:00:00    Designated Port Id: 0
Partner STP version: Dot1w
Restricted Role: Disabled
Active Role: Enabled
Edge Port Safe Guard: Disabled
Bpdu Restrict: Disabled
Participating Vlans: v1
```

The following command displays the detailed information for the ports in STPD s1 configured for MSTP:

```
show stpd s1 ports detail
```

The following is sample output from this command:

```
Stpd: s1          Port: 1 PortId: 8001    Stp: ENABLED    Path Cost: 4
Port Mode: 802.1D
Port State:  FORWARDING        Topology Change Ack: FALSE
Port Priority: 16
Designated Root:  80:00:00:04:96:1f:a8:44    Designated Cost: 0, IntCost: 0
Designated Bridge: 80:00:00:04:96:1f:a8:44    Designated Port Id: 8001
Partner STP version: MSTP
Restricted Role: Disabled
Active Role: Disabled
Edge Port Safe Guard: Disabled
maxAge: 20      msgAge: 0      fwdDelay: 15      helloTime: 2      maxHops: 20
Participating Vlans: v1
Stpd: s1          Port: 2 PortId: 8002    Stp: ENABLED    Path Cost: 4
Port Mode: 802.1D
Port State:  BLOCKING          Topology Change Ack: FALSE
Port Priority: 16
Designated Root:  80:00:00:04:96:1f:a8:44    Designated Cost: 0, IntCost: 0
Designated Bridge: 80:00:00:04:96:1f:a8:44    Designated Port Id: 8002
Partner STP version: Dot1d
Restricted Role: Enabled
Active Role: Disabled
Edge Port Safe Guard: Disabled
maxAge: 20      msgAge: 0      fwdDelay: 15      helloTime: 2      maxHops: 20
Participating Vlans: v1
```



The following is sample output from this command:

```

Port    Mode    State      Cost  Flags      Priority Port ID Designated Bridge
9       EMISTP FORWARDING 20000 eDeepw-G-- 128      8009
80:00:00:04:96:1f:a8:48
Total Ports: 1
----- Flags: -----
1:                e=Enable, d=Disable
2: (Port role)     R=Root, D=Designated, A=Alternate, B=Backup, M=Master
3: (Config type)  b=broadcast, p=point-to-point, e=edge, a=auto
4: (Oper. type)   b=broadcast, p=point-to-point, e=edge
5:                p=proposing, a=agree
6: (partner mode) d = 802.1d, w = 802.1w, m = mstp
7:                i = edgeport inconsistency
8:                S = edgeport safe guard active
s = edgeport safe guard configured but inactive
G = edgeport safe guard bpdu restrict active
g = edgeport safe guard bpdu restrict configured but inactive only dot1w, mstp
9:                B = Boundary, I = Internal
10:               r = Restricted Role, t = active role

```

History

This command was first available in ExtremeXOS 10.1.

Information about MSTP was added in ExtremeXOS 11.4.

Information about BPDU Restrict was added in ExtremeXOS 12.4.

Information about active role was added in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

show vlan stpd

```
show {vlan} vlan_name stpd
```

Description

Displays the STP configuration of the ports assigned to a specific VLAN.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
------------------	------------------------



Default

N/A.

Usage Guidelines

If you have a VLAN that spans multiple STPDs, use this command to display the STP configuration of the ports assigned to that specific VLAN.

This command displays the following:

- STPD port configuration.
- STPD port mode of operation.
- STPD path cost.
- STPD priority.
- STPD state (root bridge, and so on).
- Port role (root designated, alternate and so on).
- STPD port state (forwarding, blocking, and so on).
- Configured port link type.
- Operational port link type.

If your VLAN has the same name as another component, for example an STPD, Extreme Networks recommends that you specify the identifying keyword as well as the name. If you do not specify the vlan keyword, the switch displays an error message similar to the following:

```
%% Ambiguous command: "show Test stpd"
```

In this example, to view the STPD state of VLAN Test, enter `show vlan Test stpd`.

If you enter a VLAN name that is not associated with an STPD or does not exist, the switch displays an error message similar to the following:

```
Failed to find vlan 'vlan1' or it has no STP domains configured on it
```

If this happens, check to make sure you typed the correct name of the VLAN and that the VLAN is associated with an STPD.

If your VLAN has a name unique only to that VLAN, the keyword `vlan` is optional.

Example

The following command displays the spanning tree configurations for the VLAN Default:

```
show vlan default stpd
```

The following is sample output from this command:

```
s0(enabled) Tag: (none) Ports: 8 Root/P/C: 80:00:00:01:30:94:79:00/-----/0
Port  Mode  State      Cost  Flags  Priority Port ID Designated Bridge
1:1   802.1D LEARNING  19    eDbb-d- 16    8001
80:00:00:01:30:94:79:00
```



```

1:2      802.1D DISABLED  4      e----- 16      8002
00:00:00:00:00:00:00:00
1:3      802.1D DISABLED  4      e----- 16      8003
00:00:00:00:00:00:00:00
1:4      802.1D LEARNING  4      eDbb-d- 16      8004
80:00:00:01:30:94:79:00
1:5      802.1D LEARNING  4      eDbb-d- 16      8005
80:00:00:01:30:94:79:00
1:6      802.1D DISABLED  4      e----- 16      8006
00:00:00:00:00:00:00:00
1:7      802.1D DISABLED  4      e----- 16      8007
00:00:00:00:00:00:00:00
1:8      802.1D DISABLED  4      e----- 16      8008
00:00:00:00:00:00:00:00
----- Flags: -----
1:          e=Enable, d=Disable
2: (Port role)  R=Root, D=Designated, A=Alternate, B=Backup, M=Master,
Y=Boundary
3: (Config type) b=broadcast, p=point-to-point, e=edge, a=auto
4: (Oper. type) b=broadcast, p=point-to-point, e=edge
5:          p=proposing, a=agree
6: (partner mode) d=802.1d, w=802.1w, m=mstp
7:          i=edgeport inconsistency
8:          B = Boundary, I = Internal

```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

unconfigure mstp region

```
unconfigure mstp region
```

Description

Unconfigures the MSTP region on the switch and returns all MSTP settings to their default values.

Syntax Description

This command has no arguments or variables.

Default

N/A.



Usage Guidelines

Before you unconfigure an MSTP region, we recommend that you disable all active STPDs in the region. This includes the CIST and any active MSTIs.

After you issue this command, all of the MSTP settings return to their default values, as described below:

- **Region Name**—This indicates the name of the MSTP region. In the Extreme Networks implementation, the maximum length of the name is 32 characters and can be a combination of alphanumeric characters and underscores (_).
- **Format Selector**—This indicates a number to identify the format of MSTP BPDUs. The default is 0.
- **Revision Level**—This identifier is reserved for future use; however, the switch uses and displays a default of 3.

Example

The following command unconfigures the MSTP region on the switch:

```
unconfigure mstp region
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

unconfigure stpd

```
unconfigure stpd {stpd_name}
```

Description

Restores default STP values to a particular STPD or all STPDs.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
------------------	---------------------------------------

Default

N/A.



Usage Guidelines

If you create an STPD with a unique name, the keyword `stpd` is optional.

Use this command to restore default STP values to a particular STPD. If you want to restore default STP values on all STPDs, do not specify a spanning tree name.

Example

The following command restores default values to an STPD named `Backbone_st`:

```
unconfigure stpd backbone_st
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

unconfigure stpd ports link-type

```
unconfigure stpd stpd_name ports link-type port_list
```

Description

Returns the specified port to the factory default setting of broadcast link.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

All ports are broadcast link types.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, you must enter the `stpd` keyword to specify the STPD. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

If the switch operates in 802.1D mode, any configured port link type will behave the same as the broadcast link type.



In an MSTP environment, configure the same link types for the CIST and all MSTIs.

Example

The following command configures slot 2, ports 1 through 4 to return to the factory default of broadcast links in STPD s1:

```
unconfigure stpd s1 ports link-type 2:1-2:4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.



31 ESRP Commands

```
clear esrp counters
clear esrp neighbor
clear esrp sticky
configure esrp add elrp-poll ports
configure esrp add master
configure esrp add member
configure esrp add track-environment
configure esrp add track-iproute
configure esrp add track-ping
configure esrp add track-vlan
configure esrp aware add selective-forward-ports
configure esrp aware delete selective-forward-ports
configure esrp delete elrp-poll ports
configure esrp delete master
configure esrp delete member
configure esrp delete track-environment
configure esrp delete track-iproute
configure esrp delete track-ping
configure esrp delete track-vlan
configure esrp domain-id
configure esrp election-policy
configure esrp elrp-master-poll disable
configure esrp elrp-master-poll enable
configure esrp elrp-premaster-poll disable
configure esrp elrp-premaster-poll enable
configure esrp group
configure esrp mode
configure esrp name
configure esrp ports mode
configure esrp ports no-restart
configure esrp ports restart
configure esrp ports weight
configure esrp priority
configure esrp timer hello
configure esrp timer neighbor
configure esrp timer neutral
configure esrp timer premaster
```



```
configure esrp timer restart
create esrp
delete esrp
disable esrp
enable esrp
show esrp
show esrp aware
show esrp counters
```

This chapter describes the commands for:

- Enabling and disabling ESRP
- Performing ESRP configuration
- Enabling and disabling port restart and failure tracking for ESRP
- Displaying ESRP configuration information
- Enabling and disabling ELRP in an ESRP environment

For an introduction to ESRP, see the *ExtremeXOS Concepts Guide*.

clear esrp counters

```
clear esrp counters
```

Description

Clears the statistics gathered by ESRP for all ESRP domains on the switch.

Syntax Description

This command has no arguments or variables.

Default

None.

Usage Guidelines

Use this command to clear the state transition and the protocol packet counters gathered by ESRP.

The state transition count displays the number of times the ESRP domain entered the following states:

- **Aware**—An Extreme switch that does not participate in ESRP elections but is capable of listening to ESRP Bridge Protocol Data Units (BPDUs).
- **Master**—The master switch is the device with the highest priority based on the election algorithm. The master is responsible for responding to clients for Layer 3 routing and Layer 2 switching for the ESRP domain.



- Neutral—The neutral state is the initial state entered by the switch. In a neutral state, the switch waits for ESRP to initialize and run. A neutral switch does not participate in ESRP elections.
- PreMaster—The pre-master state is an ESRP switch that is ready to be master but is going through possible loop detection prior to transitioning to master.
- Slave—The slave switch participates in ESRP but is not elected or configured the master and does not respond to ARP requests but does exchange ESRP packets with other switches on the same VLAN. The slave switch is available to assume the responsibilities of the master switch if the master becomes unavailable or criteria for ESRP changes.

If the slave is in extended mode, it does not send ESRP hello messages; however, it sends PDUs that can trigger a change in the master switch.

For more information about configuring the ESRP mode of operation on the switch, see the `configure esrp mode [extended | standard]` command. By default, ESRP operates in extended mode.

To display information about the ESRP domain, including the previously described states, use the `show esrp { {name} | {type [vpls-redundancy | standard]} }` command.

The protocol packet count displays the number of times ESRP, ESRP-aware, and ESRP error packets were transmitted and received.

To display information about the ESRP counters, use the `show esrp {name} counters` command.

Example

The following command clears the statistics gathered by ESRP:

```
clear esrp counters
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

clear esrp neighbor

```
clear esrp esrpDomain neighbor
```

Description

Clears the neighbor information for the specified ESRP domain.



Syntax Description

<i>esrpDomain</i>	Specifies the name of an ESRP domain.
-------------------	---------------------------------------

Default

N/A.

Usage Guidelines

If you add a new switch to your ESRP domain, use this command to clear the existing neighbor information for the ESRP domain. After the switch is up, running, and configured as an ESRP-aware or ESRP-enabled device, new neighbor information is learned.

Before using this command, schedule a downtime for your network. Use this command for maintenance purposes only.

Example

The following command clears the existing neighbor information on the ESRP domain `esrp1` after adding a new switch to the ESRP domain:

```
clear esrp esrp1 neighbor
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

clear esrp sticky

```
clear esrp esrpDomain sticky
```

Description

Clears the stickiness in the ESRP domain and forces the election of the ESRP master switch.

Syntax Description

<i>esrpDomain</i>	Specifies the name of an ESRP domain.
-------------------	---------------------------------------



Default

N/A.

Usage Guidelines

Use the `clear esrp sticky` command to force the election of the ESRP master switch. Before using this command, schedule a downtime for your network.

For example, without stickiness configured, if an event causes the ESRP master to failover to the backup, the previous backup becomes the new master. If another event causes the new master to return to backup, you have experienced two network interruptions. To prevent this, use the `configure esrp election-policy` command and select stickiness as an election algorithm.

If you use sticky as an election metric, and an event causes the ESRP master to failover, ESRP assigns the new master with the highest sticky election metric of 1. Therefore, regardless of changes to the neighbor's election algorithm, the new master retains its position. Sticky is set on the master switch only.

ESRP re-election can occur if sticky is set on the master and a local event occurs. During this time, if the current master has lower election parameters, the backup can become the new master.

If you use `clear esrp esrpDomain sticky` command, it only affects the current master and can trigger ESRP re-election.

Example

The following command clears the stickiness on the ESRP domain esrp1:

```
clear esrp esrp1 sticky
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp add elrp-poll ports

```
configure esrp esrpDomain add elrp-poll ports [ports | all]
```

Description

Configures the ports of an ESRP domain where ELRP packet transmission is requested by ESRP.



Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>ports</i>	Specifies list of slots and ports.
all	Specifies all ports in the ESRP domain.

Default

All ports of an ESRP domain have ELRP transmission enabled.

Usage Guidelines

This command allows you to configure the ports in your network that might experience loops, such as ports that connect to master, slave, or ESRP-aware switches, to receive ELRP packets. You do not need to send ELRP packets to host ports.



Note

The ExtremeXOS software does not support ELRP and Network Login on the same port.

Example

The following command enables ELRP packet transmission for slot 2, ports 3-5 on ESRP domain esrp1:

```
configure esrp esrp1 add elrp-poll ports 2:3-2:5
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure esrp add master

```
configure esrp esrpDomain add master vlan_name
```

Description

Adds a master VLAN to an ESRP domain.



Syntax Description

<i>esrpDomain</i>	Specifies the name of an ESRP domain.
<i>vlan_name</i>	Specifies the name of the master VLAN.

Default

N/A.

Usage Guidelines

You must configure one master VLAN for each ESRP domain. A master VLAN can belong to one ESRP domain only. An ESRP domain contains one master and zero or more member VLANs.

The master VLAN:

- Exchanges ESRP PDUs, hello messages, and data between a pair of ESRP-enabled switches.
- Contains the total number of active physical ports that are counted when determining the master ESRP domain. The switch with the highest number of active ports takes priority.

Master VLANs can have their own set of ports, and member VLANs can have a different set of ports. The state of the ESRP device determines whether the ports in the master and member VLANs are in the forwarding or blocking state.

Example

The following command adds VLAN purple to the ESRP domain esrp1 as the master VLAN:

```
configure esrp esrp1 add master purple
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp add member

```
configure esrp esrpDomain add member vlan_name
```

Description

Adds a member VLAN to an ESRP domain.



Syntax Description

<i>esrpDomain</i>	Specifies the name of an ESRP domain.
<i>vlan_name</i>	Specifies the name of the member VLAN.

Default

N/A.

Usage Guidelines

You can configure zero or more member VLANs for each ESRP domain. An ESRP domain contains one master and zero or more member VLANs.

Master VLANs can have their own set of ports, and member VLANs can have a different set of ports. The state of the ESRP device determines whether the ports in the master and member VLANs are in the forwarding or blocking state.

Example

The following command adds VLAN green to the ESRP domain esrp1 as a member VLAN:

```
configure esrp esrp1 add member vlan green
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp add track-environment

```
configure esrp esrpDomain add track-environment failover priority
```

Description

Configures an ESRP domain to track environmental failures.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>priority</i>	Specifies a number between 0 and 254. The default priority is 255. See the following "Usage Guidelines" section for more information.



Default

No environmental tracking.

Usage Guidelines

Environmental tracking tracks power supply and chassis temperature status.

If a failure is detected, the ESRP domain priority steps to the failover-priority value specified. By setting the failover priority to be lower than the normal priority of the domain, it causes the affected domain to go into slave mode.

The range of the priority value is 0 to 254. Setting the priority to 255 configures the switch to slave mode, and to be ineligible to become the master. The switch remains in slave mode even when the VLAN fails over from the current master.

To make effective use of this feature, the normal priority of the ESRP domain must be higher than the failover priority of this command.

Example

The following command enables environmental failure tracking, and specifies that the ESRP priority for ESRP domain `esrp1` be set to 10 upon an environmental failure.

```
configure esrp esrp1 add track-environment failover 10
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp add track-iproute

```
configure esrp esrpDomain add track-iproute ipaddress/masklength
```

Description

Configures an ESRP domain to track a route entry in the system's routing table.



Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>ipaddress</i>	Specifies the IPv4 address of the route entry to be tracked.
<i>masklength</i>	Specifies the subnet of the route entry to be tracked.

Default

Disabled.

Usage Guidelines

The track-ip metric consists of the total number of tracked IPv4 routes that are up or functional.

An ESRP domain can track eight IPv4 routes.



Note

ESRP route tracking is not supported on IPv6 networks.

Example

The following command enables IPv4 route failure tracking for routes to the specified subnet:

```
configure esrp esrp1 add track-iproute 192.168.46.0/24
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp add track-ping

```
configure esrp esrpDomain add track-ping ipaddress frequency seconds miss misses
```

Description

Configures an ESRP domain to track an external gateway using ping.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>ipaddress</i>	Specifies the IPv4 address of the external gateway.



<i>seconds</i>	Specifies the interval in seconds between ping requests. The range is 1 to 600 seconds.
<i>misses</i>	Specifies the number of consecutive ping failures that initiates failover to an ESRP slave. The range is 1 to 256 pings.

Default

No ping tracking.

Usage Guidelines

The tracked-ping metric consists of the total number of stations that are successfully tracked using ping. ESRP uses an aggregate of tracked pings and traced routes to track an external gateway.

An ESRP domain can track eight stations.



Note

ESRP ping tracking is not supported on IPv6 networks.

Example

The following command enables ping tracking for the external gateway at 10.207.29.17, pinging every 10 seconds, and considering the gateway to be unreachable if no response is received to 5 consecutive pings:

```
configure esrp esrp1 add track-ping 10.207.29.17 frequency 10 miss 5
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp add track-vlan

```
configure esrp esrpDomain add track-vlan vlan_name
```

Description

Configures an ESRP domain to track port connectivity to a specified VLAN.



Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>vlan_name</i>	Specifies the VLAN to be tracked.

Default

Disabled.

Usage Guidelines

The track-vlan metric is derived from the total number of active physical ports on the VLAN being tracked by the ESRP domain.

If more than one VLAN shares a physical link, each VLAN counts the physical link.

The ESRP switch should have a higher priority number than its neighbors to ensure master election.

An ESRP domain can track one VLAN, and the tracked VLAN should not be a member of any other ESRP domain in the system.

Example

The following command enables ESRP domain esrp1 to track port connectivity to VLAN engineering:

```
configure esrp esrp1 add track-vlan engineering
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp aware add selective-forward-ports

```
configure esrp domain aware add selective-forward-ports port_list {group group number}
```

Description

Enables selective forwarding by creating an aware port list and adds additional ports to the list.



Syntax Description

<i>domain</i>	Specifies an ESRP domain name.
<i>port_list</i>	Specifies the ports to be added to the aware port list.
<i>group number</i>	Specifies the ESRP group within the given domain name

Default

The group number defaults to '0'.

Usage Guidelines

An ESRP-aware switch floods ESRP PDUs from all ports in an ESRP-aware VLAN. This flooding creates unnecessary network traffic because some ports forward ESRP PDUs to switches that are not running the same ESRP groups. You can select the ports that are appropriate for forwarding ESRP PDUs by configuring selective forwarding on an ESRP-aware VLAN and thus reduce this excess traffic. Configuring selective forwarding creates a port list of only those ports that forward to the ESRP groups that are associated with an ESRP-aware VLAN. This ESRP-aware port list is then used for forwarding ESRP PDUs.

Use this command to create or add to an existing port list for the ESRP groups associated with an ESRP-aware VLAN.

Example

The following command configures esrp domain (d1) to forward ESRP PDUs on ports 5:1, 5:2, and 6:2.

```
configure esrp d1 aware add selective-forward-ports 5:1,5:2,6:2 group 0
```

History

This command was first available in Extreme XOS 12.0.

Platform Availability

This command is available on all platforms.

configure esrp aware delete selective-forward-ports

```
configure esrp domain aware delete selective-forward-ports all | port_list { group  
group number }
```

Description

Disables all or part of selective forwarding by deleting ports from the ESRP-aware port list.



Syntax Description

<i>domain</i>	Specifies an ESRP domain name.
all	Specifies that all of the ports are to be disabled.
<i>port_list</i>	Specifies the ports to be disabled from the ESRP-aware port list.
<i>group number</i>	Specifies the ESRP group within the given domain name

Default

The group number defaults to '0'.

Usage Guidelines

By configuring selective forwarding, you create an ESRP-aware port list of only those ports that forward to the ESRP groups that are associated with an ESRP-aware VLAN. That port list is used for forwarding ESRP PDUs from the selected ports only of an ESRP-aware switch.

Use this command to delete one or more or all of the ports from an ESRP-aware port list. Deleting all of the ports puts the domain back to the default state.

Example

The following command configures esrp domain (d1) to exclude ESRP PDUs on ports 5:1, 5:2, and 6:2.

```
configure esrp d1 aware delete selective-forward-ports 5:1,5:2,6:2 group 0
```

History

This command was first available in Extreme XOS 12.0.

Platform Availability

This command is available on all platforms.

configure esrp delete elrp-poll ports

```
configure esrp esrpDomain delete elrp-poll ports [ports | all]
```

Description

Disables ELRP packet transmission on ports of an ESRP domain.



Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>ports</i>	Specifies list of slots and ports in the ESRP domain.
all	Specifies all ports in the ESRP domain.

Default

All ports of an ESRP domain have ELRP transmission enabled.

Usage Guidelines

If you have host ports on an ESRP domain, you do not need to send ELRP packets to those ports.

If you change your network configuration, and a port no longer connects to a master, slave, or ESRP-aware switch, you can disable ELRP transmission on that port.

Example

The following command disables ELRP packet transmission for slot 2, ports 3-5 on ESRP domain esrp1:

```
configure vlan esrp1 delete elrp-poll ports 2:3-2:5
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure esrp delete master

```
configure esrp esrpDomain delete master vlan_name
```

Description

Deletes the specifies master VLAN from the specified ESRP domain.

Syntax Description

<i>esrpDomain</i>	Specifies the name of an ESRP domain.
<i>vlan_name</i>	Specifies the name of the master VLAN.



Default

N/A.

Usage Guidelines

You must disable the ESRP domain before removing the master VLAN. To disable the ESRP domain, use the `disable esrp {esrpDomain}` command.

If you attempt to remove the master VLAN before disabling the ESRP domain, the switch displays an error message similar to the following:

```
ERROR: Failed to delete master vlan for domain "esrp1" ; ESRP is enabled!
```

If this happens, disable the ESRP domain and re-issue the `configure esrp delete master` command.

Example

The following command deletes the master VLAN purple from the ESRP domain esrp1:

```
configure esrp esrp1 delete master purple
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp delete member

```
configure esrp esrpDomain delete member vlan_name
```

Description

Deletes a member VLAN from the specified ESRP domain.

Syntax Description

<i>esrpDomain</i>	Specifies the name of an ESRP domain.
<i>vlan_name</i>	Specifies the name of the member VLAN.



Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the member VLAN green from the ESRP domain esrp1:

```
configure esrp esrp1 delete member vlan green
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp delete track-environment

```
configure esrp esrpDomain delete track-environment
```

Description

Disables environmental failure tracking for an ESRP domain.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
-------------------	--------------------------------

Default

No environmental tracking.

Usage Guidelines

None.



Example

The following command disables environmental failure tracking for ESRP domain esrp1:

```
configure esrp esrp1 delete track-environment
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp delete track-iproute

```
configure esrp esrpDomain delete track-iproute ipaddress/masklength
```

Description

Disables route entry tracking for an ESRP domain.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>ipaddress</i>	Specifies the IPv4 address of a tracked route entry.
<i>masklength</i>	Specifies the subnet of a tracked route entry.

Default

Disabled.

Usage Guidelines

If you disable route tracking for a failed route, the ESRP domain recovers from the forced standby state.

If you disable route tracking for a route that is up and functional, there is no impact on the ESRP state.

Example

The following command disables tracking of routes to the specified subnet for ESRP domain esrp1:

```
configure esrp esrp1 delete track-iproute 192.168.46.0/24
```



History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp delete track-ping

```
configure esrp esrpDomain delete track-ping ipaddress
```

Description

Disables the tracking of an external gateway using ping.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>ipaddress</i>	Specifies the IPv4 address of the external gateway.

Default

No ping tracking.

Usage Guidelines

If you disable ping tracking for a failed ping, the ESRP domain recovers from the forced standby state.

If you disable route tracking for a successful ping, there is no impact on the ESRP state.

Example

The following command disables ping tracking for the external gateway at 10.207.29.17:

```
configure esrp esrp1 delete track-ping 10.207.29.17
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.



configure esrp delete track-vlan

```
configure esrp esrpDomain delete track-vlan vlan_name
```

Description

Disables the tracking of port connectivity to a specified VLAN.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>vlan_name</i>	Specifies the VLAN to be tracked.

Default

Disabled.

Usage Guidelines

If you delete a VLAN that is down, the ESRP domain recovers from the forced standby state.

Example

The following command disables the tracking of port connectivity to VLAN engineering:

```
configure esrp esrp1 delete track-vlan engineering
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp domain-id

```
configure esrp esrpDomain domain-id number
```

Description

Assigns an ESRP domain ID to an ESRP domain.



Syntax Description

<i>esrpDomain</i>	Specifies the name of an ESRP domain.
<i>number</i>	Specifies the number to use for the ESRP domain ID. The user-configured ID range is 4096 through 65,535.

Default

If the master VLAN is tagged, ESRP uses that VLANid for the ESRP domain ID. If the master VLAN is untagged, you must specify the ESRP domain ID.

Usage Guidelines

Before you enable a specific ESRP domain, it must have a domain ID. A domain ID is either a user-configured number or the VLANid of the tagged master VLAN. If you do not have a domain ID, you cannot enable ESRP on that domain.

Each switch participating in ESRP for a particular domain must have the same domain ID configured.

The number parameter range for user-configured domain IDs is 4096 through 65,535.

If the master VLAN is tagged, you can use that VLANid for the ESRP domain ID. The range for VLAN tags is 2 through 4095. Tag 1 is assigned to the default VLAN.

Example

The following command assigns the domain ID 5000 to ESRP domain esrp1:

```
configure esrp esrp1 domain-id 5000
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp election-policy

```
configure esrp esrpDomain election-policy [ports > track > priority | ports > track > priority > mac | priority > mac | priority > ports > track > mac | priority > track > ports > mac | sticky > ports > track > priority | sticky > ports > track > priority > mac | sticky > ports > weight > track > priority > mac | sticky > priority > mac | sticky > priority > ports > track > mac | sticky > priority > track > ports > mac | sticky > track > ports > priority | sticky >
```



```
track > ports > priority > mac | track > ports > priority | track > ports >
priority > mac]
```

Description

Configures the election algorithm on the switch.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
ports > track > priority	Specifies that this ESRP domain should consider election factors in the following order: Active ports, tracking information, ESRP priority.
ports > track > priority > mac	Specifies that this ESRP domain should consider election factors in the following order: Active ports, tracking information, ESRP priority, MAC address. NOTE: This is the default election algorithm for standard mode.
priority > mac	Specifies that this ESRP domain should consider election factors in the following order: ESRP priority, MAC address.
priority > ports > track > mac	Specifies that this ESRP domain should consider election factors in the following order: ESRP priority, active ports, tracking information, MAC address.
priority > track > ports > mac	Specifies that this ESRP domain should consider election factors in the following order: ESRP priority, tracking information, active ports, MAC address.
sticky > ports > track > priority	Specifies that this ESRP domain should consider election factors in the following order: Stickiness, active ports, tracking information, ESRP priority.
sticky > ports > track > priority > mac	Specifies that this ESRP domain should consider election factors in the following order: Stickiness, active ports, tracking information, ESRP priority, MAC address.
sticky > ports > weight > track > priority > mac	Specifies that this ESRP domain should consider election factors in the following order: Stickiness, active ports, port weight, tracking information, ESRP priority, MAC address. NOTE: Beginning with ExtremeXOS 11.1 and later, this is the default election algorithm for extended mode.
sticky > priority > mac	Specifies that this ESRP domain should consider election factors in the following order: Stickiness, ESRP priority, MAC address.
sticky > priority > ports > track > mac	Specifies that this ESRP domain should consider election factors in the following order: Stickiness, ESRP priority, active ports, tracking information, MAC address.
sticky > priority > track > ports > mac	Specifies that this ESRP domain should consider election factors in the following order: Stickiness, ESRP priority, tracking information, active ports, MAC address.
sticky > track > ports > priority	Specifies that this ESRP domain should consider election factors in the following order: Stickiness, tracking information, active ports, ESRP priority.
sticky > track > ports > priority > mac	Specifies that this ESRP domain should consider election factors in the following order: Stickiness, tracking information, active ports, ESRP priority, MAC address.



track > ports > priority	Specifies that this ESRP domain should consider election factors in the following order: Tracking information, active ports, ESRP priority.
track > ports > priority > mac	Specifies that this ESRP domain should consider election factors in the following order: Tracking information, active ports, ESRP priority, MAC address.

Default

In extended mode, the default election algorithm is sticky > ports > weight > track > priority > mac.

In standard mode, the default election algorithm is ports > track > priority > mac.

Usage Guidelines

The election algorithm determines the order of precedence of the election factors used to determine the ESRP Master. The election factors are:

- Stickiness (sticky): the switch with the higher sticky value has higher priority. When an ESRP domain claims master, its sticky value is set to 1 (available in extended mode only).
- Active Ports (ports): the number of active ports (the switch with the highest number takes priority)
- Tracking Information (track): whether the switch is using ESRP tracking. A switch using tracking has priority.
- ESRP Priority (priority): a user-defined priority number between 0 and 254. A higher number has higher priority. The default priority setting is 0. A priority setting of 255 makes an ESRP switch a standby switch that remains in slave mode until you change the priority setting. We recommend this setting for system maintenance. A switch with a priority setting of 255 never becomes the master.
- MAC address (mac): the switch MAC address. A higher-number address has priority.
- Active port weight (weight)—The switch that has the highest port weight takes precedence. The bandwidth of the port automatically determines the port weight (available only in extended mode). ESRP does not count ports with a weight of 0 (known as don't count ports) regardless of ESRP running in extended or standard mode.

The election algorithm must be the same on all switches for a particular ESRP domain. The election algorithms that use sticky are and weight are available in extended mode only.

In ExtremeXOS 11.0, the extended mode default election algorithm is: sticky > ports > track > priority > mac > weight. This election algorithm is not supported in ExtremeXOS 11.1.

Factors to Consider

The ports-track-priority or track-ports-priority options can be used to ensure that there is no failback if the original Master recovers (the Master has the same ports, tracks and priority, but a higher MAC).

Any of the options with sticky can also be used to ensure that there is no failback if the original master recovers. With sticky, if an event causes the ESRP master to failover, ESRP assigns the new master with the sticky count of 1. After sticky is set on the master, regardless of changes to its neighbor's election algorithm, the new master retains its position. For example, adding active ports to the slave does not cause the new master to failback to the original master, even if the slave has more active ports than the



master. Sticky algorithms provide for fewer network interruptions than non-sticky algorithms. Sticky is set on the master switch only.

ESRP re-election can occur if sticky is set on the master and a local event occurs. During this time, if the current master has lower election parameters, the backup can become the new master.

Switch Behavior

If a switch is master, it actively provides Layer 3 routing services to other VLANs, and Layer 2 switching between all the ports of that VLAN. Additionally, the switch exchanges ESRP packets with other switches that are in slave mode.

If a switch is in slave mode, it exchanges ESRP packets with other switches on that same VLAN. When a switch is in slave mode, it does not perform Layer 3 routing or Layer 2 switching services for the VLAN.

Updating the Election Algorithm

ESRP uses the default election policy for extended mode. If you have an ESRP domain operating in standard mode, the domain ignores the sticky and weight algorithms. To change the election algorithm, you must first disable the ESRP domain and then configure the new election algorithm. If you attempt to change the election algorithm without disabling the domain first, an error message appears.

To disable the ESRP domain, use the following command:

```
disable esrp {esrpDomain}
```

To modify the election algorithm, use the following command:

```
configure esrp esrpDomain election-policy [ports > track > priority | ports >
track > priority > mac | priority > mac | priority > ports > track > mac |
priority > track > ports > mac | sticky > ports > track > priority | sticky >
ports > track > priority > mac | sticky > ports > weight > track > priority > mac
| sticky > priority > mac | sticky > priority > ports > track > mac | sticky >
priority > track > ports > mac | sticky > track > ports > priority | sticky >
track > ports > priority > mac | track > ports > priority | track > ports >
priority > mac]
```

If you attempt to use an election algorithm not supported by the switch, an error message similar to the following appears:

```
ERROR: Specified election-policy is not supported!
Supported Policies:
1. sticky > ports > weight > track > priority > mac
2. ports > track > priority
3. sticky > ports > track > priority
4. ports > track > priority > mac
5. sticky > ports > track > priority > mac
6. priority > mac
7. sticky > priority > mac
8. priority > ports > track > mac
9. sticky > priority > ports > track > mac
10. priority > track > ports > mac
11. sticky > priority > track > ports > mac
```



12. `track > ports > priority`
13. `sticky > track > ports > priority`
14. `track > ports > priority > mac`
15. `sticky > track > ports > priority > mac`

Example

The following command configures the election algorithm to use tracking information as the first criteria for determining the ESRP master switch for ESRP domain `esrp1`:

```
configure esrp esrp1 election-policy track > ports > priority > mac
```

History

This command was first available in ExtremeXOS 11.0.

The default election algorithm for extended mode was updated to `sticky > ports > weight > track > priority > mac`, and the weight election factor was used in ExtremeXOS 11.1. The `sticky > ports > track > priority > mac > weight` election algorithm is not supported in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure esrp elrp-master-poll disable

```
configure esrp esrpDomain elrp-master-poll disable
```

Description

Disables the use of ELRP by ESRP in the master state.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
-------------------	--------------------------------

Default

Disabled.

Usage Guidelines

Use this command to disable the use of ELRP by ESRP in the master state. When you disable ELRP, the ESRP master switch no longer transmits ELRP PDUs to detect network loops.



Example

The following command disables the use of ELRP in the master state on ESRP domain elrp1:

```
configure esrp elrp1 esrp elrp-master poll disable
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure esrp elrp-master-poll enable

```
configure esrp esrpDomain elrp-master-poll enable {interval interval}
```

Description

Enables the use of ELRP by ESRP in the master state, and configures how often the master checks for loops in the network.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>interval</i>	Specifies how often, in seconds, successive ELRP packets are sent. The default is 1 second. The range is 1 to 64 seconds.

Default

- Use of ELRP in the master state—disabled
- Interval—1 second

Usage Guidelines

Use this command to enable the use of ELRP by ESRP in the master state. When an ESRP-enabled switch is in the master state, and you enable `elrp-master-poll`, the switch periodically sends ELRP PDUs at the configured interval level. If a loop is detected in the network, the transmitted PDUs are received by the switch. The ESRP master switch then transitions to the slave state to break the network loop.

We recommend that you enable both premaster and master polling when using ELRP with ESRP. To enable premaster polling, use the `configure esrp esrpDomain elrp-premaster-poll enable {count count | interval interval}`.



If you attempt to configure master polling before premaster polling, the switch displays an error message similar to the following:

```
ERROR: Premaster-poll should be enabled before enabling master-poll!
```

If this happens, first configure premaster polling followed by master polling (if required).

Specify the interval parameter to configure how often successive ELRP PDUs are sent while in the master state. If you do not specify an interval value, the default value is used.

Example

The following command enables the use of ELRP in the master state on ESRP domain elrp1:

```
configure esrp elrp1 esrp elrp-master poll enable
```

The following command configures the ESRP master to check for loops in the network every 3 seconds:

```
configure esrp elrp1 esrp elrp-master-poll enable interval 3
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure esrp elrp-premaster-poll disable

```
configure esrp esrpDomain elrp-premaster-poll disable
```

Description

Disables the use of ELRP by ESRP in the pre-master state.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
-------------------	--------------------------------

Default

Disabled.



Usage Guidelines

Use this command to disable the use of ELRP by ESRP in the pre-master state. When you disable ELRP in the pre-master state, the ESRP pre-master switch no longer transmits ELRP PDUs to detect network loops prior to changing to the master state.

Example

The following command disables the use of ELRP in the pre-master state on the ESRP domain elrp1:

```
configure esrp elrp1 esrp elrp-premaster poll disable
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure esrp elrp-premaster-poll enable

```
configure esrp esrpDomain elrp-premaster-poll enable {count count | interval interval}
```

Description

Enables the use of ELRP by ESRP in the pre-master state, and configures how many times the switch sends ELRP PDUs and how often the switch sends ELRP PDUS in the pre-master state.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>count</i>	Specifies the number of times the switch sends ELRP PDUs. The default is 3. The range is 1 to 32.
<i>interval</i>	Specifies how often, in seconds, the ELRP PDUs are sent. The default is 1 second. The range is 1 to 32 seconds.

Default

- Use of ELRP in the pre-master state—disabled
- Count—3 times
- Interval—1 second



Usage Guidelines

Use this command to enable the use of ELRP by ESRP in the pre-master state to prevent network loops from occurring. When an ESRP-enabled switch is in the pre-master state (waiting to become the master), and you enable `elrp-premaster-poll`, the switch periodically sends ELRP PDUs at the configure level for a specified number of times. If there is a loop in the network, the transmitted PDUs are received by the switch. If this happens, the ESRP pre-master switch does not transition to the master state; rather, the switch transitions to the slave state.

We recommend that you enable both premaster and master polling when using ELRP with ESRP. To enable master polling, use the `configure esrp esrpDomain elrp-master-poll enable {interval interval}`.

If you attempt to configure master polling before premaster polling, the switch displays an error message similar to the following:

```
ERROR: Premaster-poll should be enabled before enabling master-poll!
```

If this happens, first configure premaster polling followed by master polling (if required).

If you do not specify the optional count or interval parameters, the default values are used.

If the sender does not receive packets, there is no loop in the network.

Example

The following command enables the use of ELRP—with the default settings—in the pre-master state on ESRP domain `elrp1`:

```
configure esrp elrp1 esrp elrp-premaster poll enable
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure esrp group

```
configure esrp esrpDomain group number
```

Description

Configures the group number to be used for the ESRP domain.



Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>group number</i>	Specifies the ESRP group number to which this ESRP domain should be added. The range is 0 through 31.

Default

The default group number is 0.

Usage Guidelines

Each group runs an instance of ESRP within the same VLAN or broadcast domain. A maximum of seven ESRP groups can be defined within the same networked broadcast domain. In addition, a maximum of seven distinct ESRP groups can be supported on a single ESRP switch. You can configure a maximum of 32 ESRP groups in a network.

The range for the `group_number` parameter is 0 through 31.

The most typical application for multiple ESRP groups is when two or more sets of ESRP switches are providing fast-failover protection within a common subnet for two or more groups of users. An additional use for ESRP groups is ESRP Host Attach; ESRP VLANs that share the same ESRP HA ports must be members of different ESRP groups.

You must first disable an ESRP domain before you modify an existing or add a new group number. If you try to modify the group number without disabling the ESRP domain, an error message similar to the following is displayed:

```
ERROR: can't change ESRP group for active domain "esrp1"!
```

To disable an ESRP domain, use the `disable esrp {esrpDomain}` command.

Example

The following command configures ESRP domain `esrp1` to be a member of ESRP group 2:

```
configure esrp esrp-1 group 2
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.



configure esrp mode

```
configure esrp mode [extended | standard]
```

Description

Configures the mode of operation for ESRP on the switch.

Syntax Description

extended	Specifies ESRP extended mode. This mode is compatible with devices running ExtremeXOS and is not backward compatible with devices running ExtremeWare.
standard	Specifies ESRP standard mode. This mode is backward compatible with devices running ExtremeWare.

Default

The default mode is extended.

Usage Guidelines

Use standard ESRP if your network contains a combination of switches running ExtremeXOS and ExtremeWare participating in ESRP. With standard ESRP, the switches running ExtremeXOS are backward compatible with the switches running ExtremeWare.

Use extended ESRP if your network contains switches running only ExtremeXOS; this is the default.

If your network has switches currently running ExtremeWare, and you add a BlackDiamond 8800 series switch, SummitStack, or a Summit family switch running ExtremeXOS, select standard ESRP. By selecting standard, the switch running ExtremeXOS is backward compatible with the ExtremeWare implementation of ESRP.

If you use the default mode—extended—and your ESRP domain contains a switch running ExtremeXOS that detects a neighbor switch running ExtremeWare, the mode automatically changes to standard for that domain. This action causes the switch to enter the neutral state and re-elect the ESRP master. Since you are using the default mode of operation, and the switch running ExtremeXOS detected a neighbor switch running ExtremeWare, the ExtremeXOS switch toggles to standard mode although the configured mode of operation remains as extended.

Note



ExtremeWare switches forward only those ESRP hello messages that apply to the ESRP group to which the switch belongs. ExtremeWare switches do not forward ESRP hello messages for other ESRP groups in the same VLAN. This limitation does not apply to ExtremeXOS switches operating in standard mode.



Example

The following command configures ESRP to run in standard mode:

```
configure esrp mode standard
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp name

```
configure esrp esrpDomain name new-name
```

Description

Renames an existing ESRP domain.

Syntax Description

<i>esrpDomain</i>	Specifies the current name of an ESRP domain.
<i>new-name</i>	Specifies a new name for the ESRP domain.

Default

N/A.

Usage Guidelines

The maximum length for a name is 32 characters. Names can contain alphanumeric characters and underscores (_) but cannot be any reserved keywords, for example, esrp. Names must start with an alphabetical character, for example, a, Z.

You can rename an ESRP domain regardless of its current state.

Example

The following command renames ESRP domain esrp1 to esrp3:

```
configure esrp esrp1 name esrp3
```



History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp ports mode

```
configure esrp ports ports mode [host | normal]
```

Description

Configures the ESRP port mode for ESRP host attach.

Syntax Description

<i>ports</i>	Specifies one or more ports or slots and ports that should be configured.
host	Specifies that the ports should be configured as host ports.
normal	Specifies that the ports should be configured as normal ports.

Default

The default port mode is normal.

Usage Guidelines

Ports configured as normal ports do not accept or transmit Layer 2 or Layer 3 traffic when the local ESRP device is a slave.

Ports configured as host ports allow the network to continue operation independent of ESRP status. The command sets the port to forward, allowing those ports directly attached to the slave's hosts to communicate with other hosts that are connected to the master. If you use load sharing with the ESRP HA feature, configure the load-sharing group first and then enable Host Attach on the group.

A Layer 2 connection for VLANs between ESRP switches is required.

An ESRP Host Attach port cannot be a mirroring port, software-controlled redundant port, or Netlogin port.

Example

The following command configures ports 1 through 5 on slot 3 as host ports:

```
configure esrp port 3:1-3:5 mode host
```



History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp ports no-restart

configure esrp ports *ports* no-restart

Description

Disables port restart for a port.

Syntax Description

<i>ports</i>	Specifies one or more ports or slots and ports.
--------------	---

Default

N/A.

Usage Guidelines

None.

Example

The following command disables port restart for ports 7-9 in slot 3 in the ESRP master domain:

```
configure esrp port 3:7-3:9 no-restart
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp ports restart

configure esrp ports *ports* restart



Description

Configures ESRP to restart ports if there is a state change and the downstream switch is from another vendor.

Syntax Description

<i>ports</i>	Specifies one or more ports or slots and ports.
--------------	---

Default

N/A.

Usage Guidelines

If an ESRP domain becomes a slave, ESRP disconnects member ports that have port restart enabled. The disconnection of these ports causes downstream devices to remove the ports from their FDB tables. After 3 seconds the ports re-establish connection with the ESRP-enabled device. This feature allows you to use ESRP in networks that include equipment from other vendors.

If switch becomes a slave, ESRP disconnects the physical links of member ports that have port restart enabled.

An ESRP restart port cannot be a mirroring port, software-controlled redundant port, or Netlogin port.

Example

The following command enables port restart for ports 7-9 in slot 3 on the ESRP master domain:

```
configure esrp port 3:7-3:9 restart
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp ports weight

```
configure esrp ports ports weight [auto | port-weight]
```

Description

Assigns the port weight for the specified ESRP port(s).



Syntax Description

<i>ports</i>	Specifies one or more ports or slots and ports.
auto	Specifies the switch to calculate the weight of a port based on the port's bandwidth and link speed.
<i>port-weight</i>	Specifies an ESRP port weight of 0. With a port weight of 0, the ports are not counted.

Default

The switch automatically calculates the weight of a port based on the bandwidth of the port.

Usage Guidelines

Use this command to override the automatically calculated port weight.

The *port-weight* parameter specifies a weight of 0. With this configuration, ESRP does not count host ports and normal ports as active. With a weight of 0, ESRP experiences fewer state changes due to frequent client activities like rebooting and unplugging laptops. A don't-count port cannot be a mirroring, software-controlled redundant port, or a Netlogin port.

For load shared ports, configure one master port in the load-share group with the port weight. A single command specifies the weight for the entire load shared group. You can specify any port from the load share group in the command. A load-shared port has an aggregate weight of all of its member ports. If you add or delete a member port (or trunk), the weight of the master load-shared port is updated. For more information about load sharing, see “Configuring Slots and Ports on a Switch” in the *ExtremeXOS Concepts Guide*.

Example

The following command configures port 1 on slot 3 with a weight of 0:

```
configure esrp port 3:1 weight 0
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure esrp priority

```
configure esrp esrpDomain priority number
```



Description

Configures the ESRP priority.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain number.
<i>number</i>	Specifies a number between 0 and 255.

Default

The default ESRP priority is 0.

Usage Guidelines

The ESRP priority is one of the factors used by the ESRP election algorithm in determining which switch is the Master switch.

The range of the priority value is 0 to 254, with 0 being the lowest priority, 254 being the highest. If the ESRP priority is the determining criteria for the election algorithm, the highest priority value determines which switch acts as master for a particular ESRP domain.

Setting the priority to 255 configures the switch to slave mode, and to be ineligible to become the master. The switch remains in slave mode even when the ESRP domain fails over from the current master. This feature is typically used to ensure a switch cannot become the ESRP master while it is offline for servicing.

Example

The following command configures the ESRP priority to the highest priority on ESRP domain esrp1:

```
configure esrp esrp1 priority 254
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp timer hello

```
configure esrp esrpDomain timer hello seconds
```



Description

Configures the ESRP hello timer value.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>seconds</i>	Specifies the number of seconds between keep-alive packets. The range is 1 to 255 seconds.

Default

The default hello timer is 2 seconds.

Usage Guidelines

The timer specifies the interval, in seconds, for exchanging keep-alive packets between the ESRP switches for this ESRP domain. A lower value specifies a more frequent exchange of keep-alive messages, resulting in the faster detection of a failover condition. The timer setting must be configured identically for the ESRP domain across all participating switches. To see the hello settings, use the `show esrp { {name} | {type [vpls-redundancy | standard]} }` command.

The seconds range is 1 to 255.

If your configuration contains more than 2,000 ESRP VLANs and 256,000 FDB entries, we recommend a timer setting greater than 3 seconds.

To view the hello timer settings, use the `show esrp { {name} | {type [vpls-redundancy | standard]} }` command.

In a large ESRP configuration, the slave ESRP domain might inadvertently become the master ESRP domain. This can occur when FDB entries are flushed during a master-slave transition. To avoid this we recommend the general neighbor and hello timeout guidelines listed in [Table 36: General Neighbor and Hello Timeout](#) on page 2040, which is described in the description for the `configure esrp timer neighbor` command.

Example

The following command configures the ESRP hello timer to 4 seconds for the ESRP domain esrp1:

```
configure esrp esrp1 timer hello 4
```

History

This command was first available in ExtremeXOS 11.0.



Platform Availability

This command is available on all platforms.

configure esrp timer neighbor

```
configure esrp esrpDomain timer neighbor seconds
```

Description

Configures the ESRP neighbor timeout value.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>seconds</i>	Specifies the number of seconds after which an ESRP neighbor times out. The range is 6 to 1024 seconds.

Default

The default neighbor timeout is 8 seconds (four times the hello timer).

Usage Guidelines

The neighbor timeout specifies the amount of time that ESRP waits before considering the neighbor down. The neighbor value must be at least 3 times the hello timer value. Entering a value outside of that range generates an error message similar to the following:

```
operation Failed. Valid timer relationship "neighbor timeout >= 3*hello ;
neutral timeout >= 2*hello ; premaster timeout >= 3*hello"!
```

The seconds range is 3*hello to 1024 seconds.

To view the neighbor timer settings, use the `show esrp { {name} | {type [vpls-redundancy | standard]} }` command.

In a large ESRP configuration, the slave ESRP domain might inadvertently become the master ESRP domain. This can occur when FDB entries are flushed during a master-slave transition. To avoid this we recommend the general neighbor and hello timeout guidelines listed in following table.

Table 36: General Neighbor and Hello Timeout

Number of Domains	Number of VLANs	Suggested Neighbor and Hello Timeout
64 or less	1000	Use the default timer values
64	1000 to 3000	hello >=3, neighbor >=9
128	3000	hello >=4, neighbor >=12



Example

The following command configures the ESRP neighbor timeout to 14 seconds for the ESRP domain `esrp1`:

```
configure esrp esrp1 timer neighbor 14
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp timer neutral

Configures the ESRP neutral timeout value.

```
configure esrp esrpDomain timer neutral seconds
```

Description

Configures the ESRP neutral timeout value.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>seconds</i>	Specifies the number of seconds after which an ESRP domain. The range is 4 to 1024 seconds.

Default

The default neutral timeout is 4 seconds (two times the hello timer).

Usage Guidelines

After you create, configure, and enable the ESRP domain, it enters the neutral state. The neutral timeout specifies the amount of time the ESRP domain stays in this temporary state before entering the slave state. The neutral value must be at least 2 times the hello timer value. Entering a value outside of that range generates an error message similar to the following:

```
operation Failed. Valid timer relationship "neighbor timeout >= 3*hello ;
neutral timeout >= 2*hello ; premaster timeout >= 3*hello"!
```

The seconds range is 2*hello to 1024.



To view the neutral timer settings, use the `show esrp { {name} | {type [vpls-redundancy | standard]} }` command.

Example

The following command configures the ESRP neutral timeout to 8 seconds for the ESRP domain esrp1:

```
configure esrp esrp1 timer neutral 8
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp timer premaster

```
configure esrp esrpDomain timer premaster seconds
```

Description

Configures the ESRP pre-master timeout value.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>seconds</i>	Specifies the maximum length of time, in seconds, that the transitioning master VLAN remains in the pre-master state. The range is 6 to 1024.

Default

The default timeout is 6 seconds (three times the hello timer).

Usage Guidelines

The premaster timer specifies how long the ESRP domain stays in the pre-master state. The pre-master timer expires if the neighbor agrees to be the slave. The premaster value must be at least three times the hello timer value. Entering a value outside of that range generates an error message similar to the following:

```
operation Failed. Valid timer relationship "neighbor timeout >= 3*hello ;
neutral timeout >= 2*hello ; premaster timeout >= 3*hello"!
```

The seconds range is 3*hello-1024.



To view the pre-master timer settings, use the `show esrp { {name} | {type [vpls-redundancy | standard]} } command`.



Caution

Configure the pre-master state timeout only with guidance from Extreme Networks personnel. Misconfiguration can severely degrade the performance of ESRP and your switch.

Example

The following command configures the pre-master timeout to 10 seconds for the ESRP domain `esrp1`:

```
configure esrp esrp-1 timer premaster 10
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure esrp timer restart

```
configure esrp esrpDomain timer restart seconds
```

Description

Configures the ESRP restart timer value.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>seconds</i>	Specifies the maximum length of time, in seconds, that the neighbor ESRP switch remains in its current state during an MSM hitless failover. The range is 2 to 1024.

Default

The default restart timer value is 2 seconds.

Usage Guidelines

The restart timer specifies the amount of time that the neighbor ESRP switch remains in its current state during a hitless failover. This timer prevent the slave ESRP switch from trying to become master during a hitless failover.



The seconds range is 2-1024.

To view the restart settings, use the `show esrp { {name} | {type [vpls-redundancy | standard]} } command`.

Example

The following command configures the restart timer value to 40 seconds for the ESRP domain esrp1:

```
configure esrp esrp-1 timer restart 40
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

create esrp

```
create esrp esrp_domain {type [vpls-redundancy | standard]}
```

Description

Creates an ESRP domain with the specified name on the switch.

Syntax Description

<i>esrp_domain</i>	Specifies the name of an ESRP domain to be created. Can be up to 32 characters in length.
--------------------	---

Default

The ESRP domain is disabled and in the “Aware” state.

When you create an ESRP domain, it has the following default parameters:

- Operational version—Extended
- Priority—0
- VLAN interface—none
- VLAN tag—0
- Hello timer—2 seconds
- Neighbor timer—8 seconds
- Premaster timer—6 seconds
- Neutral timer—4 seconds



- Neighbor restart timer—30 seconds
- VLAN tracking—none
- Ping tracking—none
- IP route tracking—none

Usage Guidelines

The `type` keyword specifies the type of ESRP domain when a new ESRP domain is created. The only types supported are `vpls-redundancy` and `standard`. Not specifying the optional ESRP domain type results in the creation of an ESRP domain of type `standard`. The `standard` ESRP domain is equivalent to the legacy ESRP domain type that was implicitly created. The `vpls-redundancy` domain type is only specified when redundant access to an MPLS VPLS network is desired.

An ESRP domain name must begin with an alphabetical character and may contain alphanumeric characters and underscores (`_`), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For ESRP domain name guidelines and a list of reserved names, see [Object Names](#) in the *ExtremeXOS Concepts Guide*.

Each ESRP domain name must be unique and cannot duplicate any other named ESRP domains on the switch. If you are uncertain about the ESRP names on the switch, use the `show esrp` command to view the ESRP domain names.

You can create a maximum of 128 ESRP domains.

Configuring ESRP-Aware Switches

For an Extreme Networks switch to be ESRP-aware, you must create an ESRP domain on the aware switch, add a master VLAN to that ESRP domain, add a member VLAN to that ESRP domain if configured, and configure a domain ID if necessary.

For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

Example

The following command creates ESRP domain `esrp1` on the switch:

```
create esrp esrp1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.



delete esrp

```
delete esrp esrpDomain
```

Description

Deletes the ESRP domain with the specified name.

Syntax Description

<i>esrpDomain</i>	Specifies the name of an ESRP domain to be deleted.
-------------------	---

Default

N/A.

Usage Guidelines

You must first disable an ESRP domain before you delete it. To disable an ESRP domain, use the `disable esrp` command.

You do not have to remove the master or member VLANs from an ESRP domain before you delete it. When you delete an ESRP domain, All VLANs are automatically removed from the domain.

For ESRP domains configured of type VPLS-redundancy, you need to unconfigure all associated VPLS instances from the ESRP domain using the `unconfigure vpls redundancy` command before deleting the domain.

Example

The following command deletes ESRP domain `esrp1` from the switch:

```
delete esrp esrp1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

disable esrp

```
disable esrp {esrpDomain}
```



Description

Disables ESRP for a named domain or for the entire switch.

Syntax Description

<i>esrpDomain</i>	Specifies the name of an ESRP domain.
-------------------	---------------------------------------

Default

Disabled for the entire switch.

Usage Guidelines

If you do not specify a domain name, ESRP is disabled for the entire switch.

If you disable an ESRP domain, the domain enters the Aware state, the switch notifies its neighbor that the ESRP domain is going down, and the neighbor clears its neighbor table. If the master switch receives this information, it enters the neutral state to prevent a network loop. If the slave switch receives this information, it enters the neutral state.

Example

The following command disables ESRP for the entire switch:

```
disable esrp
```

The following command disables ESRP for the domain esrp1:

```
disable esrp esrp1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms that use the Edge, Advanced Edge, or Core license. For information on the licenses available for each platform, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable esrp

```
enable esrp esrpDomain
```



Description

Enables ESRP for a named domain.

Syntax Description

<code>esrpDomain</code>	Specifies the name of an ESRP domain.
-------------------------	---------------------------------------

Default

Disabled.

Usage Guidelines

Before you enable an ESRP domain, it must have a domain ID. The ESRP domain ID is determined from one of the following user-configured parameters:

- ESRP domain number created with the `configure esrp domain-id` command
- 802.1Q tag (VLANid) of the tagged master VLAN

If you do not have a domain ID, you cannot enable ESRP on that domain. A message similar to the following appears:

```
ERROR: Cannot enable ESRP Domain "esrp1" ; No domain id configured!
```

If you add an untagged master VLAN to the ESRP domain, make sure to create an ESRP domain ID with the `configure esrp domain-id` command before you attempt to enable the domain.

Example

The following command enables ESRP for the domain esrp1:

```
enable esrp esrp1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms that use the Edge, Advanced Edge, or Core license. For information on the licenses available for each platform, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show esrp

```
show esrp { {name} | {type [vpls-redundancy | standard]} }
```



Description

Displays ESRP configuration information for one or all ESRP domains on the switch.

Syntax Description

name	Specifies an ESRP domain name.
type	Specifies whether ESRP is standard or redundant VPLS.
vpls-redundancy	Specifies redundant VPLS.
standard	Specifies standard ESRP.

Default

Shows summary ESRP information.

Usage Guidelines

This command shows information about the state of an ESRP domain and its neighbors. This includes information about tracked devices.

In addition to ESRP information, ESRP status information is also displayed. This includes information about the master and pre-master states, number of transitions to the pre-master state, and the ports where ESRP is disabled.

The output varies depending upon the configuration and the state of the switch.

Example

The following command displays summary ESRP status information for the ESRP domains on the switch:

```
show esrp
```

The following is sample output from this command:

```
ESRP:                               Enabled
Configured Version:                 Extended
# ESRP domain configuration :
-----
-- Domain Grp Ver VLAN   VID  DId   IP/IPX           State  Master MAC Address
Nbr
-----
--
ed2    0  E v2    2    2    2.2.2.3         Master  00:01:30:f9:9e:90  0
ed2    5  E v2    2    2    2.2.2.3         Aware   00:01:00:0D:9e:8a  0
ed3    0  E v3    3    3    0.0.0.0         Aware   00:01:00:0C:F0:D1  0
ed4    0  E v4    4    4    0.0.0.0         Slave  00:00:00:00:00:00  0
-----
--
# ESRP Port configuration:
```



```

-----
--
Port          Weight          Host      Restart
-----
--
6:1           0                 H
6:2           10
6:3           0                 R

```

The following command displays detailed ESRP status information for the specified ESRP domain on the switch (the election policy displayed is the default policy in extended mode):

```
show esrp ed2
```

The following is sample output from this command:

```

show esrp ed2
Domain:                ed2
Group:                 0
Operational Version:  extended
Vlan Interface:       v2
Vlan Tag:              2
Domain Id:             2
Rtif. Admin Status:   DOWN
Rtif. Virtual Mac :   00:e0:2b:00:00:80
IP Address:            2.2.2.3
Election Policy:
standby > sticky > ports > weight > track > priority > mac
-----
--
                                This System          Neighbor system
-----
--
State:                        Slave
Active Ports:                  1
Tracked Active Vlan Ports:    0
Tracked IP Routes & Pings:    0
Priority:                       255[255]
Sticky Flag:                   0
MAC:                           00:01:30:f9:9e:90
Active Ports Weight:           10
Sequence Number:               8
Hand Shake Flag:               0
Restart Flag:                   0
Timer Configuration:           Hello          2s(0)      Neighbor    6s(0)
PreMaster 4s(0) Neutral 4s(0)
NbrRestart 30s(0)
State Transition Counters: To Master 0 To Neutral 1
To PreMaster 0 To Aware 1
To Slave 1
Last State Change: Mon Apr  5 10:43:34 2004
ELRP in PreMaster:             Enabled (Interval: 1, Count: 3)
ELRP in Master:                 Enabled (Interval: 1)
Tracked Vlans: t_vlan
Tracked Pings: 40.0.1.2 / 3 second intervals / 5 misses allowed
Tracked Ip Routes: 30.4.2.16/255.255.255.0
Tracked Environment: System Power : OK

```




```

Temperature      : WARNING
-----
# Domain Member VLANs:
VLAN Name        VID      Virtual IP/IPX  State
-----
m_vlan1          1001    0.0.0.0         DOWN
m_vlan2          1002    0.0.0.0         DOWN
m_vlan3          1003    0.0.0.0         DOWN

```

History

This command was first available in ExtremeXOS 11.0.

ELRP status information was added in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

show esrp aware

```
show esrp domain aware {selective-forward-ports | statistics}
```

Description

Displays all selective forwarding information for a given ESRP-aware domain.

Syntax Description

<i>domain</i>	Specifies the name of an ESRP domain.
selective-forward-ports	Specifies that the selective-forward-port table is the only table displayed.
statistics	Specifies that the selective-forward-port statistics table is the only table displayed.

Default

Disabled.

Usage Guidelines

An ESRP-aware switch floods ESRP PDUs from all ports in an ESRP-aware VLAN. This flooding creates unnecessary network traffic because some ports forward ESRP PDUs to switches that are not running the same ESRP groups. You can select the ports that are appropriate for forwarding ESRP PDUs by configuring selective forwarding on an ESRP-aware VLAN and thus reduce this excess traffic. Configuring selective forwarding creates a port list of only those ports that forward to the ESRP groups that are associated with an ESRP-aware VLAN. This aware port list is then used for forwarding ESRP PDUs.



The first of the two tables that this command displays shows Selective Forward Ports information:

- Group—The number of an ESRP group within the given domain.
- Port Count—The number of ports in the group that are selected to forward PDUs on the master VLAN.
- Selective Forward Ports—The list of ports in the group that are selected to forward PDUs on the master VLAN.

The second of the two tables displays statistical information about the activity of the ports:

- Group—The number of an ESRP group within the given domain.
- Master MAC—The MAC address for the master of the group.
- Rx—The number of PDUs received matching the domain/group pair.
- Fwd—The number of PDUs received and forwarded matching the domain/group pair.
- FDB Flush—The number of FDB Flush requests received from the ESRP Master for this domain/group pair.
- Fwd Ports—Selective or Default.

Selective describes the group as having a configured aware port list for selective forwarding of PDUs on the Master VLAN. The list of ports is displayed in the first table above.

Default describes the group as one where all the ports on the master VLAN forward the ESRP PDUs that are received for the domain/group pair. Because there is no selective forwarding configuration for this group, there is no entry in the first table.

Example

The following command displays the ESRP aware information for the domain d1.

```
show esrp d1 aware
Domain:          d1
Vlan:           vesrp1
-----
```

Group	Port Count	Selective Forward Ports
0	5	5:1, 5:2, 7:31, 7:32: 8:1
3	2	5:1, 8:1

```
-----
```

Group	Master MAC	Rx
Fwd	FDB Flush	Fwd Ports
0	00:12:00:33:44:55	10
10	1	selective
1	00:22:00:12:21:1F	77
77		
3		default
3	00:02:00:13:11:11	99
99	3	selective

History

This command was first available in Extreme XOS 12.0.



Platform Availability

This command is available on all platforms.

show esrp counters

```
show esrp {name} counters
```

Description

Displays ESRP counter information for ESRP or for a specified ESRP domain.

Syntax Description

<i>name</i>	Specifies an ESRP domain name.
-------------	--------------------------------

Default

Displays summary ESRP counter information.

Usage Guidelines

The `show esrp counters` command displays information about the number of:

- Failed received protocol packets
- Failed sent protocol packets
- Dropped protocol packets belonging to unknown ESRP domains
- Dropped protocol packets due to invalid Extreme Encapsulation Protocol (EEP) data
- Dropped packets due to old sequence numbers
- Dropped packets due to an invalid 802.1Q tag
- Dropped packets because the packet length was truncated (packet length is less than expected)
- Dropped packets due to failed checksum verification

The `show esrp {name} counters` command displays information about the number of times ESRP, ESRP-aware, and ESRP error packets were transmitted and received.

Example

The following command displays ESRP counter information:

```
show esrp counters
```

The following is sample output from this command:

```
Current-time:                Sun Nov 16 00:25:08 2003
esrpStatsRxHelloFailed      = 0
esrpStatsTxHelloFailed      = 0
```



```

esrpStatsUnknownDomain          = 0
esrpStatsUnsupportedEEPVersion  = 0
esrpStatsInvalidEEPLength      = 0
esrpStatsNotInTimeWindow       = 0
esrpStatsInvalid8021Qtag       = 0
esrpStatsInvalidSNAPType       = 0
esrpStatsUndersizePkt          = 0
esrpStatsInvalidChecksum       = 0
esrpStatsWrongDigest           = 0

```

The following command displays counter information for ESRP domain ed5:

```
show esrp ed5 counters
```

The following is sample output from this command:

```

Domain:  ed5                Current-time:  Sun Nov 16 00:25:27 2003
Rx-Esrp-Pkts          = 628      Tx-Esrp-Pkts          = 630
Rx-Aware-Esrp-Pkts   = 112      Tx-Aware-Esrp-Pkts   = 34
Rx-Err-Pkts          = 0      Tx-Err-Pkts          = 0

```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.



32 VRRP Commands

```
clear counters vrrp
configure vrrp vlan vrid accept-mode
configure vrrp vlan vrid add ipaddress
configure vrrp vlan vrid add track-iproute
configure vrrp vlan vrid add track-ping
configure vrrp vlan vrid add virtual-link-local
configure vrrp vlan vrid add track-vlan
configure vrrp vlan vrid advertisement-interval
configure vrrp vlan vrid authentication
configure vrrp vlan vrid delete
configure vrrp vlan vrid delete track-iproute
configure vrrp vlan vrid delete track-ping
configure vrrp vlan vrid delete track-vlan
configure vrrp vlan vrid dont-preempt
configure vrrp vlan vrid preempt
configure vrrp vlan vrid priority
configure vrrp vlan vrid track-mode
configure vrrp vlan vrid version
create vrrp vlan vrid
delete vrrp vlan vrid
disable vrrp vrid
enable vrrp vrid
show vrrp
show vrrp vlan
```

This chapter describes commands for:

- Enabling and disabling Virtual Router Redundancy Protocol (VRRP).
- Performing basic VRRP configuration.
- Displaying VRRP information.

For an introduction to VRRP, see the *ExtremeXOS Concepts Guide*.

clear counters vrrp

```
clear counters vrrp {{vlan vlan_name} {vrid vridval}}
```

Description

Clears, resets all VRRP statistics and counters.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN.
<i>vridval</i>	Specifies the VRRP Router ID (VRID) for a VRRP instance. To display the configured VRRP router instances, enter the <code>show vrrp</code> command.

Default

N/A.

Usage Guidelines

Use this command to reset the VRRP statistics on the switch. Statistics are not reset when you disable and re-enable VRRP.

If you do not enter a parameter, statistics for all VRRP VLANs are cleared.

If you specify only VLAN name, statistics for all VRRP VRIDs on that VLAN are cleared.

If you specify VLAN name and VRRP VRID, only statistics for that particular VRID are cleared.

Example

The following command clears the VRRP statistics on VRRP VLAN v1:

```
clear counters vrrp vlan v1
```

The following command clears the VRRP statistics for VRID 1 on VRRP VLAN v1:

```
clear counters vrrp vlan v1 vrid 1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure vrrp vlan vrid accept-mode


```
configure vrrp vlan vlan_name vrid vridval accept-mode [on | off]
```

Description

Configures a backup VRRP router instance to accept or reject packets addressed to the IP address owner when operating as the VRRP master.

Additionally, this command provides capability for switches to configure the VRRP virtual IP as NTP server address.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN.
<i>vridval</i>	Specifies the VRID of a VRRP instance. To display the configured VRRP router instances, enter the <code>show vrrp</code> command.
on	Specifies that the VRRP instance is to accept packets addressed to the IP address owner.
off	Specifies that the VRRP instance not accept packets addressed to the IP address owner.
 Note Ping packets are accepted, regardless of the configuration for this command.	

Default

Off.

Usage Guidelines

When a backup VRRP router operates as master, it accepts VRRP traffic and routes traffic. The backup router in master mode also accepts ping packets and IPv6 neighbor solicitations and advertisements. However, because the backup router is not the IP address owner, the default configuration rejects all other traffic addressed to the IP address owner.

If your network requires that a backup VRRP router in master mode accept all traffic addressed to the IP address owner, use this command to configure accept-mode on.

In the ExtremeXOS 15.3 release, NTP VRRP Virtual IP support is added. This feature allows you to configure the VRRP virtual IP as NTP server address. The NTP server when configured on the VRRP master will listen on the physical and virtual IP address for NTP clients. For this feature to work correctly, you need to enable accept mode in VRRP. Enabling accept mode allows the switch to process non-ping packets that have a destination IP set to the virtual IP address.



Example

The following example configures a backup VRRP router in master mode to accept packets addressed to the IP address owner:

```
configure vrrp vlan vlan-1 vrid 1 accept-mode on
```

History

This command was first available in ExtremeXOS 12.7.

NTP VRRP Virtual IP support was added in ExtremXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

For NTP VRRP Virtual IP support, Summit switches configured as NTP clients need to have the following bootrom version:

- X480, X460, X440, X670 - 2.0.1.7
- NWI - 1.0.5.7

configure vrrp vlan vrid add ipaddress

```
configure vrrp vlan vlan_name vrid vridval add ipaddress
```

Description

Associates a virtual IP address with a specific VRRP instance.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN.
<i>vridval</i>	Specifies the VRID of a VRRP instance. To display the configured VRRP router instances, enter the <code>show vrrp</code> command.
<i>ipaddress</i>	Specifies a virtual IPv4 or IPv6 address to be assigned to the VRRP instance.

Default

N/A.



Usage Guidelines

Each VRRP instance is identified by an ID number, VLAN name, and virtual IP address. When two or more routers are configured with the same VRRP ID number, VLAN name, and virtual IP address, the routers with matching parameters are all part of the same VRRP instance. One router within the instance will become the VRRP instance master, and the others will become backup routers for the VRRP instance.

Most routers within a VRRP instance will have a virtual IP address that is different from the actual IP addresses configured on the router. If the virtual IP address for a VRRP instance matches an IP address configured on a host router, the VRRP instance is known as the IP address owner. On the IP address owner, the VRRP instance priority defaults to 255, and by default, the IP address owner becomes the VRRP master when VRRP is enabled.



Note

There is no requirement to configure an IP address owner within a VRRP instance.

Before each VRRP router is enabled, it must be configured with at least one virtual IPv4 or IPv6 address. You can repeat this command to add additional virtual IP addresses to the VRRP router. If a virtual IPv4 address is added to a VRRP router, you cannot later add a virtual IPv6 address. Similarly, if a virtual IPv6 address is added to a VRRP router, you cannot later add a virtual IPv4 address.

Each IPv6 VRRP instance is associated with one and only one virtual link local address, which serves as the source IP address for subsequent router announcement packets generated by the master VRRP router. The virtual link local address can be explicitly configured or generated automatically. One way to explicitly configure the virtual link local address is to add it to the virtual IP address list with this command.

Example

The following example associates virtual IPv4 address 10.1.2.3 to VRRP router instance 1:

```
configure vrrp vlan vlan-1 vrid 1 add 10.1.2.3
```

The following example associates virtual IPv6 address 2001:db8::3452/128 to VRRP router instance 2:

```
configure vrrp vlan vlan-1 vrid 2 add 2001:db8::3452/128
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 addresses was added in ExtremeXOS 12.7.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vrrp vlan vrid add track-iproute

```
configure vrrp vlan vlan_name vrid vridval add track-iproute ipaddress/masklength
```

Description

Creates a tracking entry for the specified route. When this route becomes unreachable, this entry is considered to be failing.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN.
<i>vridval</i>	Specifies the VRID of the target VRRP instance. To display the configured VRRP router instances, enter the <code>show vrrp</code> command.
<i>ipaddress</i>	Specifies the IPv4 or IPv6 prefix of the route to track.
<i>masklength</i>	Specifies the length of the route's prefix.

Default

N/A.

Usage Guidelines

The route specified in this command might not exist in the IP routing table. When you create the entry for a route, an immediate VRRP failover might occur.



Note

VRRP tracking is not supported on MPLS LSPs.

Example

The following command enables IP route failure tracking for routes to the specified subnet:

```
configure vrrp vlan vlan-1 vrid 1 add track-iproute 3.1.0.0/24
```

History

This command was first available in ExtremeXOS 10.1.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vrrp vlan vrid add track-ping

```
configure vrrp vlan vlan_name vrid vridval add track-ping ipaddress frequency
seconds miss misses
```

Description

Creates a tracking entry for the specified IP address. The entry is tracked using pings to the IP address, sent at the specified frequency.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN.
<i>vridval</i>	Specifies the VRID of the target VRRP instance. To display the configured VRRP router instances, enter the <code>show vrrp</code> command.
<i>ipaddress</i>	Specifies the IPv4 or IPv6 address to be tracked.
<i>seconds</i>	Specifies the number of seconds between pings to the target IP address. The range is 1 to 600 seconds.
<i>misses</i>	Specifies the number of misses allowed before this entry is considered to be failing. The range is 1 to 255 pings.

Default

N/A.

Usage Guidelines

Adding an entry with the same IP address as an existing entry causes the new values to overwrite the existing entry's frequency and miss number.

Example

The following command enables ping tracking for the external gateway at 3.1.0.1, pinging every 3 seconds, and considering the gateway to be unreachable if no response is received to 5 consecutive pings:

```
configure vrrp vlan vlan-1 vrid 1 add track-ping 3.1.0.1 frequency 3 miss 5
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vrrp vlan vrid add virtual-link-local

```
configure vrrp vlan vlan_name vrid vridval add virtual-link-local vll_addr
```

Description

Specifies a virtual IPv6 link local address for the VRRP router instance.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN.
<i>vridval</i>	Specifies the VRID of a VRRP instance. To display the configured VRRP router instances, enter the <code>show vrrp</code> command.
<i>vll_addr</i>	Specifies a virtual link local address to be assigned to the VRRP instance.

Usage Guidelines

Each IPv6 VRRP instance is associated with one and only one virtual link local address, which serves as the source IP address for subsequent router announcement packets generated by the master VRRP router. The virtual link local address can be explicitly configured or generated automatically.

One way to explicitly configure the virtual link local address is to add it to the virtual IP address list with this command. The new link local address must match the FE80::/64 subnet, and it must match the address in use on all other router in this VRRP instance.

If no virtual link local address is configured, an appropriate address is generated automatically.



Note

If an IPv4 address has been added to a VRRP router, you cannot later add any IPv6 address, so you cannot add a link local address.

Example

The following example associates virtual IPv6 link local address fe80::1111 to VLAN vlan-1:

```
configure vrrp vlan vlan-1 vrid 1 add virtual-link-local fe80::1111
```



History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vrrp vlan vrid add track-vlan

```
configure vrrp vlan vlan_name vrid vridval add track-vlan target_vlan_name
```

Description

Configures a VRRP VLAN to track port connectivity to a specified VLAN. When this VLAN is in the down state, this entry is considered to be failing.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN.
<i>vridval</i>	Specifies the VRID of the target VRRP instance. To display the configured VRRP router instances, enter the <code>show vrrp</code> command.
<i>target_vlan_name</i>	Specifies the name of the VLAN to track.

Default

N/A.

Usage Guidelines

Up to eight VLANs can be tracked.

Deleting a tracked VLAN does not constitute a failover event for the VRRP VLAN tracking it, and the tracking entry is deleted.

Example

The following command enables VRRP VLAN `vlan-1` to track port connectivity to VLAN `vlan-2`:

```
configure vrrp vlan vlan-1 vrid 1 add track-vlan vlan-2
```

History

This command was first available in ExtremeXOS 10.1.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vrrp vlan vrid advertisement-interval

```
configure vrrp vlan vlan_name vrid vridval advertisement-interval interval
[{seconds} | centiseconds]
```

Description

Configures the time between VRRP advertisements in seconds or centiseconds.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN.
<i>vridval</i>	Specifies the VRID of a VRRP instance. To display the configured VRRP router instances, enter the <code>show vrrp</code> command.
<i>interval</i>	Specifies an interval value for the time between advertisements. The range is 1 through 40 seconds or 10 through 4095 centiseconds.
seconds	Specifies that the interval value is in seconds. If you do not specify seconds or centiseconds, the interval value is applied as seconds.
centiseconds	Specifies that the interval value is in centiseconds.

Default

The advertisement interval is 1 second.

Usage Guidelines

The advertisement interval specifies the interval between advertisements sent by the master router to inform the backup routers that its alive. You must use whole integers when configuring the advertisement interval.

An extremely busy CPU can create a short dual master situation. To avoid this, increase the advertisement interval.

Note



The milliseconds keyword is replaced by the centiseconds keyword, but the milliseconds keyword is still recognized to support existing configurations and scripts. Any values specified in milliseconds are converted to centiseconds. All new configurations and scripts should specify the interval in either seconds or centiseconds. The maximum value for an interval specified in seconds is 40. However, the software supports older configurations and scripts that specify values up to 255, which were supported prior to ExtremeXOS Release 12.7.



To view your VRRP configuration, including the configured advertisement interval, use one of the following commands:

- `show vrrp {virtual-router {vr-name}} {detail}`
- `show vrrp vlan vlan_name {stats}`

If you enter a number that is out of the seconds or centiseconds range, the switch displays an error message. For example, if the interval value is set to 999 and the centiseconds keyword is missing, the switch displays an error message similar to the following:

```
configure vrrp blue vrid 250 advertisement-interval 999 Error:
Advertisement interval must be between 1 and 255 seconds. 999 out of
range
```

Example

The following command configures the advertisement interval for 15 seconds:

```
configure vrrp vlan vrrp-1 vrid 1 advertisement-interval 15
```

The following command configures the advertisement interval for 200 centiseconds:

```
configure vrrp vlan vrrp-1 vrid 1 advertisement-interval 200 centiseconds
```

History

This command was first available in ExtremeXOS 10.1.

The milliseconds and seconds keywords were added in ExtremeXOS 11.5.

The centiseconds keyword replaced the milliseconds keyword, and the maximum value for intervals specified in seconds was reduced to 40 in ExtremeXOS 12.7.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vrrp vlan vrid authentication

```
configure vrrp vlan vlan_name vrid vridval authentication [none | simplepassword
password]
```

Description

Configures the authentication type for a specific VRRP router.



Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN.
<i>vridval</i>	Specifies the VRID of a VRRP instance. To display the configured VRRP router instances, enter the <code>show vrrp</code> command.
<i>password</i>	Specifies the user-defined password for authentication.

Default

Authentication is set to none.

Usage Guidelines

Note



This command applies only to VRRP routers configured only for VRRPv2 (`configure vrrp vlan vrid version` command, v2 option), and is only supported for backward compatibility. If you try to enter this command in the combined VRRPv2 and VRRPv3 mode or VRRPv3 mode, an error message appears.

This command can add a modest amount of security to VRRP advertisements. All VRRP routers using the same VRID must use the same password when using this feature.

A simple password must be between 1 and 8 characters long.

Example

The following command configures authentication for VRRP VLAN vrrp-1 with the password newvrrp:

```
configure vrrp vlan vrrp-1 vrid 1 authentication simple-password newvrrp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vrrp vlan vrid delete

```
configure vrrp vlan vlan_name vrid vridval delete ipaddress
```



Description

Deletes a virtual IPv4 or IPv6 address from a specific VRRP router.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN.
<i>vridval</i>	Specifies the VRID of the VRRP instance. To display the configured VRRP router instances, enter the <code>show vrrp</code> command.
<i>ipaddress</i>	Specifies the virtual IP address to be deleted from the VRRP instance. This is common for IPv4/IPv6 addresses.

Usage Guidelines

When a VRRP router is enabled, it must have at least one virtual IP address. When the VRRP router is not enabled, there are no restrictions on deleting the IP address.

Example

The following command removes IP address 10.1.2.3 from VLAN vlan-1:

```
configure vrrp vlan vlan-1 vrid 1 delete 10.1.2.3
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vrrp vlan vrid delete track-iproute

```
configure vrrp vlan vlan_name vrid vridval delete track-iproute ipaddress/  
masklength
```

Description

Deletes a tracking entry for the specified route.



Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN.
<i>vridval</i>	Specifies the VRID of the target VRRP instance. To display the configured VRRP router instances, enter the <code>show vrrp</code> command.
<i>ipaddress</i>	Specifies the IPv4 or IPv6 prefix of the route.
<i>masklength</i>	Specifies the length of the route's prefix.

Default

N/A.

Usage Guidelines

Deleting a tracking entry while VRRP is enabled causes the VRRP VRs state to be re-evaluated for failover.

Example

The following command disables tracking of routes to the specified subnet for VLAN vlan-1:

```
configure vrrp vlan vlan-1 vrid 1 delete track-iproute 3.1.0.0/24
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vrrp vlan vrid delete track-ping

```
configure vrrp vlan vlan_name vrid vridval delete track-ping ipaddress
```

Description

Deletes a tracking entry for the specified IP address.



Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN.
<i>vridval</i>	Specifies the VRID of the target VRRP instance. To display the configured VRRP router instances, enter the <code>show vrrp</code> command.
<i>ipaddress</i>	Specifies the IP address to be tracked.

Default

N/A.

Usage Guidelines

Deleting a tracking entry while VRRP is enabled causes the VRRP VRs state to be re-evaluated for failover.

A VRRP node with a priority of 255 might not recover from a ping-tracking failure if there is a Layer 2 switch between it and another VRRP node. In cases where a Layer 2 switch is used to connect VRRP nodes, we recommend that those nodes have priorities of less than 255.

Example

The following command disables ping tracking for the external gateway at 3.1.0.1:

```
configure vrrp vlan vlan-1 vrid 1 delete track-ping 3.1.0.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vrrp vlan vrid delete track-vlan

```
configure vrrp vlan vlan_name vrid vridval delete track-vlan target_vlan_name
```

Description

Deletes the tracking of port connectivity to a specified VLAN.



Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN.
<i>vridval</i>	Specifies the VRID of the VRRP instance. To display the configured VRRP router instances, enter the <code>show vrrp</code> command.
<i>target_vlan_name</i>	Specifies the name of the tracked VLAN.

Default

N/A.

Usage Guidelines

Deleting a tracking entry while VRRP is enabled causes the VRRP VRs state to be re-evaluated for failover.

Example

The following command disables the tracking of port connectivity to VLAN vlan-2:

```
configure vrrp vlan vlan-1 vrid 1 delete track-vlan vlan-2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vrrp vlan vrid dont-preempt

```
configure vrrp vlan vlan_name vrid vridval dont-preempt
```

Description

Specifies that a higher priority backup router does not preempt a lower priority master.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN.
<i>vridval</i>	Specifies the VRID of a VRRP instance. To display the configured VRRP router instances, enter the <code>show vrrp</code> command.



Default

The default setting is preempt.

Usage Guidelines

The preempt mode controls whether a higher priority backup router preempts a lower priority master. dont-preempt prohibits preemption. The router that owns the virtual IP address always preempts, independent of the setting of this parameter.

Example

The following command disallows preemption:

```
configure vrrp vlan vlan-1 vrid 1 dont-preempt
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vrrp vlan vrid preempt

```
configure vrrp vlan vlan_name vrid vridval preempt {delay seconds}
```

Description

Specifies that a higher priority backup router preempts a lower priority master.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN.
<i>vridval</i>	Specifies the VRID for a VRRP instance. To display the configured VRRP router instances, enter the <code>show vrrp</code> command.
<i>seconds</i>	Specifies a preempt delay period in seconds. The value range is 1 to 3600 seconds, or 0, which selects the original preempt delay period.

Default

Preempt enabled.



Delay configuration: 0.

Usage Guidelines

The preempt option enables a higher-priority backup router to preempt a master with a lower priority. When a VRRP enabled router receives a lower priority VRRP advertisement and preemption is enabled, the higher-priority VRRP enabled router takes over as master. The new master starts sending VRRP advertisements and the old, lower-priority master relinquishes mastership.



Note

The router that owns the virtual IP address always preempts, independent of the setting of this parameter.

When a VRRP enabled router preempts the master, it does so in one of the following ways:

- If the preempt delay timer is configured for between 1 and 3600 seconds and the lower-priority master is still operating, the router preempts the master when the timer expires.
- If the preempt delay timer is configured for 0, the router preempts the master after 3 times the hello interval.
- If the higher priority router stops receiving advertisements from the current master for 3 times the hello interval, it takes over mastership immediately.



Note

The preempt feature can be disabled with the `configure vrrp vlan vrid dont-preempt` command.

Example

The following command allows preemption:

```
configure vrrp vlan vlan-1 vrid 1 preempt
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vrrp vlan vrid priority

```
configure vrrp vlan vlan_name vrid vridval priority priorityval
```



Description

Configures the priority value of a VRRP router instance.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN.
<i>vridval</i>	Specifies the VRID for the VRRP instance. To display the configured VRRP router instances, enter the <code>show vrrp</code> command.
<i>priorityval</i>	Specifies the priority value of the router. The default is 100. The priority range is 1-255.

Default

The default priority is 100.

Usage Guidelines

This command changes the priority of a VRRP router. If the VRRP router is the IP address owner (which means that the VRRP router IP address matches the VRRP VLAN IP address), the priority is 255 and cannot be changed. If the VRRP router is not the IP address owner, the priority can be changed to values in the range of 1 to 254.

To change the priority of the IP address owner or to make a different VRRP router the IP address owner, disable VRRP and reconfigure the affected switches to use VRRP router addresses that support the priorities you want to assign.

Example

The following command configures a priority of 150 for VLAN vrrp-1:

```
configure vrrp vlan vrrp-1 vrid 1 priority 150
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vrrp vlan vrid track-mode

```
configure vrrp vlan vlan_name vrid vridval track-mode [all | any]
```



Description

Defines the conditions under which the router automatically relinquishes master status when the tracked entities fail.

Syntax Description

<code>vlan_name</code>	Specifies the name of a VRRP VLAN.
<code>vridval</code>	Specifies the VRID for the VRRP instance. To display the configured VRRP router instances, enter the <code>show vrrp</code> command.
all	Specifies that the mastership is relinquished when one of the following events occur: All of the tracked VLANs fail All of the tracked routes fail All of the tracked PINGs fail
any	Specifies that the mastership is relinquished when any of the tracked VLANs, routes, or PINGs fail.

Default

The default setting is all.

Usage Guidelines

None.

Example

The following command configures the track mode to any:

```
configure vrrp vlan vrrp-1 vrid 1 track-mode any
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vrrp vlan vrid version

```
configure vrrp vlan vlan_name vrid vridval version [v3-v2 | v3 | v2]
```



Description

Selects the VRRP version to apply to the VRRP router instance.

Syntax Description

<code>vlan_name</code>	Specifies the name of a VRRP VLAN.
<code>vridval</code>	Specifies the VRID for the VRRP instance. To display the configured VRRP router instances, enter the <code>show vrrp</code> command.
<code>v3-v2</code>	Specifies VRRP v3 with VRRP v2 compatibility.
<code>v3</code>	Selects VRRP v3.
<code>v2</code>	Specifies VRRP v2.

Default

VRRP v3 with VRRP v2 compatibility.



Note

Configurations created by earlier ExtremeXOS software releases have an implied version of v2. If the configuration is subsequently saved, the version is explicitly set to v2.

Usage Guidelines

None.

Example

The following command configures the VRRP router instance to use VRRP v3 only:

```
configure vrrp vlan vrrp-1 vrid 1 version v3
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

create vrrp vlan vrid

```
create vrrp vlan vlan_name vrid vridval
```



Description

Creates a VRRP instance on the switch.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN.
<i>vridval</i>	Specifies a VRID for the VRRP instance. The value can be in the range of 1-255.

Default

N/A.

Usage Guidelines

VRRP Router IDs can be used across multiple VLANs. You can create multiple VRRP routers on different VLANs. VRRP router IDs need not be unique to a specific VLAN.



Note

The total number of supported VRRP router instances is dependent on the switch hardware. For more information, see the ExtremeXOS Release Notes.

Before configuring any VRRP router parameters, you must first create the VRRP instance on the switch. If you define VRRP parameters before creating the VRRP, you might see an error similar to the following:

```
Error: VRRP VR for vlan vrrp1, vrid 1 does not exist.
Please create the VRRP VR before assigning parameters.
Configuration failed on backup MSM, command execution aborted!
```

If this happens, create the VRRP instance and then configure its parameters.

Example

The following command creates a VRRP router on VLAN vrrp-1, with a VRRP router ID of 1:

```
create vrrp vlan vrrp-1 vrid 1
```

History

This command was first available in ExtremeXOS 10.1.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

delete vrrp vlan vrid

```
delete vrrp vlan vlan_name vrid vridval
```

Description

Deletes a specified VRRP instance.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN.
<i>vridval</i>	Specifies the VRID for the VRRP instance. To display the configured VRRP router instances, enter the <code>show vrrp</code> command.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the VRRP instance on the VLAN vrrp-1 identified by VRID 2:

```
delete vrrp vlan vrrp-1 vrid 2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



disable vrrp vrid

```
disable vrrp {vlan vlan_name vrid vridval}
```

Description

Disables a specific VRRP instance or all VRRP instances.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN.
<i>vridval</i>	Specifies the VRID for the VRRP instance. To display the configured VRRP router instances, enter the <code>show vrrp</code> command.

Default

N/A.

Usage Guidelines

This disables a specific VRRP instance on the switch. If no VRRP VLAN is specified, all VRRP instances on the switch are disabled.

Example

The following command disables all VRRP instances on the switch:

```
disable vrrp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable vrrp vrid

```
enable vrrp {vlan vlan_name vrid vridval}
```



Description

Enables a specific VRRP instance or all VRRP instances on the switch.

Syntax Description

<code>vlan_name</code>	Specifies the name of a VRRP VLAN.
<code>vridval</code>	Specifies the VRID for the VRRP instance to be enabled. To display the configured VRRP router instances, enter the <code>show vrrp</code> command.

Default

N/A.

Usage Guidelines

This enables a specific VRRP instance on the device. If you do not specify a VRRP instance, all VRRP instances on this device are enabled.

Example

The following command enables all VRRP instances on the switch:

```
enable vrrp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show vrrp

```
show vrrp {virtual-router {vr_name}} {detail}
```

Description

Displays VRRP configuration information for all VRRP VLANs.



Syntax Description

vr_name	Specifies a virtual router (VR) for which to display VRRP information.
detail	Specifies more detailed VRRP information.

Default

N/A.

Usage Guidelines

The following table describes the significant fields and values that can appear when you enter the different forms of this command:

Field	Description
Advertisement Interval	Indicates the configured advertisement interval.
Authentication	If configured, identifies the VRRP simple password.
IPv6 Router Advertisement Mask	IPv6 router advertisement mask.
Master Mac Address	The MAC address of the master VRRP router.
P	Indicates that the VRRP VLAN is configured for preempt mode, which controls whether a higher priority backup router preempts a lower priority master.
Preempt	Indicates that the VRRP VLAN is configured for preempt mode, which controls whether a higher priority backup router preempts a lower priority master.
Pri Priority	The priority value of the VRRP VLAN.
State	The current state of the VRRP router. The state includes the following: Init—The VRRP router is in the initial state.Backup—The VRRP router is a backup router.Master—The VRRP router is the master router.
T	Indicates the configured advertisement interval.
TP	Indicates the number of tracked pings.
TR	Indicates the number of tracked routes.
Tracked IP Routes	If configured, displays the IP address and subnet mask of the tracked route(s).
Tracked Pings	If configured, displays the: Target IP address you are pinging.Number of seconds between pings to the target IP address.Number of misses allowed before this entry is considered to be failing.
Tracked VLANs	If configured, displays the name of the tracked VLAN(s).
Tracking Mode	Indicates the VRRP tracking mode, which is either ALL or ANY.
TV	Indicates the number of tracked VLANs.
Virtual IP Addr Virtual IP Addresses	If configured, the virtual IPv4 or IPv6 address associated with the VRRP VLAN.
Virtual Link-Local Address	Virtual IPv6 link local address configured on the interface.
VLAN	The name of the VRRP VLAN.



Field	Description
VLAN Name	The name of the VRRP VLAN and whether VRRP is enabled or disabled on the VLAN. The enable/disable state appears as follows: En—VRRP is enabled on this VLAN.Ds—VRRP is disabled on this VLAN.
VRID	The VRRP Router Identification number for the VRRP instance.
VRRP	The enabled/disabled state of VRRP on the VLAN.

Example

The following example displays summary VRRP status information for the current VR context:

```
# show vrrp

Virtual
  VLAN Name VRID Pri IP Address      Master      MAC Address
TP/TR/TV/P/T
  v1001(En)      0001 100 51.1.1.1      MSTR      00:00:5e:00:01:01
0  0      0  Y  0.10
```

Authentication will not be displayed if its not configured.

```
* (Beta) DUT2.21 # sh vrrp v1001
VLAN: v1001      VRID: 1      VRRP: Enabled State: MASTER
Virtual Router: VR-Default
Priority: 100(backup) Advertisement Interval: 1 sec
Version: v2 Preempt: Yes Preempt Delay: 0 sec
Virtual IP Addresses:
  51.1.1.1
Accept mode: Off
Tracking mode: ALL
Tracked Pings: -
Tracked IP Routes: -
Tracked VLANs: -
```

```
VLAN: v1001      VRID: 2      VRRP: Enabled State: MASTER
Virtual Router: VR-Default
Priority: 100(backup) Advertisement Interval: 1 sec
Version: v3 Preempt: Yes Preempt Delay: 0 sec
Virtual Link Local Address: fe80::30 (Configured)
Virtual IP Addresses:
  3303::100
Accept mode: Off
Tracking mode: ALL
Tracked Pings: -
Tracked IP Routes: -
Tracked VLANs: -
```

If Authentication password is set , the output will be

```
* (Beta) DUT2.27 # conf vrrp v1001 vrid 1 authentication simplepassword
VRRP_NEW
* (Beta) DUT2.28 # sh vrrp v1001
VLAN: v1001      VRID: 1      VRRP: Enabled State: MASTER
Virtual Router: VR-Default
Priority: 100(backup) Advertisement Interval: 0.10 sec
Version: v2 Preempt: Yes Preempt Delay: 0 sec
Authentication: simple-password key: VRRP_NEW
```



```

Virtual IP Addresses:
  51.1.1.1
Accept mode: Off
Tracking mode: ALL
Tracked Pings: -
Tracked IP Routes: -
Tracked VLANs: -

```

The following example displays detailed VRRP status information for the current VR context:

```

show vrrp detail
VLAN: t2t1      VRID: 24      VRRP: Enabled  State: MASTER
Virtual Router: VR-Default
Priority: 255(master) Advertisement Interval: 1 sec
Preempt: Yes   Authentication: None
Virtual Link-Local Address: fe80::1
IPv6 Router Advertisement Mask: 64
Virtual IP Addresses:
10.0.0.10
Tracking mode: ALL
Tracked Pings: -
Tracked IP Routes: -
Tracked VLANs: -
VLAN: t2t1      VRID: 39      VRRP: Enabled  State: MASTER
Virtual Router: VR-Default
Priority: 255(master) Advertisement Interval: 50 centiseconds
Preempt: Yes   Authentication: None
Virtual IP Addresses:
2002::10
fe80::204:96ff:fe10:e610
Tracking mode: ALL
Tracked Pings: -
Tracked IP Routes: -
Tracked VLANs: -

```

The following example displays VRRP information for all VRs:

```

show vrrp virtual-router
VLAN Name VRID Pri Virtual IP Addr State Master Mac Address Virtual-Router
vlan_3(En) 0003 200 30.1.2.1 MSTR 00:00:5e:00:01:03 vir_3
vlan_4(En) 0003 200 40.1.2.1 MSTR 00:00:5e:00:01:03 vir_4
vlan_5(En) 0005 100 50.1.2.1 BKUP 00:00:5e:00:01:05 vir_3
vlan_6(En) 0005 100 60.1.2.1 BKUP 00:00:5e:00:01:05 vir_4

```

History

This command was first available in ExtremeXOS 10.1.

Support for virtual routers was added in ExtremeXOS 12.0.

Support for IPv6 was added in ExtremeXOS 12.7.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show vrrp vlan

```
show vrrp vlan vlan_name {stats}
```

Description

Displays VRRP information for a particular VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN.
stats	Specifies statistics for a particular VLAN.

Default

N/A.

Usage Guidelines

Depending on the software version running on your switch or your switch model, additional or different VRRP information might be displayed.

If you specify the command without the stats keyword, the following VRRP information appears:

- VLAN—The name of the VRRP VLAN.
- VRID—The VRRP Router Identification number for the VRRP instance.
- VRRP—The enabled/disabled state of VRRP on the VLAN.
- State—The current state of the VRRP router. The state includes the following:
 - Init—The VRRP router is in the initial state.
 - Backup—The VRRP router is a backup router.
 - Master—The VRRP router is the master router.
- Priority—The priority value of the VRRP VLAN.
- Advertisement Interval—Indicates the configured advertisement interval.
- Preempt—Indicates that the VRRP VLAN is configured for preempt mode, which controls whether a higher priority backup router preempts a lower priority master.
- Authentication—If configured, identifies the VRRP simple password.
- Virtual IP Addresses—If configured, the virtual IP address associated with the VRRP VLAN.
- Tracked Pings—If configured, displays the:
 - Target IP address you are pinging.
 - Number of seconds between pings to the target IP address.



- Number of misses allowed before this entry is considered to be failing.
- Tracked IP Routes—If configured, displays the IP address and subnet mask of the tracked route(s).
- Tracked VLANs—If configured, displays the name of the tracked VLAN(s).

If you specify the stats keyword, you see counter and statistics information for the specified VRRP VLAN.

Example

The following example displays configuration information for the specified VRRP VLAN:

```
show vrrp vlan blue
VLAN: blue      VRID: 2          VRRP: Disabled State: INIT
Priority: 1(backup)  Advertisement Interval: 1 sec
Preempt: Yes    Authentication: None
Virtual IP Addresses:
Tracked Pings: -
Tracked IP Routes: -
Tracked VLANs: -
* indicates a tracking condition has failed
```

The following example displays statistics for VLAN vrrp-1:

```
show vrrp vlan vrrp-1 stats
VLAN vrrp-1, VR ID 25
Chksum Err:0, Ver Err:0, VRID Err:0, Auth Mismatch:0, Pkt-len Err:0
Become Master:0, Adv recv:0, Adv Err:0, Auth Fail:0, TTL Err:0
Pri-0-recv:0, Pri-0-snt:0, Addr-List Err:0, Invalid Auth Err:0
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



33 MPLS Commands

```
clear counters l2vpn
clear counters mpls
clear counters mpls ldp
clear counters mpls rsvp-te
clear counters mpls static lsp
clear counters vpls
clear fdb vpls
configure forwarding switch-fabric protocol
configure iproute add default
configure iproute add lsp
configure iproute delete
configure iproute delete default
configure l2vpn
configure l2vpn add peer
configure l2vpn add service
configure l2vpn delete peer
configure l2vpn delete service
configure l2vpn health-check vccv
configure l2vpn peer mpls lsp
configure l2vpn peer
configure l2vpn peer mpls lsp
configure l2vpn vpls peer static-pw
configure l2vpn vpls redundancy
configure l2vpn vpls add peer ipaddress
configure l2vpn vpws add peer ipaddress
configure l2vpn vpws peer static-pw
configure mpls add vlan
configure mpls delete vlan
configure mpls exp examination
configure mpls exp replacement
configure mpls labels max-static
configure mpls ldp advertise
configure mpls ldp loop-detection
configure mpls ldp pseudo-wire
configure mpls ldp timers
configure mpls lsr-id
configure mpls rsvp-te bandwidth committed-rate
```

```
configure mpls rsvp-te lsp add path
configure mpls rsvp-te lsp change
configure mpls rsvp-te lsp delete path
configure mpls rsvp-te lsp fast-reroute
configure mpls rsvp-te lsp path use profile
configure mpls rsvp-te lsp transport
configure mpls rsvp-te metric
configure mpls rsvp-te path add ero
configure mpls rsvp-te path delete ero
configure mpls rsvp-te profile
configure mpls rsvp-te profile (fast-reroute)
configure mpls rsvp-te timers lsp rapid-retry
configure mpls rsvp-te timers lsp standard-retry
configure mpls rsvp-te timers session
configure mpls static lsp
configure mpls static lsp transport
configure vpls
configure vpls add peer
configure vpls add service
configure vpls delete peer
configure vpls delete service
configure vpls health-check vccv
configure vpls peer mpls lsp
configure vpls peer
configure vpls peer mpls lsp
configure vpls snmp-vpn-identifier
create l2vpn fec-id-type pseudo-wire
create mpls rsvp-te lsp
create mpls rsvp-te path
create mpls rsvp-te profile
create mpls rsvp-te profile fast-reroute
create mpls static lsp
create vpls fec-id-type pseudo-wire
delete l2vpn
delete mpls rsvp-te lsp
delete mpls rsvp-te path
delete mpls rsvp-te profile
delete mpls static lsp
delete vpls
disable bgp mpls-next-hop
disable iproute mpls-next-hop
disable l2vpn
```



```
disable l2vpn vpls fdb mac-withdrawal
disable l2vpn health-check vccv
disable l2vpn service
disable l2vpn sharing
disable mpls
disable mpls bfd
disable mpls exp examination
disable mpls exp replacement
disable mpls ldp
disable mpls ldp bgp-routes
disable mpls ldp loop-detection
disable mpls php
disable mpls protocol ldp
disable mpls protocol rsvp-te
disable mpls rsvp-te
disable mpls rsvp-te bundle-message
disable mpls rsvp-te fast-reroute
disable mpls rsvp-te lsp
disable mpls rsvp-te summary-refresh
disable mpls static lsp
disable mpls vlan
disable ospf mpls-next-hop
disable snmp traps l2vpn
disable snmp traps mpls
disable vpls
disable vpls fdb mac-withdrawal
disable vpls health-check vccv
disable vpls service
enable bgp mpls-next-hop
enable iproute mpls-next-hop
enable l2vpn
enable l2vpn vpls fdb mac-withdrawal
enable l2vpn health-check vccv
enable l2vpn service
enable l2vpn sharing
enable mpls
enable mpls bfd
enable mpls exp examination
enable mpls exp replacement
enable mpls ldp
enable mpls ldp bgp-routes
enable mpls ldp loop-detection
```



```
enable mpls php
enable mpls protocol ldp
enable mpls protocol rsvp-te
enable mpls rsvp-te
enable mpls rsvp-te bundle-message
enable mpls rsvp-te fast-reroute
enable mpls rsvp-te lsp
enable mpls rsvp-te summary-refresh
enable mpls static lsp
enable mpls vlan
enable ospf mpls-next-hop
enable snmp traps l2vpn
enable snmp traps mpls
enable vpls
enable vpls fdb mac-withdrawal
enable vpls health-check vccv
enable vpls service
ping mpls lsp
restart process mpls
show bandwidth pool
show ces
show l2vpn
show mpls
show mpls bfd
show mpls exp examination
show mpls exp replacement
show mpls interface
show mpls label
show mpls label usage
show mpls ldp
show mpls ldp interface
show mpls ldp label
show mpls ldp label advertised
show mpls ldp label l2vpn
show mpls ldp label l2vpn retained
show mpls ldp label lsp retained
show mpls ldp label retained
show mpls ldp lsp
show mpls ldp peer
show mpls rsvp-te
show mpls rsvp-te bandwidth
show mpls rsvp-te interface
```



```

show mpls rsvp-te lsp
show mpls rsvp-te lsp [egress | transit]
show mpls rsvp-te lsp ingress
show mpls rsvp-te neighbor
show mpls rsvp-te path
show mpls rsvp-te profile
show mpls rsvp-te profile fast-reroute
show mpls static lsp
show mpls statistics l2vpn
show vpls
traceroute mpls lsp
unconfigure l2vpn dot1q ethertype
unconfigure l2vpn vpls redundancy
unconfigure mpls
unconfigure mpls exp examination
unconfigure mpls exp replacement
unconfigure mpls vlan
unconfigure vpls dot1q ethertype
unconfigure vpls snmp-vpn-identifier

```

This chapter describes commands for configuring and managing the following Multiprotocol Label Switching (MPLS) features:

- Basic MPLS.
- Virtual Private LAN Service (VPLS) Layer-2 Virtual Private Networks (VPNs).
- Hierarchical VPLS (H-VPLS).
- Protected VPLS.
- Resource ReSerVation Protocol-Traffic Engineering (RSVP-TE).
- Time Division Multiplexing (TDM) Pseudowire Ethernet (PWE) over MPLS.

For an introduction to MPLS, see the *ExtremeXOS Concepts Guide*.



Note

The commands in this chapter operate only on the platforms listed for specific features in FOO in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

clear counters l2vpn

```
clear counters l2vpn [vpls [vpls_name | all] | vpws [vpws_name | all]]
```

Description

Clears all the specified VPLS or VPWS counters.



Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
<i>vpws_name</i>	Identifies the VPWS within the switch (character string).
all	Specifies all VPLS or VPWS VPNs.

Default

N/A.

Usage Guidelines

The l2vpn keyword was introduced in ExtremeXOS Release 12.4 and is required when clearing counters for a VPWS. For backward compatibility, the l2vpn keyword is optional when clearing counters for a VPLS. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

This example clears all VPLS counters for the specified VPLS:

```
clear counters vpls myvpls
```

This example clears all VPWS counters for the specified VPWS:

```
clear counters l2vpn vpws myvpws
```

History

This command was first available in ExtremeXOS 11.6.

The l2vpn and vpws keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support this feature as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

clear counters mpls

```
clear counters mpls {[lsp all | [{vlan} vlan_name | vlan all]]}
```



Description

Clears all packet and byte counters for all MPLS LSPs and all MPLS protocol counters for all MPLS interfaces.

Syntax Description

lsp all	Clears all MPLS protocol counters for all MPLS LSPs.
<i>vlan_name</i>	Clears all MPLS protocol counters for the MPLS interface on the specified VLAN.
vlan all	Clears all MPLS protocol counters for all MPLS interfaces.

Default

N/A.

Usage Guidelines

This command clears all packet and byte counters for all MPLS LSPs and all MPLS protocol counters for all MPLS interfaces. If the `lsp all` keywords are specified, all packet and byte counters for all MPLS LSPs are cleared. If the `vlan all` keywords are specified, all MPLS protocol counters for all MPLS interfaces are cleared. If a VLAN name is specified, all MPLS protocol counters for the MPLS interface on that VLAN are cleared.

Example

```
This example clears all MPLS counters associated with VLAN 1:
clear counters mpls vlan vlan_1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support this feature as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

clear counters mpls ldp

```
clear counters mpls ldp {{{vlan} vlan_name} | lsp all}
```

Description

Clears LDP control protocol counters and packet and byte counters associated with LDP LSPs.



Syntax Description

<i>vlan_name</i>	Clears LDP control protocol counters on the specified VLAN.
vlan all	Clears LDP control protocol counters on all MPLS interfaces.
lsp all	Clears all LDP LSP packet and byte counters.

Default

N/A.

Usage Guidelines

By default, all LDP control protocol counters are cleared for all LDP interfaces and all byte counters. Specifying the `vlan` keyword clears only the protocol counters associated with a specified LDP interface. Specifying the `lsp` keyword clears only the packet and byte counters associated with LDP LSPs.

Example

```
This example clears all LDP control protocol counters and all packet and byte
counters for all LDP LSPs:
clear counters mpls ldp
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support this feature as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

clear counters mpls rsvp-te

```
clear counters mpls rsvp-te {[lsp all | [{vlan} vlan_name | vlan all]]}
```

Description

Clears all packet and byte counters for all RSVP-TE LSPs and all RSVP-TE protocol counters for all MPLS interfaces.



Syntax Description

lsp all	Clears all packet and byte counters for all RSVP-TE LSPs.
<i>vlan_name</i>	Clears all RSVP-TE protocol counters for the MPLS interface on the specified VLAN.
vlan all	Clears all RSVP-TE protocol counters on all MPLS interfaces.

Default

By default, all RSVP-TE control protocol counters are cleared for all RSVP-TE interfaces.

Usage Guidelines

This command clears all packet and byte counters for all RSVP-TE LSPs and all RSVP-TE protocol counters for all MPLS interfaces. If the `lsp all` keywords are specified, all packet and byte counters for all RSVP-TE LSPs are cleared. If the `vlan all` keywords are specified, all RSVP-TE protocol counters for all MPLS interfaces are cleared. If a VLAN name is specified, all RSVP-TE protocol counters for the MPLS interface on that VLAN are cleared.

Example

```
This example clears the RSVP-TE protocol counters on VLAN 1 only:
clear counters mpls rsvp-te vlan vlan_1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support this feature as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

clear counters mpls static lsp

```
clear counters mpls static lsp {lsp_name | all }
```

Description

Clears the packet and byte counters for one or all static LSPs.

Syntax Description

<i>lsp_name</i>	Identifies the LSP for which counters are to be cleared.
all	Specifies that counters are to be cleared for all static LSPs on this LSR.



Default

N/A.

Usage Guidelines

None.

Example

The following command clears the counters for a static LSP:

```
clear counters mpls static lsp lsp598
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support this feature as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

clear counters vpls

```
clear counters vpls [vpls_name | all]
```

Description

Clears all VPLS counters for the specified *vpls_name*.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string)
all	Specifies all VPLS VPNs.

Default

N/A.

Usage Guidelines

This command clears all VPLS counters for the specified *vpls_name*. If the optional **all** keyword is specified, all packet and byte counters for all VPLS VPNs are cleared.



Example

This example clears all VPLS counters for the specified VPLS:

```
clear counters vpls myvpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support this feature as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

clear fdb vpls

```
clear fdb vpls {vpls_name {peer_ip_address}}
```

Description

Clears the FDB information learned for VPLS.

Syntax Description

<i>vpls_name</i>	Clears all FDB entries for the specified VPLS and its associated VLAN.
<i>peer_ip_address</i>	Clears all FDB entries for the pseudowire (PW) associated with the specified VPLS and LDP peer.

Default

N/A.

Usage Guidelines

If the command is used without keywords, every FDB entry learned from any PW is cleared. Using the keywords *vpls_name* clears every FDB entry, (both PW and front panel Ethernet port for the service VLAN) associated with the specified VPLS and the associated VLAN. If the specified VPLS is not bound to a VLAN, the following error message appears:

```
Error: vpls VPLS_NAME not bound to a vlan
```

Using the keywords *vpls_name* and *peer_ip_address* clears all FDB entries from the PW associated with the specified VPLS and LDP peer.



Once the information is cleared from the FDB, any packet destined to a MAC address that has been flushed from the hardware is flooded until the MAC address has been re-learned.

Example

This example clears the FDB information for VPLS 1:

```
clear fdb vpls vpls1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support this feature as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure forwarding switch-fabric protocol

```
configure forwarding switch-fabric protocol [standard | enhanced]
```

Description

Configures the switch-fabric protocol on a BlackDiamond 8800 series switch.

Syntax Description

standard	Specifies the standard protocol, which is supported on all switches.
enhanced	Specifies the enhanced protocol, which is supported only on BlackDiamond 8900 series modules and Summit X460, X480, and X670 series switches. The enhanced protocol is required to support MPLS.

Default

Standard.



Usage Guidelines

To support MPLS, the BlackDiamond 8800 series switch must use one or two 8900-MSM128 cards and only the following BlackDiamond 8900 series modules: 8900-G96T-c and all BlackDiamond 8900 xl- and xm-series modules.



Note

You do not need to reboot the switch to activate the protocol change, but there will be some traffic interruption during protocol activation.

If MPLS is enabled on the switch, you must disable MPLS before you can change the stacking protocol to standard.

To display the switch-fabric protocol configuration, enter the show forwarding configuration command.

Example

To configure a switch to use the enhanced protocol, enter the following command:

```
configure forwarding switch-fabric protocol enhanced
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available only on BlackDiamond 8800 series switches that have the proper hardware installed as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure iproute add default

```
configure iproute add default [{gateway {metric} {vr vr_name} {unicast-only | multicast-only}} | {lsp lsp_name {metric}}]
```

Description

Assigns a named LSP to the default route.

Syntax Description

<i>gateway</i>	Specifies a gateway.
<i>metric</i>	Specifies a cost metric.
<i>vr_name</i>	Specifies the virtual router to which the route is added.
unicast-only	Specifies only unicast traffic for the route.



multicast-only	Specifies only multicast traffic for the route.
<i>lsp_name</i>	Specifies a named MPLS LSP to be used to reach the default route.
<i>metric</i>	Specifies a cost metric.

Default

N/A.

Usage Guidelines

This command assigns a named LSP to the default route. Once configured, all IP traffic matching the configured route is forwarded over the specified LSP. For an RSVP-TE LSP, the correct label information is only associated with the default route if the LSP is active. If the RSVP-TE LSP is disabled or is inactive, the label information is removed from the route table and the default route entry is marked down. If multiple LSPs are added to the default route and ECMP is enabled using the route-sharing command, traffic is forwarded over only one LSP.



Note

IP routes can only be assigned to named LSPs in the VR in which MPLS is configured to operate.

Example

The following command specifies a default route for the switch:

```
configure iproute add default lsp lsp598
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure iproute add lsp

```
configure iproute add [ipaddress netmask | ipNetmask] lsp lsp_name {metric}
{multicast | multicast-only | unicast | unicast-only} {vr vrname}
```



Description

Assigns a specific IP route to use a named LSP.

Note



To create a static IP route that does not use a specific named LSP as an mpls-next-hop, use the following command: `configure iproute add [<ipNetmask | ip_addr mask] gateway {bfd} {metric} {multicast | multicast-only | unicast | unicast-only} {vrvrname} .`

Syntax Description

<i>ipaddress</i>	Specified an IP address.
<i>netmask</i>	Specifies an IP address/prefix length.
<i>ipNetmask</i>	Specifies an IP address/prefix length.
<i>lsp_name</i>	Specifies a named MPLS LSP to be used to reach the route.
metric	Specifies a cost metric.
multicast	Adds the specified route to the multicast routing table.
multicast-only	Adds the specified route to the multicast routing table. This option is provided for backward compatibility with releases prior to ExtremeXOS Release 12.1.
unicast	Adds the specified route to the unicast routing table.
unicast-only	Adds the specified route to the unicast routing table. This option is provided for backward compatibility with releases prior to ExtremeXOS Release 12.1.
<i>vrname</i>	Specifies the virtual router to which the route is added.

Default

N/A.

Usage Guidelines

This command assigns a named LSP to a specific IP route. Once configured, all IP traffic matching the configured route is forwarded over the specified LSP. For an RSVP-TE LSP, the correct label information is only associated with the route if the LSP is active. If the RSVP-TE LSP is disabled or is withdrawn, the label information is removed from the route table and the route entry is marked down. If multiple LSPs are added to a route and ECMP is enabled using route-sharing command, only one LSP is used to forward IP traffic.

Note



IP routes can only be assigned to named LSPs in the VR in which MPLS is configured to operate.



Example

The following command adds a static address to the routing table:

```
configure iproute add 10.1.1.0/24 lsp lsp598
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure iproute delete

```
configure iproute delete [ipaddress netmask | ipNetmask] [{gateway {vr vr_name}} | lsp lsp_name]
```

Description

Deletes a route entry from the routing table based on the configured LSP.

Syntax Description

<i>ipaddress</i>	Specified an IP address.
<i>netmask</i>	Specifies an IP address/prefix length.
<i>ipNetmask</i>	Specifies an IP address/prefix length.
<i>gateway</i>	Specifies a gateway.
<i>vr_name</i>	Specifies the virtual router from which the route is deleted.
<i>lsp_name</i>	Specifies the named MPLS LSP used to reach the route.

Default

N/A.



Usage Guidelines

This command deletes a route entry from the routing table based on the configured LSP. If the configured IP netmask and LSP name do not match any route entry, the command fails and nothing is deleted.



Note

IP routes can only be assigned to named LSPs in the VR in which MPLS is configured to operate.

Example

The following command deletes an address from the LSP:

```
configure iproute delete 10.1.1.0/24 lsp lsp598
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure iproute delete default

```
configure iproute delete default [{gateway {vr vr_name}} | lsp lsp_name]
```

Description

Deletes a default route entry from the routing table based on the configured LSP.

Syntax Description

<i>gateway</i>	Specifies a gateway.
<i>vr_name</i>	Specifies the virtual router from which the route is deleted
<i>lsp_name</i>	Specifies the named MPLS LSP used to reach the default route.

Default

N/A.



Usage Guidelines

This command deletes a default route entry from the routing table based on the configured LSP. If the specified LSP name doesn't match any configured default route entry, the command fails and nothing is deleted.



Note

IP routes can only be assigned to named LSPs in the VR in which MPLS is configured to operate.

Example

The following command deletes a default route from the LSP:

```
configure iproute delete default lsp lsp598
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure l2vpn

```
configure l2vpn [vpls vpls_name | vpws vpws_name] {dot1q [ethertype hex_number | tag [include | exclude]]} {mtu number}
```

Description

Configures VPLS or VPWS parameters.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
<i>vpws_name</i>	Identifies the VPWS within the switch (character string).
dot1q	Specifies the action the switch performs with respect to the 802.1Q ethertype or tag.
ethertype	Overwrites the ethertype value for the customer traffic sent across the PW
<i>hex_number</i>	Identifies the ethertype, uses the format of 0xN.
tag	Specifies the action the switch performs with respect to the 802.1Q tag.
include	Includes the 802.1Q tag when sending packets over the VPLS L2 VPN.
exclude	Strips the 802.1Q tag before sending packets over the VPLS L2 VPN.



mtu	Specifies the MTU value of the VPLS transport payload packet.
number	The size (in bytes) of the MTU value. The configurable MTU range is 1492 through 9216. The default VPLS MTU value is 1500.

Default

dot1q tag - excluded.

ethertype - the configured switch ethertype is used.

number (MTU) - 1500.

Usage Guidelines

This command configures the VPLS and VPWS parameters. PWs are point-to-point links used to carry VPN traffic between two devices within the VPLS. Each device must be configured such that packets transmitted between the endpoints are interpreted and forwarded to the local service correctly. The optional ethertype keyword may be used to overwrite the Ethertype value for the customer traffic sent across the PW. By default, the configured switch ethertype is used. If configured, the ethertype in the outer 802.1q field of the customer packet is overwritten using the configured ethertype value. The ethertype value is ignored on receipt.

Optionally, the switch can be configured to strip the 802.1q tag before sending packets over the VPLS or VPWS Layer 2 VPN. This capability may be required to provide interoperability with other vendor products or to emulate port mode services. The default configuration is to include the 802.1q tag.

The mtu keyword optionally specifies the MTU value of the VPLS or VPWS transport payload packet (customer packet). The MTU value is exchanged with VPLS-configured peer nodes. All VPLS peer nodes must be configured with the same MTU value. If the MTU values do not match, PWs cannot be established between peers. The MTU values are signaled during PW establishment so that endpoints can verify that MTU settings are equivalent before establishing the PW. By default the MTU is set to 1500. The configurable MTU range is 1492 through 9216. Changing the MTU setting causes established PWs to terminate. Payload packets might be dropped if the VPLS or VPWS MTU setting is greater than the MPLS MTU setting for the PW interface.



Note

The maximum MTU value supported depends on the current configuration options. For more information, see [Configure the Layer 2 VPN MTU](#) in the *ExtremeXOS Concepts Guide*.

The l2vpn keyword was introduced in ExtremeXOS Release 12.4 and is required when enabling a VPWS. For backward compatibility, the l2vpn keyword is optional when enabling a VPLS. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.



Example

The following commands change the various parameters of a particular VPLS:

```
configure vpls vpls1 dot1q ethertype 0x8508
configure vpls vpls1 dot1q ethertype 0x8509 mtu 2500
configure vpls vpls1 dot1q tag exclude mtu 2430
configure vpls vpls1 dot1q mtu 2500
```

History

This command was first available in ExtremeXOS 11.6.

The l2vpn and vpws keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure l2vpn add peer

```
configure l2vpn [vpls vpls_name | vpws vpws_name] add peer ipaddress {{core
{full-mesh | primary | secondary} | spoke}
```

Description

Configures a VPLS, H-VPLS, or VPWS peer for the node you are configuring.

Syntax Description

<i>vpls_name</i>	Specifies the VPLS for which you are configuring a peer.
<i>vpws_name</i>	Specifies the VPWS for which you are configuring a peer.
<i>ipaddress</i>	Specifies the IP address of the peer node.
core	Specifies that the peer is a core node. This option applies only to VPLS peers.
full-mesh	Specifies that the peer is a core full-mesh node. This is the default setting if neither the core or spoke options are specified. This option applies only to VPLS peers.
primary	Specifies that the peer is an H-VPLS core node and configures a primary H-VPLS connection to that core node. This option applies only to H-VPLS peers.
secondary	Specifies that the peer is an H-VPLS core node and configures a secondary H-VPLS connection to that core node. This option applies only to H-VPLS peers.
spoke	Specifies that the peer is a H-VPLS spoke node. This option applies only to H-VPLS peers.



Default

N/A.

Usage Guidelines

Each VPLS or H-VPLS node supports up to 64 peers, and each VPWS supports one peer. H-VPLS core nodes can peer with other core nodes and/or spoke nodes. H-VPLS spoke nodes can peer with core nodes but not with other spoke nodes.

VPLS core nodes must be configured in a full-mesh with other core nodes. Thus, all core nodes in the VPLS must have a configured PW to every other core node serving this VPLS. By default, the best LSP is chosen for the PW. The underlying LSP used by the PW can be configured by specifying the named LSP using the CLI command `configure l2vpn [vpls vpls_name | vpwsvpws_name] peeripaddress [add | delete] mpls lsp_lsp_name .`

H-VPLS spoke nodes establish up to two point-to-point connections to peer with core nodes. If both primary and secondary peers are defined for a spoke node, the spoke node uses one of the peers for all communications. If both peers are available, the spoke node uses the connection to the primary peer. If the primary peer connection fails, the spoke node uses the secondary peer. If the primary peer later recovers, the spoke node reverts back to using the primary peer.

VPWS nodes establish a point-to-point connection to one peer.

The `l2vpn` keyword was introduced in ExtremeXOS Release 12.4 and is required when configuring a VPWS peer. For backward compatibility, the `l2vpn` keyword is optional when configuring a VPLS peer. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The following command adds a connection from the local core switch to the core switch at 1.1.1.202:

```
configure l2vpn vpls vpls1 add peer 1.1.1.202
```

The following command adds a connection from the local core switch to the spoke switch at 1.1.1.201:

```
configure l2vpn vpls vpls1 add peer 1.1.1.201 spoke
```

The following command adds a primary connection from the local spoke switch to the core switch at 1.1.1.203:

```
configure l2vpn vpls vpls1 add peer 1.1.1.203 core primary
```

The following command adds a VPWS connection from the local node to the peer switch at 1.1.1.204:

```
configure l2vpn vpws vpws1 add peer 1.1.1.204
```



History

This command was first available in ExtremeXOS 11.6.

Support for H-VPLS was first available in ExtremeXOS 12.1.

The **l2vpn** and **vpws** keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure l2vpn add service

```
configure l2vpn [vpls vpls_name | vpws vpws_name] add service [{vlan} vlan_name | {vman} vman_name]
```

Description

Adds a VLAN or VMAN service to a VPLS or VPWS.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS interface within the switch (character string).
<i>vpws_name</i>	Identifies the VPWS interface within the switch (character string).
<i>vlan_name</i>	Logically binds the VLAN to the specified VPLS or VPWS.
<i>vman_name</i>	Adds the named VMAN to the VPLS or VPWS.

Default

N/A.

Usage Guidelines

Only one VLAN or VMAN can be configured per VPLS or VPWS.

When a VLAN service is added to a VPLS or VPWS, the VLAN ID is locally significant to the switch. Thus, each VLAN VPLS or VPWS interface within the Layer 2 VPN can have a different VLAN ID. This greatly simplifies VLAN ID coordination between metro network access points. Traffic may be switched locally between VLAN ports if more than one port is configured for the VLAN.

When a VMAN service has been configured for a VPLS or VPWS, the VMAN ID is locally significant to the switch. Thus, each VMAN VPLS or VPWS interface within the Layer 2 VPN can have a different VMAN ID, just like the VLAN service. The only difference is that the Layer 2 VPN overwrites the outer VMAN tag on Layer 2 VPN egress and leaves the inner VLAN tag unmodified. Because the inner VLAN tag is considered part of the customer packet data, the VMAN service can be used to emulate port-based services. This is accomplished by configuring the Layer 2 VPN to strip the 802.1Q tag from the



tunneled packet. Since the switch inserts the VMAN tag when the packet is received and the 802.1Q tag is stripped before the packet is sent on the VPLS or VPWS PW, all packets received on ports that are members of the VMAN are transmitted unmodified across the Layer 2 VPN. The command `configure l2vpn [vpls vpls_name | vpwsvpws_name] dot1q tag exclude` is used to configure the switch to strip the 802.1Q tag on the VPLS.

The `l2vpn` keyword was introduced in ExtremeXOS Release 12.4 and is required when adding a service to VPWS. For backward compatibility, the `l2vpn` keyword is optional when adding a service to VPLS. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The example below adds a VLAN and a VMAN to the named VPLS:

```
configure l2vpn vpls myvpls add service vlan myvlan
configure l2vpn vpls myvpls add service vman myvman
```

The following example adds a VLAN and a VMAN to the named VPWS:

```
configure l2vpn vpws myvpws add service vlan vlan2
```

History

This command was first available in ExtremeXOS 11.6.

The `l2vpn` and `vpws` keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure l2vpn delete peer

```
configure l2vpn [vpls vpls_name | vpws vpws_name] delete peer [ipaddress | all]
```

Description

Deletes the specified VPLS or VPWS peer.

Syntax Description

<code>vpls_name</code>	Identifies the VPLS within the switch (character string).
<code>vpws_name</code>	Identifies the VPWS within the switch (character string).



<i>ipaddress</i>	Specifies the IP address for the peer node that is the endpoint of the VC-LSP. This option applies only to VPLS peers.
all	Deletes all VPLS or VPWS peers. This option applies only to VPLS peers.

Default

N/A.

Usage Guidelines

When the VPLS or VPWS peer is deleted, VPN connectivity to the peer is terminated. The **all** keyword can be used to delete all peers associated with the specified Layer 2 VPN.

The **l2vpn** keyword was introduced in ExtremeXOS Release 12.4 and is required when deleting a VPWS peer. For backward compatibility, the **l2vpn** keyword is optional when deleting a VPLS peer. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The following example removes connectivity to 1.1.1.202 from VPLS1:

```
configure vpls vpls1 delete peer 1.1.1.202
```

History

This command was first available in ExtremeXOS 11.6.

The **l2vpn** and **vpws** keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure l2vpn delete service

```
configure l2vpn [vpls vpls_name | vpws vpws_name] delete service [{vlan}  
vlan_name | {vman} vman_name]
```

Description

Deletes the specified VLAN or VMAN service from the specified Layer 2 VPN.



Syntax Description

<i>vpls_name</i>	Identifies the VPLS interface within the switch (character string).
<i>vpws_name</i>	Identifies the VPWS interface within the switch (character string).
<i>vlan_name</i>	Logically binds the VLAN to the specified VPLS.
<i>vman_name</i>	Adds the named VMAN to the VPLS.

Default

N/A.

Usage Guidelines

If there are no services configured for the VPLS or VPWS, all PWs within the Layer 2 VPN are terminated from the switch.

The `l2vpn` keyword was introduced in ExtremeXOS Release 12.4 and is required when deleting a service from a VPWS. For backward compatibility, the `l2vpn` keyword is optional when deleting a service from a VPLS. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The following example removes a service interface from a VPLS:

```
configure vpls vpls1 delete vman vman1
```

History

This command was first available in ExtremeXOS 11.6.

The `l2vpn` and `vpws` keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure l2vpn health-check vccv

```
configure l2vpn [vpls [vpls_name | all] | vpws [vpws_name | all]] health-check
vccv {interval interval_seconds} {fault-multiplier fault_multiplier_number}
```



Description

Configures the Virtual Circuit Connectivity Verification (VCCV) health check test and fault notification intervals for the specified VPLS or VPWS instance.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS instance for which health check is to be configured.
<i>vpws_name</i>	Identifies the VPWS instance for which health check is to be configured.
all	Specifies that the configuration applies to all VPLS instances on the local node.
<i>interval_seconds</i>	Defines the interval between health check tests. The range is 1 to 10 seconds.
<i>fault_multiplier_number</i>	Specifies how long health check waits before a warning level message is logged. The wait period is the <i>interval_seconds</i> multiplied by the <i>fault_multiplier_number</i> . The <i>fault_multiplier_number</i> range is 2 to 6.

Default

Interval is 5 seconds.

Fault multiplier is 4.

Usage Guidelines

The VCCV health-check configuration parameters can be configured at anytime after the VPLS has been created.

The `show l2vpn {vpls {vpls_name} | vpws {vpws_name}} {peeripaddress} {detail} | summary` command displays the configured *interval_seconds* and *fault-multiplier_number* values for the VPLS or VPWS and the VCCV activity state.

The `l2vpn` keyword is introduced in ExtremeXOS Release 12.4 and is required when configuring health check for a VPWS. For backward compatibility, the `l2vpn` keyword is optional when configuring health check for a VPLS. However, this keyword will be required in a future release, so Extreme Networks recommends that you use this keyword for new configurations and scripts.

Example

The following command configures the health check feature on the VPLS instance `myvpls`:

```
configure vpls myvpls health-check vccv interval 10 fault-notification 40
```

History

This command was first available in ExtremeXOS 12.1.

The `l2vpn` and `vpws` keywords were first available in ExtremeXOS 12.4.



Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure l2vpn peer mpls lsp

```
configure l2vpn [vpls vpls_name | vpws vpws_name] peer ipaddress [add | delete]
mpls lsp lsp_name
```

Description

Adds or deletes a named LSP as a specified PW for the specified Layer 2 VPN peer.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
<i>vpws_name</i>	Identifies the VPWS within the switch (character string).
<i>ipaddress</i>	Specifies the IP address for the peer node that is the endpoint of the PW-LSP. This option applies only to VPLS peers.
add	Permits addition of up to four RSVP-TE LSPs to the VPLS peer.
delete	Removes the LSP specified by the <i>lsp_name</i> parameter from the PW-LSP aggregation list.
<i>lsp_name</i>	Removes the specified lsp.

Default

N/A.

Usage Guidelines

If all the named LSPs are deleted from the configured Layer 2 VPN peer, VPLS or VPWS attempts to use the best-routed path LSP, if one exists. The delete portion of this command cannot be used to remove a named LSP that was selected by the switch as the best LSP. If no LSPs exist to the peer, Layer 2 VPN connectivity to the peer is lost. Currently, the VPLS or VPWS PW uses only one LSP.

The `l2vpn` keyword was introduced in ExtremeXOS Release 12.4 and is required when configuring a VPWS instance. For backward compatibility, the `l2vpn` keyword is optional when configuring a VPLS instance. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.



Example

The following examples add and remove a named LSP:

```
configure l2vpn vpls vpls1 peer 1.1.1.202 add mpls lsp "to-olympic4"
configure l2vpn vpls vpls1 peer 1.1.1.202 delete mpls lsp "to-olympic4"
```

The following example adds a named LSP for a VPWS peer:

```
configure l2vpn vpws vpws1 peer 1.1.1.203 add mpls lsp "to-olympic5"
```

History

This command was first available in ExtremeXOS 11.6.

The l2vpn and vpws keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure l2vpn peer

```
configure l2vpn [vpls vpls_name | vpws vpws_name] peer ipaddress [limit-learning
number | unlimited-learning]
```

Description

Configures the maximum number of MAC SAs (Source Addresses) that can be learned for a given VPLS or VPWS peer.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
<i>vpws_name</i>	Identifies the VPWS within the switch (character string).
<i>ipaddress</i>	Specifies the IP address for the peer node that is the endpoint of the PW-LSP. This option applies only to VPLS peers.
limit-learning	Specifies a limit to the number of MAC SAs to be learned for the specified VPLS and peer.
<i>number</i>	The maximum number of MAC SAs that can be learned for the specified VPLS and peer.
unlimited-learning	Specifies no limit to the number of MAC SAs to be learned for the specified VPLS and peer.



Default

Unlimited.

Usage Guidelines

This parameter can only be modified when the specified VPLS or VPWS is disabled. The unlimited-learning keyword can be used to specify that there is no limit. The default value is unlimited-learning.

The l2vpn keyword was introduced in ExtremeXOS Release 12.4 and is required when configuring a VPWS instance. For backward compatibility, the l2vpn keyword is optional when configuring a VPLS instance. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The following example causes no more than 20 MAC addresses to be learned on VPLS1's PW to 1.1.1.202:

```
configure vpls vpls1 peer 1.1.1.202 limit-learning 20
```

History

This command was first available in ExtremeXOS 11.6.

The l2vpn and vpws keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure l2vpn peer mpls lsp

```
configure l2vpn [vpls vpls_name | vpws vpws_name] peer ipaddress [add | delete]  
mpls lsp lsp_name
```

Description

Adds or deletes a named LSP as a specified PW for the specified Layer 2 VPN peer.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
<i>vpws_name</i>	Identifies the VPWS within the switch (character string).



<i>ipaddress</i>	Specifies the IP address for the peer node that is the endpoint of the PW-LSP. This option applies only to VPLS peers.
add	Permits addition of up to four RSVP-TE LSPs to the VPLS peer.
delete	Removes the LSP specified by the <i>lsp_name</i> parameter from the PW-LSP aggregation list.
<i>lsp_name</i>	Removes the specified lsp.

Default

N/A.

Usage Guidelines

If all the named LSPs are deleted from the configured Layer 2 VPN peer, VPLS or VPWS attempts to use the best-routed path LSP, if one exists. The delete portion of this command cannot be used to remove a named LSP that was selected by the switch as the best LSP. If no LSPs exist to the peer, Layer 2 VPN connectivity to the peer is lost. Currently, the VPLS or VPWS PW uses only one LSP.

The `l2vpn` keyword was introduced in ExtremeXOS Release 12.4 and is required when configuring a VPWS instance. For backward compatibility, the `l2vpn` keyword is optional when configuring a VPLS instance. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The following examples add and remove a named LSP:

```
configure l2vpn vpls vpls1 peer 1.1.1.202 add mpls lsp "to-olympic4"
configure l2vpn vpls vpls1 peer 1.1.1.202 delete mpls lsp "to-olympic4"
```

The following example adds a named LSP for a VPWS peer:

```
configure l2vpn vpws vpws1 peer 1.1.1.203 add mpls lsp "to-olympic5"
```

History

This command was first available in ExtremeXOS 11.6.

The `l2vpn` and `vpws` keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure l2vpn vpls peer static-pw

```
configure l2vpn vpls vpls_name peer ipaddress static-pw {transmit-label
outgoing_pw_label receive-label incoming_pw_label }
```

Description

Changes the labels of a statically configured Ethernet PW for a VPLS.

Syntax Description

<i>vpls_name</i>	Specifies the VPLS for which you are configuring a peer.
peer	Specifies the peer IP address.
<i>ipaddress</i>	Specifies the IP address of the peer node.
static-pw	Specifies the static pseudowire.
transmit label	Specifies the pseudowire transmit label.
<i>outgoing_pw_label</i>	Specifies the name of the egress label.
receive-label	Specifies the static pseudowire receive label.
<i>incoming_pw_label</i>	Specifies the name of the ingress label.

Default

N/A.

Usage Guidelines

Use this command to change the labels of a statically configured Ethernet PW for a VPLS that already exists. Either or both the outgoing (MPLS ingress) and incoming (MPLS egress) PW labels can be specified. The peer must be similarly configured with a static PW that has the reverse PW label mappings. Locally, the *incoming_pw_label* must be unique and is allocated out of the static label space. The *outgoing_pw_label* must match the peer's configured incoming PW label. The CES or L2VPN can remain operational during the change; however, the PW will go down and come back up.

Example

The following command changes the VPLS name to "VPLS1":

```
configure l2vpn vpls vpls1 peer 1.1.1.202
```

History

This command was first available in ExtremeXOS 15.4.



Platform Availability

This command is available on all platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure l2vpn vpls redundancy

```
configure {l2vpn} vpls vpls_name redundancy [esrp esrpDomain | eaps | stp]
```

Description

Configures a VPLS instance to provide protected access using the EAPS redundancy type, the specified ESRP domain, or STP.

Syntax Description

vpls_name	Specifies the VPLS for which you are configuring protection.
esrpDomain	Configures a VPLS instance to provide protected access using the specified ESRP domain.
eaps	Configures a VPLS instance to use the EAPS redundancy type.
stp	Configures a VPLS instance to request an FDB relearning process on an adjacent node when STP responds to a topology change for a VLAN.

Default

Redundancy disabled.

Usage Guidelines

Only one redundancy mode can be configured at a time on a VPLS, and the VPLS must be disabled when the redundancy mode is configured. If you attempt to configure a second mode, an error appears. The current redundancy mode must be unconfigured before you configure a different redundancy mode.

The ESRP domain specified must be a valid ESRP domain of type vpls-redundancy. If not, the command is rejected with an appropriate error message. When a VPLS instance is associated with an ESRP domain, the user cannot delete the ESRP domain unless the VPLS redundancy has been unconfigured. For VPLS access protection to become fully functional, VPLS redundancy must also be configured on a second VPLS peer using the same VPLS name and ESRP domain.

Specify the redundancy type as EAPS when using redundant EAPS access rings. This configuration requires EAPS shared links to be configured between redundant VPLS nodes. This configures VPLS to use a PW between VPLS attachment nodes instead of using a customer VLAN. This configuration is only required when there is an EAPS ring on the VPLS service VLAN.



Note

The EAPS master should not be on a VPLS node.



The STP option enables VPLS interfaces to respond appropriately to STP topology changes in a VLAN. For example, if STP detects a link failure, it will flush the appropriate FDB entries to initiate relearning on the STP protected interfaces. When this option is selected and STP initiates relearning, the VPLS interfaces on the same VLAN also initiate relearning so that a new VLAN path to the VPLS core can be learned. For more information, including limitations and restrictions, see [VPLS STP Redundancy Overview](#) in the *ExtremeXOS Concepts Guide*.

The l2vpn keyword was introduced in ExtremeXOS Release 12.4. For backward compatibility, the l2vpn keyword is optional when configuring a VPLS instance. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The following command adds redundancy to the vpls1 VPLS using the esrp1 domain:

```
configure l2vpn vpls vpls1 redundancy esrp esrp1
```

The following command specifies the EAPS redundancy type for the vpls2 VPLS:

```
configure l2vpn vpls vpls2 redundancy eaps
```

The following command specifies the STP redundancy type for the vpls3 VPLS:

```
configure l2vpn vpls vpls3 redundancy STP
```

History

This command was first available in ExtremeXOS 12.1.

The l2vpn keyword and the STP option were added in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure l2vpn vpls add peer ipaddress

```
configure {l2vpn} vpls vpls_name add peer ipaddress { static-pw transmit-label outgoing_pw_label receive-label incoming_pw_label }
```

Description

Configures L2VPN VPLS service over MPLS Static PW.



Syntax Description

<i>vpls_name</i>	Specifies the VPLS for which you are configuring a peer.
<i>ipaddress</i>	Specifies the IP address of the peer node.
static-pw transmit label	Specifies the static pseudowire transmit label.
<i>outgoing_pw_label</i>	Specifies the name of the egress label.
receive-label	Specifies the static PW receive label.
<i>incoming_pw_label</i>	Specifies the name of the ingress label.

Default

N/A.

Usage Guidelines

Use this command to statically configure a new MPLS Ethernet PW for the specified VPLS. You must specify the outgoing (MPLS ingress) and incoming (MPLS egress) PW labels. Similarly, you must configure the peer with a static PW that has the reverse PW label mappings.

Locally, the *incoming_pw_label* must be unique and is allocated out of the static label space. The *outgoing_pw_label* must match the peer's configured incoming PW label.

Just like a signaled PW, a static PW can optionally be configured to use any type of tunnel LSP: LDP, RSVP-TE, or Static. In the case of RSVP-TE and LDP, those protocols must be configured and enabled and an LSP must be established before traffic can be transmitted over the static PW.

For Static LSPs, only the MPLS ingress LSP (or outgoing LSP) is specified. Unlike signaled PWs, there is no end-to-end PW communication that is used to verify that the PW endpoint is operational, and in the case of static LSPs, that the data path to the PW endpoint is viable. In the event of a network fault, if a secondary RSVP-TE LSP is configured or the routing topology changes such that there is an alternate LDP LSP, the static PW will automatically switch LSPs in order to maintain connectivity with the PW endpoint. Static LSPs can be protected proactively by configuring BFD to verify the static LSPs IP next hop connectivity. Optionally, the underlying LSP for the PW can be explicitly specified using a named LSP. When a named LSP is explicitly specified, only the specified named LSP is used to carry the PW. In the event that a specified named LSP is withdrawn, the VPLS/VPWS remains operationally down until the named LSP is restored.

Since VC Status signaling is not supported, the VC Status "standby" bit cannot be used to allow support for PW redundancy and H-VPLS. Consequently, only "core full-mesh" PWs are allowed to have statically configured labels.

Example

The following command configures a new MPLS ethernet pseudowire for vpls1 :

```
configure vpls vpls1 add peer 1.1.1.202
```



History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure l2vpn vpws add peer ipaddress

```
configure l2vpn vpws vpws_name add peer ipaddress ipaddress {static-pw transmit-label outgoing_pw_label receive-label incoming_pw_label }
```

Description

Configures L2VPN VPWS service over MPLS Static PW.

Syntax Description

<i>vpws_name</i>	Specifies the VPWS for which you are configuring a peer.
ipaddress	Specifies the peer IP address.
<i>ipaddress</i>	Specifies the IP address of the peer node.
static-pw transmit-label	Specifies the static pseudowire transmit label.
<i>outgoing_pw_label</i>	Specifies the name of the egress label.
receive-label	Specifies the static PW receive label.
<i>incoming_pw_label</i>	Specifies the name of the ingress label.

Default

N/A.

Usage Guidelines

Use this command to statically configure a new MPLS Ethernet PW for the specified VPWS. You must specify the outgoing (MPLS ingress) and incoming (MPLS egress) PW labels. Similarly, you must configure the peer with a static PW that has the reverse PW label mappings.

Locally, the *incoming_pw_label* must be unique and is allocated out of the static label space. The *outgoing_pw_label* must match the peer's configured incoming PW label.

Just like a signaled PW, a static PW can optionally be configured to use any type of tunnel LSP: LDP, RSVP-TE, or Static. In the case of RSVP-TE and LDP, those protocols must be configured and enabled and an LSP must be established before traffic can be transmitted over the static PW.



For Static LSPs, only the MPLS ingress LSP (or outgoing LSP) is specified. Unlike signaled PWs, there is no end-to-end PW communication that is used to verify that the PW endpoint is operational, and in the case of static LSPs, that the data path to the PW endpoint is viable. In the event of a network fault, if a secondary RSVP-TE LSP is configured or the routing topology changes such that there is an alternate LDP LSP, the static PW will automatically switch LSPs in order to maintain connectivity with the PW endpoint. Static LSPs can be protected proactively by configuring BFD to verify the static LSPs IP next hop connectivity. Optionally, the underlying LSP for the PW can be explicitly specified using a named LSP. When a named LSP is explicitly specified, only the specified named LSP is used to carry the PW. In the event that a specified named LSP is withdrawn, the VPLS/VPWS remains operationally down until the named LSP is restored.

Since VC Status signaling is not supported, the VC Status “standby” bit cannot be used to allow support for PW redundancy and H-VPLS. Consequently, only “core full-mesh” PWs are allowed to have statically configured labels.

Example

The following command configures VPWS service for VPWS1 on peer 1.1.1.202:

```
configure vpws vpws1 add peer 1.1.1.202
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure l2vpn vpws peer static-pw

```
configure l2vpn vpws vpws_name peer ipaddress static-pw {transmit-label  
outgoing_pw_label receive-label incoming_pw_label }
```

Description

Changes the labels of a statically configured Ethernet pseudowire for a VPWS.

Syntax Description

<i>vpws_name</i>	Specifies the VPWS for which you are configuring a peer.
peer	Specifies the peer IP address.
<i>ipaddress</i>	Specifies the IP address of the peer node.
static-pw	Specifies the static pseudowire.



transmit label	Specifies the pseudowire transmit label.
<i>outgoing_pw_label</i>	Specifies the name of the egress label.
receive-label	Specifies the static pseudowire receive label.
<i>incoming_pw_label</i>	Specifies the name of the ingress label.

Default

N/A.

Usage Guidelines

Use this command to change the labels of a statically configured Ethernet pseudowire for a VPWS that already exists. Either or both the outgoing (MPLS ingress) and incoming (MPLS egress) PW labels can be specified. The peer must be similarly configured with a static PW that has the reverse PW label mappings. Locally, the *incoming_pw_label* must be unique and is allocated out of the static label space. The *outgoing_pw_label* must match the peer's configured incoming PW label. The CES or L2VPN can remain operational during the change; however, the PW will go down and come back up.

Example

The following command changes the VPWS name to "vpws1":

```
configure l2vpn vpws vpws1 peer 1.1.1.202
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure mpls add vlan

```
configure mpls add {vlan} vlan_name
```

Description

Adds an MPLS interface to the specified VLAN.

Syntax Description

<i>vlan_name</i>	Identifies the VLAN where the MPLS interface is added.
------------------	--



Default

VLANs are not configured with an MPLS interface.

Usage Guidelines

An MPLS interface must be configured on a VLAN in order to transmit or receive MPLS packets on that interface. By default, MPLS, LDP, and RSVP-TE are disabled for the MPLS interface. The specified VLAN should have an IP address configured and should have IP forwarding enabled. The MPLS interface on the VLAN does not become active until these two conditions are met. Also, if the IP address is unconfigured from the VLAN or IP forwarding is disabled for the VLAN, the MPLS interface goes down. The MPLS interface state is viewed using the `show mpls interface` command.

The VLAN must be operational for the MPLS interface to be up. This means that at least one port in the VLAN must be active or the VLAN must be enabled for loopback mode.

It is recommended that when you configure MPLS on an OSPF interface that can be used to reach a given destination, you should configure MPLS on all OSPF interfaces that can be used to reach that destination. (You should enable MPLS on all of the VLANs connected to the backbone network).

Example

The following example adds MPLS to the VLAN `vlan_usa`:

```
configure mpls add vlan vlan_usa
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mpls delete vlan

```
configure mpls delete [{vlan} vlan_name | vlan all]
```

Description

Removes an MPLS interface from the specified VLAN.



Syntax Description

<i>vlan_name</i>	Identifies the VLAN for which the MPLS interface is deleted.
vlan all	Deletes the MPLS interface from all VLANs that have MPLS configured.

Default

VLANs are not configured with an MPLS interface.

Usage Guidelines

An MPLS interface must be configured on a VLAN in order to transmit or receive MPLS packets on that interface. If the MPLS interface is deleted, all configuration information associated with the MPLS interface is lost. Issuing this command brings down all LDP neighbor sessions and all LSPs that are established through the specified VLAN interface. When the all VLANs option is selected, the MPLS interface for all MPLS configured VLANs is deleted.

Example

The following example deletes MPLS from the VLAN `vlan_k`:

```
configure mpls delete vlan vlan_k
```

History

This command was first available in ExtremeXOS 11.6

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mpls exp examination

```
configure mpls exp examination {value} value {qosprofile} qosprofile
```

Description

Configures the QoS profile that is used for the EXP value when EXP examination is enabled.

Syntax Description

<i>value</i>	Specifies the value that is used for the EXP value.
<i>qosprofile</i>	Specifies the QoS profile that is used for the EXP value.



Default

The QoS profile matches the EXP value + 1.

Usage Guidelines

This command configures the QoS profile that is used for the EXP value when EXP examination is enabled. By default, the QoS profile matches the EXP value + 1. That is, EXP value of 0 is mapped to QoS profile qp1, EXP value of 1 is mapped to QoS profile qp2, etc. This configuration has switch-wide significance. The EXP value must be a valid number from 0 through 7 and the qosprofile must match one of the switch's QoS profiles.



Note

EXP examination must be enabled using the “enable mpls exp examination” command before the configured EXP value to QoS profile mapping is actually used to process packets.

Example

The following command sets QoS profile q5 to be used for EXP value 7:

```
configure mpls exp examination value 7 qosprofile 5
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure mpls exp replacement

```
configure mpls exp replacement {qosprofile} qosprofile {value} value
```

Description

Configures the EXP value that is used for the specified QoS profile when EXP replacement is enabled.

Syntax Description

<i>qosprofile</i>	Specifies the QoS profile that is used for the EXP value.
<i>value</i>	Specifies the value that is used for the EXP value.



Default

The EXP value matches the QoS profile -1.

Usage Guidelines

This command configures the EXP value that is used for the QoS profile when EXP replacement is enabled. By default, the EXP value matches the QoS profile - 1. That is, QoS profile qp1 is mapped to EXP value of 0, QoS profile qp2 is mapped to EXP value of 1, etc. This configuration has switch-wide significance. The qosprofile must match one of the switch's QoS profiles and the EXP value must be a valid number from 0 through 7.



Note

EXP replacement must be enabled using the “enable mpls exp replacement” command before the configured EXP value to QoS profile mapping is actually used to process packets.

Example

The following command sets EXP value 2 to be used with QoS profile 4:

```
configure mpls exp replacement qosprofile qp4 value 2
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure mpls labels max-static

```
configure mpls labels max-static max_static_labels
```

Description

Configures the number of labels that are reserved for specifying the incoming label for static LSPs and static pseudowires.



Syntax Description

labels	Specifies that labels are reserved to specify the incoming label for static LSPs and static pseudowires.
max-static	Specifies the number of labels that are reserved to specify the incoming label for static LSPs and static PWs.
<i>max-static-labels</i>	Specifies the value for the maximum number of static labels.

Default

The default static label range size is 100.

Usage Guidelines

Use this command to configure the number of labels that are reserved for specifying the incoming label for static LSPs and static PWs. The static label range generally starts at 16 and the default static label range size is 100. This means that the default static label range is 16 through 115 and can be allocated for either incoming (both transit and egress) static LSPs, or incoming static PWs. The maximum static label_range_size is equal to the incoming label table size - 100 labels for signaling. For example, the Summit X480 has a hardware incoming label capacity of 8176 (8k-16) labels for egress LSPs. 960 labels are reserved for L3VPNs. The maximum number of labels available for static configuration is 7116, since at least 100 of those labels are reserved for dynamic signaling.

Since these values vary per-platform, use the `show mpls label usage` command to see details about label usage and platform capability. The minimum static label range size is 0.

Note



MPLS must be disabled when issuing this command. If MPLS is enabled, an error message is displayed and the command has no affect. All other labels, including outgoing labels for static LSPs and PWs and signaled labels used by RSVP-TE and LDP, are allocated out of the dynamic label space.

Example

The following example illustrates how to configure MPLS max-static labels, and how to display them:

```
Summit1.2 # show mpls lab usage
```

```
Label Type           Size      Label Range
-----
Supported            1048576   0x00000 - 0xfffff (0 - 1048575)
Reserved              16        0x00000 - 0x0000f (0 - 15)
Static                100       0x00010 - 0x00073 (16 - 115)
L3VPN                 960       0x00074 - 0x00433 (116 - 1075)
Dynamic               7116     0x00434 - 0x01fff (1076 - 8191)
Internal Use          0         0x00000 - 0x00000 (0 - 0)
...
```

```
Summit1.3 # disable mpls
```

```
* Summit1.4 # conf mpls lab max-static 7117
```



Error: There must be at least 100 dynamic labels remaining for MPLS signalling protocols.

```
* Summit1.5 # conf mpls lab max-static 7116
```

```
* Summit1.6 # show mpls lab usage
```

Label Type	Size	Label Range
Supported	1048576	0x00000 - 0xfffff (0 - 1048575)
Reserved	16	0x00000 - 0x0000f (0 - 15)
Static	7116	0x00010 - 0x01bdb (16 - 7131)
L3VPN	960	0x01bdc - 0x01f9b (7132 - 8091)
Dynamic	100	0x01f9c - 0x01fff (8092 - 8191)
Internal Use	0	0x00000 - 0x00000 (0 - 0)
...		

History

This command was first available in ExtremeXOS 15.4

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure mpls ldp advertise

```
configure mpls ldp advertise [{direct [all | lsr-id | none]} | {rip [all | none]} | {static [all | none]}]
```

Description

Configures a filter to be used by LDP when originating unsolicited label mapping advertisements to LDP neighbors.

Syntax Description

direct	Specifies that the advertisement filter is applied to the associated FECs with directly-attached routing interfaces.
rip	Specifies that the advertisement filter is applied to FECs associated with RIP routes exported by OSPF.
static	Specifies that the advertisement filter is applied to FECs associated with static routes.
all	Specifies that unsolicited label mapping advertisements are originated for all routes of the specified type.
lsr-id	Specifies that an unsolicited label advertisement is originated for a direct route that matches the MPLS LSR ID.
none	Specifies that no unsolicited label mapping advertisements are originated for the specified route type.



Default

None—the default setting for RIP and static routing methods.

lsr-id—the default setting for direct routes.

Usage Guidelines

You can configure how the advertisement filter is applied, as follows:

- **direct**—The advertisement filter is applied to the FECs associated with directly-attached routing interfaces.
- **rip**—The advertisement filter is applied to the FECs associated with RIP routes exported by OSPF.
- **static**—The advertisement filter is applied to the FECs associated with static routes.

You can configure the advertisement filter, as follows:

- **all**—Label mappings are originated for all routes of the specified type.
- **none**—No label mappings are originated for all routes of the specified type. This is the default setting for RIP and static routes.
- **lsr-id**—A label mapping is originated for a direct route that matches the MPLS LSR ID. This is the default setting for direct routes.

Advertising labels for a large number of routes may increase the required number of labels that must be allocated by LSRs. Take care to ensure that the number of labels advertised by LERs does not overwhelm the label capacity of the LSRs.

Example

The following command configures LDP to originate labels for all local IP interfaces:

```
configure mpls ldp advertise direct all
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mpls ldp loop-detection

```
configure mpls ldp loop-detection [{hop-count hop_count_limit} {path-vector path_vector_limit}]
```



Description

Configures the loop-detection parameters used by LDP.

Syntax Description

hop-count	Configures the number of LSRs that the label message can traverse.
<i>hop_count_limit</i>	Specifies the hop count limit. The valid configuration range is from 1 to 255.
path-vector	Configures the maximum number of LSR IDs that can be propagated in the label message.
<i>path_vector_limit</i>	Specifies the path vector limit. The valid configuration range is from 1 to 255.

Default

The default for the hop-count and path-vector limits is 255.

Usage Guidelines

Configuration changes are only applicable to newly created LDP sessions. Disabling and enabling LDP forces all the LDP sessions to be recreated. LDP loop detection must first be enabled for these configuration values to be used.

Example

This command sets the LDP hop count loop detection value to 10. The configured path vector value remains at 255.

```
configure mpls ldp loop-detection hop-count 10
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mpls ldp pseudo-wire

```
configure mpls ldp pseudo-wire parm-mismatch-recovery ces [auto | none]
```

Description

Enables or disables automatic recovery from parameter mismatch.



Syntax Description

auto	Allow automatic recovery from parameter mismatch.
none	Disable automatic recovery from parameter mismatch.

Default

The default is auto.

Usage Guidelines

If interface parameters do not match during pseudowire signaling, the remote side may release our VC label which will prevent the pseudowire from coming up. If this is detected, the system will attempt automatic recovery to bring the pseudowire up if this configuration setting is auto. This automatic recovery can be disabled by setting the parameter mismatch recovery to none.

Example

To disable automatic recovery:

```
config mpls ldp pseudo-wire parm-mismatch-recovery ces none
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on the cell site routers (E4G-200 and E4G-400).

configure mpls ldp timers

```
configure mpls ldp timers [targeted | link] [{hello-time hello_hold_seconds}
{keep-alive-time keep_alive_hold_seconds}]
```

Description

Configures LDP peer session timers for the switch.

Syntax Description

targeted	Specifies targeted LDP sessions.
link	Specifies link LDP sessions.



<i>hello_hold_seconds</i>	The amount of time (in seconds) that a hello message received from a neighboring LSR remains valid. The rate at which Hello messages are sent is 1/3 the configured hello-time. If a Hello message is not received from a particular neighboring LSR within the specified <i>hello_hold_seconds</i> , then the hello-adjacency is not maintained with that neighboring LSR. The range is 6 to 65,534 seconds.
<i>keep_alive_hold_seconds</i>	The time (in seconds) during which an LDP message must be received for the LDP session with a particular peer LSR to be maintained. If an LDP PDU is not received within the specified session <i>keep_alive_hold_seconds</i> , the corresponding LDP session is torn down. The range is 6 to 65,534 seconds.

Default

link *hello_hold_seconds* - 15 seconds

targeted *hello_hold_seconds* - 45 seconds

link *keep_alive_hold_seconds* - 40 seconds

targeted *keep_alive_hold_seconds* - 60 seconds

Usage Guidelines

The LDP peer hello-adjacency timers are separately configurable for link and targeted LDP sessions. The hello timer parameter specifies the amount of time (in seconds) that a Hello message received from a neighboring LSR remains valid. The rate at which Hello messages are sent is 1/3 the configured hello-time. If a Hello message is not received from a particular neighboring LSR within the specified *hello_hold_seconds*, then the hello-adjacency is not maintained with that neighboring LSR.

The session *keep_alive_hold_seconds* parameter specifies the time (in seconds) during which an LDP message must be received for the LDP session to be maintained. The rate at which Keep Alive messages are sent, provided there are no LDP messages transmitted, is 1/6 the configured keep-alive-time. If an LDP PDU is not received within the specified session *keep_alive_hold_seconds* interval, the corresponding LDP session is torn down. The minimum and maximum values for hold timers are 6 and 65,534, respectively.

Changes to targeted timers only affect newly created targeted sessions. Disabling and then enabling VPLS or LDP causes all current targeted sessions to be re-created. The default values for the various times are as follows: link *hello_hold_seconds* (15), link *keep_alive_hold_seconds* (40), targeted *hello_hold_seconds* (45), and targeted *keep_alive_hold_seconds* (60). Changes to the link keep-alive timers do not take effect until the LDP session is cycled.

Example

The following command configures link-level LDP hello adjacency hold time to 30 seconds and the keep alive time to 10 seconds:

```
configure mpls ldp timers link hello-time 30 keep-alive-time 10
```



History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mpls lsr-id

```
configure mpls lsr-id ipaddress
```

Description

Configures the MPLS LSR ID for the switch.

Syntax Description

<i>ipaddress</i>	Specifies an IP address to identify the MPLS LSR for the switch. The MPLS LSR-ID should be configured to the same IP address as the OSPF Router ID.
------------------	---

Default

No LSR ID is configured by default.

Usage Guidelines

LDP, RSVP-TE, and L2 VPNs all use the LSR ID. It is normally set to the OSPF Router ID.

The LSR ID must be configured before MPLS can be enabled. The LSR ID cannot be changed while MPLS is enabled. It is highly recommended that an IP address be configured on a OSPF enabled loopback VLAN that matches the configured LSR ID and OSPF ID. If an LSR ID loopback IP address is configured, OSPF automatically advertises the LSR ID as a routable destination for setting up LSPs. The LSR ID remains active if an interface goes down if the LSR-ID is configured as an IP address on a loopback VLAN, as recommended. This significantly enhances network stability and operation of an MPLS network.

Example

The following command configures the LSR ID to 192.168.50.5:

```
configure mpls lsr-id 192.168.50.5
```



History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mpls rsvp-te bandwidth committed-rate

```
configure mpls rsvp-te bandwidth committed-rate committed_bps [Kbps | Mbps | Gbps] [{vlan} vlan_name | vlan all] {receive | transmit | both}
```

Description

Specifies the maximum amount of Committed Information Rate (CIR) bandwidth which can be used by RSVP-TE LSP reservations.

Syntax Description

<i>committed_bps</i>	Specifies a bitrate for the bandwidth to be reserved.
Kbps	Specifies the designated bitrate in kilobits per second.
Mbps	Specifies the designated bitrate in megabits per second.
Gbps	Specifies the designated bitrate in gigabits per second.
vlan	Specifies that the bandwidth is to be reserved for a specific VLAN.
<i>vlan_name</i>	Identifies the VLAN for which the bandwidth is reserved.
vlan all	Specifies that the bandwidth is reserved for all VLANs that have MPLS configured.
receive	Specifies that the bandwidth is reserved for ingress traffic only.
transmit	Specifies that the bandwidth is reserved for egress traffic only.
both	Specifies that the bandwidth is reserved for both ingress and egress traffic.

Default

The default is zero, which means no RSVP-TE LSP bandwidth reservations are accepted.

If bandwidth is specified without specifying traffic direction, the default is both directions.

Usage Guidelines

This command specifies the maximum amount of Committed Information Rate (CIR) bandwidth which can be used by dynamic RSVP-TE LSP bandwidth reservations. By sub-allocating reserveable



bandwidth for RSVP-TE from the VLAN's available bandwidth, the switch can guarantee that as LSPs are established, a minimum amount of CIR bandwidth is available for other traffic.

Note



Beginning with ExtremeXOS Release 12.2.1, CIR bandwidth for the receive direction is not tracked by TE IGPs, such as OSPF-TE, and configuring it is not required. Configuring CIR bandwidth for the receive direction does not prevent an LSP from going operational due to lack of receive bandwidth; however, it can be useful for tracking and informational purposes. An Info level log (MPLS.RSVPTTE.IfRxBwdthExcd) is generated if the setup of a TE LSP requires receive bandwidth greater than that which is currently available for the receive direction on a particular interface. This generally happens only when TE LSPs with different previous hops ingress the switch on the same interface (for example, from a multi-access link) and egress the switch on different interfaces.

The keyword `both` configures the reserved bandwidth for both ingress and egress LSP CIR reservations and overwrites any previous receive or transmit settings.

Example

The following command reserves 25 Mbps of CIR bandwidth for all RSVP-TE CIR reservations on the specified VLAN:

```
configure mpls rsvp-te bandwidth committed-rate 25 Mbps vlan vlan_10
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mpls rsvp-te lsp add path

```
configure mpls rsvp-te lsp lsp_name add path [path_name | any] {profile profile_name} {primary {frr_profile_name} | secondary}
```

Description

Adds a configured path to the specified RSVP-TE LSP.



Syntax Description

<i>lsp_name</i>	Specifies the name of the LSP you are configuring.
<i>path_name</i>	Specifies the name of the path to be used by the specified LSP.
any	Configures the specified LSP to use any path.
<i>profile_name</i>	Specifies a profile to be applied to the specified LSP. If the profile name is omitted, the profile named default is used.
primary	Designates the specified path as the primary path. Only one primary path can be configured for an RSVP-TE LSP. If this option is omitted and no primary path has been specified, the specified path is added as a primary path. If not specified and a primary path has already been added, the path is added as a secondary path.
secondary	Designates the specified path as a secondary path.
<i>frr_profile_name</i>	Specifies a fast reroute (FRR) profile to be applied to the detour LSP that backs up the specified LSP.

Default

N/A.

Usage Guidelines

The LSP is not signaled until a path is added to the LSP.

If you want fast reroute protection for the LSP, use the primary option and specify the fast reroute profile name you want to use. To specify the default fast reroute profile, enter default-frr.

The switch chooses the local MPLS VLAN interface from which to signal the LSP. To force an LSP to use a specific local MPLS interface, configure the local interface IP address as the first ERO in the associated path.

Example

This command adds the path sydney-bypass to the LSP named aus as a secondary path:

```
configure mpls rsvp-te lsp aus add path sydney-bypass secondary
```

History

This command was first available in ExtremeXOS 11.6.

The fast reroute capability was added in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure mpls rsvp-te lsp change

```
configure mpls rsvp-te lsp lsp_name change [path_name | any] use profile
[{standard_profile_name} {frr_profile_name}]
```

Description

Changes the configuration that has been configured with the `configure mpls rsvp-te lsp lsp_name add path [path_name | any] {profile_name} {primary {frr_profile_name} | secondary}` command.

Syntax Description

<i>lsp_name</i>	Specifies the name of the LSP you are changing.
<i>path_name</i>	Specifies the name of the path to be used by the specified LSP.
any	Configures the specified LSP to use any path.
<i>standard_profile_name</i>	Specifies a profile to be applied to the specified LSP. If the profile name is omitted, the profile named default is used.
<i>frr_profile_name</i>	Specifies a fast reroute (FRR) profile to be applied to the detour LSP that backs up the specified LSP.

Default

N/A.

Usage Guidelines

None.

Example

This command changes the LSP named `aus` to use any available path:

```
configure mpls rsvp-te lsp aus change any
```

History

This command was first available in ExtremeXOS 11.6.

The fast reroute capability was added in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.



configure mpls rsvp-te lsp delete path

```
configure mpls rsvp-te lsp lsp_name delete path [path_name | any | all]
```

Description

Deletes a path from the specified RSVP-TE LSP.

Syntax Description

<i>lsp_name</i>	Specifies a name for the RSVP-TE LSP.
<i>path_name</i>	Specifies a name for the path to be deleted from the RSVP-TE LSP.
any	Configures the specified LSP to use any path.
all	Deletes all added paths from the specified RSVP-TE LSP.

Default

N/A.

Usage Guidelines

This command deletes a path from the specified RSVP-TE LSP. All the added paths can be deleted by specifying the all keyword. If the active path is deleted, then one of the other configured paths becomes the active path for the LSP. If there are no other defined paths, then the LSP is marked down and cannot be used to forward IP or VPN traffic.

Example

The following command deletes the path called through-knightsbridge for the LSP london:

```
configure mpls rsvp-te lsp london delete path through-knightsbridge
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure mpls rsvp-te lsp fast-reroute

```
configure mpls rsvp-te lsp lsp_name fast-reroute [enable | disable]
```



Description

Enables or disables fast-reroute protection for the specified LSP.

Syntax Description

<code>lsp_name</code>	Specifies the name of the LSP you are configuring.
-----------------------	--

Default

Disabled.

Usage Guidelines

To signal the fast-reroute protected LSP, use the `enable mpls rsvp-te lsp [lsp_name | all]` command. Similarly, to disable the fast-reroute protected LSP, use the `disable mpls rsvp-te lsp [lsp_name | all]` command.

Example

This command enables fast-reroute protection on LSP aus:

```
configure mpls rsvp-te lsp aus fast-reroute enable
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure mpls rsvp-te lsp path use profile

```
configure mpls rsvp-te lsp lsp_name path [path_name | any] use profile profile_name
```

Description

Changes the profile that the configured LSP path uses.



Syntax Description

<i>lsp_name</i>	Specifies the RSVP-TE LSP.
<i>path_name</i>	Specifies the configured RSVP-TE LSP path.
<i>profile_name</i>	Specifies a profile to be applied to the configured LSP path.

Default

N/A.

Usage Guidelines

This command changes the profile that the configured LSP path uses.



Note

Changing the profile while an LSP is active may cause the LSP to be torn down and re-signaled.

Example

The following command configures the switch to apply the LSP profile gold-class to the LSP path sydney-bypass for the LSP aus:

```
configure mpls rsvp-te lsp aus path sydney-bypass use profile gold-class
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure mpls rsvp-te lsp transport

```
configure mpls rsvp-te lsp lsp_name transport [ip-traffic [allow | deny] | vpn-traffic [allow {all | assigned-only} | deny]]
```

Description

Configures the type of traffic that may be transported across a named LSP.



Syntax Description

<i>lsp_name</i>	Specifies the RSVP-TE LSP.
ip-traffic	Controls the forwarding of routed IP traffic across the specified LSP.
vpn-traffic	Controls the forwarding of VPN traffic over the LSP.
allow	Allows transport of the specified traffic across the LSP.
deny	Denies transport of the specified traffic across the LSP.
allow	Allows all VPLS VPN traffic to be transported across the LSP.
all	Allows the transmission of all VPN traffic over the LSP.
assigned-only	Limits the transport of VPN traffic to VPLS instances that are explicitly configured to use the specified LSP name.

Default

The default behavior is to allow RSVP-TE LSPs to transport all types of traffic without restriction.

Usage Guidelines

This command configures the type of traffic that may be transported across a named LSP. By default, both IP traffic and VPN traffic are set to allow transport for a newly created LSP. The `ip-traffic` keyword is used to allow or deny forwarding of routed IP traffic across the specified LSP. If allowed, the LSP label information is inserted into the routing table and the switch forwards traffic over the LSP that matches the IP route entry to which this LSP is associated. If denied, the LSP label information is removed from the routing table and the switch does not use the LSP to transport IP traffic. The `vpn-traffic` keyword controls the transmission of VPN traffic over the LSP. When denied, the LSP is not used as a transport for PWs or other VPN related traffic. These transport configuration options are independent. For example, if `vpn-traffic` is set to allow and `ip-traffic` is set to deny, then no routed IP traffic is transported across the LSP, but the LSP may still be used to transport VPN traffic.

The optional `assigned-only` keyword limits the transport of VPN traffic to only those VPLS instances that are explicitly configured to use the specified LSP name.

Example

The following command prevents the switch from using LSP `aus` to forward IP traffic:

```
configure mpls rsvp-te lsp aus transport ip-traffic deny
```

History

This command was first available in ExtremeXOS 11.6.



Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mpls rsvp-te metric

```
configure mpls rsvp-te metric [value | use-igp] {vlan} vlan_name
```

Description

Configures the TE metric value for the RSVP-TE interface specified by the *vlan_name* argument.

Syntax Description

<i>value</i>	Specifies a value for the RSVP-TE metric.
vlan	Specifies that the RSVP-TE metric is configured for a specific VLAN.
<i>vlan_name</i>	Identifies the VLAN for which the RSVP-TE metric is configured.

Default

The associated default IGP metric.

Usage Guidelines

The TE metric can be any unsigned non-zero 32-bit integer. The default value for the RSVP-TE interface is to use the associated default IGP metric. The TE metric is exchanged between OSPF routers and is used in the calculation of the CSPF topology graph.

Example

The following command configures an RSVP-TE metric of 220 on the specified VLAN:

```
configure mpls rsvp-te metric 220 vlan vlan_10
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure mpls rsvp-te path add ero

```
configure mpls rsvp-te path path_name add ero ipNetmask [strict|loose] {order
number}
```

Description

Adds an IP address to the Explicit Route Object (ERO) for the specified path name.

Syntax Description

<i>path_name</i>	Specifies the path to which the IP address is added.
<i>ipNetmask</i>	Specifies an IP prefix.
strict	Specifies that the subobject must be topologically adjacent to the previous subobject in the ERO list.
loose	Specifies that the subobject need not be topologically adjacent to the previous subobject in the ERO list.
<i>number</i>	Specifies the LSR path order.

Default

The order value defaults to 100 if the path has no EROs configured or a value 100 more than the highest order number configured for the path.

Usage Guidelines

This command adds an IP address to the Explicit Route Object (ERO) for the specified path name. The RSVP-TE routed path may be described by a configured sequence of the LSRs and/or subnets that the path traverses. Each defined LSR or subnet represents an ERO subobject. Up to 64 subobjects can be added to each path name. The ERO keyword identifies an LSR using an IP prefix, which may represent an LSR's Router ID, loopback address, or direct router interface. Each IP prefix is included in the ERO as an IPv4 subobject.

If the ERO is specified as strict, the strict subobject must be topologically adjacent¹¹ to the previous subobject as listed in the ERO. If the ERO is specified as loose, the loose subobject is not required to be topologically adjacent to the previous subobject as listed in the ERO. If the specified IP prefix matches the OSPF router ID or a configured loopback IP address, the ERO must be configured as loose.

The LSR path order is optionally specified using the order keyword. The order number parameter is an integer value from 1 to 65535. IP prefixes with a lower number are sequenced before IP prefixes with a higher number. Thus, the LSP path follows the configured path of IP prefixes with a number value from low to high. If the order keyword is not specified, the number value for the LSR defaults to a value equal to the current highest number value plus 100. If the list of IP prefixes added to the path does not reflect an actual path through the network topology, the path message is returned with an error from a downstream LSR and the LSP is not established.

¹¹ "Topologically adjacent" indicates that the router next hop matches either the interface IP address or OSPF router ID of an immediate peer LSR.



The order of a configured subobject cannot be changed. The ERO subobject must be deleted and re-added with a different order. If a subobject is added to or deleted from the ERO while the associated LSP is established, the path is torn down and is re-signaled using the new ERO. Duplicate ERO subobjects are not allowed.

Defining an ERO for the path is optional. If no ERO is configured, the path is signaled along the best available path and the ERO is not included in the path message. When the last subobject in the ERO of the path message is reached and the egress IP node of the path has not been reached, the remaining path to the egress node is signaled along the best available path. If the next subobject in the ERO is loose, the best available path to the next subobject is chosen. Configuring EROs could lead an LSP to take an undesirable path through the network, so care should be taken when specifying EROs.

Example

The following example adds the IP interface address 197.57.30.7/24 as a loose ERO to the path sydney-bypass:

```
configure mpls rsvp-te path sydney-bypass add ero 197.57.30.7/24 loose
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mpls rsvp-te path delete ero

```
configure mpls rsvp-te path path_name delete ero [all | ipNetmask | order number]
```

Description

Deletes a subobject from the Explicit Route Object (ERO) for the specified path name.

Syntax Description

<i>path_name</i>	Specifies the path from which the ERO is deleted.
all	Specifies that the entire ERO should be deleted from the named path.
<i>ipNetmask</i>	Specifies the ERO subobject to be deleted.
<i>number</i>	Specifies the order number of the ERO subobject to be deleted.



Default

N/A.

Usage Guidelines

This command deletes a subobject from the Explicit Route Object (ERO) for the specified path name. The ERO subobject is specified using an IP prefix or order number. If a subobject is deleted from an ERO while the associated LSP is established, the path is torn down and is re-signaled using a new ERO. The all keyword may be used to delete the entire ERO from the path name. When there is no configured ERO, the path is no longer required to take an explicit routed path. The path is then signaled along the best available path and no ERO is included in the path message.

Example

The following command deletes all the configured EROs from the path sydney-bypass:

```
configure mpls rsvp-te path sydney-bypass delete ero all
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure mpls rsvp-te profile

```
configure mpls rsvp-te profile profile_name {bandwidth [best-effort |
[committed-rate committed_bps [Kbps | Mbps | Gbps]] {max-burst-size burst_size
[Kb | Mb]} {peak-rate peak_bps [Kbps | Mbps | Gbps]]}] {hold-priority
hold_priority} {mtu [number | use-local-interface]} {record [enabled {route-only}
| disabled]} {setup-priority setup_priority}{[Cosprofile cosprofile | default-
cosprofile]}
```

Description

Configures an RSVP-TE profile with the specified profile name.

Syntax Description

<i>profile_name</i>	Specifies the LSP profile.
bandwidth	Specifies bandwidth reservation.
best-effort	Indicates no bandwidth reservation.



<i>committed_bps</i>	Specifies the committed bandwidth to be reserved across the MPLS network, in bits per second. The range is from 64 Kbps to 10 Gbps.
<i>peak_bps</i>	Specifies the maximum bandwidth signaled in bits per second. The range is from 64 Kbps to 10 Gbps.
Kbps	Specifies the designated bitrate in kilobits per second.
Mbps	Specifies the designated bitrate in megabits per second.
Gbps	Specifies the designated bitrate in gigabits per second.
<i>burst_size</i>	Specifies the maximum number of bytes (specified in bits) that the LSP is allowed to burst above the specified peak-rate. The range is from 0 to 1000 Mb.
Kb	Kilobits
Mb	Megabits
<i>hold_priority</i>	Specifies the priority of the LSP. Lower numbers indicate higher priority. The range is from 0 to 7.
<i>setup_priority</i>	Specifies the priority of the LSP. Lower numbers indicate higher priority. The range is from 0 to 7.
<i>number</i>	Specifies the MTU value for the LSP. The range is from 296 to 9216/
use-local-interface	Specifies that the MTU value is inherited from the local egress VLAN interface.
record	Configures hop-by-hop path recording.
enabled route-only	Causes the Record Route Object (RRO) to be inserted into the path message. The enabled option enables recording of hops and labels. The enabled route-only option records only hops.
disabled	Specifies that no RRO is inserted into the path message.
cosprofile	This command configures the Class of Service for a RSVP-TE LSP, and specifies the name of the COS profile CP1, CP2, ... CP8. Default is CP1 and represents Best Effort class of Service. CP8 is the highest COS profile.

Default

Bandwidth: Newly-created profiles are configured as best-effort.

Setup-priority: 7 (lowest).

Hold-priority: 0 (highest).

Path recording: disabled.

MTU: use-local-interface.

If the **Cosprofile** is not specified for an LSP, it is assumed to be in the default CosProfile CP1.

Usage Guidelines

A profile is a set of attributes that are applied to the LSP when the LSP is configured using the `configure mpls rsvp-te lsp` command. A default profile is provided which cannot be deleted, but may be



applied to any TE LSP. The *profile_name* for the default profile is default. The default profile parameter values are initially set to their respective default values. The maximum number of configurable profiles is 1000.

LSPs may signal reserved bandwidth. By default, newly created profiles are configured to not signal bandwidth requirements and thus are classified as best-effort. If bandwidth needs to be reserved across the MPLS network, the bandwidth parameters specify the desired reserved bandwidth for the LSP. The committed-rate specifies the mean bandwidth and the peak-rate specifies the maximum bandwidth signaled. The peak-rate must be equal to or greater than the committed-rate. If the peak-rate is not specified, traffic is not clipped above the committed-rate setting. The rates are specified in bps and must be qualified by Kbps, Mbps, or Gbps. The minimum and maximum bandwidth rates are 64 Kbps and 10 Gbps, respectively. The max-burst-size specifies the maximum number of bytes (specified in bits) that the LSP is allowed to burst above the specified peak-rate. The minimum burst size is 0 and the maximum burst size is 1000 Mb.

The setup-priority and hold-priority are optional parameters indicating the LSP priority. During path set up, if the requested bandwidth cannot be reserved through the LSR, the setup-priority parameter is compared to the hold-priority of existing LSPs to determine if any of the existing LSPs need to be preempted to allow a higher priority LSP to be established. Lower numerical values represent higher priorities. The setup-priority range is 0 to 7 and the default value is 7 (lowest). The hold-priority range is also 0 to 7 and the default value is 0 (highest). If bandwidth is requested for the LSP, the CSPF calculation uses the available bandwidth associated with the CoS as specified by the hold-priority.

The bandwidth, hold-priority, and setup-priority values are signaled in the path message. If the bandwidth setting is changed, all LSPs using this profile are re-signaled. If the bps setting is decreased, a new path message is sent along the LSP indicating the new reservation. If the bps setting is increased, the LSP is torn down and resignaled using the new bandwidth reservations.

The record command is used to enable hop-by-hop path recording. The enabled keyword causes the Record Route Object (RRO) to be inserted into the path message. The RRO is returned in the RESV Message and contains a list of IPv4 subobjects that describe the RSVP-TE path. Path recording by default is disabled. When disabled, no RRO is inserted into the path message.

The mtu keyword optionally specifies the MTU value for the LSP. By default, this value is set to use-local-interface. In the default configuration, the MTU value is inherited from the local egress VLAN interface. The minimum MTU value is 296 and the maximum value is 9216. Path MTU information is carried in the Integrated Services or Null Service RSVP objects and is used by RSVP to perform path MTU identification.

Note



Changing any of the profile parameters causes LSPs using the profile to be torn down and re-signaled. There is no guarantee that the re-signaled LSP will be successfully established. Future ExtremeXOS implementations may support the make-before-break LSP concept.

To view a profile configuration, enter the following command:

```
show mpls rsvp-te profile {profile_name} {detail}
```

To view LSP recorded route information, enter one of the following commands:

```
show mpls rsvp-te lsp [ingress {fast-reroute} | ingress_lsp_name |
ingressingress_lsp_name | ingress [destination | origin]ipaddress] {[all-paths |
```




```
detail] | summary | down-paths {detail}} show mpls rsvp-te lsp [egress | transit]
{fast-reroute} {{lsp_name} [[destination | origin]ipaddress} {detail} | summary}
```

Example

The following command configures the RSVP-TE profile gold-class with a committed bandwidth of 100 Mbps and the setup and hold priorities are both set to 0 (highest priority):

```
configure mpls rsvp-te profile gold-class bandwidth committed-rate 100 mbps
hold-priority 0 setup-priority 0
```

The following example configures the Cosprofile to CP3:

```
configure mpls rsvp-te profile gold-class bandwidth committed-rate 100 mbps
holdpriority 0 setup-priority 0 Cosprofile CP3
```

History

This command was first available in ExtremeXOS 11.6.

The **cosprofile** keyword was added int ExtremeXOS 15.3.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure mpls rsvp-te profile (fast-reroute)

```
configure mpls rsvp-te profile frr_profile_name {bandwidth bandwidth_rate_bps
bandwidth_rate_unit} {detour {hop-limit hop_limit_value} {bandwidth-protection
[enabled | disabled]} {node-protection [enabled | disabled]}} {hold-priority
hold_priority_value} {setup-priority setup_priority_value}
```

Description

Configures the specified RSVP-TE FRR profile.

Syntax Description

<i>frr_profile_name</i>	Specifies the FRR LSP profile to configure.
<i>bandwidth_rate_bps</i>	Specifies the bandwidth requirement for the FRR LSP. This should be set to match the options chosen for the protected LSP. Otherwise, a mismatch between the bandwidth settings for the detour and protected LSPs can impact service.



<i>bandwidth_rate_unit</i>	Specifies the units for the bandwidth rate. Valid entries are Kbps, Mbps, and Gbps.
detour	Specifies the detour method of fast reroute. This is the only method supported in this release.
<i>hop_limit_value</i>	Specifies the maximum number of hops that the detour path is allowed to take from the current node or point of local repair (PLR) to a merge point (MP) node. If set to 0, only link protection is provided.
bandwidth-protection	When enabled, this option specifies that the signaled bandwidth on the detour path must be guaranteed. If this option is disabled, the detour path might not support the bandwidth needed for the protected LSP.
node-protection	When enabled, the this option indicates to the PLRs along a protected path that a detour path that bypasses at least the next node of the protected LSP is desired. If this option is disabled, the backup path might or might not bypass the next node, in which case the user might or might not have next-node protection.
hold-priority	Specifies the hold priority of the LSP. Lower numbers indicate higher priority. The range is from 0 to 7. Hold priority is used when deciding whether a session can be preempted by another session. This works exactly the same as the hold-priority set in the standard profile that is valid for the protected LSP and for standard LSPs.
setup-priority	Specifies the setup priority of the LSP. Lower numbers indicate higher priority. The range is from 0 to 7. The setup priority is used when deciding whether the detour LSP can preempt another session. This works exactly the same as the setup-priority set in the standard profile that is valid for the protected LSP and standard LSPs.

Default

Bandwidth: Newly-created profiles are configured as best-effort. Setup-priority: 7 (lowest) Hold-priority: 0 (highest) Hop-limit: 3 Protect-bandwidth: enabled Protect-node: enabled

Usage Guidelines

A FRR profile is a set of attributes that are applied to the detour and protected LSPs when a protected LSP is configured. A default profile (*frr-default*) is provided which cannot be deleted, but can be applied to any protected LSP. The maximum number of configurable profiles is 1000.



Note

Changing any of the profile parameters causes LSPs using the profile to be torn down and re-signaled. There is no guarantee that the re-signaled LSP will be successfully established. Future ExtremeXOS implementations may support the make-before-break LSP concept.

Example

The following command configures the FRR profile *frrprofile* for 100 Mbps bandwidth:

```
configure mpls rsvp-te profile frrprofile bandwidth 100 Mbps
```



History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mpls rsvp-te timers lsp rapid-retry

```
configure mpls rsvp-te timers lsp rapid-retry {decay-rate percent} {delay-interval milliseconds} {retry-limit [number]}
```

Description

Configures the timers associated with rapidly retrying failed LSPs.

Syntax Description

<i>percent</i>	Specifies a percent increase in the interval allowed before each subsequent attempt to re-signal an LSP. The valid range is from 0 to 100 percent.
<i>milliseconds</i>	Specifies the time (in milliseconds) to wait before attempting to re-signal the LSP.
retry-limit	Specifies the maximum allowed attempts to establish an LSP.
<i>number</i>	Specifies a maximum number of allowed attempts to establish an LSP. The valid number range is from zero to 255.

Default

Delay interval: 500 milliseconds.

Decay rate: 50%.

Retry limit: 10.

Usage Guidelines

This command configures the timers associated with rapidly retrying failed LSPs. If an LSP fails to establish, the switch attempts to rapidly retry the setup by sending additional path messages based on the rapid-retry timers. The delay-interval timer specifies the time (in milliseconds) to wait before sending another path message. If the LSP fails to establish itself on subsequent attempts, the delay-interval time is incremented based on the decay-rate setting. The decay operation multiplies the delay-interval time by the decay rate, and adds the result to the current delay-interval time.

For example, if the decay-rate is set to 50 percent and the current delay-interval time is 500 milliseconds, a path message is retransmitted in 750 milliseconds. If the LSP fails to establish on the next attempt, a path message is retransmitted after a further decayed delay interval of 1125 milliseconds (1.125 seconds). A per-LSP delay-interval time is maintained for each LSP until the LSP is



established. This process of decaying the retry time continues until the LSP is established or the retry-limit expires. If the retry-limit is reached, attempts to rapidly retry the LSP are suspended.

When the switch starts the process of re-signaling the LSP based on the standard-retry timers, the LSP's rapid-retry timers return to the initial configuration settings. If the standard-retry delay-interval time is reached before all of the rapid-retry attempts have completed, the standard-retry mechanisms take over.

The default rapid-retry LSP timer parameter values are 500 milliseconds for the delay-interval, 50 percent for the decay-rate, and a retry-limit of 10. The valid range for delay-interval is 10 to 1000 milliseconds. The valid decay-rate range is 0 to 100 percent. The valid retry-limit is 0 to 100. A value of 0 indicates that the LSP is not re-signaled using the rapid-retry timers.

When summary-refresh or bundle-message is enabled, the rapid-retry timer values are used for resending any message that is not acknowledged.



Note

RSVP-TE must be disabled before these parameters can be modified.

Example

The following command sets the maximum number of rapid retries to five:

```
configure mpls rsvp-te timers lsp rapid-retry retry-limit 5
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mpls rsvp-te timers lsp standard-retry

```
configure mpls rsvp-te timers lsp standard-retry {decay-rate percent} {delay-interval seconds} {retry-limit [number | unlimited]}
```

Description

Configures the timers associated with the establishment of an LSP.



Syntax Description

<i>percent</i>	Specifies a percent increase in the interval allowed before each subsequent attempt to re-signal an LSP. The valid range is from 0 to 100 percent.
<i>seconds</i>	Specifies the time (in seconds) to wait before attempting to re-signal the LSP.
retry-limit	Specifies the maximum allowed attempts to establish an LSP.
<i>number</i>	Specifies a maximum number of allowed attempts to establish an LSP. The valid number range is from zero to 255.
unlimited	Allows unlimited attempts to establish an LSP.

Default

Delay interval: 30 seconds.

Decay rate: 0%.

Retry limit: unlimited.

Usage Guidelines

This command configures the timers associated with the establishment of an LSP. If an LSP fails to establish, the LSP is re-signaled based on the configuration of these timers. The delay-interval timer specifies the time (in seconds) to wait before attempting to re-signal the LSP. If the LSP fails to establish itself on subsequent attempts, the delay-interval time is incremented based on the decay-rate setting. The decay operation multiplies the delay-interval time by the decay rate, and adds the result to the current delay-interval time. For example, if the decay-rate is set to 50 percent and the current delay-interval time is 30 seconds, the LSP is re-signaled in 45 seconds. If the LSP failed to establish on the next attempt, the delay interval would be further decayed to 67 seconds.

A per-LSP delay-interval time is maintained for each LSP until the LSP is established. This operation of decaying the retry time continues until the LSP is established or the retry-limit expires. If the retry-limit is reached, attempts to establish the LSP are suspended.

Disabling and enabling the LSP resets the LSP's delay-interval time and retry-limit to the initial configuration settings and LSP establishment attempts resume. The default LSP timer parameter values are 30 seconds for delay-interval, with a 0 percent decay-rate, and retry-limit of unlimited. The valid range for delay-interval is 1 to 60 seconds. The valid decay-rate range is 0 to 100 percent. The valid retry-limit is 0 to 255 or unlimited. A value of 0 indicates that the LSP is not re-signaled.



Note

RSVP-TE must be disabled before these parameters can be modified.

Example

The following command allows unlimited retries for establishing MPLS RSVP-TE LSPs:

```
configure mpls rsvp-te timers lsp standard-retry retry-limit unlimited
```



History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mpls rsvp-te timers session

```
configure mpls rsvp-te timers session [{bundle-message-time
bundle_message_milliseconds} {hello-keep-multiplier hello_keep_number} {hello-
time hello_interval_seconds} {refresh-keep-multiplier refresh_keep_number}
{refresh-time refresh_seconds} {summary-refresh-time
summary_refresh_milliseconds}] [{vlan} vlan_name | vlan all]
```

Description

Configures the RSVP-TE protocol parameters for the specified VLAN.

Syntax Description

<i>bundle_message_milliseconds</i>	Specifies the maximum time a transmit buffer is held to allow multiple RSVP messages to be bundled into a single PDU. The valid range is from 50 to 3000 milliseconds.
<i>hello_keep_number</i>	Specifies the number of hello-time intervals that can elapse before an RSVP-TE peer is declared unreachable. The range is from one to 255.
<i>hello_interval_seconds</i>	Specifies the RSVP Hello packet transmission interval. The valid range is from 1 to 60 seconds.
<i>refresh_keep_number</i>	Specifies a factor to be used in calculating the maximum allowed interval without an RSVP refresh message before an RSVP session is torn down. The range is from one to 255.
<i>refresh_seconds</i>	Specifies the interval for sending refresh path messages. The range is from 1 to 600 seconds.
<i>summary_refresh_milliseconds</i>	Specifies the interval for sending summary refresh messages. The valid range is from 50 (1/20 second) to 10000 (10 seconds).
vlan	Specifies that the configured protocol parameters are for a specific VLAN.
<i>vlan_name</i>	Identifies a particular VLAN for which the protocol parameters are configured.
vlan all	indicates that the protocol configuration parameters apply to all RSVP-TE enabled VLANs.

Default

Bundle-message-time: 1000 milliseconds (1 second).

Hello-keep-multiplier value: 3.



Hello-time: 3 seconds.

Refresh-keep-multiplier value: 3.

Refresh-time: 30 seconds.

Summary-refresh-time: 3000 milliseconds (3 seconds).

Usage Guidelines

This command configures the RSVP-TE protocol parameters for the specified VLAN. The VLAN keyword all indicates that the configuration changes apply to all VLANs that have been added to MPLS.

The hello-time value specifies the RSVP hello packet transmission interval. The RSVP hello packet enables the switch to detect when an RSVP-TE peer is no longer reachable. If an RSVP hello packet is not received from a peer within the configured interval, the peer is declared down and all RSVP sessions to and from that peer are torn down. The formula for calculating the maximum allowed interval is: $[\text{hello-time} * \text{hello-keep-multiplier}]$. The default hello-interval time is 3 seconds with a valid range from 1 to 60 seconds. The default hello-keep-multiplier value is three with a range from one to 255.

The refresh-time specifies the interval for sending refresh path messages. RSVP refresh messages provide “soft state” link-level keep-alive information for previously established paths and enable the switch to detect when an LSP is no longer active. Path messages are used to refresh the LSP if summary refresh is disabled. If summary refresh is enabled, summary refresh messages are sent in place of sending individual path messages for every LSP. The default refresh-time is 30 seconds. The minimum and maximum refresh-time values are one and 600 (or 10 minutes) respectively.

If summary refresh is enabled, summary refresh messages are sent at intervals represented by the configured summary-refresh-time. The configurable summary-refresh-time range is 50 milliseconds (one twentieth of a second) to 10000 milliseconds (10 seconds). The default setting for summary-refresh-time is 3000 milliseconds (3 seconds). RSVP sessions are torn down if an RSVP refresh message is not received from a peer within the configured interval. The formula for calculating the maximum allowed interval is: $[(\text{refresh-keep-multiplier} + 0.5) * 1.5 * (\text{refresh-time or summary-refresh-time})]$. The default refresh-keep-multiplier value is three. The minimum and maximum refresh-keep-multiplier values are one and 255 respectively.

The bundle-message-time, specified in milliseconds, indicates the maximum time a transmit buffer is held to allow multiple RSVP messages to be bundled into a single PDU. The default bundle-message-time is 1000 milliseconds (one second). The bundle-message-time value may be set to any value between 50 milliseconds and 3000 milliseconds (or 3 seconds). Message bundling is only attempted when it is enabled.



Note

Summary refresh must be enabled using the “enable mpls rsvp-te summary-refresh” command for a configured summary-refresh-time to actually be used.



Example

The following command sets the RSVP-TE hello time to 5 seconds on all MPLS interfaces:

```
configure mpls rsvp-te timers session hello-time 5 vlan all
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mpls static lsp

```
configure mpls static lsp lsp_name [{egress [egress_label | implicit-null]
egress-vlan evlan_name next-hop ipaddress} {ingress ingress_label {ingress-vlan
ivlan_name}}]
```

Description

Configures the ingress and egress segments of a static LSP.

Syntax Description

<i>lsp_name</i>	Identifies the static LSP to be configured.
<i>egress_label</i>	Specifies the egress label for the LSP. The supported range is x7FC00 to x803FF. The egress label should match the corresponding ingress label of the next hop. There is no egress label at the egress LSR of a static LSP.
egress implicit-null	If PHP is supported, an LSR can be configured to use the implicit-null label for LSPs that terminate at the next-hop LER.
<i>evlan_name</i>	Specifies the egress VLAN for the LSP.
<i>ipaddress</i>	Specifies the IP address for the next-hop router along the static LSP.
<i>ingress_label</i>	Identifies the ingress label for this LSP. The supported range is x7FC00 to x7FFFF at transit LSRs and 0x80000 to 0x803FF at destination LSRs. The ingress label should match the corresponding egress label of the previous hop. There is no ingress label at the ingress LSR of a static LSP.
<i>ivlan_name</i>	When an ingress label is specified, this argument optionally specifies the ingress VLAN for the LSP.

Default

N/A.



Usage Guidelines

The ingress and egress segments can be configured any time before enabling the LSP. At the ingress LER, only the egress segment is configured and at the egress LER, only the ingress segment is configured. For LSPs that transit an LSR, it is mandatory to configure both ingress and egress segments. On any given LSR, the ingress label, if present, must match the egress label on the upstream LSR and the egress label must match the ingress label of the downstream LSR. Once configured, any change to the ingress or egress segments requires administratively disabling the LSP first. If the next-hop IP address is not within the subnet as defined by the interface VLAN name, the configuration is rejected.

Example

The following command configures a static LSP on an ingress LSR:

```
configure mpls static lsp lsp1 egress 0x7fc01 egress-vlan v50 next-hop
50.0.0.2
```

The following command configures a static LSP on a transit LSR:

```
configure mpls static lsp lsp1 egress 0x80001 egress-vlan v100 next-hop
100.0.0.2 ingress 0X7FC01 ingress-vlan v50
```

The following command configures a static LSP on an egress LSR:

```
configure mpls static lsp lsp1 ingress 0x80001 ingress-vlan v100
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mpls static lsp transport

```
configure mpls static lsp lsp_name transport [ip-traffic [allow | deny] | vpn-  
traffic [allow {all | assigned-only} | deny]]
```

Description

Configures the type of traffic that can be transported across a static ingress LSP.



Syntax Description

<i>lsp_name</i>	Identifies the static LSP to be configured.
ip-traffic [allow deny]	Specifies whether IP traffic is to be allowed or denied access to the LSP.
vpn-traffic [allow { all assigned-only } deny]	Specifies whether VPN traffic is to be allowed or denied access to the LSP. The optional assigned-only keyword limits the transport of VPN traffic to only those VPLS instances that are explicitly configured to use the specified LSP.

Default

N/A.

Usage Guidelines

This command has no effect if the named LSP is a transit or egress LSP. By default, IP traffic and VPN traffic are set to deny for a newly created static LSP. The transport configuration options are independent. For example, if VPN traffic is set to allow and IP traffic is set to deny, then no routed IP traffic is transported across the LSP, but the LSP can still transport VPN traffic. When configured to deny for IP traffic, the specified LSP cannot be configured as an IP next hop for a default or static route.

Example

The following command configures a static LSP to transport IP traffic and all VPN traffic:

```
configure mpls static lsp lsp598 transport ip-traffic allow vpn-traffic allow all
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure vpls

```
configure vpls vpls_name {dot1q [ethertype hex_number | tag [include | exclude]]}
{mtu number}
```

Note



This command has been replaced with the following command: `configure l2vpn [vpls vpls_name | vpwsvpws_name] {dot1q [ethertypehex_number | tag [include | exclude]]} {mtunumber}`. This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Configures VPLS parameters.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
dot1q	Specifies the action the switch performs with respect to the 802.1Q ethertype or tag.
ethertype	Overwrites the ethertype value for the customer traffic sent across the PW.
<i>hex_number</i>	Identifies the ethertype, uses the format of 0xN.
tag	Specifies the action the switch performs with respect to the 802.1Q tag.
include	Includes the 802.1Q tag when sending packets over the VPLS L2 VPN.
exclude	Strips the 802.1Q tag before sending packets over the VPLS L2 VPN.
mtu	Specifies the MTU value of the VPLS transport payload packet.
<i>number</i>	The size (in bytes) of the MTU value. The configurable MTU range is 1492 through 9216. The default VPLS MTU value is 1500.

Default

dot1q tag - excluded.

ethertype - the configured switch ethertype is used.

number (MTU) - 1500.

Usage Guidelines

This command configures the VPLS parameters. PWs are point-to-point links used to carry VPN traffic between two devices within the VPLS. Each device must be configured such that packets transmitted between the endpoints are interpreted and forwarded to the local service correctly. The optional ethertype keyword may be used to overwrite the Ethertype value for the customer traffic sent across the PW. By default, the configured switch ethertype is used. If configured, the ethertype in the outer



802.1q field of the customer packet is overwritten using the configured ethertype value. The ethertype value is ignored on receipt.

Optionally, the switch can be configured to strip the 802.1q tag before sending packets over the VPLS L2 VPN. This capability may be required to provide interoperability with other vendor products or to emulate port mode services. The default configuration is to include the 802.1q tag.

The `mtu` keyword optionally specifies the MTU value of the VPLS transport payload packet (customer packet). The MTU value is exchanged with VPLS-configured peer nodes. All VPLS peer nodes must be configured with the same MTU value. If the MTU values do not match, PWs cannot be established between VPLS peers. The MTU values are signaled during PW establishment so that endpoints can verify that MTU settings are equivalent before establishing the PW. By default the VPLS MTU is set to 1500. The configurable MTU range is 1492 through 9216. Changing the MTU setting causes established PWs to terminate. VPLS payload packets may be dropped if the VPLS MTU setting is greater than the MPLS MTU setting for the PW interface.



Note

The maximum MTU value supported depends on the current configuration options. For more information, see [Configure the Layer 2 VPN MTU](#) in the *ExtremeXOS Concepts Guide*.

Example

The following commands change the various parameters of a particular VPLS:

```
configure vpls vpls1 dot1q ethertype 0x8508
configure vpls vpls1 dot1q ethertype 0x8509 mtu 2500
configure vpls vpls1 dot1q tag exclude mtu 2430
configure vpls vpls1 dot1q mtu 2500
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure vpls add peer

```
configure vpls vpls_name add peer ipaddress {core {full-mesh | primary | secondary} | spoke}
```

Note



This command has been replaced with the following command: `configure l2vpn [vpls vpls_name | vpws vpws_name] add peer ipaddress {core {full-mesh | primary | secondary} | spoke}`. This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Configures a VPLS or H-VPLS peer for the node you are configuring.

Syntax Description

<i>vpls_name</i>	Specifies the VPLS for which you are configuring a peer.
<i>ipaddress</i>	Specifies the IP address of the peer node.
core	Specifies that the peer is a core node.
full-mesh	Specifies that the peer is a core full-mesh node. This is the default setting if neither the core or spoke options are specified.
primary	Specifies that the peer is an H-VPLS core node and configures a primary H-VPLS connection to that core node.
secondary	Specifies that the peer is an H-VPLS core node and configures a secondary H-VPLS connection to that core node.
spoke	Specifies that the peer is a H-VPLS spoke node.

Default

N/A.

Usage Guidelines

Up to 32 core nodes can be configured for each VPLS. H-VPLS spoke nodes can peer with core nodes. Nodes can belong to multiple VPLS instances. The *ipaddress* parameter identifies the VPLS node that is the endpoint of the VPLS PW.

Core nodes must be configured in a full-mesh with other core nodes. Thus, all core nodes in the VPLS must have a configured PW to every other core node serving this VPLS. By default, the best LSP is chosen for the PW. The underlying LSP used by the PW can be configured by specifying the named LSP using the CLI command `configure l2vpn [vpls vpls_name | vpws vpws_name] peer ipaddress [add | delete] mpls lsp_lsp_name .`



Spoke nodes establish up to two point-to-point connections to peer with core nodes. If both primary and secondary peers are defined for a spoke node, the spoke node uses one of the peers for all communications. If both peers are available, the spoke node uses the connection to the primary peer. If the primary peer connection fails, the spoke node uses the secondary peer. If the primary peer later recovers, the spoke node reverts back to using the primary peer.

Example

The following command adds a connection from the local core switch to the core switch at 1.1.1.202:

```
configure vpls vpls1 add peer 1.1.1.202
```

The following command adds a connection from the local core switch to the spoke switch at 1.1.1.201:

```
configure vpls vpls1 add peer 1.1.1.201 spoke
```

The following command adds a primary connection from the local spoke switch to the core switch at 1.1.1.203:

```
configure vpls vpls1 add peer 1.1.1.203 core primary
```

History

This command was first available in ExtremeXOS 11.6.

Support for H-VPLS was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vpls add service

```
configure vpls vpls_name add service [{vlan} vlan_name | {vman} vman_name]
```

Note



This command has been replaced with the following command: `configure l2vpn [vpls vpls_name | vpws vpws_name] add service [{vlan} vlan_name | {vman} vman_name]`. This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Configures service for VPLS.



Syntax Description

<i>vpls_name</i>	Identifies the VPLS interface within the switch (character string)
<i>vlan_name</i>	Logically binds the VLAN to the specified VPLS.
<i>vman_name</i>	Adds the named VMAN to the VPLS.

Default

N/A.

Usage Guidelines

This command configures the VPLS service for the specified *vpls_name*. The VPLS service may be a customer VLAN or a customer VMAN. Specifying the *vlan_name* logically binds the VLAN to the specified VPLS. Only one VLAN or VMAN may be configured per VPLS.

When a VLAN service has been configured for a VPLS, the VLAN is added to the VPLS specified by the *vpls_name*. The VLAN ID is locally significant to the switch. Thus, each VLAN VPLS interface within the VPLS network may have a different VLAN ID service bound to the VPLS. This greatly simplifies VLAN ID coordination between metro network access points. Traffic may be switched locally between VLAN ports if more than one port is configured for the VLAN.

When a VMAN service has been configured for a VPLS, the VMAN is added to the VPLS specified by *vpls_name*. The VMAN ID is locally significant to the switch. Thus, each VMAN VPLS interface within the VPLS network may have a different VMAN ID, just like the VLAN service. The only difference is that the VPLS network overwrites the outer VMAN tag on VPLS egress and leaves the inner VLAN tag unmodified. Because the inner VLAN tag is considered part of the customer packet data, the VMAN service can be used to emulate port-based services. This is accomplished by configuring the VPLS to strip the 802.1Q tag from the tunneled packet. Since the switch inserts the VMAN tag when the packet is received and the 802.1Q tag is stripped before the packet is sent on the VPLS PW, all packets received on ports that are members of the VMAN are transmitted unmodified across the VPLS. The command `configure vpls vpls_name dot1q tag exclude` is used to configure the switch to strip the 802.1Q tag on the VPLS.

Example

The example below adds a VLAN and a VMAN to the named VPLS:

```
configure vpls myvpls add service vlan myvlan
configure vpls myvpls add service vman myvman
```

History

This command was first available in ExtremeXOS 11.6.



Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vpls delete peer

```
configure vpls vpls_name delete peer [ipaddress | all]
```

Note



This command has been replaced with the following command: `configure l2vpn [vpls vpls_name | vpws vpws_name] delete peer [ipaddress | all]`. This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Deletes a VPLS peer from the specified *vpls_name*.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
<i>ipaddress</i>	Specifies the IP address for the peer node that is the endpoint of the VC-LSP.
all	Deletes all VPLS peers.

Default

N/A.

Usage Guidelines

This command deletes a VPLS peer from the specified *vpls_name*. When the VPLS peer is deleted, VPN connectivity to the VPLS peer is terminated. The **all** keyword may be used to delete all peers associated with the specified VPLS.

Example

The following example removes connectivity to 1.1.1.202 from VPLS1:

```
configure vpls vpls1 delete peer 1.1.1.202
```

History

This command was first available in ExtremeXOS 11.6.



Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vpls delete service

```
configure vpls vpls_name delete service [{vlan} vlan_name | {vman} vman_name]
```

Note



This command has been replaced with the following command: `configure l2vpn [vpls vpls_name | vpws vpws_name] delete service [{vlan} vlan_name | {vman} vman_name]`. This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Deletes local VPLS service from the specified *vpls_name*.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS interface within the switch (character string).
<i>vlan_name</i>	Logically binds the VLAN to the specified VPLS.
<i>vman_name</i>	Adds the named VMAN to the VPLS.

Default

N/A.

Usage Guidelines

This command deletes the local VPLS service from the specified *vpls_name*. Specifying the *vlan_name* or *vman_name* deletes the service from the VPLS. If there are no services configured for the VPLS, all PWs within the VPLS are terminated from the switch.

Example

The following example removes a service interface from a VPLS:

```
configure vpls vpls1 delete vman vman1
```

History

This command was first available in ExtremeXOS 11.6.



Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vpls health-check vccv

```
configure vpls [vpls_name | all] health-check vccv {interval interval_seconds}
{fault-multiplier fault_multiplier_number}
```

Note



This command has been replaced with the following command: `configure l2vpn [vpls [vpls_name | all] | vpws [vpws_name | all]] health-check vccv {interval interval_seconds} {fault-multiplier fault_multiplier_number}`. This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Configures the VCCV health check test and fault notification intervals for the specified VPLS instance.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS instance for which health check is to be configured.
all	Specifies that the configuration applies to all VPLS instances on the local node.
<i>interval_seconds</i>	Defines the interval between health check tests. The range is 1 to 10 seconds.
<i>fault_multiplier_number</i>	Specifies how long health check waits before a warning level message is logged. The wait period is the <i>interval_seconds</i> multiplied by the <i>fault_multiplier_number</i> . The <i>fault_multiplier_number</i> range is 2 to 6.

Default

Interval is 5 seconds.

Fault multiplier is 4.

Usage Guidelines

The VCCV health-check configuration parameters can be configured at anytime after the VPLS has been created.

The `show l2vpn {vpls {{vpls_name}} | vpws {{vpws_name}} {peeripaddress} {detail} | summary}` command displays the configured `interval_seconds` and `fault-multiplier_number` values for the VPLS and the VCCV activity state.



Example

The following command configures the health check feature on the VPLS instance myvpls:

```
configure vpls myvpls health-check vccv interval 10 fault-notification 40
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vpls peer mpls lsp

```
configure vpls vpls_name peer ipaddress [add | delete] mpls lsp lsp_name
```

Note



This command has been replaced with the following command: `configure l2vpn [vpls vpls_name | vpwsvpws_name] peer ipaddress [add | delete] mpls lsp lsp_name`. This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Configures a named LSP to be used for the PW to the specified VPLS peer.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
<i>ipaddress</i>	Specifies the IP address for the peer node that is the endpoint of the PW-LSP.
add	Permits addition of up to four RSVP-TE LSPs to the VPLS peer.
delete	Removes the LSP specified by the <i>lsp_name</i> parameter from the PW-LSP aggregation list.
<i>lsp_name</i>	Removes the specified lsp.

Default

N/A.



Usage Guidelines

This command configures a named LSP to be used for the PW to the specified VPLS peer. The delete keyword removes the LSP specified by the `lsp_name`. If all the named LSPs are deleted to the configured VPLS peer, VPLS attempts to use the best-routed path LSP, if one exists. The delete portion of this command cannot be used to remove a named LSP that was selected by the switch as the best LSP. If no LSPs exist to the VPLS peer, VPN connectivity to the VPLS peer is lost. Currently, the VPLS PW uses only one LSP.

In ExtremeXOS 15.4, this command is modified to display an informational message when multiple transport LSPs are configured for a VPLS PW, when LSP sharing is not enabled. This message is only displayed once per switch boot.

Example

The following examples add and remove a named LSP:

```
configure vpls vpls1 peer 1.1.1.202 add mpls lsp "to-olympic4"  
configure vpls vpls1 peer 1.1.1.202 delete mpls lsp "to-olympic4"  
  
configure vpls vpls1 peer 20.20.20.83 add mpls lsp lsp2
```



Note

To share LSP's in HW, use the `enable l2vpn sharing` command.

History

This command was first available in ExtremeXOS 11.6.

This command was modified, in ExtremeXOS 15.4, to display an informational message when multiple transport LSPs are configured for a VPLS PW, and LSP sharing is not enabled.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure vpls peer

```
configure vpls vpls_name peer ipaddress [limit-learning number | unlimited-learning]
```

Note



This command has been replaced with the following command: `configure l2vpn [vpls vpls_name | vpwsvpws vpwsvpws_name] peer ipaddress [limit-learning number | unlimited-learning]`. This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Configures the maximum number of MAC SAs (Source Addresses) that can be learned for a given VPLS and peer.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
<i>ipaddress</i>	Specifies the IP address for the peer node that is the endpoint of the PW-LSP.
limit-learning	Specifies a limit to the number of MAC SAs to be learned for the specified VPLS and peer.
<i>number</i>	The maximum number of MAC SAs that can be learned for the specified VPLS and peer.
unlimited-learning	Specifies no limit to the number of MAC SAs to be learned for the specified VPLS and peer.

Default

Unlimited.

Usage Guidelines

This command configures the maximum number of MAC SAs (Source Addresses) that can be learned for a given VPLS and peer. This parameter can only be modified when the specified VPLS is disabled. The `unlimited-learning` keyword can be used to specify that there is no limit. The default value is `unlimited-learning`.

Example

The following example causes no more than 20 MAC addresses to be learned on VPLS1's PW to 1.1.1.202:

```
configure vpls vpls1 peer 1.1.1.202 limit-learning 20
```



History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vpls peer mpls lsp

```
configure vpls vpls_name peer ipaddress [add | delete] mpls lsp lsp_name
```

Note



This command has been replaced with the following command: `configure l2vpn [vpls vpls_name | vpwsvpws_name] peer ipaddress [add | delete] mpls lsp_name`. This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Configures a named LSP to be used for the PW to the specified VPLS peer.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
<i>ipaddress</i>	Specifies the IP address for the peer node that is the endpoint of the PW-LSP.
add	Permits addition of up to four RSVP-TE LSPs to the VPLS peer.
delete	Removes the LSP specified by the <i>lsp_name</i> parameter from the PW-LSP aggregation list.
<i>lsp_name</i>	Removes the specified lsp.

Default

N/A.

Usage Guidelines

This command configures a named LSP to be used for the PW to the specified VPLS peer. The delete keyword removes the LSP specified by the *lsp_name*. If all the named LSPs are deleted to the configured VPLS peer, VPLS attempts to use the best-routed path LSP, if one exists. The delete portion of this command cannot be used to remove a named LSP that was selected by the switch as the best LSP. If no LSPs exist to the VPLS peer, VPN connectivity to the VPLS peer is lost. Currently, the VPLS PW uses only one LSP.



In ExtremeXOS 15.4, this command is modified to display an informational message when multiple transport LSPs are configured for a VPLS PW, when LSP sharing is not enabled. This message is only displayed once per switch boot.

Example

The following examples add and remove a named LSP:

```
configure vpls vpls1 peer 1.1.1.202 add mpls lsp "to-olympic4"
configure vpls vpls1 peer 1.1.1.202 delete mpls lsp "to-olympic4"

configure vpls vpls1 peer 20.20.20.83 add mpls lsp lsp2
```



Note

To share LSP's in HW, use the `enable l2vpn sharing` command.

History

This command was first available in ExtremeXOS 11.6.

This command was modified, in ExtremeXOS 15.4, to display an informational message when multiple transport LSPs are configured for a VPLS PW, and LSP sharing is not enabled.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vpls snmp-vpn-identifier

```
configure vpls vpls_name snmp-vpn-identifier identifier
```

Description

Configures a SNMP VPN identifier for traps from the specified VLPLS.

Syntax Description

<i>vpls_name</i>	Specifies the VPLS for which you are configuring the identification string.
<i>identifier</i>	Specifies a text string to identify the VPLS in SNMP traps.

Default

N/A.



Usage Guidelines

None.

Example

The following command configures the identifier `vpls1trap` for SNMP VPN traps on VPLS `vpls1`:

```
configure vpls vpls1 snmp-vpn-identifier vpls1trap
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

create l2vpn fec-id-type pseudo-wire

```
create l2vpn [vpls vpls_name | vpws vpws_name] fec-id-type pseudo-wire pwid
```

Description

Creates a Layer 2 VPN, which can be either a VPLS or VPWS.

Syntax Description

<code>vpls_name</code>	Identifies the VPLS within the switch (character string). The <code>vpls_name</code> string must begin with an alphabetic character, and may contain up to 31 additional alphanumeric characters.
<code>vpws_name</code>	Identifies the VPWS within the switch (character string). The <code>vpws_name</code> string must begin with an alphabetic character, and may contain up to 31 additional alphanumeric characters.
<code>pwid</code>	Specifies a PW ID. Must be a non-zero 32-bit value that has network-wide significance.

Default

For the VPLS dot1q tag, the default value is `exclude`.

Usage Guidelines

Each VPLS or VPWS is a member of a single VPN, and each VPN can have only one associated VPLS or VPWS per switch. External to the switch, each VPN has an identifier.

A VPLS or VPWS name must begin with an alphabetical character and may contain alphanumeric characters and underscores (`_`), but it cannot contain spaces. The maximum allowed length for a name



is 32 characters. For name creation guidelines and a list of reserved names, see [Object Names](#) in the *ExtremeXOS Concepts Guide*.

Any non-zero 32-bit value that has network-wide significance can be specified for the identifier. This `pwid` is used on all pseudo-wires in the VPLS.

The `l2vpn` keyword is introduced in ExtremeXOS Release 12.4 and is required when creating a VPWS. For backward compatibility, the `l2vpn` keyword is optional when creating a VPLS. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

**Note**

The switch's LSR ID must be configured before a VPLS or VPWS can be created.

Example

This example creates a VPLS with 99 as the PW ID:

```
create vpls vpls1 fec-id-type pseudo-wire 99
```

The following example creates a VPWS with 101 as the PW ID:

```
create l2vpn vpws vpws1 fec-id-type pseudo-wire 101
```

History

This command was first available in ExtremeXOS 11.6.

The `l2vpn` and `vpws` keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

create mpls rsvp-te lsp

```
create mpls rsvp-te lsp lsp_name destination ipaddress
```

Description

Creates internal resources for an RSVP-TE LSP.



Syntax Description

<i>lsp_name</i>	Specifies a name for the LSP you are creating. The character string must begin with an alphabetic character, and may contain up to 31 additional alphanumeric characters.
<i>ipaddress</i>	Specifies the endpoint of the LSP.

Default

N/A.

Usage Guidelines

This command creates internal resources for an RSVP-TE LSP.

The LSP name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see [Object Names](#) in the *ExtremeXOS Concepts Guide*.

The *ipaddress* specifies the endpoint of the LSP. The LSP is not signaled until a path is specified for the LSP using the `configure mpls rsvp-te lsp lsp_name add path` command. When multiple LSPs are configured to the same destination, IP traffic is load-shared across active LSPs that have IP transport enabled. The maximum number of RSVP-TE LSPs that can be created is 1024.



Note

The LSP must be created before it can be configured.

Example

The following command creates an RSVP-TE LSP:

```
create mpls rsvp-te lsp lsp598 destination 11.100.100.8
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

create mpls rsvp-te path

```
create mpls rsvp-te path path_name
```



Description

Creates an RSVP-TE routed path resource.

Syntax Description

<i>path_name</i>	Identifies the path within the switch. The character string must begin with an alphabetic character, and may contain up to 31 additional alphanumeric characters.
------------------	---

Default

N/A.

Usage Guidelines

This command creates an RSVP-TE path resource.

The *path_name* parameter must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see [Object Names](#) in the *ExtremeXOS Concepts Guide*.

The maximum number of configurable paths is 255.



Note

The RSVP-TE LSP is not signaled along the path until an LSP is created and then configured with the specified *path_name*.

Example

The following example creates an RSVP-TE path:

```
create mpls rsvp-te path path598
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

create mpls rsvp-te profile

```
create mpls rsvp-te profile profile_name {standard}
```



Description

Creates configured RSVP-TE profile with the specified profile name.

Syntax Description

<i>profile_name</i>	Identifies the RSVP-TE profile. The character string must begin with an alphabetic character and may contain up to 31 additional alphanumeric characters.
standard	The standard option differentiates this command version from the command that creates a fast-reroute profile. If you do not specify an option, a standard RSVP-TE profile is created.

Default

N/A.

Usage Guidelines

This command creates a configured RSVP-TE profile with the specified profile name. The default profile cannot be deleted. If a profile is associated with a configured LSP, the profile cannot be deleted.

A profile name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see [Object Names](#) in the *ExtremeXOS Concepts Guide*.

Example

The following command creates an RSVP-TE profile:

```
create mpls rsvp-te profile prof598
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

create mpls rsvp-te profile fast-reroute

```
create mpls rsvp-te profile profile_name fast-reroute
```



Description

Creates an LSP container to hold FRR configuration parameters.

Syntax Description

<i>profile_name</i>	Specifies a name for the new RSVP-TE fast-reroute profile. The character string must begin with an alphabetic character and may contain up to 31 additional alphanumeric characters.
---------------------	--

Default

N/A.

Usage Guidelines

A profile name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see [Object Names](#) in the *ExtremeXOS Concepts Guide*.

Example

The following command creates a new FRR profile named frrprofile:

```
create mpls rsvp-te profile frrprofile fast-reroute
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

create mpls static lsp

```
create mpls static lsp lsp_name destination ipaddress
```

Description

Creates internal resources for a static LSP and assigns a name to the LSP.



Syntax Description

<i>lsp_name</i>	Identifies the LSP to be created.
<i>ipaddress</i>	Specifies the endpoint of the LSP.

Default

N/A.

Usage Guidelines

An LSP name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see *Object Names* in the *ExtremeXOS Concepts Guide*.

Example

The following command creates a static LSP:

```
create mpls static lsp lsp598 destination 11.100.100.8
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support MPLS as described in *Feature License Requirements* in the *ExtremeXOS Concepts Guide*.

create vpls fec-id-type pseudo-wire

```
create vpls vpls_name fec-id-type pseudo-wire pwid
```

Note



This command has been replaced with the following command: `create l2vpn [vpls vpls_name | vpwsvpws_name] fec-id-type pseudo-wire pwid`. This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Creates a VPLS instance with the specified *vpls_name*.



Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string). The <i>vpls_name</i> string must begin with an alphabetic character, and may contain up to 31 additional alphanumeric characters.
<i>pwid</i>	Specifies a PW ID. Must be a non-zero 32-bit value that has network-wide significance.

Default

For the VPLS dot1q tag, the default value is exclude.

Usage Guidelines

This command creates a VPLS instance with the specified *vpls_name*. Each VPLS represents a separate virtual switch instance (VSI).

The *vpls_name* parameter must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see [Object Names](#) in the *ExtremeXOS Concepts Guide*.

Each VPLS is a member of a single VPN and each VPN may have only one associated VPLS per switch. External to the switch, each VPN has an identifier.

Any non-zero 32-bit value that has network-wide significance can be specified for the identifier. This *pwid* is used on all pseudo-wires in the VPLS.



Note

The switch's LSR ID must be configured before a VPLS can be created.

Example

This example creates a VPLS with 99 as the PW ID:

```
create vpls vpls1 fec-id-type pseudo-wire 99
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

delete l2vpn

```
delete l2vpn [vpls [vpls_name | all] | vpws [vpws_name | all]]
```



Description

Deletes the specified VPLS or VPWS.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
<i>vpws_name</i>	Identifies the VPWS within the switch (character string).
all	Specifies all VPLS or VPWS instances.

Default

N/A.

Usage Guidelines

All PWs established to VPLS or VPWS peers are terminated.

The **l2vpn** keyword is introduced in ExtremeXOS Release 12.4 and is required when deleting a VPWS. For backward compatibility, the **l2vpn** keyword is optional when deleting a VPLS. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

This commands deletes the VPLS myvpls:

```
delete vpls myvpls
```

This commands deletes the VPWS myvpws:

```
delete l2vpn vpws myvpws
```

History

This command was first available in ExtremeXOS 11.6.

The l2vpn and vpws keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



delete mpls rsvp-te lsp

```
delete mpls rsvp-te lsp [lsp_name | all]
```

Description

Deletes internal resources for the specified RSVP-TE LSP.

Syntax Description

<i>lsp_name</i>	Specifies the LSP within the switch to be deleted.
all	Deletes all RSVP-TE configured LSPs.

Default

N/A.

Usage Guidelines

This command deletes internal resources for the specified RSVP-TE LSP. The LSP is first withdrawn if it is currently active. Deleting an LSP may cause a PW to fail. Any static routes configured to a deleted LSP are also removed.

Example

The following command deletes the configured RSVP-TE LSP named lsp598:

```
delete mpls rsvp-te lsp lsp598
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

delete mpls rsvp-te path

```
delete mpls rsvp-te path [path_name | all]
```

Description

Deletes a configured RSVP-TE routed path with the specified path name.



Syntax Description

<i>path_name</i>	Specifies a path within the switch to be deleted.
all	Deletes all paths not associated with an LSP.

Default

N/A.

Usage Guidelines

This command deletes a configured RSVP-TE routed path with the specified name. All associated configuration information for the specified path is deleted. If the all keyword is specified, all paths not associated with an LSP are deleted.



Note

A path cannot be deleted as long as the path name is associated with an LSP.

Example

The following command deletes the configured RSVP-TE path named path598:

```
delete mpls rsvp-te path path598
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

delete mpls rsvp-te profile

```
delete mpls rsvp-te profile [profile_name | all]
```

Description

Deletes a configured RSVP-TE profile with the specified profile name.

Syntax Description

<i>profile_name</i>	Specifies a configured RSVP-TE profile to be deleted.
all	Deletes all profiles not associated with an LSP, except the default profile.



Default

N/A.

Usage Guidelines

This command deletes a configured RSVP-TE profile with the specified profile name. If the all keyword is specified, all profiles not associated with an LSP are deleted (except for the default profile).



Note

A profile cannot be deleted as long as the profile name is associated with a configured LSP. The **default** profile cannot be deleted.

Example

The following command deletes the configured RSVP-TE profile named prof598:

```
delete mpls rsvp-te profile prof598
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

delete mpls static lsp

```
delete mpls static lsp [lsp_name | all]
```

Description

Deletes internal resources for one or all static LSPs.

Syntax Description

<i>lsp_name</i>	Identifies the LSP to be deleted.
all	Specifies that all LSPs are to be deleted.

Default

N/A.



Usage Guidelines

All resources associated with the specified LSPs are released. Static LSPs cannot be deleted when the LSP is configured for an IP route or VPLS configuration.

Example

The following command deletes a static LSP:

```
delete mpls static lsp lsp598
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

delete vpls

```
delete vpls [vpls_name | all]
```

Note



This command has been replaced with the following command: `delete l2vpn [vpls [vpls_name | all] | vpws [vpws_name | all]]`. This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Deletes the VPLS with the specified `vpls_name`.

Syntax Description

<code>vpls_name</code>	Identifies the VPLS within the switch (character string).
<code>all</code>	Specifies all VPLS.

Default

N/A.



Usage Guidelines

This command deletes the VPLS with the specified *vpls_name*. All PWs established to VPLS peers are terminated. The **all** keyword may be used to indicate that all VPLS instances are to be deleted.

Example

This commands deletes the VPLS myvpls:

```
delete vpls myvpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

disable bgp mpls-next-hop

```
disable bgp mpls-next-hop
```

Description

Disables IP forwarding over calculated MPLS LSPs to subnets learned via BGP.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables IP forwarding over calculated MPLS LSPs to subnets learned via BGP. (Calculated refers to an LSP that only reaches part of the way to the destination). By default, IP forwarding over MPLS LSPs to subnets learned via BGP is disabled.



Example

The following command disables BGP's use of MPLS LSPs to reach BGP routes:

```
disable bgp mpls-next-hop
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable iproute mpls-next-hop

```
disable iproute mpls-next-hop
```

Description

Disables IP forwarding over MPLS LSPs for the default VR.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables IP forwarding over MPLS LSPs for the default VR. When disabled, any route with an MPLS LSP as its next hop becomes inactive and is not used to tunnel IP traffic across the MPLS network. By default, IP forwarding over MPLS LSPs is disabled.

Example

This command disables IP forwarding over MPLS LSPs.

```
disable iproute mpls-next-hop
```



History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable l2vpn

```
disable l2vpn [vppls [vppls_name | all] | vpws [vpws_name | all]]
```

Description

Disables the specified VPLS or VPWS.

Syntax Description

<i>vppls_name</i>	Identifies the VPLS within the switch (character string).
<i>vpws_name</i>	Identifies the VPWS within the switch (character string).
all	Specifies all VPLS or VPWS instances.

Default

All newly created VPLS instances are enabled.

Usage Guidelines

When a VPLS or VPWS instance is disabled, all sessions to its configured peers are terminated. Any locally attached service VLAN/VMAN is immediately isolated from other devices residing in the VPN. If this is an H-VPLS core node, then all spoke nodes connected to this peer are isolated unless redundant core access is configured.

The **l2vpn** keyword is introduced in ExtremeXOS Release 12.4 and is required when disabling a VPWS. For backward compatibility, the **l2vpn** keyword is optional when disabling a VPLS. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The following command disables the VPLS named myvpls:

```
disable vpls myvpls
```



The following command disables the VPWS named myvpws:

```
disable l2vpn vpws myvpws
```

History

This command was first available in ExtremeXOS 11.6.

The **l2vpn** and **vpws** keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable l2vpn vpls fdb mac-withdrawal

```
disable l2vpn vpls fdb mac-withdrawal
```

Description

Disables the MAC address withdrawal capability.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

When disabled, the switch does not send MAC address withdrawal messages. If a MAC address withdrawal message is received from another VPLS peer, the local peer processes the message and withdraws the specified MAC addresses from its FDB, regardless of the MAC address withdrawal configuration.

Example

The following command disables MAC address withdrawal:

```
disable l2vpn vpls fdb mac-withdrawal
```



History

This command was first available in ExtremeXOS 12.1.

The **l2vpn** keyword was added in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable l2vpn health-check vccv

```
disable l2vpn [vpls [vpls_name | all] | vpws [vpws_name | all]] health-check vccv
```

Description

Disables the VCCV health check feature on the specified VPLS or VPWS instances.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS for which health check is to be disabled.
<i>vpws_name</i>	Identifies the VPWS for which health check is to be disabled.
all	Specifies that health check is to be disabled on all VPLS instances on the local node.

Default

Health check is disabled.

Usage Guidelines

The **l2vpn** keyword was introduced in ExtremeXOS Release 12.4 and is required when disabling health check for a VPWS instance. For backward compatibility, the **l2vpn** keyword is optional when disabling health check for VPLS instance. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The following command disables the health check feature on the VPLS instance myvpls:

```
disable l2vpn vpls myvpls health-check vccv
```

History

This command was first available in ExtremeXOS 12.1.



The **l2vpn** and **vpws** keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable l2vpn service

```
disable l2vpn [vpls [vpls_name | all] | vpws [vpws_name | all]] service
```

Description

Disables the configured services for the specified VPLS or VPWS.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
<i>vpws_name</i>	Identifies the VPWS within the switch (character string).
all	Specifies all VPLS or VPWS instances.

Default

Enabled.

Usage Guidelines

When services are disabled, the VPLS or VPWS is removed from all peer sessions. The keyword **all** disables services for all VPLS instances.

The **l2vpn** keyword was introduced in ExtremeXOS Release 12.4 and is required when disabling a service for a VPWS peer. For backward compatibility, the **l2vpn** keyword is optional when disabling a service for a VPLS peer. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The following command disables the configured services for VPLS myvpls:

```
disable l2vpn vpls myvpls service
```

History

This command was first available in ExtremeXOS 11.6.



The **l2vpn** and **vpws** keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable l2vpn sharing

disable l2vpn sharing

Description

Disables LSP sharing for Layer 2 VPN pseudo-wires.

Syntax Description

This command has no keywords or arguments.

Default

Disabled.

Usage Guidelines

This command disables LSP sharing for L2VPN PWs. When LSP sharing is disabled, only 1 named LSP is used for a PW. When LSP sharing is enabled, up to 16 named LSPs are used for a PW.

If LSP Sharing is disabled, and more than 1 Transport LSP is programmed into HW, all but 1 Transport LSP is removed from HW, and the configuration is preserved. If LSP Sharing is enabled, and more than 1 Transport LSP was previously configured, the remaining LSPs is programmed into HW as they become available for use.

Example

The following command disables LSP sharing for L2VPN PWs:

```
disable l2vpn sharing
```

History

This command was first available in ExtremeXOS 15.4.



Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable mpls

disable mpls

Description

Disables MPLS on the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

When MPLS is disabled, no label traffic is received or transmitted, and all MPLS-related protocol peer sessions are terminated.

Example

The following command globally disables MPLS on the switch:

```
disable mpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable mpls bfd

```
disable mpls bfd [vlan all | {vlan} vlan_name] {delete-sessions}
```



Description

Disables the Bidirectional Forwarding Detection (BFD) client for MPLS on the specified VLAN or on all VLANs.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN on which to disable the MPLS BFD client.
delete-sessions	Specifies to delete all MPLS BFD sessions.

Default

Keep existing MPLS BFD sessions.

Usage Guidelines

This command instructs MPLS to cease the establishment of new BFD sessions with neighbors as LSPs are established with those neighbors. The default behavior retains the existing BFD sessions and ignores status updates from those existing sessions. The **delete-sessions** option instructs MPLS to request the deletion of existing sessions. Whether the sessions are deleted or not, the link state presented to the upper MPLS layers reverts to the normal link operational status.

Note



Deleting existing sessions can result in a neighbor DOWN indication from BFD to MPLS on the other end of the session (the peer switch) and a subsequent interface DOWN indication presented to the upper layers of MPLS on that peer switch. These actions can cause MPLS to reroute or fail the affected LSPs.

To disable the MPLS BFD client and delete all BFD sessions without disrupting the LSPs between two switches, do the following:

- Log into switch A as an admin user and issue the command: `disable mpls bfd vlanx`.
- Log into switch B as an admin user and issue the command: `disable mpls bfd vlanx delete-sessions`

Example

The following command disables the MPLS BFD client on VLAN vlan1:

```
disable mpls bfd vlan1
```

History

This command was first available in ExtremeXOS 12.4.



Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable mpls exp examination

```
disable mpls exp examination
```

Description

Disables assigning an MPLS packet to a QoS profile based on the MPLS packet's EXP value.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables assigning an MPLS packet to a QoS profile based on the MPLS packet's EXP value.

When disabled, all received MPLS packets are assigned to QoS profile **qp1**.

Example

The following command disables the assignment of an MPLS packet to a QoS profile based on the MPLS packet's EXP value:

```
disable mpls exp examination
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



disable mpls exp replacement

disable mpls exp replacement

Description

Disables setting an MPLS packet's EXP value based on the packet's QoS profile.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables setting an MPLS packet's EXP value based on the packet's QoS profile. The QoS profiles to EXP value mappings are configured using the **configure mpls exp replacement** command.

When disabled, all MPLS packets are transmitted with an EXP value of zero.

Example

The following command disables the setting of an MPLS packet's EXP value based on the packet's QoS profile:

```
disable mpls exp replacement
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

disable mpls ldp

disable mpls ldp [{**vlan**} *vlan_name* | **vlan all**]



Description

Disables LDP for the specified MPLS-configured VLANs.

Syntax Description

vlan	Disables LDP for one or more specific VLANs.
<i>vlan_name</i>	Disables LDP on the specified VLAN.
vlan all	Disables LDP for all VLANs that have been added to MPLS.

Default

Disabled.

Usage Guidelines

When LDP is disabled, all LDP-advertised labels are withdrawn and all LDP peer sessions are terminated on the specified VLAN(s). By default, LDP is disabled for all VLANs. Specifying the optional **all** keyword disables LDP for all VLANs that have been added to MPLS.

Example

The following command disables LDP for all VLANs:

```
disable mpls ldp vlan all
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

disable mpls ldp bgp-routes

```
disable mpls ldp bgp-routes
```

Description

Disables LDP's use of IP prefixes learned from BGP when establishing LDP LSPs.



Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command disables LDP's establishment of LSPs to routes learned via BGP, thus reducing the internal resources used by LDP. Note that MPLS LSPs can still be used to transport packets to routes learned via BGP through the use of the `enable bgp mpls-next-hop` command.

When enabled, LDP uses routes learned via BGP when establishing LDP LSPs. As each established LSP consumes internal resources, it is recommended that this setting be used only in BGP environments where the number of BGP routes is controlled.

Example

The following command disables the use of BGP routes by LDP:

```
disable mpls ldp bgp-routes
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable mpls ldp loop-detection

```
disable mpls ldp loop-detection
```

Description

Disables LDP loop detection on the switch.

Syntax Description

This command has no arguments or variables.



Default

Disabled.

Usage Guidelines

Loop detection provides a mechanism for finding looping LSPs and for preventing Label Request messages from looping in the presence of non-merge-capable LSRs. The mechanism makes use of Path Vector and Hop Count TLVs carried by Label Request and Label Mapping messages.

When LDP loop detection is disabled, LDP does not attempt to detect routing loops.

Example

The following command globally disables LDP loop detection on the switch:

```
disable mpls ldp loop-detection
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable mpls php

```
disable mpls php [{vlan} vlan_name | vlan all]
```

Description

Disables penultimate hop popping (PHP) on the specified VLAN. When enabled, PHP is requested on all LSPs advertised over that VLAN for which the switch is the egress LSR.

Syntax Description

vlan	Disables PHP for one or more specific VLANs.
<i>vlan_name</i>	Disables PHP on the specified VLAN.
vlan all	Disables PHP for all VLANs that have been added to MPLS.

Default

Disabled



Usage Guidelines

When PHP is disabled on a VLAN, penultimate hop popping is not requested on any LSPs advertised over that VLAN for which the switch is the egress LSR. Therefore, the Implicit Null Label is not used for any advertised mapping. Extreme's MPLS implementation always performs penultimate hop popping when requested to do so by a peer LSR. When the all VLANs option is selected, PHP is disabled on all existing MPLS interfaces.

Note



PHP is sometimes used to reduce the number of MPLS labels in use. If PHP is enabled on any MPLS interface, a unique MPLS label is consumed for every label advertised over that interface. Therefore, if PHP is being disabled to reduce label consumption, it should be done on all interfaces for minimal label consumption.

In ExtremeXOS, this command can be executed while MPLS is enabled.

Example

The following command disables penultimate hop popping (PHP) on the specified VLAN:

```
disable mpls php vlan vlan1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable mpls protocol ldp

```
disable mpls protocol ldp
```

Description

Disables LDP for the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.



Usage Guidelines

When LDP is disabled, all advertised LDP labels are withdrawn and LDP peer sessions are terminated. Note that this includes any LDP peer sessions established for L2 VPNs. By default, LDP is globally disabled. While LDP is transitioning to the enabled state, only the MPLS show commands are accepted.

Example

The following command globally disables LDP on the switch:

```
disable mpls protocol ldp
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

disable mpls protocol rsvp-te

```
disable mpls protocol rsvp-te
```

Description

Disables RSVP-TE for the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

When RSVP-TE is disabled, all TE LSPs are released and TE LSPs cannot be established or accepted. While RSVP-TE is transitioning to the disabled state, only the MPLS show commands are accepted.



Example

The following command globally disables RSVP-TE on the switch:

```
disable mpls protocol rsvp-te
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable mpls rsvp-te

```
disable mpls rsvp-te te [{vlan} vlan_name | vlan all]
```

Description

Disables RSVP-TE for the specified MPLS-configured VLAN.

Syntax Description

vlan	Specifies that RSVP-TE is to be disabled on a specific VLAN.
<i>vlan_name</i>	Specifies the VLAN for which RSVP-TE is disabled.
vlan all	Disables RSVP-TE on all VLANs that have been added to MPLS.

Default

Disabled.

Usage Guidelines

This command disables RSVP-TE for the specified MPLS configured VLANs. When RSVP-TE is disabled, all TE LSPs are released and TE LSPs cannot be established or accepted. By default, RSVP-TE is disabled for all MPLS configured VLANs. Specifying the optional **all** keyword disables RSVP-TE for all VLANs that have been added to MPLS.



Example

The following command disables RSVP-TE on the named VLAN:

```
disable mpls rsvp-te vlan vlan_10
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable mpls rsvp-te bundle-message

```
disable mpls rsvp-te bundle-message [{vlan} vlan_name | vlan all]
```

Description

Disables the bundling of RSVP-TE messages for the specified VLAN interface.

Syntax Description

vlan	Specifies that message-bundling is to be disabled on a specific VLAN.
<i>vlan_name</i>	Identifies the VLAN interface on which message bundling is disabled.
vlan all	Specifies that message bundling is disabled on all VLAN interfaces that have been added to MPLS.

Default

Disabled.

Usage Guidelines

This command disables the bundling of RSVP-TE messages for the VLAN specified interface. By default, message bundling is disabled. Specifying the **all** keyword disables message bundling on all VLANs that have been added to MPLS.



Example

The following command disables message bundling on the specified VLAN:

```
disable mpls rsvp-te bundle-message vlan vlan_1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable mpls rsvp-te fast-reroute

```
disable mpls rsvp-te fast-reroute
```

Description

Disables the MPLS RSVP-TE fast reroute (FRR) protection feature.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

When FRR is disabled on the LSR, all established FRR LSPs on the local LSR are torn down, and only standard LSPs can be signaled and processed. The configuration for any existing FRR LSPs is retained, but it is not used until the FRR protection feature is enabled. This command can be used to test the performance of an LSR without the FRR functionality or when the LSR doesn't behave as expected for either standard or FRR LSPs.

Example

The following command disables FRR protection on the local switch:

```
disable mpls rsvp-te fast-reroute
```



History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable mpls rsvp-te lsp

```
disable mpls rsvp-te lsp [lsp_name | all]
```

Description

Disables an RSVP-TE LSP.

Syntax Description

<i>lsp_name</i>	Specifies the LSP within the switch to be disabled.
all	Disables all RSVP-TE configured LSPs.

Default

Enabled.

Usage Guidelines

This command disables an RSVP-TE LSP. When an RSVP-TE LSP is disabled, the switch terminates the LSP by signaling the destination by sending a PATH_TEAR message. If there are other LSPs configured to the same destination, traffic may continue to be transmitted to the destination over another LSP. Disabling an LSP does not otherwise change its configuration.

Example

The following command disables the LSP named lsp598:

```
disable mpls rsvp-te lsp lsp598
```

History

This command was first available in ExtremeXOS 11.6.



Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable mpls rsvp-te summary-refresh

```
disable mpls rsvp-te summary-refresh [{vlan} vlan_name | vlan all]
```

Description

Disables the sending of summary refresh messages, instead of path messages, to refresh RSVP-TE path state for the specified VLAN interface.

Syntax Description

vlan	Specifies that summary refresh messages cannot refresh the RSVP-TE path state on one or more VLAN interfaces.
<i>vlan_name</i>	Specifies the VLAN interface for which RSVP-TE summary refresh messages are to be disabled.
vlan all	Specifies that summary refresh messages are to be disabled on all VLAN interfaces that have been added to MPLS.

Default

Disabled.

Usage Guidelines

This command disables the sending of summary refresh messages to refresh RSVP-TE path state for the specified VLAN interface. By default, summary refresh is disabled. Specifying the **all** keyword disables summary refresh on all VLANs that have been added to MPLS.

Example

The following command disables summary refresh on the specified VLAN:

```
disable mpls rsvp-te summary-refresh vlan vlan_1
```

History

This command was first available in ExtremeXOS 11.6.



Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable mpls static lsp

```
disable mpls static lsp { lsp_name | all }
```

Description

Administratively disables one or all static LSPs.

Syntax Description

<i>lsp_name</i>	Identifies an LSP to be disabled.
all	Specifies that all static LSPs on this LSR are to be disabled.

Default

N/A.

Usage Guidelines

On executing this command, the software de-activates the specified LSPs by setting the administrative state of each LSP to down.

Example

The following command disables a static LSP:

```
disable mpls static lsp lsp598
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable mpls vlan

```
disable mpls [{vlan} vlan_name | vlan all]
```



Description

Disables the MPLS interface for the specified VLAN(s).

Syntax Description

vlan	Disables an MPLS interface for one or more specific VLANs.
<i>vlan_name</i>	Disables an MPLS interface on the specified VLAN.
vlan all	Disables an MPLS interface for all VLANs that have been added to MPLS.

Default

The MPLS interface is disabled for a VLAN.

Usage Guidelines

Disabling MPLS causes all LSPs to be released and all LDP and RSVP-TE peer sessions to be terminated on the specified VLAN(s).

Example

The following command disables an MPLS interface for the specified VLAN:

```
disable mpls vlan vlan-nyc
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

disable ospf mpls-next-hop

```
disable ospf mpls-next-hop {vr vrf_name}
```

Description

Disables IP forwarding over calculated MPLS LSPs to subnets learned via OSPF.

Syntax Description

<i>vrf_name</i>	Specifies OSPF on a particular VRF.
-----------------	-------------------------------------



Default

Disabled.

Usage Guidelines

This command disables IP forwarding over calculated MPLS LSPs to subnets learned via OSPF. (Calculated refers to an LSP that only reaches part of the way to the destination). By default, IP forwarding over MPLS LSPs to subnets learned via OSPF is disabled.

In order to disable OSPF on a particular VRF, you must supply the optional `vr vr-name` CLI parameter.

Example

The following command disables OSPF's use of MPLS LSPs to reach OSPF routes:

```
disable ospf mpls-next-hop
```

History

This command was first available in ExtremeXOS 11.6.

The `vr` keyword and `vrf_name` variable were added in ExtremeXOS 15.3.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable snmp traps l2vpn

```
disable snmp traps l2vpn
```

Description

Disables SNMP traps associated with Layer 2 VPNs for all MPLS configured VLANs.

Syntax Description

This command has no arguments or variables.

Default

All Layer 2 VPN traps are disabled.



Example

The following command disables SNMP traps associated with Layer 2 VPNs:

```
disable snmp traps l2vpn
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable snmp traps mpls

```
disable snmp traps mpls
```

Description

Disables SNMP traps associated with MPLS for all MPLS configured VLANs.

Syntax Description

This command has no arguments or variables.

Default

All MPLS traps are disabled.

Example

The following command disables SNMP traps associated with MPLS:

```
disable snmp traps mpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



disable vpls

```
disable vpls [vpls_name | all]
```

Note



This command has been replaced with the following command: `disable l2vpn [vpls [vpls_name | all] | vpws [vpws_name | all]]`. This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Disables the VPLS instance specified by `vpls_name`.

Syntax Description

<code>vpls_name</code>	Identifies the VPLS within the switch (character string)
<code>all</code>	Specifies all VPLS.

Default

All newly created VPLS instances are enabled.

Usage Guidelines

This command disables the VPLS instance specified by `vpls_name`. When a VPLS instance is disabled, all sessions to its configured peers are terminated. Any locally attached service VLAN/VMAN is immediately isolated from other devices residing in the VPN. If this is an H-VPLS core node, then all spoke nodes connected to this peer are isolated unless redundant core access is configured.

Example

The following command disables the VPLS named `myvpls`:

```
disable vpls myvpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



disable vpls fdb mac-withdrawal

```
disable vpls fdb mac-withdrawal
```

Note



This command has been replaced with the following command: `disable l2vpn vpls fdb mac-withdrawal`. This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Disables the VPLS MAC address withdrawal capability.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

When disabled, the switch does not send MAC address withdrawal messages. If a MAC address withdrawal message is received from another VPLS peer, the local VPLS peer processes the message and withdraws the specified MAC addresses from its FDB, regardless of the MAC address withdrawal configuration.

Example

The following command disables MAC address withdrawal:

```
disable vpls fdb mac-withdrawal
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support MPLS as described in *Feature License Requirements* in the *ExtremeXOS Concepts Guide*.



disable vpls health-check vccv

```
disable vpls [vpls_name | all] health-check vccv
```

Note



This command has been replaced with the following command: `disable l2vpn [vpls [vpls_name | all] | vpws [vpws_name | all]] health-check vccv`. This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Disables the VCCV health check feature on one or all VPLS instances on the local node.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS for which health check is to be disabled.
all	Specifies that health check is to be disabled on all VPLS instances on the local node.

Default

Health check is disabled.

Usage Guidelines

None.

Example

The following command disables the health check feature on the VPLS instance myvpls:

```
disable vpls myvpls health-check vccv
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



disable vpls service

```
disable vpls [vpls_name | all] service
```

Note



This command has been replaced with the following command: `disable l2vpn [vpls [vpls_name | all] | vpws [vpws_name | all]] service`. This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Disables the configured VPLS services for the specified *vpls_name*.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
all	Specifies all VPLS.

Default

Enabled.

Usage Guidelines

When services are disabled, the VPLS is removed from all peer sessions. The keyword **all** disables services for all VPLS instances.

Example

The following command disables the configured VPLS services for the specified VPLS:

```
disable vpls myvpls service
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.



enable bgp mpls-next-hop

```
enable bgp mpls-next-hop
```

Description

Enables IP forwarding over calculated MPLS LSPs to subnets learned via BGP.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables IP forwarding over calculated MPLS LSPs to subnets learned via BGP. (Calculated refers to an LSP that only reaches part of the way to the destination). IP forwarding over MPLS LSPs must be enabled to forward over calculated LSPs. By default, IP forwarding over MPLS LSPs to subnets learned via BGP is disabled.

Example

The following command enables BGP's use of MPLS LSPs to reach BGP routes:

```
enable bgp mpls-next-hop
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable iproute mpls-next-hop

```
enable iproute mpls-next-hop
```

Description

Enables IP forwarding over MPLS LSPs for the default VR.



Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables IP forwarding over MPLS LSPs for the default VR. When enabled, LSP next hops can be used to tunnel IP traffic across the MPLS network. By default, IP forwarding over MPLS LSPs is disabled.



Note

You can enable the use of LSP next hops, or you can enable DHCP/BOOTP relay. The software does not support both features at the same time.

Example

The following command enables IP forwarding over MPLS LSPs:

```
enable iproute mpls-next-hop
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable l2vpn

```
enable l2vpn [vpls [vpls_name | all] | vpws [vpws_name | all]]
```

Description

Enables the specified VPLS or VPWS.



Syntax Description

<code>vpls_name</code>	Identifies the VPLS within the switch (character string).
<code>vpws_name</code>	Identifies the VPWS within the switch (character string).
all	Specifies all VPLS or VPWS instances.

Default

All newly created VPLS or VPWS instances are enabled.

Usage Guidelines

When enabled, VPLS or VPWS attempts to establish sessions between all configured peers. Services must be configured and enabled for sessions to be established successfully.

The **12vpn** keyword is introduced in ExtremeXOS Release 12.4 and is required when enabling a VPWS. For backward compatibility, the **12vpn** keyword is optional when enabling a VPLS. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The following command enables the VPLS instance myvpls:

```
enable vpls myvpls
```

The following command enables the VPWS instance myvpws:

```
enable 12vpn vpws myvpws
```

History

This command was first available in ExtremeXOS 11.6.

The **12vpn** and **vpws** keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable 12vpn vpls fdb mac-withdrawal

```
enable 12vpn vpls fdb mac-withdrawal
```



Description

Enables the Layer 2 VPN MAC address withdrawal capability.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Use this command to enable FDB MAC withdrawal after it has been disabled.

Example

The following command enables MAC address withdrawal:

```
enable l2vpn vpls fdb mac-withdrawal
```

History

This command was first available in ExtremeXOS 12.1.

The `l2vpn` keyword was added in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable l2vpn health-check vccv

```
enable l2vpn [vpls vpls_name | vpws vpws_name] health-check vccv
```

Description

Enables the VCCV health check feature on the specified VPLS or VPWS.

Syntax Description

<code>vpls_name</code>	Identifies the VPLS for which health check is to be enabled.
<code>vpws_name</code>	Identifies the VPWS for which health check is to be enabled.



Default

Health check is disabled.

Usage Guidelines

Health check must be enabled on both ends of a PW to verify connectivity between two VPLS or VPWS peers. Both VCCV peers negotiate capabilities at PW setup. A single VCCV session monitors a single PW. Therefore, a VPLS with multiple PWs will have multiple VCCV sessions to multiple peers.

VCCV in EXOS uses LSP ping to verify connectivity.

The **l2vpn** keyword was introduced in ExtremeXOS Release 12.4 and is required when enabling health check for a VPWS instance. For backward compatibility, the **l2vpn** keyword is optional when enabling health check for a VPLS instance. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The following command enables the health check feature on the VPLS instance myvpls:

```
enable l2vpn vpls myvpls health-check vccv
```

History

This command was first available in ExtremeXOS 12.1.

The **l2vpn** and **vpws** keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable l2vpn service

```
enable l2vpn [vpls [vpls_name | all] | vpws [vpws_name | all]] service
```

Description

Enables the configured services for the specified VPLS or VPWS.



Syntax Description

<code>vpls_name</code>	Identifies the VPLS within the switch (character string).
<code>vpws_name</code>	Identifies the VPWS within the switch (character string).
all	Specifies all VPLS or VPWS instances.

Default

Enabled.

Usage Guidelines

When services are disabled, the VPLS or VPWS is withdrawn from all peer sessions. The keyword `all` enables services for all VPLS or VPWS instances.

The **l2vpn** keyword was introduced in ExtremeXOS Release 12.4 and is required when enabling services for a VPWS instance. For backward compatibility, the **l2vpn** keyword is optional when enabling services for a VPLS instance. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The following command enables the configured VPLS services for the specified VPLS instance:

```
enable l2vpn vpls myvpls service
```

History

This command was first available in ExtremeXOS 11.6.

The **l2vpn** and **vpws** keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable l2vpn sharing

```
enable l2vpn sharing
```

Description

Enables LSP sharing for Layer 2 VPN pseudo-wires .



Syntax Description

This command has no keywords or arguments.

Default

Disabled.

Usage Guidelines

This command enables LSP sharing for L2VPN PWs. When LSP sharing is enabled, up to 16 named LSPs are used for a PW. When LSP sharing is disabled, only 1 named LSP is used for a PW.

If LSP Sharing is disabled, and more than 1 Transport LSP is programmed into HW, all but 1 Transport LSP is removed from HW, and the configuration is preserved. If LSP Sharing is enabled, and more than 1 Transport LSP was previously configured, the remaining LSPs is programmed into HW as they become available for use.

Example

The following command enables LSP sharing for L2VPN PWs:

```
enable l2vpn sharing
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the ExtremeXOS Concepts Guide.

enable mpls

```
enable mpls
```

Description

Enables MPLS on the switch.

Syntax Description

This command has no arguments or variables.



Default

Disabled.

Usage Guidelines

Enabling MPLS allows MPLS processing to begin for any enabled MPLS protocols (RSVP-TE and/or LDP).

While MPLS is transitioning to the enabled state, only the MPLS show commands are accepted.

Before you can enable MPLS on BlackDiamond 8800 series switches, the switch must meet the following requirements:

- Each switch requires the specific software and hardware listed in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.
- You must configure the enhanced stacking protocol on the switch.
- If unsupported modules are installed in the switch, we recommend that you remove them and unconfigure the appropriate slots using the `unconfigure slot` command before enabling MPLS. You cannot enable MPLS until the appropriate slots are empty or disabled with the `disable slot` command.



Note

When MPLS is enabled on a BlackDiamond 8800 series switch, you cannot enable any modules using the `enable slot` command that are incompatible with MPLS.

Before you can enable MPLS on a SummitStack, the stack must meet the following requirements:

- Each stack switch must meet the software and hardware requirements listed in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.
- You must configure the enhanced stacking protocol on each Summit family switch.
- Although you can mix Summit X460 switches with Summit X480, X670, and X770 series switches in a SummitStack, we recommend that you do not mix these switch types if the desired routing table exceeds the supported limit for the Summit X460 switch, which is 12,256 IPv4 LPM routes.



Note

When MPLS is enabled on a stack, you can only add a MPLS-compatible Summit family switches to the stack.

Example

The following command globally enables MPLS on the switch:

```
enable mpls
```

History

This command was first available in ExtremeXOS 11.6.



Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable mpls bfd

```
enable mpls bfd [{vlan} vlan_name | vlan all]
```

Description

Enables the Bidirectional Forwarding Detection (BFD) client for MPLS on the specified VLAN or all VLANs.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN on which to enable the MPLS BFD client.
vlan all	Enables the MPLS BFD client on all VLANs.

Default

Disabled.

Usage Guidelines

This command causes MPLS to request a BFD session to each next-hop peer reachable through the named interface. BFD sessions are triggered by the establishment of an LSP over the interface. If this command is issued after LSPs are established, then the list of active LSPs is searched for next-hop peers associated with the named interface, and a BFD session is requested for each neighbor which does not already have a session. This command also instructs MPLS to begin to consider BFD neighbor session state updates as part of the effective interface link state reported to the MPLS upper layer protocols.



Note

BFD must be enabled on the interface before sessions can be established. To enable BFD, use the command: `[enable | disable] bfd vlan vlan_name .`

Example

The following command enables the MPLS BFD client on VLAN `vlan1`:

```
enable mpls bfd vlan1
```

History

This command was first available in ExtremeXOS 12.4.



Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable mpls exp examination

enable mpls exp examination

Description

Enables assigning an MPLS packet to a QoS profile based on the MPLS packet's EXP value.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables assigning an MPLS packet to a QoS profile based on the MPLS packet's EXP value. The EXP values to QoS profile mappings are configured using the `configure mpls exp examination` command.

When disabled, all received MPLS packets are assigned to QoS profile qp1.

Example

The following command enables assignment of an MPLS packet to a QoS profile based on the MPLS packet's EXP value:

```
enable mpls exp examination
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



enable mpls exp replacement

enable mpls exp replacement

Description

Enables setting an MPLS packet's EXP value based on the packet's QoS profile.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables setting an MPLS packet's EXP value based on the packet's QoS profile. The QoS profiles to EXP value mappings are configured using the **configure mpls exp replacement** command.

When disabled, all MPLS packets are transmitted with an EXP value of zero.

Example

The following command enables the setting of an MPLS packet's EXP value based on the packet's QoS profile:

```
enable mpls exp replacement
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

enable mpls ldp

enable mpls ldp [{vlan} *vlan_name* | **vlan all**]



Description

Enables LDP for the specified MPLS configured VLANs.

Syntax Description

vlan	Enables LDP for one or more specific VLANs.
<i>vlan_name</i>	Enables LDP on the specified VLAN.
vlan all	Enables LDP for all VLANs that have been added to MPLS.

Default

Disabled.

Usage Guidelines

When LDP is enabled, LDP attempts to establish peer sessions with neighboring routers on the enabled VLAN. Once a peer session is established, LDP advertises labels for local IP interfaces and for routes learned from other neighboring routers. By default, LDP is disabled for all VLANs that have been added to MPLS. Specifying the optional **all** keyword enables LDP for all MPLS configured VLANs.

Example

The following command enables LDP for all VLANs that have been added to MPLS:

```
enable mpls ldp vlan all
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable mpls ldp bgp-routes

```
enable mpls ldp bgp-routes
```

Description

Enables LDP to use IP prefixes learned from BGP when establishing LDP LSPs.



Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command allows LDP to use routes learned via BGP when establishing LDP LSPs. Because each established LSP consumes internal resources, it is recommended that this setting be used only in BGP environments where the number of BGP routes is controlled.

When disabled, LDP does not establish LSPs to routes learned via BGP, thus reducing the internal resources used by LDP. Note that MPLS LSPs can still be used to transport packets to routes learned via BGP through the use of the `enable bgp mpls-next-hop` command.

Example

The following command enables the use of BGP routes by LDP:

```
enable mpls ldp bgp-routes
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable mpls ldp loop-detection

```
enable mpls ldp loop-detection
```

Description

Enables LDP loop detection on the switch.

Syntax Description

This command has no arguments or variables.



Default

Disabled.

Usage Guidelines

Loop detection provides a mechanism for finding looping LSPs and for preventing Label Request messages from looping in the presence of non-merge capable LSRs. The mechanism makes use of Path Vector and Hop Count TLVs carried by Label Request and Label Mapping messages.

Example

The following command globally enables LDP loop detection on the switch:

```
enable mpls ldp loop-detection
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable mpls php

When enabled, PHP is requested on all LSPs advertised over that VLAN for which the switch is the egress LSR.

```
enable mpls php [ {vlan} vlan_name | vlan all ]
```

Description

Enables penultimate hop popping (PHP) on the specified VLAN.

Syntax Description

vlan	Enables PHP for one or more specific VLANs.
<i>vlan_name</i>	Enables PHP on the specified VLAN.
vlan all	Enables PHP for all VLANs that have been added to MPLS.

Default

Disabled.



Usage Guidelines

Penultimate hop popping is requested by assigning the Implicit Null Label in an advertised mapping. Extreme's MPLS implementation always performs penultimate hop popping when requested to do so by a peer LSR. When the all VLANs option is selected, PHP is enabled on all configured VLANs that have been added to MPLS.



Note

ExtremeWare always used the Implicit NULL Label in conjunction with routes exported into OSPF, regardless of the setting of the PHP configuration parameter. This method of operation is not utilized in ExtremeXOS.

Example

The following command enables penultimate hop popping (PHP) on the specified VLAN:

```
enable mpls php vlan vlan1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable mpls protocol ldp

```
enable mpls protocol ldp
```

Description

Enables LDP for the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.



Usage Guidelines

When LDP is enabled, LDP attempts to establish peer sessions with neighboring routers on VLAN interfaces where LDP has been enabled (see 6.3.5 page 77). Once a peer session is established, LDP can advertise labels for local IP interfaces and for routes learned from other neighboring routers. While LDP is transitioning to the enabled state, only the MPLS show commands are accepted.

Note that the LDP protocol must be enabled to establish VPLS pseudo-wires even if the transport LSPs are being established using RSVP-TE.

Example

The following command globally enables LDP on the switch:

```
enable mpls protocol ldp
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable mpls protocol rsvp-te

```
enable mpls protocol rsvp-te
```

Description

Enables RSVP-TE for the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

When RSVP-TE is enabled, configured LSPs begin the process of TE LSP establishment and VLAN interfaces that have RSVP-TE enabled begin processing RSVP path/reserve messages. By default, RSVP-TE is disabled. While RSVP-TE is transitioning to the enabled state, only the MPLS show commands are accepted.



Example

The following command globally enables RSVP-TE on the switch:

```
enable mpls protocol rsvp-te
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable mpls rsvp-te

```
enable mpls rsvp-te [{vlan} vlan_name | vlan all]
```

Description

Enables RSVP-TE for the specified MPLS-configured VLAN.

Syntax Description

vlan	Specifies that RSVP-TE is to be enabled on one or more VLANs.
<i>vlan_name</i>	Identifies a specific VLAN on which RSVP-TE is to be enabled.
vlan all	Enables RSVP-TE on all VLANs that have been added to MPLS.

Default

Disabled.

Usage Guidelines

When RSVP-TE is enabled, TE LSP establishment for configured LSPs begins and the processing of RSVP path/reserve messages from peer LSRs is permitted. By default, RSVP-TE is disabled for all MPLS-configured VLANs. Specifying the optional all keyword enables RSVP-TE for all VLANs that have been added to MPLS.



Example

The following command enables RSVP-TE on the specified VLAN:

```
enable mpls rsvp-te vlan vlan_1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable mpls rsvp-te bundle-message

```
enable mpls rsvp-te bundle-message [{vlan} vlan_name | vlan all]
```

Description

Enables the bundling of RSVP-TE messages for the specified VLAN interface.

Syntax Description

vlan	Specifies that message-bundling is to be enabled on one or more VLAN interfaces.
<i>vlan_name</i>	Identifies a VLAN interface for which message bundling is to be enabled.
vlan all	Specifies that message bundling is to be enabled on all VLANs that have been added to MPLS.

Default

Disabled.

Usage Guidelines

Enabling message bundling can improve control plane scalability by allowing the switch to bundle multiple RSVP-TE messages into a single PDU. Not all devices support bundled messages. If the switch determines that a peer LSR, connected to a specific interface, does not support message bundling, the switch reverts to sending separate PDUs for each message on that interface. By default, message bundling is disabled. Specifying the **all** keyword enables message bundling on all MPLS-configured VLANs.



Example

The following command enables message bundling on the specified VLAN:

```
enable mpls rsvp-te bundle-message vlan vlan_1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable mpls rsvp-te fast-reroute

```
enable mpls rsvp-te fast-reroute
```

Description

Enables the MPLS RSVP-TE fast reroute (FRR) protection feature.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

You can configure FRR LSPs only when FRR is enabled on the LSR. Enabling FRR protection on the LSR automatically enables the point-of-local-repair and merge-point capabilities on the LSR. FRR should be enabled on all LSRs along each FRR LSP path.

Example

The following command enables FRR protection on the local switch:

```
enable mpls rsvp-te fast-reroute
```



History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable mpls rsvp-te lsp

```
enable mpls rsvp-te lsp [lsp_name | all]
```

Description

Enables one or more RSVP-TE LSPs.

Syntax Description

<i>lsp_name</i>	Specifies the ingress LSP within the switch to be enabled.
all	Enables all RSVP-TE configured ingress LSPs.

Default

Enabled.

Usage Guidelines

When an RSVP-TE LSP is enabled, the switch attempts to set up the LSP by signaling the destination by sending a path message using the assigned path and profile. By default, all newly created LSPs are enabled and can become active when the LSP has been configured. Note that an LSP must be configured with at least one path before it can be signaled.

Example

The following command enables all RSVP-TE-configured LSPs:

```
enable mpls rsvp-te lsp all
```

History

This command was first available in ExtremeXOS 11.6.



Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable mpls rsvp-te summary-refresh

```
enable mpls rsvp-te summary-refresh [{vlan} vlan_name | vlan all]
```

Description

Enables the sending of summary refresh messages, instead of path messages, to refresh RSVP-TE path state for the specified VLAN interface.

Syntax Description

vlan	Specifies that summary refresh messages are to refresh the RSVP-TE path state on one or more VLAN interfaces.
<i>vlan_name</i>	Identifies a VLAN interface on which RSVP-TE summary refresh messages are to refresh the RSVP-TE path state.
vlan all	Specifies that summary refresh messages are to refresh the RSVP-TE path state on all VLANs that have been added to MPLS.

Default

Disabled.

Usage Guidelines

Enabling summary refresh can improve control plane scalability by refreshing multiple LSPs in a single message. Not all devices support summary refresh. If the switch determines that a peer LSR, connected to a specific interface, does not support summary refresh, the switch reverts to using path messages on that interface. By default, summary refresh is disabled. Specifying the **all** keyword enables summary refresh on all MPLS-configured VLANs.

Example

The following command enables summary refresh on the specified VLAN:

```
enable mpls rsvp-te summary-refresh vlan vlan_1
```

History

This command was first available in ExtremeXOS 11.6.



Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable mpls static lsp

```
enable mpls static lsp {lsp_name | all }
```

Description

Administratively enables one or all static LSPs.

Syntax Description

<i>lsp_name</i>	Identifies the LSP to be enabled.
all	Specifies that all static LSPs on this LSR are to be enabled.

Default

N/A.

Usage Guidelines

On executing this command, the software tries to activate the static LSP by programming the LSP in hardware. Static LSPs are not enabled by default. You need to explicitly enable LSPs after the ingress and egress segments have been configured.

Example

The following command enables a static LSP:

```
enable mpls static lsp lsp598
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



enable mpls vlan

```
enable mpls [{vlan}vlan_name|vlan all]
```

Description

Enables the MPLS interface for the specified VLAN.

Syntax Description

vlan	Enables an MPLS interface for one or more specific VLANs.
<i>vlan_name</i>	Enables an MPLS interface on the specified VLAN.
vlan all	Enables an MPLS interface for all VLANs that have been added to MPLS.

Default

Disabled.

Example

The following command enables an MPLS interface for the specified VLAN:

```
enable mpls vlan vlan-nyc
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

enable ospf mpls-next-hop

```
enable ospf mpls-next-hop {vr vrf_name}
```

Description

Enables IP forwarding over calculated MPLS LSPs to subnets learned through OSPF.

Syntax Description

<i>vrf_name</i>	Specifies OSPF on a particular VRF.
-----------------	-------------------------------------



Default

Disabled.

Usage Guidelines

This command enables IP forwarding over calculated MPLS LSPs to subnets learned through OSPF. (Calculated refers to an LSP that only reaches part of the way to the destination). By default, IP forwarding over MPLS LSPs to subnets learned via OSPF is disabled.

In order to configure OSPF on a particular VRF, you must supply the optional *vr* *vr-name* CLI parameter.

Example

The following command enables OSPF's use of MPLS LSPs to reach OSPF routes:

```
enable ospf mpls-next-hop
```

History

This command was first available in ExtremeXOS 11.6.

The **vr** keyword and *vrf_name* variable were added in ExtremeXOS 15.3.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

enable snmp traps l2vpn

```
enable snmp traps l2vpn
```

Description

Enables SNMP traps associated with Layer 2 VPNs for all MPLS configured VLANs.

Syntax Description

This command has no arguments or variables.

Default

All Layer 2 VPN traps are disabled.



Example

The following command enables SNMP traps associated with Layer 2 VPNs:

```
enable snmp traps l2vpn
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable snmp traps mpls

```
enable snmp traps mpls
```

Description

Enables SNMP traps associated with MPLS for all MPLS configured VLANs.

Syntax Description

This command has no arguments or variables.

Default

All MPLS traps are disabled.

Example

The following command enables SNMP traps associated with MPLS:

```
enable snmp traps mpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



enable vpls

```
enable vpls [vpls_name | all]
```

Note



This command has been replaced with the following command: `enable l2vpn [vpls [vpls_name | all] | vpws [vpws_name | all]]`. This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Enables the VPLS instance specified by *vpls_name*.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
all	Specifies all VPLS.

Default

All newly created VPLS instances are enabled.

Usage Guidelines

This command enables the VPLS instance specified by *vpls_name*. By default, all newly created VPLS instances are enabled. When enabled, VPLS attempts to establish sessions between all configured peers. Services must be configured and enabled for sessions to be established successfully.

Example

The following command enables the VPLS instance myvpls:

```
enable vpls myvpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.



enable vpls fdb mac-withdrawal

```
enable vpls fdb mac-withdrawal
```

Note



This command has been replaced with the following command: `enable l2vpn vpls fdb mac-withdrawal`. This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Enables the VPLS MAC address withdrawal capability.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Use this command to enable FDB MAC withdrawal after it has been disabled.

Example

The following command enables MAC address withdrawal:

```
enable vpls fdb mac-withdrawal
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



enable vpls health-check vccv

```
enable vpls vpls_name health-check vccv
```

Note



This command has been replaced with the following command: `enable l2vpn [vpls vpls_name | vpws vpws_name] health-check vccv`. This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Enables the VCCV health check feature on the specified VPLS.

Syntax Description

<code>vpls_name</code> Identifies the VPLS for which health check is to be enabled.

Default

Health check is disabled.

Usage Guidelines

Health check must be enabled on both ends of a PW to verify connectivity between two VPLS peers. Both VCCV peers negotiate capabilities at PW setup. A single VCCV session monitors a single PW. Therefore, a VPLS with multiple PWs will have multiple VCCV sessions to multiple peers.

VCCV in EXOS uses LSP ping to verify connectivity.

Example

The following command enables the health check feature on the VPLS instance myvpls:

```
enable vpls myvpls health-check vccv
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.



enable vpls service

```
enable vpls [vpls_name | all] service
```

Note



This command has been replaced with the following command: `enable l2vpn [vpls [vpls_name | all] | vpws [vpws_name | all]] service`. This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Enables the configured VPLS services for the specified *vpls_name*.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
all	Specifies all VPLS.

Default

Enabled.

Usage Guidelines

This command enables the configured VPLS services for the specified *vpls_name*. When services are disabled, the VPLS is withdrawn from all peer sessions. The keyword **all** enables services for all VPLS instances.

Example

The following command enables the configured VPLS services for the specified VPLS instance:

```
enable vpls myvpls service
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.



ping mpls lsp

```
ping mpls lsp [lsp_name | any host | prefix ipNetmask] {reply-mode [ip | ip-router-alert]} {continuous | count count} {interval interval} {start-size start-size {end-size end-size}} {ttl ttl} {{from from} {next-hop hopaddress}}
```

Description

Sends an MPLS ping packet to a FEC over an LSP.

Syntax Description

<i>lsp_name</i>	Specifies the LSP on which to send the MPLS echo request.
any	Allows the echo request to be sent over any available LSP.
<i>host</i>	Specifies the FEC using an ipaddress or hostname.
prefix	Specifies a prefix.
<i>ipNetmask</i>	Specifies the prefix address.
reply-mode	Specifies the return path for the MPLS echo response.
ip	Requests an IP UDP reply packet. This is the default mode.
ip-router-alert	Requests an IP UDP reply packet with the IP Router Alert option. If the reply is sent in an LSP, the router-alert label is inserted at the top of the label stack.
continuous	Sends pings continuously until the user intervenes.
<i>count</i>	Determines whether the size of the packet increments by one byte for each new MPLS echo request sent.
<i>interval</i>	Specifies the time interval (in seconds) between pings.
<i>start-size</i>	The number of payload data bytes in the MPLS ping packet. The range is from 1 - 1518 (if jumbo frames are disabled) and from 1 - the configured jumbo packet size (if jumbo frames are enabled). The default is 8 bytes.
<i>end-size</i>	Specifies that the size of the packet increments by one byte for each new MPLS echo request sent, up to the specified maximum size for the MPLS ping packet.
<i>ttl</i>	Sets the time-to-live value in the ping packet
<i>from</i>	Specifies the source IP address of the packet.
<i>hopaddress</i>	Specifies the next-hop address.

Default

Destination IP address for MPLS echo request - random, from the 127 and 128 IP address space
 IP TTL - 1
 TTL value in MPLS echo request - 255
 Destination UDP port - 3503
 Payload data packet size - 8 bytes
 Number of pings sent - 4



Usage Guidelines

This command sends an MPLS ping packet to a FEC over an LSP. The ping command, with `mpls` keyword option, can be used to verify data plane connectivity across an LSP. This is useful because not all failures can be detected using the MPLS control plane. The `lsp` keyword and `lsp_name` parameter may be used to specify the LSP on which to send the MPLS echo request. The `lsp` keyword along with the any keyword allows the echo request to be sent over any available LSP that terminates at host, specified as an ipaddress or hostname. If no LSP exists to the specified host, the ping command fails even though an IP routed path may exist. If the optional `next-hop` is specified, the MPLS echo request is sent along the LSP that traverses the specified node. This option is useful for specifying an LSP when multiple LSPs exist to the specified FEC. For RSVP-TE LSPs, the FEC is implied from the LSP configuration. The TTL value in the MPLS Echo Request is set to 255.

By default, the destination IP address of the MPLS echo request is randomly chosen from the 127/8 IP address space and the IP TTL is set to 1. The destination UDP port is 3503 and the sender chooses the source UDP port.

The optional `start-size` keyword specifies the number of bytes to include as payload data in the MPLS ping packet. If no `start-size` parameter is specified, the size of the payload data is eight bytes. The minimum valid `start-size` value is one. The maximum `start-size` value is variable, depending on the type of MPLS ping packet sent, but the total size of the MPLS ping packet cannot exceed the configured jumbo packet size, if jumbo frames are enabled, or 1518 if jumbo frames are disabled. If the `end-size` keyword is specified, the size of the packet increments by one byte for each new MPLS echo request sent. The next MPLS echo request is not sent until the MPLS echo response for the previous packet is received. This is useful for detecting interface MTU mismatch configurations between LSRs. The switch ceases sending MPLS echo requests when the specified `end-size` value is reached, the MPLS ping is user interrupted, or an MPLS echo response is not received after four successive retries.

The optional `reply-mode` keyword is used to specify the reply mode for the MPLS echo response. When the `ip` option is specified, the MPLS echo reply is routed back to the sender in a normal IPv4 packet. When the `ip-router-alert` option is specified, the MPLS echo reply is routed back to the sender in an IPv4 packet with the Router Alert IP option set. Additionally, if the `ip-router-alert` option is specified and the reply route is through an LSP, the Router Alert Label is pushed onto the top of the label stack. If the reply-mode is not specified, the `reply-mode ip` option applies.

Example

The following example shows a ping command and the resulting display:

```
ping mpls lsp prefix 11.100.100.212/32
Ping(MPLS) : 4 packets, 8 data bytes, interval 1 second(s).
98 bytes from 11.100.100.212: mpls_seq=0 ttl=64 time=6.688 ms
98 bytes from 11.100.100.212: mpls_seq=1 ttl=64 time=6.036 ms
98 bytes from 11.100.100.212: mpls_seq=2 ttl=64 time=6.218 ms
98 bytes from 11.100.100.212: mpls_seq=3 ttl=64 time=6.467 ms
--- ping statistics ---
4 packets transmitted, 4 received, 0% loss
round-trip min/avg/max = 6/6/6/ms
```



History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

restart process mpls

```
restart process mpls {msm slot}
```

Description

Restarts the MPLS process when it does not respond to the CLI commands.

Syntax Description

<i>slot</i>	Specifies the MSM/MM where process should be terminated and restarted. "A" specifies the MSM/MM installed in slot A, and "B" specifies the MSM/MM installed in slot B.
-------------	--

Default

N/A.

Usage Guidelines

None.

Example

The following command restarts the MPLS process:

```
restart process mpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



show bandwidth pool

```
show bandwidth pool [ingress | egress | duplex] vlan vlan_name
```

Description

Displays the configured bandwidth pool settings for the specified VLAN.

Syntax Description

ingress	Displays configured bandwidth pool settings for incoming traffic only.
egress	Displays configured bandwidth pool settings for outgoing traffic only.
duplex	Displays configured bandwidth pool settings for traffic in both directions.
<i>vlan_name</i>	Displays configured bandwidth pool settings only for the specified VLAN.

Default

N/A.

Usage Guidelines

This command displays the configured bandwidth pool settings for a VLAN. Values displayed include the VLAN, maximum reserveable bandwidth (both ingress and egress), and bandwidth reserved by application and by priority level.

Example

The following command displays bandwidth pool settings and accepted bandwidth reservations for all ports:

```
show bandwidth pool duplex vlan vlan_1
* BD-10K.1 # show bandwidth pool duplex vlan vlan_1
(mbps) Rsvd CIRBW Cmnt Cmnt CIRBW
Vlan      Dir  Phy   BE Limit Pools Total Avail
-----
vlan_1    Rx  1000    0  1000   300   300   700
Tx  1000    0  1000   500   500   500
-----
(mbps)  CIRBW Available in Pool (per priority level)
Appl Dir Pool    0    1    2    3    4    5    6    7
-----
mpls  Rx   300   300   300   290   290   290   290   290   290
      Tx   500   500   491   491   491   491   491   491   491
(Rx)-Receive, (Tx)-Transmit (BE)-Best Effort
```



History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show ces

```
show ces {ces_name} {detail}
```

Description

Displays CES services created for use with MPLS.

Syntax Description

<i>ces_name</i>	Specifies the name of the CES.
detail	Specifies to display additional status information about the interface.

Default

N/A.

Usage Guidelines

For CES services created for use with MPLS, the “Type: Static/Signaled” line in the CES section of the output will show “N/A” until a PW has been configured, since this the PW type is not known until the peer is added to the CES.

The PW section of the output includes a “PW Signaling” line that will display “LDP” or “None (Static)”, depending on the PW configuration. This is in keeping with the PW display in the `show l2vpn detail` command. Since the configured labels can be changed while the current labels are in-use, there is a small window where the configured labels and in-use labels are different. If the `show ces detail` command is issued during this window, an extra line is output to indicate this extra information. The configured labels will be noted as “pending” in this case.

Example

The following command displays CES information:

```
# show ces jces detail

CES PW Name   : jces
  Type                : Static
  PW ID              : 900
Admin State    : Enabled
```



```

PSN Transport          : MPLS          Oper State   : Enabled

Transport Type        : E1, Structured-agnostic
Signaling              : None
Service Name          : mysvc
Payload Size          : 256 bytes
Packet Latency        : 1000 us
Jitter Buffer          : 3000 us (max: 6000 us)
Clock Recovery        : None
LOPS Entry Thresh     : 8
LOPS Exit Thresh     : 8
Filler Pattern        : 255 (0xFF)
QOS Profile           : QP1
PW Label TTL          : 4
Peer Address          : 11.100.100.101
  PW Admin State      : Enabled
  PW State            : Up
  PW Signaling        : None (Static)
  PW Uptime           : 0d:0h:46m:19s
  PW Installed        : TRUE
  Local PW Status     : Fault, Att-Rx, PW-Rx
  Remote PW Status    : Fault, Att-Rx
  Transport LSP       : lsprsvp02 (Not Configured)
    Next Hop I/F      : e2-s4vlan1
    Next Hop Addr     : 9.21.1.243      Tx Label      :
0x23F
  PW Rx Label         : 0x9A          PW Tx Label     :
0x9F
  PW Rx Label (pend) : 0x9B          PW Tx Label (pend) :
0x9E
  Signaled Parameters :
    Payload Size      : Local=256          Remote=N/A
    Bit-Rate          : Local=32           Remote=N/A
    Fragmentation     : Local=Disabled     Remote=N/A

    Signaling         : Local=Disabled     Remote=N/A

    RTP Header        : Local=Disabled     Remote=N/A

    Signaling Error   : N/A
  PW Rx Pkts         : 0                  PW Tx Pkts     :
13907583

```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on the following platforms:

- E4G-200
- E4G-400



show l2vpn

```
show [ {l2vpn} vpls {vpls_name} | l2vpn vpws {vpws_name} | l2vpn ] {peer
ipaddress} {detail} | summary sharing }
```

Description

Displays Layer 2 VPN configuration and status information.

Syntax Description

l2vpn	Displays specified Layer 2 VPN information.
<i>vpls_name</i>	Displays information for the specified VPLS.
<i>vpws_name</i>	Displays information for the specified VPWS.
<i>ipaddress</i>	Specifies a Layer 2 VPN peer for which to display information.
detail	Displays additional information in comprehensive detail format.
summary	Displays summary information about all VPLS or VPWS instances.
sharing	Displays L2VPN Ethernet Pseudo Wire Tunnel LSP sharing configuration.

Default

N/A.

Usage Guidelines

The `show l2vpn` command (without any optional parameters) displays all currently configured Layer 2 VPN instances for the switch. The summarized list of Layer 2 VPN instances is displayed in alphabetical order based on the Layer 2 VPN name. Peers are displayed in the reverse of the order they were added.

When you specify a Layer 2 VPN peer, the display includes a list of all PWs established to the peer, the PW status and PW ID, and information about each Layer 2 VPN to which this peer belongs.

The following table describes the display fields that appear when this command is entered with the **detail** option.

Table 37: Selected show l2vpn Field Definitions

Field	Definition
VPLS Name	VPLS instance or domain name.
VPN ID	Virtual Private Network identifier.
Source Address	Source IP address.
VCCV Status	Virtual Circuit Connectivity Verification (VCCV) feature status, which is either Enabled or Disabled .
VCCV Interval Time	Displays the configured VCCV interval time.



Table 37: Selected show l2vpn Field Definitions (continued)

Field	Definition
VCCV Fault Multiplier	Displays the configured VCCV fault multiplier.
Redundancy Type	Displays the configured VPLS redundancy type, which is EAPS , ESRP , STP , or None .
Service Interface	Displays a VLAN or VMAN interface name.
Admin State	Displays the administrative state of the VPLS, which is either Enabled or Disabled .
Oper State	Displays the operational state of the VPLS, which is either Enabled or Disabled .
MTU	Displays the maximum transmission unit (MTU) size for the VPLS.
Ethertype	Displays the ethertype for the service interface.
.Iq tag	Displays the 802.Iq priority tag for the VPLS.
Peer IP	Displays the IP address for the VPLS peer.
PW State	<p>PW State represents the state, or status, of a PW. The possible PW state values are:</p> <ul style="list-style-type: none"> • UP—The PW is fully operational and installed in hardware. Traffic is forwarded over PW and VPLS service VLAN/VMAN. • Down—The PW is not operational and is not installed in hardware. This only happens when the VPLS instance is disabled, VPLS service is disabled, or there is no service VLAN assigned to the VPLS. No traffic is forwarded. • Sgnl—The PW is in a signalling state. The PW is not operational, and no traffic is forwarded. This can occur for a number of reasons, including: No LDP adjacency to peer, No transport LSP to peer, No VC LSP to peer. • Remote peer not configuredReady—The PW has been signalled, but it has not been installed in hardware. Traffic is not forwarded. The PW can be in a Ready state for a number of reasons, including: <ul style="list-style-type: none"> • The VPLS instance is configured for EAPS redundancy, and the EAPS shared port associated with this VPLS instance is Connected. • The VPLS instance is configured for ESRP redundancy, and the ESRP domain associated with this VPLS instance is Slave. • The service VLAN associated with this VPLS instance is down. • The remote peer has signalled that it has a fault (remote PW status). The remote peer may have a fault due to its service VLAN being down.
PW Uptime	PW Uptime is the elapsed time that the PW has been in the UP state.
PW Installed	PW Installed is a flag to indicate whether the PW is installed in hardware or not. If the PW is in the UP state, this field is True, otherwise, this field is False.



Table 37: Selected show l2vpn Field Definitions (continued)

Field	Definition
Local PW Status	Local PW Status displays the VC status of the local PW. The values are: <ul style="list-style-type: none"> • No Faults—No faults detected. • PW-Tx—Local PSN-facing PW transmit fault. This is set if there is a problem with the VPLS transport LSP. • PW-Rx—Local PSN-facing PW receive fault. This is set if there is a problem with the VPLS transport LSP. • Att-Tx—Local attachment circuit transmit fault. This is set if there is a problem with the VPLS service VLAN. • Att-Rx—Local attachment circuit receive fault. This is set if there is a problem with the VPLS service VLAN. • Not Forwarding—The local PW is not forwarding. Look for more information in the PW State field. For example, if VPLS is configured for EAPS redundancy, the Local PW Status is Not Forwarding and the PW State is Ready whenever the EAPS Shared Port state is Connected.
Remote PW Status	Remote PW Status is the VC status of the remote PW. The values for this field are the same values as for Local PW Status.
PW Mode	PW Mode describes how the PW was configured. The values are: <ul style="list-style-type: none"> • Core-to-Core—This VPLS instance is a core node, and the other end of the PW connects to a core node. • Core-to-Spoke—This VPLS instance is a core node, and the other end of the PW connects to a spoke node. This is for HVPLS. • Spoke-to-Core—This VPLS instance is a spoke node, and the other end of the PW connects to a core node. This is for HVPLS.
Transport LSP	Transport LSP is the LSP that is used to forward frames over the PW. When an LDP LSP is used as a transport, the display shows LDP LSP (Not configured). If an RSVP LSP is used, the name of the RSVP LSP being used as a transport LSP is displayed. An RSVP LSP can be specified as the LSP to use during VPLS configuration.
Next Hop I/F	Displays the interface name for the next hop router.
Next Hop Addr	Displays the interface IP address for the next hop router.
PW Rx Label	Receive label for the VPLS PW.
PW Rx Pkts	Total packets received on the VPLS PW.
PW Rx Bytes	Total bytes received on the VPLS PW.
Tx Label	Transmit label for the LSP.
PW Tx Label	Transmit label for the VPLS PW.
PW Tx Pkts	Total packets transmitted on the VPLS PW.
PW Tx Bytes	Total bytes transmitted on the VPLS PW.

The **l2vpn** keyword is introduced in ExtremeXOS Release 12.4 and is required when displaying VPWS information. For backward compatibility, the **l2vpn** keyword is optional when displaying VPLS information. However, this keyword will be required in a future release, so Extreme Networks recommends that you use this keyword for new configurations and scripts.



Example

The following example shows the display that appears when you enter the `show l2vpn` command without any options:

```
Switch.1 # show l2vpn
```

L2VPN Name	VPN ID	Flags	Services Name	Peer IP	State	Flags
Pws-3344	20	EAX--W	NONE			
jwcvpls	99	EAX--L	torix	11.100.100.219	Up	C--NV-
keeper	90	EAX--L	NONE			
pws-1	2009	EAX--W	pwserve	11.100.100.219	Up	----V-
pws-10	70	EAX--W	NONE			
pws-2	2008	EAX--W	pw2serve	11.100.100.219	Up	---NV-
pws-3	2007	EAX--W	NONE			
sarsparilla	80	EAX--W	NONE			
whoopwoo	100	EAX--L	NONE	11.100.100.219	Down	C--N--

VPN Flags: (E) Admin Enabled, (A) Oper Active, (I) Include Tag,
(X) Exclude Tag, (T) Ethertype Configured,
(V) VCCV HC Enabled, (W) VPN Type VPWS, (L) VPN Type VPLS
Peer Flags: (C) Core Peer, (S) Spoke Peer, (A) Active Core,
(p) Configured Primary Core, (s) Configured Secondary Core,
(N) Named LSP Configured, (V) VCCV HC Capabilities Negotiated,
(F) VCCV HC Failed

```
-----
Total number of configured L2VPNs:      9
Total number of active L2VPNs:         3
Total number of configured PWs:        4
Total number of active PWs:            3
PWs auto-selecting transport LSP:      1
PWs configured with a transport LSP:    3
PWs using LDP for transport:            0
PWs using RSVP for transport:          4
PWs using static for transport:         0
```

The following command shows summary Layer 2 VPN information for VPLS peer 2.2.2.2:

```
Switch.451 # show l2vpn vpls peer 2.2.2.2
```

L2VPN Name	VPN ID	Flags	Services Name	Peer IP	State	Flags
vs1	105	EAX--L	cust1	2.2.2.2	UP	CAp-V-

VPN Flags: (E) Admin Enabled, (A) Oper Active, (I) Include Tag,
(X) Exclude Tag, (T) Ethertype Configured,
(V) VCCV HC Enabled, (W) VPN Type VPWS, (L) VPN Type VPLS
Peer Flags: (C) Core Peer, (S) Spoke Peer, (A) Active Core,
(p) Configured Primary Core, (s) Configured Secondary Core,
(N) Named LSP Configured, (V) VCCV HC Capabilities Negotiated,
(F) VCCV HC Failed



The following command shows detailed Layer 2 VPN information for VPLS peer 11.100.100.210:

```
BD-12804.452 # show l2vpn vpls peer 11.100.100.210 detail

VPLS Name : vpls10

VPN ID           : 10                Admin State      : Enabled
Source Address   : 11.100.100.212    Oper State       : Enabled
VCCV Status      : Disabled          MTU              : 1500
VCCV Interval Time : 5 sec.         Ethertype        : 0x8100
VCCV Fault Multiplier : 4           .1q tag         : exclude
L2VPN Type       : VPLS              Redundancy       : None
Service Interface : vlan10

Peer IP : 11.100.100.210
PW State      : Up
PW Uptime     : 18d:0h:28m:26s
PW Installed  : True
Local PW Status : No Faults
Remote PW Status : No Faults
PW Mode       : Core-to-Core
Transport LSP : LDP LSP (Not Configured)
Next Hop I/F  : o6vlan1
Next Hop Addr : 12.182.0.216          Tx Label        : 0x00010
PW Rx Label   : 0x80405              PW Tx Label     : 0x80401
PW Rx Pkts    : 3806161633          PW Tx Pkts     : 4294967296
PW Rx Bytes   : 912385942           PW Tx Bytes    : 4294967296
MAC Limit     : No Limit
VCCV HC Status : Not Sending (VCCV Not Enabled For This VPLS)
CC Type       : Rtr Alert             Total Pkts Sent : 0
CV Type       : LSP Ping              Total Pkts Rcvd : 0
Send Next Pkt : --

Total Failures : 0                Pkts During Last Failure : 0
Last Failure Tm : --
```

The following command shows the Layer 2 sharing information:

```
show l2vpn sharing
L2VPN Sharing      : Enabled
```

The non-detail version of this command now includes a peer flag that indicates the signaling protocol, if any, for a Pseudo Wire/peer. An “L” indicates LDP is used to signal the PW. A “T” indicates that no signaling is done, and therefore, this is a static PW.

The detail version of this command now displays a “PW Signaling” line that will display “LDP” or “None (Static)”, depending on the PW configuration. The “Local PW Status” will show “--” instead of “Not Signaled”, since the PW status is not currently signaled; however, for informational purposes, any local faults are still shown. The “Remote PW Status” and “Remote I/F MTU” will always show “--”.

Since the configured labels can be changed while the current labels are in-use, there is a small window where the configured labels and in-use labels are different. If you issue the `show l2vpn detail`



command during this window, an extra line will be output to indicate this extra information. The configured labels are noted as “pending” in this case.

```

show l2vpn vpls red04 detail
L2VPN Name: red04
  VPN ID           : 104           Admin State      :
Enabled
  Source Address   : 11.100.100.102 Oper State       :
Enabled
  VCCV Status      : Disabled      MTU              : 1500
  VCCV Interval Time : 5 sec.      Ethertype        :
0x8100
  VCCV Fault Multiplier : 4           .lq tag         :
exclude
  L2VPN Type       : VPLS          Redundancy       : None
  Service Interface : red04svc

Peer IP: 11.100.100.244
  PW State         : Up
  PW Signaling     : None (Static)
  PW Uptime        : 0d:0h:57m:32s
  PW Installed     : True
  Local PW Status  : No Faults
  Remote PW Status : --
  Remote I/F MTU   : --
  PW Mode          : Core-to-Core
  Transport LSP    : LDP LSP (Not Configured)
    Next Hop I/F   : e2-s4vlan1
    Next Hop Addr  : 9.21.1.243           Tx Label        :
0x00248
  PW Rx Label     : 0x0002c           PW Tx Label     :
0x000c2
  PW Rx Label (pend) : 0x0002b           PW Tx Label (pend) :
0x000b2
  PW Rx Pkts      : 0                 PW Tx Pkts      : 0
  PW Rx Bytes     : 0                 PW Tx Bytes     : 0
  MAC Limit       : No Limit
  VCCV HC Status  : Not Sending (VCCV Not Supported For Static
PWs)
  CC Type         : --                 Total Pkts Sent :
0
  CV Type         : --                 Total Pkts Rcvd :
0
  Send Next Pkt   : --
  Total Failures  : 0                 Pkts During Last Failure :
0
  Last Failure Tm : --

```

History

This command was first available in ExtremeXOS 11.6.

This command was updated to display flags for H-VPLS spoke nodes and protected VPLS and H-VPLS in ExtremeXOS 12.1.

The output for this command was modified in ExtremeXOS 12.2.2.



The `l2vpn` and `vpws` keywords were first available in ExtremeXOS 12.4.

The `sharing` keyword was added in ExtremeXOS 15.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show mpls

show mpls

Description

Displays MPLS configuration information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The `show mpls` command displays the current values for all the MPLS configuration parameters that apply to the entire switch. The parameters displayed include:

- MPLS and MPLS protocol (RSVP-TE and LDP) status.
- SNMP traps configuration.
- EXP examination/replacement configuration.
- the MPLS LSR ID.
- the list of the VLANs which have been added to MPLS.

Example

The following command displays the MPLS configuration parameters for the switch:

```
show mpls
* BD-10K.30 # show mpls
MPLS System
MPLS System           : User VR (Green)
MPLS Admin Status     : Enabled
MPLS Oper Status      : Enabled
RSVP-TE Admin Status  : Enabled
RSVP-TE Oper Status   : Enabled
```



```

LDP Admin Status      : Enabled
LDP Oper Status       : Enabled
SNMP Traps            : Disabled
L2VPN SNMP Traps     : Enabled
EXP Examination       : Disabled
EXP Replacement       : Disabled
LSR ID                 : 11.100.100.20

```

History

This command was first available in ExtremeXOS 11.6

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show mpls bfd

```
show mpls bfd [{vlan} vlan_name | ip_addr]
```

Description

Displays MPLS BFD client information for a VLAN or interface.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN for which to display MPLS BFD client information.
ip_addr	Specifies the IP address of an interface for which to display MPLS client information.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the MPLS BFD client information for all next-hop peers:

```

Switch.1 # show mpls bfd
Next Hop IP      Count  Flags  Admin  Oper  IfName
-----
192.84.86.2      13   ASIU   Up     Up    vlan1

```



```
192.84.93.12      13  ASIU      Up    Up    vlan2
Flags: A=Session added to BFD server, S=BFD Server synced,
I=Session Init complete, U=State Updates accepted
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show mpls exp examination

```
show mpls exp examination
```

Description

Displays MPLS EXP value to QoS profile mappings and whether MPLS EXP examination is enabled or disabled.

Syntax Description

This command has no arguments or keywords.

Default

N/A.

Usage Guidelines

This command displays MPLS EXP value to QoS profile mappings and the status of MPLS EXP examination (enabled or disabled). These values are set using the `configure mpls exp examination qosprofile` command and can be reset using the `unconfigure mpls exp examination` command.

Example

The following is an example of the output of the `show mpls exp examination` command:

```
* BD-10K.29 # show mpls exp examination
EXP --> QoS Profile mapping:
00 --> QP1
01 --> QP2
02 --> QP3
03 --> QP4
```



```
04 --> QP5
05 --> QP6
06 --> QP7
07 --> QP8
EXP Examination is disabled
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

show mpls exp replacement

show mpls exp replacement

Description

Displays the MPLS QoS profile to EXP value mappings and the status of MPLS EXP replacement (enabled or disabled).

Syntax Description

This command has no arguments or keywords.

Default

N/A.

Usage Guidelines

This command displays MPLS QoS profile to EXP value mappings and the status of MPLS EXP replacement (enabled or disabled). These values are set using the `configure mpls exp replacement qosprofile` command and can be reset using the `unconfigure mpls exp replacement` command.

Example

The following is an example of the output of the `show mpls exp replacement` command:

```
* BD-10K.28 # show mpls exp replacement
QoS Profile --> EXP mapping:
QP1 --> 00
QP2 --> 01
```



```

QP3 --> 02
QP4 --> 03
QP5 --> 04
QP6 --> 05
QP7 --> 06
QP8 --> 07
EXP Replacement is disabled

```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

show mpls interface

```
show mpls interface {{vlan} vlan_name} {detail}
```

Description

Displays the MPLS interface information. Information is displayed in tabular format for all VLANs that have been added to MPLS.

Syntax Description

vlan	Specifies to display information for one VLAN.
<i>vlan_name</i>	Specifies the name of the VLAN.
detail	Specifies to display additional status information about the interface.

Default

N/A.

Usage Guidelines

Information displayed includes:

- a list of all VLANs added to MPLS.
- MTU size.
- Local interface IP address.
- Number of RSVP-TE neighbors.
- Number of LDP adjacencies.
- RSVP-TE and LDP uptimes.
- MPLS protocols and capabilities configured on each VLAN.



Specifying the optional detail keyword displays the information in verbose form and also includes the operational state for RSVP-TE and LDP. Specifying a VLAN limits the output to that of the individual VLAN.

Example

The following command display MPLS interface information:

```
Switch # show mpls interface
Local
VLAN Name          RSVP-TE          LDP
IP Address          MTU    UpTm #Nbr UpTm #Adj  Flags
-----
loopb               11.100.100.218  1500   3h   0   3h   0   MRL-I-U
toratora           192.84.86.1     1500   3h   0   3h   1   MRL-IBU
tordoze            192.84.93.1     1500   3h   0   3h   1   MRL-IbU
torfour            192.84.83.1     1500   --   0   --   0   MRL-I--
torinasp           192.84.85.1     1500   --   0   --   0   MRL-I--
Flags: (M) MPLS Enabled, (R) RSVP-TE Enabled, (L) LDP Enabled,
(P) PHP Enabled, (I) IP Forwarding Enabled, (B) BFD Enabled,
(b) BFD Disabled(Sessions Exist) (U) MPLS Operational
```

The following command displays detailed MPLS information for VLAN 1:

```
* BD-10K.27 # show mpls interface vlan1 detail
VLAN Name : vlan1
Local IP Address      : 192.84.86.1
IP Forwarding         : Enabled
MPLS I/F MTU         : 1500
PHP Status            : Disabled
BFD Status            : Enabled
MPLS Admin Status    : Enabled
MPLS Oper Status     : Disabled
RSVP-TE Admin Status : Enabled
Oper Status          : Disabled
UpTime               : 0d:0h:0m:0s
# Neighbors          : 0
LDP Admin Status     : Enabled
Oper Status          : Disabled
UpTime               : 0d:0h:0m:0s
# Link Adjacencies   : 0
```

History

This command was first available in ExtremeXOS 11.6.

The BFD flags and status were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



show mpls label

```
show mpls {rsvp-te | static} label {summary | label_num | [advertised | received]
{label_num} | received implicit-null}
```

Description

Displays label information for all label types and protocols.

Syntax Description

rsvp-te	Specifies that only RSVP-TE LSP labels are displayed.
static	Specifies that only static LSP labels are displayed.
summary	Specifies that summary information for the labels is displayed.
<i>label_num</i>	Specifies that only labels that match the specified hexadecimal label number are displayed.
advertised	Specifies that only labels advertised to other routers are displayed.
received	Specifies that only labels received from other routers are displayed.
implicit-null	Specifies that only implicit-null labels are displayed.

Default

N/A.

Usage Guidelines

If no options are specified, tabular information for all labels, except advertised implicit-null labels, is displayed. The following table describes the display fields that appear when this command is entered.

Table 38: show mpls label Field Definitions

Field	Definition
Advertised Label	Advertised Label is the label advertised to other routers.
Destination Mapping	Destination Mapping displays the destination IP address and mask.
LSP Flags	LSP Flags lists the flags for each label, which are described in the key at the bottom of the display.
Peer Label	The peer label that replaces the received label of incoming packets during a label swap on a transit LSP. The peer label appears only for transit LSPs.



Table 38: show mpls label Field Definitions (continued)

Field	Definition
NHop Type	NHop Type helps define the handling of a packet arriving with the specified advertised label. Possible values are: IP—The label on the packet is swapped and the packet forwarded to the IP address shown in the NextHop field.loc1—Indicates that the tunnel destination is local to this switch.VLAN—The label on the packet is stripped and is IP routed according to the Destination Mapping field.VRF—For advertised labels, the Next Hop column contains the name of the virtual router to which packets with the givenLayer 3 VPNlabel will be forwarded. For received labels, the Next Hop column displays the router ID of the BGP peer, and the Name column displays the name of the VR using this label.VPLS—The label on the packet is stripped and forwarded according to the service configuration of the VPLS specified in the Name field.
Next Hop	Displays an interface ID for the next hop router, which can be a VLAN name, VMAN name, or IP address.
Name	Displays a VPLS or LSP name.
Received Label	Received Label displays the label received from other routers.

**Note**

Unsupported labels will contain no information.

Example

The following command displays all labels except received implicit-null labels:

```
* Switch.1 # show mpls label
```

Advertised Label	Destination Mapping	LSP Flags	Peer Label	NHop Type	NextHop	Name
-						
0x8082c	3.3.3.3/32	-LE	--	VLAN	lpbk	--
0x8082d	5.5.5.5/32	2LE	--	VPLS	--	
extreme-501						
0x8082e	5.5.5.5/32	2LE	--	VPLS	--	
extreme-503						
0x8082f	5.5.5.5/32	2LE	--	VPLS	--	
extreme-504						
0x80830	5.5.5.5/32	2LE	--	VPLS	--	
extreme-505						
0x80831	5.5.5.5/32	2LE	--	VPLS	--	
extreme-506						
0x80832	5.5.5.5/32	2LE	--	VPLS	--	
extreme-507						
0x80833	5.5.5.5/32	2LE	--	VPLS	--	
extreme-508						
0x80834	5.5.5.5/32	2LE	--	VPLS	--	
extreme-509						
0x8082a	3.3.3.3/32	-RE	--	loc1	3.3.3.3	lsp5to3-2
0x8082b	101.0.0.1/32	-RE	--	loc1	101.0.0.1	lsp5to3
0x80400	--	3-E	--	VRF	blue-vr	--
0x80401	--	3-E	--	VRF	red-vr	--



```
0x80402          --      3-E          -- VRF   white-vr          --
```

Received Label	Destination Mapping	LSP Flags	NextHop	Name
0x8082d	5.5.5.5/32	2LI	5.5.5.5	extreme-501
0x8082f	5.5.5.5/32	2LI	5.5.5.5	extreme-503
0x80830	5.5.5.5/32	2LI	5.5.5.5	extreme-504
0x80831	5.5.5.5/32	2LI	5.5.5.5	extreme-505
0x80832	5.5.5.5/32	2LI	5.5.5.5	extreme-506
0x80833	5.5.5.5/32	2LI	5.5.5.5	extreme-507
0x80834	5.5.5.5/32	2LI	5.5.5.5	extreme-508
0x80835	5.5.5.5/32	2LI	5.5.5.5	extreme-509
0x8082a	101.0.0.2/32	-RI	101.0.0.2	lsp3to5
0x8082b	5.5.5.5/32	-RI	101.0.0.2	lsp3to5-2
0x80401	--	3-I	5.5.5.5	red-vr
0x80400	--	3-I	5.5.5.5	red-vr
0x80401	--	3-I	5.5.5.5	blue-vr
0x80400	--	3-I	5.5.5.5	blue-vr
0x80402	--	3-I	5.5.5.5	white-vr

Flags:(3) L3VPN,(2) L2VPN, (L) LDP, (R) RSVP-TE, (S) Static
 (T) Transit LSP, (I) Ingress to LSP, (E) Egress from LSP,
 (M) Multiple Next Hops

Summary of Labels	Advertised	Received
Total number of RSVP-TE LSP labels	2	2
Total number of LDP LSP labels	1	0
Total number of Static LSP labels	0	0
Total number of L2VPN Labels	8	8
Total number of L3VPN Labels	3	5

The following command displays all rsvp-te labels except received implicit-null labels:

```
* Switch.2 # show mpls rsvp-te label
```

Advertised Label	Destination Mapping	Label Flags	Peer Label	NHop Type	NextHop	Name
-						
0x80834	101.0.0.1/32	-RE	--	loc1	101.0.0.1	lsp5to3
0x80835	3.3.3.3/32	-RE	--	loc1	3.3.3.3	lsp5to3-2

Received Label	Destination Mapping	Label Flags	NextHop	Name
0x8082a	101.0.0.2/32	-RI	101.0.0.2	lsp3to5
0x8082b	5.5.5.5/32	-RI	101.0.0.2	lsp3to5-2

Flags:(3) L3VPN,(2) L2VPN, (L) LDP, (R) RSVP-TE, (S) Static,
 (T) Transit LSP, (I) Ingress to LSP, (E) Egress from LSP,
 (M) Multiple Next Hops

Summary of Labels	Advertised	Received



Total number of RSVP-TE LSP labels	2	2
------------------------------------	---	---

History

This command was first available in ExtremeXOS 11.6.

This command was modified to display only RSVP-TE and static label information in ExtremeXOS 12.5. Additional commands were added to display LDP Layer 2 VPN and Layer 3 VPN labels.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show mpls label usage

```
show mpls label usage {static}
```

Description

Displays the label ranges on the current running system, including configurable and non-configurable ranges.

Syntax Description

usage	Displays MPLS label ranges on the current running system, including configurable and non-configurable ranges.
static	Displays MPLS label usage of static labels.

Default

N/A.

Usage Guidelines

With the addition of the static PW configuration, there is the need to configure static labels and display more detailed label information. This command displays the label ranges on the current running system, including configurable and non-configurable ranges. The output also includes hardware resource usage to provide better information about MPLS hardware utilization and capacity.



Example

The following command displays label ranges in-use, information about the static label usage, and some of the label hardware usage:

```
* show mpls lab usage
Label Type           Size      Label
Range
-----
Supported           1048576   0x00000 - 0xfffff (0 -
1048575)
Reserved              16       0x00000 - 0x0000f (0 -
15)
Static                300       0x00010 - 0x0013b (16 -
315)
L3VPN                 255       0x0013c - 0x0023a (316 -
570)
Dynamic              7365       0x0023b - 0x01eff (571 -
7935)
Internal Use         256       0x01f00 - 0x01fff (7936 -
8191)

Static Label Configuration Usage
-----
                In-Use  Avail  Total  %Avail
                -----  -----  -----  -----
Total                24    276    300    92%
  Ingress LSP         2
  Egress LSP          0
  Transit LSP         2
  L2VPN               17
  CES                  3

Label Hardware Resource Usage
-----
                                Incoming
Outgoing
-----
                In-Use  Avail  Total  %Avail  In-Use  Avail  Total
%Avail
-----
Static Ingress LSP    -      -      -      -      2      0
0      0%
Static Transit LSP    0     278    300    92%    0      0
0      0%
Static Egress LSP     2     278    300    92%    -      -
-
Static L2VPN          17    278    300    92%    0      0
0      0%
Static CES             3     278    300    92%    2      0
0      0%
RSVP-TE Ingress LSP  -      -      -      -      5      0
0      0%
RSVP-TE Transit LSP  0     7342   7365   99%    0      0
0      0%
RSVP-TE Egress LSP   1     7342   7365   99%    -      -
```



```

-      -
LDP Ingress LSP      -      -      -      -      5      0
0      0%
LDP Transit LSP     5      7342     7365     99%     5      0
0      0%
LDP Egress LSP      1      7342     7365     99%     -      -
-      -
LDP L2VPN           14      7342     7365     99%     14     0
0      0%
LDP CES             2      7342     7365     99%     2      0
0      0%
L3VPN              0      255      255     100%    0      0
0      0%

```

Here is an example of the **static** option:

```

Summit1.2 # show mpls label usage static
Inbound Label      Name                               Peer/Dest                Flags
-----
0x22 (34)          green01                           11.100.100.241          2
0x23 (35)          green01                           11.100.100.243          2
0x24 (36)          green01                           11.100.100.244          2
0x27 (39)          green03                           11.100.100.241          2
0x28 (40)          green03                           11.100.100.243          2
0x31 (49)          red02                             11.100.100.101          2
0x32 (50)          red02                             11.100.100.102          2
0x34 (52)          red02                             11.100.100.241          2
0x35 (53)          red02                             11.100.100.242          2
0x36 (54)          red02                             11.100.100.243          2
0x37 (55)          red02                             11.100.100.244          2
0x3d (61)          red05                             11.100.100.241          2
0x3f (63)          red05                             11.100.100.244          2
0x81 (129)         red04                             11.100.100.101          2
0x82 (130)         red04                             11.100.100.102          2
0x84 (132)         red04                             11.100.100.241          2
0x85 (133)         red04                             11.100.100.242          2
0x86 (134)         red04                             11.100.100.243          2
0x87 (135)         red04                             11.100.100.244          2
0xd1 (209)         red05                             11.100.100.101          2
0x104 (260)        lspstat02                         11.100.100.101          L
0x10d (269)        lspstat04                         11.100.100.102          L

```

Flags: (2) L2VPN, (C) CES, (L) Static LSP

Static Label Range: 0x10 - 0x13b (16 - 315)
Static Label Statistics: 300 total, 22 in-use, 278 (92%) available

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



show mpls ldp

show mpls ldp

Description

Displays summary configuration and status information for LDP. Global status of LDP, LDP session timer configuration, loop detection, and label advertisement status are included in the display output.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the global status of LDP, LDP session timer configuration, loop detection, label advertisement, and LDP-enabled VLANs.

The following table describes the display fields that appear when this command is entered.

Table 39: show mpls ldp Field Definitions

Field	Definition
LDP Admin Status	LDP Admin Status shows whether LDP has been administratively Enabled or Disabled .
LDP Oper Status	LDP Oper Status shows whether LDP is operating (Enabled) or not (Disabled).
Protocol Version	Protocol Version specifies the LDP protocol version number.
Label Retention Mode	Label Retention Mode is either Conservative or Liberal , as described in RFC 3036, LDP Specification. The ExtremeXOS software only supports Liberal mode. In Liberal mode, a label switch router maintains all received label-to-FEC mapping advertisements.
Label Distribution Method	Label Distribution Method is either Downstream Unsolicited or Downstream On Demand , as described in RFC 3036, LDP Specification. The ExtremeXOS software only supports Downstream Unsolicited mode. In Downstream Unsolicited mode, a label switch router distributes label bindings to peer label switch routers without waiting for the bindings to be requested.
Label Distribution Control Mode	Label Distribution Control Mode is either Independent or Ordered , as described in RFC 3036, LDP Specification. The ExtremeXOS software only supports Ordered mode. In Ordered control mode, a label switch router only creates a binding of a label-to-FEC when it is the egress router or it has received a label binding for that FEC from the next hop router for that FEC.



Table 39: show mpls ldp Field Definitions (continued)

Field	Definition
LDP BGP LSPs	LDP BGP LSPs shows whether LDP uses BGP routes. When the displayed value is Enabled , LDP accepts BGP routes and stores them in the LDP internal routing table. When the displayed value is Disabled , LDP does not accept BGP routes, which reduces memory requirements when LSPs based on BGP routes are not desired. Note that when Disabled is displayed, no LDP LSPs are established to prefixes for which BGP is the preferred routing protocol.
LDP Loop Detection	LDP Loop Detection displays the LDP loop detection configuration.
LDP Targeted Timers	LDP Targeted Timers displays the LDP timer configuration used for LDP targeted adjacencies and sessions.
LDP Link Timers	LDP Link Timers displays the LDP timer configuration used for LDP link adjacencies and sessions.

Example

The following command displays summary configuration and status information for LDP:

```
* BD-10K.25 # show mpls ldp
LDP Admin Status           : Enabled
LDP Oper Status            : Enabled
Protocol Version           : v1*
Label Retention Mode       : Liberal*
Label Distribution Method   : Downstream Unsolicited*
Label Distribution Control Mode : Ordered*
LDP BGP LSPs               : Enabled
LDP Loop Detection
Status                      : Disabled
Hop-Count Limit            : 255
Path-Vector Limit         : 255
LDP Targeted Timers
Hello Hold                  : 45 seconds
Keep Alive Hold            : 60 seconds
LDP Link Timers
Hello Hold                  : 15 seconds
Keep Alive Hold            : 40 seconds
Label Advertisement
Direct : Matching LSR-ID Only
Rip    : None
Static : None
LDP VLANs : vlan1
          : vlan2
          : vlan3
* Indicates parameters that cannot be modified
E4G-400.1 # show mpls ldp
LDP Admin Status           : Enabled
LDP Oper Status            : Disabled
Protocol Version           : v1*
Label Retention Mode       : Liberal*
Label Distribution Method   : Downstream Unsolicited*
Label Distribution Control Mode : Ordered*
LDP BGP LSPs               : Disabled
```




```

LDP Loop Detection
Status          : Disabled
Hop-Count Limit : 255
Path-Vector Limit : 255
LDP Targeted Timers
Hello Hold      : 45 seconds
Keep Alive Hold : 60 seconds
LDP Link Timers
Hello Hold      : 15 seconds
Keep Alive Hold : 40 seconds
Label Advertisement
Direct : Matching LSR-ID Only
Rip    : None
Static : None
LDP VLANs : karen
        : lb
CES Pseudo Wire Parameter Mismatch Recovery : Auto
* Indicates parameters that cannot be modified

```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show mpls ldp interface

```
show mpls ldp interface {{vlan} vlan_name} {detail | counters}
```

Description

Displays LDP information about MPLS interfaces. Summary information is displayed in tabular format for all VLANs that are configured for MPLS.

Syntax Description

vlan	Displays LDP interface information for one VLAN.
<i>vlan_name</i>	Specifies the name of the VLAN.
detail	Displays LDP interface information including LDP control packet counts.
counters	Displays only the LDP control protocol packet counts.

Default

N/A.



Usage Guidelines

If the optional **detail** keyword is specified, the information is shown in verbose form and LDP control packet counts are displayed. If the optional **counters** keyword is specified, only the LDP control protocol packet counts are shown. The counters are described in RFC 3036, LDP Specification.

Example

The following command displays detailed LDP information for the interface associated with VLAN 1:

```
* BD-10K.22 # show mpls ldp interface vlan1 detail
VLAN Name : vlan1
Local IP Address      : 11.121.96.20
MPLS Admin Status    : Enabled
MPLS Oper Status     : Enabled
LDP Admin Status     : Enabled
LDP Oper Status      : Enabled
LDP UpTime           : 0d:1h:59m:56s
Current Adjacencies  : 1
Negotiated Hello Hold Time : 15000 ms
Time to Send Next Hello : 4060 ms
Link      Targeted
Counter                               Adjacencies  Adjacencies
-----
Shutdown Notifications (Rcvd)          0             0
Shutdown Notifications (Sent)          0             0
Failed Session Attempts (NAKs)         0             0
Hello Errors                           0             0
Parameters Advertised Errors           0             0
Max PDU Length Errors                  0             0
Label Range Errors                     0             0
Bad LDP ID Errors                      0             0
Bad PDU Length Errors                  0             0
Bad Msg Length Errors                  0             0
Bad TLV Length Errors                  0             0
Bad TLV Value Errors                   0             0
Keep-Alive Timeout Errors              0             0
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

show mpls ldp label

```
show mpls {ldp} label {lsp} {summary | label_num | [advertised | received]
{label_num} | received implicit-null}
```



Description

Displays LDP LSP label information.

Syntax Description

ldp	Specifies that only LDP labels are displayed.
summary	Specifies that summary information for the labels is displayed.
<i>label_num</i>	Specifies that only labels that match the specified hexadecimal label number are displayed.
advertised	Specifies that only labels advertised to other routers are displayed.
received	Specifies that only labels received from other routers are displayed.
implicit-null	Specifies that only implicit-null labels are displayed.

Default

N/A.

Usage Guidelines

If no options are specified, tabular information for all LDP labels, except advertised implicit-null labels, is displayed. The following table describes the display fields that appear when this command is entered.

Table 40: show mpls ldp label Field Definitions

Field	Definition
Advertised Label	Advertised Label is the label advertised to other routers.
Destination Mapping	Destination Mapping displays the destination IP address and mask.
LSP Flags	LSP Flags lists the flags for each label, which are described in the key at the bottom of the display.
Peer Label	The peer label that replaces the received label of incoming packets during a label swap on a transit LSP. The peer label appears only for transit LSPs.
NHop Type	NHop Type helps define the handling of a packet arriving with the specified advertised label. Possible values are: IP—The label on the packet is swapped and the packet forwarded to the IP address shown in the NextHop field. vlan—The label on the packet is stripped and is IP routed according to the Destination Mapping field. locl—Indicates that the tunnel destination is local to this switch.
Next Hop	Displays an interface ID for the next hop router, which can be a VLAN name, VMAN name, or IP address.
Name	Displays an LSP name.
Received Label	Received Label displays the label received from other routers.



Example

The following command displays all LDP labels except received implicit-null labels:

```
* Switch.1 # show mpls ldp label
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show mpls ldp label advertised

```
show mpls ldp label {lsp} advertised implicit-null {ipNetmask}
```

Description

Displays advertised LDP LSP implicit-null label information.

Syntax Description

<i>ipNetmask</i>	Specifies an IP network mask.
------------------	-------------------------------

Default

N/A.

Usage Guidelines

If no options are specified, tabular information for all LDP advertised implicit-null labels is displayed. The following table describes the display fields that appear when this command is entered.

Table 41: show mpls ldp label Field Definitions

Field	Definition
Advertised Label	Advertised Label is the label advertised to other routers.
Destination Mapping	Destination Mapping displays the destination IP address and mask.
LSP Flags	LSP Flags lists the flags for each label, which are described in the key at the bottom of the display.
Peer Label	The peer label that replaces the received label of incoming packets during a label swap on a transit LSP. The peer label appears only for transit LSPs.



Table 41: show mpls ldp label Field Definitions (continued)

Field	Definition
NHop Type	NHop Type helps define the handling of a packet arriving with the specified advertised label. Possible values are: IP—The label on the packet is swapped and the packet forwarded to the IP address shown in the NextHop field.vlan—The label on the packet is stripped and is IP routed according to the Destination Mapping field.locl—Indicates that the tunnel destination is local to this switch.
Next Hop	Displays an interface ID for the next hop router, which can be a VLAN name, VMAN name, or IP address.
Name	Displays a VPLS or LSP name.
Received Label	Received Label displays the label received from other routers.

Example

The following command displays all advertised LDP implicit-null labels:

```
* Switch.1 # show mpls ldp label advertised implicit-null
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

show mpls ldp label l2vpn

```
show mpls {ldp} label l2vpn {summary | label_num | [advertised | received]
{label_num}}
```

Description

Displays LDP Layer 2 VPN label information.

Syntax Description

ldp	Specifies that only LDP labels are displayed.
summary	Specifies that summary information for the labels is displayed.
<i>label_num</i>	Specifies that only labels that match the specified hexadecimal label number are displayed.



advertised	Specifies that only labels advertised to other routers are displayed.
received	Specifies that only labels received from other routers are displayed.

Default

N/A.

Usage Guidelines

If no options are specified, tabular information for all LDP Layer 2 VPN labels is displayed. The following table describes the display fields that appear when this command is entered.

Table 42: show mpls ldp label l2vpn Field Definitions

Field	Definition
Advertised Label	Advertised Label is the label advertised to other routers.
Destination Mapping	Destination Mapping displays the destination IP address and mask.
LSP Flags	LSP Flags lists the flags for each label, which are described in the key at the bottom of the display.
Peer Label	The peer label that replaces the received label of incoming packets during a label swap on a transit LSP. The peer label appears only for transit LSPs.
NHop Type	NHop Type helps define the handling of a packet arriving with the specified advertised label. Possible values are: VPLS—The label on the packet is stripped and forwarded according to the service configuration of the VPLS specified in the Name field.
Next Hop	Displays an interface ID for the next hop router, which can be a VLAN name, VMAN name, or IP address.
Name	Displays a VPLS or LSP name.
Received Label	Received Label displays the label received from other routers.

Example

The following command displays all Layer 2 VPN labels:

```
* Switch.1 # show mpls ldp label l2vpn
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



show mpls ldp label l2vpn retained

```
show mpls ldp label l2vpn retained {ipaddress}
```

Description

Displays Layer 2 VPN liberally retained labels received from a peer.

Syntax Description

<i>ipaddress</i>	Specifies an IP address.
------------------	--------------------------

Default

N/A.

Usage Guidelines

If no options are specified, tabular information for all Layer 2 VPN liberally-retained labels is displayed. The following table describes the display fields that appear when this command is entered.

Table 43: show mpls ldp label l2vpn retained Field Definitions

Field	Definition
Advertised Label	Advertised Label is the label advertised to other routers.
Destination Mapping	Destination Mapping displays the destination IP address and mask.
LSP Flags	LSP Flags lists the flags for each label, which are described in the key at the bottom of the display.
Peer Label	The peer label that replaces the received label of incoming packets during a label swap on a transit LSP. The peer label appears only for transit LSPs.
NHop Type	NHop Type helps define the handling of a packet arriving with the specified advertised label. Possible values are: VPLS—The label on the packet is stripped and forwarded according to the service configuration of the VPLS specified in the Name field.
Next Hop	Displays an interface ID for the next hop router, which can be a VLAN name, VMAN name, or IP address.
Name	Displays a VPLS name.
Received Label	Received Label displays the label received from other routers.

Example

The following command displays liberally-retained Layer 2 VPN labels received from peers:

```
* Switch.1 # show mpls ldp label l2vpn retained
```



History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show mpls ldp label lsp retained

```
show mpls ldp label lsp retained {ipNetmask}
```

Description

Displays LSP liberally retained labels received from a peer.

Syntax Description

<i>ipNetmask</i>	Specifies an IP network mask.
------------------	-------------------------------

Default

N/A.

Usage Guidelines

If no options are specified, tabular information for all LSP liberally-retained labels is displayed. The following table describes the display fields that appear when this command is entered.

Table 44: show mpls ldp label lsp retained Field Definitions

Field	Definition
Advertised Label	Advertised Label is the label advertised to other routers.
Destination Mapping	Destination Mapping displays the destination IP address and mask.
LSP Flags	LSP Flags lists the flags for each label, which are described in the key at the bottom of the display.
Peer Label	The peer label that replaces the received label of incoming packets during a label swap on a transit LSP. The peer label appears only for transit LSPs.
NHop Type	NHop Type helps define the handling of a packet arriving with the specified advertised label. Possible values are: IP—The label on the packet is swapped and the packet forwarded to the IP address shown in the NextHop field. vlan—The label on the packet is stripped and is IP routed according to the Destination Mapping field. locl—Indicates that the tunnel destination is local to this switch.
Next Hop	Displays an interface ID for the next hop router, which can be a VLAN name, VMAN name, or IP address.



Table 44: show mpls ldp label lsp retained Field Definitions (continued)

Field	Definition
Name	Displays an LSP name.
Received Label	Received Label displays the label received from other routers.

Example

The following command displays liberally-retained LSP labels received from peers:

```
* Switch.1 # show mpls ldp label lsp retained lsp
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show mpls ldp label retained

```
show mpls ldp label retained [l2vpn {ipaddress} | lsp {ipNetmask}]
```

Description

Displays liberally-retained labels received from a peer for either the Layer 2 VPN protocol or LSP protocol.

Syntax Description

<i>ipaddress</i>	Specifies an IP address.
<i>ipNetmask</i>	Specifies an IP network mask.

Default

N/A.

Usage Guidelines

If no options are specified, tabular information for all liberally-retained labels is displayed. The following table describes the display fields that appear when this command is entered.



Table 45: show mpls label retained Field Definitions

Field	Definition
Advertised Label	Advertised Label is the label advertised to other routers.
Destination Mapping	Destination Mapping displays the destination IP address and mask.
LSP Flags	LSP Flags lists the flags for each label, which are described in the key at the bottom of the display.
Peer Label	The peer label that replaces the received label of incoming packets during a label swap on a transit LSP. The peer label appears only for transit LSPs.
NHop Type	NHop Type helps define the handling of a packet arriving with the specified advertised label. Possible values are: IP—The label on the packet is swapped and the packet forwarded to the IP address shown in the NextHop field. local—Indicates that the tunnel destination is local to this switch. VLAN—The label on the packet is stripped and is IP routed according to the Destination Mapping field. VPLS—The label on the packet is stripped and forwarded according to the service configuration of the VPLS specified in the Name field.
Next Hop	Displays an interface ID for the next hop router, which can be a VLAN name, VMAN name, or IP address.
Name	Displays a VPLS or LSP name.
Received Label	Received Label displays the label received from other routers.

Example

The following command displays liberally-retained LSP labels received from peers:

```
* Switch.1 # show mpls ldp label retained lsp
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show mpls ldp lsp

```
show mpls ldp lsp {prefix ipNetmask} {ingress | egress | transit} {detail}
```

Description

Displays the LSP information associated with LDP that is used to forward packets within the MPLS network. If no options are specified, summary information for all LSPs is displayed.



Syntax Description

prefix	Displays information for a single FEC that matches the prefix.
<i>ipNetmask</i>	Designates the FEC for which to display information.
ingress	Displays information for LSPs that originate from the switch into the MPLS network.
egress	Displays information for LSPs that terminate at the switch from the MPLS network.
transit	Displays information for LSPs that traverse the switch.
detail	Display detailed LSP information.

Default

N/A.

Usage Guidelines

If no options are specified, this command displays summary information for all LSPs.

Optionally, the LSPs displayed can be further qualified by the keywords **ingress**, **egress**, and **transit**. These keywords qualify the LSPs displayed from the perspective of the switch. Ingress LSPs originate from the switch into the MPLS network. Egress LSPs terminate at the switch from the MPLS network. Transit LSPs traverse the switch. If the optional **prefix** keyword is specified, only the LSP information associated with the FEC that matches the prefix is displayed.

If the **detail** keyword is specified, information is displayed in verbose form and includes received packet and byte counts.

Example

The following command displays LDP information for an ingress LSP:

```
* BD-10K.5 # show mpls ldp lsp 11.100.100.59/32 ingress detail
FEC IP/Prefix: 11.100.100.59/32
Next Hop I/F      : m5vlan1
Next Hop Addr    : 12.224.0.55
Advertised Label : n/a                Received Label : 0x80403 (525315)
Rx Packets       : n/a                Tx Packets    : 61
Rx Bytes        : n/a                Tx Bytes     : 4294967296
```

The following command displays LDP information for a transit LSP:

```
* BD-10K.5 # show mpls ldp lsp 11.100.100.55/32 transit detail
FEC IP/Prefix: 11.100.100.55/32
Next Hop I/F      : m5vlan1
Next Hop Addr    : 12.224.0.55
Advertised Label : 0x11 (17)         Received Label : 0x80403 (525315)
Rx Packets       : 61                Tx Packets    : 61
Rx Bytes        : 4294967296        Tx Bytes     : 4294967296
```



The following command displays LDP information for an egress LSP:

```
* BD-10K.5 # show mpls ldp lsp 11.100.100.30/32 egress detail
FEC IP/Prefix: 11.100.100.30/32
Direct VLAN      : loop
Advertised Label : 0x80400 (525312)   Received Label : n/a
Rx Packets       : 61                  Tx Packets      : n/a
Rx Bytes         : 4294967296         Tx Bytes       : n/a
```

History

This command was first available in ExtremeXOS 11.6.

The output for this command was modified in ExtremeXOS 12.2.2.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show mpls ldp peer

```
show mpls ldp peer {ipaddress} {detail}
```

Description

Displays information about the status of the LDP sessions and hello adjacencies for all LDP peers.

Syntax Description

<i>ipaddress</i>	Display session and hello adjacency information for a single LDP peer.
detail	Display additional detailed information related to the session and adjacencies.

Default

N/A.

Usage Guidelines

Specifying the LDP peer's *ipaddress* displays session and hello adjacency information for a single LDP peer. When the **detail** keyword is specified, additional detailed information related to the session and adjacencies is displayed.

Table 38: [show mpls label Field Definitions](#) on page 2259 describes the display fields that appear when this command is entered.



Table 46: show mpls ldp peer Field Definitions

Field	Definition
IP Address	Local IP address, which is used as the LSR-ID.
LDP Peer	LDP identifier of LDP peer.
State	Displays the state of the Initialization State Machine as described in RFC 3036, LDP Specification. The states are: NonExistent, Initialized, OpenRec, OpenSent, and Operational.
Uptime	Displays the total time the session has been operational.
Adjacencies	Displays the number of active adjacencies with the LDP peer.
Index	The Entity Index used in the LDP Entity Table MIB.
Targeted Peer	The IP address of the peer used in Extended Discovery. If this is not a targeted peer (Basic Discovery was used), this field displays Not Targeted.
Attempted Sessions	This and other counters are described in RFC 3036, LDP Specification.
Shutdown Notifications	This and other counters are described in RFC 3036, LDP Specification.
Peer	IP address of peer.
Peer Label Space	The label space as given by the peer and derived from the LDP Identifier. A zero value represents the global or per-platform label space. A non-zero value represents a per interface label space.
Session State	Displays the state of the Initialization State Machine as described in RFC 3036, LDP Specification. The states are: NonExistent, Initialized, OpenRec, OpenSent, and Operational.
Session Uptime	Displays the total time the session has been operational.
Discontinuity Time	The system uptime for the most recent period after which one or more of the session's counters suffered a discontinuity.
Keep Alive Hold Timer	Displays the configured keep alive hold timer value and the remaining hold time.
Label Distribution Method	Label Distribution Method is either DU (Downstream Unsolicited) or DOD (Downstream On Demand), as described in RFC 3036, LDP Specification.
Max PDU Length	The maximum allowable length for LDP PDUs (Protocol Data Units) for this session.
Unknown Msg Type Errors	The number of messages received for this session without a recognized message type and without the ignore bit set.
Unknown TLV Errors	The number of TLVs received for this session without a recognized TLV type and without the ignore bit set.
Next Hop Addr(s)	A list of next hop addresses received from the peer through LDP address messages.

Example

The following command displays MPLS LDP session information for the LDP entity 11.100.100.30:

```
Mariner3.59 # show mpls ldp peer
IP Address      LDP Peer      State      Uptime      Adjacencies
11.100.100.30  11.100.100.55:0  Operational  0d:14h:51m:53s  1
11.100.100.30  14.4.0.99:15   Operational  0d:1h:0m:43s   1
11.100.100.30  14.4.0.99:16   Operational  0d:0h:34m:51s  1
```



```

Adjacencies:
Index                : 1                Attempted
Sessions             : 1
Targeted Peer       : 11.100.100.210:0   Shutdown
Notifications      : Sent 0 Rcvd 0
Mariner3.32 # show mpls ldp peer detail
Peer: 11.100.100.55   Peer label space: 0 (global)
Session State       : Operational
Session Uptime      : 0d:0h:13m:41s
...
Peer: 14.4.0.99      Peer label space: 15
Session State       : Operational
Session Uptime      : 0d:0h:57m:4s
...
Peer: 14.4.0.99      Peer label space: 16
Session State       : Operational
Session Uptime      : 0d:0h:31m:12s
...
* DUT65.2 # show mpls ldp peer detail
Peer: 11.100.100.210 Peer Label Space: 0 (global)
Session State       : Operational
Session Uptime      : 0d:0h:6m:30s
Discontinuity Time  : 34677
Keep Alive Hold Timer : 40 (remaining: 37.86)
Label Distribution Method : DU
Max PDU Length      : 4096
Unknown Msg Type Errors : 0
Unknown TLV Errors   : 0
Next Hop Addr(s)    : 11.100.100.210 12.20.20.210
Adjacencies:
Index                : 1                Attempted
Sessions             : 0
Targeted Peer       : 11.100.100.210:0   Shutdown
Notifications      : Sent 0 Rcvd 0
Admin Status        : Enabled           No Hello
Errors              : 0
Operational Status  : Up               Advertisement
Errors              : 0
Label Retention Mode : Liberal           Max PDU
Errors              : 0
Hop Count Limit     : Disabled         Bad LDP Identifier
Errors : 0
Path Vector Limit   : Disabled         Bad PDU Length
Errors              : 0
Hello Hold Timer    : 45 (remaining: 38) Bad TLV Length
Errors              : 0
Malformed TLV Errors : 0
Bad Message Length Errors : 0
Session Rejected Errors : 0
Keep Alive Expired Errors : 0
Index              : 6                Attempted
Sessions           : 0
Targeted Peer      : Not Targeted     Shutdown
Notifications     : Sent 0 Rcvd 0
Admin Status       : Enabled           No Hello
Errors             : 0
Operational Status : Up               Advertisement
Errors             : 0

```



```

Label Retention Mode      : Liberal                               Max PDU
Errors                    : 0
Hop Count Limit          : Disabled                             Bad LDP Identifier
Errors                   : 0
Path Vector Limit        : Disabled                             Bad PDU Length
Errors                   : 0
Hello Hold Timer         : 15 (remaining: 10)                 Bad TLV Length
Errors                   : 0
VLAN                     : v1                                  Malformed TLV
Errors                   : 0
Interface address        : 12.20.20.182                       Bad Message Length
Errors                   : 0
Session Rejected Errors  : 0
Keep Alive Expired Errors : 0

```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show mpls rsvp-te

```
show mpls rsvp-te
```

Description

Displays displays summary configuration and status information for RSVP-TE.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays summary configuration and status information for RSVP-TE. The parameters displayed include:

- Global status of RSVP-TE.
- Configured standard LSP timer values.
- Configured rapid-retry LSP timer values.
- RSVP-TE VLANs.



Example

The following command shows the summary configuration and status information for RSVP-TE:

```
* BD-10K.16 # show mpls rsvp-te
RSVP-TE Admin Status : Enabled
RSVP-TE Oper Status  : Enabled
LSP Standard-Retry Timers
Delay-Interval      : 30 seconds
Decay-Rate          : 50 %
Retry-Limit         : unlimited
LSP Rapid-Retry Timers
Delay-Interval      : 500 milliseconds
Decay-Rate          : 50 %
Retry-Limit         : 10
RSVP-TE VLANs      : vlan1
                   : vlan2
                   : vlan3
```

History

This command was first available in ExtremeXOS 11.6

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show mpls rsvp-te bandwidth

```
show mpls rsvp-te bandwidth {{vlan} vlan_name} {detail}
```

Description

Displays the reserved bandwidth for each TE LSP by interface.

Syntax Description

<i>vlan_name</i>	Displays the reserved bandwidth for each TE LSP associated with the specified VLAN.
detail	Displays the path information in verbose format.

Default

All TE LSPs for all RSVP-TE-enabled interfaces are shown.



Usage Guidelines

This command displays the reserved bandwidth for each TE LSP by interface. By default, all TE LSPs for all RSVP-TE enabled interfaces are shown.

Note



Beginning with ExtremeXOS Release 12.2.1, the receive bandwidth can only be used for tracking. If the configured receive bandwidth is exceeded, the available bandwidth shown might be negative. In this case, “Ovr” is displayed to indicate that the link is oversubscribed in the receive direction. The **detail** option can be used to show the actual LSPs using this bandwidth.

The optional **vlan** keyword limits the display to only those LSPs that have bandwidth reservations against the specified VLAN. Only committed-rate bandwidth is displayed. Bandwidth is displayed as either received or transmitted bandwidth with respect to the switch.

LSPs are listed using the configured or signaled LSP name. If the LSP name was not included in the setup control messages (which can only occur when using OEM vendor equipment), the LSP is uniquely identified using a concatenated string that includes the tunnel ID and source IP address. Per VLAN, each LSP is listed in descending priority order. That is, the LSPs listed at the top of each VLAN have the highest bandwidth priority and are less likely to be preempted. Bandwidth priority is determined by the signaled hold-priority and the uptime. The TE LSP with a hold-priority of zero and the highest uptime has the highest bandwidth priority and the TE LSP with a hold-priority of seven and the lowest uptime has the lowest bandwidth priority.

Use the **detail** keyword to display detailed information.

Example

The following command displays bandwidth reservation information for the specified VLAN:

```
show mpls rsvp-te bandwidth vlan vlan_1 detail
* BD-10K.2 # show mpls rsvp-te bandwidth vlan_1 detail
Vlan          Dir  Pool      CIR (per priority level)
LSP           0    1    2    3    4    5    6    7
-----
vlan_1              Rx  300
vlalsp1            -   -   10   -   -   -   -   -
-----
Available         300  300  290  290  290  290  290  290
Tx  500
vlalsp2            -   -    9   -   -   -   -   -
-----
Available         500  500  491  491  491  491  491  491
(Rx)Receive Bandwidth  (Tx)Transmit Bandwidth
```

History

This command was first available in ExtremeXOS 11.6.



Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show mpls rsvp-te interface

```
show mpls rsvp-te interface {{vlan} vlan_name} {detail | counters}
```

Description

Displays RSVP-TE information about MPLS interfaces.

Syntax Description

vlan	Display information for one VLAN.
<i>vlan_name</i>	Specifies the name of the VLAN.
detail	Display RSVP-TE information including the interface up time and LDP control packet counts.
counters	Display only the RSVP-TE control protocol packet counts.

Default

N/A.

Usage Guidelines

This command displays RSVP-TE information about MPLS interfaces. Summary information is displayed in tabular format for all VLANs that are configured for MPLS. The following information is displayed:

- VLAN name.
- Bandwidth reserved.
- TE metric.
- Hello interval time.
- Refresh interval time.
- Summary refresh time.
- Bundle message time.
- Uptime.
- Number of neighbors.
- RSVP-TE state information

When the optional **detail** keyword is specified, additional RSVP-TE information is displayed. This additional information includes:

- RSVP-TE hello keep multiplier.
- RSVP-TE refresh keep multiplier.



- RSVP-TE available bandwidth per priority level.
- RSVP-TE control protocol packet counts.

When the optional **counters** keyword is specified, only the RSVP-TE control protocol packet counts are shown.

Example

The following command displays detailed RSVP-TE information for the interfaces associated with VLAN 1:

```
* BD-10K.15 # show mpls rsvp-te interface vlan1 det
VLAN Name : vlan1
Local IP Address      : 11.121.96.20
MPLS Admin Status    : Enabled
MPLS Oper Status     : Enabled
RSVP-TE Admin Status : Enabled
RSVP-TE Oper Status  : Enabled
RSVP-TE Up-Time      : 0d:1h:19m:46s
# Neighbors          : 1
Receive CIR          : 50000 Kbps
Transmit CIR         : 50000 Kbps
TE Metric            : Use IGP Cost/Metric
Hello Interval       : 3 seconds
Refresh Time         : 30 seconds
Hello Keep Multiplier : 3
Refresh Keep Multiplier : 3
Summary Refresh     : Disabled
Summary Refresh Time : 3000 milliseconds
Bundle Message       : Disabled
Bundle Message Time  : 1000 milliseconds
-----
```

Dir	Pool	CIR Available (per priority level)							
		0	1	2	3	4	5	6	7
Rx	50	50	50	50	50	50	50	50	50
Tx	50	50	50	50	50	50	50	50	40

```
(Rx)Receive Bandwidth (Tx)Transmit Bandwidth
Message                Sent      Received
-----
```

PATH	165	165
PATH_TEAR	2	3
PATH_ERR	4	18
RESV	160	145
RESV_TEAR	0	2
RESV_ERR	0	0
RESV_CONFIRM	4	1
SUMMARY_REFRESH	0	0
BUNDLE	0	0
HELLO	42	30

History

This command was first available in ExtremeXOS 11.6.



Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show mpls rsvp-te lsp

```
show mpls rsvp-te lsp {[destination | origin] ipaddress} {fast-reroute} {detail}
| summary}
```

Description

Displays complete or filtered information for all RSVP-TE LSPs.

Syntax Description

destination	Displays only those LSPs that terminate at the specified IP address.
origin	Displays only those LSPs that originate at the specified IP address.
<i>ipaddress</i>	Specifies an IP address for an LSP origin or destination.
fast-reroute	Limits the display to only those LSPs with fast reroute protection.
detail	Displays the LSP information in verbose format.
summary	Displays only the LSP summary statistics section of the normal display.

Default

N/A.

Usage Guidelines

If no options are specified, information for all RSVP-TE LSPs is displayed.

You can limit the display to ingress, transit, or egress LSPs with the following commands:

```
show mpls rsvp-te lsp [ingress {fast-reroute} | ingress_lsp_name |
ingress ingress_lsp_name | ingress [destination | origin]ipaddress] {[all-paths |
detail] | summary | down-paths {detail}} show mpls rsvp-te lsp [egress | transit]
{fast-reroute} {[lsp_name} {[destination | origin] ipaddress} {detail} | summary}
```

When label recording is enabled for an LSP, labels are displayed only for the ingress node (the egress label from the previous node always matches the ingress node label).

Example

The following command example displays information about all RSVP-TE LSPs:

```
# show mpls rsvp-te lsp
Ingress LSP Name Path Name          Destination      Next Hop I/F      UpTm  Flags
```



```

-----
-----
frrlspltom2      any          11.100.100.50  m2vlan1      5m UEF---
IV
lsptom2         pathtom2    11.100.100.50  m2vlan1      5m UES--
OIV
tom5            any          11.100.100.55  m5vlan1      5m UEP---
IV
Egress LSP Name Source IP      Destination    Prev Hop I/F   UpTm
-----
tom3            11.100.100.55 11.100.100.30  m5vlan1        5m
frrlspltom3    11.100.100.50 11.100.100.30  m2vlan1        5m
tom3            11.100.100.50 11.100.100.30  m2vlan1        5m
Transit LSP Name Source IP      Destination    Prev Hop I/F   Next Hop I/F
UpTm
-----
-----
tom2            11.100.100.55 11.100.100.50  m5vlan1
m2vlan1        5m
frrlspltom5    11.100.100.50 11.100.100.55  m2vlan1
m5vlan1        5m
lsptom5        11.100.100.50 11.100.100.55  m2vlan1
m5vlan1        5m
Flags: (U) Up, (E) Enabled, (P) Primary LSP, (S) Secondary LSP,
(F) FRR Primary LSP (R) Redundant Paths, (B) Bandwidth Requested,
(O) ERO Specified, (I) IP Traffic Allowed, (V) VPN Traffic Allowed,
(v) VPN Assigned Traffic Allowed
Summary of RSVP-TE LSPs
Ingress LSPs (Enabled/Disabled)      : 3 (3/0)
Ingress LSPs with no configured path : 0
Ingress LSP Paths (Up/Down)         : 3 (3/0)
Detour LSP Paths (Up/Down)          : 0 (0/0)
Transit LSPs                         : 3
Egress LSPs                          : 3

```

The next command example displays only the summary information for all RSVP-TE LSPs:

```

# show mpls rsvp-te lsp summary
Summary of RSVP-TE LSPs
Ingress LSPs (Enabled/Disabled)      : 3 (3/0)
Ingress LSPs with no configured path : 0
Ingress LSP Paths (Up/Down)         : 3 (3/0)
Detour LSP Paths (Up/Down)          : 0 (0/0)
Transit LSPs                         : 3
Egress LSPs                          : 3

```

The following command example limits the display to RSVP-TE LSPs with fast reroute protection:

```

# show mpls rsvp-te lsp fast-reroute
Ingress LSP Name Path Name      Destination    Next Hop I/F   UpTm Flags
-----
frrlspltom2      any          11.100.100.50  m2vlan1      5m UEF---
IV
Egress LSP Name Source IP      Destination    Prev Hop I/F   UpTm

```



```

-----
tom3          11.100.100.55  11.100.100.30  m5vlan1      5m
frrlspltom3  11.100.100.50  11.100.100.30  m2vlan1      5m
tom3          11.100.100.50  11.100.100.30  m2vlan1      5m
Transit LSP Name Source IP      Destination    Prev Hop I/F  Next Hop I/F
UpTm
-----
-----
tom2          11.100.100.55  11.100.100.50  m5vlan1
m2vlan1      5m
frrlspltom5  11.100.100.50  11.100.100.55  m2vlan1
m5vlan1      5m
lsptom5      11.100.100.50  11.100.100.55  m2vlan1
m5vlan1      5m
Flags: (U) Up, (E) Enabled, (P) Primary LSP, (S) Secondary LSP,
(F) FRR Primary LSP (R) Redundant Paths, (B) Bandwidth Requested,
(O) ERO Specified, (I) IP Traffic Allowed, (V) VPN Traffic Allowed,
(v) VPN Assigned Traffic Allowed
Summary of RSVP-TE LSPs
Ingress FRR LSPs (Enabled/Disabled)          : 1 (1/0)
Ingress FRR LSPs with no configured path    : 0
Ingress FRR LSP Paths (Up/Down)             : 1 (1/0)
Detour LSP Paths (Up/Down)                   : 0 (0/0)
Transit FRR LSPs                             : 3
Egress FRR LSPs                              : 3

```

The following command example displays detailed information for RSVP-TE LSPs that originate at IP address 11.100.100.50:

```

# show mpls rsvp-te lsp origin 11.100.100.50 detail
Egress LSP Name: frrlspltom3
Tunnel ID       : 1                Ext Tunnel ID  : 11.100.100.50
LSP ID          : 0                UpTime        : 0d:0h:5m:25s
Source IP       : 11.100.100.50    Destination IP : 11.100.100.30
Previous Hop I/F : 12.220.0.30 - m2vlan1
Advertised Label : 0x80402         Received Label : n/a
Rx Packets      : 0                Tx Packets     : n/a
Rx Bytes       : 0                Tx Bytes      : n/a
Record Route    : Indx  IP Address
: 1 12.220.0.50
Detour LSP:
Bandwidth Protection : Enabled
Egress LSP Name: tom3
Tunnel ID       : 11               Ext Tunnel ID  : 11.100.100.50
LSP ID          : 0                UpTime        : 0d:0h:5m:25s
Source IP       : 11.100.100.50    Destination IP : 11.100.100.30
Previous Hop I/F : 12.220.0.30 - m2vlan1
Advertised Label : 0x80404         Received Label : n/a
Rx Packets      : 0                Tx Packets     : n/a
Rx Bytes       : 0                Tx Bytes      : n/a
Record Route    : Indx  IP Address
: 1 12.220.0.50
Transit LSP Name: frrlspltom5
Tunnel ID       : 2                Ext Tunnel ID  : 11.100.100.50
LSP ID          : 0                UpTime        : 0d:0h:5m:15s
Source IP       : 11.100.100.50    Destination IP : 11.100.100.55
Previous Hop I/F : 12.220.0.30 - m2vlan1

```



```

Next Hop I/F      : 12.224.0.30 - m5vlan1
NextHop Addr     : 12.224.0.55
Advertised Label : 0x70004                Received Label : 0x804c9
Rx Packets       : 0                      Tx Packets      : n/a
Rx Bytes         : 0                      Tx Bytes       : n/a
Record Route     : Empty
Detour LSP:
Bandwidth Protection : Enabled      Node/Link Protection : Enabled(Node)
LSP Available       : False         LSP in use        : False
Transit LSP Name: lsptom5
Tunnel ID          : 5              Ext Tunnel ID    : 11.100.100.50
LSP ID             : 0              UpTime           : 0d:0h:5m:25s
Source IP          : 11.100.100.50  Destination IP   : 11.100.100.55
Previous Hop I/F  : 12.220.0.30 - m2vlan1
Next Hop I/F      : 12.224.0.30 - m5vlan1
NextHop Addr     : 12.224.0.55
Advertised Label : 0x70005                Received Label : 0x80598
Rx Packets       : 0                      Tx Packets      : n/a
Rx Bytes         : 0                      Tx Bytes       : n/a
Record Route     : Empty

```

History

This command was first available in ExtremeXOS 11.6.

This command and its output were modified, and the summary option was added in ExtremeXOS 12.0.

The fast-reroute feature was first available in ExtremeXOS 12.1.

This command and its output were modified in ExtremeXOS 12.2.2.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

show mpls rsvp-te lsp [egress | transit]

```

show mpls rsvp-te lsp [egress | transit] {fast-reroute} {{lsp_name} {[destination
| origin] ipaddress} {detail} | summary}

```

Description

Displays complete or filtered information for one or all egress or transit RSVP-TE LSPs.

Syntax Description

egress	Limits the display to only the LSPs that terminate at this switch.
transit	Limits the display to only the LSPs that transit this switch.
fast-reroute	Limits the display to only those LSPs with fast reroute protection.



<i>lsp_name</i>	When either the transit or the egress option is specified, this variable specifies a name for a single LSP for which information is displayed.
destination	Displays only those LSPs that terminate at the specified IP address.
origin	Displays only those LSPs that originate at the specified IP address.
<i>ipaddress</i>	Specifies an IP address for an LSP origin or destination.
detail	Displays the LSP information in verbose format.
summary	Displays only the LSP summary statistics section of the normal display.

Default

N/A.

Usage Guidelines

You can limit the display to ingress LSPs with the following command:

```
show mpls rsvp-te lsp [ingress {fast-reroute} | ingress_lsp_name |
ingress_lsp_name | ingress [destination | origin]ipaddress] {[all-paths |
detail] | summary | down-paths {detail}}
```

You can display information for all LSPs with the following command:

```
show mpls rsvp-te lsp {[destination | origin] ipaddress} {fast-reroute} {detail}
| summary}
```

When label recording is enabled for an LSP, labels are displayed only for the ingress node (the egress label from the previous node always matches the ingress node label).

Example

The following command example displays RSVP-TE LSPs that terminate at this switch and at IP address 11.100.100.30:

```
# show mpls rsvp-te lsp egress destination 11.100.100.30
Egress LSP Name  Source IP      Destination    Prev Hop I/F    UpTm
-----
tom3             11.100.100.55  11.100.100.30  m5vlan1         5m
frrlspltom3     11.100.100.50  11.100.100.30  m2vlan1         5m
tom3             11.100.100.50  11.100.100.30  m2vlan1         5m
Summary of Egress RSVP-TE LSPs to destination 11.100.100.30
Egress LSPs                : 3
Egress Protected LSPs     : 0
```

History

This command was first available in ExtremeXOS 11.6.

This command and its output were modified, and the summary option was added in ExtremeXOS 12.0.



The fast-reroute feature was first available in ExtremeXOS 12.1.

This command and its output were modified in ExtremeXOS 12.2.2.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show mpls rsvp-te lsp ingress

```
show mpls rsvp-te lsp [ingress {fast-reroute} | ingress_lsp_name | ingress
ingress_lsp_name | ingress [destination | origin] ipaddress] [{all-paths |
detail} | summary | down-paths {detail}]
```

Description

Displays information for the specified ingress RSVP-TE LSP.

Syntax Description

fast-reroute	Limits the display to only those LSPs with fast reroute protection.
<i>ingress_lsp_name</i>	Identifies the ingress LSP for which you want to display information.
destination	Displays only those LSPs that terminate at the specified IP address.
origin	Displays only those LSPs that originate at the specified IP address.
<i>ipaddress</i>	Specifies an IP address for an LSP origin or destination.
all-paths	Specifies that the display include all redundant paths.
detail	Displays the LSP information in verbose format.
summary	Displays only the LSP summary statistics section of the normal display.
down-paths	Specifies that the display include only those paths that are operationally down.

Default

N/A.

Usage Guidelines

You can limit the display to egress or transit LSPs with the following command:

```
show mpls rsvp-te lsp [egress | transit] {fast-reroute} {{lsp_name} {[destination |
origin]ipaddress} {detail} | summary}
```

You can display information for all LSPs with the following command:

```
show mpls rsvp-te lsp {[destination | origin] ipaddress} {fast-reroute} {detail} | summary}
```



When label recording is enabled for an LSP, labels are displayed only for the ingress node (the egress label from the previous node always matches the ingress node label).

Example

Use the following command to display information about a specific LSP:

```
# show mpls rsvp-te lsp jefflsp1
Ingress LSP Name Path Name          Destination      Next Hop I/F      UpTm  Flags
-----
jefflsp1          jeffpath1       11.100.100.204  n/a              0 --PR--
IV
Flags: (U) Up, (E) Enabled, (P) Primary LSP, (S) Secondary LSP,
(F) FRR Primary LSP (R) Redundant Paths, (B) Bandwidth Requested,
(O) ERO Specified, (I) IP Traffic Allowed, (V) VPN Traffic Allowed,
(v) VPN Assigned Traffic Allowed
Summary of Ingress RSVP-TE LSPs named jefflsp1
Ingress LSPs (Enabled/Disabled)      : 1 (0/1)
Ingress LSPs with no configured path : 0
Ingress LSP Paths (Up/Down)         : 2 (0/2)
```

Use the following command to display detailed information about a specific ingress LSP:

```
* BD-10K.7 # show mpls rsvp-te lsp ingress "lsp598" detail
Ingress LSP Name: lsp598
Destination      : 11.100.100.8          Admin Status : Enabled
IP Traffic       : Allow              #VPLS Cfgd   : 0
VPN Traffic      : Allow              #VPLS In-Use : 0
Path Name: path598
Profile Name : prof598
Tunnel ID    : 1                  Ext Tunnel ID : 11.100.100.20
LSP ID       : 0                  State Changes : 5
Oper Status  : Enabled            Bandwidth Cfgd : False
LSP Type     : Primary
Activity     : Active
Failures     : 2                  Retries-since last failure : 0
Retries-Total : 12
Rcv Label    : 0x0052e           UpTime       : 0d:0h:3m:44s
Next Hop     : 11.121.96.5
Tx I/F       : 11.121.96.20 - vlan1
Record Route : Indx  IP Address
: 1 11.121.96.5
: 2 11.95.96.9
: 3 11.98.96.8
```

Use the following command to display detailed information about all paths for an LSP:

```
* BD-10K.5 # show mpls rsvp-te lsp jefflsp1 all-paths
Ingress LSP Name Path Name          Destination      Transmit I/F      UpTm  Flags
-----
jefflsp1          jeffpath0       11.100.100.204  n/a              0 --SR--
```



```

IV
jefflsp1          jeffpath1          11.100.100.204  n/a          0 --PR--
IV
Flags: (U) Up, (E) Enabled, (P) Primary LSP, (S) Secondary LSP,
(F) FRR Primary LSP (R) Redundant Paths, (B) Bandwidth Requested,
(O) ERO Specified, (I) IP Traffic Allowed, (V) VPN Traffic Allowed,
(v) VPN Assigned Traffic Allowed
Summary of Ingress RSVP-TE LSPs named jefflsp1
Ingress LSPs (Enabled/Disabled)          : 1 (0/1)
Ingress LSPs with no configured path    : 0
Ingress LSP Paths (Up/Down)              : 2 (0/2)

```

Use the following command to display information about all ingress down paths:

```

* BD-10K.5 # show mpls rsvp-te lsp ingress down-paths
Ingress LSP Name Path Name          Destination      Transmit I/F      UpTm  Flags
-----
jefflsp1          jeffpath0          11.100.100.204  n/a          0 --SR--
IV
jefflsp1          jeffpath1          11.100.100.204  n/a          0 --PR--
IV
jefflsp2          jeffpath2          11.100.100.203  n/a          0 -EPR-
OIV
Flags: (U) Up, (E) Enabled, (P) Primary LSP, (S) Secondary LSP,
(F) FRR Primary LSP (R) Redundant Paths, (B) Bandwidth Requested,
(O) ERO Specified, (I) IP Traffic Allowed, (V) VPN Traffic Allowed,
(v) VPN Assigned Traffic Allowed
Summary of Ingress RSVP-TE LSPs
Ingress LSP Paths that are Down          : 3

```

History

This command was first available in ExtremeXOS 11.6.

This command and its output were modified, and the summary option was added in ExtremeXOS 12.0.

The command output was modified in ExtremeXOS 12.2.2.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show mpls rsvp-te neighbor

```
show mpls rsvp-te neighbor {{vlan} vlan_name | ipaddress} {detail}
```

Description

Displays all recognized RSVP-TE neighbors.



Syntax Description

<i>vlan_name</i>	Displays only the neighbors for the specified VLAN.
<i>ipaddress</i>	Displays only the neighbor with the specified ipaddress.

Default

N/A.

Usage Guidelines

This command displays all recognized RSVP-TE neighbors. The IP address of each neighbor is displayed along with the VLAN name for the MPLS interface. For each neighbor, the following information is displayed:

- Number of RSVP-TE LSPs.
- Number of hello periods that have elapsed without receiving a valid hello.
- Remaining time before next hello is sent.
- Remaining time before next bundle message is sent.
- Neighbor up time.
- Neighbor supports RSVP hello.
- RSVP hello state.
- Neighbor supports refresh reduction.

If *vlan_name* is specified, only neighbors for the matching VLAN are shown. If *ipaddress* is specified, only the neighbor with that IP address is shown. If the **detail** keyword is specified, the information is shown in a verbose manner.

Example

The following command displays all recognized RSVP-TE neighbors:

```
* BD-10K.5 # show mpls rsvp-te neighbor
NeighborIP      VLAN Name      #LSPs #Miss  NxtHello  NxtBundl  Flag  UpTm
-----
11.121.96.5     vlan1          2      0      2870      0 UTH-  15m
11.122.96.8     vlan2          2      0      1770      0 UTH-  15m
Flags: (U) Hello Session Up, (T) Two Way Hello, (H) Neighbor Supports Hello,
(R) Neighbor Supports Refresh Reduction
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



show mpls rsvp-te path

```
show mpls rsvp-te path {path_name} {detail}
```

Description

Displays the configuration and usage information for MPLS RSVP-TE routed paths.

Syntax Description

<i>path_name</i>	Displays configuration and usage information for the specified MPLS RSVP-TE path.
detail	Displays the path information in verbose format.

Default

N/A.

Usage Guidelines

This command displays the configuration and usage information for MPLS RSVP-TE paths. Information is listed in tabular format and includes:

- Path name.
- Number of configured ERO objects.
- Number of LSPs configured to use this path.
- List of EROs and their type.

Specifying the optional **detail** keyword displays the path information in verbose format. If **detail** is specified, all LSPs that are configured to use the path are also displayed.

Example

The following command displays configuration and status information for the specified MPLS RSVP-TE paths:

```
* BD-10K.2 # show mpls rsvp-te path path598
Path Name          #LSP #ERO Ord#  ERO IP Netmask      Type
-----
path598            1    3   100 11.100.100.5/32     loose
200 11.100.100.9/32      loose
300 11.100.100.8/32      loose
```

History

This command was first available in ExtremeXOS 11.6.



Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show mpls rsvp-te profile

```
show mpls rsvp-te profile {profile_name} {detail}
```

Description

Displays the configuration for the specified profile.

Syntax Description

<i>profile_name</i>	Displays configuration and usage information for the specified profile, which can be either a standard or a fast-reroute profile.
detail	Displays the profile information in verbose format, and displays all LSPs that are configured to use the specified profile.

Default

N/A.

Usage Guidelines

If the *profile_name* argument is omitted, the profile parameter values for all profiles are displayed.

Example

The following command displays configuration information for all defined profiles:

```
* BD-10K.13 # show mpls rsvp-te profile
Profile Name      Peak      Committed  Max Burst  SPri  HPri  RRO  MTU  #LSP
-----
default          0         0          0         7    0    Off  i/f    0
prof598          0         0          0         7    0    On   1500  1
FRR Profile Name Mode   Bandwidth SPri  HPri  HopLmt  P-BW  P-Node #LSP
-----
Prfl_frr         Detour    0         7     0         3  Ena   Ena     2
```

The following command displays configuration information for a specific fast-reroute profile:

```
* BD-10K.14 # show mpls rsvp-te profile prfl_frr detail
Profile Name : prfl_frr
Profile type  : Fast Reroute / Standard
Peak Rate    : 0 Kbps
Committed Rate : 0 Kbps
```



```

Max Burst Size      : 0 Kb
Setup Priority      : 7
Hold Priority       : 4
Hop Limit          : 1
Protected BW       : Disabled
Protect Node       : Enabled
#LSP References    : 2
LSP / Path         : pc10_lsp / p1

```

History

This command was first available in ExtremeXOS 11.6.

The fast-reroute feature was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

show mpls rsvp-te profile fast-reroute

```
show mpls rsvp-te profile fast-reroute {detail}
```

Description

Displays the configuration for all fast-reroute profiles.

Syntax Description

detail	Displays the profile information in verbose format.
---------------	---

Default

N/A.

Usage Guidelines

None.

Example

The following command displays summary configuration information for all fast-reroute profiles:

```

* BD-10K.13 # show mpls rsvp-te profile fast-reroute
FRR Profile Name Mode   Bandwidth SPri HPri HopLmt P-BW P-Node #LSP
-----

```



default_frr	Detour	10	7	0	3 Dis	Ena	4
prfl_frr	Detour	0	7	0	3 Dis	Ena	2

The following command displays detailed configuration information for all fast-reroute profiles:

```
* BD-10K.14 # show mpls rsvp-te profile fast-reroute detail
Profile Name : prfl_frr
Peak Rate      : 0 Kbps
Committed Rate : 0 Kbps
Max Burst Size : 0 Kb
Setup Priority  : 7
Hold Priority   : 4
Hop Limit      : 1
Protected BW   : Off
Protect Node   : On
#LSP References : 2
LSP / Path     : pc10_lsp / p1
Profile Name : default_frr
Peak Rate      : 0 Kbps
Committed Rate : 0 Kbps
Max Burst Size : 0 Kb
Setup Priority  : 4
Hold Priority   : 0
Hop Limit      : 2
Protected BW   : On
Protect Node   : On
#LSP References : 4
LSP / Path     : pc_frr_lsp / p2
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

show mpls static lsp

```
show mpls static lsp {summary | {lsp_name} {detail}}
```

Description

Displays the configuration of one or all static LSPs.



Syntax Description

summary	Displays only static LSP summary information.
<i>lsp_name</i>	Identifies the LSP to be displayed.
detail	Displays additional information about the static LSPs, including packet and byte counts.

Default

N/A.

Usage Guidelines

If no command options are specified, all defined static LSPs are displayed in tabular format. The information displayed includes the configured ingress label, egress label, next-hop router IP address, and the MPLS interface status for the egress path. The summarized list of static LSPs is displayed in alphabetical order based on the LSP name.

Example

The following command displays detailed information about an ingress static LSP:

```
show mpls static lsp statlsp1 detail
Static LSP Name : statlsp1           LSP Type      : Ingress-Lsp
Destination    : 11.100.100.55      Next Hop Addr  : 12.220.0.30
Ingress I/F    : **None**           Egress I/F    : m3vlan1
Admin Status   : Disabled            Oper Status    : Disabled
Ing-Label      : **None**           Eg-Label      : 0x7FF00
IP Traffic     : Allow               VPN Traffic    : Allow
Rx Packets     : n/a                 Tx Packets     : 0
Rx Bytes       : n/a                 Tx Bytes       : 0
```

The following command displays detailed information about a transit static LSP:

```
show mpls static lsp statlsp1 detail
Static LSP Name : statlsp1           LSP Type      : Transit-Lsp
Destination    : 11.100.100.55      Next Hop Addr  : 12.224.0.55
Ingress I/F    : m2vlan1            Egress I/F    : m5vlan1
Admin Status   : Disabled            Oper Status    : Disabled
Ing-Label      : 0x7FF00            Eg-Label      : 0x80300
IP Traffic     : Not Applicable      VPN Traffic    : Not Applicable
Rx Packets     : 0                   Tx Packets     : 0
Rx Bytes       : 0                   Tx Bytes       : 0
```

The following command displays detailed information about an egress static LSP:

```
show mpls static lsp statlsp1 detail
Static LSP Name : statlsp1           LSP Type      : Egress-Lsp
Destination    : 11.100.100.55      Next Hop Addr  : **None**
```



```

Ingress I/F   : m3vlan1           Egress I/F   : **None**
Admin Status  : Disabled          Oper Status   : Disabled
Ing-Label     : 0x80300           Eg-Label     : **None**
IP Traffic    : Not Applicable    VPN Traffic   : Not Applicable
Rx Packets    : 0                 Tx Packets    : n/a
Rx Bytes      : 0                 Tx Bytes      : n/a

```

History

This command was first available in ExtremeXOS 12.1.

The output for this command was modified in ExtremeXOS 12.2.2.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show mpls statistics l2vpn

```
show mpls statistics l2vpn {vpls_name | vpws_name } {detail}
```

Description

Displays MPLS statistics for one or all Layer 2 VPNs.

Syntax Description

<i>vpls_name</i>	Specifies VPLS for which to display statistics.
<i>vpws_name</i>	Specifies VPWS for which to display statistics.
detail	Displays additional information about the Layer 2 VPNs.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays statistics for all Layer 2 VPNs:

```

Switch.1 # show mpls statistics l2vpn
VPN ID      Peer IP           RxPackets      RxBytes      TxPackets

```



```

TxBytes
-----
-
99          11.100.100.219          32866      4967734      14005
2394407
11.100.100.218          398          8235          577          10583
2009        11.100.100.219          0            0            0
0
2008        11.100.100.219          0            0            15
688

```

The following command displays detailed statistics for all Layer 2 VPNs:

```

Switch.2 # (debug) 12.2 # show mpls statistics l2vpn detail
VPNID (L2VPN Name)
Peer IP          State RxLabel TxLabel LSPTxLabel NextHopI/F
RxPackets        RxBytes      TxPackets      TxBytes
-----
-
99 (jwcvpls)
11.100.100.219 Up      x80402 x80402 x00010      tordoze
32866          4967734          14005          2394407
11.100.100.218 Up      x80407 x80405 x00011      tornext
398            8235            577            10583
2009 (pws-1)
11.100.100.219 Up      x80403 x80403 x00010      tordoze
0              0              0              0
2008 (pws-2)
11.100.100.219 Up      x80404 x80404 x00010      tordoze
0              0              15             688

```

The following show output is for ExtremeXOS 15.4:

```

L2VPN Name
Peer IP/LSP      Flags      RxPackets      RxBytes      TxPackets      TxBytes
-----
vpls1
  1.1.1.1        U          0              0              0              0
  lsp1           +          0              0              0              0
  lsp2           +          0              0              0              0

```

Flags: (U) Up, (D) Down, (R) Ready, (S) Signaling, (+) In-use,
(-) Not In-use

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



show vpls

```
show vpls {{vpls_name}} {peer ipaddress} {detail} | summary
```

Note



This command has been replaced with the following command: `show l2vpn {vpls {{vpls_name}} | vpws {{vpws_name}} {peeripaddress} {detail} | summary}`. This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Displays VPLS and H-VPLS configuration and status information.

Syntax Description

<i>vpls_name</i>	Displays information for the specified vpls.
<i>ipaddress</i>	Specifies a VPLS peer for which to display information.
detail	Displays additional information in comprehensive detail format.

Default

N/A.

Usage Guidelines

The `show vpls` command (without any optional parameters) displays all currently configured VPLS instances for the switch. The summarized list of VPLS instances is displayed in alphabetical order based on the *vpls_name*. Peers are displayed in the reverse of the order they were added.

When you specify a VPLS peer, the display includes a list of all PWs established to the peer, the PW status and PW ID, and information about each VPLS to which this peer belongs.

The `show l2vpn` command describes the display fields that appear when this command is entered with the **detail** option.

Table 47: Selected show vpls Field Definitions

Field	Definition
VPLS Name	VPLS instance or domain name.
VPN ID	Virtual Private Network identifier.
Source Address	Source IP address.
VCCV Status	Virtual Circuit Connectivity Verification (VCCV) feature status, which is either Enabled or Disabled .
VCCV Interval Time	Displays the configured VCCV interval time.



Table 47: Selected show vpls Field Definitions (continued)

Field	Definition
VCCV Fault Multiplier	Displays the configured VCCV fault multiplier.
Redundancy Type	Displays the configured VPLS redundancy type, which is EAPS , ESRP , or None .
Service Interface	Displays a VLAN or VMAN interface name.
Admin State	Displays the administrative state of the VPLS, which is either Enabled or Disabled .
Oper State	Displays the operational state of the VPLS, which is either Enabled or Disabled .
MTU	Displays the maximum transmission unit (MTU) size for the VPLS.
Ethertype	Displays the ethertype for the service interface.
.Iq tag	Displays the 802.Iq priority tag for the VPLS.
Peer IP	Displays the IP address for the VPLS peer.
PW State	<p>PW State represents the state, or status, of a PW. The possible PW state values are:</p> <ul style="list-style-type: none"> • UP—The PW is fully operational and installed in hardware. Traffic is forwarded over PW and VPLS service VLAN/VMAN. • Down—The PW is not operational and is not installed in hardware. This only happens when the VPLS instance is disabled, VPLS service is disabled, or there is no service VLAN assigned to the VPLS. No traffic is forwarded. • Sgnl—The PW is in a signalling state. The PW is not operational, and no traffic is forwarded. This can occur for a number of reasons, including: No LDP adjacency to peer, No transport LSP to peer, No VC LSP to peer. • Remote peer not configuredReady—The PW has been signalled, but it has not been installed in hardware. Traffic is not forwarded. The PW can be in a Ready state for a number of reasons, including: <ul style="list-style-type: none"> • The VPLS instance is configured for EAPS redundancy, and the EAPS shared port associated with this VPLS instance is Connected. • The VPLS instance is configured for ESRP redundancy, and the ESRP domain associated with this VPLS instance is Slave. • The service VLAN associated with this VPLS instance is down. • The remote peer has signalled that it has a fault (remote PW status). The remote peer may have a fault due to its service VLAN being down.
PW Uptime	PW Uptime is the elapsed time that the PW has been in the UP state.
PW Installed	PW Installed is a flag to indicate whether the PW is installed in hardware or not. If the PW is in the UP state, this field is True, otherwise, this field is False.



Table 47: Selected show vpls Field Definitions (continued)

Field	Definition
Local PW Status	Local PW Status displays the VC status of the local PW. The values are: <ul style="list-style-type: none"> • No Faults—No faults detected. • PW-Tx—Local PSN-facing PW transmit fault. This is set if there is a problem with the VPLS transport LSP. • PW-Rx—Local PSN-facing PW receive fault. This is set if there is a problem with the VPLS transport LSP. • Att-Tx—Local attachment circuit transmit fault. This is set if there is a problem with the VPLS service VLAN. • Att-Rx—Local attachment circuit receive fault. This is set if there is a problem with the VPLS service VLAN. • Not Forwarding—The local PW is not forwarding. Look for more information in the PW State field. For example, if VPLS is configured for EAPS redundancy, the Local PW Status is Not Forwarding and the PW State is Ready whenever the EAPS Shared Port state is Connected.
Remote PW Status	Remote PW Status is the VC status of the remote PW. The values for this field are the same values as for Local PW Status.
PW Mode	PW Mode describes how the PW was configured. The values are: <ul style="list-style-type: none"> • Core-to-Core—This VPLS instance is a core node, and the other end of the PW connects to a core node. • Core-to-Spoke—This VPLS instance is a core node, and the other end of the PW connects to a spoke node. This is for HVPLS. • Spoke-to-Core—This VPLS instance is a spoke node, and the other end of the PW connects to a core node. This is for HVPLS.
Transport LSP	Transport LSP is the LSP that is used to forward frames over the PW. When an LDP LSP is used as a transport, the display shows LDP LSP (Not configured). If an RSVP LSP is used, the name of the RSVP LSP being used as a transport LSP is displayed. An RSVP LSP can be specified as the LSP to use during VPLS configuration.
Next Hop I/F	Displays the interface name for the next hop router.
Next Hop Addr	Displays the interface IP address for the next hop router.
PW Rx Label	Receive label for the VPLS PW.
PW Rx Pkts	Total packets received on the VPLS PW.
PW Rx Bytes	Total bytes received on the VPLS PW.
Tx Label	Transmit label for the LSP.
PW Tx Label	Transmit label for the VPLS PW.
PW Tx Pkts	Total packets transmitted on the VPLS PW.
PW Tx Bytes	Total bytes transmitted on the VPLS PW.



Example

The following example shows the display that appears when you enter the show vpls command without the **detail** option:

```
Switch.38 # show vpls
L2VPN Name      VPN ID  Flags  Services Name  Peer IP      State  Flags
-----
Pws-3344        20     EAX--W  NONE
jwcvpls         99     EAX--L  torix          11.100.100.219  Up     C--NV-
keeper          90     EAX--L  NONE
pws-1           2009   EAX--W  pwserve        11.100.100.219  Up     ----V-
pws-10          70     EAX--W  NONE
pws-2           2008   EAX--W  pw2serve       11.100.100.219  Up     ---NV-
pws-3           2007   EAX--W  NONE
sarsparilla     80     EAX--W  NONE
whoopwoo        100    EAX--L  NONE           11.100.100.219  Down   C--N--
VPN Flags:  (E) Admin Enabled, (A) Oper Active, (I) Include Tag,
(X) Exclude Tag, (T) Ethertype Configured,
(V) VCCV HC Enabled, (W) VPN Type VPWS, (L) VPN Type VPLS
Peer Flags: (C) Core Peer, (S) Spoke Peer, (A) Active Core,
(p) Configured Primary Core, (s) Configured Secondary Core,
(N) Named LSP Configured, (V) VCCV HC Capabilities Negotiated,
(F) VCCV HC Failed
-----
Total number of configured L2VPNs:    9
Total number of active L2VPNs:       3
Total number of configured PWs:      4
Total number of active PWs:          3
PWs auto-selecting transport LSP:    1
PWs configured with a transport LSP:  3
PWs using LDP for transport:         0
PWs using RSVP for transport:        4
PWs using static for transport:      0
```

The following command shows summary L2 VPN information for the specified VPLS peer:

```
Switch.451 # sh vpls peer 2.2.2.2
L2VPN Name      VPN ID  Flags  Services Name  Peer IP      State  Flags
-----
vs1             105    EAX--L  cust1         2.2.2.2      UP     CAp-V-
VPN Flags:  (E) Admin Enabled, (A) Oper Active, (I) Include Tag,
(X) Exclude Tag, (T) Ethertype Configured,
(V) VCCV HC Enabled, (W) VPN Type VPWS, (L) VPN Type VPLS
Peer Flags: (C) Core Peer, (S) Spoke Peer, (A) Active Core,
(p) Configured Primary Core, (s) Configured Secondary Core,
(N) Named LSP Configured, (V) VCCV HC Capabilities Negotiated,
(F) VCCV HC Failed
```

The following command shows detailed L2 VPN information for the specified VPLS peer:

```
Switch.452 # sh vpls peer 11.100.100.210 detail
VPLS Name : vpls10
VPN ID      : 10          Admin State : Enabled
```



```

Source Address      : 11.100.100.212      Oper State   : Enabled
VCCV Status        : Disabled            MTU          : 1500
VCCV Interval Time : 5 sec.              Ethertype    : 0x8100
VCCV Fault Multiplier : 4                .1q tag     : exclude
L2VPN Type         : VPLS                Redundancy   : None
Service Interface  : vlan10
Peer IP : 11.100.100.210
PW State           : Up
PW Uptime          : 18d:0h:28m:26s
PW Installed       : True
Local PW Status    : No Faults
Remote PW Status   : No Faults
PW Mode            : Core-to-Core
Transport LSP      : LDP LSP (Not Configured)
Next Hop I/F      : o6vlan1
Next Hop Addr     : 12.182.0.216          Tx Label     : 0x00010
PW Rx Label       : 0x80405              PW Tx Label  : 0x80401
PW Rx Pkts        : 3806161633          PW Tx Pkts   : 4294967296
PW Rx Bytes       : 912385942           PW Tx Bytes  : 4294967296
MAC Limit         : No Limit
VCCV HC Status    : Not Sending (VCCV Not Enabled For This VPLS)
CC Type           : Rtr Alert             Total Pkts Sent : 0
CV Type           : LSP Ping              Total Pkts Rcvd : 0
Send Next Pkt     : --
Total Failures    : 0                    Pkts During Last Failure : 0
Last Failure Tm   : --

```

The following command shows the status of L2VPN Sharing configuration. If L2VPN Sharing is enabled, and more than 1 Transport LSP is configured, the output will display the status of each Transport LSP. The following example shows VPLS instance "vpls1" with LSP sharing enabled, with 2 Transport LSPs configured and programmed into HW.

```

* dut1.11 # show vpls detail
L2VPN Name: vpls1
  VPN ID                : 91                Admin State   : Enabled
  Source Address        : 20.20.20.91       Oper State    : Enabled
  VCCV Status           : Disabled           MTU           : 1500
  VCCV Interval Time    : 5 sec.            Ethertype     : 0x8100
  VCCV Fault Multiplier : 4                 .1q tag      : exclude
  L2VPN Type            : VPLS              Redundancy    : None
  Service Interface     : vlan91
  Peer IP: 20.20.20.83
  PW State              : Up
  PW Uptime             : 2d:22h:28m:44s
  PW Installed          : True
  Local PW Status       : No Faults
  Remote PW Status      : No Faults
  Remote I/F MTU        : 1500
  PW Mode               : Core-to-Core
  Transport LSP         : lsp1 (Configured)
    Next Hop I/F        : vlan21
    Next Hop Addr       : 10.0.21.92        Tx Label     : 0x00074

```




```

Transport LSP      : lsp2 (Configured)
  Next Hop I/F    : vlan21
  Next Hop Addr   : 10.0.21.92           Tx Label   : 0x00077

PW Rx Label       : 0x00075             PW Tx Label  : 0x00085

PW Rx Pkts       : 0                   PW Tx Pkts   : 0
PW Rx Bytes      : 0                   PW Tx Bytes   : 0
MAC Limit        : No Limit
VCCV HC Status   : Not Sending (VCCV Not Enabled For This L2VPN)

  CC Type         : Rtr Alert           Total Pkts Sent : 0
  CV Type         : LSP Ping           Total Pkts Rcvd : 0
  Send Next Pkt  : --
  Total Failures : 0                   Pkts During Last Failure : 0
  Last Failure Tm: --

```

```

-----
Total number of configured L2VPNs: 1
Total number of active L2VPNs: 1
Total number of configured PWs: 1
Total number of active PWs: 1
Total number of ready PWs: 0
PWs auto-selecting transport LSP: 0
PWs configured with a transport LSP: 1
PWs using LDP for transport: 0
PWs using RSVP for transport: 1
PWs using static for transport: 0

```

* - Indicates LSP is configured, but not installed in HW because L2VPN Sharing is disabled.

The following example shows VPLS instance "vpls1" with LSP sharing disabled, with 2 Transport LSPs configured, but only 1 LSP programmed into HW (because sharing is disabled).

```

* dut1.11 # show vpls detail
L2VPN Name: vpls1
  VPN ID           : 91                 Admin State      : Enabled
  Source Address   : 20.20.20.91       Oper State       : Enabled
  VCCV Status      : Disabled           MTU              : 1500
  VCCV Interval Time : 5 sec.          Ethertype        : 0x8100
  VCCV Fault Multiplier : 4            .1q tag          : exclude
  L2VPN Type       : VPLS               Redundancy       : None
  Service Interface : vlan91

Peer IP: 20.20.20.83
  PW State         : Up
  PW Uptime        : 2d:22h:28m:44s
  PW Installed     : True
  Local PW Status  : No Faults
  Remote PW Status : No Faults
  Remote I/F MTU   : 1500
  PW Mode          : Core-to-Core
  Transport LSP    : lsp1 (Configured)
  Next Hop I/F    : vlan21

```



```

Next Hop Addr : 10.0.21.92                Tx Label : 0x00074
Transport LSP : lsp2 (Configured*)
Next Hop I/F  : vlan21
Next Hop Addr : 10.0.21.92                Tx Label : 0x00077

PW Rx Label   : 0x00075                    PW Tx Label : 0x00085

PW Rx Pkts    : 0                          PW Tx Pkts  : 0
PW Rx Bytes   : 0                          PW Tx Bytes : 0
MAC Limit     : No Limit
VCCV HC Status : Not Sending (VCCV Not Enabled For This L2VPN)

CC Type       : Rtr Alert                    Total Pkts Sent : 0
CV Type       : LSP Ping                     Total Pkts Rcvd : 0
Send Next Pkt : --
Total Failures : 0                          Pkts During Last Failure : 0
Last Failure Tm : --

```

```

-----
Total number of configured L2VPNs: 1
Total number of active L2VPNs: 1
Total number of configured PWs: 1
Total number of active PWs: 1
Total number of ready PWs: 0
PWs auto-selecting transport LSP: 0
PWs configured with a transport LSP: 1
PWs using LDP for transport: 0
PWs using RSVP for transport: 1
PWs using static for transport: 0

```

* - Indicates LSP is configured, but not installed in HW because L2VPN Sharing is disabled.

This example show VPLS instance “vpls1” with LSP sharing enabled, 2 LSPs configured, but only 1 LSP programmed into HW. (LSP2 is disabled).

```

* dut1.13 # show vpls detail    L2VPN
Name: vpls1
  VPN ID           : 91                Admin State   : Enabled
  Source Address   : 20.20.20.91       Oper State    : Enabled
  VCCV Status      : Disabled           MTU           : 1500
  VCCV Interval Time : 5 sec.           Ethertype     : 0x8100
  VCCV Fault Multiplier : 4             .lq tag       : exclude
  L2VPN Type       : VPLS               Redundancy    : None
  Service Interface : vlan91

Peer IP: 20.20.20.83      PW
  State              : Up
  PW Uptime          : 2d:22h:30m:13s
  PW Installed       : True
  Local PW Status    : No Faults
  Remote PW Status   : No Faults
  Remote I/F MTU     : 1500
  PW Mode            : Core-to-Core
  Transport LSP      : lsp1 (Configured)

```



```

Next Hop I/F      : vlan21
Next Hop Addr    : 10.0.21.92                Tx Label   : 0x00074

Transport LSP    : lsp2 (Configured)
Next Hop I/F    : --
Next Hop Addr    : --                        Tx Label   : --

PW Rx Label      : 0x00075                    PW Tx Label : 0x00085

PW Rx Pkts      : 0                          PW Tx Pkts  : 0
PW Rx Bytes     : 0                          PW Tx Bytes  : 0
MAC Limit       : No Limit
VCCV HC Status  : Not Sending (VCCV Not Enabled For This L2VPN)

CC Type         : Rtr Alert                    Total Pkts Sent : 0

CV Type         : LSP Ping                     Total Pkts Rcvd : 0

Send Next Pkt   : --
Total Failures  : 0                          Pkts During Last Failure : 0

Last Failure Tm: --

```

```

-----
Total number of configured L2VPNs: 1
Total number of active L2VPNs: 1
Total number of configured PWs: 1
Total number of active PWs: 1
Total number of ready PWs: 0
PWs auto-selecting transport LSP: 0
PWs configured with a transport LSP: 1
PWs using LDP for transport: 0
PWs using RSVP for transport: 1
PWs using static for transport: 0

```

* - Indicates LSP is configured, but not installed in HW because L2VPN Sharing is disabled.

The following example shows VPLS instance "vpls1" with LSP Sharing enabled, with 5 Transport LSPs configured and programmed into HW.

```

* dut1.19 # show vpls detail
L2VPN Name: vpls1
VPN ID      : 91                      Admin State : Enabled
Source Address : 20.20.20.91          Oper State  : Enabled
VCCV Status   : Disabled              MTU        : 1500
VCCV Interval Time : 5 sec.          Ethertype   : 0x8100
VCCV Fault Multiplier : 4            .lq tag    : exclude
L2VPN Type    : VPLS                  Redundancy  : None
Service Interface : vlan91

Peer IP: 20.20.20.83
PW State      : Up
PW Uptime    : 2d:22h:31m:54s
PW Installed  : True

```



```

Local PW Status      : No Faults
Remote PW Status    : No Faults
Remote I/F MTU      : 1500
PW Mode             : Core-to-Core
Transport LSP       : lsp1 (Configured)
  Next Hop I/F      : vlan21
  Next Hop Addr     : 10.0.21.92                Tx Label : 0x00074

Transport LSP       : lsp2 (Configured)
  Next Hop I/F      : vlan21
  Next Hop Addr     : 10.0.21.92                Tx Label : 0x00077
Transport LSP       : lsp3 (Configured)
  Next Hop I/F      : vlan21
  Next Hop Addr     : 10.0.21.92                Tx Label : 0x00075
Transport LSP       : lsp4 (Configured)
  Next Hop I/F      : vlan21
  Next Hop Addr     : 10.0.21.92                Tx Label : 0x00076
Transport LSP       : lsp5 (Configured)
  Next Hop I/F      : vlan21
  Next Hop Addr     : 10.0.21.92                Tx Label : 0x00078

PW Rx Label         : 0x00075                    PW Tx Label : 0x00085

PW Rx Pkts          : 0                          PW Tx Pkts  : 0
PW Rx Bytes         : 0                          PW Tx Bytes : 0
MAC Limit           : No Limit
VCCV HC Status     : Not Sending (VCCV Not Enabled For This L2VPN)

  CC Type           : Rtr Alert                    Total Pkts Sent : 0
  CV Type           : LSP Ping                      Total Pkts Rcvd : 0
  Send Next Pkt    : --
  Total Failures   : 0                            Pkts During Last Failure : 0
  Last Failure Tm : --

```

```

-----
Total number of configured L2VPNs: 1
Total number of active L2VPNs: 1
Total number of configured PWs: 1
Total number of active PWs: 1
Total number of ready PWs: 0
PWs auto-selecting transport LSP: 0
PWs configured with a transport LSP: 1
PWs using LDP for transport: 0
PWs using RSVP for transport: 1
PWs using static for transport: 0

```

* - Indicates LSP is configured, but not installed in HW because L2VPN Sharing is disabled.

History

This command was first available in ExtremeXOS 11.6.

This command was updated to display flags for H-VPLS spoke nodes and protected VPLS and H-VPLS in ExtremeXOS 12.1.



The output for this command was modified in ExtremeXOS 12.2.2.

The output for this command was modified in ExtremeXOS 15.4 to display LSP sharing information.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

traceroute mpls lsp

```
traceroute mpls lsp [lsp_name | any host | prefix ipNetmask] {reply-mode [ip | ip-router-alert]} {{from from} {ttl ttl} {next-hop hopaddress}}
```

Description

Traces the path an LSP takes for the specified FEC.

Syntax Description

<i>lsp_name</i>	Specifies the LSP on which to send the MPLS echo request.
any	Allows the echo request to be sent over any available LSP.
<i>host</i>	Specifies the FEC using an ipaddress or hostname.
prefix	Specifies a prefix.
<i>ipNetmask</i>	Specifies the prefix address.
reply-mode	Specifies the reply mode for the MPLS echo response.
ip	Requests an IP UDP reply packet. This is the default mode.
ip-router-alert	Requests an IP UDP reply packet with the IP Router Alert option.
<i>from</i>	Specifies the IP address to be used as the source address in the MPLS echo request.
<i>ttl</i>	Specifies the starting TTL hop value. The range is from 1 - 30. The default is 1.
<i>hopaddress</i>	Specifies the next-hop address.

Default

The maximum time-to-live value is 30 seconds.

Usage Guidelines

This command traces the path an LSP takes for the specified FEC. The `traceroute` command, with the `mpls` keyword option, works by repeatedly sending an MPLS echo request (or “MPLS Ping”). The TTL value is incremented for each successive MPLS echo request sent. The sending LSR waits 5 seconds before sending the next MPLS echo request. This operation continues until either the egress LSR for the FEC is reached, the maximum TTL value is reached, or the operation is interrupted. For each response received, the following information is displayed on the console:



- IP address of the replying LSR
- Return code
- Indication of an MPLS echo reply timeout if no response was received

The FEC can be specified using the `ipaddress` or `hostname` via the `host` parameter. If the optional `next-hop` is specified, the MPLS echo request is sent along the LSP that traverses the specified node. This option is useful for tracing a specific LSP when multiple LSPs exist to the specified FEC. The `lsp` keyword may be used to specify a named LSP to trace. The selected LSP is specified by the `lsp_name` parameter. The `any` keyword indicates that the switch can trace any available LSP to the specified host.

The optional `reply-mode` keyword is used to specify the reply mode for the MPLS echo response. When the `ip` option is specified, the MPLS echo reply is routed back to the sender in a normal IPv4 packet. When the `ip-router-alert` option is specified, the MPLS echo reply is routed back to the sender in an IPv4 packet with the Router Alert IP option set. Additionally, if the `ip-router-alert` option is specified and the reply route is via an LSP, the Router Alert Label is pushed onto the top of the label stack. If the `reply-mode` is not specified, the `reply-mode ip` option applies.

The optional `ttl` keyword specifies the starting TTL value in the MPLS echo request packet. Within each router along the path, the TTL value is decremented. When the TTL value reaches zero, the LSR drops the packet and replies with a TTL-expired Internet Control Message Protocol (ICMP) message. The originating LSR responds by displaying the hop for which the TTL expired. To discover all hops to a destination, the originating router repeats the MPLS echo request and increments the TTL start value by one each time until the destination is reached. The maximum TTL is 30, so the `traceroute` command terminates if the destination is not reached in 30 hops.

If the `ttl` keyword is omitted, the starting TTL value is 1. If you specify a larger starting TTL value, initial hops are excluded from the `traceroute` display. For example, if you specify a start TTL value of 5, the TTL value does not decrement to 0 at the first four routers, so the fifth hop router is the first to appear in the `traceroute` command display.

The `from` keyword is used to specify the source IP address used in the MPLS echo request. This is the IP address used by the target LSR to send the MPLS echo reply. If not specified, the OSPF router ID is used.

Example

The following example shows a sample display for the `traceroute` command:

```
BD-10K.5 # traceroute mpls lsp prefix 11.100.100.10/32
traceroute to 11.100.100.10, 30 hops max
 1  11.100.100.8                5 ms          5 ms          2 ms
 2  11.100.100.10              2 ms          1 ms          2 ms
BD-10K.6 #
BD-10K.6 # traceroute mpls lsp lsp598
traceroute to lsp598, 30 hops max
 1  11.100.100.5                6 ms          1 ms          5 ms
 2  11.100.100.9                3 ms          2 ms          2 ms
 3  11.100.100.8                3 ms          4 ms          3 ms
BD-10K.7 #
```



History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure l2vpn dot1q ethertype

```
unconfigure l2vpn [vp1s vp1s_name | vpws vpws_name] dot1q ethertype
```

Description

Resets the ethertype setting for the specified VPLS or VPWS.

Syntax Description

<i>vp1s_name</i>	Identifies the VPLS within the switch (character string).
<i>vpws_name</i>	Identifies the VPWS within the switch (character string).

Default

N/A.

Usage Guidelines

The setting is changed back to the value displayed in the `show dot1q` command.

The `l2vpn` keyword is introduced in ExtremeXOS Release 12.4 and is required when resetting the ethertype for a VPWS. For backward compatibility, the `l2vpn` keyword is optional when resetting the ethertype for a VPLS. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The following command changes the ethertype setting for the specified VPLS to the value displayed in the `show dot1q` command:

```
unconfigure l2vpn vp1s my_vp1s dot1q ethertype
```

History

This command was first available in ExtremeXOS 11.6.



The **l2vpn** and **vpws** keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure l2vpn vpls redundancy

```
unconfigure {l2vpn} vpls vpls_name redundancy [eaps | esrp | stp]
```

Description

Disassociates the VPLS instance from EAPS, an ESRP domain, or STP.

Syntax Description

<i>vpls_name</i>	Specifies the VPLS for which you are removing protection.
eaps	Disassociates the VPLS instance from EAPS.
esrp	Disassociates the VPLS instance from the ESRP domain.
stp	Disassociates the VPLS instance from STP.

Default

Redundancy disabled.

Usage Guidelines

The **l2vpn** keyword is introduced in ExtremeXOS Release 12.4. For backward compatibility, the **l2vpn** keyword is optional. However, this keyword will be required in a future release, so Extreme Networks recommends that you use this keyword for new configurations and scripts.

Example

The following command disassociates the VPLS instance from ESRP:

```
unconfigure l2vpn vpls vpls1 redundancy esrp
```

History

This command was first available in ExtremeXOS 12.1.

The **l2vpn** keyword and the **STP** option were added in ExtremeXOS 12.4.



Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure mpls

unconfigure mpls

Description

Resets MPLS configuration parameters to the default settings.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command deletes all VLANs from MPLS and resets all MPLS configuration parameters to their default values. The parameters that are reset include the LSR ID, all LDP-specific settings, all RSVP-TE-specific settings, all RSVP-TE reserved bandwidth, and all EXP Qos Profile mappings. MPLS must be disabled to unconfigure MPLS.



Note

MPLS must be disabled to globally unconfigure MPLS.

Example

The following command resets MPLS configuration parameters to the default settings:

```
unconfigure mpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



unconfigure mpls exp examination

```
unconfigure mpls exp examination [all | {{value} value}]
```

Description

Resets the QoS profile assigned to the EXP value back to the default QoS profile.

Syntax Description

<code>value</code>	Specifies the EXP value whose QoS profile is reset.
--------------------	---

Default

The QoS profile matches the EXP value + 1.

Usage Guidelines

This command resets the QoS profile assigned to the EXP value (defined by the value keyword and argument) back to the default QoS profile. If the **all** option is specified, all EXP values are reset back to their default QoS profiles. By default, the QoS profile matches the EXP value + 1. That is, EXP value of 0 is mapped to QoS profile qp1, EXP value of 1 is mapped to QoS profile qp2, etc. This configuration has switch-wide significance.

Example

Use the following command to restore the QoS profile of EXP value 5 to its default setting:

```
unconfigure mpls exp examination value 5
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

unconfigure mpls exp replacement

```
unconfigure mpls exp replacement [all | {{qosprofile} qosprofile}]
```



Description

Resets the EXP value assigned to the QoS profile back to the default EXP value.

Syntax Description

<i>qosprofile</i>	Specifies the QoS profile whose EXP replacement value is unconfigured.
-------------------	--

Default

The EXP value matches the QoS profile -1.

Usage Guidelines

This command resets the EXP value assigned to the QoS defined by *qosprofile* back to the default EXP value. If the **all** option is specified, all QoS profiles are reset back to their default EXP values. By default, the EXP value matches the QoS profile - 1. That is, QoS profile qp1 is mapped to EXP value of 0, QoS profile qp2 is mapped to EXP value of 1, etc. This configuration has switch-wide significance.

Example

Use the following command to restore all EXP values to their default setting:

```
unconfigure mpls exp replacement all
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure mpls vlan

```
unconfigure mpls {{vlan} vlan_name}
```

Description

Resets MPLS configuration parameters to the default settings. This command does not delete the VLAN from MPLS.

Syntax Description

<i>vlan_name</i>	Specifies the VLANs for which MPLS is unconfigured.
------------------	---



Default

N/A.

Usage Guidelines

This command resets all MPLS configuration parameters for the specified VLAN to their default values. It does not delete the VLAN from MPLS. These parameters include the enable state for LDP and RSVP-TE, the bandwidth reserved for RSVP-TE LSPs, RSVP-TE timers, and the RSVP-TE metric. MPLS does not have to be disabled to unconfigure a specific VLAN.

Example

The following command resets MPLS configuration parameters to the default settings for a single VLAN:

```
unconfigure mpls vlan boone
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure vpls dot1q ethertype

```
unconfigure vpls vpls_name dot1q ethertype
```

Note



This command has been replaced with the following command: `unconfigure l2vpn [vpls vpls_name | vpwsvpws_name] dot1q ethertype`. This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Unconfigures the ethertype setting for the VPLS specified by `vpls_name`.

Syntax Description

<code>vpls_name</code>	Identifies the VPLS within the switch (character string)
<code>e</code>	



Default

N/A

Usage Guidelines

This command unconfigures the ethertype setting for the VPLS specified by *vpls_name*. The setting is changed back to the value displayed in the `show dot1q` command.

Example

The following command changes the ethertype setting for the specified VPLS to the value displayed in the `show dot1q` command:

```
unconfigure vpls my_vpls dot1q ethertype
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

unconfigure vpls snmp-vpn-identifier

```
unconfigure vpls vpls_name snmp-vpn-identifier
```

Description

Removes an SNMP VPN identifier for traps from the specified VLPLS.

Syntax Description

<i>vpls_name</i>	Specifies the VPLS for which you are removing the identification string.
------------------	--

Default

N/A.

Usage Guidelines

None.



Example

The following command removes the identifier for SNMP VPN traps on VPLS vpls1:

```
unconfigure vpls vpls1 snmp-vpn-identifier
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



34 IP Unicast Commands

```
clear ip dad
clear iparp
configure bootprelay add
configure bootprelay delete
configure bootprelay dhcp-agent information check
configure bootprelay dhcp-agent information circuit-id port-information
configure bootprelay dhcp-agent information circuit-id vlan-information
configure bootprelay dhcp-agent information option
configure bootprelay dhcp-agent information policy
configure bootprelay dhcp-agent information remote-id
configure bootprelay include-secondary
configure bootprelay vlan include-secondary
configure forwarding sharing
configure ip dad
configure iparp add
configure iparp add proxy
configure iparp delete
configure iparp delete proxy
configure iparp distributed-mode
configure iparp max_entries
configure iparp max_pending_entries
configure iparp max_proxy_entries
configure iparp timeout
configure ipforwarding originated-packets
configure iproute add (IPv4)
configure iproute add blackhole
configure iproute add blackhole ipv4 default
configure iproute add default
configure iproute delete
configure iproute delete blackhole
configure iproute delete blackhole ipv4 default
configure iproute delete default
configure iproute priority
configure iproute reserved-entries
configure iproute sharing hash-algorithm crc
configure iproute sharing max-gateways
configure irdp
```

```
configure vlan add secondary-ipaddress
configure vlan delete secondary-ipaddress
configure vlan subvlan
configure vlan subvlan-address-range
configure vlan delete secondary-ipaddress
disable bootp vlan
disable bootprelay
disable icmp address-mask
disable icmp parameter-problem
disable icmp port-unreachables
disable icmp redirects
disable icmp time-exceeded
disable icmp timestamp
disable icmp unreachable
disable icmp userredirects
disable iparp checking
disable iparp refresh
disable ipforwarding
disable ip-option loose-source-route
disable ip-option record-route
disable ip-option record-timestamp
disable ip-option router-alert
disable ip-option strict-source-route
disable iproute bfd
disable iproute compression
disable iproute sharing
disable irdp
disable subvlan-proxy-arp vlan
disable udp-echo-server
enable bootp vlan
enable bootprelay
enable icmp address-mask
enable icmp parameter-problem
enable icmp port-unreachables
enable icmp redirects
enable icmp time-exceeded
enable icmp timestamp
enable icmp unreachable
enable icmp userredirects
enable iparp checking
enable iparp refresh
enable ipforwarding
```




```
enable ip-option record-route
enable ip-option record-timestamp
enable ip-option strict-source-route
enable ip-option router-alert
enable iproute bfd
enable iproute compression
enable iproute sharing
enable irdp
enable subvlan-proxy-arp vlan
enable udp-echo-server
rtlookup
run ip dad
show bootprelay
show bootprelay configuration
show bootprelay dhcp-agent information circuit-id port-information
show bootprelay dhcp-agent information circuit-id vlan-information
show bootprelay configuration ipv4
show bootprelay ipv6
show ip dad
show iparp
show iparp distributed-mode statistics
show iparp proxy
show iparp security
show iparp stats
show ipconfig
show iproute
show iproute mpls
show iproute mpls origin
show iproute origin
show iproute reserved-entries
show iproute reserved-entries statistics
show ipstats
show udp-profile
unconfigure bootprelay dhcp-agent information check
unconfigure bootprelay dhcp-agent information circuit-id port-information
unconfigure bootprelay dhcp-agent information circuit-id vlan-information
unconfigure bootprelay dhcp-agent information option
unconfigure bootprelay dhcp-agent information policy
unconfigure bootprelay dhcp-agent information remote-id
unconfigure bootprelay include-secondary
unconfigure icmp
unconfigure iparp
```



[unconfigure iproute priority](#)
[unconfigure irdp](#)
[unconfigure vlan subvlan-address-range](#)
[unconfigure vlan udp-profile](#)

This chapter describes commands for configuring and managing the following IP protocols and functions:

- DHCP and BOOTP relay
- IP ARP
- IP routing
- IP multinetting
- IP broadcast handling
- IP Duplicate Address Detection (DAD)
- Broadcast UDP packet forwarding
- Static routes
- ICMP
- IRDP
- VLAN aggregation

For an introduction to these IP protocols and functions, see the *ExtremeXOS Concepts Guide*.

clear ip dad

```
clear ip dad {{vr} vr_name {ip_address} | vr all | {vlan} vlan_name} {counters}
```

Description

Clears the counters for the DAD feature on the specified IP interfaces.

Syntax Description

<i>vr_name</i>	Specifies a VR for which to clear the counters.
<i>ip_address</i>	Specifies an IP address for which to clear the counters.
<i>vlan_name</i>	Specifies a VLAN for which to clear the counters.

Default

If you do not specify a VR or VRF, the command applies to the current VR context.

Usage Guidelines

The `vr all` option clears the DAD counters for all IP interfaces on the switch.



Example

The following command clears the DAD counters for all IP interfaces in all VRs:

```
clear ip dad vr all
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on BlackDiamond X8, BlackDiamond 8800 series switches and Summit family switches.

clear iparp

```
clear iparp {ip_addr {vr vr_name} | vlan vlan_name | vr vr_name}
```

Description

Removes dynamic entries in the IP ARP table.

Syntax Description

<i>ip_addr</i>	Specifies an IP address.
<i>vlan_name</i>	Specifies a VLAN name.
<i>vr_name</i>	Specifies a Virtual Router (VR) or Virtual Router Forwarding instance (VRF) name.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

Permanent IP ARP entries are not affected.

This command is specific to a single VR or VRF, and it applies to the current VR context if you do not specify a VR or VRF.



Example

The following command removes a dynamically created entry from the IP ARP table:

```
clear iparp 10.1.1.5
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure bootprelay add

```
configure bootprelay add ip_address {vr vrid}
```

Description

Configures the addresses to which BOOTP requests should be directed.

Syntax Description

<i>ip_address</i>	Specifies an IP address.
<i>vrid</i>	Specifies a VR name.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

After IP unicast routing has been configured, you can configure the switch to forward Dynamic Host Configuration Protocol (DHCP) or BOOTP requests coming from clients on subnets being serviced by the switch and going to hosts on different subnets.

To configure the relay function, follow these steps:

- Configure VLANs and IP unicast routing.
- Configure the addresses to which DHCP or BOOTP requests should be directed, using the following command:

```
configure bootprelay add ip_address
```

- Enable the DHCP or BOOTP relay function using the following command:

```
enable bootprelay
```



Example

The following command configures BOOTP requests to be directed to 123.45.67.8:

```
configure bootprelay add 123.45.67.8
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure bootprelay delete

```
configure bootprelay delete [ip_address | all] {vr vrid}
```

Description

Removes one or all IP destination addresses for forwarding BOOTP packets.

Syntax Description

<i>ip_address</i>	Specifies an IP address.
all	Specifies all IP address entries.
<i>vrid</i>	Specifies a VR name.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

None.

Example

The following command removes the destination address:

```
configure bootprelay delete 123.45.67.8
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure bootrelay dhcp-agent information check

configure bootrelay dhcp-agent information check

Description

Enables the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82) checking.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

In some instances, a DHCP server may not properly handle a DHCP request packet containing a relay agent option. Use this command to prevent DHCP reply packets with invalid or missing relay agent options from being forwarded to the client.

To disable this check, use the following command:

```
unconfigure bootrelay dhcp-agent information check
```

Example

The following command configures the DHCP relay agent option check:

```
configure bootrelay dhcp-agent information check
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.



configure bootrelay dhcp-agent information circuit-id port-information

```
configure bootrelay dhcp-agent information circuit-id port-information port_info
port port
```

Description

Configures the circuit ID sub-option that identifies the port for an incoming DHCP request.

Syntax Description

<i>port_info</i>	Specifies a text string that becomes the circuit ID sub-option for the specified port. Specify a text string composed of 1 to 32 characters.
<i>port</i>	Specifies the port to which the circuit ID sub-option is assigned.

Default

The default *port_info* is encoded as ((slot_number * 1000) + port_number/portlindex). For example, if the DHCP request is received on port 3:12, the default circuit ID *port_info* value is 3012. On standalone switches, the slot number is one, so the default circuit ID *port_info* value is (1000 + port_number/portlindex). For example, the default *port_info* for port 3 on a standalone switch is 1003.

Usage Guidelines

The full circuit ID string uses the format *vlan_info-port_info* . To configure the *vlan_info* portion of the circuit ID string, use the following command:

```
configure bootrelay dhcp-agent information circuit-id vlan-information vlan_info
{vlan} [vlan_name | all]
```

To display the *port_info* information, use the following command:

```
show bootrelay dhcp-agent information circuit-id port-information ports all
```

Example

The following command configures the circuit ID *port_info* value slot1port3 for port 1:3:

```
configure bootrelay dhcp-agent information circuit-id port-information
slot1port3 port 1:3
```

History

This command was first available in ExtremeXOS 12.4.



Platform Availability

This command is available on all platforms.

configure bootprelay dhcp-agent information circuit-id vlan-information

```
configure bootprelay dhcp-agent information circuit-id vlan-information vlan_info
{vlan} [vlan_name|all]
```

Description

Configures the circuit ID sub-option that identifies the VLAN for an incoming DHCP request.

Syntax Description

<i>vlan_info</i>	Specifies a text string that becomes the circuit ID sub-option for the specified VLAN. Specify a text string composed of 1 to 32 characters.
<i>vlan_name</i>	Specifies the VLAN to which the circuit ID sub-option is assigned.
all	Specifies that the <i>vlan_info</i> entered is to be used in the circuit ID sub-option for all VLANs.

Default

The default *vlan_info* for each VLAN is the VLAN ID or tag.

Usage Guidelines

The full circuit ID string uses the format *vlan_info-port_info* . To configure the *port_info* portion of the circuit ID string, use the following command:

```
configure bootprelay dhcp-agent information circuit-id port-information port_info
portport
```

To display the *vlan_info* information, use the following command:

```
show bootprelay dhcp-agent information circuit-id vlan-information
```

Example

The following command configures the circuit ID *vlan_info* value VLANblue for VLAN blue:

```
configure bootprelay dhcp-agent information circuit-id vlan-information
VLANblue blue
```



History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

configure bootrelay dhcp-agent information option

configure bootrelay dhcp-agent information option

Description

Enables the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82).

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

After IP unicast routing has been configured, you can configure the switch to forward DHCP or BOOTP requests coming from clients on subnets being serviced by the switch and going to hosts on different subnets.

To configure the relay function, follow these steps:

- Configure VLANs and IP unicast routing.
- Enable the DHCP or BOOTP relay function, using the following command:

```
enable bootrelay {{vlan} [vlan_name] | {{vr}vr_name} | all [{vr}vr_name]}
```

- Configure the addresses to which DHCP or BOOTP requests should be directed, using the following command:

```
configure bootrelay add ip_address {vr vrid}
```

```
configure bootrelay dhcp-agent information option
```

Configure the DHCP relay agent option (option 82), using the following command:

To disable the DHCP relay agent option (option 82), use the following command:

```
unconfigure bootrelay dhcp-agent information option
```



Example

The following command configures the DHCP relay agent option:

```
configure bootprelay dhcp-agent information option
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure bootprelay dhcp-agent information policy

```
configure bootprelay dhcp-agent information policy [drop | keep | replace]
```

Description

Configures the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82) policy.

Syntax Description

drop	Specifies to drop the packet.
keep	Specifies to keep the existing option 82 information in place.
replace	Specifies to replace the existing data with the switch's own data.

Default

Replace.

Usage Guidelines

Use this command to set a policy for the relay agent. Packets can be dropped, the option 82 information can be replaced (the default), or the packet can be forwarded with the information unchanged.

Example

The following command configures the DHCP relay agent option 82 policy to keep:

```
configure bootprelay dhcp-agent information policy keep
```



History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure bootprelay dhcp-agent information remote-id

```
configure bootprelay dhcp-agent information remote-id [remote_id | system-name]
{vr vrid}
```

Description

Configures the remote ID sub-option that identifies the relaying switch for DHCP requests and replies.

Syntax Description

<i>remote_id</i>	Specifies a text string that becomes the remote ID sub-option for the switch. Specify a text string composed of 1 to 32 characters.
system-name	Specifies that the switch name is used as the remote ID sub-option for the switch.
<i>vrid</i>	Specifies the VR on which to configure the remote ID sub-option.

Default

The switch MAC address.

Usage Guidelines

To display the remote-ID, use the following command:

```
show bootprelay
```

Example

The following command configures the remote ID sub-option to specify the switch name in DHCP requests and replies:

```
configure bootprelay dhcp-agent information remote-id system-name
```

History

This command was first available in ExtremeXOS 12.4.



Platform Availability

This command is available on all platforms.

configure bootprelay include-secondary

Short reference description.

```
configure bootprelay {ipv4 | ipv6} include-secondary {sequential | parallel | off} {vr vr_name
```

Description

Configures DHCP smart relay mode, and includes a secondary IP address as the giaddr at the VR level.

Syntax Description

ipv4	Specifies DHCPv4 BOOTP Relay service (default).
ipv6	Specifies DHCPv6 BOOTP Relay service.
include-secondary	(Optional) Use both primary and secondary address(es) of the client VLAN as gateway address.
sequential	Use primary and secondary address(es) of client VLAN in sequence after 3 retries (default if include-secondary is on).
parallel	Use primary and secondary address(es) of client VLAN in parallel.
off	Disable use of both primary and secondary address(es) of client VLAN as gateway address (default).
vr	Specifies a virtual router ID.
<i>vr_name</i>	Specifies the virtual router.

Default

IPv4 is the default relay service.

off is the default value, but **sequential** is the default if **include-secondary** is on.

Usage Guidelines

Use this command to configure DHCP smart relay mode, and to include a secondary IP address as giaddr at the VR level.



Example

The following command configures DHCPv4 BOOTP Relay service to use both primary and secondary address(es) of the client VLAN as the gateway address. By default, the command specifies that you use the primary and secondary address(es) of the client VLAN in sequence after 3 retries:

```
configure bootprelay ipv4 include-secondary sequential
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

configure bootprelay vlan include-secondary

```
configure bootprelay {ipv4 | ipv6} {vlan vlan_name} include-secondary {sequential
| parallel | off}
```

Description

Configures DHCP smart relay mode to include secondary IP address as giaddr at VLAN level.

Syntax Description

ipv4	Specifies DHCPv4 BOOTP Relay service (default).
ipv6	Specifies DHCPv6 BOOTP Relay service.
vlan	Configure BOOTP relay for this VLAN, and overrides the VR level configuration.
<i>vlan_name</i>	Specifies the VLAN name.
include-secondary	(Optional) Use both primary and secondary address(es) of the client VLAN as gateway address.
sequential	Use primary and secondary address(es) of client VLAN in sequence after 3 retries (default if include-secondary is on).
parallel	Use primary and secondary address(es) of client VLAN in parallel.
off	Disable use of both primary and secondary address(es) of client VLAN as gateway address (default).

Default

IPv4 is the default relay service.

off is the default value, but **sequential** is the default if **include-secondary** is on.



Usage Guidelines

Use this command to configure DHCP smart relay mode to include the secondary IP address as giaddr at the VLAN level.

Example

The following command configures DHCPv4 BOOTP Relay service for the "vlan_100" VLAN, and uses both primary and secondary address(es) of the client VLAN as gateway address. This overrides the VR level configuration.

```
configure bootprelay ipv4 vlan vlan_100 include-secondary
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

configure forwarding sharing

```
configure forwarding sharing [L3 | L3_L4]
```

Description

Identifies the fields that are used to select ECMP routes and load-sharing group ports.

Syntax Description

L3	Uses only Layer 3 IP addresses to select ECMP routes and load-sharing ports.
L3_L4	Uses Layer 3 IP addresses and Layer4 TCP/UDP port numbers, if present, to select ECMP routes and load-sharing ports.

Default

L3_L4.

Usage Guidelines

This command configures the criteria used to select ECMP routes and load-sharing group ports.



For ECMP routes, the configured criteria selects the next hop gateway. The L3 option uses only the source and destination IP addresses to select the next hop gateway. The L3_L4 option uses the Layer4 TCP or UDP port and the source and destination IP addresses to select the next hop gateway.

For load-sharing groups (link aggregation groups), the configured criteria selects the load-sharing group port. The load-sharing groups can be configured to use the following address-based algorithms:

- L2—Specifies port selection based on Layer 2 information.
- L3—Specifies port selection based on Layer 3 information.
- L3_L4—Specifies port selection based on Layer 3 and Layer4 information.

This command affects all the load-sharing groups that use either the L3 or L3_L4 link aggregation algorithm. If the L3 option is specified, all the load-sharing groups that are configured with either the L3 or the L3_L4 address-based link aggregation algorithm use just the Layer 3 IP addresses for the egress port selection. Similarly if the L3_L4 option is specified, all the load-sharing groups that are configured with either L3 or L3_L4 address-based link aggregation algorithm use the Layer 3 IP addresses and Layer4 port number for the egress port selection.

Selecting the L3 option over L3_L4 can be useful in a network where IP fragments are present, since only the first fragment contains the Layer4 TCP or UDP port number. If the L3 option is selected, all IP fragments in a given TCP or UDP session use the same ECMP gateway or load-sharing group port, potentially avoiding inefficient packet reordering by the destination. If IP fragments are not prevalent, better traffic distribution can be achieved by selecting L3_L4.

To display the forwarding sharing feature configuration, enter the command:

```
show forwarding configuration
```

Example

The following command modifies the sharing selection criteria to use just the Layer 3 IP addresses:

```
configure forwarding sharing L3
```

The following command modified the sharing selection criteria to use the Layer 3 and Layer4 information:

```
configure forwarding sharing L3_L4
```

History

This command was first available in ExtremeXOS 11.6.4.

Platform Availability

This command is available only on the BlackDiamond X8 series switches, BlackDiamond 8800 series switch, SummitStack, and the Summit family of switches.



configure ip dad

```
configure ip dad [off | on | {on} attempts max_solicitations] [{vr} vr_name | vr all]
```

Description

Configures the operation of the duplicate address detection (DAD) feature on the specified VR.

Syntax Description

<i>max_solicitations</i>	Specifies the number of times the DAD feature tests for a duplicate address. The range is 1 to 10, and the default value is 1.
<i>vr_name</i>	Specifies a VR on which to enable this feature.

Default

DAD status: Off on VR-Default.

Maximum solicitations: 1 on VR-Default.

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

This command can be entered in a configuration file for execution at switch startup, or it can be entered at the CLI prompt.

Changes to the number of solicitations configuration take affect the next time the DAD check is run.

By default, this command applies to the current VR context, if no VR name is specified. If `vr all` is specified, the command applies to all user VRs and VR-Default.

Example

The following command enables the DAD feature on all user VRs and VR-Default:

```
configure ip dad on vr all
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on BlackDiamond X8, BlackDiamond 8800 series switches and Summit family switches.



configure iparp add

```
configure iparp add ip_addr {vrvr_name} mac
```

Description

Adds a permanent entry to the ARP table.

Specify the IP address and MAC address of the entry.

Syntax Description

<i>ip_addr</i>	Specifies an IP address.
<i>vr_name</i>	Specifies a VR name.
<i>mac</i>	Specifies a MAC address.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

None.

Example

The following command adds a permanent IP ARP entry to the switch for IP address 10.1.2.5:

```
configure iparp add 10.1.2.5 00:11:22:33:44:55
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure iparp add proxy

```
configure iparp add proxy [ipNetmask | ip_addr {mask}] {vr vr_name} {mac | vrrp}  
{always}
```



Description

Configures the switch to respond to ARP Requests on behalf of devices that are incapable of doing so.

Syntax Description

<i>ipNetmask</i>	Specifies an IP address/mask length.
<i>ip_addr</i>	Specifies an IP address.
<i>mask</i>	Specifies a subnet mask.
<i>vr_name</i>	Specifies a VR name.
<i>mac</i>	Specifies a MAC address to use in the ARP reply.
vrrp	Specifies a MAC address to use in the ARP reply. For VLANs running VRRP, the switch replies with the VRRP virtual MAC. For non-VRRP VLANs, the switch replies with the switch MAC.
always	Specifies that the switch responds regardless of the VLAN that the request arrives from.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

When *mask* is not specified, an address with the mask 255.255.255.255 is assumed. When neither *mac* nor *vrrp* is specified, the MAC address of the switch is used in the ARP Response. When *always* is specified, the switch answers ARP Requests without filtering requests that belong to the same subnet of the receiving router interface.

After IP ARP is configured, the system responds to ARP requests on behalf of the device as long as the following conditions are satisfied:

- The valid IP ARP request is received on a router interface.
- The target IP address matches the IP address configured in the proxy ARP table.
- The source IP address is not on the same subnet as the target address (unless the *always* flag is set).

After all the proxy ARP conditions have been met, the switch formulates an ARP response using the configured MAC address in the packet.

The default maximum number of proxy entries is 256, but can be increased to 4096 by using the following command:

```
configure iparp max_proxy_entries {vr vr_name}max_proxy_entries
```



Example

The following command configures the switch to answer ARP requests for all devices with the address range of 100.101.45.1 to 100.101.45.255:

```
configure iparp add proxy 100.101.45.0/24
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure iparp delete

```
configure iparp delete ip_addr {vr vr_name}
```

Description

Deletes an entry from the ARP table. Specify the IP address of the entry to delete.

Syntax Description

<i>ip_addr</i>	Specifies an IP address.
<i>vr_name</i>	Specifies a VR name.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

Removes any IP ARP entry (dynamic or permanent) from the table.

Example

The following command deletes an IP address entry from the ARP table:

```
configure iparp delete 10.1.2.5
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure iparp delete proxy

```
configure iparp delete proxy [ipNetmask | ip_addr {mask}] {vr vr_name} | all
```

Description

Deletes one or all proxy ARP entries.

Syntax Description

<i>ipNetmask</i>	Specifies an IP address/mask length.
<i>ip_addr</i>	Specifies an IP address.
<i>mask</i>	Specifies a subnet mask.
<i>vr_name</i>	Specifies a VR name.
all	Specifies all ARP entries.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

When the mask is not specified, the software assumes a host address (that is, a 32-bit mask).

Example

The following command deletes the IP ARP proxy entry 100.101.45.0/24:

```
configure iparp delete proxy 100.101.45.0/24
```

History

This command was first available in ExtremeXOS 10.1.



Platform Availability

This command is available on all platforms.

configure iparp distributed-mode

```
configure iparp distributed-mode [on | off]
```

Description

Configures the distributed IP ARP feature as either on or off.

Syntax Description

on	Configures the distributed IP ARP feature to become active after the next switch restart.
off	Configures the distributed IP ARP feature to become inactive after the next switch restart.

Default

Off.

Usage Guidelines

The distributed IP ARP feature provides higher IP ARP scaling by distributing IP ARP forwarding information to only the I/O module to which each IP host is connected. This feature is off by default to match the operation in ExtremeXOS releases prior to 12.5. When this feature is off, complete IP ARP information for all destinations is stored on all modules, reducing the available space for unique destinations.

If the sum of the values required for the `configure iparp max_entries` command for all VRs is greater than 8000, ExtremeNetworks recommends that you configure the distributed IP ARP feature as on, save the configuration, and activate the new configuration by restarting the switch.



Note

To activate or deactivate the distributed IP ARP feature, you must change the configuration, save the configuration, and restart the switch.

The distributed IP ARP feature imposes the following limitations:

- The number of unique load share groups in a BlackDiamond 8800 series switch is reduced from 128 to 64 groups. On BlackDiamond X8, the number of load share groups is not reduced.
- A load share group must include ports from only BlackDiamond 8000 c-, xl-, and xm-series modules or from modules that are not BlackDiamond 8000 c-, xl-, or xm-series modules. The switch does not support load share groups that mix ports from the two module groups when the distributed IP ARP feature is active. On BlackDiamond X8, the limitation is not applicable; load share group ports may be from any module type.



- Load distribution of IPv4 unicast packets across ports of a load share group is affected if the ingress port and any load share member ports reside on the same I/O module and port cluster within that module. Packets are distributed as expected if the ingress port is on a different module or different port cluster than all ports in the destination's load share group.

To view the configured and current operational state of the distributed IP ARP feature, use the `show iparp` command. To display distributed IP ARP statistics by slot when this feature is enabled, use the `show iparp distributed-mode statistics` command.

Example

The following command configures the distributed IP ARP feature to become active after the next switch restart:

```
configure iparp distributed-mode on
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on BlackDiamond X8 series switches, BlackDiamond 8000 c-, xl-, and xm-series modules.

configure iparp max_entries

```
configure iparp max_entries {vr vr_name} max_entries
```

Description

Configures the maximum allowed IP ARP entries.

Syntax Description

<i>vr_name</i>	Specifies a VR name.
<i>max_entries</i>	Specifies the maximum number of IP ARP entries. The range is 1 to x, where x is the number listed for the appropriate platform in table below.

Default

8192 for the Default VR and all user-created VRs.

4096 for the VR-Mgmt VR.



If you do not specify a VR or VRF, the current VR context is used.



Note

The default value for the Default and user-created VRs changed from 4,096 to 8,192 in ExtremeXOS Release 12.2.

Usage Guidelines

The maximum IP ARP entries include dynamic, static, and incomplete IP ARP entries. The range for the `max_entries` parameter is 1 to x, where x is the number listed for the appropriate platform in the following table.

Table 48: Maximum IP ARP Entries for each Platform

Platform	Maximum Entries	Distributed IP ARP Feature Configuration	
		Off (Default)	On
BlackDiamond 8806	Depends on distributed IP ARP configuration.	20480	130,000 ¹²
BlackDiamond 8810	Depends on distributed IP ARP configuration.	20480	260,000a
All others	20480	N/A.	N/A

The switch hardware supports only the maximum number of entries listed in the table above. If the hardware limit is reached, the switch displays a message and can no longer store additional ARP entries.

Example

The following command sets the maximum IP ARP entries to 2000 entries:

```
configure iparp max_entries 2000
```

History

This command was first available in ExtremeXOS 10.1.

Support for up to 32,768 ARP entries was first available in ExtremeXOS 12.4.

The support for more than 20,480 ARP entries on BlackDiamond 8800 series switches was added with the distributed IP ARP feature in ExtremeXOS 12.5.

¹² The distributed IP ARP feature must be configured as on before the switch will accept any value above 20480. The switch cannot support values above 20480 until after the distributed IP ARP mode has been configured as on and the switch has been restarted.



Platform Availability

This command is available on all platforms.

configure iparp max_pending_entries

```
configure iparp max_pending_entries {vrvr_name} max_pending_entries
```

Description

Configures the maximum allowed incomplete IP ARP entries.

Syntax Description

<i>vr_name</i>	Specifies a VR name.
<i>max_pending_entries</i>	Specifies a number of maximum IP ARP entries.

Default

256.

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

Range: 1 - 4096.

Example

The following command sets the maximum pending IP ARP entries to 500 entries:

```
configure iparp max_pending_entries 500
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure iparp max_proxy_entries

```
configure iparp max_proxy_entries {vr vr_name} max_proxy_entries
```



Description

Configures the maximum allowed IP ARP proxy entries.

Syntax Description

<i>vr_name</i>	Specifies a VR name.
<i>max_proxy_entries</i>	Specifies maximum number of IP ARP proxy entries.

Default

256.

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

Range: 0 - 4096.

Example

The following command sets the maximum IP ARP proxy entries to 500 entries:

```
configure iparp max_proxy_entries 500
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure iparp timeout

```
configure iparp timeout {vr vr_name} minutes
```

Description

Configures the IP ARP timeout period.

Syntax Description

<i>vr_name</i>	Specifies which VR or VRF IP ARP setting to change.
<i>minutes</i>	Specifies a time in minutes.



Default

20 minutes.

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

The range is 0-32,767. A setting of 0 disables timeout.

When the switch learns an ARP entry, it begins the timeout for that entry. When the timer reaches 0, the entry is aged out, unless IP ARP refresh is enabled. If ARP refresh is enabled, the switch sends an ARP request for the address before the timer expires. If the switch receives a response, it resets the timer for that address.

Example

The following command sets the IP ARP timeout period to 10 minutes:

```
configure iparp timeout 10
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure ipforwarding originated-packets

```
configure ipforwarding originated-packets [require-ipforwarding | dont-require-ipforwarding]
```

Description

Configures whether IP forwarding must be enabled on a VLAN before transmitting IP packets originated by the switch on that VLAN to a gateway.

Syntax Description

require-ipforwarding	Specifies that IP forwarding must be enabled on a VLAN before IP packets that originate on the switch can be transmitted to a gateway.
dont-require-ipforwarding	Specifies that all IP packets that originate on the switch can be transmitted, regardless of the IP forwarding configuration to the gateway.



Default

dont-require-ipforwarding.

Usage Guidelines

To display the current setting for this command, use the `show ipconfig` command.

Example

The following command configures the switch to transmit switch-originated packets to gateways only on those VLANs for which IP forwarding is enabled:

```
configure ipforwarding originated-packets require-ipforwarding
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms that use the Edge, Advanced Edge, or Core license. For information on the licenses available for each platform, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure iproute add (IPv4)

```
configure iproute add [ipNetmask | ip_addr mask] gateway {bfd} {metric}  
{multicast | multicast-only | unicast | unicast-only} {vr vrname}
```

Description

Adds a static route to the specified routing table.

Syntax Description

<i>ipNetmask</i>	Specifies an IP address/mask length.
<i>ip_addr</i>	Specifies an IP address.
<i>mask</i>	Specifies a subnet mask.
<i>gateway</i>	Specifies a gateway IP address.
bfd	Enables Bidirectional Forwarding Detection (BFD) protection for the route.
metric	Specifies a cost metric.
multicast	Adds the specified route to the multicast routing table.



multicast-only	Adds the specified route to the multicast routing table. This option is provided for backward compatibility with releases prior to ExtremeXOS Release 12.1.
unicast	Adds the specified route to the unicast routing table.
unicast-only	Adds the specified route to the unicast routing table. This option is provided for backward compatibility with releases prior to ExtremeXOS Release 12.1.
<i>vrname</i>	Specifies the VR or VRF to which the route is added.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

Use a mask value of 255.255.255.255 to indicate a host entry.

The gateway address must be present on a directly attached subnet, or the following message appears:

ERROR: Gateway is not on directly attached subnet

The gateway address must be different from the VLAN address, or the following message appears:

ERROR: Gateway cannot be own address (x.x.x.x) #where x.x.x.x is the IP address specified

Note



Although dynamic unicast routes can be captured in the multicast routing table, unicast static routes cannot be captured in the multicast routing table. To create a static route for the multicast routing table, you must specify the multicast option.

This command can add BFD protection to a link only when the BFD client at each end of the link is enabled (`configure iproute add (IPv4)` command). Once the BFD session is established, the operational status of the route reflects the operational status of the BFD session. To remove BFD protection for a static route, enter this command without the BFD keyword.

Example

The following command adds a static address to the routing table:

```
configure iproute add 10.1.1.0/24 123.45.67.1 5
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.



configure iproute add blackhole

```
configure iproute add blackhole [ipNetmask | ip_address mask] {multicast | multicast-only | unicast | unicast-only} {vr vrname}
```

Description

Adds a blackhole address to the routing table. All traffic destined for a configured blackhole IP address is silently dropped, and no Internet Control Message Protocol (ICMP) message is generated.

Syntax Description

<i>ipNetmask</i>	Specifies an IP address/mask length.
<i>ip_address</i>	Specifies an IP address.
<i>mask</i>	Specifies a subnet mask.
multicast	Adds the blackhole route to the multicast routing table.
multicast-only	Adds the blackhole route to the multicast routing table. This option is provided for backward compatibility with releases prior to ExtremeXOS Release 12.1.
unicast	Adds the blackhole route to the unicast routing table.
unicast-only	Adds the blackhole route to the unicast routing table. This option is provided for backward compatibility with releases prior to ExtremeXOS Release 12.1.
<i>vrname</i>	Specifies the VR or VRF to which the route is added.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

A blackhole entry configures packets with the specified destination IP subnet to be discarded. Blackhole entries are useful as a security measure or in special circumstances where a specific destination IP subnet must be discarded. Blackhole entries are treated like permanent entries in the event of a switch reset or power off/on cycle.

Example

The following command adds a blackhole address to the routing table for packets with a destination address of 100.101.145.4:

```
configure iproute add blackhole 100.101.145.4/32
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure iproute add blackhole ipv4 default

```
configure iproute add blackhole ipv4 default {multicast | multicast-only | unicast | unicast-only} {vr vrname}
```

Description

Adds a default blackhole route to the routing table. All traffic destined for an unknown IP destination is silently dropped, and no Internet Control Message Protocol (ICMP) message is generated.

Syntax Description

multicast	Adds the default blackhole route to the multicast routing table.
multicast-only	Adds the default blackhole route to the multicast routing table. This option is provided for backward compatibility with releases prior to ExtremeXOS Release 12.1.
unicast	Adds the default blackhole route to the unicast routing table.
unicast-only	Adds the default blackhole route to the unicast routing table. This option is provided for backward compatibility with releases prior to ExtremeXOS Release 12.1.
<i>vrname</i>	Specifies the VR or VRF to which the route is added.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

While a default route is for forwarding traffic destined to an unknown IP destination, and a blackhole route is for discarding traffic destined to a specified IP destination, a default blackhole route is for discarding traffic to the unknown IP destination.

Using this command, all traffic with an unknown destination is discarded.

The default blackhole route is treated like a permanent entry in the event of a switch reset or power off/on cycle. The default blackhole route's origin is "b" or "blackhole" and the gateway IP address for this route is 0.0.0.0.



Example

The following command adds a blackhole default route into the routing table:

```
configure iproute add blackhole default
```

History

This command was first available in ExtremeXOS 10.1.

The ipv4 keyword was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure iproute add default

```
configure iproute add default gateway [metric] {multicast | multicast-only | unicast | unicast-only} {vr vrname}
```

Description

Adds a default gateway to the routing table.

Syntax Description

<i>gateway</i>	Specifies a VLAN gateway.
<i>metric</i>	Specifies a cost metric. If no metric is specified, the default of 1 is used.
multicast	Adds the default route to the multicast routing table.
multicast-only	Adds the default route to the multicast routing table. This option is provided for backward compatibility with releases prior to ExtremeXOS Release 12.1.
unicast	Adds the default route to the unicast routing table.
unicast-only	Adds the default route to the unicast routing table. This option is provided for backward compatibility with releases prior to ExtremeXOS Release 12.1.
<i>vrname</i>	Specifies the VR or VRF to which the route is added.

Default

If no metric is specified, the default metric of 1 is used. If you do not specify a VR or VRF, the current VR context is used.



Usage Guidelines

Default routes are used when the router has no other dynamic or static route to the requested destination. A default gateway must be located on a configured IP interface. Use the unicast-only or multicast-only options to specify a particular traffic type. If not specified, both unicast and multicast traffic uses the default route.

Example

The following command configures a default route for the switch:

```
configure iproute add default 123.45.67.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure iproute delete

```
configure iproute delete [ipNetmask | ip_address mask] gateway {multicast | multicast-only | unicast | unicast-only} {vr vrname}
```

Description

Deletes a static address from the routing table.

Syntax Description

<i>ipNetmask</i>	Specifies an IP address/mask length.
<i>ip_address</i>	Specifies an IP address.
<i>mask</i>	Specifies a subnet mask.
<i>gateway</i>	Specifies a VLAN gateway.
multicast	Specifies a multicast route to delete.
multicast-only	Specifies a multicast route to delete.
unicast	Specifies a unicast route to delete.
unicast-only	Specifies a unicast route to delete.
<i>vrname</i>	Specifies the VR or VRF from which the route is deleted.



Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

Use a value of 255.255.255.255 or /32 for mask to indicate a host entry.

Example

The following command deletes an address from the gateway:

```
configure iproute delete 10.101.0.0/24 10.101.0.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure iproute delete blackhole

```
configure iproute delete blackhole [ipNetmask | ip_address mask] {multicast | multicast-only | unicast | unicast-only} {vr vrname}
```

Description

Deletes a blackhole address from the routing table.

Syntax Description

<i>ipNetmask</i>	Specifies an IP address/mask length.
<i>ip_address</i>	Specifies an IP address.
<i>mask</i>	Specifies a netmask.
multicast	Specifies a blackhole multicast route to delete.
multicast-only	Specifies a blackhole multicast-only route to delete. This option is provided for backward compatibility with releases prior to ExtremeXOS Release 12.1.
unicast	Specifies a blackhole unicast route to delete.
unicast-only	Specifies a blackhole unicast-only route to delete. This option is provided for backward compatibility with releases prior to ExtremeXOS Release 12.1.
<i>vrname</i>	Specifies the VR or VRF from which the route is deleted.



Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

None.

Example

The following command removes a blackhole address from the routing table:

```
configure iproute delete blackhole 100.101.145.4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure iproute delete blackhole ipv4 default

```
configure iproute delete blackhole ipv4 default {multicast | multicast-only | unicast | unicast-only} {vr vrname}
```

Description

Deletes a default blackhole route from the routing table.

Syntax Description

multicast	Specifies a default blackhole multicast route to delete.
multicast-only	Specifies a default blackhole multicast route to delete. This option is provided for backward compatibility with releases prior to ExtremeXOS Release 12.1.
unicast	Specifies a default blackhole unicast route to delete.
unicast-only	Specifies a default blackhole unicast-only route to delete. This option is provided for backward compatibility with releases prior to ExtremeXOS Release 12.1.
<i>vrname</i>	Specifies a VR or VRF name.

Default

If you do not specify a VR or VRF, the current VR context is used.



Usage Guidelines

None.

Example

The following command deletes a blackhole default route from the routing table:

```
configure iproute delete blackhole default
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure iproute delete default

```
configure iproute delete default gateway {multicast | multicast-only | unicast | unicast-only} {vr vrname}
```

Description

Deletes a default gateway from the routing table.

Syntax Description

<i>gateway</i>	Specifies a VLAN gateway.
multicast	Specifies a default multicast route to delete.
multicast-only	Specifies a default multicast route to delete. This option is provided for backward compatibility with releases prior to ExtremeXOS Release 12.1.
unicast	Specifies a default unicast route to delete.
unicast-only	Specifies a default unicast route to delete. This option is provided for backward compatibility with releases prior to ExtremeXOS Release 12.1.
<i>vrname</i>	Specifies the VR or VRF from which the route is deleted.

Default

If you do not specify a VR or VRF, the current VR context is used.



Usage Guidelines

Default routes are used when the router has no other dynamic or static route to the requested destination. A default gateway must be located on a configured IP interface.

Example

The following command deletes a default gateway:

```
configure iproute delete default 123.45.67.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

configure iproute priority

```
configure iproute {ipv4} priority [blackhole | bootp | ebgp | ibgp | icmp | isis
| isis-level-1 | isis-level-1-external | isis-level-2 | isis-level-2-external |
mpls | ospf-as-external | ospf-extern1 | ospf-extern2 | ospf-inter | ospf-intra |
rip | static] priority {vr vrname}
```

Description

Changes the priority for all routes from a particular route origin.

Syntax Description

blackhole	Specifies the blackhole route.
bootp	Specifies BOOTP.
ebgp	Specifies E-BGP routes
ibgp	Specifies I-BGP routes
icmp	Specifies ICMP.
isis	Specifies IS-IS and applies only to blackhole routes installed for summary addresses.
isis-level-1	Specifies ISIS Level 1 routing.
isis-level-1-external	Specifies ISIS Level 1 External routing.
isis-level-2	Specifies ISIS Level 2 routing.
isis-level-2-external	Specifies ISIS Level 2 External routing.



mpls	Specifies MPLS routing.
ospf-as-external	Specifies OSPF as External routing.
ospf-extern1	Specifies OSPF External 1 routing.
ospf-extern2	Specifies OSPF External 2 routing.
ospf-inter	Specifies OSPFInter routing.
ospf-intra	Specifies OSPFIntra routing.
rip	Specifies RIP.
static	Specifies static routes.
<i>priority</i>	Specifies a priority number in the range of 11 to 65534.
<i>vrname</i>	Specifies a VR or VRF name.

Default

The following table lists the relative priorities assigned to routes depending upon the learned source of the route.

Table 49: Relative Route Priorities

Route Origin	Priority
Direct	10
MPLS	20
Blackhole	50
Static	1100
ICMP	1200
EBGP	1700
IBGP	1900
OSPFIntra	2200
OSPFInter	2300
IS-IS	2350
IS-IS L1	2360
IS-IS L2	2370
RIP	2400
OSPFAsExt	3100
OSPF External 1	3200
OSPF External 2	3300
IS-IS L1 Ext	3400
IS-IS L2 Ext	3500
BOOTP	5000



Usage Guidelines

Although these priorities can be changed, do not attempt any manipulation unless you are expertly familiar with the possible consequences. If you change the route priority, you must save the configuration and reboot the system.



Note

The priority for a blackhole route cannot overlap with the priority of any other route origin.

Example

The following command sets IP route priority for static routing to 1200:

```
configure iproute priority static 1200
```

History

This command was first available in ExtremeXOS 10.1.

The route priority restrictions were added in ExtremeXOS 11.1.

The ipv4 keyword was added in ExtremeXOS 11.2.

The vr option was added in ExtremeXOS 12.1.2.

Platform Availability

This command is available on all platforms.

configure iproute reserved-entries

```
configure iproute reserved-entries [ num_routes_needed | maximum | default ] slot
[all | slot_num]
```

Description

Reserves storage space for IPv4 and IPv6 routes in the Longest Prefix Match (LPM) hardware tables, allowing individual local and remote IPv4 unicast hosts to occupy the unused portions of the tables.

Syntax Description

<i>num_routes_needed</i>	Specifies a specific number of routes to reserve.
maximum	Reserves the maximum amount of space for IP route entries. No IPv4 hosts are stored in the LPM and External tables.
default	Reserves the default amount of space for IP route entries.



all	For BlackDiamond X8, BlackDiamond 8800 series and SummitStack switches only, this option applies the reservation to all applicable slots.
<i>slot_num</i>	For BlackDiamond X8, BlackDiamond 8800 series and SummitStack switches only, this option applies the reservation to the specified slot.

Default

The default values are as follows:

- For a BlackDiamond 8000 e-series module: 464.
- For a BlackDiamond X8, BlackDiamond 8000 c-, or xm-series module or a Summit X460 or X670 switch: 12240.
- For a BlackDiamond 8900 xl-series module or a Summit X480 switch: 245728. Note that the default value for a BlackDiamond 8900 xl-series module or a Summit X480 switch depends on the value configured with the command: `configure forwarding external-tables`.
- For a Summit X440 switch: 16.

Usage Guidelines

Demand on the Layer 3 Hash table can be reduced by allowing IPv4 hosts to be stored in the LPM tables instead. This command allows you to reserve a portion of the LPM tables for routes, and this creates an unreserved portion that can be used to store IPv4 hosts. For more information, see the [Extended IPv4 Host Cache](#) section in the Extreme XOS Concepts Guide.

The default setting can support most networks, but if more than a few hundred local IP hosts and IP multicast entries are present, you can improve switch performance by calculating and configuring the reserved space for route entries to allow unreserved space for IPv4 hosts. Changing the number of reserved route entries does not require a reboot of the affected slots or switches.

You can view the current LPM hardware table usage by entering the `show iproute reserved-entries statistics` command. The LPM table statistics are in the columns under the In HW Route Table heading.

If the switch contains fewer routes than the capacity of the LPM tables, the number of route entries to reserve for a slot or switch should be the number of routes currently used in the hardware tables, plus an additional cushion for anticipated growth. Because each IPv6 route takes up the space of two IPv4 routes, the number of route entries to reserve is two times the value in the IPv6 routes column, plus the



value in the IPv4 routes column, plus room for anticipated growth. For example, if you want to reserve space for 100 IPv4 routes and 20 IPv6 routes, the required number of route entries is 140 (100 + 2*20).

Note



For a BlackDiamond 8900 xl-series module or a Summit X480 switch, IPv6 routes are not included in the calculation for the number of reserved route entries if the `configure forwarding external-tables` command is set to include IPv4 routes (for example, `I3-only`, `I3-only ipv4`, `I3-only ipv4-and-ipv6`, `I2-and-I3`, and `I2-and-I3-and-acl`). When the external tables are configured for IPv4 routes, IPv6 routes occupy the entire Internal LPM table, or in the case of `I3-only ipv4-and-ipv6`, IPv6 routes occupy a separate partition within the External LPM table.

On a BlackDiamond X8, BlackDiamond 8000 c-, or xm-series module, Summit X460, X670, or X770 switch, the capacity of the LPM table is 4,096 higher than the capacity for local IPv4 or IPv6 hosts. Therefore, on such hardware, there is no need to configure fewer than 4096 reserved route entries.

The maximum value for `num_routes_needed` is as follows:

- For a Summit X440 switch: 32
- For a BlackDiamond 8000 e-series module: 480
- For a BlackDiamond 8000 c-series module or a Summit X460 switch: 12256
- For a BlackDiamond X8, BlackDiamond 8900 xm-series module or a Summit X670 switch: 16352
- For a BlackDiamond 8900 xl-series module or a Summit X480 switch, the maximum value depends on the value configured with the `configure forwarding external-tables` command as follows:
 - `I2-only`, `acl-only`, `I3-only ipv6` and `none` options: 16352 (default 12240)
 - `I2-and-I3-and-acl` option: 131040 (default 114656)
 - `I2-and-I3` option: 262112 (default 245728)
 - `I3-only {ipv4}` option: 524256 (default 507872)
 - `I3-only ipv4-and-ipv6` option: 475104 (default 458720)
 - `I2-and-I3-and-ipmc` option: 131040 (default 114656)

The maximum values shown above apply to Summit family switches operating independently or as part of a SummitStack. The maximum option can be used to specify the maximum values.

When maximum is specified, IPv4 hosts do not occupy LPM table space. Note that when maximum is specified, software forwarding can result, depending on the utilization and addresses in the Layer 3 Hash table, and is therefore not recommended.

If the switch contains more routes than the capacity of the LPM tables, say 700 routes on a BlackDiamond 8000 e-series module, a trade-off can be made. You can choose to reserve 400 iproute entries, for example. The 400 IPv4 routes with the longest length network masks will be installed in the LPM table, and the remainder of the LPM table can be used for cache space for local and remote hosts. The remote host entries are only required for IPv4 addresses matching one of the 300 routes not installed in the LPM table. Since not all 700 routes can be stored on a BlackDiamond 8000 e-series module anyway, leaving appropriate room for individual remote hosts can result in more fast-path forwarding.



Depending on the actual routes present, IP route compression can be enabled to reduce the number of routes required in the LPM tables. For more information, see the description for the following command:

```
enable iproute compression {vr vrname}
```

Example

The following command reserves up to 140 IPv4 routes or 70 IPv6 routes, or any combination in between, on all BlackDiamond X8, BlackDiamond 8000 series modules, or on all Summit X440, X460, X480, and X670 switches in a SummitStack:

```
configure iproute reserved-entries 140 slot all
```

For details on the configuration changes, see the command descriptions for the following commands:

```
show iproute reserved-entries
show iproute reserved-entries statistics
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on BlackDiamond X8, BlackDiamond 8000 series modules and on Summit X440, X460, X480, X670, and X770 switches when operating independently or in a SummitStack.



configure iproute sharing hash-algorithm crc

```
configure iproute sharing hash-algorithm crc [lower | upper]
```

Description

This command is used to configure the hash algorithm used to choose a gateway when hardware forwards an IPv4 or IPv6 unicast packet to a route with multiple equal-cost multipath gateways. The values within the IP unicast packet that are considered in the hash calculation depend on the setting of another command, “configure forwarding sharing [L3 | L3_L4]”. With the default, L3_L4, the hash calculation includes Source and Destination IP addresses, and the Source and Destination Layer 4 Port numbers. Or, if “configure forwarding sharing L3” is configured, the hash calculation only includes Source and Destination IP addresses. The distribution of packets among multiple gateways based on the IP Route Sharing lower or upper hash algorithm will depend on network traffic. The command will not result in traffic loss and takes effect immediately.



Syntax Description

iproute	IP routing module.
sharing	Configure settings for equal cost multipath routing";capability="route_sharing.
hash-algorithm	Configure hardware forwarding hash algorithm used to select among ECMP gateways for an IPv4 or IPv6 destination";capability="pib".
crc	Cyclic Redundancy Check (CRC).
lower	Lower bits of CRC32 hash calculation of source and destination packet criteria, used to select an ECMP gateway (Default).
upper	Upper bits of CRC32 hash calculation. May improve distribution when source and destination IP and ports do not vary much.

Default

Lower.

Usage Guidelines

Use this command to configure the hash algorithm used to choose a gateway when hardware forwards an IPv4 or IPv6 unicast packet to a route with multiple equal-cost multipath gateways. The values within the IP unicast packet that are considered in the hash calculation depend on the setting of another command, "configure forwarding sharing [L3 | L3_L4]". With the default, L3_L4, the hash calculation includes Source and Destination IP addresses, and the Source and Destination Layer 4 Port numbers. Or, if "configure forwarding sharing L3" is configured, the hash calculation only includes Source and Destination IP addresses. The distribution of packets among multiple gateways based on the IP Route Sharing lower or upper hash algorithm will depend on network traffic. The command will not result in traffic loss and takes effect immediately.

Example

```
# configure iproute sharing hash-algorithm upper
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

configure iproute sharing max-gateways

```
configure iproute sharing max-gateways max_gateways
```



Description

Specifies the maximum number of gateways in each gateway set in the equal-cost multipath (ECMP) hardware table.

Syntax Description

<code>max_gateways</code>	Specifies the maximum number of ECMP gateways in a gateway. The only values allowed are 2, 4, 8, 16 and 32.
---------------------------	---

Default

4 gateways.

Usage Guidelines

When IP route sharing is enabled, the maximum number of gateways value represents the maximum number of next-hop gateways that can be used for communications with a destination subnet. Each gateway represents an alternative path to a subnet. The gateways can be defined with static routes, or they can be learned through the OSPF, BGP, or IS-IS protocols.

The *ExtremeXOS Release Notes* lists the total number of route destinations and the total combinations of gateway sets that each platform can support with the different max-gateways option selections. For more information on selecting the maximum number of gateways and how this affects different platforms, see [ECMP Hardware Table](#) in the *ExtremeXOS Concepts Guide*.

You must save the configuration and reboot the switch for the new value to take effect. To see the current and configured value, use the following command:

```
show ipconfig
```

To see the current and configured value, use the following command:

```
show ipconfig ipv6
```

Example

The following command changes the maximum number of ECMP gateways per subnet or gateway set to eight:

```
configure iproute sharing max-gateways 8
```

History

This command was first available in ExtremeXOS 11.4.

The value 2 was first available in ExtremeXOS 12.0.2.

Support for shared gateway sets in the ECMP table was added in ExtremeXOS 12.4.



The values 16 and 32 were first available in ExtremeXOS 15.3.

This command first applied to IPv6 routes in ExtremeXOS 15.3.

Platform Availability

This command is available only on BlackDiamond X8, BlackDiamond 8000 series modules, SummitStack, and Summit X460, X480, X670, and X770 series switches.

configure irdp

```
configure irdp [multicast | broadcast | mininterval maxinterval lifetime preference]
```

Description

Configures the destination address of the router advertisement messages.

Syntax Description

multicast	Specifies multicast setting.
broadcast	Specifies broadcast setting.
<i>mininterval</i>	Specifies the minimum time between advertisements.
<i>maxinterval</i>	Specifies the maximum time between advertisements. Default is 600.
<i>lifetime</i>	Specifies the lifetime of the advertisement. Default is 1800.
<i>preference</i>	Specifies the router preference level. Default is 0.

Default

Broadcast (255.255.255.255). The default mininterval is 450.

Usage Guidelines

ICMP Router Discovery Protocol allows client machines to determine what default gateway address to use. The switch sends out IP packets at the specified intervals identifying itself as a default router. IRDP enabled client machines use this information to determine which gateway address to use for routing data packets to other networks.

Example

The following command sets the address of the router advertiser messages to multicast:

```
configure irdp multicast
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms that use the Edge, Advanced Edge, or Core license. For information on the licenses available for each platform, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vlan add secondary-ipaddress

```
configure vlan vlan_name add secondary-ipaddress [ip_address {netmask} | ipNetmask]
```

Description

Configures secondary IP addresses on a VLAN to support multinetting.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>ip_address</i>	Specifies an IP address.
<i>netmask</i>	Specifies a network mask.
<i>ipNetmask</i>	Specifies an IP address with network mask.

Default

N/A.

Usage Guidelines

Adding a secondary IP address to a VLAN enables multinetting. Secondary addresses are added to support legacy stub IP networks.

Once you have added a secondary IP address to a VLAN, you cannot unconfigure the primary IP address of that VLAN until you delete all the secondary addresses. Delete secondary address with the following command:

```
configure vlan vlan_name delete secondary-ipaddress [ip_address | all]
```



Example

The following command configures the VLAN multi to support the 10.1.1.0/24 subnet in addition to its primary subnet:

```
configure vlan multi add secondary-ipaddress 10.1.1.1/24
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure vlan delete secondary-ipaddress

```
configure vlan vlan_name delete secondary-ipaddress [ip_address | all]
```

Description

Removes secondary IP addresses on a VLAN that were added to support multinetting.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>ip_address</i>	Specifies an IP address.
all	Specifies all secondary IP addresses.

Default

N/A.

Usage Guidelines

Once you have added a secondary IP address to a VLAN, you cannot unconfigure the primary IP address of that VLAN until you delete all the secondary addresses. Use the all keyword to delete all the secondary IP addresses from a VLAN.

Example

The following command removes the 10.1.1.0 secondary IP address from the VLAN multi:

```
configure vlan multi delete secondary-ipaddress 10.1.1.1
```



History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

configure vlan subvlan

```
configure vlan vlan_name [add | delete] subvlan sub_vlan_name
```

Description

Adds or deletes a subVLAN to a superVLAN.

Syntax Description

<i>vlan_name</i>	Specifies a superVLAN name.
add	Specifies to add the subVLAN to the superVLAN.
delete	Specifies to delete the subVLAN from the superVLAN.
<i>sub_vlan_name</i>	Specifies a subVLAN name.

Default

N/A.

Usage Guidelines

The following properties apply to VLAN aggregation operation:

- All broadcast and unknown traffic remains local to the subVLAN and does not cross the subVLAN boundary. All traffic within the subVLAN is switched by the subVLAN, allowing traffic separation between subVLANs (while using the same default router address among the subVLANs).
- Hosts can be located on the superVLAN or on subVLANs. Each host can assume any IP address within the address range of the superVLAN router interface. Hosts on the subVLAN are expected to have the same network mask as the superVLAN and have their default router set to the IP address of the superVLAN.
- All IP unicast traffic between subVLANs is routed through the superVLAN. For example, no ICMP redirects are generated for traffic between subVLANs, because the superVLAN is responsible for subVLAN routing. Unicast IP traffic across the subVLANs is facilitated by the automatic addition of an ARP entry (similar to a proxy ARP entry) when a subVLAN is added to a superVLAN. This feature can be disabled for security purposes.



Example

The following command adds the subVLAN `vsub1` to the superVLAN `vsuper`:

```
configure vlan vsuper add subvlan vsub1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

configure vlan subvlan-address-range

```
configure vlan vlan_name subvlan-address-range ipaddress1 ipaddress2
```

Description

Configures subVLAN address ranges on each subVLAN to prohibit the entry of IP addresses from hosts outside of the configured range.

Syntax Description

<i>vlan_name</i>	Specifies a subVLAN name.
<i>ipaddress1</i>	Specifies an IP address.
<i>ipaddress2</i>	Specifies another IP address.

Default

N/A.

Usage Guidelines

There is no error checking to prevent the configuration of overlapping subVLAN address ranges between multiple subVLANs. Doing so can result in unexpected behavior of ARP within the superVLAN and associated subVLANs.



Example

The following command configures the subVLAN vsuper to prohibit the entry of IP addresses from hosts outside of the configured range of IP addresses:

```
configure vlan vsuper subvlan-address-range 10.1.1.1 - 10.1.1.255
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

configure vlan delete secondary-ipaddress

```
configure vlan vlan_name delete secondary-ipaddress [ip_address | all]
```

Description

Removes secondary IP addresses on a VLAN that were added to support multinetting.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>ip_address</i>	Specifies an IP address.
all	Specifies all secondary IP addresses.

Default

N/A.

Usage Guidelines

Once you have added a secondary IP address to a VLAN, you cannot unconfigure the primary IP address of that VLAN until you delete all the secondary addresses. Use the all keyword to delete all the secondary IP addresses from a VLAN.

Example

The following command removes the 10.1.1.0 secondary IP address from the VLAN multi:

```
configure vlan multi delete secondary-ipaddress 10.1.1.1
```



History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

disable bootp vlan

```
disable bootp vlan [vlan | all]
```

Description

Disables the generation and processing of BOOTP packets on a VLAN to obtain an IP address for the VLAN from a BOOTP server.

Syntax Description

<i>vlan</i>	Specifies a VLAN name.
all	Specifies all VLANs.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables the generation and processing of BOOTP packets on a VLAN named accounting:

```
disable bootp vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.



disable bootprelay

```
disable bootprelay {{vlan} [vlan_name] | {{vr} vr_name} | all [{vr} vr_name]}
```

Description

Disables the BOOTP relay function on one or all VLANs for the specified VR or VRF.

Syntax Description

<i>vlan_name</i>	Specifies a single VLAN on which to disable the BOOTP relay feature.
<i>vr_name</i>	Specifies a single VR on which to disable the BOOTP relay feature.
all	Specifies that BOOTP relay is to be disabled for all VLANs on the specified VR or VRF.

Default

The BOOTP relay function is disabled on all VLANs and VRs.

Usage Guidelines

Because VLAN names are unique on the switch, you can specify only a VLAN name (and omit the VR name) to disable BOOTP relay. When you disable BOOTP relay on a VR or VRF, BOOTP relay is disabled on all VLANs for that VR. If you enter the command without specifying a VLAN or a VR, the functionality is disabled for all VLANs in the current VR context.

Example

The following command disables the forwarding of BOOTP requests on all VLANs in the current VR context:

```
disable bootprelay
```

You can use either of the following commands to disable the forwarding of BOOTP requests on VLAN unit2:

```
disable bootprelay unit2
disable bootprelay vlan unit2
```

You can use any one of the following commands to disable the forwarding of BOOTP requests on all VLANs in VR zone3:

```
disable bootprelay zone3
disable bootprelay vr zone3
```



```
disable bootprelay all zone3
disable bootprelay all vr zone3
```

History

This command was first available in ExtremeXOS 10.1.

The capability to disable BOOTP relay on a VLAN was added in ExtremeXOS 12.4.2.

Platform Availability

This command is available on all platforms.

disable icmp address-mask

```
disable icmp address-mask {vlan name}
```

Description

Disables the generation of an ICMP address-mask reply on one or all VLANs.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

Disabled.

Usage Guidelines

Disables the generation of an ICMP address-mask reply (type 18, code 0) when an ICMP address mask request is received. The default setting is disabled. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command disables the generation of an ICMP address-mask reply on VLAN accounting:

```
disable icmp address-mask vlan accounting
```



History

This command was first available in ExtremeXOS 10.1.

The default was changed to disabled in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

disable icmp parameter-problem

```
disable icmp parameter-problem {vlan name}
```

Description

Disables the generation of an ICMP parameter-problem message on one or all VLANs.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

Enabled.

Usage Guidelines

Disables the generation of an ICMP parameter-problem message (type 12) when the switch cannot properly process the IP header or IP option information. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command disables the generation of an ICMP parameter-problem message on VLAN accounting:

```
disable icmp parameter-problem vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.



Platform Availability

This command is available on all platforms.

disable icmp port-unreachables

```
disable icmp port-unreachables {vlan name}
```

Description

Disables the generation of ICMP port unreachable messages on one or all VLANs.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

Enabled.

Usage Guidelines

Disables the generation of ICMP port unreachable messages (type 3, code 3) when a TCP or UDP request is made to the switch and no application is waiting for the request, or an access policy denies the request. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command disables ICMP port unreachable messages on VLAN accounting:

```
disable icmp port-unreachables vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable icmp redirects

```
disable icmp redirects {ipv4} {vlan all | {vlan} {name}}
```



Description

Disables the generation of ICMP redirect messages on one or all VLANs.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

Enabled.

Usage Guidelines

Disables the generation of ICMP redirects (Type 5) to hosts who direct routed traffic to the switch where the switch detects that there is another router in the same subnet with a better route to the destination.

Example

The following command disables ICMP redirects from VLAN accounting:

```
disable icmp redirects vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable icmp time-exceeded

```
disable icmp time-exceeded {vlan name}
```

Description

Disables the generation of ICMP time exceeded messages on one or all VLANs.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------



Default

Enabled.

Usage Guidelines

Disables the generation of an ICMP time exceeded message (type 11) when the TTL field expires during forwarding. IP multicast packets do not trigger ICMP time exceeded messages. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command disables the generation of ICMP time exceeded messages on VLAN accounting:

```
disable icmp time-exceeded vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable icmp timestamp

```
disable icmp timestamp {vlan name}
```

Description

Disables the generation of an ICMP timestamp response on one or all VLANs.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

Enabled.



Usage Guidelines

Disables the generation of an ICMP timestamp response (type 14, code 0) when an ICMP timestamp request is received. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command disables the generation of an ICMP timestamp response on VLAN accounting:

```
disable icmp timestamp vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable icmp unreachableables

```
disable icmp unreachableables {vlan name}
```

Description

Disables the generation of ICMP unreachable messages on one or all VLANs.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

Enabled.

Usage Guidelines

Disables the generation of an ICMP timestamp response (type 3, code 0) when an ICMP timestamp request is received. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.



Example

The following command disables the generation of ICMP unreachable messages on all VLANs:

```
disable icmp unreachable
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable icmp useredirects

```
disable icmp useredirects
```

Description

Disables the modification of route table information when an ICMP redirect message is received.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This option only applies to the switch when the switch is not in routing mode.

If the switch has a route to a destination network, the switch uses that router as the gateway to forward the packets to. If that router knows about a better route to the destination, and the next hop is in the same subnet as the originating router, the second router sends an ICMP redirect message to the first router. If ICMP useredirects is disabled, the switch disregards these messages and continues to send the packets to the second router.

Example

The following command disables the changing of routing table information:

```
disable icmp useredirects
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable iparp checking

```
disable iparp {vr vr_name} checking
```

Description

Disable checking if the ARP request source IP address is within the range of the local interface or VLAN domain.

Syntax Description

<i>vr_name</i>	Specifies a VR or VRF.
----------------	------------------------

Default

Enabled.

Usage Guidelines

If you do not specify a VR or VRF, the command applies to the current VR context.

Example

The following command disables IP ARP checking:

```
disable iparp checking
```

History

This command was first available in ExtremeXOS 10.1.

The vr option was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.



disable iparp refresh

```
disable iparp {vr vr_name} refresh
```

Description

Disables IP ARP to refresh its IP ARP entries before timing out.

Syntax Description

<i>vr_name</i>	Specifies a VR or VRF.
----------------	------------------------

Default

Enabled.

Usage Guidelines

The purpose of disabling ARP refresh is to reduce ARP traffic in a high node count Layer 2 switching only environment.

If you do not specify a VR or VRF, the command applies to the current VR context.

Example

The following command disables IP ARP refresh:

```
disable iparp refresh
```

History

This command was first available in ExtremeXOS 10.1.

The vr option was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

disable ipforwarding

```
disable ipforwarding {broadcast} {vlan vlan_name}
```



Description

Disables routing (or routing of broadcasts) for one or all VLANs. If no argument is provided, disables routing for all VLANs.

Syntax Description

broadcast	Specifies broadcast IP forwarding.
<i>vlan_name</i>	Specifies a VLAN name.

Default

Disabled.

Usage Guidelines

Disabling IP forwarding also disables broadcast forwarding. Broadcast forwarding can be disabled without disabling IP forwarding. When new IP interfaces are added, IP forwarding (and IP broadcast forwarding) is disabled by default.

Other IP related configuration is not affected.

Example

The following command disables forwarding of IP broadcast traffic for a VLAN named accounting:

```
disable ipforwarding broadcast vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

The ignore-broadcast and fast-direct-broadcast keywords were added in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms that use the Edge, Advanced Edge, or Core license. For information on the licenses available for each platform, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable ip-option loose-source-route

```
disable ip-option loose-source-route
```



Description

Disables processing of the loose source route IP option in the IPv4 packet header.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Disables the switch from forwarding IP packets with the IP option for loose source routing turned on. Packets with the loose-source-route option enabled are dropped by the switch.

Example

The following command disables processing of the loose source route IP option:

```
disable ip-option loose-source-route
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable ip-option record-route

```
disable ip-option record-route
```

Description

Disables processing of the record route IP option in the IPv4 packet header.

Syntax Description

This command has no arguments or variables.

Default

Enabled.



Usage Guidelines

Disables the switch from adding itself into the IP options header when the record route IP option is enabled in a packet that is transiting the switch.

Example

The following command disables processing of the record route IP option:

```
disable ip-option record-route
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable ip-option record-timestamp

```
disable ip-option record-timestamp
```

Description

Disables processing of the record timestamp IP option in the IPv4 packet header.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Disables the switch from adding a timestamp into the IP options header when it receives a packet with the record timestamp IP option.

Example

The following command disables processing of the record timestamp IP option:

```
disable ip-option record-timestamp
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable ip-option router-alert

disable ip-option router-alert

Description

Disables processing of the router alert IP option in IPv4 packet headers.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

None.

Example

The following command disables processing of the router alert IP option:

```
disable ip-option router-alert
```

History

This command was first available in EXOS 10.1.

Platform Availability

This command is available on all platforms.

disable ip-option strict-source-route

disable ip-option strict-source-route



Description

Disables processing the strict source route IP option in the IPv4 packet header.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Disables the switch from forwarding IP packets that have the strict source routing IP option turned on. The switch drops packets that have the strict source routing IP option enabled.

Example

The following command disables processing of the strict source route IP option:

```
disable ip-option strict-source-route
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

disable iproute bfd

```
disable iproute bfd {gateway} ip_addr {vr vrname}
```

Description

Disables BFD client services for IPv4 static routes.

Syntax Description

<i>ip_addr</i>	Specifies the IPv4 address of a neighbor for which BFD services are to be stopped.
<i>vrname</i>	Specifies the VR or VRF name for which BFD services are being disabled.



Default

Disabled.

Usage Guidelines

When the BFD client is disabled, BFD services for all static IP routes terminates. This command does not disable services for other BFD clients (such as the MPLS BFD client).

Example

The following example disables BFD client protection for communications with neighbor 10.10.10.1:

```
disable iproute bfd 10.10.10.1
```

History

This command was first available in ExtremeXOS 12.5.3.

Platform Availability

This command is available on all platforms.

disable iproute compression

```
disable iproute compression {vr vrname}
```

Description

Disables IPv4 route compression.

Syntax Description

<i>vrname</i>	VR or VRF name for which the IP route compression is being disabled. If the VR or VRF name is not specified, route compression is disabled for the VR context from which CLI command is issued.
---------------	---

Default

Disabled.

Usage Guidelines

Disables IPv4 route compression for a specified VR or VRF.



Example

The following example disables IP route compression:

```
disable iproute compression
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

disable iproute sharing

```
disable iproute {ipv4} sharing {{{vr} vrname} | { {vr} all}}
```

Description

Disables IPv4 route sharing.

Syntax Description

vrname	VR or VRF name for which IP route sharing is being disabled.
---------------	--

Default

Disabled.

Usage Guidelines

If a VR is not specified, this command disables IP route sharing in the current VR context.

Example

The following command disables load sharing for multiple routes:

```
disable iproute sharing
```

History

This command was first available in ExtremeXOS 12.1.

The vr option was added in ExtremeXOS 12.5.



Platform Availability

This command is available on all platforms.

disable irdp

```
disable irdp {vlan name}
```

Description

Disables the generation of ICMP router advertisement messages on one or all VLANs.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

Disabled.

Usage Guidelines

If no optional argument is specified, all the IP interfaces are affected.

Example

The following command disables IRDP on VLAN accounting:

```
disable irdp vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms that use the Edge, Advanced Edge, or Core license. For information on the licenses available for each platform, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable subvlan-proxy-arp vlan

```
disable subvlan-proxy-arp vlan [vlan-name | all]
```



Description

Disables the automatic entry of subVLAN information in the proxy ARP table.

Syntax Description

<i>vlan-name</i>	Specifies a superVLAN name.
all	Specifies all VLANs.

Default

Enabled.

Usage Guidelines

To facilitate communication between subVLANs, by default, an entry is made in the IP ARP table of the superVLAN that performs a proxy ARP function. This allows clients on one subVLAN to communicate with clients on another subVLAN. In certain circumstances, intra-subVLAN communication may not be desired for isolation reasons.



Note

The isolation option works for normal, dynamic, ARP-based client communication.

Example

The following command disables the automatic entry of subVLAN information in the proxy ARP table of the superVLAN vsuper:

```
disable subvlan-proxy-arp vlan vsuper
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

disable udp-echo-server

```
disable udp-echo-server {vr vrid}
```

Description

Disables UDP echo server support.



Syntax Description

<i>vrid</i>	Specifies a VR or VRF.
-------------	------------------------

Default

Disabled.

Usage Guidelines

UDP Echo packets are used to measure the transit time for data between the transmitting and receiving end.

Example

The following command disables UDP echo server support:

```
disable udp-echo-server
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable bootp vlan

```
enable bootp vlan [vlan | all]
```

Description

Enables the generation and processing of BOOTP packets on a VLAN to obtain an IP address for the VLAN from a BOOTP server.

Syntax Description

<i>vlan</i>	Specifies a VLAN name.
all	Specifies all VLANs.

Default

Disabled.



Usage Guidelines

None.

Example

The following command enables the generation and processing of BOOTP packets on a VLAN named accounting:

```
enable bootp vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable bootprelay

```
enable bootprelay [{vlan} [vlan_name] | {vr} vr_name] | all [{vr} vr_name]
```

Description

Enables the BOOTP relay function on one or all VLANs for the specified VR or VRF.

Syntax Description

<i>vlan_name</i>	Specifies a single VLAN on which to enable the BOOTP relay feature.
<i>vr_name</i>	Specifies a single VR on which to enable the BOOTP relay feature.
all	Specifies that BOOTP relay is to be enabled for all VLANs on the specified VR or VRF.

Default

The BOOTP relay function is disabled on all VLANs and VRs.

Usage Guidelines

Because VLAN names are unique on the switch, you can specify only a VLAN name (and omit the VR name) to enable BOOTP relay on a particular VLAN. When you enable BOOTP relay on a VR or VRF,



BOOTP relay is enabled on all VLANs for that VR. If you enter the command without specifying a VLAN or a VR, the functionality is enabled for all VLANs in the current VR context.

Note

If DHCP/BOOTP Relay is enabled on a per VLAN basis, make sure it is enabled on both the client-side and server-side VLANs.

You can enable the use of LSP next hops, or you can enable DHCP/BOOTP relay. The software does not support both features at the same time.

Example

The following command enables the forwarding of BOOTP requests for all VLANs in the current VR context:

```
enable bootprelay
```

You can use either of the following commands to enable the forwarding of BOOTP requests for VLAN client1:

```
enable bootprelay "client1"  
enable bootprelay vlan "client1"
```

You can use any one of the following commands to enable the forwarding of BOOTP requests for all VLANs on VR zone3:

```
enable bootprelay zone3  
enable bootprelay vr zone3  
enable bootprelay all zone3  
enable bootprelay all vr zone3
```

History

This command was first available in ExtremeXOS 10.1.

The capability to enable BOOTP relay on a VLAN was added in ExtremeXOS 12.4.2.

Platform Availability

This command is available on all platforms.

enable icmp address-mask

```
enable icmp address-mask {vlan name}
```



Description

Enables the generation of an ICMP address-mask reply on one or all VLANs.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

Disabled.

Usage Guidelines

Enables the generation of an ICMP address-mask reply (type 18, code 0) when an ICMP address mask request is received. The default setting is disabled. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command enables the generation of an ICMP address-mask reply on VLAN accounting:

```
enable icmp address-mask vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

The default was changed to disabled in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

enable icmp parameter-problem

```
enable icmp parameter-problem {vlan name}
```

Description

Enables the generation of an ICMP parameter-problem message on one or all VLANs.



Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

Enabled.

Usage Guidelines

Enables the generation of an ICMP parameter-problem message (type 12) when the switch cannot properly process the IP header or IP option information. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command enables the generation of an ICMP parameter-problem message on VLAN accounting:

```
enable icmp parameter-problem vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable icmp port-unreachables

```
enable icmp port-unreachables {vlan name}
```

Description

Enables the generation of ICMP port unreachable messages on one or all VLANs.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------



Default

Enabled.

Usage Guidelines

Enables the generation of ICMP port unreachable messages (type 3, code 3) when a TCP or UDP request is made to the switch and no application is waiting for the request, or when an access policy denies the request. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command enables ICMP port unreachable messages on VLAN accounting:

```
enable icmp port-unreachables vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable icmp redirects

```
enable icmp redirects {ipv4} {vlan all | {vlan} {name}}
```

Description

Enables the generation of ICMP redirect messages on one or all VLANs.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

Enabled.

Usage Guidelines

This option only applies to the switch when the switch is in routing mode.



ICMP redirects are used in the situation where there are multiple routers in the same subnet. If a host sends a packet to one gateway, the gateway router looks at its route table to find the best route to the destination. If it sees that the best route is through a router in the same subnet as the originating host, the switch sends an ICMP redirect (type 5) message to the host that originated the packet, telling it to use the other router with the better route. The switch also forwards the packet to the destination.

ICMP redirects are only generated for IPv4 unicast packets that are "slowpath" forwarded by the CPU. That is, IPv4 packets that contain IP Options, or packets whose Destination IP is not in the Layer 3 forwarding hardware table.

Example

The following command enables the generation of ICMP redirect messages on all VLANs:

```
enable icmp redirects
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable icmp time-exceeded

```
enable icmp time-exceeded {vlan name}
```

Description

Enables the generation of ICMP time exceeded messages on one or all VLANs.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

Enabled.

Usage Guidelines

Enables the generation of an ICMP time exceeded message (type 11) when the TTL field expires during forwarding. IP multicast packets do not trigger ICMP time exceeded messages. If a VLAN is not specified, the command applies to all IP interfaces.



This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command enables the generation of ICMP time exceeded messages on VLAN accounting:

```
enable icmp time-exceeded vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable icmp timestamp

```
enable icmp timestamp {vlan name}
```

Description

Enables the generation of an ICMP timestamp response on one or all VLANs.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

Enabled.

Usage Guidelines

Enables the generation of an ICMP timestamp response (type 14, code 0) when an ICMP timestamp request is received. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.



Example

The following command enables the generation of an ICMP timestamp response on VLAN accounting:

```
enable icmp timestamp vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable icmp unreachablees

```
enable icmp unreachablees {vlan name}
```

Description

Enables the generation of ICMP unreachable messages on one or all VLANs.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

Enabled.

Usage Guidelines

Enables the generation of an ICMP timestamp response (type 3, code 0) when an ICMP timestamp request is received. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command enables the generation of ICMP unreachable messages on all VLANs:

```
enable icmp unreachablees
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable icmp useredirects

enable icmp useredirects

Description

Enables the modification of route table information when an ICMP redirect message is received.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

If the switch has a route to a destination network, the switch uses that router as the gateway to forward the packets to. If that router knows about a better route to the destination, and the next hop is in the same subnet as the originating router, the second router sends an ICMP redirect message to the originating router. If ICMP useredirects is enabled, the switch adds a route to the destination network using the third router as the next hop and starts sending the packets to the third router.

Example

The following command enables the modification of route table information:

```
enable icmp useredirects
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.



enable iparp checking

```
enable iparp {vr vr_name} checking
```

Description

Enables checking if the ARP request source IP address is within the range of the local interface or VLAN domain.

Syntax Description

<i>vr_name</i>	Specifies a VR or VRF.
----------------	------------------------

Default

Enabled.

Usage Guidelines

If you do not specify a VR or VRF, the command applies to the current VR context.

Example

The following command enables IP ARP checking:

```
enable iparp checking
```

History

This command was first available in ExtremeXOS 10.1.

The vr option was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

enable iparp refresh

```
enable iparp {vr vr_name} refresh
```

Description

Enables IP ARP to refresh its IP ARP entries before timing out.



Syntax Description

<i>vr_name</i>	Specifies a VR or VRF.
----------------	------------------------

Default

Enabled.

Usage Guidelines

If ARP refresh is enabled, the switch resends ARP requests for the host at 3/4 of the configured ARP timer value.

For example: If the ARP timeout is set to 20 minutes, the switch attempts to resend an ARP request for the host when the host entry is at 15 minutes. If the host replies, the ARP entry is reset back to 0, and the timer starts again.

If you do not specify a VR or VRF, the command applies to the current VR context.

Example

The following command enables IP ARP refresh:

```
enable iparp refresh
```

History

This command was first available in ExtremeXOS 10.1.

The vr option was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

enable ipforwarding

```
enable ipforwarding {ipv4 | broadcast} {vlan vlan_name}
```

Description

Enables IPv4 routing or IPv4 broadcast forwarding for one or all VLANs. If no argument is provided, enables IPv4 routing for all VLANs that have been configured with an IP address on the current VR or VRF.



Syntax Description

ipv4	Specifies IPv4 forwarding.
broadcast	Specifies broadcast IP forwarding.
<i>vlan_name</i>	Specifies a VLAN name.

Default

Disabled.

Usage Guidelines

IP forwarding must first be enabled before IP broadcast forwarding can be enabled. When new IP interfaces are added, IP forwarding (and IP broadcast forwarding) is disabled by default. Currently, Extreme Networks switches only have a single hardware control per VLAN for IP forwarding of IPv4 and IPv6 unicast packets. Therefore, enabling IPv4 forwarding on a VLAN also enables IPv6 hardware forwarding on that VLAN. Future switches may have independent controls per-VLAN for forwarding of IPv4 and IPv6 unicast packets.

The broadcast, ignore-broadcast, and fast-directbroadcast options each prompt with a warning message when executed while the IP forwarding on the corresponding VLAN is disabled. The hardware and software are NOT programmed until IP forwarding is enabled on the VLAN.

The fast-direct-broadcast and ignore-broadcast options cannot be enabled simultaneously. These are mutually exclusive.

The broadcast option can be enabled in conjunction with fast-direct-broadcast and ignore-broadcast.

Example

The following command enables forwarding of IP traffic for all VLANs in the current VR context with IP addresses:

```
enable ipforwarding
```

The following command enables forwarding of IP broadcast traffic for a VLAN named accounting:

```
enable ipforwarding broadcast vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

The `ipv4` keyword was added in ExtremeXOS 11.2.

The `ignore-broadcast` and the `fast-direct-broadcast` keywords were added in ExtremeXOS 12.0.



Platform Availability

This command is available on all platforms that use the Edge, Advanced Edge, or Core license. For information on the licenses available for each platform, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable ip-option record-route

enable ip-option record-route

Description

Enables processing of the record route IP option in the IPv4 packet header.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

IP option record-route (IP option 7) means that each router along the path should add its IP address into the options data.

Enabling means that the switch adds itself into the IP options header when the record route IP option is enabled in a packet that is transiting the switch.

Example

The following command enables processing of the record route IP option:

```
enable ip-option record-route
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.



enable ip-option record-timestamp

enable ip-option record-timestamp

Description

Enables processing of the record timestamp IP option in the IPv4 packet header.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Enables the switch to use the timestamp IP option (0x44). When the switch receives an IP packet with the timestamp option turned on, it inserts the timestamp into the IP options header before forwarding the packet to the destination.

Example

The following command enables processing of the record timestamp IP option:

```
enable ip-option record-timestamp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable ip-option strict-source-route

enable ip-option strict-source-route

Description

Enables processing of the strict source route IP option in the IPv4 packet header.



Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This enables the switch to forward IP packets that have the strict source route IP option (0x89) enabled.

Source routing is used when a sending host specifies the router interfaces that the packet must traverse on its way to its destination.

When strict source routing is used, it means that the packet must use the exact path of routers that lie in the designated router path.

With strict source routing enabled, the switch forwards IP packets with the strict source route option enabled, only if the switch's IP is in the designated list and as long as the next hop in the list is directly attached to one of the router's interfaces.

Example

The following command enables processing of the strict source route IP option:

```
enable ip-option strict-source-route
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable ip-option router-alert

```
enable ip-option router-alert
```

Description

Enables processing of the router alert IP option in IPv4 packet headers.



Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

None.

Example

The following command enables processing of the router alert IP option:

```
enable ip-option router-alert
```

History

This command was first available in EXOS 10.1.

Platform Availability

This command is available on all platforms.

enable iproute bfd

```
enable iproute bfd {gateway} ip_addr {vr vrname}
```

Description

Enables the BFD client to provide services for IPv4 static routes.

Syntax Description

<i>ip_addr</i>	Specifies the IPv4 address of a neighbor to which BFD services are to be provided.
<i>vrname</i>	Specifies the VR or VRF name for which BFD services are being enabled.

Default

Disabled.



Usage Guidelines

To enable BFD services to an IPv4 neighbor, you must do the following:

- Execute this command on the switches at both ends of the link.
- Enable BFD for specific IPv4 static routes with the `configure iproute add (IPv4)` command.

Once a BFD session is established between two neighbors, BFD notifies the Route Manager process of the BFD session status and any changes. If other BFD clients (such as the MPLS BFD client) are configured between the same neighbors, the clients share a single session between the neighbors.

Example

The following example enables BFD client protection for communications with neighbor 10.10.10.1:

```
enable iproute bfd 10.10.10.1
```

History

This command was first available in ExtremeXOS 12.5.3.

Platform Availability

This command is available on all platforms.

enable iproute compression

```
enable iproute compression {vr vrname}
```

Description

Enables IPv4 route compression.

Syntax Description

<i>vrname</i>	VR or VRF name for which the IP route compression is being enabled.
---------------	---

Default

Disabled.

Usage Guidelines

Enables IPv4 route compression for the specified VR or VRF. If the VR name is not specified, route compression is enabled for the VR context from which the CLI command is issued.



The command applies a compression algorithm on each of the IP prefixes in the routing table. Essentially, routes with longer network masks might not be necessary if they are a subset of other routes with shorter network masks using the same gateway(s). When IP route compression is enabled, these unnecessary routes are not provided to the Forwarding Information Base (FIB).

Example

The following example enables IP route compression:

```
enable iproute compression
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

enable iproute sharing

```
enable iproute {ipv4} sharing {{{vr} vrname} | { {vr} all}}
```

Description

Enables load sharing if multiple routes to the same destination are available. When multiple routes to the same destination are available, load sharing can be enabled to distribute the traffic to multiple destination gateways. Only paths with the same lowest cost are shared.

Syntax Description

vrname	VR or VRF name for which IP route sharing is being enabled.
---------------	---

Default

Disabled.

Usage Guidelines

IP route sharing allows multiple equal-cost routes to be used concurrently. IP route sharing can be used with static routes or with OSPF, BGP, or IS-IS routes. In OSPF, BGP, and IS-IS, this capability is referred to as equal cost multipath (ECMP) routing.

Configure static routes and OSPF, BGP, or IS-IS as you would normally. The ExtremeXOS software supports route sharing across up to 32 ECMP static routes or up to 8 ECMP routes for OSPF, BGP, or IS-IS. However, on the BlackDiamond 8800 family, SummitStack, and Summit family switches, by default,



up to 4 routes are supported. To support 2, 4, 8, 16 or 32 routes on these switches, use the following command:

```
configure iproute sharing max-gateways max_gateways
```

If a VR is not specified, this command enables IP route sharing in the current VR context.

Example

The following command enables load sharing for multiple routes:

```
enable iproute sharing
```

History

This command was first available in ExtremeXOS 11.1.

The vr option was added in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms, except Summit X440. If present in a SummitStack, any Summit X440 nodes ingressing IPv4 unicast packets will still perform unipath IPv4 forwarding in hardware.

enable irdp

```
enable irdp {vlan name}
```

Description

Enables the generation of ICMP router advertisement messages on one or all VLANs.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

Disabled.

Usage Guidelines

ICMP Router Discovery Protocol allows client machines to determine what default gateway address to use. The switch sends out IP packets at the specified intervals identifying itself as a default router. IRDP enabled client machines use this information to determine which gateway address to use for routing data packets to other networks.



If no optional argument is specified, all the IP interfaces are affected.

Example

The following command enables IRDP on VLAN accounting:

```
enable irdp vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms that use the Edge, Advanced Edge, or Core license. For information on the licenses available for each platform, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable subvlan-proxy-arp vlan

```
enable subvlan-proxy-arp vlan [vlan-name | all]
```

Description

Enables the automatic entry of subVLAN information in the proxy ARP table.

Syntax Description

<i> vlan-name </i>	Specifies a superVLAN name.
all	Specifies all VLANs.

Default

Enabled.

Usage Guidelines

To facilitate communication between subVLANs, by default, an entry is made in the IP ARP table of the superVLAN that performs a proxy ARP function. This allows clients on one subVLAN to communicate with clients on another subVLAN. In certain circumstances, intra-subVLAN communication may not be desired for isolation reasons.



Note

The isolation option works for normal, dynamic, ARP-based client communication.



Example

The following command enables the automatic entry of subVLAN information in the proxy ARP table of the superVLAN vsuper:

```
enable subvlan-proxy-arp vlan vsuper
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

enable udp-echo-server

```
enable udp-echo-server {vr vrid}{udp-port port}
```

Description

Enables UDP echo server support.

Syntax Description

<i>vrid</i>	Specifies the VR or VRF.
<i>port</i>	Specifies the UDP port.

Default

Disabled.

Usage Guidelines

UDP Echo packets are used to measure the transit time for data between the transmitting and receiving ends.

Example

The following command enables UDP echo server support:

```
enable udp-echo-server
```



History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

rtlookup

```
rtlookup [ipaddress | ipv6address] { unicast | multicast | vr vr_name}
```

Description

Looks up and displays routes to the specified IP address.

Syntax Description

<i>ipaddress</i>	Specifies an IPv4 address.
<i>ipv6address</i>	Specifies an IPv6 address.
unicast	Displays the routes from the unicast routing table in the current router context.
multicast	Selects the multicast routing table for the search.
<i>vr_name</i>	Displays the available routes in the specified router context.

Default

N/A.

Usage Guidelines

When IP Route sharing is enabled, the `rtlookup` command displays all ECMP routes for the specified IP address.

When IP route sharing is disabled and there are multiple ECMP routes for the specified IP address, the `rtlookup` command displays only one route, which is the route with lowest value gateway IP address.

Example

The following command looks up IP address 10.0.0.0 in the VR-Mgmt router and displays the available routes:

```
BD-12804.4 # rtlookup 66.6.6.6
Ori Destination      Gateway      Mtr  Flags      VLAN      Duration
#s  66.6.6.6/32      80.1.10.58  1    UG---S-um--f v8      0d:0h:
18m:58s
```



```

Origin(Ori): (b) BlackHole, (be) EBGp, (bg) BGP, (bi) IBGP, (bo) BOOTP
(ct) CBT, (d) Direct, (df) DownIF, (dv) DVMRP, (e1) ISISL1Ext
(e2) ISISL2Ext, (h) Hardcoded, (i) ICMP, (i1) ISISL1 (i2) ISISL2
(is) ISIS, (mb) MBGP, (mbe) MBGPEExt, (mbi) MBGPInter, (mp) MPLS Lsp
(mo) MOSPF (o) OSPF, (o1) OSPFExt1, (o2) OSPFExt2
(oa) OSPFIntra, (oe) OSPFAsExt, (or) OSPFInter, (pd) PIM-DM, (ps) PIM-SM
(r) RIP, (ra) RtAdvrt, (s) Static, (sv) SLB_VIP, (un) UnKnown
(*) Preferred unicast route (@) Preferred multicast route
(#) Preferred unicast and multicast route
Flags: (B) BlackHole, (D) Dynamic, (G) Gateway, (H) Host Route
(L) Matching LDP LSP, (l) Calculated LDP LSP, (m) Multicast
(P) LPM-routing, (R) Modified, (S) Static, (s) Static LSP
(T) Matching RSVP-TE LSP, (t) Calculated RSVP-TE LSP, (u) Unicast, (U) Up
(f) Provided to FIB (c) Compressed Route

```

History

This command was first available in ExtremeXOS 10.1.

The xhostname option was removed in ExtremeXOS 11.0.

Support for IPv6 was added in ExtremeXOS 11.2.

The c flag was added in ExtremeXOS 12.0.

The unicast and multicast options were added in ExtremeXOS 12.1.

The f flag was added in ExtremeXOS 12.2.2.

Platform Availability

This command is available on all platforms.

run ip dad

```
run ip dad [{vlan} vlan_name | {{vr} vr_name} ip_address]
```

Description

Runs the DAD check on the specified IP interface for which the DAD feature is enabled.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN on which to run the test.
<i>vr_name</i>	Specifies a VR on which to run the test.
<i>ip_address</i>	Specifies an IP address for which to run the test.



Default

If you do not specify a VR or VRF, the command applies to the current VR context.

Usage Guidelines

To run the check, you must specify a VLAN name or a specific IP address. To support the check, the DAD feature must be enabled on the parent VR for the specified VLAN or IP interface.

This command is ignored for the following conditions:

- The specified IP address is in tentative state
- DAD is configured to be off
- The host VLAN is disabled
- Loopback mode is enabled on the host VLAN
- DAD is already running due to interface initialization or a previous issue of this command
- The host VLAN belongs to virtual router VR-Mgmt

If a duplicate address is detected during the check, the event is logged and the address remains valid if it was already valid. A valid address is an address that previously passed a DAD check (no duplicate address detected) and displays the U flag (Interface up) when the `show ip dad` command is entered. In this situation, the `show ip dad` command displays the D (duplicate address detected), E (Interface enabled), and U flags.

If the address was not already valid, the event is logged, the duplicate address transitions to duplicate address detected state, and the duplicate address displays the D flag when the `show ip dad` command is entered. If no duplicate IP address is detected, the specified IP interface transitions to or remains in the Interface up (U flag) state.

Example

The following command runs the DAD check on the IP interfaces in VLAN vlan1:

```
run ip dad vlan1
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on BlackDiamond X8, BlackDiamond 8800 series switches and Summit family switches.

show bootprelay

```
show bootprelay
```



Description

Displays the DHCP/BOOTP relay statistics and the configuration for the VRs.

Syntax Description

This command has no arguments or variables.

Default

None.

Usage Guidelines

The fields displayed in the DHCP Information Option 82 section depend on the configuration defined by the `configure bootprelay dhcp-agent information policy [drop | keep | replace]` command. If the policy configured is keep, the Requests unmodified counter appears. If the policy configured is replace, the Requests replaced counter appears. And if the drop policy is configured, the Requests dropped counter appears.

The Opt82 added to Requests counter indicates the number of DHCP requests to which the bootprelay agent (the switch) has added its own option 82 information.

Example

The following example displays the DHCP/BOOTP relay statistics for existing VRs:

```
Switch.1 # show bootprelay
Bootprelay : Enabled on virtual router "VR-Default"
DHCP Relay Agent Information Option : Enabled on virtual router "VR-Default"
DHCP Relay Agent Information Check : Enabled on virtual router "VR-Default"
DHCP Relay Agent Information Policy : Replace
DHCP Relay Agent Information Remote-ID : "default"
Bootprelay servers for virtual router "VR-Default":
Destination: 10.127.8.1
DHCP/BOOTP relay statistics for virtual router "VR-Default"
Received from client =          2 Received from server =          2
Requests relayed    =          2 Responses relayed    =          2
DHCP Discover       =          1 DHCP Offer           =          1
DHCP Request        =          1 DHCP Ack          =          1
DHCP Decline        =          0 DHCP NACK         =          0
DHCP Release        =          0
DHCP Inform         =          0
DHCP Information Option 82 packets statistics for virtual router "VR-Default"
Received from client =          0 Received from server =          2
Requests replaced    =          0 Responses dropped    =          0
Opt82 added to Requests =          2
Note: Default Remote-ID : System MAC Address
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show bootprelay configuration

```
show bootprelay configuration [{vlan} vlan_name | {vr} vr_name]
```

Description

Displays the enabled/disabled configuration of BOOTP relay on one or all VLANs for the specified VR.

Syntax Description

<i>vlan_name</i>	Specifies a single VLAN for which to display BOOTP relay configuration information.
<i>vr_name</i>	Specifies a single VR for which to display BOOTP relay configuration information.

Default

None.

Usage Guidelines

If a VR is not specified, this command displays the specified VLANs for the current VR context.

Example

The following example displays the BOOTP relay configuration for all VLANs on the VR-Default virtual router:

```
Switch.88 # show bootprelay configuration vr "VR-Default"
BOOTP Relay : Enabled on virtual router "VR-Default"
VLAN                               BOOTP Relay
-----
Default                             Disabled
client1                             Enabled
serv                                 Enabled
```

The following example displays the BOOTP relay configuration for all VLANs in the current VR context:

```
Switch.95 # show bootprelay configuration
```




```

BOOTP Relay : Enabled on virtual router "VR-Default"
VLAN
-----
Default
client1
serv
BOOTP Relay
-----
Disabled
Disabled
Disabled

```

The following example displays the BOOTP relay configuration for VLAN client1:

```

Switch.87 # show bootprelay configuration vlan "client1"
BOOTP Relay : Enabled on virtual router "VR-Default"
VLAN
-----
client1
BOOTP Relay
-----
Disabled

```

History

This command was first available in ExtremeXOS 12.4.2.

Platform Availability

This command is available on all platforms.

show bootprelay dhcp-agent information circuit-id port-information

```
show bootprelay dhcp-agent information circuit-id port-information ports all
```

Description

Displays the circuit ID sub-option that identifies the port for an incoming DHCP request.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.



Example

The following command displays the circuit ID port_info value for all ports:

```
Switch.12 # show bootprelay dhcp-agent information circuit-id port-
information ports all
Port                Circuit-ID Port information string
-----
1                   1001
2                   1002
3                   extreme1
4                   1004
5                   1005
6                   1006
7                   1007
8                   1008
9                   1009
10                  1010
:
:
11                  1011
12                  1012
:
:
48                  1048
49                  1049
50                  1050
Note: The full Circuit ID string has the form '<Vlan Info>-<Port Info>'
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

show bootprelay dhcp-agent information circuit-id vlan- information

```
show bootprelay dhcp-agent information circuit-id vlan-information
```

Description

Displays the circuit ID sub-options that identify the VLANs on the switch.

Syntax Description

This command has no arguments or variables.



Default

N/A.

Usage Guidelines

None.

Example

The following command displays the circuit ID `vlan_info` for all VLANs:

```
XSummit # show bootprelay dhcp-agent information circuit-id vlan-information
Vlan          Circuit-ID  vlan information string
----          -
Default      1
Mgmt         4095
v1           4094
v2           extreme123
Note: The full Circuit ID string has the form '<Vlan Info>-<Port Info>'
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

show bootprelay configuration ipv4

Short reference description.

show bootprelay

Description

Displays various bootprelay configuration details.

Syntax Description

ipv4	Specifies unconfiguring the DHCPv4 BOOTP Relay service (default).
ipv6	Specifies unconfiguring the DHCPv6 BOOTP Relay service.
vlan	Unconfigures BOOTP relay for a specified VLAN.
<i>vlan_name</i>	Specifies the VLAN name.
include-secondary	Removes the include-secondary configuration for the specified VLAN.



Default

Usage Guidelines

Use this command to display ...

Example

The following command :

```
X480-48t.4 # show bootprelay configuration ipv4
DHCPv4 BOOTP Relay : Enabled on virtual router "VR-Default"
  Include Secondary : Enabled (parallel)
  BOOTP Relay Servers :
    DHCP Relay Agent Information Option: Enabled
    DHCP Relay Agent Information Check : Disabled
    DHCP Relay Agent Information Policy: Replace
VLAN                                DHCPv4 BOOTP Relay
-----
VLAN "Default":
  BOOTP Relay                        : Enabled
  Include Secondary                  : Enabled (sequential)
  BOOTP Relay Servers                : 192.168.1.1
  DHCP Relay Agent Information Option: Enabled
  DHCP Relay Agent Information Check : Disabled
  DHCP Relay Agent Information Policy: Replace
VLAN "v1":
  BOOTP Relay                        : Enabled
  Include Secondary                  : Enabled (sequential)
  BOOTP Relay Servers                : 10.1.1.1
  DHCP Relay Agent Information Option: Enabled
  DHCP Relay Agent Information Check : Disabled
  DHCP Relay Agent Information Policy: Replace
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

show bootprelay ipv6

Short reference description.

show bootprelay



Description

Displays various bootprelay configuration details.

Syntax Description

ipv4	Specifies unconfiguring the DHCPv4 BOOTP Relay service (default).
ipv6	Specifies unconfiguring the DHCPv6 BOOTP Relay service.
vlan	Unconfigures BOOTP relay for a specified VLAN.
<i>vlan_name</i>	Specifies the VLAN name.
include-secondary	Removes the include-secondary configuration for the specified VLAN.

Default

Usage Guidelines

Use this command to display ...

Example

The following command :

```
X480-48t.3 # show bootprelay ipv6
BOOTP Relay: DHCPv6 BOOTP Relay enabled on virtual router "VR-Default"
Include Secondary : Enabled (parallel)
VLAN "Default"      :
  BOOTP Relay       : Enabled
  Include Secondary : Enabled (sequential)
  Interface ID      : 1(Default)
VLAN "v1"           :
  BOOTP Relay       : Enabled
  Include Secondary : Disabled
  Interface ID      : MyVLAN
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.



show ip dad

```
show ip dad [{vr} vr_name {ip_address} | vr all | {vlan} vlan_name] {tentative
| valid | duplicate} | [{vr} vr_name] ip_address}]
```

Description

Displays the configuration and run time status for the DAD feature on the specified IP interface.

Syntax Description

<i>vr_name</i>	Specifies a VR for which to display the DAD information.
<i>ip_address</i>	Specifies an IP address for which to display the DAD information.
<i>vlan_name</i>	Specifies a VLAN for which to display the DAD information.
tentative	Displays information for IP interfaces for which the status is up and the DAD check is incomplete.
valid	Displays information for IP interfaces that the DAD check has declared valid.
duplicate	Displays information for IP interfaces for which a duplicate IP address was detected.

Default

If you do not specify a VR or VRF, the command applies to the current VR context.

Usage Guidelines

The `vr all` option displays DAD information for all IP interfaces on the switch.

The DAD check takes very little time, so you might not see the T flag (tentative address) during normal operation.

A valid IP address displays the E (interface enabled) and U (interface up) flags, and switch processes can use the IP address. If a duplicate IP address is detected after an IP address is declared valid, the D flag (duplicate address detected) also appears.

An invalid IP address does not show the U flag and might not show the E flag. If a duplicate address was detected for the invalid address, the D flag appears.

The L flag indicates an IP address for a loopback VLAN. The DAD check does not run on loopback VLANs and always marks a loopback VLAN as valid.

Example

The following command displays the DAD feature status for all IP interfaces in the current VR context:

```
IPv4 Duplicate Address Detection
DAD Status : On
```



```

Max Solicitation Attempts : 1
Virtual Router      Interface      Flags      IP Address
Conflict MAC       Failures
-----
--
VR-Default          loop158      -ELU      81.70.100.35
00:00:00:00:00:00      0
VR-Default          exL-v188     -E-U      13.224.90.90
00:04:96:12:ae:60      1
Flags : (D) Duplicate address detected, (T) Tentative address,
(E) Interface enabled, (L) Loopback enabled, (U) Interface up

```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on BlackDiamond X8, BlackDiamond 8800 series switches and Summit family switches.

show iparp

```
show iparp {ip_addr | mac | vlan vlan_name | permanent} {vr vr_name}
```

Description

Displays the IP Address Resolution Protocol (ARP) table. You can filter the display by IP address, MAC address, VLAN, or permanent entries.

Syntax Description

<i>ip_addr</i>	Specifies an IP address.
<i>mac</i>	Specifies a MAC address.
<i>vlan_name</i>	Specifies a VLAN name.
permanent	Specifies permanent entries.
<i>vr_name</i>	Specifies a VR or VRF.

Default

Show all entries, except for proxy entries.

Usage Guidelines

Displays the IP ARP table, including:



- IP address
- MAC address
- Aging timer value
- VLAN name, VLAN ID and port number
- Flags

If you do not specify a VR or VRF, the command applies to the current VR context.

The show output displays the following information:

ARP address check	Whether IP ARP checking is enabled or disabled. IP ARP checking verifies if the ARP request's source address is in the receiving interface's subnet.
ARP refresh	Whether ARP refresh is enabled or disabled. ARP refresh is performed when an ARP entry's age is three-fourths of the timeout value.
Distributed mode	Displays the configured and current states for the distributed IP ARP feature, which applies only to BlackDiamond 8800 series switches. The Configured field shows the feature state that will apply the next time the switch is restarted. The Current state shows the current state that was applied the last time the switch was restarted.
Dup IP Addr	IP addresses that have been used by other hosts on the network.
Dynamic entries	The number of dynamic (learned ARP) entries in the table.
Failed requests	The number of failed ARP requests sent (by this VR or VRF).
In Request	The number of ARP request packets received (by this VR or VRF).
In Response	The number of ARP reply packets received (by this VR or VRF).
Max ARP entries	Maximum ARP table size for the VR or VRF (each VR has its own ARP table).
Max ARP pending entries	Maximum number of incomplete (pending) ARP entries allowed in the table.
Out Request	The number of ARP request packets sent (by this VR or VRF).
Out Response	The number of ARP reply packets sent (by this VR or VRF).
Pending entries	The number of sent ARP requests that have not yet received a response.
Proxy Answered	The number of ARP requests answered by the ARP proxy.
RX Error	The number of incorrect ARP request and reply packets received. The malformed packets include the following errors: incorrect ARP op code, hardware address type is not Ethernet, the protocol address is not IP, and similar errors.
Static entries	The number of configured (static ARP) entries in the table.
Rejected Count	The number of rejected ARP request packets.
Rejected I/F	The VLAN on which the last rejected ARP request packet arrived.
Rejected IP	The source address for the last rejected ARP request. An example reason for an ARP request packet to be rejected is if the source address of the packet is not in the subnet.
Rejected Port	The port on which the last rejected ARP request packet arrived.
Timeout	Timeout value for a dynamic (learned) ARP entry.



Example

The following command displays the IP ARP table for the current VR or VRF context:

```
show iparp
```

The following is sample output for the command:

VR VID Port	Destination	Mac	Age	Static	VLAN
VR-Default 4092 1	10.10.10.6	00:04:96:1f:a5:71	8	NO	bluered
VR-Default 4095	10.128.32.1	00:01:30:ba:6a:a0	0	NO	Default
VR-Default 4095	10.128.32.2	00:01:03:1c:ae:b0	5	NO	Default
VR-Default 4095	10.128.32.4	00:d0:59:17:74:83	3	NO	Default
VR-Default 4095	10.128.32.5	00:02:a5:c2:5c:dd	0	NO	Default
VR-Default 4095	10.128.32.6	00:12:3f:1c:f8:fb	5	NO	Default
VR-Default 4095	10.128.32.7	00:11:11:80:9c:b9	7	NO	Default
VR-Default 4095	10.128.32.8	00:11:43:53:8e:f1	0	NO	Default
VR-Default 4095	10.128.32.9	00:02:a5:bf:ac:70	7	NO	Default
VR-Default 4095	10.128.32.10	00:11:43:44:18:68	10	NO	Default
VR-Default 4095	10.128.32.11	00:12:3f:1c:e9:f2	0	NO	Default
VR-Default 4095	10.128.32.12	00:02:a5:bf:af:79	8	NO	Default
VR-Default 4095	10.128.32.13	00:11:43:40:89:91	0	NO	Default
VR-Default 4095	10.128.32.16	00:0f:1f:c9:2d:80	2	NO	Default
VR-Default 4095	10.128.32.17	00:06:5b:b1:6a:91	1	NO	Default
VR-Default 4095	10.128.32.19	00:11:43:3a:96:1d	10	NO	Default
VR-Default 4095	10.128.32.20	00:08:02:d5:c5:b7	6	NO	Default
VR-Default 4095	10.128.32.24	00:12:3f:0a:44:92	14	NO	Default
VR-Default 4095	10.128.32.26	00:50:04:ad:36:5e	6	NO	Default
VR-Default 4095	10.128.32.30	00:b0:d0:23:f2:9a	11	NO	Default
VR-Default 4095	10.128.32.54	00:b0:d0:59:e4:e2	6	NO	Default
VR-Default 4095	10.128.32.55	00:a0:c9:0c:41:de	3	NO	Default
VR-Default 4095	10.128.32.59	00:b0:d0:7c:d6:07	14	NO	Default



```

VR-Default      10.128.32.99      00:04:96:05:00:03  13      NO  Default
4095
VR-Default      10.128.32.101     00:04:96:1f:a8:48   0       NO  Default
4095
VR-Default      10.128.32.104     00:30:48:41:ed:45   0       NO  Default
4095
VR-Default      10.128.32.105     00:30:48:41:ed:97   0       NO  Default
4095
VR-Default      10.128.32.106     00:01:30:23:c1:00   0       NO  Default
4095
VR-Default      10.128.32.108     00:04:96:1f:a5:71   0       NO  Default
4095
VR-Default      10.128.32.116     00:04:96:1f:a4:0e   0       NO  Default
4095
Dynamic Entries :           1           Static
Entries         :           0
Pending Entries :           0
In Request      :           111          In
Response        :           :           3
Out Request     :           110          Out Response      :
111
Failed Requests :           0
Proxy Answered  :           0
Rx Error        :           0           Dup IP
Addr           :           0.0.0.0
Rejected Count  :                               Rejected IP      :
Rejected Port  :                               Rejected I/F     :
Max ARP entries :           4096          Max ARP pending entries :
256
ARP address check: Enabled          ARP refresh      :
Enabled
Timeout         :   Configured: 10 minutes   Current: 30 minutes
Distributed mode :   Configured: On           Current: Off
Max ARP entries :   Configured: 259000 (Distributed mode was turned "on"
after boot)
ARP Sender-Mac Learning : Disabled

```

History

This command was first available in ExtremeXOS 10.1.

The `vr` option was added in ExtremeXOS 11.0.

For BlackDiamond 8800 series switches only, additional display information was added to show the configured and current state of the distributed IP ARP feature in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.

show iparp distributed-mode statistics

```
show iparp distributed-mode statistics { slot [ slot | all ] }
```



Description

Displays distributed IP ARP statistics when this feature is activated on BlackDiamond X8, BlackDiamond 8000 c-, xl-, and xm-series modules.

Syntax Description

<code>slot</code>	Displays distributed IP ARP statistics for the specified slot.
<code>all</code>	Displays distributed IP ARP statistics for all slots.

Default

Displays distributed IP ARP statistics for all slots.

Usage Guidelines

When distributed IP ARP mode is active on a switch, this command allows you to view how many IP ARP entries are assigned to the ports on supported I/O modules. To turn on distributed IP ARP mode, use the `configure iparp distributed-mode` command.

You can use the information from this show command to determine if the total entries in use is close to or equal to the total entries supported by the I/O module hardware. When the total entries in use reaches the hardware limit, new entries must replace existing entries, and IP forwarding temporarily takes place using the slow path through software, rather than the fast path through hardware.

If the total entries in use gets close to the maximum entries supported by the hardware, consider moving IP hosts to a port or port group that is not so close to the hardware limit.

Example

The following command displays distributed IP ARP statistics for all slots:

```
Switch # show iparp distributed-mode statistics
System- || # ARPs      GW's,
wide    || Native     IPv6,
Slot Type
Port List
-----
1
2      8900-G48T-x1      162592
2:1-2:48                16273      0      16273!    16273
3      8900-G48T-x1(P)   162592
3:1-3:48                16273      0      16273!    16273
4      8900-G48X-x1        162592
4:1-4:48                16273      0      16273!    16273
5      G8Xc                 20476
5:1-5:8                  1          0          1          8081
6      10G1Xc               20476
6:1                      1          0          1          8081
7      8900-G48X-x1        162592
```



```

7:1-7:48                16273      0  16273!  16273
8      8900-10G8X-x1    162592
8:1-8:4                16273      0  16273!  16273
8:5-8:8                16213      0  16213   16273
9      8900-10G8X-x1    162592
9:1-9:4                16273      0  16273!  16273
9:5-9:8                16273      0  16273!  16273
10     8900-10G8X-x1    162592
10:1-10:4              16230      0  16230   16273
10:5-10:8              16236      0  16236   16273

```

Flags: (!) Indicates all hardware entries in use.

BD-X8.12 # show iparp distributed-mode statistics

```

System- || # ARPs      GW's,
wide    || Attached  IPv6,
Slot Type
Port List          In HW      ARPs    || To This  PBR,      Total
                || Port List MPLS    In Use   HW Max
-----
-----
1      BDXA-10G48X      14032
1:1-1:12, 1:25-1:36      1      1      2      15965
1:13-1:24, 1:37-1:48    4654    1      4655   15965
2
                n/a
3      BDXA-40G24X      14032
3:49, 3:53, 3:57, 3:61, 3:65, 3:69      0      1      1      15965
3:73, 3:77, 3:81, 3:85, 3:89, 3:93      0      1      1      15965
3:1, 3:5-3:9, 3:13, 3:17, 3:21      9376    1      9377   15965
3:25, 3:29, 3:33, 3:37, 3:41, 3:45      0      1      1      15965
4
                n/a
5
                n/a
6
                n/a
7
                n/a
8
                n/a

```

Flags: (!) Indicates all hardware entries in use.

The following command displays distributed IP ARP statistics for slot 2:

```

* Switch # show iparp distributed-mode statistics slot 2
System- || # ARPs      GW's,
wide    || Native     IPv6,
Slot Type
Port List          In HW      ARPs    || To This  PBR,      Total
                || Port List MPLS    In Use   HW Max
-----
-----
2      8900-G48T-x1    162592
2:1-2:48          16273      0  16273!  16273

```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available only on BlackDiamond X8, BlackDiamond 8000 c-, xl-, and xm-series modules.



show iparp proxy

```
show iparp proxy {[ipNetmask | ip_addr mask]} {vr vr_name}
```

Description

Displays the proxy ARP table.

Syntax Description

<i>ipNetmask</i>	Specifies an IP address/mask length.
<i>ip_addr</i>	Specifies an IP address.
<i>mask</i>	Specifies a subnet mask.
<i>vr_name</i>	Specifies a VR or VRF.

Default

N/A.

Usage Guidelines

If no argument is specified, then all proxy ARP entries are displayed.

If you do not specify a VR or VRF, the command applies to the current VR context.

Example

The following command displays the proxy ARP table:

```
show iparp proxy 10.1.1.5/24
```

History

This command was first available in ExtremeXOS 10.1.

The vr option was added in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show iparp security

```
show iparp security [{vlan} vlan_name]
```



Description

Displays the IP ARP security violation information for one or all VLANs.

Syntax Description

<i>vlan_name</i>	Specifies a single VLAN for which to display security violation information.
------------------	--

Default

Shows security violation information for all VLANs except Mgmt.

Usage Guidelines

None.

Example

The following command displays IP ARP security violation information for all VLANs:

```
Switch.4 # show iparp security
Most Recent Violation
=====
Vlan          Security  Violations  Type   IP address
MAC          Port
=====
Default      ----
test        ----
Security Setting: (G) Gratuitous ARP Protection
Violation Type  : (g) Gratuitous ARP Violation
```

The following command displays IP ARP security violation information for VLAN Default:

```
Switch.5 # show iparp security "Default"
Most Recent Violation
=====
Vlan          Security  Violations  Type   IP address
MAC          Port
=====
Default      ----
Security Setting: (G) Gratuitous ARP Protection
Violation Type  : (g) Gratuitous ARP Violation
```

History

This command was first available in ExtremeXOS 10.1.

The vr option was added in ExtremeXOS 11.0.



Platform Availability

This command is available on all platforms.

show iparp stats

```
show iparp stats [[ vr_name | vr {all | vr_name} ] {no-refresh} | {vr} summary]
show iparp stats [vlan {all {vr vr_name}} | {vlan} vlan_name] {no-refresh}
show iparp stats ports {all | port_list} {no-refresh}
```

Description

Displays the IP ARP statistics for one or more VRs, VLANs, or ports.

Syntax Description

<i>vr_name</i>	Specifies a VR or VRF for which to display statistics.
<i>vlan_name</i>	Specifies a VLAN for which to display statistics.
<i>port_list</i>	Specifies a list of ports for which to display statistics.
no-refresh	Requests a static display for statistics.

Default

Shows all VLAN ARP statistics in a dynamic display.

Usage Guidelines

VLAN statistics and totals are displayed for a single VR. When you display IPARP statistics for one or all VLANs, the display includes the specified VLANs for the specified VR. If you do not specify a VR for a VLAN report, the display includes the specified VLANs for the current VR context.

Counters displayed under Pending, Failed, ARP Unneeded are per VR and 'show iparp stats' includes VR-Management as well which typically skews the ARP statistics considerably. The Total entries counter reflects the total number of entries that are currently allocated and not freed. Hence they also include Failed entries as well as ARP unneeded entries. Dynamic Entries counter indicates the reachable entries. Periodically as part of cleanup, failed entries will go down and hence the total entries goes down. In certain scenarios they may help detect a problem (e.g. memory leak or an attack). Since the counters together accurately convey the state of the system, we choose to display these entries. "show iparp stats" shows totals across all VRs* including VR-Mgmt.

Example

The following command displays ARP table statistics for all VRs and VRFs:

```
Switch1 # show iparp stats vr all
IP ARP VR Statistics                               Wed Apr 07 15:30:49 2010
ARP Total      Dynamic      Static      Pending      Unneeded      Failed (Rejected)
```



```

=====
VR-Default
96          89          5          0          0          2          0
VR-Mgmt
4           2           2           0           0           0           0
VR-SV_PPPOE
287        286        1           0           0           0           2785
VR-NV_PPPOE
19         19         0           0           0           0           0
chicago2
50         44         5           0           1           0           0
Total for all VRs
456        440        13          0           1           2           2785
=====
U->page up  D->page down  ESC->exit

```

The following command displays ARP table statistics for all VLANs in the current VR context:

```

Switch.2 # show iparp stats vlan all
IP ARP VLAN Statistics                                     Wed Apr 07 15:30:49
2010
VLAN          ARP Total          Dynamic
Static
=====
=
VLAN_06-AAR          94          89
5
VLAN_07-AAR          122         121
1
VLAN_02-BOT          43          42
1
=====
=
Totals for VR U3c-South.          Total Entries  :
455
Dynamic :          440   Static   :          13   Pending  :          0
Failed  :           2   Unneeded :           0   (Rejected): 5639
Last Rejected ARP :
IP: 10.66.118.243   Port: 1:23   Vlan: VLAN_02-BOT
=====
=
U->page up  D->page down  ESC->exit

```

The following command displays ARP table statistics for ports 1:1 to 1:17:

```

Switch.3 # show iparp stats ports 1:1-1:17
IP ARP Port Statistics                                     Wed Apr 07 15:30:49
2010
Port          Link State          ARP Total          Dynamic
Static
=====
=
1:1          A          94          89
5
1:2          A          37          37

```




```

0
1:3          A          122          121
1
1:4          R          0           0
0
1:5          R          0           0
0
1:6          A          43          43
0
1:7          A          118         118
0
1:8          R          0           0
0
1:9          R          0           0
0
1:10         A          8           8
0
1:11         A          8           6
2
1:12         A          41          41
0
1:13         A          17          17
0
1:14         R          0           0
0
1:15         R          0           0
0
1:16         A          8           8
0
1:17         A          8           6
2

```

```

=====
=
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
U->page up  D->page down  ESC->exit

```

History

This command was first available in ExtremeXOS 12.4.2.

Platform Availability

This command is available on all platforms.

show ipconfig

```
show ipconfig {ipv4} {vlan vlan_name}
```

Description

Displays configuration information for one or more VLANs in the current VR context.



Syntax Description

<code>vlan_name</code>	Specifies a VLAN name.
------------------------	------------------------

Default

N/A.

Usage Guidelines

If no VLAN information is specified, then global IP configuration is displayed. Otherwise, specific VLAN information is displayed.

Example

```
* switch # sho ipconfig
  Use Redirects : Disabled
  IpOption LSRR : Enabled
  IpOption SSRR : Enabled
  IpOption RR : Enabled
  IpOption TS : Enabled
  IpOption RA : Enabled
  Route Sharing : Disabled
  Originated Packets : Don't require ipforwarding
  IP Fwding into LSP : Disabled
  Max Shared Gateways : Current: 4  Configured: 4

  IRDP:
    Advertisement Address: 255.255.255.255
    Maximum Interval: 600
    Minimum Interval: 450      Lifetime: 1800      Preference: 0
  Interface  IP Address      Flags              nSIA
  data2      200.0.0.1       /24 EUf---MPuRX-----  0
  inet       1.1.1.2         /24 EUf---MPuRX-----  0
  mytun      2.0.0.2         /24 EU---MPuRX-----   0
```

History

This command was first available in ExtremeXOS 10.1.

The `ipv4` keyword was added in ExtremeXOS 11.2.

This command changed to display information for the current VR context in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms.



show iproute

```
show iproute {ipv4} {priority | vlan vlan_name | permanent | ip_address netmask |
summary} {multicast | unicast} {vr vrname}}
```

Description

Displays the contents of the IP routing table or the route origin priority.

Syntax Description

priority	Displays the priority values for each route origin type.
<i>vlan_name</i>	Specifies a VLAN name.
permanent	Specifies permanent routing.
<i>ip_address</i>	Specifies an IP address.
<i>netmask</i>	Specifies a subnet mask.
summary	Displays summary information.
multicast	Displays information for IPv4 multicast routes only.
unicast	Displays information for IPv4 unicast routes only.
<i>vrname</i>	Specifies a VR or VRF.

Default

N/A.

Usage Guidelines

A *c* flag in the Flags column indicates a compressed route resulting from enabling compression using the `enable iproute compression` command. The total number of compressed routes is also shown.

All routes that are provided to the FIB display the *f* flag.

If you do not specify a VR or VRF, the command applies to the current VR context.

Example

The following command example displays detailed information about all IP routing:

```
Switch.3 # show iproute
Ori Destination      Gateway      Mtr  Flags      VLAN      Duration
d   2.2.0.0/16        2.2.2.3     1    -----um--- v2        2d:10h:
17m:41s
#d  3.2.2.0/24        3.2.2.23    1    U-----um--f jim_igmp  1d:19h:
49m:49s
d   3.3.3.0/24        3.3.3.1     1    -----um--- v3        2d:10h:
17m:49s
```



```

d 4.4.4.0/24          4.4.4.2          1  -----um--- v4          2d:10h:
17m:5s
Origin(Ori): (b) BlackHole, (be) EBGp, (bg) BGP, (bi) IBGP, (bo) BOOTP
(ct) CBT, (d) Direct, (df) DownIF, (dv) DVMRP, (e1) ISISL1Ext
(e2) ISISL2Ext, (h) Hardcoded, (i) ICMP, (i1) ISISL1 (i2) ISISL2
(is) ISIS, (mb) MBGP, (mbe) MBGPExt, (mbi) MBGPInter, (mp) MPLS Lsp
(mo) MOSPF (o) OSPF, (o1) OSPFExt1, (o2) OSPFExt2
(oa) OSPFIntra, (oe) OSPFAsExt, (or) OSPFInter, (pd) PIM-DM, (ps) PIM-SM
(r) RIP, (ra) RtAdvrt, (s) Static, (sv) SLB_VIP, (un) UnKnown
(*) Preferred unicast route (@) Preferred multicast route
(#) Preferred unicast and multicast route
Flags: (B) BlackHole, (b) BFD protection requested, (c) Compressed, (D)
Dynamic
(f) Provided to FIB, (G) Gateway, (H) Host Route, (L) Matching LDP LSP
(l) Calculated LDP LSP, (3) L3VPN Route, (m) Multicast, (P) LPM-routing
(p) BFD protection active, (R) Modified, (S) Static, (s) Static LSP
(T) Matching RSVP-TE LSP, (t) Calculated RSVP-TE LSP, {u} Unicast, (U) Up
Mask distribution:
1 routes at length 16          3 routes at length 24
Route Origin distribution:
4 routes from Direct
Total number of routes = 4
Total number of compressed routes = 0

```

History

This command was first available in ExtremeXOS 10.1.

The `ipv4` keyword was added in ExtremeXOS 11.2.

The `c` flag was added in ExtremeXOS 12.0.

The `f` flag was added in ExtremeXOS 12.2.2.

Platform Availability

This command is available on all platforms.

show iproute mpls

```

show iproute mpls {lsp lsp_name | vlan vlan_name | permanent | ip_address netmask
| summary} {unicast} {vr vrname}

```

Description

Displays the MPLS contents of the IP routing table.

Syntax Description

<code>lsp_name</code>	Specifies an LSP name.
<code>vlan_name</code>	Specifies a VLAN name.



permanent	Specifies permanent routing.
<i>ip_address</i>	Specifies an IP address.
<i>netmask</i>	Specifies a subnet mask.
unicast	Displays unicast routes.
<i>vrname</i>	Specifies a VR or VRF.

Default

N/A.

Usage Guidelines

A *c* flag in the Flags column indicates a compressed route resulting from enabling compression using the `enable iproute compression` command. The total number of compressed routes is also shown.

All routes that are provided to the FIB display the *f* flag.

If you do not specify a VR or VRF, the command applies to the current VR context.

Example

The following command example displays detailed information about all IP routing:

```
Switch.3 # show iproute mpls
```

History

This command was first available in ExtremeXOS 12.2.2.

Platform Availability

This command is available on all platforms.

show iproute mpls origin

```
show iproute mpls origin [bgp | blackhole | bootp | direct | ebgp | ibgp | icmp |
isis | isis-level-1 | isis-level-1-external | isis-level-2 | isis-level-2-
external | mpls {signaling-protocol [ldp | rsvp-te | static]} | ospf | ospf-
extern1 | ospf-extern2 | ospf-inter | ospf-intra | rip | static ] {unicast} {vr
vrname}
```

Description

Displays the MPLS contents of the IP routing table for routes with the specified origin.



Syntax Description

bgp	Specifies BGP routes.
blackhole	Specifies blackhole routes.
bootp	Specifies BOOTP routes.
direct	Specifies direct routes.
ebgp	Specifies E-BGP routes.
ibgp	Specifies I-BGP routes.
icmp	Specifies ICMP routes.
isis	Specifies IS-IS routes.
isis-level-1	Specifies ISIS Level 1 routing.
isis-level-1-external	Specifies ISIS Level 1 External routing.
isis-level-2	Specifies ISIS Level 2 routing.
isis-level-2-external	Specifies ISIS Level 2 External routing.
mpls	Specifies MPLS routes. This option is available only on platforms that support the MPLS feature pack, which is described in the Feature Pack Features section of Feature License Requirements
signaling-protocol [ldp rsvp-te static]	Specifies an MPLS signaling protocol. This option is available only on platforms that support the MPLS feature pack, which is described in the Feature Pack Features section of Feature License Requirements
ospf	Specifies OSPF routes.
ospf-extern1	Specifies OSPF External 1 routing.
ospf-extern2	Specifies OSPF External 2 routing.
ospf-inter	Specifies OSPFInter routing.
ospf-intra	Specifies OSPFIntra routing.
rip	Specifies RIP routes.
static	Specifies static routes.
unicast	Displays unicast routes with the specified origin.
<i>vrname</i>	Specifies a VR or VRF.

Default

N/A.

Usage Guidelines

None.



Example

The following command displays all the MPLS routes that originate from BGP:

```
show iproute mpls origin bgp
```

History

This command was first available in ExtremeXOS 12.2.2.

Platform Availability

This command is available on all platforms that support MPLS.

show iproute origin

```
show iproute origin [bgp | blackhole | bootp | direct | ebgp | embgp | ibgp |
icmp | imbgp | isis | isis-level-1 | isis-level-1-external | isis-level-2 | isis-
level-2-external | mbgp | mpls {signaling-protocol [ldp | rsvp-te | static]} |
ospf | ospf-extern1 | ospf-extern2 | ospf-inter | ospf-intra | rip | static ]
{unicast} {vr vrname}
```

Description

Displays the contents of the IP routing table for routes with the specified origin.

Syntax Description

bgp	Specifies BGP routes.
blackhole	Specifies blackhole routes.
bootp	Specifies BOOTP routes.
direct	Specifies direct routes.
ebgp	Specifies E-BGP routes.
embgp	Specifies EMBGP routes.
ibgp	Specifies I-BGP routes.
icmp	Specifies ICMP routes.
imbgp	Specifies IMBGP routes.
isis	Specifies IS-IS routes.
isis-level-1	Specifies ISIS Level 1 routing.
isis-level-1-external	Specifies ISIS Level 1 External routing.
isis-level-2	Specifies ISIS Level 2 routing.
isis-level-2-external	Specifies ISIS Level 2 External routing.



mbgp	Specifies MBGP routes.
mpls	Specifies MPLS routes. This option is available only on platforms that support the MPLS feature pack, which is described in the Feature Pack Features section of Feature License Requirements .
signaling-protocol [ldp rsvp-te static]	Specifies an MPLS signaling protocol. This option is available only on platforms that support the MPLS feature pack, which is described in the Feature Pack Features section of Feature License Requirements .
ospf	Specifies OSPF routes.
ospf-extern1	Specifies OSPF External 1 routing.
ospf-extern2	Specifies OSPF External 2 routing.
ospf-inter	Specifies OSPF Inter routing.
ospf-intra	Specifies OSPF Intra routing.
rip	Specifies RIP routes.
static	Specifies static routes.
unicast	Displays unicast routes with the specified origin.
vrname	Specifies a VR or VRF.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays all the BGP routes:

```
show iproute origin bgp
```

History

This command was first available in ExtremeXOS 10.1.

The ipv4 keyword was added in ExtremeXOS 11.2.

The embgp, mbgp, and mpls options are added and the ipv4 and multicast options removed in ExtremeXOS 12.2.2.

Platform Availability

This command is available on all platforms.



show iproute reserved-entries

```
show iproute reserved-entries {slot slot_num}
```

Description

Displays the configured number of IPv4 and IPv6 routes reserved in the Longest Prefix Match (LPM) hardware table.

Syntax Description

<i>slot_num</i>	For BlackDiamond X8 and 8800 series switches and SummitStack only, this option displays the reservations for the specified slot.
-----------------	--

Default

N/A.

Usage Guidelines

The IPv4 Routes column in the command output shows whether IPv4 routes are stored in internal or external LPM tables.

The "(or IPv6)" column in the command output shows whether IPv6 routes are stored in internal ("int.") or external ("ext.") LPM tables.

Use the following command to modify the configuration that the show iproute reserved-entries command displays:

```
configure iproute reserved-entries [ num_routes_needed | maximum | default ] slot
[all |slot_num]
```

Example

The following command displays the reserved space for IP routes:

```
# show iproute reserved-entries
IPv4      # Reserved Routes      Minimum #
Slot  Type                Routes      IPv4      (or IPv6)                IPv4 Hosts
-----
1     G48Pe                Internal    464      ( 232) [default]                16
2     G48Pe                Internal    480      ( 240) [maximum]                0
3     G48Ta                Internal   12240    ( 6120) [default]                16
4     G48Ta                Internal   12256    ( 6128) [maximum]                0
5     G8X                  Internal    n/a
6
7     10G4Xa              Internal   10000    ( 5000)                2256
8     10G4Xa              Internal    9000    ( 4500)                3256
9     G48Xa               Internal    9000    ( 4500)                3256
10    8900-10G8X-x1      External   507872   ( int.) [default]            16384
```



```

Maximum supported # IPv4 (or IPv6) Reserved Routes:
"e"-series      Internal      480  ( 240)
"a" and "c"     Internal     12256 ( 6128)
"xl"-series     External    524256 ( int.)
IPv6: "xl"-series "int." indicates external table, up to 8192 IPv6 routes.
IPv6: "xl"-series "ext." indicates external table.
Note: IPv4 Hosts can occupy unused HW Route table space,
unless # Reserved Routes is "maximum".

```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on BlackDiamond X8, BlackDiamond 8000 series modules, SummitStack, and on Summit X460, X480, X670, and X770 switch families.

show iproute reserved-entries statistics

```
show iproute reserved-entries statistics { slot slot_num }
```

Description

Displays the current usage statistics of the Longest Prefix Match (LPM) hardware table and the Layer 3 hardware hash table by resource type.

Syntax Description

<i>slot_num</i>	For BlackDiamond 8800 series switches and SummitStack only, this option displays the statistics for the specified slot.
-----------------	---

Default

N/A.

Usage Guidelines

This command shows the current number of IP routes and local and remote IPv4 hosts in the LPM hardware table. It also shows the number of IPv4 unicast, multicast, and IPv6 unicast entries in the Layer 3 hardware hash table. The theoretical maximums for each individual resource type are shown at the bottom of the output. These maximum values cannot all be achieved simultaneously, and individual values might not be reached depending on the addresses or routes in use.

The ExtremeXOS software supports the coexistence of higher- and lower-capacity hardware in the same BlackDiamond 8800 chassis or Summit family switch stack. To allow for coexistence and increased hardware forwarding, when the number of IPv4 routes exceeds 25,000, the lower-capacity hardware automatically transitions from using LPM routing to forwarding of individual remote hosts,



also known as IP Forwarding Database (IP FDB) mode. Higher-capacity hardware continues using LPM routing. Lower capacity hardware operating in IP FDB mode will be indicated with a d flag in the output of `show iproute reserved-entries statistics` command, indicating that only direct routes are installed. For more information, see [Coexistence of Higher- and Lower-Capacity Hardware](#) in the *ExtremeXOS Concepts Guide*.

Example

The following command displays usage statistics for the LPM and Layer 3 hardware tables:

```
# show iproute reserved-entries statistics
|-----In HW Route Table-----|  |--In HW L3 Hash Table--|
# Used Routes  # IPv4 Hosts  IPv4  IPv4  IPv6  IPv4
Slot  Type                IPv4  IPv6   Local Remote  Local  Rem.  Loc.
MCast
-----
1     G48Pe                100   20>   372   0     613   0    34
12
2     G48Pe                100   20>   -     -     637   0    34
12
3     G48Ta                100   20>   386   0     2114  0    34
12
4     G48Ta                100   20>   -     -     2253  0    34
12
5     G8X                  100   -     -     -     2253  0    -
12
6     -                    -     -     -     -     -     -     -
-
7     10G4Xa              100   20>   2288  0     212   0    34
12
8     10G4Xa              100   20>   2500  0     0     0    34
12
9     G48Ta                -     -     -     -     -     -     -
-
10    8900-10G8X-xl      100   20>   2500  0     0     0    34
12
Theoretical maximum for each resource type:
"e"-series          480    240    512    512    2045   2048   1024  *2048
"a"-series         12256   6128   8189  12288   8189   8192   4096  *3000
"c"-series         12256   6128   12256  12288   8189   8192   4096  *6000
"xl"-series        262112  8192   260000 40960   16381  16384  8192  *6000
original           12256   n/a    n/a    n/a    8189   8188   n/a   *3000
Flags: (!) Indicates all reserved route entries in use.
(d) Indicates only direct IPv4 routes are installed.
(>) Some IPv6 routes with mask > 64 bits are installed and do not use
entries in the internal HW Route Table.
(R) IPv6 hosts in external HW Route Table.
(*) Assumes IP Multicast compression is on.
```

History

This command was first available in ExtremeXOS 12.1.



Support for additional IPv4 local host entries was added with the distributed IP ARP feature for BlackDiamond X8, BlackDiamond 8000 c- and xl-series modules in ExtremeXOS 12.5.

Platform Availability

This command is available on BlackDiamond X8, BlackDiamond 8000 series modules, SummitStack, and on Summit X440, X460, X480, and X670 switch families.

show ipstats

```
show ipstats {ipv4} {vlan name | vr vrname}
```

Description

Displays IP statistics for the switch CPU or for a particular VLAN.

Syntax Description

<i>name</i>	Specifies a VLAN name.
<i>vrname</i>	Specifies a VR or VRF.

Default

N/A.

Usage Guidelines

This command only shows statistics of the CPU-handled packets. Not all packets are handled by the CPU.

If you do not specify a VR or VRF, the command applies to the current VR context.

The fields displayed in the show ipstats command are defined in the following tables.

Table 50: Global IP Statistics Field Definitions

Field	Definition
InReceives	Total number of incoming IP packets processed by the CPU.
InUnicast	Total number of unicast IP packets processed by the CPU.
InBcast	Total number of broadcast IP packets processed by the CPU.
InMcast	Total number of multicast IP packets processed by the CPU.
InHdrEr	Total number of packets with an IP Header Error forwarded to the CPU.
Bad vers	Total number of packets with a version other than IPv4 in the IP version field.
Bad chksum	Total number of packets with a bad IP checksum forwarded to the CPU.
Short pkt	IP packets that are too short.



Table 50: Global IP Statistics Field Definitions (continued)

Field	Definition
Short hdr	IP packets with a header that is too short.
Bad hdrlen	IP packets with a header length that is less than the length specified.
Bad length	IP packets with a length less than that of the header.
InDelivers	IP packets passed to upper layer protocols.
Bad Proto	IP packets with unknown (not standard) upper layer protocol.
OutRequest	IP packets sent from upper layers to the IP stack.
OutDiscard	IP packets that are discarded due to lack of buffer space or the router interface being down, or broadcast packets with broadcast forwarding disabled.
OutNoRoute	IP packets with no route to the destination.
Forwards	ForwardOK and Fwd Err aggregate count.
ForwardOK	Total number of IP packets forwarded correctly.
Fwd Err	Total number of IP packets that cannot be forwarded.
NoFwding	Aggregate number of IP packets not forwarded due to errors.
Redirects	IP packets forwarded on the same network.
No route	Not used.
Bad TTL	IP packets with a bad time-to-live.
Bad MC TTL	IP packets with a bad multicast time-to-live.
Bad IPdest	IP packets with an address that does not comply with the IPv4 standard.
Blackhole	IP packets with a destination that is a blackhole entry.
Output err	Not used. This is the same as Fwd Err.
MartianSrc	IP packets with an invalid source address.

Table 51: Global ICMP Statistics Field Definitions

Field	Definition
OutResp	Echo replies sent from the CPU.
OutError	Redirect from broadcast or multicast source addresses.
InBadcode	Incoming ICMP packets with an invalid CODE value.
InTooshort	Incoming ICMP packets that are too short.
Bad chksum	Incoming ICMP packets with checksum errors.
In Badlen	Incoming ICMP packets with length errors.
echo reply (In/Out):	ICMP "echo reply" packets that are received and transmitted.
destination unreachable (In/Out):	ICMP packets with destination unreachable that are received and transmitted.



Table 51: Global ICMP Statistics Field Definitions (continued)

Field	Definition
port unreachable (In/Out):	ICMP packets with port unreachable that are received and transmitted.
echo (In/Out):	ICMP echo packets that are received and transmitted.

Table 52: Global IGMP Statistics Field Definitions

Field	Definition
Out Query	Number of IGMP query messages sent by the router.
Out Report	Number of reports sent on an active multicast route interface for reserved multicast addresses and for regular IGMP reports forwarded by the query router.
Out Leave	Number of IGMP out leave messages forwarded for IP multicast router interfaces.
In Query	Number of IGMP query messages received.
In Report	Number of IGMP report messages received (mostly from hosts).
In Leave	Number of IGMP leave messages received (mostly from hosts).
In Error	Number of IGMP packets with bad header fields or checksum failures.

Table 53: Router Interface Statistics Field Definitions

Field	Definition
Packets IN/OUT	Total number of IP packets received or transmitted on a VLAN router interface.
Octets IN/OUT	Total number of octets received or transmitted on a VLAN router interface.
Mcast packets IN/OUT	Total number of multicast packets received or transmitted on a VLAN router interface.
Bcast packets IN/OUT	Total number of broadcast packets received or transmitted on a VLAN router interface.
Errors IN/OUT	Total number of IP packets with errors received or transmitted on a VLAN router interface.
Discards IN/OUT	Total number of IP packets that cannot travel up to the CPU due to lack of buffer space.
Unknown Protocols IN/OUT	Total number of IP packets with unknown upper layer protocols received by the router interface.

Example

The following command displays IP statistics for the VLAN accounting:

```
show ipstats vlan accounting
```



History

This command was first available in ExtremeXOS 10.1.

The keyword `ipv4` was added in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show udp-profile

```
show udp-profile {vlan vlan-name | {policy} policy-name}
```

Description

Displays UDP forwarding profiles.

Syntax Description

<i>vlan-name</i>	Specifies a VLAN.
<i>policy-name</i>	Specifies a UDP forwarding profile.

Default

If no VLAN or policy is specified, all configured profiles are displayed.

Usage Guidelines

UDP profiles can also be displayed by using the policy manager command `show policy {policy-name | detail}`. However, the format of the policy display is different than that for this command.

Example

The following command displays all the configured UDP forwarding profiles on the switch:

```
show udp-profile
```

The following is sample output:

```
UDP Profile Name: move_to7
Number of datagram forwarded: 181
Dest UDP Port: 67 Fwd to IP Addr: 20.0.0.5
Dest UDP Port: 67 Fwd to VLAN: to7
Applied to incoming traffic on VLANS:
to-mariner
```



History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms that use the Edge, Advanced Edge, or Core license. For information on the licenses available for each platform, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure bootprelay dhcp-agent information check

unconfigure bootprelay dhcp-agent information check

Description

Disables Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82) checking.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

In some instances, a DHCP server may not properly handle a DHCP request packet containing a relay agent option. Use this command to disable the switch from preventing DHCP reply packets with invalid or missing relay agent options from being forwarded to the client.

To enable this check, use the following command:

```
configure bootprelay dhcp-agent information check
```

Example

The following command disables the DHCP relay agent option check:

```
unconfigure bootprelay dhcp-agent information check
```

History

This command was first available in ExtremeXOS 11.1.



Platform Availability

This command is available on all platforms.

unconfigure bootprelay dhcp-agent information circuit-id port-information

```
unconfigure bootprelay dhcp-agent information circuit-id port-information ports  
[port_list | all]
```

Description

Configures the circuit ID sub-option that identifies the specified ports to use the default value.

Syntax Description

<i>port_list</i>	Specifies a list of one or more ports that are to be configured to use the default value.
all	Specifies that all ports are to be configured to use the default value.

Default

The port_info is encoded as ((slot_number * 1000) + port_number). For example, if the DHCP request is received on port 3:12, the default circuit ID port_info value is 3012. On non-slot-based switches, the default circuit ID port_info value is simply the port number.

Usage Guidelines

None.

Example

The following command configures port 1:3 to use the default circuit ID port information value:

```
unconfigure bootprelay dhcp-agent information circuit-id port-information  
ports 1:3
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.



unconfigure bootprelay dhcp-agent information circuit-id vlan-information

```
unconfigure bootprelay dhcp-agent information circuit-id vlan-information {vlan}
[vlan_name | all]
```

Description

Configures the circuit ID sub-option that identifies the specified VLANs to use the default value.

Syntax Description

<i>vlan_name</i>	Names a VLAN to be configured to use the default value.
all	Specifies that all VLANs are to be configured to use the default value.

Default

N/A.

Usage Guidelines

None.

Example

The following command configures VLAN blue to use the default VLAN information for the circuit ID sub-option:

```
unconfigure bootprelay dhcp-agent information circuit-id vlan-information blue
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

unconfigure bootprelay dhcp-agent information option

```
unconfigure bootprelay dhcp-agent information option
```



Description

Disables the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82).

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

To enable the DHCP relay agent option (option 82), use the following command:

```
configure bootprelay dhcp-agent information option
```

Example

The following command disables the DHCP relay agent option:

```
unconfigure bootprelay dhcp-agent information option
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

unconfigure bootprelay dhcp-agent information policy

```
unconfigure bootprelay dhcp-agent information policy
```

Description

Unconfigures the Dynamic Host Configuration Protocol (DHCP) relay agent option (option 82) policy.

Syntax Description

This command has no arguments or variables.

Default

Replace.



Usage Guidelines

Use this command to unconfigure the policy for the relay agent.

Example

The following command unconfigures the DHCP relay agent option 82 policy:

```
unconfigure bootrelay dhcp-agent information policy
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

unconfigure bootrelay dhcp-agent information remote-id

```
unconfigure bootrelay dhcp-agent information remote-id {vr vrid}
```

Description

Configures the remote ID sub-option to the default value.

Syntax Description

<i>vrid</i>	Specifies the VR on which to configure the remote ID sub-option to the default value.
-------------	---

Default

The switch MAC address.

Usage Guidelines

None.

Example

The following command configures the remote ID sub-option to use the default value on the current VR:

```
configure bootrelay dhcp-agent information remote-id
```



History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

unconfigure bootprelay include-secondary

Short reference description.

```
unconfigure bootprelay {ipv4 | ipv6} {vlan} vlan_name include-secondary
```

Description

Removes the specified smart relay configuration that was specified at the VLAN level.

Syntax Description

ipv4	Specifies unconfiguring the DHCPv4 BOOTP Relay service (default).
ipv6	Specifies unconfiguring the DHCPv6 BOOTP Relay service.
vlan	Unconfigures BOOTP relay for a specified VLAN.
<i>vlan_name</i>	Specifies the VLAN name.
include-secondary	Removes the include-secondary configuration for the specified VLAN.

Default

IPv4 is the default relay service.

Usage Guidelines

Use this command to remove the specified smart relay configuration that was specified at the VLAN level.

Example

The following command removes DHCPv4 BOOTP Relay service, and removes the **include-secondary** configuration for the *vlan 100* VLAN:

```
unconfigure bootprelay ipv4 vlan vlan_100 include-secondary
```

History

This command was first available in ExtremeXOS 15.4.



Platform Availability

This command is available on all platforms.

unconfigure icmp

unconfigure icmp

Description

Resets all ICMP settings to the default values.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command resets all ICMP settings to the default values.

```
unconfigure icmp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

unconfigure iparp

unconfigure iparp

Resets the following to their default values:

- IP ARP timeout
- Maximum ARP entries



- Maximum ARP pending entries
- ARP checking
- ARP refresh

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command resets IP ARP timeout to its default value:

```
unconfigure iparp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

unconfigure iproute priority

```
unconfigure iproute {ipv4} priority [all | blackhole | bootp | ebgp | ibgp | icmp
| isis | isis-level-1 | isis-level-1-external | isis-level-2 | isis-level-2-
external | mpls | ospf-as-external | ospf-extern1 | ospf-extern2 | ospf-inter |
ospf-intra | rip | static] {vr vrname}
```

Description

Unconfigures the priority for all IP routes from one or all route origin types.

Syntax Description

all	Specifies all route origins.
blackhole	Specifies the blackhole route.



bootp	Specifies BOOTP.
ebgp	Specifies E-BGP routes.
ibgp	Specifies I-BGP routes.
icmp	Specifies ICMP.
isis	Specifies IS-IS and applies only to blackhole routes installed for summary addresses.
isis-level-1	Specifies ISIS Level 1 routing.
isis-level-1-external	Specifies ISIS Level 1 External routing.
isis-level-2	Specifies ISIS Level 2 routing.
isis-level-2-external	Specifies ISIS Level 2 External routing.
mpls	Specifies MPLS routing.
ospf-as-external	Specifies OSPF as External routing.
ospf-extern1	Specifies OSPF External 1 routing.
ospf-extern2	Specifies OSPF External 2 routing.
ospf-inter	Specifies OSPFInter routing.
ospf-intra	Specifies OSPFIntra routing.
rip	Specifies RIP.
static	Specifies static routes.
vrname	Specifies a VR or VRF name.

Default

N/A.

Usage Guidelines

Default Route Priorities

The following table lists the default priorities that apply after you enter this command.

Route Origin	Priority
Direct	10
MPLS	20
Blackhole	50
Static	1100
ICMP	1200
EBGP	1700
IBGP	1900
OSPFIntra	2200



Route Origin	Priority
OSPFInter	2300
IS-IS	2350
IS-IS L1	2360
IS-IS L2	2370
RIP	2400
OSPFAExt	3100
OSPF External 1	3200
OSPF External 2	3300
IS-IS L1 Ext	3400
IS-IS L2 Ext	3500
BOOTP	5000

Example

The following command returns the IP route priority for all route origins to the default values:

```
unconfigure iproute priority all
```

History

This command was first available in ExtremeXOS 12.1.2.

Platform Availability

This command is available on all platforms.

unconfigure irdp

```
unconfigure irdp
```

Syntax Description

This command has no arguments or variables.

Description

Resets all router advertisement settings to the default values.

Default

N/A.



Usage Guidelines

None.

Example

The following command resets all router advertisement settings to the default values.

```
unconfigure irdp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms that use the Edge, Advanced Edge, or Core license. For information on the licenses available for each platform, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure vlan subvlan-address-range

```
unconfigure vlan vlan_name subvlan-address-range
```

Description

Unconfigures subVLAN address ranges on each subVLAN to prohibit the entry of IP addresses from hosts outside of the configured range.

Syntax Description

<i>vlan_name</i>	Specifies a subVLAN name.
------------------	---------------------------

Default

N/A.

Usage Guidelines

This command removes a subVLAN address range. There is no error checking to prevent the configuration of overlapping subVLAN address ranges between multiple subVLANs. Doing so can result in unexpected behavior of ARP within the superVLAN and associated subVLANs.

History

This command was first available in ExtremeXOS 12.1.



Platform Availability

This command is available on all platforms.

unconfigure vlan udp-profile

```
unconfigure vlan vlan_name udp-profile
```

Description

Removes any UDP forwarding profile from a VLAN.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
------------------	------------------------

Default

No UDP profiles are associated with the VLAN.

Usage Guidelines

None.

Example

The following command removes any UDP forwarding profile from the VLAN to-sales:

```
unconfigure vlan to-sales udp-profile
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms that use the Edge, Advanced Edge, or Core license. For information on the licenses available for each platform, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



35 IPv6 Unicast Commands

```
clear ipv6 dad
clear neighbor-discovery cache
configure iproute add (IPv6)
configure iproute add blackhole
configure iproute add blackhole ipv6 default
configure iproute add default
configure iproute delete
configure iproute delete blackhole
configure iproute delete blackhole ipv6 default
configure iproute delete default
configure iproute ipv6 priority
configure iproute sharing max-gateways
configure ipv6 dad
configure ipv6 hop-limit
configure neighbor-discovery cache add
configure neighbor-discovery cache delete
configure neighbor-discovery cache max_entries
configure neighbor-discovery cache max_pending_entries
configure neighbor-discovery cache timeout
configure vlan router-discovery add prefix
configure vlan router-discovery delete prefix
configure vlan router-discovery default-lifetime
configure vlan router-discovery link-mtu
configure vlan router-discovery managed-config-flag
configure vlan router-discovery max-interval
configure vlan router-discovery min-interval
configure vlan router-discovery other-config-flag
configure vlan router-discovery reachable-time
configure vlan router-discovery retransmit-time
configure vlan router-discovery set prefix
configure tunnel ipaddress
create tunnel 6to4
create tunnel gre destination source
create tunnel ipv6-in-ipv4
delete tunnel
disable icmp redirects ipv6 fast-path
disable ipforwarding ipv6
```

```
disable iproute ipv6 compression
disable iproute ipv6 sharing
disable neighbor-discovery refresh
disable router-discovery
disable tunnel
enable icmp redirects ipv6 fast-path
enable ipforwarding ipv6
enable ipforwarding
enable iproute ipv6 compression
enable iproute ipv6 sharing
enable neighbor-discovery refresh
enable router-discovery
enable tunnel
rtlookup
rtlookup rpf
run ipv6 dad
show ipconfig ipv6
show iproute ipv6
show iproute ipv6 origin
show ipstats ipv6
show ipv6 dad
show neighbor-discovery cache ipv6
show router-discovery
show tunnel
unconfigure iproute ipv6 priority
unconfigure neighbor-discovery cache
unconfigure vlan router-discovery
unconfigure vlan router-discovery default-lifetime
unconfigure vlan router-discovery hop-limit
unconfigure vlan router-discovery link-mtu
unconfigure vlan router-discovery managed-config-flag
unconfigure vlan router-discovery max-interval
unconfigure vlan router-discovery min-interval
unconfigure vlan router-discovery other-config-flag
unconfigure vlan router-discovery reachable-time
unconfigure vlan router-discovery retransmit-time
unconfigure tunnel
```

This chapter describes commands for configuring and managing the following IPv6 features:

- IPv6 unicast routing
- Duplicate Address Detection (DAD)
- Route sharing



- Route compression
- IPv6 multinetting

For an introduction to these IPv6 features, see the ExtremeXOS Concepts Guide.

clear ipv6 dad

```
clear ipv6 dad {{vr} vr_name {ipaddress} | vr all | {vlan} vlan_name} {counters}
```

Description

Clears the counters for the DAD feature.

Syntax Description

<i>vr_name</i>	Specifies a VR for which to clear the counters.
<i>ipaddress</i>	Specifies an IPv6 address for which to clear the counters.
<i>vlan_name</i>	Specifies a VLAN for which to clear the counters.

Default

If you do not specify a VR or VRF, the command applies to the current VR context.

Usage Guidelines

The `vr all` option clears the DAD counters for all IPv6 interfaces on the switch.

This command clears the DAD failure counters and removes the MAC for the conflicting IPv6 address after the duplicate address condition has been resolved. The DAD counters and saved MAC addresses are not automatically cleared; they must be cleared with this command.

Example

The following command clears the DAD counters for all IPv6 interfaces in all VRs:

```
clear ipv6 dad vr all
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on BlackDiamond X8, BlackDiamond 8800 series switches and Summit family switches.



clear neighbor-discovery cache

```
clear neighbor-discovery cache ipv6 {ipv6address {vr vr_name} | vlan vlan_name |
vr vr_name}
```

Description

Deletes a dynamic entry from the neighbor cache.

Syntax Description

<i>vr_name</i>	Specifies a VR or VRF.
<i>ipv6address</i>	Specifies an IPv6 address.
<i>vlan_name</i>	Specifies an IPv6 configured VLAN.

Default

N/A.

Usage Guidelines

This command clears dynamic entries from the neighbor cache. The *vr* option is used to specify the VR or VRF on which the operation is performed. When this option is omitted it applies to current VR context.

When the *ipv6address* or *vlan* options are specified, only the entries with matching IPv6 addresses or that correspond to that VLAN are cleared.

Example

The following command clears all entries from the neighbor cache:

```
clear neighbor-discovery cache
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure iproute add (IPv6)

```
configure iproute add ipv6Netmask [ipv6Gateway | ipv6ScopedGateway] {metric} {vr_name} {multicast | multicast-only | unicast | unicast-only}
```

Description

Adds an IPv6 static route to the routing table.

Syntax Description

<i>ipv6Netmask</i>	Specifies an IPv6 address/prefix length.
<i>ipv6Gateway</i>	Specifies a gateway.
<i>ipv6ScopedGateway</i>	Specifies a scoped gateway.
<i>metric</i>	Specifies a cost metric.
<i>vr_name</i>	Specifies the VR or VRF to which the route is added.
multicast	Adds the specified route to the multicast routing table.
multicast-only	Adds the specified route to the multicast routing table. This option is provided for backward compatibility with releases prior to ExtremeXOS Release 12.1.
unicast	Adds the specified route to the unicast routing table.
unicast-only	Adds the specified route to the unicast routing table. This option is provided for backward compatibility with releases prior to ExtremeXOS Release 12.1.

Default

If you do not specify a VR or VRF, the current VR context is used. If you do not specify a metric, then the default metric of 1 is used.

Usage Guidelines

Use a prefix length of 128 to indicate a host entry.



Note

Although dynamic unicast routes can be captured in the multicast routing table, unicast static routes cannot be captured in the multicast routing table. To create a static route for the multicast routing table, you must specify the multicast option.

Example

The following command adds a static route to the routing table:

```
configure iproute add 2001:db8:0:1111::/64 fe80::1111%default
```



History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure iproute add blackhole

```
configure iproute add blackhole {ipv6} [ipv6Netmask] {vr vr_name} {multicast-only
| unicast-only}
```

Description

Adds a blackhole address to the routing table. All traffic destined for an unknown IPv6 destination is silently dropped.

Syntax Description

<i>ipv6Netmask</i>	Specifies an IPv6 address/prefix length.
<i>vr_name</i>	Specifies the VR or VRF to which the route is added.
multicast-only	Specifies only multicast traffic for the route.
unicast-only	Specifies only unicast traffic for the route.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

A blackhole entry directs packets with a matching specified address prefix to be discarded. Blackhole entries are useful as a security measure or in special circumstances where a specific destination address must be discarded. Blackhole entries are treated like permanent entries in the event of a switch reset or power off/on cycle.

The packets are silently discarded. In other words, no ICMP message is sent to indicate that the packets are discarded.



Example

The following command causes packets with a destination address of 2001:db8::3452 to be silently discarded:

```
configure iproute add blackhole 2001:db8::3452/128
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure iproute add blackhole ipv6 default

```
configure iproute add blackhole ipv6 default {vr vr_name} {multicast-only | unicast-only}
```

Description

Adds a default blackhole route to the routing table. All traffic destined for an unknown IPv6 destination is silently dropped.

Syntax Description

<i>vr_name</i>	Specifies the VR or VRF to which the route is added.
multicast-only	Specifies only multicast traffic for the route.
unicast-only	Specifies only unicast traffic for the route.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

While a default route is for forwarding traffic destined to an unknown IPv6 destination, and a blackhole route is for discarding traffic destined to a specified IPv6 destination, a default blackhole route is for discarding traffic to the unknown IPv6 destination.

Using this command, all traffic with an unknown destination is discarded.



The default blackhole route is treated like a permanent entry in the event of a switch reset or power off/on cycle. The default blackhole route's origin is "b" or "blackhole" and the gateway IPv6 address for this route is ::.

The packets are silently discarded. In other words, no ICMP message is sent to indicate that the packets are discarded.

Example

The following command adds a blackhole default route into the routing table:

```
configure iproute add blackhole ipv6 default
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure iproute add default

```
configure iproute add default [ipv6Gateway | ipv6ScopedGateway] {metric} {vr  
vr_name} {multicast-only | unicast-only}
```

Description

Adds a default gateway to the routing table.

Syntax Description

<i>ipv6Gateway</i>	Specifies a VLAN gateway IPv6 address.
metric	Specifies a cost metric. If no metric is specified, the default of 1 is used.
<i>ipv6ScopedGateway</i>	Specifies a scoped gateway.
<i>vr_name</i>	Specifies the VR or VRF to which the route is added.

Default

If no metric is specified, the default metric of 1 is used. If you do not specify a VR or VRF, the current VR context is used.



Usage Guidelines

Default routes are used when the router has no other dynamic or static route to the requested destination. A default gateway must be located on a configured IPv6 interface. Use the unicast-only or multicast-only options to specify a particular traffic type. If not specified, both unicast and multicast traffic uses the default route.

Example

The following command configures a default route for the switch:

```
configure iproute add default 2001:db8::1234:5678
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure iproute delete

```
configure iproute delete ipv6Netmask [ipv6Gateway | ipv6ScopedGateway] {vr  
vr_name}
```

Description

Deletes an IPv6 static route from the routing table.

Syntax Description

<i>ipv6Netmask</i>	Specifies an IPv6 address/prefix length.
<i>ipv6Gateway</i>	Specifies a gateway.
<i>ipv6ScopedGateway</i>	Specifies a scoped gateway.
<i>vr_name</i>	Specifies the VR or VRF from which the route is deleted.

Default

If you do not specify a VR or VRF, the current VR context is used.



Usage Guidelines

Use a prefix length of 128 to indicate a host entry.

Example

The following command deletes a static address from the routing table:

```
configure iproute delete 2001:db8:0:1111::/64 fe80::1111
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure iproute delete blackhole

```
configure iproute delete blackhole [ipv6Netmask] {vr vr_name}
```

Description

Deletes a blackhole route from the routing table.

Syntax Description

<i>ipv6Netmask</i>	Specifies an IPv6 address/prefix length.
<i>vr_name</i>	Specifies the VR or VRF from which the route is deleted.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

A blackhole entry directs packets with a specified destination address to be discarded. Blackhole entries are useful as a security measure or in special circumstances where a specific destination address must be discarded. Blackhole entries are treated like permanent entries in the event of a switch reset or power off/on cycle.



Example

The following command deletes a blackhole route from the routing table for packets with a destination address of 2001:db8::3452, so the packets are no longer discarded:

```
configure iproute delete blackhole 2001:db8::3452/128
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure iproute delete blackhole ipv6 default

```
configure iproute delete blackhole ipv6 default {vr vr_name}
```

Description

Deletes a default blackhole route from the routing table.

Syntax Description

<code>vr_name</code>	Specifies the VR or VRF from which the route is deleted.
----------------------	--

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

While a default route is for forwarding traffic destined to an unknown IPv6 destination, and a blackhole route is for discarding traffic destined to a specified IPv6 destination, a default blackhole route is for discarding traffic to the unknown IPv6 destination.

Using this command, all traffic with an unknown destination is discarded.

The default blackhole route is treated like a permanent entry in the event of a switch reset or power off/on cycle. The default blackhole route's origin is "b" or "blackhole" and the gateway IPv6 address for this route is ::.



Example

The following command deletes a blackhole default route from the routing table:

```
configure iproute delete blackhole default
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure iproute delete default

```
configure iproute delete default [ipv6Gateway | ipv6ScopedGateway] {vr vr_name}
```

Description

Deletes a default gateway from the routing table.

Syntax Description

<i>ipv6Gateway</i>	Specifies a VLAN gateway IPv6 address.
<i>ipv6ScopedGateway</i>	Specifies a scoped gateway.
<i>vr_name</i>	Specifies the VR or VRF from which the route is deleted.

Default

If no metric is specified, the default metric of 1 is used. If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

Default routes are used when the router has no other dynamic or static route to the requested destination. A default gateway must be located on a configured IPv6 interface.



Example

The following command deletes a default route from the switch:

```
configure iproute delete default 2001:db8::1234:5678
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure iproute ipv6 priority

```
configure iproute ipv6 priority [ripng | blackhole | icmp | static | ospfv3-intra
| ospfv3-inter | ospfv3-as-external | ospfv3-extern1 | ospfv3-extern2 | isis-
level-1 | isis-level-2 | isis-level-1-external | isis-level-2-external] priority
{vr vr_name}
```

Description

Changes the priority for all routes from a particular route origin.

Syntax Description

ripng	Specifies RIPng.
blackhole	Specifies the blackhole route.
icmp	Specifies ICMP.
static	Specifies static routes.
ospfv3-intra	Specifies OSPFv3 Intra routing.
ospfv3-inter	Specifies OSPFv3 Inter routing.
ospfv3-as-external	Specifies OSPFv3 AS External routing.
ospfv3-extern1	Specifies OSPFv3 External 1 routing.
ospfv3-extern2	Specifies OSPFv3 External 2 routing.
isis-level-1	Specifies ISIS Level 1 routing.
isis-level-2	Specifies ISIS Level 2 routing.
isis-level-1-external	Specifies ISIS Level 1 External routing.
isis-level-2-external	Specifies ISIS Level 2 External routing.



<i>priority</i>	Specifies a priority number in the range of 11 to 65534.
<i>vr_name</i>	Specifies a VR or VRF name.

Default

The following table lists the relative priorities assigned to routes depending upon the learned source of the route.

Table 54: Route Priorities

Route Origin	Priority
Direct	10
BlackHole	50
Static	1100
ICMP	1200
OSPF3Intra	2200
OSPF3Inter	2300
IS-IS L1	2360
IS-IS L2	2370
RIPg	2400
OSPFv3 ASExt	3100
OSPFv3 Extern1	3200
OSPFv3 Extern2	3300
IS-IS L1 Ext	3400
IS-IS L2 Ext	3500

Usage Guidelines

Although these priorities can be changed, do not attempt any manipulation unless you are expertly familiar with the possible consequences. If you change the route priority, you must save the configuration and reboot the system.



Note

The priority for a blackhole route can not overlap with the priority of any other route origin.

Example

The following command sets the IPv6 route priority for static routing to 1200:

```
configure iproute ipv6 priority static 1200
```



History

This command was first available in ExtremeXOS 11.2.

The `vr` option was added in ExtremeXOS 12.1.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure iproute sharing max-gateways

```
configure iproute sharing max-gateways max_gateways
```

Description

Specifies the maximum number of gateways in each gateway set in the equal-cost multipath (ECMP) hardware table.

Syntax Description

<i>max_gateways</i>	Specifies the maximum ECMP gateways per IP destination supported in the switch hardware. The only values allowed are 2, 4, 8, 16 and 32.
---------------------	--

Default

Four gateways.

Usage Guidelines

When IPv6 route sharing is enabled, the maximum number of `gateways` value represents the maximum number of next-hop gateways that can be used for communications with a destination subnet. Each gateway represents an alternative path to a subnet. The gateways can be defined with static routes, or they can be learned through the OSPFv3, or BGP protocols. The ExtremeXOS Release Notes lists the total number of route destinations and the total combinations of gateway sets that each platform can support with the different `max-gateways` option selections. For more information on selecting the maximum number of gateways and how this affects different platforms, see ECMP Hardware Table in the ExtremeXOS Concepts Guide. The value for `max-gateways` applies to both IPv4 and IPv6 on all VRs. The maximum number of gateways in each IPv4 or IPv6 gateway set can be 2, 4, 16, or 32.



Note

You must save the configuration and reboot the switch for the new value to take effect.



Example

The following command changes the maximum number of ECMP gateways per subnet or gateway set to eight: configure

```
iproute sharing max-gateways 8
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on the platforms listed for the IPv4 or IPv6 unicast routing feature in the Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ipv6 dad

```
configure ipv6 dad [off | on | {on} attempts max_solicitations] {{vr} vr_name | vr all}
```

Description

Configures the operation of the duplicate address detection (DAD) feature on the specified VR.

Syntax Description

<i>max_solicitations</i>	Specifies the number of times the DAD feature tests for a duplicate address. The range is 1 to 10, and the default value is 1.
<i>vr_name</i>	Specifies a VR on which to enable this feature.

Default

DAD status: On on VR-Default.

Maximum solicitations: 1 for VR-Default.

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

When the DAD feature is enabled, the switch checks for duplicate IPv6 addresses on the specified VR when an IPv6 interface is initialized, or when a DAD check is initiated with a CLI command. After initialization, and when this feature is off, the switch does not start DAD checks.

Changes to the number of solicitations configuration take affect the next time the DAD check is run.



By default, this command applies to the current VR context, if no VR name is specified. If `vr all` is specified, the command applies to all user VRs and VR-Default.

The DAD feature does not run on loopback VLANs.

Example

The following command enables the DAD feature on all user VRs and VR-Default:

```
configure ipv6 dad on vr all
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on BlackDiamond X8, BlackDiamond 8800 series switches and Summit family switches.

configure ipv6 hop-limit

```
configure ipv6 hop-limit hop_limit {dont-specify-in-ra} {{vr} vr_name | {vlan}  
vlan_name | vlan all}
```

Description

This command allows you to configure the ipv6 hop-limit. This hop-limit is used in all originated IPv6 packets, and (if router discovery is enabled) in outgoing Router Advertisement packets as well.

Syntax Description

<i>hop_limit</i>	Hop limit for all originated IPv6 packets, and the advertised hop-limit for Router Advertisements. Hop limit value between 1 and 255. Default is 64.
dont-specify-in-ra	Sets the advertised hop-limit in Router Advertisements to zero.
vr	Virtual router.
vlan	VLAN.
all	All VLANs.

Default

64.



Usage Guidelines

Use this command to configure the ipv6 hop-limit. The hop-limit is used in all originated IPv6 packets, and (if router discovery is enabled) in outgoing Router Advertisement packets as well.

0 is a special value used only in outgoing Router Advertisements to convey to the receiving hosts that the router has not specified a hop-limit value to be used when originating ipv6 packets. This can be configured by specifying the optional 'dont-specify-in-ra' keyword. The hop-limit can be configured for a vlan, all vlans in a Virtual Router, or all vlans in the system. By default, the hop-limit is configured for all vlans in the current Virtual Router context of the CLI.

Example

Example output not yet available and will be provided in a future release.

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on all platforms.

configure neighbor-discovery cache add

```
configure neighbor-discovery cache {vr vr_name} add [ipv6address | scoped_link_local] mac
```

Description

Adds a static entry to the neighbor cache.

Syntax Description

<i>vr_name</i>	Specifies a VR or VRF.
<i>ipv6address</i>	Specifies an IPv6 address.
<i>scoped_link_local</i>	Specifies a scoped, link-local address.
<i>mac</i>	Specifies a MAC address.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

This command adds static entries to the neighbor cache.



Example

The following command adds a static entry to the neighbor cache:

```
configure neighbor-discovery cache add fe80::2315%default 00:11:22:33:44:55
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure neighbor-discovery cache delete

```
configure neighbor-discovery cache {vr vr_name} delete [ipv6address | scoped_link_local]
```

Description

Deletes a static entry from the neighbor cache.

Syntax Description

<i>vr_name</i>	Specifies a VR or VRF.
<i>ipv6address</i>	Specifies an IPv6 address.
<i>scoped_link_local</i>	Specifies a scoped, link-local address.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

This command deletes static entries from the neighbor cache.

Example

The following command deletes a static entry from the neighbor cache:

```
configure neighbor-discovery cache delete fe80::2315%default
```



History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure neighbor-discovery cache max_entries

```
configure neighbor-discovery cache {vr vr_name} max_entries max_entries
```

Description

Configures the maximum allowed IPv6 neighbor entries.

Syntax Description

<i>vr_name</i>	Specifies a VR or VRF.
<i>max_entries</i>	Specifies the maximum allowed IPv6 neighbor entries. The range is 1 to 20480.

Default

4096.

Usage Guidelines

None.

Example

The following command sets the maximum allowed IPv6 neighbor entries to 512:

```
configure neighbor-discovery cache max_entries 512
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure neighbor-discovery cache max_pending_entries

```
configure neighbor-discovery cache {vr vr_name} max_pending_entries
max_pending_entries
```

Description

Configures the maximum number of pending IPv6 neighbor entries.

Syntax Description

<i>vr_name</i>	Specifies a VR or VRF.
<i>max_pending_entries</i>	Specifies the maximum number of pending IPv6 neighbor entries. The range is 1 to 4096.

Default

1024.

Usage Guidelines

None.

Example

The following command sets the maximum number of pending IPv6 neighbor entries to 2056:

```
configure neighbor-discovery cache max_pending_entries 2056
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure neighbor-discovery cache timeout

```
configure neighbor-discovery cache {vr vr_name} timeout timeout
```

Description

Configures a timeout value for entries in the neighbor cache.



Syntax Description

<i>vr_name</i>	Specifies a VR or VRF.
<i>timeout</i>	Specifies a timeout value for neighbor cache entries. The range is 1 to 32767 minutes.

Default

20 minutes.

Usage Guidelines

None.

Example

The following command configures the neighbor cache timeout for 30 minutes:

```
configure neighbor-discovery cache timeout 30
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vlan router-discovery add prefix

```
configure vlan vlan_name router-discovery {ipv6} add prefix prefix
```

Description

Adds a prefix to the router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>prefix</i>	Specifies the prefix to add.

Default

N/A.



Usage Guidelines

This command adds a prefix to the router advertisement messages for the VLAN. Prefixes defined with this command are only included in the router advertisement messages and have no operational impact on VLANs.

To configure the parameters for this prefix, use the following command:

```
configure vlan vlan_name router-discovery {ipv6} set prefix prefix [autonomous-  
flag auto_on_off | onlink-flag onlink_on_off | preferred-lifetime preflife | valid-  
lifetime validlife]
```

Example

The following command adds the prefix 2001:db8:3456::/64 for the VLAN top_floor:

```
configure vlan top_floor router-discovery add prefix 2001:db8:3456::/64
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure vlan router-discovery delete prefix

```
configure vlan vlan_name router-discovery {ipv6} delete prefix [prefix | all]
```

Description

Deletes prefixes from the router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>prefix</i>	Specifies the prefix to delete.
all	Specifies to delete all prefixes.

Default

N/A.



Usage Guidelines

This command deletes previously defined router advertisement prefixes.

Example

The following command deletes the prefix 2001:db8:3161::/64 for the VLAN top_floor:

```
configure vlan top_floor router-discovery delete 2001:db8:3161::/64
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vlan router-discovery default-lifetime

```
configure vlan vlan_name router-discovery {ipv6} default-lifetime defaultlifetime
```

Description

Configures the router lifetime value sent in router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>defaultlifetime</i>	Specifies the router lifetime. Range is 0, max-interval to 9000 seconds.

Default

1800 seconds.

Usage Guidelines

This command configures the router lifetime value to be included in the router advertisement messages.

The value is specified in seconds and is either 0, or between max-interval and 9000 seconds. A value of 0 indicates that the router is not to be used as a default router.



After a host sends a router solicitation, and receives a valid router advertisement with a non-zero router lifetime, the host must desist from sending additional solicitations on that interface, until an event such as re-initialization takes place.

Example

The following command configures the default-lifetime to be 3600 seconds for the VLAN `top_floor`:

```
configure vlan top_floor router-discovery default-lifetime 3600
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vlan router-discovery link-mtu

```
configure vlan vlan_name router-discovery {ipv6} link-mtu linkmtu
```

Description

Configures the link MTU value sent in router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>linkmtu</i>	Specifies the link MTU. Range is 0 to 9216.

Default

0, meaning that no link MTU information is sent.

Usage Guidelines

This command configures the link MTU placed into the router advertisement messages. Advertisement of the MTU helps ensure use of a consistent MTU by hosts on the VLAN.

The minimum value is 0. The maximum value is 9216. The default value is 0, which means that no link MTU information is included in the router discovery messages.



Example

The following command configures the link MTU to be 5126 for the VLAN `top_floor`:

```
configure vlan top_floor router-discovery link-mtu 5126
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vlan router-discovery managed-config-flag

```
configure vlan vlan_name router-discovery {ipv6} managed-config-flag on_off
```

Description

Configures the managed address configuration flag value sent in router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>on_off</i>	Specifies setting the flag to on or off.

Default

Off.

Usage Guidelines

This command configures the contents of the managed address configuration flag in the router advertisement messages.

A value of `on` tells hosts to use the administered (stateful) protocol (DHCP) for address autoconfiguration in addition to any addresses autoconfigured using stateless address autoconfiguration. A value of `off` tells hosts to use stateless address autoconfiguration. If this command is not entered, the default value is `off`.



Example

The following command configures the managed address configuration flag to be on for the VLAN `top_floor`:

```
configure vlan top_floor router-discovery managed-config-flag on
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure vlan router-discovery max-interval

```
configure vlan vlan_name router-discovery {ipv6} max-interval maxinterval
```

Description

Configures the maximum time between unsolicited router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>maxinterval</i>	Specifies the maximum time between advertisements, in seconds. Range is 4 to 1800

Default

600 seconds.

Usage Guidelines

This command configures the maximum amount of time before an unsolicited router advertisement message is advertised over the links corresponding to the VLAN.

Example

The following command configures the `max-interval` to be 300 seconds for the VLAN `top_floor`:

```
configure vlan top_floor router-discovery max-interval 300
```



History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vlan router-discovery min-interval

```
configure vlan vlan_name router-discovery {ipv6} min-interval mininterval
```

Description

Configures the minimum time between unsolicited router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>mininterval</i>	Specifies the minimum time between advertisements, in seconds. Range is 3 to 1350 (see guidelines).

Default

200 seconds, or max-interval * .33 (see guidelines).

Usage Guidelines

This command configures the minimum amount of time before an unsolicited router advertisement message is advertised over the links corresponding to the VLAN.

The minimum value is 3 seconds. The maximum time is (.75 * max-interval) seconds. If you do not explicitly set this value, the min-interval value is reset whenever the max-interval is configured. Min-interval will then be dynamically adjusted to .33 times the max-interval.

Example

The following command configures the min-interval to be 300 seconds for the VLAN top_floor:

```
configure vlan top_floor router-discovery min-interval 300
```

History

This command was first available in ExtremeXOS 11.2.



Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vlan router-discovery other-config-flag

```
configure vlan vlan_name router-discovery {ipv6} other-config-flag on_off
```

Description

Configures the other stateful configuration flag value sent in router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>on_off</i>	Specifies setting the flag to on or off.

Default

Off.

Usage Guidelines

This command configures the contents of the other stateful configuration flag in the router advertisement messages.

When set to on, hosts use the administered (stateful) protocol (DHCP) for autoconfiguration of other (non-address) information. If this command is not entered, the default value is off.

Example

The following command configures the other stateful configuration flag to be on for the VLAN `top_floor`:

```
configure vlan top_floor router-discovery other-config-flag on
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure vlan router-discovery reachable-time

```
configure vlan vlan_name router-discovery {ipv6} reachable-time reachbletime
```

Description

Configures the reachable time value in router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>reachbletime</i>	Specifies the reachable time value in advertisements, in milliseconds. Range is 0 to 3,600,000 (one hour).

Default

30,000 milliseconds.

Usage Guidelines

The reachable time is the time, in milliseconds, that a node assumes a neighbor is reachable after having received a reachability confirmation. A value of 0 means the time is unspecified by this router. The maximum value is 3,600,000 (1 hour).

Example

The following command configures the reachable time to be 3,600,000 milliseconds for the VLAN `top_floor`:

```
configure vlan top_floor router-discovery reachable-time 3600000
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vlan router-discovery retransmit-time

```
configure vlan vlan_name router-discovery {ipv6} retransmit-time retransmittime
```



Description

Configures the retransmit time value in router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>retransmittime</i>	Specifies the reachable time value in advertisements, in milliseconds. Range is 0 to 4,294,967,295 (approximately 50 days).

Default

1,000 milliseconds.

Usage Guidelines

This command configures the retransmit time value in the router advertisement messages.

The retransmit time, in milliseconds, is the time between retransmitted neighbor solicitation messages. A value of 0 means the value is unspecified by this router. The maximum value is 4,294,967,295.

Example

The following command configures the retransmit time to be 604,800,000 milliseconds (one week) for the VLAN `top_floor`:

```
configure vlan top_floor router-discovery retransmit-time 604800000
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure vlan router-discovery set prefix

```
configure vlan vlan_name router-discovery {ipv6} set prefix prefix [autonomous-flag auto_on_off | onlink-flag onlink_on_off | preferred-lifetime preflife | valid-lifetime validlife]
```

Description

Sets the parameters for a prefix in the router discovery advertisements on the VLAN.



Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>prefix</i>	Specifies which prefix's parameters to set.
<i>auto_on_off</i>	Specifies the autonomous flag.
<i>onlink_on_off</i>	Specifies the on link flag.
<i>preflife</i>	Specifies the preferred lifetime in seconds. Maximum value is 4,294,967,295.
<i>validlife</i>	Specifies the valid lifetime in seconds. Maximum value is 4,294,967,295.

Default

The prefix parameter defaults are:

- Valid lifetime—2,592,000 seconds (30 days)
- On-link flag—on
- Preferred lifetime—604,800 seconds (7 days)
- Autonomous flag—on

Usage Guidelines

This command configures the attributes associated with the specified prefix.

The autonomous flag option modifies the autonomous flag of the prefix. The autonomous flag value specifies whether the prefix can be used for autonomous address configuration (on) or not (off).

The onlink flag option modifies the on link flag of the prefix. The on link flag specifies whether the prefix can be used for on link determination (on) or not (off). The default value of the on link flag is on.

The preferred lifetime option modifies the preferred lifetime of a prefix. The preferred lifetime value is the time (from when the packet is sent) that addresses generated from the prefix via stateless address autoconfiguration remain preferred. The maximum value is 4,294,967,295. The default value is 604,800 seconds (7 days).

The valid lifetime option modifies the valid lifetime of a prefix. The valid lifetime value is the time (from when the packet was sent) that the prefix is valid for the purpose of on-link determination. The maximum value is a 4,294,967,295. The default value is 2,592,000 seconds (30 days).

Example

The following command sets the on link parameter of the prefix 2001:db8:3161::/64 to off, for the VLAN top_floor:

```
configure vlan top_floor router-discovery set prefix 2001:db8:3161::/64
onlink-flag off
```



History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure tunnel ipaddress

```
configure tunnel tunnel_name ipaddress [ipv6-link-local | {eui64}
ipv6_address_mask ]
```

Description

Configures an IPv6 address/prefix on a tunnel.

Syntax Description

<i>tunnel_name</i>	Specifies an IPv6 tunnel.
ipv6-link-local	Specifies the link-local address for a tunnel.
eui64	Specifies an EUI64 interface identifier for the lower 64 bits of the address.
<i>ipv6_address_mask</i>	Specifies an IPv6 address / IPv6 prefix length.

Default

N/A.

Usage Guidelines

This command will configure an IPv6 address/prefix route on the specified tunnel.

6to4 tunnels must follow the standard address requirement. The address must be of the form `2002:IPv4_source_endpoint::/16`, where *IPv4_source_endpoint* is replaced by the IPv4 source address of the endpoint, in hexadecimal, colon separated form. For example, for a tunnel endpoint located at IPv4 address 10.20.30.40, the tunnel address would be `2002:a14:1e28::/16`. In hex, 10 is a, 20 is 14, 30 is 1e and 40 is 28.

6in4 tunnels have no restrictions on their address format or prefix allocations.



Note

This command does not work for GRE tunnels. The following error message is displayed:
Error: IPv6 addresses can not be configured on GRE type tunnels!



Example

The following command configures the 6in4 tunnel link39 with the IPv6 link-local address:

```
configure tunnel link39 ipaddress ipv6-link-local
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 interworking feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

create tunnel 6to4

```
create tunnel tunnel_name 6to4 source source-address
```

Description

Creates an IPv6-to-IPv4 (6to4) tunnel.

Syntax Description

<i>tunnel_name</i>	Specifies an IPv6 tunnel.
<i>source-address</i>	Specifies an IPv4 address for the tunnel.

Default

N/A.

Usage Guidelines

This command will create a new IPv6-to-IPv4 (also known as a 6to4 tunnel), and add it to the system. A maximum of 1 6to4 tunnel can be configured on any particular VR.

The tunnel name must be unique and cannot overlap the same name space as VLANs, other tunnels, or VRs. The name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see [Object Names](#) in the *ExtremeXOS Concepts Guide*.

The source address of the tunnel must be one of the IPv4 addresses already configured on the switch. You cannot remove an IPv4 address from the switch if a tunnel that uses it still exists.



Example

The following command creates the 6to4 tunnel link35 with source address 192.168.10.1:

```
create tunnel link35 6to4 source 192.168.10.1
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 interworking feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

create tunnel gre destination source

```
create tunnel tunnel_name gre destination destination-address source source-address
```

Description

Allows switch administrators to add a GRE tunnel. This command is in-line with adding an ipv6-in-ipv4 tunnel.

Syntax Description

gre	Generic Routing Encapsulation tunnel.
<i>destination-address</i>	IPv4 destination address of the tunnel.
<i>source-address</i>	IPv4 source address of the tunnel.

Default

No GRE tunnels exist in the system.

Usage Guidelines

Use this command to add a GRE tunnel.

Example

```
create tunnel myGREtunnel gre destination 10.0.0.2 source 10.0.0.1
```



History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X460, X480, X670, X770 and E4G, SummitStack, BD8800 (G96T-c, 10G24X-c, G48T-XL, G48X-XL, 10G8X-XL, and 40G6x-xm), and BlackDiamond X8 (all I/O cards).

create tunnel ipv6-in-ipv4

create tunnel *tunnel_name* ipv6-in-ipv4 destination *destination-address* source *source-address*

Description

Creates an IPv6-in-IPv4 (6in4) tunnel.

Syntax Description

<i>tunnel_name</i>	Specifies an IPv6 tunnel.
<i>source-address</i>	Specifies an IPv4 address for the tunnel.

Default

N/A.

Usage Guidelines

This command will create a new IPv6-in-IPv4 (otherwise known as a configured tunnel or a 6in4 tunnel) and add it to the system. A maximum of 255 tunnels (including one 6to4 tunnel) can be configured on the system.

The tunnel name must be unique and cannot overlap the same name space as VLANs, other tunnels, or VRs. The name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see [Object Names](#) in the ExtremeXOS Concepts Guide.

The source address of the tunnel must be one of the IPv4 addresses already configured on the switch. You cannot remove an IPv4 address from the switch if a tunnel is still exists that uses it.



Example

The following command creates the 6in4 tunnel link39 with destination address 10.10.10.10 and source address 192.168.10.15:

```
create tunnel link39 ipv6-in-ipv4 destination 10.10.10.10 source 192.168.10.15
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 interworking feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

delete tunnel

```
delete tunnel tunnel_name
```

Description

Deletes an IPv6 tunnel.

Syntax Description

<i>tunnel_name</i>	Specifies an IPv6 tunnel.
--------------------	---------------------------

Default

N/A.

Usage Guidelines

This command will destroy a previously created tunnel. The command acts on either a 6to4 or a 6in4 tunnel. When the tunnel interface is removed, all dynamic routes through that interface are purged from the system. The configured static routes are removed from the hardware tables and become inactive.

Example

The following command deletes the tunnel link39:

```
delete tunnel link39
```



History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 interworking feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable icmp redirects ipv6 fast-path

disable icmp redirects ipv6 fast-path

Description

When disabled (default), only slow path packets (packets that cannot be forwarded by hardware) may trigger ICMP redirects.

Syntax Description

fast-path	Only slow path packets (packets that cannot be forwarded by hardware) may trigger ICMP redirects.
------------------	---

Default

Disabled.

Usage Guidelines

Use this command so that only slow path packets (packets that cannot be forwarded by hardware) may trigger ICMP redirects.

Example

The enabled or disabled setting is displayed in the CLI command `show ipconfig ipv6`.

```
BD-8810.1 # show ipconfig ipv6
Route Sharing           : Disabled
ICMP Redirect for Fast Path : Enabled
Max Shared Gateways    : Current: 4   Configured: 4

Interface              IPv6 Prefix
Flags
v1                     2001::1/24                -Euf---R-
v1                     fe80::204:96ff:fe1e:ec00%v1/64 -EufP--R-
Flags : D - Duplicate address detected on VLAN, T - Tentative address
E - Interface enabled, U - Interface up, f - IPv6 forwarding enabled,
i - Accept received router advertisements enabled,
R - Send redirects enabled, r - Accept redirects enabled
```



P - Prefix address
BD-8810.2 #

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on all platforms.

disable ipforwarding ipv6

```
disable ipforwarding ipv6 {vlan vlan_name | tunnel tunnel_name | vr vr_name}
```

Description

Disables routing for one or all interfaces. If no argument is provided, disables routing for all interfaces on the current VR or VRF.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
<i>vr_name</i>	Specifies a VR or VRF.

Default

Disabled.

Usage Guidelines

When new IPv6 interfaces are added, IPv6 forwarding is disabled by default.

Extreme Networks switches have a single hardware control per VLAN for IPv6 forwarding of IPv4 and IPv6 unicast packets. Therefore, enabling IPv6 forwarding on a VLAN also enables IPv4 hardware forwarding on that VLAN.

Example

The following command disables forwarding of IPv6 traffic for a VLAN named accounting:

```
disable ipforwarding ipv6 vlan accounting
```



History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the *Feature License Requirements* in the *ExtremeXOS Concepts Guide*.

disable iproute ipv6 compression

```
disable iproute ipv6 compression {vr vr_name}
```

Description

This command disables IPv6 route compression.

Syntax Description

<code>vr_name</code>	Specifies a VR or VRF.
----------------------	------------------------

Default

Vrname—current CLI context VR

By default, IPv6 route compression is disabled for all address families and VRs.

Usage Guidelines

This command disables IPv6 route compression for the IPv6 address family and VR. This command decompresses previously compressed prefixes in the IPv6 prefix database.

Example

The following example disables IPv6 route compression for the IPv6 address family and the VR of the current CLI context.

```
disable iproute ipv6 compression
```

History

This command was first available in ExtremeXOS 12.0.



Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable iproute ipv6 sharing

```
disable iproute ipv6 sharing {{{vr} vr_name} | {{{vr} all}}}
```

Description

This command disables IPv6 route sharing.

Syntax Description

<code>vr_name</code>	Specifies a VR or VRF.
<code>all</code>	Specifies all VR or VRF.

Default

`vrname`—current CLI context VR

By default, IPv6 route sharing is disabled.

Usage Guidelines

This command disables IPv6 route sharing for the IPv6 address family and VR.

Example

The following example disables IPv6 route sharing for the IPv6 address family and the VR of the current CLI context.

```
disable iproute ipv6 sharing
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

The ability to enable and disable ECMP for IPv6 is now supported for all Summit and BlackDiamond, except for the Summit X440.



disable neighbor-discovery refresh

```
disable neighbor-discovery {vr vr_name} refresh
```

Description

Prevents the IPv6 neighbor cache from refreshing an entry before the timeout period expires.

Syntax Description

<code>vr_name</code>	Specifies a VR or VRF.
----------------------	------------------------

Default

Enabled.

Usage Guidelines

None.

Example

The following command disables the refresh of neighbor discovery cache entries:

```
disable neighbor-discovery refresh
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

disable router-discovery

```
disable router-discovery {ipv6} vlan vlan_name
```

Description

Disables router discovery advertisements on the VLAN and the processing of router discovery messages.



Syntax Description

<code>vlan_name</code>	Specifies an IPv6 configured VLAN.
------------------------	------------------------------------

Default

N/A.

Usage Guidelines

None.

Example

The following command disables router discovery for the VLAN `top_floor`:

```
disable router-discovery vlan top_floor
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable tunnel

```
disable {tunnel} tunnel_name
```

Description

Allows GRE tunnels to be disabled.

Syntax Description

<code>tunnel_name</code>	GRE tunnel name.
--------------------------	------------------

Default

Enabled.



Usage Guidelines

Use this command to disable GRE tunnels.

Example

```
disable myGREtunnel
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X460, X480, X670, X770, and E4G, SummitStack, BD8800 (G96T-c, 10G24X-c, G48T-XL, G48X-XL, 10G8X-XL, and 40G6x-xm), and BlackDiamond X8 (all I/O cards).

enable icmp redirects ipv6 fast-path

```
enable icmp redirects ipv6 fast-path
```

Description

When enabled, IPv6 packets forwarded by hardware (fast path) may trigger ICMP redirects.

Syntax Description

fast-path	IPv6 packets forwarded by hardware may trigger ICMP redirects
------------------	---

Default

Disabled.

Usage Guidelines

Use this command to trigger ICMP redirects when IPv6 packets are forwarded by hardware (fast-path).

Example

The enabled or disabled setting is displayed in the CLI command `show ipconfig ipv6`.

```
BD-8810.1 # show ipconfig ipv6
Route Sharing           : Disabled
ICMP Redirect for Fast Path : Enabled
Max Shared Gateways    : Current: 4   Configured: 4
```



```

Interface          IPv6 Prefix
Flags
v1                 2001::1/24                -Euf---R-
v1                 fe80::204:96ff:fe1e:ec00%v1/64 -EufP--R-
Flags : D - Duplicate address detected on VLAN, T - Tentative address
E - Interface enabled, U - Interface up, f - IPv6 forwarding enabled,
i - Accept received router advertisements enabled,
R - Send redirects enabled, r - Accept redirects enabled
P - Prefix address
BD-8810.2 #

```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on all platforms.

enable ipforwarding ipv6

```
enable ipforwarding ipv6 {vlan vlan_name | tunnel tunnel_name | vr vr_name}
```

Description

Enables IPv6 routing VLANs. If no argument is provided, enables IPv6 routing for all VLANs and tunnels that have been configured with an IPv6 address on the current VR or VRF.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
<i>vr_name</i>	Specifies a VR or VRF.

Default

Disabled.

Usage Guidelines

When new IPv6 interfaces are added, IPv6 forwarding is disabled by default.

Extreme Networks switches have a single hardware control per VLAN for forwarding of IPv4 and IPv6 unicast packets. Therefore, enabling IPv6 forwarding on a VLAN also enables IPv4 hardware forwarding on that VLAN.



Example

The following command enables forwarding of IPv6 traffic for all VLANs in the current VR context with IPv6 addresses:

```
enable ipforwarding ipv6
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

enable ipforwarding

enable ipforwarding

Description

Allows switch administrators to enable IPv4 forwarding on a GRE tunnel. This command is in-line with the existing command to configure IPv4 forwarding on a VLAN.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Use this command to enable IPv4 forwarding on a GRE tunnel.

Example

```
enable ipforwarding ipv4 myGREtunnel
```

History

This command was first available in ExtremeXOS 15.3.



Platform Availability

This command is available on Summit X460, X480, X670, X770, and E4G, SummitStack, BD8800 (G96T-c, 10G24X-c, G48T-XL, G48X-XL, 10G8X-XL, and 40G6x-xm), and BlackDiamond X8 (all I/O cards).

enable iproute ipv6 compression

This command enables IPv6 route compression.

```
enable iproute ipv6 compression {vr vrname}
```

Syntax Description

<i>vrname</i>	Specifies a VR or VRF.
---------------	------------------------

Default

Vrname—current CLI context VR.

Usage Guidelines

This command enables IPv6 route compression for the VR. This command applies a compression algorithm to each IPv6 prefix in the IPv6 prefix database.

Example

The following example enables IPv6 route compression in the current VR context.

```
enable iproute ipv6 compression
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

enable iproute ipv6 sharing

```
enable iproute ipv6 sharing {{{vr} vr_name} | { {vr} all}}
```



Description

This command enables IPv6 route sharing.

Syntax Description

<i>vr_name</i>	Specifies a VR or VRF.
all	Specifies all VR or VRF.

Default

vrname—current CLI context VR.

By default, IPv6 route sharing is disabled.

Usage Guidelines

This command enables IPv6 route sharing for the IPv6 address family and VR.

Example

The following example enables IPv6 route sharing for the IPv6 address family and the VR of the current CLI context.

```
enable iproute ipv6 sharing
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

The ability to enable and disable ECMP for IPv6 is now supported for all Summit and BlackDiamond, except for the Summit X440. If present in a SummitStack, any Summit X440 nodes ingressing IPv6 unicast packets will still perform unipath IPv6 forwarding in hardware.

enable neighbor-discovery refresh

```
enable neighbor-discovery {vr vr_name} refresh
```

Description

Enables the IPv6 neighbor cache to refresh each entry before the timeout period expires.



Syntax Description

<code>vr_name</code>	Specifies a VR or VRF.
----------------------	------------------------

Default

Enabled.

Usage Guidelines

None.

Example

The following command enables the refresh of neighbor discovery cache entries:

```
enable neighbor-discovery refresh
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable router-discovery

```
enable router-discovery {ipv6} vlan vlan_name
```

Description

Enables router discovery advertisements on the VLAN and the processing of router discovery messages.

Syntax Description

<code>vlan_name</code>	Specifies an IPv6 configured VLAN.
------------------------	------------------------------------

Default

N/A.



Usage Guidelines

This command is only valid when the specified VLAN has an IPv6 address associated with it. After IPv6 Router Discovery is enabled on a VLAN, router advertisement messages are regularly sent on all ports associated with the VLAN.

Example

The following command enables router discovery for the VLAN `top_floor`:

```
enable router-discovery vlan top_floor
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

enable tunnel

```
enable {tunnel} tunnel_name
```

Description

Allows GRE tunnels to be enabled.

Syntax Description

<i>tunnel_name</i>	GRE tunnel name
--------------------	-----------------

Default

Enabled.

Usage Guidelines

Use this command to enable GRE tunnels.

Example

```
enable myGREtunnel
```



History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on Summit X460, X480, X670, and E4G, SummitStack, BD8800 (G96T-c, 10G24X-c, G48T-XL, G48X-XL, 10G8X-XL, and 40G6x-xm), and BlackDiamond X8 (all I/O cards).

rtlookup

```
rtlookup [ipaddress | ipv6address] { unicast | multicast | vr vr_name}
```

Description

Displays the available routes to the specified IPv6 address.

Syntax Description

<i>ipaddress</i>	Specifies an IPv4 address.
<i>ipv6address</i>	Specifies an IPv6 address.
unicast	Displays the routes from the unicast routing table in the current router context.
multicast	Displays the routes from the multicast routing table in the current router context.
<i>vr_name</i>	Specifies the VR or VRF for which to display the route.

Default

N/A.

Usage Guidelines

None.

Example

The following command performs a look up in the route table to determine the best way to reach the specified IPv6 address:

```
rtlookup 2001:db8::ef80:2525:1023:5213 unicast
```

History

This command was first available in ExtremeXOS 10.1.



The xhostname option was removed in ExtremeXOS 11.0.

Support for IPv6 was added in ExtremeXOS 11.2.

The unicast and multicast options were added in ExtremeXOS 12.1.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

rtlookup rpf

```
rtlookup [ipaddress | ipv6address] rpf {vr vr_name}
```

Description

Displays the RPF for a specified multicast source.

Syntax Description

<i>ipaddress</i>	Specifies an IPv4 address.
<i>ipv6address</i>	Specifies an IPv6 address.
rpf	Selects the RPF for the specified multicast source.
<i>vr_name</i>	Specifies the VR or VRF for which to display the route.

Default

vr_name is the VR of the current CLI context.

Usage Guidelines

None.

Example

The following example displays the RPF lookup for a multicast source through VR-Default:

```
rtlookup 2001db8::ef80:2525:1023:5213 rpf vr vr-default
```

History

This command was first available in ExtremeXOS 12.1.



Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

run ipv6 dad

```
run ipv6 dad [{vlan} vlan_name | {{vr} vr_name} ipaddress]
```

Description

Runs the DAD check on the specified IPv6 interface for which the DAD feature is enabled.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN on which to run the test.
<i>vr_name</i>	Specifies a VR on which to run the test.
<i>ipaddress</i>	Specifies an IPv6 address for which to run the test.

Default

If you do not specify a VR or VRF, the command applies to the current VR context.

Usage Guidelines

To run the test, you must specify a VLAN name or a specific IPv6 address. To support the test, the DAD feature must be enabled on the parent VR for the specified VLAN or IPv6 interface.

This command is ignored for the following conditions:

- The specified IP address is in tentative state
- DAD is configured to be off
- The host VLAN is disabled
- Loopback mode is enabled on the host VLAN
- DAD is already running due to interface initialization or a previous issue of this command
- The host VLAN belongs to virtual router VR-Mgmt

If a duplicate address is detected during the test, the event is logged and the address remains valid if it was already valid. If the address was not already valid, the event is logged and the duplicate address transitions to duplicate state. If no duplicate IPv6 address is detected, the specified IPv6 interface transitions to or remains in the up state.



Example

The following command runs the DAD check on the IPv6 interfaces in VLAN vlan1:

```
run ipv6 dad vlan1
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on BlackDiamond X8, BlackDiamond 8800 series switches and Summit family switches.

show ipconfig ipv6

```
show ipconfig ipv6 {vlan vlan_name | tunnel tunnel_name}
```

Description

Displays configuration information for one or more interfaces in the current VR context.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel_name</i>	Specifies an IPv6 tunnel.

Default

N/A.

Usage Guidelines

If no interface is specified, then global IPv6 configuration is displayed. Otherwise, specific interface(s) will be displayed. Global IPv6 configuration information includes:

- IPv6 address/netmask/etc.
- IPv6 forwarding information / IPv6 multicast forwarding information

Example

The following command displays configuration information on a VLAN named accounting:

```
show ipconfig ipv6 vlan accounting
```



The current and configured values for **max-gateways** now apply to IPv6 gateway sets as well as IPv4, so these values are added to the output of `show ipconfig ipv6`.

```
show ipconfig ipv6
Route Sharing           : Disabled
ICMP Redirect for Fast Path : Disabled
Max Shared Gateways    : Current: 32  Configured: 32
```

History

This command was first available in ExtremeXOS 11.2.

This command changed to display information for the current VR context in ExtremeXOS 12.5.

The show output for **max-gateways** was added in ExtremeXOS 15.3.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show iproute ipv6

```
show iproute ipv6 {priority | vlan vlan_name | tunnel tunnel_name | ipv6Netmask |
summary {multicast | unicast}} {vr vr_name}}
```

Description

Displays the contents of the IPv6 routing table.

Syntax Description

priority	Displays the priority values for each route origin type.
<i>vlan_name</i>	Specifies a VLAN name.
<i>tunnel_name</i>	Specifies a tunnel name.
<i>ipv6Netmask</i>	Specifies an IPv6 address/prefix length.
summary	Specifies summary information.
<i>vr_name</i>	Specifies a VR or VRF.

Default

N/A.

Usage Guidelines

If you do not specify a VR or VRF, the command applies to the current VR context.



Example

The following command displays detailed information about all IPv6 routing:

```
Switch.18 # show iproute ipv6
Ori Destination          Mtr  Flags          Duration
Gateway                 Interface
#d 2001:db8::/64        1    U-----um--f 0d:0h:5m:
31s
2001:db8::52           ixia
#or 2001:db8:2:78::/64 50    UG-D---um--f 0d:0h:0m:
1s
fe80::200:40ff:feba:a38e ixia
#or 2001:db8:2:79::/64 50    UG-D---um--f 0d:0h:0m:
1s
fe80::200:40ff:feba:a38e ixia
#or 2001:db8:2:7a::/64 50    UG-D---um--f 0d:0h:0m:
1s
fe80::200:40ff:feba:a38e ixia
#or 2001:db8:2:7b::/64 50    UG-D---um--f 0d:0h:0m:
1s
fe80::200:40ff:feba:a38e ixia
#or 2001:db8:2:7c::/64 50    UG-D---um--f 0d:0h:0m:
1s
fe80::200:40ff:feba:a38e ixia
#or 2001:db8:2:7d::/64 50    UG-D---um--f 0d:0h:0m:
1s
fe80::200:40ff:feba:a38e ixia
#or 2001:db8:2:7e::/64 50    UG-D---um--f 0d:0h:0m:
1s
fe80::200:40ff:feba:a38e ixia
#or 2001:db8:2:7f::/64 50    UG-D---um--f 0d:0h:0m:
1s
fe80::200:40ff:feba:a38e ixia
#or 2001:db8:2:80::/64 50    UG-D---um--f 0d:0h:0m:
1s
fe80::200:40ff:feba:a38e ixia
#or 2001:db8:2:81::/64 50    UG-D---um--f 0d:0h:0m:
1s
#d fe80::%ixia/64      1    U-----um--f 0d:0h:5m:
31s
fe80::204:96ff:fe27:8697 ixia
Origin(Ori): (b) BlackHole, (be) EBGp, (bg) BGP, (bi) IBGP, (bo) BOOTP
(ct) CBT, (d) Direct, (df) DownIF, (dv) DVMRP, (e1) ISISL1Ext
(e2) ISISL2Ext, (h) Hardcoded, (i) ICMP, (i1) ISISL1 (i2) ISISL2
(is) ISIS, (mb) MBGP, (mbe) MBGPExt, (mbi) MBGPInter, (ma) MPLSIntra
(mr) MPLSInter, (mo) MOSPF (o) OSPFv3, (o1) OSPFv3Ext1, (o2) OSPFv3Ext2
(oa) OSPFv3Intra, (oe) OSPFv3AsExt, (or) OSPFv3Inter, (pd) PIM-DM, (ps) PIM-SM
(r) RIPng, (ra) RtAdvrt, (s) Static, (sv) SLB_VIP, (un) UnKnown
(*) Preferred unicast route (@) Preferred multicast route
(#) Preferred unicast and multicast route
Flags: (B) BlackHole, (D) Dynamic, (G) Gateway, (H) Host Route
(L) Matching LDP LSP, (l) Calculated LDP LSP, (m) Multicast
(P) LPM-
routing, (R) Modified, (S) Static, (s) Static LSP
(T) Matching RSVP-TE LSP (t) Calculated RSVP-TE LSP, (u) Unicast, (U) Up
```



```
(f) Provided to FIB (c) Compressed Route
Mask distribution:
12 routes at length 64
Route Origin distribution:
2 routes from Direct          10 routes from OSPFv3Inter
Total number of routes = 12
Total number of compressed routes = 0
```

The following command displays the IPv6 route origin priority:

```
Switch.4 # show iproute ipv6 priority
Direct          10
Blackhole       50
Static          1100
ICMP            1200
OSPFv3Intra    2200
OSPFv3Inter    2300
Isis            2350
IsisL1         2360
IsisL2         2370
RIPng          2400
OSPFv3AsExt    3100
OSPFv3Ext1     3200
OSPFv3Ext2     3300
IsisL1Ext      3400
IsisL2Ext      3500
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

show iproute ipv6 origin

```
show iproute ipv6 origin [direct | static | blackhole | ripng | ospfv3 | ospfv3-
intra | ospfv3-inter | ospfv3-extern1 | ospfv3-extern2 | isis | isis-level-1 |
isis-level-2 | isis-level-1-external | isis-level-2-external] {vr vr_name}
```

Description

Displays the contents of the IPv6 routing table for routes with the specified origin.



Syntax Description

direct	Specifies direct routes.
static	Specifies static routes.
blackhole	Specifies blackhole routes.
ripng	Specifies RIPng routes.
ospfv3	Specifies OSPFv3 routes.
ospfv3-intra	Specifies OSPFv3 Intra routing.
ospfv3-inter	Specifies OSPFv3 Inter routing.
ospfv3-extern1	Specifies OSPFv3 External 1 routing.
ospfv3-extern2	Specifies OSPFv3 External 2 routing.
isis	Specifies IS-IS routes.
isis-level-1	Specifies ISIS Level 1 routing.
isis-level-2	Specifies ISIS Level 2 routing.
isis-level-1-external	Specifies ISIS Level 1 External routing.
isis-level-2-external	Specifies ISIS Level 2 External routing.
<i>vr_name</i>	Specifies a VR or VRF.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the RIPng routes:

```
show iproute ipv6 origin ripng
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



show ipstats ipv6

```
show ipstats ipv6 {vlan name | tunnel tunnel_name | vr vr_name}
```

Description

Displays IPv6 statistics for the CPU for the switch or for a particular VLAN.

Syntax Description

<i>name</i>	Specifies a VLAN name.
<i>tunnel_name</i>	Specifies a tunnel name.
<i>vr_name</i>	Specifies a VR or VRF.

Default

N/A.

Usage Guidelines

This command only shows statistics of the CPU-handled packets. Not all packets are handled by the CPU. For example, packets forwarded in hardware do not increment the statistics counters.

If you do not specify a VR or VRF, the command applies to the current VR context.

Example

The following command displays IPv6 statistics for the VLAN accounting:

```
show ipstats ipv6 vlan accounting
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show ipv6 dad

```
show ipv6 dad [{vr vr_name ip_address} | vr all | {vlan} vlan_name]
{tentative | valid | duplicate} | [{vr vr_name} ipaddress] | {tunnel}
tunnel_name}
```



Description

Displays the configuration and run time status for the DAD feature on the specified IPv6 interface.

Syntax Description

<i>vr_name</i>	Specifies a VR for which to display the DAD information.
<i>ip_address</i>	Specifies an IPv6 address for which to display the DAD information.
<i>vlan_name</i>	Specifies a VLAN for which to display the DAD information.
tentative	Displays information for IPv6 interfaces for which the status is up and the DAD check is incomplete.
valid	Displays information for IPv6 interfaces for which the status is up, the DAD check is complete, and no duplicate IPv6 addresses were detected.
duplicate	Displays information for IPv6 interfaces for which the status is down because a duplicate IPv6 address was detected.

Default

If you do not specify a VR or VRF, the command applies to the current VR context.

Usage Guidelines

The `vr all` option displays DAD information for all IPv6 interfaces on the switch.

Example

The following command displays the DAD feature status for all interfaces in the current VR context:

```
show ipv6 dad
IPv6 Duplicate Address Detection
DAD Status           : On
Max Solicitation Attempts : 1
Virtual Router       Interface      Flags    IP Address
Conflict MAC        Failures
-----
--
VR-Default          exG-v154      -E-U     2001:db8:a::123/64
00:00:00:00:00:00    0
VR-Default          exL-v188      DE-U     2001:db8:b::123/64
00:04:96:12:ae:60   1
Flags : (D) Duplicate address detected, (T) Tentative address,
(E) Interface enabled, (L) Loopback enabled, (U) Interface up
```

History

This command was first available in ExtremeXOS 12.6.



Platform Availability

This command is available on BlackDiamond X8, BlackDiamond 8800 series switches and Summit family switches.

show neighbor-discovery cache ipv6

```
show neighbor-discovery {cache {ipv6}} {[ipv6_addr | mac | permanent] {vr
vr_name}} | vlan vlan_name | vr vr_name}
```

Description

This command displays all the entries from the neighbor cache.

Syntax Description

<i>ipv6_addr</i>	Specifies an IPv6 address.
<i>mac</i>	Specifies a MAC address.
permanent	Specifies static entries.
<i>vr_name</i>	Specifies a VR or VRF.
<i>vlan_name</i>	Specifies an IPv6 configured VLAN.

Default

N/A.

Usage Guidelines

This command displays the entries present in the neighbor cache.

The entries displayed can be filtered by IPv6 address, MAC address, or by VLAN. The permanent keyword filters the output to display static entries.

The *vr_name* indicates the VR or VRF on which the operation is performed. In its absence, the operation applies to VR-Default.

Example

The following command shows all entries from the neighbor cache:

```
show neighbor-discovery cache ipv6
```

The following is sample output:

```
VR          Destination
Mac         Age  Static  VLAN          VID  Port
```




```

VR-Default    2001:db8:100::7
00:01:30:00:6b:00    0    NO    gtag100    100    1:2
VR-Default    2001:db8:100::99
00:01:02:33:33:33    0    YES   gtag100    100
VR-Default    2001:db8:99::99
00:01:02:01:01:01    0    YES   gtag99     99
Total Entries   :    0
Dynamic Entries :    0                Static Entries   :    0
Pending Entries :    0
Max Entries     :   1024                Max Pending entries :   1024
Timeout        :    20 minutes          Refresh         :    Enable

```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show router-discovery

```
show router-discovery {ipv6} {vlan vlan_name}
```

Description

Displays the router discovery settings.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
------------------	------------------------------------

Default

N/A.

Usage Guidelines

If no VLAN is specified, the settings are displayed for all IPv6 configured VLANs.

Example

The following command displays router discovery settings for the VLAN `top_floor`:

```
show router-discovery vlan top_floor
```



The following is sample output:

```
Router Advertisements disabled on v1
Minimum/Maximum Interval: 200 / 600
Managed / Other Info Flags: Off / Off
Link MTU: 0
Reachable Time: 0
Retrans Timer: 0
Current Hop Limit: 64
Default Lifetime: 1800
Number of Prefixes: 1, Prefix List:
Valid          Preferred
Prefix Lifetime Auto Lifetime OnLink
2001:db8::1/64
2592000        On    604800    On
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show tunnel

```
show [ {tunnel} {tunnel_name} ]
```

Description

Displays system tunnel information for a specified tunnel or for all tunnels.

Syntax Description

<i>tunnel_name</i>	Specifies an IPv6-in-IPv4 or IPv6-to-IPv4 tunnel name.
--------------------	--

Default

N/A.

Usage Guidelines

The tunnel keyword is optional only when you specify a valid IPv6-in-IPv4 or IPv6-to-IPv4 tunnel name. The Total tunnels count in the display represents all tunnels on the switch.



Example

The following command displays system tunnel information for all tunnels:

```
Switch.1 # show tunnel
Name                               Type                               Flags
tunfour                            6in4 10.20.30.40 => 10.10.10.10    U
mytun                               GRE 1.1.1.2 => 1.1.1.1
Utunfive2                          6to4 10.20.30.40 => *.*.*.*      D
Total tunnels: 3
Flags: (U) Up / (D) Down / (a) Administratively Disabled
       (S) System Disabled (incompatible hardware)
```

The following command displays system tunnel information for tunnel tunfour:

```
Switch.3 # show "tunfour"
Name                               Type                               Flags
tunfour                            6in4 10.20.30.40 => 10.10.10.10    U
Total tunnels: 2
Flags: (U) Up / (D) Down
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 interworking feature in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure iproute ipv6 priority

```
unconfigure iproute ipv6 priority [all | blackhole | icmp | isis | isis-level-1 |
isis-level-1-external | isis-level-2 | isis-level-2-external | ospfv3-as-external
| ospfv3-extern1 | ospfv3-extern2 | ospfv3-inter | ospfv3-intra | ripng | static]
{vr vr_name}
```

Description

Resets the priority for all IPv6 routes from one or all route origin types to the default values.

Syntax Description

all	Specifies all route origins.
blackhole	Specifies the blackhole route.
icmp	Specifies ICMP.



isis	Specifies IS-IS and applies only to blackhole routes installed for summary addresses.
isis-level-1	Specifies ISIS Level 1 routing.
isis-level-1-external	Specifies ISIS Level 1 External routing.
isis-level-2	Specifies ISIS Level 2 routing.
isis-level-2-external	Specifies ISIS Level 2 External routing.
ospfv3-as-external	Specifies OSPF as External routing.
ospfv3-extern1	Specifies OSPF External 1 routing.
ospfv3-extern2	Specifies OSPF External 2 routing.
ospf3-inter	Specifies OSPFInter routing.
ospf3-intra	Specifies OSPFIntra routing.
ripng	Specifies RIP.
static	Specifies static routes.
<i>vr_name</i>	Specifies a VR or VRF name.

Default

N/A.

Usage Guidelines

The following table lists the default values that apply after you enter this command.

Default Route Priorities

Route Origin	Priority
Direct	10
BlackHole	50
Static	1100
ICMP	1200
OSPF3Intra	2200
OSPF3Inter	2300
IS-IS L1	2360
IS-IS L2	2370
RIPg	2400
OSPFv3 AExt	3100
OSPFv3 Extern1	3200
OSPFv3 Extern2	3300



Route Origin	Priority
IS-IS L1 Ext	3400
IS-IS L2 Ext	3500

Example

The following command returns the IPv6 route priority for all route origins to the default values:

```
unconfigure iproute ipv6 priority all
```

History

This command was first available in ExtremeXOS 12.1.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure neighbor-discovery cache

```
unconfigure neighbor-discovery cache {vr vr_name}
```

Description

Resets the neighbor-discovery cache configuration parameters to their default values.

Syntax Description

<i>vr_name</i>	Specifies a VR or VRF.
----------------	------------------------

Default

IPv6 neighbor timeout: 20 minutes

Maximum IPv6 neighbor entries: 1024

Maximum IPv6 neighbor pending entries: 1024

IPv6 neighbor refresh: Enabled

Usage Guidelines

None.



Example

The following command resets the neighbor-discovery cache configuration:

```
unconfigure neighbor-discovery cache
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure vlan router-discovery

```
unconfigure vlan vlan_name router-discovery {ipv6}
```

Description

Unconfigures all the router-discovery parameters and resets them to their respective default values.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
------------------	------------------------------------

Default

N/A.

Usage Guidelines

Each of the router-discovery parameters is set to the default value. For example, the default-lifetime parameter is set to 1800 seconds. The default value for each of the router-discovery parameters is listed in the corresponding configure vlan router-discovery command description.

Example

The following command unconfigures all the router-discovery parameters for the VLAN top_floor:

```
unconfigure vlan top_floor router-discovery
```



History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure vlan router-discovery default-lifetime

```
unconfigure vlan vlan_name router-discovery {ipv6} default-lifetime
```

Description

Unconfigures the router lifetime value sent in router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
------------------	------------------------------------

Default

N/A.

Usage Guidelines

This command sets the default-lifetime parameter to the default value of 1800 seconds.

Example

The following command unconfigures the default-lifetime for the VLAN `top_floor`:

```
unconfigure vlan top_floor router-discovery default-lifetime
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



unconfigure vlan router-discovery hop-limit

```
unconfigure vlan vlan_name router-discovery {ipv6} hop-limit
```

Description

Unconfigures the current hop limit value sent in router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
------------------	------------------------------------

Default

N/A.

Usage Guidelines

This command sets the hop-limit parameter to the default value of 64.

Example

The following command unconfigures the current hop limit for the VLAN `top_floor`:

```
unconfigure vlan top_floor router-discovery hop-limit
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

unconfigure vlan router-discovery link-mtu

```
unconfigure vlan vlan_name router-discovery {ipv6} link-mtu
```

Description

Unconfigures the link MTU value sent in router discovery advertisements on the VLAN.



Syntax Description

<code>vlan_name</code>	Specifies an IPv6 configured VLAN.
------------------------	------------------------------------

Default

N/A.

Usage Guidelines

This command sets the link-mtu parameter to the default value of 0.

Example

The following command unconfigures the link MTU for the VLAN `top_floor`:

```
unconfigure vlan top_floor router-discovery link-mtu
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure vlan router-discovery managed-config-flag

```
unconfigure vlan vlan_name router-discovery {ipv6} managed-config-flag
```

Description

Unconfigures the managed address configuration flag value sent in router discovery advertisements on the VLAN.

Syntax Description

<code>vlan_name</code>	Specifies an IPv6 configured VLAN.
------------------------	------------------------------------

Default

N/A.



Usage Guidelines

This command sets the managed-config-flag parameter to the default value off.

Example

The following command unconfigures the managed address configuration flag for the VLAN top_floor:

```
unconfigure vlan top_floor router-discovery managed-config-flag
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

unconfigure vlan router-discovery max-interval

```
unconfigure vlan vlan_name router-discovery {ipv6} max-interval
```

Description

Unconfigures the maximum time between unsolicited router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
------------------	------------------------------------

Default

N/A.

Usage Guidelines

This command sets the max-interval parameter to the default value of 600 seconds.

Example

The following command unconfigures the max-interval for the VLAN top_floor:

```
unconfigure vlan top_floor router-discovery max-interval
```



History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure vlan router-discovery min-interval

```
unconfigure vlan vlan_name router-discovery {ipv6} min-interval
```

Description

Unconfigures the minimum time between unsolicited router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
------------------	------------------------------------

Default

N/A.

Usage Guidelines

This command sets the min-interval parameter to the default value of (max-interval * .33 seconds).

Example

The following command unconfigures the min-interval for the VLAN top_floor:

```
unconfigure vlan top_floor router-discovery min-interval
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



unconfigure vlan router-discovery other-config-flag

```
unconfigure vlan vlan_name router-discovery {ipv6} other-config-flag
```

Description

Unconfigures the other stateful configuration flag value sent in router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
------------------	------------------------------------

Default

N/A.

Usage Guidelines

This command sets the other-config-flag parameter to the default value off.

Example

The following command unconfigures the other stateful configuration flag for the VLAN top_floor:

```
unconfigure vlan top_floor router-discovery other-config-flag
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure vlan router-discovery reachable-time

```
unconfigure vlan vlan_name router-discovery {ipv6} reachable-time
```

Description

Unconfigures the reachable time value in router discovery advertisements on the VLAN.



Syntax Description

<code>vlan_name</code>	Specifies an IPv6 configured VLAN.
------------------------	------------------------------------

Default

N/A.

Usage Guidelines

This command sets the reachable-time parameter to the default value of 30,000 milliseconds.

Example

The following command unconfigures the reachable time for the VLAN `top_floor`:

```
unconfigure vlan top_floor router-discovery reachable-time
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure vlan router-discovery retransmit-time

```
unconfigure vlan vlan_name router-discovery {ipv6} retransmit-time
```

Description

Unconfigures the retransmit time value in router discovery advertisements on the VLAN.

Syntax Description

<code>vlan_name</code>	Specifies an IPv6 configured VLAN.
------------------------	------------------------------------

Default

N/A.



Usage Guidelines

This command sets the retransmit-time parameter to the default value of 1000 milliseconds.

Example

The following command unconfigures the retransmit time for the VLAN top_floor:

```
unconfigure vlan top_floor router-discovery retransmit-time
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

unconfigure tunnel

```
unconfigure tunnel tunnel_name ipaddress ipv6_address_mask
```

Description

Unconfigures an IPv6 address/prefix route from a tunnel.

Syntax Description

<i>tunnel_name</i>	Specifies an IPv6 tunnel.
<i>ipv6_address_mask</i>	Specifies an IPv6 address / IPv6 prefix length.

Default

N/A.

Usage Guidelines

Use this command to unconfigure an IPv6 address/prefix route from the specified tunnel.



Example

The following command unconfigures the 6in4 tunnel link39 with the address 2001:db8::1111/64

```
unconfigure tunnel link39 2001:db8::1111/64
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 interworking feature in Feature License Requirements in the *ExtremeXOS Concepts Guide*.



36 RIP Commands

```
clear rip counters
configure rip add vlan
configure rip delete vlan
configure rip garbage-time
configure rip import-policy
configure rip route-timeout
configure rip update-time
configure rip vlan cost
configure rip vlan route-policy
configure rip vlan rx-mode
configure rip vlan trusted-gateway
configure rip vlan tx-mode
disable rip
disable rip aggregation
disable rip export
disable rip poison-reverse
disable rip split-horizon
disable rip trigger-updates
disable rip use-ip-router-alert
enable rip
enable rip aggregation
enable rip export
enable rip originate-default cost
enable rip poison-reverse
enable rip split-horizon
enable rip trigger-updates
enable rip use-ip-router-alert
show rip
show rip interface
show rip interface vlan
show rip memory
show rip routes
unconfigure rip
```

This chapter describes commands used for the interior gateway protocol RIP. Routing Information Protocol (RIP) is an Interior Gateway Protocol (IGP) first used in computer routing in the Advanced Research Projects Agency Network (ARPAnet) as early as 1969. It is primarily intended for use in homogeneous networks of moderate size.

To determine the best path to a distant network, a router using RIP always selects the path that has the least number of hops. Each router that data must traverse is considered to be one hop.

The routing table in a router using RIP contains an entry for every known destination network. Each routing table entry contains the following information:

- IP address of the destination network.
- Metric (hop count) to the destination network.
- IP address of the next router.
- Timer that tracks the amount of time since the entry was last updated.

The router exchanges an update message with each neighbor every 30 seconds (default value), or if there is a change to the overall routed topology (also called triggered updates). If a router does not receive an update message from its neighbor within the route timeout period (180 seconds by default), the router assumes the connection between it and its neighbor is no longer available.

A new version of RIP, called RIP version 2 (RIPv2), expands the functionality of RIP version 1 to include:

- Variable-Length Subnet Masks (VLSMs).
- Next-hop addresses.
- Support for next-hop addresses allows for optimization of routes in certain environments.
- Multicasting.

If you are using RIP with supernetting/Classless Inter-Domain Routing (CIDR), you must use RIPv2 only, and RIP route aggregation must be turned off.



Note

RIP is available on all platforms with an Edge, Advanced Edge, or Core license. For licensing information, refer to [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*,

clear rip counters

```
clear rip counters
```

Description

Clears the RIP counters (statistics).

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.



Example

The following command clears the RIP statistics counters:

```
clear rip counters
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

configure rip add vlan

```
configure rip add vlan [vlan_name | all]
```

Description

Configures RIP on an IP interface.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

When an IP interface is created, RIP configuration is disabled on the interface by default. When the RIP interface is disabled, the parameters are not reset to default automatically.

Example

The following command configures RIP on the VLAN finance:

```
configure rip add finance
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

configure rip delete vlan

```
configure rip delete vlan [vlan_name | all]
```

Description

Disables RIP on an IP interface.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

When an IP interface is created, RIP configuration is disabled on the interface by default. When the RIP interface is disabled by this command, the parameters are not reset to default automatically.

Example

The following command deletes RIP on a VLAN named finance:

```
configure rip delete finance
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.



configure rip garbagetime

```
configure rip garbagetime {seconds}
```

Description

Configures the RIP garbage time.

Syntax Description

<i>seconds</i>	Specifies a time in seconds.
----------------	------------------------------

Default

120 seconds.

Usage Guidelines

None.

Example

The following command configures the RIP garbage time to have a 60-second delay:

```
configure rip garbagetime 60
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

configure rip import-policy

```
configure rip import-policy [policy-name | none]
```

Description

Configures the import policy for RIP.



Syntax Description

<i>policy-name</i>	Specifies the policy.
--------------------	-----------------------

Default

No policy.

Usage Guidelines

An import policy is used to modify route attributes while adding RIP routes to the IP route table. The import policy cannot be used to determine the routes to be added to the routing table.

Use the none option to remove an import policy.

Example

The following example applies the policy campuseast to RIP routes:

```
configure rip import-policy campuseast
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

configure rip routetimeout

```
configure rip routetimeout seconds
```

Description

Configures the route timeout period.

Syntax Description

<i>seconds</i>	Specifies a time in seconds.
----------------	------------------------------

Default

180 seconds.



Usage Guidelines

If a router does not receive an update message from its neighbor within the route timeout period (180 seconds by default), the router assumes the connection between it and its neighbor is no longer available.

Example

The following example sets the route timeout period to 120 seconds:

```
configure rip routetimeout 120
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

configure rip updatetime

```
configure rip updatetime seconds
```

Description

Specifies the time interval in seconds within which RIP sends update packets.

Syntax Description

<i>seconds</i>	Specifies a time in seconds. The range is 10 to 180.
----------------	--

Default

30 seconds.

Usage Guidelines

The router exchanges an update message with each neighbor every 30 seconds (default value) or if there is a change to the overall routed topology (also called triggered updates). The timer granularity is 10 seconds. Timer minimum is 10 seconds and maximum is 180 seconds.



Example

The following command sets the update timer to 60 seconds:

```
configure rip updatetime 60
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

configure rip vlan cost

```
configure rip vlan [vlan_name | all] cost cost
```

Description

Configures the cost (metric) of the interface.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
all	Specifies all VLANs.
<i>cost</i>	Specifies a cost metric.

Default

The default setting is 1.

Usage Guidelines

The specified interface cost is added to the cost of the route received through this interface.

Example

The following command configures the cost for the VLAN finance to a metric of 3:

```
configure rip vlan finance cost 3
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

configure rip vlan route-policy

```
configure rip vlan [vlan_name | all] route-policy [in | out] [policy-name | none]
```

Description

Configures RIP to ignore certain routes received from its neighbor, or to suppress certain routes when performing route advertisements.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
all	Specifies all VLANs.
<i>policy-name</i>	Specifies a policy.
none	Removes any policy from the VLAN.

Default

N/A.

Usage Guidelines

Use the **in** option to configure an input route policy, which determines which RIP routes are accepted as valid routes. This policy can be combined with the trusted neighbor policy to accept selected routes only from a set of trusted neighbors.

Use the **out** option to configure an output route policy, which determines which RIP routes are advertised on the VLAN.

Example

The following command configures the VLAN backbone to accept selected routes from the policy nosales:

```
configure rip vlan backbone route-policy in nosales
```



The following command uses the policy nosales to determine which RIP routes are advertised into the VLAN backbone:

```
configure rip vlan backbone route-policy out nosales
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

configure rip vlan rxmode

```
configure rip [vlan vlan_name | all] rxmode [none | v1only | v2only | any]
```

Description

Syntax Description

<i>vlan_name</i>	Specifies to apply settings to specific VLAN name.
all	Specifies all VLANs.
none	Specifies to drop all received RIP packets.
v1only	Specifies to accept only RIP version 1 format packets.
v2only	Specifies to accept only RIP version 2 format packets.
any	Specifies to accept RIP version 1 and RIP version 2 packets.

Default

N/A.

Usage Guidelines

None.

Example

The following command configures the receive mode for the VLAN finance to accept only RIP version 1 format packets:

```
configure rip finance rxmode v1only
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

configure rip vlan trusted-gateway

```
configure rip vlan [vlan_name | all] trusted-gateway [policy-name | none]
```

Description

Configures a trusted neighbor policy to determine trusted RIP router neighbors for the VLAN on the switch running RIP.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
all	Specifies all VLANs.
<i>policy-name</i>	Specifies a policy.
none	Removes any trusted-gateway policy from the VLAN.

Default

N/A.

Usage Guidelines

Use this command to set a policy to determine trusted neighbors. A neighbor is defined by its IP address. Only the RIP control packets from trusted neighbors will be processed.

Example

The following command configures RIP to use the policy `nointernet` to determine from which RIP neighbor to receive (or reject) the routes to the VLAN backbone:

```
configure rip vlan backbone trusted-gateway nointernet
```

History

This command was first available in ExtremeXOS 10.1.



Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

configure rip vlan txmode

```
configure rip [vlan vlan_name | all] txmode [none | v1only | v1comp | v2only]
```

Description

Changes the RIP transmission mode for one or all VLANs.

Syntax Description

<i>vlan_name</i>	Specifies to apply settings to a specific VLAN name.
all	Specifies all VLANs.
none	Specifies to not transmit any packets on this interface.
v1only	Specifies to transmit RIP version 1 format packets to the broadcast address.
v1comp	Specifies to transmit RIP version 2 format packets to the broadcast address.
v2only	Specifies to transmit RIP version 2 format packets to the RIP multicast address.

Default

N/A.

Usage Guidelines

None.

Example

The following command configures the transmit mode for the VLAN finance to transmit version 2 format packets to the broadcast address:

```
configure rip finance txmode v1comp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.



disable rip

disable rip

Description

Disables RIP for the whole router.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

RIP has a number of limitations that can cause problems in large networks, including:

- A limit of 15 hops between the source and destination networks.
- A large amount of bandwidth taken up by periodic broadcasts of the entire routing table.
- Slow convergence.
- Routing decisions based on hop count; no concept of link costs or delay.
- Flat networks; no concept of areas or boundaries.

Example

The following command disables RIP for the whole router:

```
disable rip
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

disable rip aggregation

disable rip aggregation



Description

Disables the RIP aggregation of subnet information on a RIP version 2 (RIPv2) router.

Syntax Description

This command has no arguments or variables.

Default

RIP aggregation is disabled by default.

Usage Guidelines

The disable RIP aggregation command disables the RIP aggregation of subnet information on a switch configured to send RIPv2-compatible traffic. The switch summarizes subnet routes to the nearest class network route. The following rules apply when using RIP aggregation:

- Within a class boundary, no routes are aggregated.
- If aggregation is disabled, subnet routes are never aggregated, even when crossing a class boundary.

Example

The following command disables RIP aggregation on the interface:

```
disable rip aggregation
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

disable rip export

```
disable rip export [bgp | direct | e-bgp | i-bgp | ospf | ospf-extern1 | ospf-extern2 | ospf-inter | ospf-intra | static | isis | isis-level-1 | isis-level-1-external | isis-level-2 | isis-level-2-external ]
```

Description

Disables RIP from redistributing routes from other routing protocols.



Syntax Description

bgp	Specifies BGP routes.
direct	Specifies interface routes (only interfaces that have IP forwarding enabled are exported).
e-bgp	Specifies external BGP routes.
i-bgp	Specifies internal BGP routes.
ospf	Specifies all OSPF routes.
ospf-extern1	Specifies OSPF external route type 1.
ospf-extern2	Specifies OSPF external route type 2.
ospf-inter	Specifies OSPF-inter area routes.
ospf-intra	Specifies OSPF-intra area routes.
static	Specifies static routes.
isis	Specifies IS-IS routes.
isis-level-1	Specifies ISIS Level 1 routes.
isis-level-1-external	Specifies ISIS Level 1 External routes.
isis-level-2	Specifies ISIS Level 2 routes.
isis-level-2-external	Specifies ISIS Level 2 External routes.

Default

Disabled.

Usage Guidelines

This command disables the exporting of BGP, static, direct, and OSPF-learned routes into the RIP domain.

Example

The following command disables RIP from redistributing any routes learned from OSPF:

```
disable rip export ospf
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.



disable rip poisonreverse

disable rip poisonreverse

Description

Disables poison reverse algorithm for RIP.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Like split horizon, poison reverse is a scheme for eliminating the possibility of loops in the routed topology. In this case, a router advertises a route over the same interface that supplied the route, but the route uses a hop count of 16, defining it as unreachable.

Example

The following command disables the split horizon with poison reverse algorithm for RIP:

```
disable rip poisonreverse
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

disable rip splithorizon

disable rip splithorizon

Description

Disables the split horizon algorithm for RIP.



Syntax Description

This command has no arguments or variable.

Default

Enabled.

Usage Guidelines

Split horizon is a scheme for avoiding problems caused by including routes in updates sent to the router from which the route was learned. Split horizon omits routes learned from a neighbor in updates sent to that neighbor.

Example

The following command disables the split horizon algorithm for RIP:

```
disable rip splithorizon
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

disable rip triggerupdates

disable rip triggerupdates

Description

Disables the trigger update mechanism. Triggered updates are a mechanism for immediately notifying a router's neighbors when the router adds or deletes routes or changes their metric.

Syntax Description

This command has no arguments or variables.

Default

Enabled.



Usage Guidelines

Triggered updates occur whenever a router changes the metric for a route and it is required to send an update message immediately, even if it is not yet time for a regular update message to be sent. This will generally result in faster convergence, but may also result in more RIP-related traffic.

Example

The following command disables the trigger update mechanism:

```
disable rip triggerupdate
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

disable rip use-ip-router-alert

```
disable rip use-ip-router-alert
```

Description

Disables router alert IP option in outgoing RIP control packets.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables the RIP router alert IP option:

```
disable rip use-ip-router-alert
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

enable rip

enable rip

Description

Enables RIP for the whole router.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

RIP has a number of limitations that can cause problems in large networks, including:

- A limit of 15 hops between the source and destination networks.
- A large amount of bandwidth taken up by periodic broadcasts of the entire routing table.
- Slow convergence.
- Routing decisions based on hop count; no concept of link costs or delay.
- Flat networks; no concept of areas or boundaries.

Example

The following command enables RIP for the whole router:

```
enable rip
```

History

This command was first available in ExtremeXOS 10.1.



Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

enable rip aggregation

enable rip aggregation

Description

Enables the RIP aggregation of subnet information on a RIP version 2 (RIPv2) interface.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

The enable (disable) rip aggregation command enables (disables) the RIP aggregation of subnet information on an interface configured to send RIPv1 or RIPv2-compatible traffic. The switch summarizes subnet routes to the nearest class network route. The following rules apply when using RIP aggregation:

- Subnet routes are aggregated to the nearest class network route when crossing a class boundary.
- Within a class boundary, no routes are aggregated.
- If aggregation is enabled, the behavior is the same as in RIPv1.
- If aggregation is disabled, subnet routes are never aggregated, even when crossing a class boundary.

Example

The following command enables RIP aggregation on the interface:

```
enable rip aggregation
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.



enable rip export

```
enable rip export [bgp | direct | e-bgp | i-bgp | ospf | ospf-extern1 | ospf-extern2 | ospf-inter | ospf-intra | static | isis | isis-level-1 | isis-level-1-external | isis-level-2 | isis-level-2-external ] [cost number {tag number} | policy policy-name]
```

Description

Enables RIP to redistribute routes from other routing functions.

Syntax Description

bgp	Specifies BGP routes.
direct	Specifies interface routes (only interfaces that have IP forwarding enabled are exported).
e-bgp	Specifies E-BGP routes.
i-bgp	Specifies I-BGP routes.
ospf	Specifies all OSPF routes.
ospf-extern1	Specifies OSPF external route type 1.
ospf-extern2	Specifies OSPF external route type 2.
ospf-inter	Specifies OSPF-inter area routes.
ospf-intra	Specifies OSPF-intra area routes.
static	Specifies static routes.
isis	Specifies IS-IS routes.
isis-level-1	Specifies ISIS Level 1 routes.
isis-level-1-external	Specifies ISIS Level 1 External routes.
isis-level-2	Specifies ISIS Level 2 routes.
isis-level-2-external	Specifies ISIS Level 2 External routes.
cost number	Specifies the cost metric, from 0-15. If set to 0, RIP uses the route metric obtained from the route origin.
tag number	Specifies a tag number.
policy-name	Specifies a policy.

Default

Disabled.



Usage Guidelines

This command enables the exporting of BGP, static, direct, and OSPF-learned routes into the RIP domain. You can choose which types of OSPF routes are injected, or you can simply choose `ospf`, which will inject all learned OSPF routes regardless of type.

The cost metric is inserted for all RIP-learned, static, and direct routes injected into RIP. If the cost metric is set to 0, the cost is inserted from the route. For example, with BGP, the cost could be the MED or the length of the BGP path. The tag value is used only by special routing applications. Use 0 if you do not have specific requirements for using a tag.

Each protocol can have a policy associated with it to control or modify the exported routes.

Example

The following command enables RIP to redistribute routes from all OSPF routes:

```
enable rip export ospf cost 0
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

enable rip originate-default cost

```
enable rip originate-default {always} cost number {tag number}
```

Description

Configures a default route to be advertised by RIP.

Syntax Description

always	Specifies to always advertise the default route.
<i>cost number</i>	Specifies a cost metric. The range is 1 - 15.
<i>tag number</i>	Specifies a tag number.

Default

Disabled.



Usage Guidelines

If `always` is specified, RIP always advertises the default route to its neighbors. If `always` is not specified, RIP advertises a default route only if a reachable default route is in the system route table.

The default route advertisement is filtered using the out policy.

The cost metric is inserted for all RIP-learned, static, and direct routes injected into RIP. The tag value is used only by special routing applications.

Example

The following command configures a default route to be advertised by RIP if there is a default route in the system routing table:

```
enable rip originate-default cost 7
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

enable rip poisonreverse

```
enable rip poisonreverse
```

Description

Enables poison reverse algorithm for RIP.

Syntax Description

Enables poison reverse algorithm for RIP.

Default

Enabled.

Usage Guidelines

Like split horizon, poison reverse is a scheme for eliminating the possibility of loops in the routed topology. In this case, a router advertises a route over the same interface that supplied the route, but the route uses a hop count of 16, defining it as unreachable.



Example

The following command enables the split horizon with poison reverse algorithm for RIP:

```
enable rip poisonreverse
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

enable rip splithorizon

enable rip splithorizon

Description

Enables the split horizon algorithm for RIP.

Syntax Description

Enables the split horizon algorithm for RIP.

Default

Enabled.

Usage Guidelines

Split horizon is a scheme for avoiding problems caused by including routes in updates sent to the router from which the route was learned. Split horizon omits routes learned from a neighbor in updates sent to that neighbor.

Example

The following command enables the split horizon algorithm for RIP:

```
enable rip splithorizon
```

History

This command was first available in ExtremeXOS 10.1.



Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

enable rip triggerupdates

Triggered updates are a mechanism for immediately notifying a router's neighbors when the router adds or deletes routes or changes their metric.

enable rip triggerupdates

Description

Enables the trigger update mechanism.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Triggered updates occur whenever a router changes the metric for a route and it is required to send an update message immediately, even if it is not yet time for a regular update message to be sent. This will generally result in faster convergence, but may also result in more RIP-related traffic.

Example

The following command enables the trigger update mechanism:

```
enable rip triggerupdate
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

enable rip use-ip-router-alert

enable rip use-ip-router-alert



Description

Enables the router alert IP option in the outgoing RIP control packets.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables the RIP router alert IP option:

```
enable rip use-ip-router-alert
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

show rip

show rip

Description

Displays RIP specific configuration.

Syntax Description

This command has no arguments or variables.

Default

N/A.



Usage Guidelines

None.

Example

The following command displays RIP specific configuration:

```
show rip
```

The following is sample output from this command:

```
# show rip
RIP Routing      : Disabled      Operational status: Down
Split Horizon    : Enabled      Poison Reverse    : Enabled
Triggered Updates: Enabled      Aggregation      : Disabled
Update Interval  : 30          Route Timeout     : 180
Garbage Timeout  : 120        Router Alert      : Disabled
Originate Default: Disabled
Sys Import-Policy: None
Redistribute:
Protocol  Status    Cost Tag Policy
-----
Direct    Disabled  0    0    none
Static    Disabled  0    0    none
OSPFIntra Disabled  0    0    none
OSPFInter Disabled  0    0    none
OSPFExt1  Disabled  0    0    none
OSPFExt2  Disabled  0    0    none
E-BGP     Disabled  0    0    none
I-BGP     Disabled  0    0    none
ISISL1    Disabled  0    0    none
ISISL2    Disabled  0    0    none
ISISL1Ext Disabled  0    0    none
ISISL2Ext Disabled  0    0    none
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

show rip interface

```
show rip interface {detail}
```



Description

Displays RIP-specific configuration and statistics for all VLANs.

Syntax Description

detail	Specifies detailed display.
---------------	-----------------------------

Default

Show summary output for all interfaces.

Usage Guidelines

Summary includes the following information per interface:

- VLAN name.
- IP address and mask.
- interface status.
- packets transmitted.
- packets received.
- number of triggered updates.
- cost.

Detail includes the following per interface:

- VLAN name.
- IP address and mask.
- tx mode.
- rx mode.
- cost.
- peer information (for each peer).
 - age.
 - version.
 - received packets.
 - received updates.
 - received bad packets.
 - received bad routes.
- in policy.
- out policy.
- trusted gateway policy.
- packets transmitted.
- sent triggered updates.
- packets received.
- bad packets received.
- bad routes received.



Example

The following command displays the RIP configuration for all VLANs:

```
show rip interface
```

The following is sample output from this command:

```
# show rip interface
VLAN      IP Address      Flags Sent      Rcvd      Triggered Cost
Packets  Packets  Updates
Flags: (f) Interface Forwarding Enabled, (i) Interface RIP Enabled
(n) Multinetted VLAN, (r) Router RIP Enabled
```

The following command displays RIP-specific statistics for all VLANs:

```
show rip interface detail
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

show rip interface vlan

```
show rip interface vlan vlan_name
```

Description

Displays RIP specific statistics and configuration for a VLAN in detail.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
------------------	------------------------

Default

N/A.

Usage Guidelines

None.



Example

The following command displays RIP specific statistics for the VLAN accounting:

```
show rip interface vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

show rip memory

```
show rip memory {detail | memoryType}
```

Description

Displays RIP specific memory usage.

Syntax Description

detail	Displays detail information.
<i>memoryType</i>	Specifies the memory type usage to display.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays RIP specific memory for TEST:

```
show rip memory test
```

The following is sample output from this command:

```
# show rip memory test
RIP Memory Information
```



```

-----
Bytes Allocated: 0          AllocFailed: 0
Current Memory Utilization Level: GREEN
Memory Utilization Statistics
-----
Memory Statistics for TEST
-----
Size      64      80      96      256     384     512     768     1024     2048     40
96 18432
-----
-----
Alloced      0      0      0      0      0      0      0      0      0      0
0      0
AllocedPeak  0      0      0      0      0      0      0      0      0      0
0      0
AllocSuccess  0      0      0      0      0      0      0      0      0      0
0      0
FreeSuccess  0      0      0      0      0      0      0      0      0      0
0      0
AllocFail    0      0      0      0      0      0      0      0      0      0
0      0
FreeFail    0      0      0      0      0      0      0      0      0      0
0      0

```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

show rip routes

```
show rip routes {detail} {network ripNetworkPrefix}
```

Description

Displays routes advertised by RIP.

Syntax Description

detail	Displays all available information from the RIP routing table.
<i>ripNetworkPrefix</i>	Specifies the route prefix for the routes to show.

Default

N/A.



Usage Guidelines

The routes displayed include all routes advertised by RIP, including routes exported from the system routing table and originated by other protocols, for example BGP.

Example

The following command displays a summary of RIP specific routes for the networks 10.0.0.0/8:

```
show rip routes network 10.0.0.0/8
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

unconfigure rip

```
unconfigure rip {vlan vlan-name | all}
```

Description

Resets all RIP parameters to the default for all VLANs or for the specified VLAN.

Syntax Description

<i>vlan-name</i>	Specifies a VLAN name.
------------------	------------------------

Default

All.

Usage Guidelines

Does not change the enable/disable state of the RIP settings.

Example

The following command resets the RIP configuration to the default for the VLAN finance:

```
unconfigure rip finance
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.



37 RIPng Commands

```
clear ripng counters
configure ripng add
configure ripng cost
configure ripng delete
configure ripng garbagetime
configure ripng import-policy
configure ripng route-policy
configure ripng routetimeout
configure ripng trusted-gateway
configure ripng updatetime
disable ripng
disable ripng export
disable ripng originate-default
disable ripng poisonreverse
disable ripng splithorizon
disable ripng triggerupdate
enable ripng
enable ripng export
enable ripng originate-default
enable ripng poisonreverse
enable ripng splithorizon
enable ripng triggerupdates
show ripng
show ripng interface
show ripng routes
unconfigure ripng
```

This chapter describes commands used for the IPv6 interior gateway protocol RIPng.

To determine the best path to a distant network, a router using RIPng always selects the path that has the least number of hops. Each router that data must traverse is considered to be one hop.

The routing table in a router using RIPng contains an entry for every known destination network. Each routing table entry contains the following information:

- IP address and prefix length of the destination network.
- Metric (hop count) to the destination network.
- IP address of the next hop router, if the destination is not directly connected.
- Interface for the next hop.
- Timer that tracks the amount of time since the entry was last updated.

- A flag that indicates if the entry is a new one since the last update.
- The source of the route, for example, static, RIPng, OSPFv3, etc.

The router exchanges an update message with each neighbor every 30 seconds (default value), or if there is a change to the overall routed topology (also called triggered updates). If a router does not receive an update message from its neighbor within the route timeout period (180 seconds by default), the router assumes the connection between it and its neighbor is no longer available.



Note

RIPng is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

clear ripng counters

```
clear ripng counters {vlan vlan-name | tunnel tunnel-name}
```

Description

Clears the RIPng global or interface-specific counters (statistics).

Syntax Description

<i>vlan-name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel-name</i>	Specifies an IPv6 tunnel.

Default

N/A.

Usage Guidelines

None.

Example

The following command clears the RIPng statistics counters:

```
clear ripng counters
```

History

This command was first available in ExtremeXOS 11.2.



Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ripng add

```
configure ripng add [vlan vlan-name | tunnel tunnel-name | [vlan | tunnel] all]
```

Description

Configures RIPng on an IP interface.

Syntax Description

<i>vlan-name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel-name</i>	Specifies an IPv6 tunnel.
all	Specifies all IPv6 configured VLANs or tunnels.

Default

N/A.

Usage Guidelines

For RIPng to be active on the interface, it must also be globally enabled using the command `disable ripng export [direct | ospfv3 | ospfv3-extern1 | ospfv3-extern2 | ospfv3-inter | ospfv3-intra | static | isis | isis-level-1 | isis-level-1-external | isis-level-2 | isis-level-2-external | bgp]`. If the keyword `all` is specified, all IPv6 configured VLANs or tunnels will be configured for RIPng.

Example

The following command configures RIPng on the VLAN `finance`:

```
configure ripng add finance
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure ripng cost

```
configure ripng [vlan vlan-name | tunnel tunnel-name] cost metric
```

Description

Configures the cost (metric) of the interface..

Syntax Description

<i>vlan-name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel-name</i>	Specifies an IPv6 tunnel.
<i>metric</i>	Specifies a cost metric. Range is 1 to 15.

Default

The default setting is 1.

Usage Guidelines

The specified interface cost is added to the cost of the route received through this interface.

Example

The following command configures the cost for the VLAN finance to a metric of 3:

```
configure ripng vlan finance cost 3
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ripng delete

```
configure ripng delete [vlan vlan-name | tunnel tunnel-name | [vlan | tunnel]  
all]
```

Default

Removes an interface from RIPng routing.



Syntax Description

<i>vlan-name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel-name</i>	Specifies an IPv6 tunnel.
all	Specifies all IPv6 configured VLANs or tunnels.

Default

N/A.

Usage Guidelines

This command removes an interface from RIPng routing. However, the RIPng-specific interface configuration will be preserved, even if RIPng is unconfigured on the interface. The interface configuration information is removed only when the IPv6 interface itself gets deleted by, for example, by unconfiguring all the IPv6 addresses on the interface.

Example

The following command removes the VLAN finance from RIPng routing:

```
configure ripng delete finance
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ripng garbagetime

```
configure ripng garbagetime {seconds}
```

Description

Configures the RIPng garbage time.

Syntax Description

<i>seconds</i>	Specifies a time in seconds. Range is 10 to 2400 seconds.
----------------	---



Default

120 seconds.

Usage Guidelines

This command configures the time interval after which a route in the RIPng routing database that has expired will be removed. The value is rounded off to nearest multiple of 10.

Example

The following command configures the RIPng garbage time to have a 60-second delay:

```
configure ripng garbagetime 60
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ripng import-policy

```
configure ripng import-policy [policy-name | none]
```

Description

Configures the import policy for RIPng.

Syntax Description

<i>policy-name</i>	Specifies the policy.
--------------------	-----------------------

Default

No policy.

Usage Guidelines

Use this command to configure the policy to be applied to RIPng routes installed into the system routing table from the RIPng routing process. This policy can be used to modify parameters associated with routes installed into the routing table. The import policy cannot be used to determine the routes to be added to the routing table.



Use the none option to remove the import policy.

The following is a sample policy file that can be used with RIPng. It changes the metric to 12 for any routes from the subnets 2001:db8:2ccc::/64 and 2001:db8:2ccd::/64:

```
entry filter_routes {
  if match any {
    nlri 2001:db8:2ccc:: /64;
    nlri 2001:db8:2ccd:: /64;
  }
  then {
    cost 12;
  }
}
```

Example

The following example applies the policy campuseast to RIPng routes:

```
configure ripng import-policy campuseast
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ripng route-policy

```
configure ripng [vlan vlan-name | tunnel tunnel-name] route-policy [in | out]
[policy-name | none]
```

Description

Configures RIPng to ignore or modify certain routes received from its neighbors, or to suppress certain routes when performing route advertisements.

Syntax Description

<i>vlan-name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel-name</i>	Specifies an IPv6 tunnel.
<i>policy-name</i>	Specifies a policy.
none	Removes any policy from the VLAN.



Default

N/A.

Usage Guidelines

Use the `in` option to configure an input route policy, which determines which RIPng routes are accepted as valid routes from RIPng neighbors. This policy can be combined with the trusted neighbor policy to accept selected routes only from a set of trusted neighbors.

Use the `out` option to configure an output route policy, which determines which RIPng routes are advertised to other RIPng neighbors.

The following is a sample policy file that could be used with RIPng. It will drop any routes from the subnets 2001:db8:2ccc::/64 and 2001:db8:2ccd::/64:

```
entry filter_routes {
  if match any {
    nlri 2001:db8:2ccc:: /64;
    nlri 2001:db8:2ccd:: /64;
  }
  then {
    deny;
  }
}
```

Example

The following command configures the VLAN backbone to accept routes from its neighbor as specified by the policy `nosales`:

```
configure ripng vlan backbone route-policy in nosales
```

The following command uses the policy `nosales` to determine which RIP routes are advertised into the VLAN backbone:

```
configure rip vlan backbone route-policy out nosales
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure ripng routetimeout

```
configure ripng routetimeout seconds
```

Description

Configures the route timeout period for RIPng.

Syntax Description

<i>seconds</i>	Specifies a time in seconds. Range is 10 to 3600.
----------------	---

Default

180 seconds.

Usage Guidelines

If a router does not receive an update message from its neighbor within the route timeout period (180 seconds by default), the router assumes the connection between it and its neighbor is no longer available.

The configured value is rounded off to the nearest multiple of 10.

Example

The following example sets the route timeout period to 120 seconds:

```
configure ripng routetimeout 120
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ripng trusted-gateway

```
configure ripng [vlan vlan-name | tunnel tunnel-name] trusted-gateway [policy-name | none]
```



Description

Configures a trusted neighbor policy to determine trusted RIPng router neighbors for the interfaces on the switch running RIPng.

Syntax Description

<i>vlan-name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel-name</i>	Specifies an IPv6 tunnel.
<i>policy-name</i>	Specifies a policy.
none	Removes any trusted-gateway policy from the VLAN.

Default

None. Control packets from all of the neighbors are processed.

Usage Guidelines

Use this command to set a policy to determine trusted neighbors. A neighbor is defined by its IP address. Only the RIPng control packets from trusted neighbors will be processed.

The following policy designates neighbors from the fe80:202:b3ff:fe4a:6ada:: /64 subnet and the neighbor at fe80:203::b3ff:fe4a:6ada as trusted gateways:

```
entry filter_gateways {
  if match any {
    nlri fe80:202:b3ff:fe4a:6ada:: /64;
    nlri fe80:203::b3ff:fe4a:6ada:: /64;
  }
  then {
    permit;
  }
}
```

Example

The following command configures RIPng to use the policy nointernet to determine from which RIPng neighbor to receive (or reject) the routes to the VLAN backbone:

```
configure ripng vlan backbone trusted-gateway nointernet
```

History

This command was first available in ExtremeXOS 11.2.



Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ripng updatetime

```
configure ripng updatetime seconds
```

Description

Specifies the time interval in seconds within which RIPng sends update packets.

Syntax Description

<i>seconds</i>	Specifies a time in seconds. The range is 10 to 3600.
----------------	---

Default

30 seconds.

Usage Guidelines

The router exchanges an update message with each neighbor every 30 seconds (default value), or if there is a change to the overall routed topology (also called triggered updates). The timer granularity is 10 seconds. Timer minimum is 10 second and maximum is 3600 seconds.

Example

The following command sets the update timer to 60 seconds:

```
configure ripng updatetime 60
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable ripng

```
disable ripng
```



Description

Disables RIPng for the whole router.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables RIPng for the whole router:

```
disable ripng
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable ripng export

```
disable ripng export [direct | ospfv3 | ospfv3-extern1 | ospfv3-extern2 | ospfv3-  
inter | ospfv3-intra | static | isis | isis-level-1 | isis-level-1-external |  
isis-level-2 | isis-level-2-external | bgp]
```

Description

Disables RIPng from redistributing routes from other routing protocols.

Syntax Description

direct	Specifies directly reachable subnets from the router (only interfaces that have IP forwarding enabled are exported).
ospfv3	Specifies all OSPFv3 routes.



ospfv3-extern1	Specifies OSPFv3 external route type 1.
ospfv3-extern2	Specifies OSPFv3 external route type 2.
ospfv3-inter	Specifies OSPFv3-inter area routes.
ospfv3-intra	Specifies OSPFv3-intra area routes.
static	Specifies user configured static routes.
isis	Specifies IS-IS routes.
isis-level-1	Specifies IS-IS Level 1 routes.
isis-level-1-external	Specifies IS-IS Level 1 External routes.
isis-level-2	Specifies IS-IS Level 2 routes.
isis-level-2-external	Specifies IS-IS Level 2 External routes.
bgp	Specifies BGP IPv6 routes

Default

Disabled.

Usage Guidelines

This command disables the exporting of static, direct, IS-IS, and OSPF-learned routes from the switch routing table into the RIPng domain.

Example

The following command disables RIPng from redistributing any routes learned from OSPFv3:

```
disable ripng export ospfv3
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable ripng originate-default

```
disable ripng originate-default
```



Description

Disables the advertisement of a default route to the neighbors.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command unconfigures a default route to be advertised by RIPng if no other default route is advertised:

```
disable ripng originate-default
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable ripng poisonreverse

```
disable ripng poisonreverse {vlan vlan-name | tunnel tunnel_name | [vlan |  
tunnel] all}
```

Description

Disables poison reverse algorithm for RIPng.



Syntax Description

<i>vlan-name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
all	Specifies all interfaces.

Default

Enabled.

Usage Guidelines

Like split horizon, poison reverse is a scheme for eliminating the possibility of loops in the routed topology. In this case, a router advertises a route over the same interface that supplied the route, but the route uses a hop count of 16, defining it as unreachable.

Example

The following command disables the split horizon with poison reverse algorithm for RIPng:

```
disable ripng poisonreverse
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable ripng splithorizon

```
disable ripng splithorizon {vlan vlan-name | tunnel tunnel_name | [vlan | tunnel]
all}
```

Description

Disables the split horizon algorithm for RIPng.



Syntax Description

<i>vlan-name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
all	Specifies all interfaces.

Default

Enabled.

Usage Guidelines

Split horizon is a scheme for avoiding problems caused by including routes in updates sent to the router from which the route was learned. Split horizon omits routes learned from a neighbor in updates sent to that neighbor.

Example

The following command disables the split horizon algorithm for RIPng:

```
disable rip splithorizon
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable ripng triggerupdate

```
disable ripng triggerupdate {vlan vlan-name | tunnel tunnel_name | [vlan |  
tunnel] all}
```

Description

Triggered updates are a mechanism for immediately notifying a router's neighbors when the router adds or deletes routes or changes their metric. This command disables the trigger update mechanism.



Syntax Description

<code>vlan-name</code>	Specifies an IPv6 configured VLAN.
<code>tunnel_name</code>	Specifies an IPv6 tunnel.
all	Specifies all interfaces.

Default

Enabled.

Usage Guidelines

Triggered updates occur whenever a router changes the metric for a route and it is required to send an update message immediately, even if it is not yet time for a regular update message to be sent. This will generally result in faster convergence, but may also result in more RIPng-related traffic.

When this feature is disabled, any metric change on the interface, or an interface going down will not be communicated until the next periodic update. To configure how often periodic updates are sent, use the following command:

```
configure ripng updatetime
```

Example

The following command disables the trigger update mechanism:

```
disable ripng triggerupdate
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable ripng

```
enable ripng
```

Description

Enables RIPng for the whole router.



Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Although RIPng is useful in small networks, it has a number of limitations that can cause problems in large networks, including:

- A limit of 15 hops between the source and destination networks.
- A large amount of bandwidth taken up by periodic broadcasts of the entire routing table.
- Slow convergence.
- Routing decisions based on hop count; no concept of link costs or delay.
- Flat networks; no concept of areas or boundaries.

For larger networks, consider OSPFv3 as an alternative IGP.

Example

The following command enables RIPng for the whole router:

```
enable ripng
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable ripng export

```
enable ripng export [direct | ospfv3 | ospfv3-extern1 | ospfv3-extern2 | ospfv3-inter | ospfv3-intra | static | isis | isis-level-1 | isis-level-1-external | isis-level-2 | isis-level-2-external | bgp] [cost number {tag number} | policy policy-name]
```

Description

Enables RIPng to redistribute routes from other routing functions.



Syntax Description

direct	Specifies interface routes (only interfaces that have IP forwarding enabled are exported).
ospfv3	Specifies all OSPFv3 routes.
ospfv3-extern1	Specifies OSPFv3 external route type 1.
ospfv3-extern2	Specifies OSPFv3 external route type 2.
ospfv3-inter	Specifies OSPFv3-inter area routes.
ospfv3-intra	Specifies OSPFv3-intra area routes.
static	Specifies static routes.
isis	Specifies IS-IS routes.
isis-level-1	Specifies IS-IS Level 1 routes.
isis-level-1-external	Specifies IS-IS Level 1 External routes.
isis-level-2	Specifies IS-IS Level 2 routes.
isis-level-2-external	Specifies IS-IS Level 2 External routes.
bgp	Specifies BGP IPv6 routes
cost <i>number</i>	Specifies the cost metric, from 0-15. If set to 0, RIPng uses the route metric obtained from the route origin.
tag <i>number</i>	Specifies a tag number.
<i>policy-name</i>	Specifies a policy.

Default

Disabled. However, direct routes will always be advertised for all the interfaces where RIPng is enabled. For those interfaces where RIPng is not enabled, the corresponding direct route could be redistributed if direct route export is enabled through this command.

Default tag is 0.

Usage Guidelines

This command enables the exporting of static, direct, IS-IS, and OSPFv3-learned routes from the routing table into the RIPng domain. You can choose which types of IS-IS or OSPFv3 routes are injected, or you can simply choose isis or ospfv3, which will inject all learned routes (of all types) for the selected protocol.

The cost metric is inserted for all RIPng-learned, static, and direct routes injected into RIPng. If the cost metric is set to 0, the cost is inserted from the route table. The tag value is used only by special routing applications. Use 0 if you do not have specific requirements for using a tag.



Each protocol can have a policy associated with it to control or modify the exported routes. The following is sample policy file which modifies the cost of redistributed routes from OSPFv3 and statically configured routes:

```
entry filter_rt {
  If match any {
    Route-origin ospfv3;
    Route-origin static;
  }
  then {
    cost 10;
  }
}
```

Example

The following command enables RIPng to redistribute routes from all OSPFv3 routes:

```
enable ripng export ospfv3 cost 0
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable ripng originate-default

```
enable ripng originate-default {always} cost metric {tag number}
```

Description

Configures a default route to be advertised by RIPng.

Syntax Description

always	Specifies to advertise the default route in addition to learned default route.
cost metric	Specifies a cost metric. The range is 1 - 15.
tag number	Specifies a tag number.



Default

Disabled.

Usage Guidelines

If `always` is specified, RIPng always advertises the default route to its neighbors. If `always` is not specified, RIPng advertises a default route only if a reachable default route is in the system route table (the route is learned from other neighbors).

The default route advertisement is filtered using the out policy. Use the command, `configure ripng route-policy`, to specify the out policy.

The cost metric is inserted for all RIPng-learned, static, and direct routes injected into RIPng. The tag value is used only by special routing applications.

Example

The following command configures a default route to be advertised by RIPng if there is a default route in the system routing table:

```
enable ripng originate-default cost 7
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable ripng poisonreverse

```
enable ripng poisonreverse {vlan vlan-name | tunnel tunnel_name | [vlan | tunnel]
all}
```

Description

Enables the split horizon with poison reverse algorithm for RIPng on specified interfaces.

Syntax Description

<i>vlan-name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
all	Specifies all interfaces.



Default

Enabled.

Usage Guidelines

Used with split horizon, poison reverse is a scheme for eliminating the possibility of loops in the routed topology. In this case, a router advertises a route over the same interface that supplied the route, but the route uses a hop count of 16, defining it as unreachable.

If both split horizon and poison reverse are enabled, poison reverse takes precedence.

Example

The following command enables split horizon with poison reverse for RIPng on all IPv6 interfaces in the virtual router:

```
enable ripng poisonreverse
```

The following command enables split horizon with poison reverse for all the IPv6 configured VLANs in the virtual router:

```
enable ripng poisonreverse vlan all
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable ripng splithorizon

```
enable ripng splithorizon {vlan vlan-name | tunnel tunnel_name | [vlan | tunnel]  
all}
```

Description

Enables the split horizon algorithm for RIPng.



Syntax Description

<i>vlan-name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
all	Specifies all interfaces.

Default

Enabled.

Usage Guidelines

Split horizon is a scheme for avoiding problems caused by including routes in updates sent to the router from which the route was learned. Split horizon omits routes learned from a neighbor in updates sent to that neighbor.

Example

The following command enables the split horizon algorithm for RIPng on all IPv6 configured interfaces:

```
enable ripng splithorizon
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable ripng triggerupdates

```
enable ripng triggerupdates {vlan vlan-name | tunnel tunnel_name | [vlan |  
tunnel] all}
```

Description

Enables the trigger update mechanism. Triggered updates are a mechanism for immediately notifying a router's neighbors when the router adds or deletes routes or changes their metric.



Syntax Description

<i>vlan-name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
all	Specifies all interfaces.

Default

Enabled.

Usage Guidelines

Triggered updates occur whenever a router changes the metric for a route and it is required to send an update message immediately, even if it is not yet time for a regular update message to be sent. This will generally result in faster convergence, but may also result in more RIPng-related traffic.

Example

The following command enables the trigger update mechanism on all IPv6 configured interfaces:

```
enable ripng triggerupdate
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show ripng

```
show ripng
```

Description

Displays RIPng global configuration and runtime information.

Syntax Description

This command has no arguments or variables.



Default

N/A.

Usage Guidelines

None.

Example

The following command displays RIPng global configuration and runtime information:

```
show ripng
```

The following is sample output from this command:

```
RIPng Routing      : Disabled
Aggregation       : Disabled
Update Interval   : 30           Route Timeout      : 180
Garbage Timeout   : 120
Originate Default: Disabled
Sys Import-Policy: None
Redistribute:
Protocol          Status    Cost Tag    Policy
-----
Direct            Disabled  0    0    none
Static            Disabled  0    0    none
Ospf3-intra       Disabled  0    0    none
Ospf3-inter       Disabled  0    0    none
Ospf3-extern1     Disabled  0    0    none
Ospf3-extern2     Disabled  0    0    none
IsisL1            Disabled  0    0    none
IsisL2            Disabled  0    0    none
IsisL1Ext         Disabled  0    0    none
IsisL2Ext         Disabled  0    0    none
bgp                Disabled  0    0    none
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show ripng interface

```
show ripng interface {detail | vlan vlan-name | tunnel tunnel-name}
```



Description

Displays RIPng-specific configuration and statistics for the specified interface.

Syntax Description

detail	Specifies detailed display.
<i>vlan-name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel-name</i>	Specifies an IPv6 tunnel.

Default

Show summary output for all interfaces.

Usage Guidelines

Displays the RIPng interface configuration and runtime information. If no interface is specified, only the summary data for all the configured interfaces is displayed. If an interface is specified, only the data for that interface is displayed in detail. If the keyword detail is specified, detailed data for all interfaces is displayed.

Example

The following command displays the RIPng configuration summary for all interfaces:

```
show ripng interface
```

The following is sample output from this command:

VLAN	IP Address	Flags	Sent Packets	Rcvd Packets	Triggered Updates	Cost
v1	22cc::3	/64 rif-pst	106349	106349	3	15
v2	22bb::1	/64 rif-pst	106349	106095	3	1
v3	2abc::1	/120 rif-pst	106351	0	4	1
v4	3ffe::1	/64 rif-pst	106349	139124	3	1

```
Flags: (f) Interface Forwarding Enabled, (i) Interface RIPng Enabled
        (n) Multinetted Interface, (r) Router RIPng Enabled
        (p) Poison Reverse Enabled, (s) Split Horizon Enabled
        (t) Triggerred Update Enabled.
```

The following command displays RIPng-specific statistics for the VLAN v1:

```
show ripng interface v1
```



The following is sample output from this command:

```

VLAN                : v1                Interface           : 22cc::3/64
Router RIPng        : Enabled            Cost                : 15
Input Policy        : None               Output Policy       : None
Trusted GW Policy   : gw6               Poison Reverse      : Enabled
Split Horizon       : Enabled            Triggered Updates   : Enabled
Rcvd Packets        : 106358             Sent Packets        : 106358
Sent Trig. Updates  : 3                  Rcvd Bad Packets   : 0
Rcvd Bad Routes     : 0
Neighbor Addresses  : fe80::201:30ff:fe94:f400
Interface Addresses : 22cc::3/64, fe80::280:c8ff:feb9:2855/64

```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show ripng routes

```
show ripng routes {detail} {network ripngNetworkPrefix}
```

Description

Displays all matching routes in the RIPng routing database.

Syntax Description

detail	Displays all available information from the RIPng routing table.
ipv6-prefix	Specifies the route prefix for the routes to show.
prefix-length	Specifies the address mask of the IPv6 prefix.

Default

N/A.

Usage Guidelines

The routes displayed include all routes advertised by RIPng, including routes exported from the system routing table and originated by other protocols, for example OSPFv3 (also called redistributed routes).



Example

The following command displays a summary of RIPng specific routes:

```
show ripng routes
```

The following is sample output from this command:

Network	Next Hop	Mtr	VLAN
*> 2aaa::/64	fe80::201:30ff:fef4:5ca0%v1	2	v1
*	fe80::201:30ff:fe94:f400%v2	2	v2
*> 2bbb::/64	fe80::201:30ff:fef4:5ca0%v1	2	v1
*	fe80::201:30ff:fe94:f400%v2	3	v2
*> 2ccc::/64	(local)	1	(direct)
*	fe80::201:30ff:fef4:5ca0%v1	2	v1
*	fe80::201:30ff:fe94:f400%v2	3	v2
*> 2ddd::/64	(local)	1	(direct)
*	fe80::201:30ff:fe94:f400%v2	2	v2

The following command displays the detailed RIPng route information:

```
show ripng routes detail
```

The following is sample output from this command:

```
IPv6 RIPng routing table entry for 2aaa::/64
Paths: (2 available, best #1)
  fe80::201:30ff:fef4:5ca0%v1 from fe80::201:30ff:fef4:5ca0%v1 (v1)
    Metric 2, tag 0, timeout in 02:38, valid, best
  fe80::201:30ff:fe94:f400%v2 from fe80::201:30ff:fe94:f400%v2 (v2)
    Metric 2, tag 0, timeout in 02:44, valid
IPv6 RIPng routing table entry for 2bbb::/64
Paths: (2 available, best #1)
  fe80::201:30ff:fef4:5ca0%v1 from fe80::201:30ff:fef4:5ca0%v1 (v1)
    Metric 2, tag 0, timeout in 02:38, valid, best
  fe80::201:30ff:fe94:f400%v2 from fe80::201:30ff:fe94:f400%v2 (v2)
    Metric 3, tag 0, timeout in 02:44, valid
IPv6 RIPng routing table entry for 2ccc::/64
Paths: (3 available, best #1)
  Local from direct
    Metric 1, tag 0, no timeout, valid, best
  fe80::201:30ff:fef4:5ca0%v1 from fe80::201:30ff:fef4:5ca0%v1 (v1)
    Metric 2, tag 0, timeout in 02:38, valid
  fe80::201:30ff:fe94:f400%v2 from fe80::201:30ff:fe94:f400%v2 (v2)
    Metric 3, tag 0, timeout in 02:44, valid
```

History

This command was first available in ExtremeXOS 11.2.



Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure ripng

```
unconfigure ripng {vlan vlan-name | tunnel tunnel-name | vlan all | tunnel all}
```

Description

Resets RIPng parameters to the default value.

Syntax Description

<i>vlan-name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel-name</i>	Specifies an IPv6 tunnel.
all	Specifies either all IPv6 configured VLANs or all IPv6 tunnels.

Default

N/A.

Usage Guidelines

Issuing the command `unconfigure ripng` resets all the interfaces and the global configuration to the defaults, and disables RIPng, as that is the default.

Example

The following command resets the RIPng configuration to the default for the VLAN `finance`:

```
unconfigure rip finance
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



38 OSPF Commands

Licensing

OSPF Edge Mode

clear ospf counters

configure ospf add virtual-link

configure ospf add vlan area

configure ospf add vlan area link-type

configure ospf area external-filter

configure ospf area interarea-filter

configure ospf area add range

configure ospf area normal

configure ospf area nssa stub-default-cost

configure ospf area stub stub-default-cost

configure ospf area timer

configure ospf ase-limit

configure ospf ase-summary add

configure ospf ase-summary delete

configure ospf authentication

configure ospf bfd

configure ospf cost

configure ospf delete virtual-link

configure ospf delete vlan

configure ospf import-policy

configure ospf lsa-batch-interval

configure ospf metric-table

configure ospf priority

configure ospf restart

configure ospf restart grace-period

configure ospf restart-helper

configure ospf routerid

configure ospf spf-hold-time

configure ospf virtual-link timer

configure ospf vlan area

configure ospf vlan neighbor add

configure ospf vlan neighbor delete

configure ospf vlan timer

create ospf area

delete ospf area

```
disable ospf
disable ospf capability opaque-lsa
disable ospf export
disable ospf originate-default
disable ospf restart-helper-lsa-check
disable ospf use-ip-router-alert
disable snmp traps ospf
enable ospf
enable ospf capability opaque-lsa
enable ospf export
enable ospf originate-default
enable ospf restart-helper-lsa-check
enable ospf use-ip-router-alert
enable snmp traps ospf
show ospf
show ospf area
show ospf ase-summary
show ospf interfaces
show ospf interfaces detail
show ospf lsdb
show ospf memory
show ospf neighbor
show ospf virtual-link
unconfigure ospf
```

This chapter describes commands used for the interior gateway protocol OSPF.

Open Shortest Path First (OSPF) is a link-state protocol that distributes routing information between routers belonging to a single IP domain, also known as an autonomous system (AS). In a link-state routing protocol, each router maintains a database describing the topology of the autonomous system. Each participating router has an identical database maintained from the perspective of that router.

From the link-state database (LSDB), each router constructs a tree of shortest paths, using itself as the root. The shortest path tree provides the route to each destination in the autonomous system. When several equal-cost routes to a destination exist, traffic can be distributed among them. The cost of a route is described by a single metric.

OSPF allows parts of a network to be grouped together into areas. The topology within an area is hidden from the rest of the autonomous system. Hiding this information enables a significant reduction in link-state advertisement (LSA) traffic, and reduces the computations needed to maintain the LSDB. Routing within the area is determined only by the topology of the area.

The three types of routers defined by OSPF are as follows:

- **Internal Router (IR)**—An internal router has all of its interfaces within the same area.
- **Area Border Router (ABR)**—An ABR has interfaces belonging to two or more areas. It is responsible for exchanging summary advertisements with other ABRs.



- **Autonomous System Border Router (ASBR)**—An ASBR acts as a gateway between OSPF and other routing protocols, or other autonomous systems.

Each switch that is configured to run OSPF must have a unique router ID. It is recommended that you manually set the router ID of the switches participating in OSPF, instead of having the switch automatically choose its router ID based on the highest interface IP address. Not performing this configuration in larger, dynamic environments could result in an older LSDB remaining in use.



Note

Do not set the router ID to 0.0.0.0.

Licensing

See *Feature License Requirements* in the *Extreme XOS Concepts Guide* for information about licensing requirements.

OSPF Edge Mode

OSPF Edge Mode is a subset of OSPF available on platforms with an Advanced Edge license. There are two restrictions on OSPF Edge Mode:

- At most, four Active OSPF VLAN interfaces are permitted. There is no restriction on the number of Passive interfaces.
- The OSPF Priority on VLANs is zero, and is not configurable. This prevents the system from acting as a DR or BDR.

clear ospf counters

```
clear ospf counters { interfaces [all | vlan vlan_name | area area-identifier] |
area [all | area-identifier] | virtual-link [all | router-identifier area-
identifier] | neighbor [all | routerid [ip-address {ip-mask} | ipNetmask] | vlan
vlan_name] | system}
```

Description

Clears the OSPF counters (statistics).

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>router-identifier</i>	Specifies a router interface number.
<i>area-identifier</i>	Specifies an OSPF area.
<i>ip-address</i>	Specifies an IP address
<i>ip-mask</i>	Specifies a subnet mask.



<i>ipNetmask</i>	Specifies IP address / Netmask.
system	Specifies the OSPF system counters.

Default

N/A.

Usage Guidelines

The global command `clear counters` also clears all OSPF counters. This global command is the equivalent of `clear ospf counters` for OSPF.

Example

The following command clears the OSPF counters for area 1.1.1.1:

```
clear ospf counters area 1.1.1.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ospf add virtual-link

```
configure ospf add virtual-link router-identifier area-identifier
```

Description

Adds a virtual link connected to another ABR.

Syntax Description

<i>router-identifier</i>	Specifies the router ID of the other end of the link.
<i>area-identifier</i>	Specifies an OSPF area.

Default

N/A.



Usage Guidelines

A virtual link provides a logical path between the ABR of the disconnected area and the ABR of the normal area that connects to the backbone. A virtual link must be established between two ABRs that have a common area, with one ABR connected to the backbone. Specify the following:

- `router-identifier`—Far-end router interface number.
- `area-identifier`—Transit area used for connecting the two end-points. The transit area cannot have the area identifier 0.0.0.0. and cannot be a stub area or an NSSA.

Example

The following command configures a virtual link between the two interfaces:

```
configure ospf add virtual-link 10.1.2.1 10.1.0.0
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ospf add vlan area

```
configure ospf add vlan [vlan-name | all] area area-identifier {passive}
```

Description

Enables OSPF on one or all VLANs (router interfaces).

Syntax Description

<i>vlan-name</i>	Specifies a VLAN name.
all	Specifies all VLANs.
<i>area-identifier</i>	Specifies the area to which the VLAN is assigned.
passive	Specifies to stop sending and receiving hello packets on this interface.

Default

Disabled.



Usage Guidelines

Not applicable.

Example

The following command enables OSPF on a VLAN named accounting:

```
configure ospf add vlan accounting area 0.0.0.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospf add vlan area link-type

```
configure ospf add vlan vlan-name area area-identifier link-type [auto | broadcast | point-to-point] {passive}
```

Description

Configures the OSPF link type.

Syntax Description

<i>vlan-name</i>	Specifies a VLAN name.
<i>area-identifier</i>	Specifies the area to which the VLAN is assigned.
auto	Specifies to automatically determine the OSPF link type based on the interface type.
broadcast	Specifies a broadcast link, such as Ethernet. Routers must elect a DR and a BDR during synchronization.
point-to-point	Specifies a point-to-point link type, such as PPP.
passive	Specifies to stop sending and receiving packets on this interface.

Default

Auto.



Usage Guidelines

The passive parameter indicates that the router only synchronizes and listens, and does not originate or send any new information on the interface.

Example

The following command configures the OSPF link type as automatic on a VLAN named accounting:

```
configure ospf add vlan accounting area 0.0.0.1 link-type auto
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ospf area external-filter

```
configure ospf area area-identifier external-filter [policy-map | none]
```

Description

Configures an external filter policy.

Syntax Description

<i>area-identifier</i>	Specifies the OSPF target area.
<i>policy-map</i>	Specifies a policy.
none	Specifies not to apply an external filter (removes the existing policy, if any).

Default

N/A.

Usage Guidelines

For switches configured to support multiple OSPF areas (an ABR function), a policy can be applied to an OSPF area that filters a set of OSPF external routes from being advertised into that area.

Using the none mode specifies that no external filter is applied.



Example

The following command configures an external filter policy, nosales:

```
configure ospf area 1.2.3.4 external-filter nosales
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospf area interarea-filter

```
configure ospf area area-identifier interarea-filter [policy-map | none]
```

Description

Configures a global inter-area filter policy.

Syntax Description

<i>area-identifier</i>	Specifies the OSPF target area.
<i>policy-map</i>	Specifies a policy.
none	Specifies not to apply an interarea filter.

Default

N/A.

Usage Guidelines

For switches configured to support multiple OSPF areas (an ABR function), a policy can be applied to an OSPF area that filters a set of OSPF inter-area routes from being sourced from any other areas.

Example

The following command configures an inter-area filter policy, nosales:

```
configure ospf area 0.0.0.6 interarea-filter nosales
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospf area add range

```
configure ospf area area-identifier add range [ip-address ip-mask | ipNetmask]
[advertise | noadvertise] [type-3 | type-7]
```

Description

Configures a range of IP addresses in an OSPF area to be aggregated.

Syntax Description

<i>area-identifier</i>	Specifies an OSPF area.
<i>ip-address</i>	Specifies an IP address
<i>ip-mask</i>	Specifies a subnet mask.
<i>ipNetmask</i>	Specifies IP address / Netmask.
advertise	Specifies to advertise the aggregated range of IP addresses.
noadvertise	Specifies not to advertise the aggregated range of IP addresses.
type-3	Specifies type 3 LSA, summary LSA.
type-7	Specifies type 7 LSA, NSSA external LSA.

Default

N/A.

Usage Guidelines

If advertised, the aggregated IP range is exported as a single LSA by the ABR.

Example

The following command is used to summarize a certain range of IP addresses within an area and export them out as a single address:

```
configure ospf area 1.2.3.4 add range 10.1.2.0/24 advertise type-3
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospf area normal

```
configure ospf area area-identifier normal
```

Description

Configures an OSPF area as a normal area.

Syntax Description

<i>area-identifier</i>	Specifies an OSPF area.
------------------------	-------------------------

Default

Normal.

Usage Guidelines

A normal area is an area that is not any of the following:

- Stub area.
- NSSA.

Virtual links can be configured through normal areas. External routes can be distributed into normal areas.

Example

The following command configures an OSPF area as a normal area:

```
configure ospf area 10.1.0.0 normal
```

History

This command was first available in ExtremeXOS 10.1.



Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospf area nssa stub-default-cost

```
configure ospf area area-identifier nssa [summary | nosummary] stub-default-cost
cost {translate}
```

Description

Configures an OSPF area as an NSSA.

Syntax Description

<i>area-identifier</i>	Specifies an OSPF area.
summary	Specifies that type-3 can be propagated into the area.
nosummary	Specifies that type-3 cannot be propagated into the area.
<i>cost</i>	Specifies a cost metric.
translate	Specifies whether type-7 LSAs are translated into type-5 LSAs.

Default

N/A.

Usage Guidelines

NSSAs are similar to the existing OSPF stub area configuration option, but have the following two additional capabilities:

- External routes originating from an ASBR connected to the NSSA can be advertised within the NSSA.
- External routes originating from the NSSA can be propagated to other areas, including the backbone area, if translated to type 5 LSAs.

When configuring an OSPF area as an NSSA, the translate option should only be used on NSSA border routers, where translation is to be enforced. If translate is not used on any NSSA border router in a NSSA, one of the ABRs for that NSSA is elected to perform translation (as indicated in the NSSA specification). The option should not be used on NSSA internal routers. Doing so inhibits correct operation of the election algorithm.



Example

The following command configures an OSPF area as an NSSA:

```
configure ospf area 10.1.1.0 nssa summary stub-default-cost 10 translate
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospf area stub stub-default-cost

```
configure ospf area area-identifier stub [summary | nosummary] stub-default-cost cost
```

Description

Configures an OSPF area as a stub area.

Syntax Description

<i>area-identifier</i>	Specifies an OSPF area.
summary	Specifies that type-3 can be propagated into the area.
nosummary	Specifies that type-3 cannot be propagated into the area.
<i>cost</i>	Specifies a cost metric.

Default

N/A.

Usage Guidelines

A stub area is connected to only one other area. The area that connects to a stub area can be the backbone area. External route information is not distributed into stub areas. Stub areas are used to reduce memory and computation requirements on OSPF routers.



Example

The following command configures an OSPF area as a stub area:

```
configure ospf area 0.0.0.6 stub nosummary stub-default-cost 10
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospf area timer

```
configure ospf area area-identifier timer retransmit-interval transit-delay  
hello-interval dead-interval [wait-timer-interval]
```

Description

Configures the timers for all interfaces in the same OSPF area.

Syntax Description

<i>area-identifier</i>	Specifies an OSPF area.
<i>retransmit-interval</i>	Specifies the length of time that the router waits before retransmitting an LSA that is not acknowledged. The range is 1- 3,600 seconds.
<i>transit-delay</i>	Specifies the length of time it takes to transmit an LSA packet over the interface. The range is 0 - 3,600 seconds.
<i>hello-interval</i>	Specifies the interval at which routers send hello packets. The range is 1 - 65,535 seconds.
<i>dead-interval</i>	Specifies the interval after which a neighboring router is declared down due to the fact that hello packets are no longer received from the neighbor. The range is 1 - 2,147,483,647 seconds.
<i>wait-timer-interval</i>	Specifies the interval between the interface coming up and the election of the DR and BDR. Usually equal to the dead timer interval.

Default

- *retransmit interval*—Default: 5
- *transit delay*—Default: 1
- *hello interval*—Default: 10



- *dead interval*—Default: 40
- *wait timer interval*—Default: dead interval

Usage Guidelines

Configuring OSPF timers on a per-area basis is a shorthand for applying the timers and authentication to each VLAN in the area at the time of configuration. If you add more VLANs to the area, you must configure the timers and authentication for the new VLANs explicitly.

Specify the following:

- *retransmit interval*—If you set an interval that is too short, unnecessary retransmissions will result.
- *transit delay*—The transit delay must be greater than 0.
- *hello interval*—Smaller times allow routers to discover each other more quickly, but also increase network traffic.
- *dead interval*—This interval should be a multiple of the hello interval.
- *wait timer interval*—This interval is required by the OSPF standard to be equal to the router dead interval. Under some circumstances, setting the wait interval to smaller values can help OSPF routers on a broadcast network to synchronize more quickly at the expense of possibly electing an incorrect DR or BDR. This value should not be set to less than the hello interval. The default value is equal to the router dead interval.

Example

The following command sets the timers in area 0.0.0.2:

```
configure ospf area 0.0.0.2 timer 10 1 20 200
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ospf ase-limit

```
configure ospf ase-limit number {timeout seconds}
```

Description

Configures the AS-external LSA limit and overflow duration associated with OSPF database overflow handling.



Syntax Description

<i>number</i>	Specifies the number of external routes that can be held in a link-state database.
<i>seconds</i>	Specifies a duration for which the system has to remain in the overflow state.

Default

The default for timeout is 0, which indicates that once the router goes into overflow state, it stays there until OSPF is disabled and then re-enabled.

Usage Guidelines

Not applicable.

Example

The following command configures the AS-external LSA limit and overflow duration:

```
configure ospf ase-limit 50000 timeout 1800
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospf ase-summary add

```
configure ospf ase-summary add [ipaddress ip-mask | ipNetmask] cost cost {tag number}
```

Description

Aggregates AS-external routes in a specified address range.

Syntax Description

<i>ipaddress</i>	Specifies an IP address.
<i>ip-mask</i>	Specifies a subnet mask.
<i>ipNetmask</i>	Specifies IP address / Netmask.



<i>cost</i>	Specifies a metric that will be given to the summarized route.
tag	Specifies an OSPF external route tag.

Default

N/A.

Usage Guidelines

This command is only valid on an ASBR.

Example

The following command summarizes AS-external routes:

```
configure ospf ase-summary add 175.1.0.0/16 cost 10
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospf ase-summary delete

```
configure ospf ase-summary delete [ip-address ip-mask | ipNetmask]
```

Description

Deletes an aggregated OSPF external route.

Syntax Description

<i>ip-address</i>	Specifies an IP address.
<i>ip-mask</i>	Specifies a subnet mask.
<i>ipNetmask</i>	Specifies IP address / Netmask.

Default

N/A.



Usage Guidelines

This command is only valid on an ASBR.

Example

The following command deletes the aggregated AS-external route:

```
configure ospf ase-summary delete 175.1.0.0/16
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospf authentication

```
configure ospf [vlan [vlan-name | all ] | area area-identifier | virtual-link
router-identifier area-identifier] authentication [{encrypted} simple-password
simple-password | {encrypted} md5 md5_key_id md5_key| none]
```

Description

Specifies the authentication password (up to eight characters) or RSA Data Security, Inc. MD5 Message-Digest Algorithm key for one or all interfaces in a specific area or a virtual link.

Syntax Description

<i>vlan-name</i>	Specifies a VLAN name.
all	Specifies all VLANs
<i>router-identifier</i>	Specifies the router ID of the remote router.
<i>area-identifier</i>	Specifies an OSPF area.
encrypted	Indicates that the password (or key) is already encrypted (do not use this option).
<i>simple-password</i>	Specifies an authentication password (up to 8 ASCII characters).
<i>md5_key_id</i>	Specifies a RSA Data Security, Inc. MD5 Message-Digest Algorithm key, from 0-255.
<i>md5_key</i>	Specifies a numeric value from 0-65,536. Can also be alphanumeric, up to 26 characters.
none	Disables authentication.



Default

N/A.

Usage Guidelines

The `md5_key` is a numeric value with the range 0 to 65,536 or alphanumeric. When the OSPF area is specified, authentication information is applied to all OSPF interfaces within the area.

The `encrypted` option is used by the switch when generating a configuration file and when parsing a switch-generated configuration file. Do not select the `encrypted` option in the CLI.

Example

The following command configures RSA Data Security, Inc. MD5 Message-Digest Algorithm authentication on the VLAN subnet_26:

```
configure ospf vlan subnet_26 authentication md5 32 test
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ospf bfd

```
configure ospf vlan vlan-name bfd on | off
```

Description

Configures BFD for OSPFv2.

Syntax Description

bfd	Bidirectional forwarding detection.
on	Turn on BFD for OSPF interface.
off	Turn off BFD for OSPF interface.

Default

Off.



Usage Guidelines

Use this command to turn BFD protection on or off on a specific OSPF interface.

The following example configures BFD protection on for VLAN 1:

Example

```
configure ospf vlan1 bfd on
```

History

This command was first available in ExtremeXOS 15.3.2.

Platform Availability

This command is available on all platforms.

configure ospf cost

```
configure ospf [area area-identifier | vlan [vlan-name | all]] cost [automatic | cost]
```

Description

Configures the cost metric of one or all interface(s) or an area.

Syntax Description

<i>area-identifier</i>	Specifies an OSPF area.
<i>vlan-name</i>	Specifies a VLAN name.
all	Specifies all VLANs.
automatic	Determine the advertised cost from the OSPF metric table.
<i>cost</i>	Specifies the cost metric.

Default

The default cost is automatic.

Usage Guidelines

The range is 1 through 65535.



Example

The following command configures the cost metric of the VLAN accounting:

```
configure ospf vlan accounting cost 10
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospf delete virtual-link

```
configure ospf delete virtual-link router-identifier area-identifier
```

Description

Removes a virtual link.

Syntax Description

<i>router-identifier</i>	Specifies the router ID of the other end of the link.
<i>area-identifier</i>	Specifies an OSPF area.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes a virtual link:

```
configure ospf delete virtual-link 10.1.2.1 10.1.0.0
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospf delete vlan

```
configure ospf delete vlan [vlan-name | all]
```

Description

Disables OSPF on one or all VLANs (router interfaces).

Syntax Description

<i>vlan-name</i>	Specifies a VLAN name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

Not applicable.

Example

The following command disables OSPF on VLAN accounting:

```
configure ospf delete vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.



configure ospf import-policy

```
configure ospf import-policy [policy-map | none]
```

Description

Configures the import policy for OSPF.

Syntax Description

<i>policy-map</i>	Specifies the policy.
-------------------	-----------------------

Default

No policy.

Usage Guidelines

An import policy is used to modify route attributes while adding OSPF routes to the IP route table. The import policy cannot be used to determine the routes to be added to the routing table.

Use the **none** option to remove an import policy.

Example

The following example applies the policy *campuseast* to OSPF routes:

```
configure ospf import-policy campuseast
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ospf lsa-batch-interval

```
configure ospf lsa-batch-interval seconds
```

Description

Configures the OSPF LSA batching interval.



Syntax Description

<i>seconds</i>	Specifies a time in seconds.
----------------	------------------------------

Default

The default setting is 30 seconds.

Usage Guidelines

The range is between 0 (disabled) and 600 seconds, using multiples of 5 seconds. The LSAs added to the LSDB during the interval are batched together for refresh or timeout.

Example

The following command configures the OSPF LSA batch interval to a value of 100 seconds:

```
configure ospf lsa-batch-interval 100
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ospf metric-table

```
configure ospf metric-table 10M cost_10m 100M cost_100m 1G cost_1g {10G cost_10g} {40G cost_40g}
```

Description

Configures the automatic interface costs for 10 Mbps, 100 Mbps, and 1 Gbps interfaces, and optionally, the 10 Gbps and 40 Gbps interfaces.

Syntax Description

<i>cost</i>	Specifies the interface cost for the indicated interfaces.
-------------	--



Default

- 10 Mbps—The default cost is 10.
- 100 Mbps—The default cost is 5.
- 1 Gbps—The default cost is 4.
- 10 Gbps—The default cost is 2.
- 40 Gbps—The default cost is 2.

Usage Guidelines

Not applicable.

Example

The following command configures the automatic interface costs for 10 Mbps, 100 Mbps, and 1 Gbps interfaces:

```
configure ospf metric-table 10m 20 100m 10 1g 2
```

History

This command was first available in ExtremeXOS 10.1.

The 40 Gbps parameter was added in ExtremeXOS 12.6.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ospf priority

```
configure ospf [area area-identifier | vlan [vlan-name | all]] priority priority
```

Description

Configures the priority used in the designated router and backup designated router election algorithm for one or all OSPF interface(s) or for all the interfaces within the area.

Syntax Description

<i>area-identifier</i>	Specifies an OSPF area.
<i>vlan-name</i>	Specifies a VLAN name.
all	Specifies all VLANs.
<i>priority</i>	Specifies a priority range. The range is 0 through 255.



Default

The default setting is 1.

Usage Guidelines

The range is 0 through 255, and the default setting is 1. Setting the value to 0 ensures that the router is never selected as the designated router or backup designated router.

Example

The following command sets all the interfaces in area 1.2.3.4 to not be selected as the designated router:

```
configure ospf area 1.2.3.4 priority 0
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ospf restart

```
configure ospf restart [none | planned | unplanned | both]
```

Description

Configures the router as a graceful OSPF restart router.

Syntax Description

none	Do not act as a graceful OSPF restart router.
planned	Only act as a graceful OSPF restart router for planned restarts.
unplanned	Only act as a graceful OSPF restart router for unplanned restarts.
both	Act as a graceful OSPF restart router for both planned and unplanned restarts.

Default

The default is none.



Usage Guidelines

This command configures the router as a graceful OSPF router. When configured for planned restarts, it will advertise Grace-LSAs before restarting (for example, during an upgrade of the OSPF module). When configured for unplanned restarts, it will advertise Grace-LSAs after restarting but before sending any Hellos. When configured for both, the router will advertise restarting regardless of whether the restart was planned or unplanned.

Example

The following command configures a router to perform graceful OSPF restarts only for planned restarts:

```
configure ospf restart planned
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ospf restart grace-period

```
configure ospf restart grace-period seconds
```

Description

Configures the grace period sent out in Grace-LSAs and used by a restarting router.

Syntax Description

<i>seconds</i>	Grace period, in seconds. The default value is 120 seconds. Range is 1 to 1800 seconds.
----------------	---

Default

The default is 120 seconds.

Usage Guidelines

This command configures the grace period sent out to helper neighbor routers and used by the restarting router. The value of the grace period must be greater than the dead interval, and less than the LSA refresh time.



Example

The following command configures a router to send LSAs with a 240 second grace period during graceful OSPF restarts:

```
configure ospf restart grace-period 240
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospf restart-helper

```
configure ospf [vlan [all | vlan-name] | area area-identifier | virtual-link router-identifier area-identifier] restart-helper [none | planned | unplanned | both]
```

Description

Configures the router as a graceful OSPF restart helper router.

Syntax Description

all	Specifies all VLANs
<i>vlan-name</i>	Specifies a VLAN name.
<i>area-identifier</i>	Specifies an OSPF area.
<i>router-identifier</i>	Specifies the router ID of the remote router of the virtual link.
none	Do not act as a graceful OSPF restart helper router.
planned	Only act as a graceful OSPF restart helper router for planned restarts.
unplanned	Only act as a graceful OSPF restart helper router for unplanned restarts.
both	Act as a graceful OSPF restart helper router for both planned and unplanned restarts.

Default

The router default is none.



Usage Guidelines

This command configures the router as a graceful OSPF restart helper router for a single or multiple routers. When the router is acting as a helper, it will continue to advertise the restarting router as if it was fully adjacent.

One OSPF interface may not help more than one restarting router. An OSPF interface may not enter helper mode when the router is performing a graceful restart. All the interfaces to a neighbor router must be configured as graceful restart helpers, or the router will not support graceful restart for its neighbor.

Example

The following command configures a router to be a graceful OSPF helper router for planned restarts for all routers in area 10.20.30.40:

```
configure ospf area 10.20.30.40 restart-helper planned
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospf routerid

```
configure ospf routerid [automatic | router-identifier]
```

Description

Configures the OSPF router ID. If automatic is specified, the switch uses the highest IP interface address as the OSPF router ID.

Syntax Description

automatic	Specifies to use automatic addressing.
<i>router-identifier</i>	Specifies a router address.

Default

Automatic.



Usage Guidelines

Each switch that is configured to run OSPF must have a unique router ID. It is recommended that you manually set the router ID of the switches participating in OSPF, instead of having the switch automatically choose its router ID based on the highest interface IP address. Not performing this configuration in larger, dynamic environments could result in an older link-state database remaining in use.



Note

Do not set the router ID to 0.0.0.0.

Example

The following command sets the router ID:

```
configure ospf routerid 10.1.6.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospf spf-hold-time

```
configure ospf spf-hold-time seconds
```

Description

Configures the minimum number of seconds between Shortest Path First (SPF) recalculations.

Syntax Description

<i>seconds</i>	Specifies a time in seconds. The range is 0 to 300 seconds.
----------------	---

Default

3 seconds.

Usage Guidelines

Not applicable.



Example

The following command configures the minimum number of seconds between Shortest Path First (SPF) recalculations:

```
configure ospf spf-hold-time 6
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospf virtual-link timer

```
configure ospf virtual-link router-identifier area-identifier timer retransmit-  
interval transit-delay hello-interval dead-interval
```

Description

Configures the timers for a virtual link.

Syntax Description

<i>router-identifier</i>	Specifies the router ID of the other end of the link.
<i>area-identifier</i>	Specifies an OSPF area.
<i>retransmit-interval</i>	Specifies the length of time that the router waits before retransmitting an LSA that is not acknowledged. The range is 1 - 3,600 seconds.
<i>transit-delay</i>	Specifies the length of time it takes to transmit an LSA packet over the interface. The range is 0 - 3,600 seconds.
<i>hello-interval</i>	Specifies the interval at which routers send hello packets. The range is 1 - 65,535 seconds.
<i>dead-interval</i>	Specifies the interval after which a neighboring router is declared down due to the fact that hello packets are no longer received from the neighbor. The range is 1 - 2,147,483,647 seconds.

Default

- *retransmit interval*—Default: 5
- *transit delay*—Default: 1
- *hello interval*—Default: 10



- *dead interval*—Default: 40
- *wait timer interval*—Default: dead interval

Usage Guidelines

Configuring OSPF timers on a per-area basis is a shorthand for applying the timers and authentication to each VLAN in the area at the time of configuration. If you add more VLANs to the area, you must configure the timers and authentication for the new VLANs explicitly.

Example

The following command sets the timers on the virtual link in area 0.0.0.2 and remote router ID 6.6.6.6:

```
configure ospf virtual-link 6.6.6.6 0.0.0.2 timer 10 1 20 200
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospf vlan area

```
configure ospf vlan vlan-name area area-identifier
```

Description

Associates a VLAN (router interface) with an OSPF area. By default, all router interfaces are associated with area 0.0.0.0.

Syntax Description

<i>vlan-name</i>	Specifies a VLAN name.
<i>area-identifier</i>	Specifies an OSPF area.

Default

Area 0.0.0.0



Usage Guidelines

Any OSPF network that contains more than one area is required to have an area configured as area 0, also called the backbone. All areas in an autonomous system must be connected to the backbone. When designing networks, you should start with area 0, and then expand into other areas.

The backbone allows summary information to be exchanged between ABRs. Every ABR hears the area summaries from all other ABRs. The ABR then forms a picture of the distance to all networks outside of its area by examining the collected advertisements, and adding in the backbone distance to each advertising router.

When a VLAN is configured to run OSPF, by default you must assign it to an area.

Example

The following command associates the VLAN accounting with an OSPF area:

```
configure ospf vlan accounting area 0.0.0.6
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ospf vlan neighbor add

```
configure ospf vlan vlan-name neighbor add ip-address
```

Description

Configures the IP address of a point-to-point neighbor.

Syntax Description

<i>vlan-name</i>	Specifies a VLAN name.
<i>ip-address</i>	Specifies an IP address.

Default

N/A.



Usage Guidelines

None.

Example

The following command configures the IP address of a point-to-point neighbor:

```
configure ospf vlan accounting neighbor add 10.0.0.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospf vlan neighbor delete

```
configure ospf vlan vlan-name neighbor delete ip-address
```

Description

Deletes the IP address of a point-to-point neighbor.

Syntax Description

<i>vlan-name</i>	Specifies a VLAN name.
<i>ip-address</i>	Specifies an IP address.

Default

N/A.

Usage Guidelines

None.



Example

The following command deletes the IP address of a point-to-point neighbor:

```
configure ospf vlan accounting neighbor delete 10.0.0.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospf vlan timer

```
configure ospf vlan [vlan-name | all] timer retransmit-interval transit-delay  
hello-interval dead-interval {wait-timer-interval}
```

Description

Configures the OSPF wait interval for a VLAN or all VLANs.

Syntax Description

<i>vlan-name</i>	Specifies a VLAN name.
all	Specifies all VLANs.
<i>retransmit-interval</i>	Specifies the length of time that the router waits before retransmitting an LSA that is not acknowledged. The range is 1 - 3,600.
<i>transit-delay</i>	Specifies the length of time it takes to transmit an LSA packet over the interface. The range is 0 - 3,600 seconds.
<i>hello-interval</i>	Specifies the interval at which routers send hello packets. The range is 1 - 65,535 seconds.
<i>dead-interval</i>	Specifies the interval after which a neighboring router is declared down due to the fact that hello packets are no longer received from the neighbor. The range is 1 - 2,147,483,647.
<i>wait-timer-interval</i>	Specifies the interval between the interface coming up and the election of the DR and BDR. Usually equal to the dead timer interval.

Default

- *retransmit interval*—5 seconds.
- *transit delay*—1 second.
- *hello interval*—10 seconds.



- *dead interval*—40 seconds.
- *wait timer interval*—dead interval.

Usage Guidelines

Specify the following:

- *retransmit interval*—If you set an interval that is too short, unnecessary retransmissions will result.
- *transit delay*—The transit delay must be greater than 0.
- *hello interval*—Smaller times allow routers to discover each other more quickly, but also increase network traffic.
- *dead interval*—This interval should be a multiple of the hello interval.
- *wait timer interval*—This interval is required by the OSPF standard to be equal to the router dead interval. Under some circumstances, setting the wait interval to smaller values can help OSPF routers on a broadcast network to synchronize more quickly at the expense of possibly electing an incorrect DR or BDR. This value should not be set to less than the hello interval. The default value is equal to the router dead interval.

Example

The following command configures the OSPF wait interval on the VLAN accounting:

```
configure ospf vlan accounting timer 10 15 20 60 60
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

create ospf area

```
create ospf area area-identifier
```

Description

Creates an OSPF area.

Syntax Description

<i>area-identifier</i>	Specifies an OSPF area.
------------------------	-------------------------



Default

Area 0.0.0.0.

Usage Guidelines

Area 0.0.0.0 does not need to be created. It exists by default.

Example

The following command creates an OSPF area:

```
create ospf area 1.2.3.4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

delete ospf area

```
delete ospf area [area-identifier | all]
```

Description

Deletes an OSPF area or all OSPF areas.

Syntax Description

<i>area-identifier</i>	Specifies an OSPF area.
all	Specifies all areas.

Default

N/A.

Usage Guidelines

An OSPF area cannot be deleted if it has an associated interface. Also, area 0.0.0.0 cannot be deleted.



Example

The following command deletes an OSPF area:

```
delete ospf area 1.2.3.4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable ospf

disable ospf

Description

Disables the OSPF process for the router.

Syntax Description

This command has no keywords or arguments.

Default

N/A.

Usage Guidelines

Not applicable.

Example

The following command disables the OSPF process for the router:

```
disable ospf
```

History

This command was first available in ExtremeXOS 10.1.



Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable ospf capability opaque-lsa

```
disable ospf capability opaque-lsa
```

Description

Disables opaque LSAs across the entire system.

Syntax Description

This command has no keywords or arguments.

Default

Enabled.

Usage Guidelines

Opaque LSAs are a generic OSPF mechanism used to carry auxiliary information in the OSPF database. Opaque LSAs are most commonly used to support OSPF traffic engineering.

Normally, support for opaque LSAs is auto-negotiated between OSPF neighbors. In the event that you experience interoperability problems, you can disable opaque LSAs.

If your network uses opaque LSAs, all routers on your OSPF network should support opaque LSAs. Routers that do not support opaque LSAs do not store or flood them. At minimum a well-interconnected subsection of your OSPF network needs to support opaque LSAs to maintain reliability of their transmission.

On an OSPF broadcast network, the designated router (DR) must support opaque LSAs or none of the other routers on that broadcast network will reliably receive them. You can use the OSPF priority feature to give preference to an opaque-capable router, so that it becomes the elected DR.

For transmission to continue reliably across the network, the backup designated router (BDR) must also support opaque LSAs.

Example

The following command disables opaque LSAs across the entire system:

```
disable ospf capability opaque-lsa
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

disable ospf export

```
disable ospf export [bgp | direct | e-bgp | i-bgp | rip | static | isis | isis-level-1 | isis-level-1-external | isis-level-2 | isis-level-2-external]
```

Description

Disables redistribution of routes to OSPF.

Syntax Description

bgp	Specifies BGP routes.
direct	Specifies direct routes.
e-bgp	Specifies E-BGP routes.
i-bgp	Specifies I-BGP routes.
rip	Specifies RIP routes.
static	Specifies static routes.
isis	Specifies IS-IS routes.
isis-level-1	Specifies ISIS Level 1 routes.
isis-level-1-external	Specifies ISIS Level 1 External routes.
isis-level-2	Specifies ISIS Level 2 routes.
isis-level-2-external	Specifies ISIS Level 2 External routes.

Default

The default setting is disabled.

Usage Guidelines

Use this command to stop OSPF from exporting routes derived from other protocols.



Example

The following command disables OSPF to export BGP-related routes to other OSPF routers:

```
disable ospf export bgp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

disable ospf originate-default

```
disable ospf originate-default
```

Syntax Description

There are no keywords or variables for this command.

Default

Not applicable.

Usage Guidelines

Not applicable.

Example

The following command disables generating a default external LSA:

```
disable ospf originate-default
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.



disable ospf restart-helper-lsa-check

```
disable ospf [vlan [all | vlan-name] | area area-identifier | virtual-link
router-identifier area-identifier] restart-helper-lsa-check
```

Description

Disables the restart helper router from terminating graceful OSPF restart when received LSAs would affect the restarting router.

Syntax Description

all	Specifies all VLANs.
<i>vlan-name</i>	Specifies a VLAN name.
<i>router-identifier</i>	Specifies the router ID of the remote router of the virtual link.
<i>area-identifier</i>	Specifies an OSPF area.

Default

The default is enabled.

Usage Guidelines

This command disables the restart helper router from terminating graceful OSPF restart when received LSAs would affect the restarting router.

Example

The following command disables a router from terminating graceful OSPF restart for all routers in area 10.20.30.40 if it receives an LSA that would affect routing:

```
disable ospf area 10.20.30.40 restart-helper-lsa-check
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.



disable ospf use-ip-router-alert

```
disable ospf use-ip-router-alert
```

Description

Disables the router alert IP option in outgoing OSPF control packets.

Syntax Description

This command has no keywords or arguments.

Default

Disabled.

Usage Guidelines

Not applicable.

Example

The following command disables the OSPF router alert IP option:

```
disable ospf use-ip-router-alert
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable snmp traps ospf

```
disable snmp traps ospf
```

Description

Disables the OSPF module from sending traps on various OSPF events.



Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables the OSPF process:

```
disable snmp traps ospf
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable ospf

enable ospf

Description

Enables the OSPF process for the router.

Syntax Description

This command has no keywords or arguments.

Default

N/A.

Usage Guidelines

Not applicable.



Example

The following command enables the OSPF process for the router:

```
enable ospf
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

enable ospf capability opaque-lsa

```
enable ospf capability opaque-lsa
```

Description

Enables opaque LSAs across the entire system.

Syntax Description

This command has no keywords or variables.

Default

Enabled.

Usage Guidelines

Opaque LSAs are a generic OSPF mechanism used to carry auxiliary information in the OSPF database. Opaque LSAs are most commonly used to support OSPF traffic engineering.

Normally, support for opaque LSAs is auto-negotiated between OSPF neighbors. In the event that you experience interoperability problems, you can disable opaque LSAs.

If your network uses opaque LSAs, all routers on your OSPF network should support opaque LSAs. Routers that do not support opaque LSAs do not store or flood them. At minimum a well-interconnected subsection of your OSPF network needs to support opaque LSAs to maintain reliability of their transmission.

On an OSPF broadcast network, the designated router (DR) must support opaque LSAs or none of the other routers on that broadcast network will reliably receive them. You can use the OSPF priority feature to give preference to an opaque-capable router, so that it becomes the elected DR.



For transmission to continue reliably across the network, the backup designated router (BDR) must also support opaque LSAs.

Example

The following command enables opaque LSAs across the entire system:

```
enable ospf capability opaque-lsa
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable ospf export

```
enable ospf export [bgp | direct | e-bgp | i-bgp | rip | static | isis | isis-level-1 | isis-level-1-external | isis-level-2 | isis-level-2-external] [cost cost type [ase-type-1 | ase-type-2] {tag number} | policy-map]
```

Description

Enables redistribution of routes to OSPF.

Syntax Description

bgp	Specifies BGP routes.
direct	Specifies direct routes.
e-bgp	Specifies E-BGP routes.
i-bgp	Specifies I-BGP routes.
rip	Specifies RIP routes.
static	Specifies static routes.
isis	Specifies IS-IS routes.
isis-level-1	Specifies ISIS Level 1 routes.
isis-level-1-external	Specifies ISIS Level 1 External routes.
isis-level-2	Specifies ISIS Level 2 routes.
isis-level-2-external	Specifies ISIS Level 2 External routes.
cost	Specifies a cost metric.



ase-type-1	Specifies AS-external type 1 routes.
ase-type-2	Specifies AS-external type 2 routes.
<i>number</i>	Specifies a tag value.
<i>policy-map</i>	Specifies a policy.

Default

The default tag number is 0. The default setting is disabled.

Usage Guidelines

After OSPF export is enabled, the OSPF router is considered to be an ASBR. Interface routes that correspond to the interface that has OSPF enabled are ignored.

The cost metric is inserted for all BGP, IS-IS, RIP-learned, static, and direct routes injected into OSPF. If the cost metric is set to 0, the cost is inserted from the route. The tag value is used only by special routing applications. Use 0 if you do not have specific requirements for using a tag. The tag value in this instance has no relationship with 802.1Q VLAN tagging.

The same cost, type, and tag values can be inserted for all the export routes, or a policy can be used for selective insertion. When a policy is associated with the export command, the policy is applied on every exported route. The exported routes can also be filtered using a policy.

Example

The following command enables OSPF to export BGP-related routes using LSAs to other OSPF routers:

```
enable ospf export bgp cost 1 ase-type-1 tag 0
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable ospf originate-default

```
enable ospf originate-default {always} cost cost type [ase-type-1 | ase-type-2]
{tag number}
```



Description

Enables a default external LSA to be generated by OSPF, if no other default route is originated by OSPF by way of RIP and static route re-distribution.

Syntax Description

always	Specifies for OSPF to always advertise the default route.
<i>cost</i>	Specifies a cost metric.
ase-type-1	Specifies AS-external type 1 routes.
ase-type-2	Specifies AS-external type 2 routes.
<i>number</i>	Specifies a tag value.

Default

N/A.

Usage Guidelines

If **always** is specified, OSPF always advertises the default route. If **always** is not specified, OSPF adds the default LSA if a reachable default route is in the route table.

Example

The following command generates a default external type-1 LSA:

```
enable ospf originate-default cost 1 ase-type-1 tag 0
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable ospf restart-helper-lsa-check

```
enable ospf [vlan [all | vlan-name] | area area-identifier | virtual-link router-  
identifier area-identifier] restart-helper-lsa-check
```



Description

Enables the restart helper router to terminate graceful OSPF restart when received LSAs would affect the restarting router.

Syntax Description

all	Specifies all VLANs
<i>vlan-name</i>	Specifies a VLAN name.
<i>router-identifier</i>	Specifies the router ID of the remote router of the virtual link.
<i>area-identifier</i>	Specifies an OSPF area.

Default

The default is enabled.

Usage Guidelines

This command configures the restart helper router to terminate graceful OSPF restart when received LSAs would affect the restarting router. This will occur when the restart-helper receives an LSA that will be flooded to the restarting router or when there is a changed LSA on the restarting router's retransmission list when graceful restart is initiated.

Example

The following command configures a router to terminate graceful OSPF restart for all routers in area 10.20.30.40 if it receives an LSA that would affect routing:

```
enable ospf area 10.20.30.40 restart-helper-lsa-check
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

enable ospf use-ip-router-alert

```
enable ospf use-ip-router-alert
```



Description

Enables the generation of the OSPF router alert IP option.

Syntax Description

This command has no keywords or arguments.

Default

Disabled.

Usage Guidelines

Not applicable.

Example

The following command enables the OSPF router alert IP option:

```
enable ospf use-ip-router-alert
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable snmp traps ospf

```
enable snmp traps ospf [all | trap-map bit-map]
```

Description

Enables the OSPF module to send traps on various OSPF events.

Syntax Description

all	Sets RFC1850 ospfSetTrap to 0x1ffff.
trap-map	Specifies the ospfSetTrap as defined in RFC1850.
<i>bit-map</i>	Specifies the ospfSetTrap value in HEX (for example, 0x1ffff for all traps).



Default

The default is disabled.

Usage Guidelines

This command enables the OSPF module to send traps on various OSPF events.

Example

The following command sets `ospfSetTrap` for all traps:

```
enable snmp traps ospf all
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show ospf

show ospf

Description

Displays global OSPF information.

Syntax Description

This command has no keywords or arguments.

Default

N/A.

Usage Guidelines

Not applicable.



Example

The following command displays global OSPF information:

```
show ospf
```

The following is sample output from this command:

```

OSPF                : Disabled                MPLS LSP as Next-Hop: No
RouterId            : 0.0.0.0                RouterId Selection  : Automatic
ASBR                : No                    ABR                : No
ExtLSA              : 0                    ExtLSAChecksum     : 0x0
OriginateNewLSA    : 0                    ReceivedNewLSA     : 0
SpfHoldTime        : 3                    Lsa Batch Interval : 30s
CapabilityOpaqueLSA : Enabled
10M Cost            : 10                    100M Cost          : 5
1000M Cost (1G)    : 4                    10000M Cost (10G) : 2
40000M Cost (40G) :                        : 1
Router Alert       : Disabled                Import Policy File  :
ASExternal LSA Limit : Disabled                Timeout (Count)    : Disabled (0)
Originate Default  : Enabled                Always : Yes Type: 2 Cost: 10 Tag: 0
SNMP Traps         : Enabled                SNMP Trap Bit Map  : 0xffff
Redistribute:
Protocol            Status  cost  Type  Tag      Policy
direct              Disabled 0     0     0        None
static              Disabled 0     0     0        None
rip                 Disabled 0     0     0        None
e-bgp               Disabled 0     0     0        None
i-bgp               Disabled 0     0     0        None
isis-level-1        Disabled 0     0     0        None
isis-level-2        Disabled 0     0     0        None
isis-level-1-external Disabled 0     0     0        None
isis-level-2-external Disabled 0     0     0        None

```

History

This command was first available in ExtremeXOS 10.1.

The SNMP Traps and 40G parameters were added in ExtremeXOS 12.6.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

show ospf area

```
show ospf area {detail | area-identifier}
```



Description

Displays information about the OSPF area.

Syntax Description

detail	Specifies to display the information in detailed format.
<i>area-identifier</i>	Specifies an OSPF area.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays information about OSPF areas:

```
show ospf area
```

The following is sample output from this command:

```
BD-8810Rack3.2 # show ospf area
AREA ID      Type Summ  Def   Num  Num  SPF  Num  LSA
Metric ABR  ASBR Runs LSAs  Checksum
0.0.0.0      NORM ---- - 0    0    0    0    0x0
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

show ospf ase-summary

```
show ospf ase-summary
```

Description

Displays the OSPF external route aggregation configuration.



Syntax Description

this command has no keywords or arguments.

Default

N/A.

Usage Guidelines

Not applicable.

Example

The following command displays the OSPF external route aggregation configuration:

```
show ospf ase-summary
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show ospf interfaces

```
show ospf interfaces {vlan vlan-name | area area-identifier | enabled}
```

Description

Displays information about one or all OSPF interfaces.

Syntax Description

<i>vlan-name</i>	Specifies a VLAN name.
<i>area-identifier</i>	Specifies an OSPF area.
enabled	Displays only OSPF enabled interfaces.

Default

If no argument is specified, all OSPF interfaces are displayed.



Usage Guidelines

Not applicable.

Example

The following command displays information about one or all OSPF interfaces on the VLAN accounting:

```
show ospf interfaces vlan accounting
```

The following output displays BFD protection configuration information:

```
# show ospf interfaces

VLAN      IP Address      AREA ID      Flags      Cost State
Neighbors
HQ_10_0_2  10.0.2.2        /24 0.0.0.0   -rifb--    4/A DR     1
HQ_10_0_5  10.0.5.2        /24 0.0.0.0   -rif---    4/A BDR     1

Flags:  b - BFD protection configured, D - Duplicate address detected on
VLAN,
        f - Interface Forwarding Enabled, i - Interface OSPF Enabled,
        n - Multinetted VLAN, p - Passive Interface, r - Router OSPF
Enable,
        T - Tentative address.
Cost:   A - Automatic Cost, C - Configured Cost.

Total number of interfaces: 2
```

The following output displays the BFD session state:

```
Interface(rif1000027): 10.0.2.2/24 Vlan: HQ_10_0_2 OSPF: ENABLED Router:
ENABLED
AreaId: 0.0.0.0 RtId: 10.0.2.2 Link Type: broadcast(auto) Passive: No
Cost: 4/A Priority: 10 Transit Delay: 1 DAD State:Valid
Hello Interval: 10s Rtr Dead Time: 40s Retransmit Interval: 5s
Wait Timer: 40s
Authentication: NONE
State: DR Number of events: 1
DR RtId: 10.0.2.2 DR IP addr: 10.0.2.2 BDR IP addr: 10.0.2.1
Num Neighbor State Change to FULL : 1
BFD Protection: On

Neighbors:
  RtrId: 10.0.3.1 IpAddr: 10.0.2.1 Pri: 5 Type: Auto
  State: FULL Dr: 10.0.2.2 BDR: 10.0.2.1 Dead Time: 00:00:00:03
  Options (0x42): Opaque LSA: Yes
  BFD Session State: Active
```

History

This command was first available in ExtremeXOS 10.1.



The enabled option was added in ExtremeXOS 12.2.

BFD display output was added in 15.3.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show ospf interfaces detail

show ospf interfaces detail

Description

Displays detailed information about all OSPF interfaces.

Syntax Description

detail	Specifies to display the information in detailed format.
---------------	--

Default

N/A.

Usage Guidelines

Not applicable.

Example

The following command displays information about all OSPF interfaces:

```
show ospf interfaces detail
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



show ospf lsdb

```
show ospf lsdb {detail | stats} {area [area-identifier | all]} {{lstype} [lstype
| all]} {lsid lsid-address{lsid-mask}} {routerid routerid-address {routerid-
mask}} {interface[[ip-address{ip-mask} | ipNetmask] | vlan vlan-name]}
```

Description

Displays a table of the current Link-State Database (LSDB).

Syntax Description

detail	Specifies to display all fields of matching LSAs in a multi-line format.
stats	Specifies to display the number of matching LSAs, but not any of their contents.
<i>area-identifier</i>	Specifies an OSPF area.
all	Specifies all OSPF areas.
<i>lstype</i>	Specifies an LS type.
lsid	Specifies an LS ID.
<i>lsid-mask</i>	Specifies an LS ID mask.
<i>routerid-address</i>	Specifies a LSA router ID address.
interface	Specifies to display interface types.
<i>vlan-name</i>	Specifies a VLAN name.

Default

Display in summary format.

Usage Guidelines

ExtremeXOS provides several filtering criteria for the show ospf lsdb command. You can specify multiple search criteria and only the results matching all of the criteria are displayed. This allows you to control the displayed entries in large routing tables.

A common use of this command is to omit all optional parameters, resulting in the following shortened form:

```
show ospf lsdb
```

The shortened form displays all areas and all types in a summary format.

You can filter the display using either the area ID, the remote router ID, or the link-state ID. The default setting is all with no detail. If detail is specified, each entry includes complete LSA information.



Example

The following command displays all areas and all types in a summary format:

```
show ospf lsdb
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

show ospf memory

```
show ospf memory {detail | memoryType}
```

Description

Displays OSPF specific memory usage.

Syntax Description

detail	Displays detail information.
<i>memoryType</i>	Specifies the memory type usage to display.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays OSPF specific memory for all types:

```
show ospf memory detail
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

show ospf neighbor

```
show ospf neighbor {routerid [ip-address {ip-mask} | ipNetmask]} {vlan vlan-name}
{detail}
```

Description

Displays information about an OSPF neighbor.

Syntax Description

<i>ip-address</i>	Specifies an IP address.
<i>ip-mask</i>	Specifies a subnet mask.
<i>ipNetmask</i>	Specifies IP address / Netmask.
<i>vlan-name</i>	Specifies a VLAN name.
detail	Specifies detail information.

Default

If no argument is specified, all OSPF neighbors are displayed.

Usage Guidelines

Not applicable.

Example

The following command displays information about the OSPF neighbors on the VLAN accounting:

```
show ospf neighbor vlan accounting
```

The following command output displays BFD protection status of all OSPF neighbors:

```
# show ospf neighbor
```

```
Neighbor ID      Pri State                Up/Dead Time
Address          Interface
```



```

          BFD Session State
=====
=====
160.26.26.2      10 FULL      /BDR      10:16:42:57/00:00:00:00  160.26.26.2
CHI_160_26_26
          Disabled
10.0.2.2        10 FULL      /BDR      07:17:55:29/00:00:00:09  10.0.2.2
HQ_10_0_2
          Active
10.0.3.2        10 FULL      /BDR      07:17:54:56/00:00:00:03  10.0.3.2
HQ_10_0_3
          Error (Session Limit Exceeded)

```

Total number of neighbors: 3 (All neighbors in Full state)

```
# show ospf neighbor {vlan} <vlan-name>
```

```

Neighbor ID      Pri State          Up/Dead Time
Address          Interface
          BFD Session State
=====
=====
10.0.3.2         1 FULL      /BDR      00:11:13:06/00:00:00:04
12.0.2.2         v2
          Active

```

Total number of neighbors: 1 (All neighbors in Full state)

```
# show ospf neighbor detail
```

```

Neighbor 10.0.3.2, interface address 12.0.2.2
  In the area 0.0.0.0 via interface v2
  Neighbor priority is 1, State is INIT,38 state changes
  DR is 12.0.2.1 BDR is 12.0.2.2
  Options is 0x42
  Neighbor is up for 00:11:04:05
  Time since last Hello 00:00:00:00
  Retransmission queue length is 0
  BFD Session State: None

```

```
# show ospfv3 neighbor
```

```

Neighbor ID      Pri State          Up/Dead Time
Interface        InstanceID
          BFD Session State
=====
=====
1.1.1.1          1  FULL      /BDR      00:03:40:45/00:00:38 HQ_10_0_4      0
          Active

```

```
# show ospfv3 neighbor detail
```

```

Neighbor 1.1.1.1, Interface address fe80::201:30ff:fe10:3ae6
  In the area 0.0.0.0 via interface HQ_10_0_4
  Neighbor priority is 1, State is FULL, 1338 events, 6 state changes

  DR is 2.2.2.2 BDR is 1.1.1.1
  Options is 0x13 (-|R|-|-|E|V6)

```




```
Neighbor is up for 00:03:42:17
Neighbor will be dead in 00:00:37
Retransmission queue length is 0
BFD Session State: Active
```

```
# show ospfv3 neighbor {vlan} <vlan-name>
```

```
Neighbor ID      Pri State           Up/Dead Time      Interface
InstanceID
      BFD Session State
=====
=====
1.1.1.1          1  FULL   /BDR             00:20:37:17/00:00:39 HQ_10_0_4      0
      Active
3.3.3.3          1  FULL   /DR              00:20:37:17/00:00:39 HQ_10_0_4      0
      Active
4.4.4.4          1  2WAY   /DOTHER          00:20:37:17/00:00:39 HQ_10_0_4      0
      None
```

History

This command was first available in ExtremeXOS 10.1.

BFD output was added in ExtremeXOS 15.3.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show ospf virtual-link

```
show ospf virtual-link {router-identifier area-identifier}
```

Description

Displays virtual link information about a particular router or all routers.

Syntax Description

<i>router-identifier</i>	Specifies a router interface number.
<i>area-identifier</i>	Specifies an OSPF area.

Default

N/A.



Usage Guidelines

The area-identifier refer to the transit area used for connecting the two end-points. The transit area cannot have an area identifier of 0.0.0.0 and cannot be a stub or NSSA area.

Example

The following command displays virtual link information about a particular router:

```
show ospf virtual-link 1.2.3.4 10.1.6.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure ospf

```
unconfigure ospf {vlan vlan-name | area area-identifier}
```

Description

Resets one or all OSPF interfaces to the default settings.

Syntax Description

<i>vlan-name</i>	Specifies a VLAN name.
<i>area-identifier</i>	Specifies an OSPF area.

Default

N/A.

Usage Guidelines

ExtremeXOS OSPF allows you to change certain configurable OSPF parameters on the fly. This command selectively resets the configurable parameters to their default values. Following is the list of parameters whose values will be reset to their defaults:

Interface

- Hello interval.
- Dead interval.



- Transmit delay.
- Retransmit interval.
- Priority.
- Cost.
- OSPF graceful restart helper mode.

Area

- All the parameters of interfaces associated with this area.
- Inter-Area-Prefix_LSA Filter.
- AS-External-LSA Filter.

OSPF Global

- All parameters of all areas in this OSPF domain.
- SPF delay interval.
- Interface cost metric table.
- Route redistribution.
- OSPF graceful restart.

Example

The following command resets the OSPF interface to the default settings on the VLAN accounting:

```
unconfigure ospf accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#), in the *ExtremeXOS Concepts Guide*.



39 OSPFv3 Commands

Licensing

OSPF Edge Mode

clear ospfv3 counters

configure ospfv3 add interface

configure ospfv3 add interface all

configure ospfv3 add virtual-link

configure ospfv3 area add range

configure ospfv3 area cost

configure ospfv3 area delete range

configure ospfv3 area external-filter

configure ospfv3 area interarea-filter

configure ospfv3 area normal

configure ospfv3 area priority

configure ospfv3 area stub

configure ospfv3 area timer

configure ospfv3 bfd

configure ospfv3 delete interface

configure ospfv3 delete virtual-link

configure ospfv3 import-policy

configure ospfv3 interface area

configure ospfv3 interface cost

configure ospfv3 interface priority

configure ospfv3 interface timer

configure ospfv3 metric-table

configure ospfv3 routerid

configure ospfv3 spf-hold-time

configure ospfv3 virtual-link timer

create ospfv3 area

delete ospfv3 area

disable ospfv3

disable ospfv3 export

enable ospfv3

enable ospfv3 export

show ospfv3

show ospfv3 area

show ospfv3 interfaces

show ospfv3 lsdb

```
show ospfv3 lsdb stats
show ospfv3 memory
show ospfv3 neighbor
show ospfv3 virtual-link
unconfigure ospfv3
```

This chapter describes commands used for the IPv6 interior gateway protocol OSPFv3. Open Shortest Path First (OSPFv3) is a link-state protocol that distributes routing information between routers belonging to a single IP domain, also known as an autonomous system (AS). In a link-state routing protocol, each router maintains a database describing the topology of the autonomous system. Each participating router in an area has an identical database maintained from the perspective of that router.

OSPFv3 supports IPv6, and uses commands only slightly modified from that used to support IPv4. OSPFv3 has retained the use of the four-byte, dotted decimal numbers for router IDs, LSA IDs, and area IDs.

From the link-state database (LSDB), each router constructs a tree of shortest paths, using itself as the root. The shortest path tree provides the route to each destination in the autonomous system. When several equal-cost routes to a destination exist, traffic can be distributed among them. The cost of a route is described by a single metric.

OSPFv3 allows parts of a network to be grouped together into areas. The topology within an area is hidden from the rest of the autonomous system. Hiding this information enables a significant reduction in link-state advertisement (LSA) traffic, and reduces the computations needed to maintain the LSDB. Routing within the area is determined only by the topology of the area.

The three types of routers defined by OSPFv3 are as follows:

- **Internal Router (IR)**—An internal router has all of its interfaces within the same area.
- **Area Border Router (ABR)**—An ABR has interfaces belonging to two or more areas. It is responsible for exchanging summary advertisements with other OSPFv3 routers.
- **Autonomous System Border Router (ASBR)**—An ASBR acts as a gateway between OSPFv3 and other routing protocols, or other autonomous systems.

Each switch that is configured to run OSPFv3 must have a unique router ID. It is recommended that you manually set the router ID of the switches participating in OSPFv3, instead of having the switch automatically choose its router ID based on the highest interface IPv4 address, since your router may not have an IPv4 address. Not performing this configuration in larger, dynamic environments could result in an older LSDB remaining in use.

**Note**

Do not set the router ID to 0.0.0.0.

Licensing

See [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide* for information about licensing requirements.



OSPF Edge Mode

OSPF Edge Mode is a subset of OSPF available on platforms with an Advanced Edge license. There are two restrictions on OSPF Edge Mode:

- At most, four Active OSPF VLAN interfaces are permitted. There is no restriction on the number of Passive interfaces.
- The OSPF Priority on VLANs is zero, and is not configurable. This prevents the system from acting as a DR or BDR.

clear ospfv3 counters

```
clear ospfv3 {domain domainName} counters { interfaces [[vlan | tunnel] all |
vlan vlan_name | tunnel tunnel_name | area area_identifier] | area [all |
area area_identifier] | virtual-link [all | {routerid} router-identifier {area}
area area_identifier] | neighbor [all | routerid router-identifier | vlan vlan_name |
tunnel tunnel_name]| system}
```

Description

Clears the OSPFv3 counters (statistics).

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
all	Specifies all VLANs, tunnels, areas, neighbors, or virtual-links.
<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
<i>router-identifier</i>	Specifies a router identifier, a four-byte, dotted decimal number.
<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
system	Specifies the OSPFv3 system/global counters.

Default

N/A.

Usage Guidelines

The global command clear counters also clears all OSPFv3 counters. This global command is the equivalent of clear ospfv3 counters for OSPFv3.

This command can be used to clear various OSPFv3 counters (Interface, Area, Virtual-Link, System etc.). The following is the list of various counters that would be reset to zero by this command:

- Neighbor specific counters:



- Number of state changes.
- Number of events.
- Interface/VLAN/Virtual-link/Tunnel specific counters:
 - Number of Hellos rxed.
 - Number of Hellos txed.
 - Number of DB Description rxed.
 - Number of DB description txed.
 - Number of LS request rxed.
 - Number of LS request txed.
 - Number of LS update rxed.
 - Number of LS update txed.
 - Number of LS ack rxed.
 - Number of LS ack txed.
 - Number of rxed OSPFv3 packet discarded.
 - Number of state changes.
 - Number of events.
- Area Specific counters:
 - All counters of interfaces associated with an area.
 - Number of SPF runs.
- Domain (global)/system specific counters:
 - Number of self originated LSAs.
 - Number of received LSAs.

Example

The following command clears the OSPFv3 counters for area 1.1.1.1:

```
clear ospfv3 counters area 1.1.1.1
```

The following command clears all the OSPFv3 counters for the neighbor 192.168.0.1 in the domain ospf-core:

```
clear ospfv3 domain ospf-core counters neighbor routerid 192.168.0.1
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure ospfv3 add interface

```
configure ospfv3 {domain domainName} add [vlan vlan_name | tunnel tunnel_name]
{instance-id instanceId} area area_identifier link-type [auto | broadcast |
point-to-point] {passive}
```

Syntax Description

Enables OSPFv3 on an interface.

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
<i>instanceId</i>	Specifies the instance ID for this interfaces. Range is 0 to 255.
<i>area_identifier</i>	Specifies the area to which the VLAN is assigned.
auto	Specifies to automatically determine the OSPFv3 link type based on the interface type.
broadcast	Specifies a broadcast link, such as Ethernet. Routers must elect a DR and a BDR during synchronization. (This option is not currently supported.)
point-to-point	Specifies a point-to-point link type, such as PPP. (This option is not currently supported.)
passive	Specifies to stop sending and receiving hello packets on this interface.

Default

The default link-type is Auto.

The default instance ID is 0.

Usage Guidelines

This command is used to enable the OSPFv3 protocol on an IPv6 configured VLAN or an IPv6 tunnel. The instance ID is used to control the selection of other routers as neighbors. The router will become a neighbor only with routers that have the same instance ID.

An interface can have only one instance ID associated with it in one OSPFv3 domain. However, the same interface can be associated with another OSPFv3 domain with a different instance ID. An interface associated with two OSPFv3 domains cannot have same instance ID.

To change the instance ID associated with an interface, you must first remove the interface from the OSPFv3 area and then add it back with a different instance ID.

The passive parameter indicates that the router only synchronizes and listens, and does not originate or send any new information on the interface.



Enable IPv6 forwarding before enabling OSPFv3, otherwise, you will receive a warning message.



Note

Configuration of the link-type parameter is not supported. OSPFv3 will always consider the link-type to be broadcast.

Example

The following command adds the VLAN accounting (enabling OSPFv3 on the interface), to the area 0.0.0.1 with an instance ID of 2:

```
configure ospfv3 add vlan accounting instance-id 2 area 0.0.0.1 link-type auto
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ospfv3 add interface all

```
configure ospfv3 {domain domainName} add [vlan | tunnel] all {instance-id instanceId} area area_identifier {passive}
```

Description

Enables OSPFv3 on all VLANs or all tunnels (router interfaces).

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
all	Specifies all IPv6 configured VLANs or all IPv6 tunnels.
<i>instanceId</i>	Specifies the instance ID for these interfaces. Range is 0 to 255.
<i>area_identifier</i>	Specifies the area to which the interfaces are assigned.
passive	Specifies to stop sending and receiving hello packets on this interface.

Default

OSPFv3 is disabled on the interfaces.



The default instance ID is 0.

Usage Guidelines

This command is used to enable the OSPFv3 protocol on all IPv6 configured VLANs or all IPv6 tunnels. The instance ID is used to control the selection of other routers as neighbors. The router will become a neighbor only with routers that have the same instance ID.

An interface can have only one instance ID associated with it in one OSPFv3 domain. However, the same interface can be associated with another OSPFv3 domain with a different instance ID. An interface associated with two OSPFv3 domains cannot have same instance ID.

To change the instance ID associated with an interface, you must first remove the interface from the OSPFv3 area and then add it back with a different instance ID.

The passive parameter indicates that the router only synchronizes and listens, and does not originate or send any new information on the interface.

Example

The following command enables OSPFv3 on all IPv6 tunnels:

```
configure ospfv3 add tunnel all area 0.0.0.1
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ospfv3 add virtual-link

```
configure ospfv3 {domain domainName} add virtual-link {routerid}  
router_identifier {area} area_identifier
```

Description

Adds a virtual link connected to another ABR.



Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>router_identifier</i>	Specifies the router ID of the other end of the link.
<i>area_identifier</i>	Specifies the transit area identifier, a four-byte, dotted decimal number.

Default

N/A.

Usage Guidelines

A virtual link provides a logical path between the ABR of the disconnected area and the ABR of the normal area that connects to the backbone. A virtual link must be established between two ABRs that have a common area, with one ABR connected to the backbone. Specify the following:

- *router_identifier*—Far-end router identifier, a four-byte, dotted decimal number.
- *area_identifier*—Transit area used for connecting the two end-points. The transit area cannot have the area identifier 0.0.0.0. and cannot be a stub area or an NSSA.

Example

The following command configures a virtual link with router ID 10.1.2.1 through the transit area 10.1.0.0:

```
configure ospfv3 add virtual-link 10.1.2.1 10.1.0.0
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ospfv3 area add range

```
configure ospfv3 {domain domainName} area area_identifier add range ipv6netmask  
[advertise | noadvert] inter-prefix
```

Description

Configures a range of IP addresses in an OSPFv3 area to be aggregated.



Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
<i>ipv6netmask</i>	Specifies an IPv6 address / prefix length.
advertise	Specifies to advertise the aggregated range of IP addresses.
noadvert	Specifies not to advertise the aggregated range of IP addresses.
inter-prefix	Specifies aggregate, inter-area-prefix LSAs.

Default

No OSPFv3 inter-area-prefix LSAs are configured.

Usage Guidelines

If advertised, the aggregated IP range is exported as a single LSA by the ABR.

Example

The following command is used to summarize a certain range of IP addresses within an area and export them out as a single address to area 0.0.0.1:

```
configure ospfv3 area 0.0.0.1 add range 2aaa:456:3ffe::/64 advertise inter-
prefix
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ospfv3 area cost

```
configure ospfv3 {domain domainName} area area_identifier cost [automatic | cost]
```

Description

Configures the cost of sending a packet to all interfaces belonging to an area.



Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
automatic	Determine the advertised cost from the OSPFv3 metric table.
<i>cost</i>	Specifies the cost metric. Range is 1 to 65535.

Default

The default cost is automatic. The default domain is OSPF-Default.

Usage Guidelines

Use this command to set the cost of the links belonging to area manually, if the default cost needs to be overwritten. The interface cost is advertised as the link cost in router-LSA.

Example

The following command configures the cost of area 0.0.0.1 to 10. All the links of this area will inherit the area's cost value of 10.

```
configure ospfv3 domain ospf-enterprise area 0.0.0.1 cost 10
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospfv3 area delete range

```
configure ospfv3 {domain domainName} area area_identifier delete range  
ipv6netmask
```

Description

Removes a range of IP addresses in an OSPFv3 area to be aggregated.



Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
<i>ipv6netmask</i>	Specifies an IPv6 address / prefix length.

Default

No OSPFv3 inter-area-prefix LSAs are configured.

Usage Guidelines

If you attempt to delete a range that was not configured, you will receive an error message.

Example

The following command is used to delete a summary network from area 0.0.0.1:

```
configure ospfv3 area 0.0.0.1 delete range 2aaa:456:3ffe::/64
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospfv3 area external-filter

```
configure ospfv3 {domain domainName} area area_identifier external-filter  
[policy_map | none]
```

Description

Configures an external filter policy.

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>area_identifier</i>	Specifies the OSPFv3 target area.
<i>r</i>	



<i>policy_map</i>	Specifies a policy.
none	Specifies not to apply an external filter (removes the existing policy, if any).

Default

N/A.

Usage Guidelines

For switches configured to support multiple OSPFv3 areas (an ABR function), a policy can be applied to an OSPFv3 area that filters a set of OSPFv3 external routes from being advertised into that area, in other words, filtering some of the inbound AS-external-LSAs.

OSPFv3 routers that do not have enough memory to hold the entire AS-external-LSAa should configure an external area filter to drop part of the external-LSAs. Configuring this policy will enable routers with limited resources to be put into an OSPFv3 network.

Using the none mode specifies that no external filter is applied.

Policy files for this command will only recognize the following policy attributes:

- Match attributes:
 - *nlri IPv6-address/mask-len*
- Action (set) attributes
 - permit
 - deny

Any other policy attribute will not be recognized and will be ignored.

The following is an example of an external area policy file:

```
entry one {
    if match any{
        nlri 2001:db8:3e5c::/48;
        nlri 2001:db8:2146:2341::/64;
    } then {
        deny;
    }
}
```

Example

The following command configures an external filter policy, nosales for area 1.2.3.4:

```
configure ospfv3 area 1.2.3.4 external-filter nosales
```



History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospfv3 area interarea-filter

```
configure ospfv3 {domain domainName} area area_identifier interarea-filter
[policy_map | none]
```

Description

Configures an inter-area filter policy.

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>area_identifier</i>	Specifies the OSPFv3 target area.
<i>policy_map</i>	Specifies a policy.
none	Specifies not to apply an inter-area filter (removes the existing policy, if any).

Default

N/A.

Usage Guidelines

ExtremeXOS OSPFv3 can apply an inter-area policy to filter some inter-area-prefix-LSAs and inter-area-router-LSAs from other areas. This can reduce the size of link state database of routers belonging to the area.

Using the none mode specifies that no external filter is applied.

Policy files for this command will only recognize the following policy attributes:

- Match attributes:
 - nlri *IPv6-address/mask-len*
- Action (set) attributes:
 - permit
 - deny

Any other policy attribute will not be recognized and will be ignored.



The following is an example of an external area policy file:

```
entry one {
  if match any{
    nlri 2001:db8:3e5c::/48;
    nlri 2001:db8:2146:2341::/64;
  } then {
    deny;
  }
}
entry two {
  if match any{
    nlri 2001:db8:444::/48;
    nlri 2001:db8:541f:65bd::/64;
  } then {
    permit;
  }
}
```

Example

The following command configures an inter-area filter policy, nosales for area 1.2.3.4:

```
configure ospfv3 area 1.2.3.4 interarea-filter nosales
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospfv3 area normal

```
configure ospfv3 {domain domainName} area area_identifier normal
```

Description

Configures an OSPFv3 area as a normal area.

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.



Default

Normal.

Usage Guidelines

A normal area is an area that is not any of the following:

- Stub area
- NSSA

Virtual links can be configured through normal areas. External routes can be distributed into normal areas.

Example

The following command configures an OSPFv3 area as a normal area:

```
configure ospfv3 area 10.1.0.0 normal
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ospfv3 area priority

```
configure ospfv3 {domain domainName} area area_identifier priority priority
```

Description

Configures the priority used in the designated router and backup designated router election algorithm for all the interfaces within the area.

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
<i>priority</i>	Specifies a priority range. The range is 0 through 255.



Default

The default setting is 1.

Usage Guidelines

When two routers are attached to a network, both attempt to become the designated router. The one with the higher priority takes precedence. If there is a tie, the router with the higher router ID takes precedence. Setting the value to 0 ensures that the router is never selected as the designated router or backup designated router.

Example

The following command sets all the interfaces in area 1.2.3.4 to not be selected as the designated router:

```
configure ospfv3 area 1.2.3.4 priority 0
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospfv3 area stub

```
configure ospfv3 {domain domainName} area area_identifier stub [summary | nosummary] stub-default-cost cost
```

Description

Configures an OSPFv3 area as a stub area.

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
summary	Specifies that inter-area LSAs can be propagated into the area.
nosummary	Specifies that inter-area LSAs cannot be propagated into the area.
<i>cost</i>	Specifies a cost metric.



Default

N/A.

Usage Guidelines

A stub area is connected to only one other area. The area that connects to a stub area can be the backbone area. External route information is not distributed into stub areas. Stub areas are used to reduce memory consumption requirements on OSPFv3 routers.

Example

The following command configures an OSPFv3 area as a stub area:

```
configure ospfv3 area 0.0.0.6 stub nosummary stub-default-cost 10
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ospfv3 area timer

```
configure ospfv3 {domain domainName} area area_identifier timer {retransmit-interval} retransmit_interval {transit-delay} transit_delay {hello-interval} hello_interval {dead-interval} dead_interval
```

Description

Configures the timers for all interfaces in the same OSPFv3 area.

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
<i>retransmit_interval</i>	Specifies the length of time that the router waits before retransmitting an LSA that is not acknowledged. The range is 1 to 3600 seconds.
<i>transit_delay</i>	Specifies the length of time it takes to transmit an LSA packet over the interface. The range is 1 to 3600 seconds.



<i>hello_interval</i>	Specifies the interval at which routers send hello packets. The range is 1 to 65535 seconds.
<i>dead_interval</i>	Specifies the interval after which a neighboring router is declared down due to the fact that hello packets are no longer received from the neighbor. The range is 1 to 65535 seconds.

Default

- Retransmit interval—Default: 5 seconds
- Transit delay—Default: 1 second
- Hello interval—Default: 10 seconds
- Dead interval—Default: 40 seconds

Usage Guidelines

Configuring OSPFv3 timers on a per-area basis is a shorthand for applying the timers to each VLAN and tunnel in the area at the time of configuration. If you add more VLANs or tunnels to the area, you must configure the timers for them explicitly.

Specify the following:

- Retransmit interval—If you set an interval that is too short, unnecessary retransmissions will result.
- Transit delay—The transit delay must be greater than 0.
- Hello interval—Smaller times allow routers to discover each other more quickly, but also increase network traffic.
- Dead interval—This interval should be a multiple of the hello interval.

The value of the dead interval and the hello interval must be same for all OSPFv3 routers connected to a common link. The value of the dead interval and the hello interval are advertised by OSPFv3 in Hello packets. The shorter the hello interval, the earlier topological changes will be detected, but more routing traffic will ensue.

The retransmit interval must be greater than the expected round trip delay between any two routers on the attached network. The setting of this parameter must be conservative, or needless retransmission will result.



Note

The wait interval for the interface is not separately configurable. It is always equal to the dead interval.

Example

The following command sets the timers in area 0.0.0.2:

```
configure ospfv3 area 0.0.0.2 timer 10 1 20 200
```



History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospfv3 bfd

```
configure ospfv3 vlan vlan-name bfd on | off
```

Description

Configures BFD for OSPFv3.

Syntax Description

bfd	Bidirectional forwarding detection
on	Turn on BFD for OSPFv3 interface.
off	Turn off BFD for OSPFv3 interface.

Default

Off.

Usage Guidelines

Use this command to turn on or off BFD protection on a specific OSPFv3 interface.

The following example configures BFD protection on for VLAN 1:

Example

```
configure ospfv3 vlan1 bfd on
```

History

This command was first available in ExtremeXOS 15.3.2.

Platform Availability

This command is available on all platforms.



configure ospfv3 delete interface

```
configure ospfv3 {domain domainName} delete [vlan vlan_name | tunnel tunnel_name
| [vlan | tunnel] all]
```

Description

Disables OSPFv3 on one or all VLANs or tunnels (router interfaces).

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
all	Specifies all VLANs, or tunnels.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables OSPFv3 on VLAN accounting:

```
configure ospfv3 delete vlan accounting
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospfv3 delete virtual-link

```
configure ospfv3 {domain domainName} delete virtual-link {routerid}
router_identifier {area} area_identifier
```



Description

Deletes a virtual link connected to another ABR.

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>router_identifier</i>	Specifies the router ID of the other end of the link.
<i>area_identifier</i>	Specifies the transit area identifier, a four-byte, dotted decimal number.

Default

N/A.

Usage Guidelines

A virtual link provides a logical path between the ABR of the disconnected area and the ABR of the normal area that connects to the backbone. A virtual link must be established between two ABRs that have a common area, with one ABR connected to the backbone. Specify the following:

- Router-identifier—Far-end router identifier, a four-byte, dotted decimal number.
- Area-identifier—Transit area used for connecting the two end-points. The transit area cannot have the area identifier 0.0.0.0. and cannot be a stub area or an NSSA.

Example

The following command deletes a virtual link with router ID 10.1.2.1 through the transit area 10.1.0.0:

```
configure ospfv3 delete virtual-link 10.1.2.1 10.1.0.0
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ospfv3 import-policy

```
configure ospfv3 {domain domainName} import-policy [policy_map | none]
```



Description

Configures the import policy for OSPFv3.

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>policy_map</i>	Specifies the policy.

Default

No policy.

Usage Guidelines

An import policy is used to modify route attributes while adding OSPFv3 routes to the IP route table. The import policy cannot be used to determine the routes to be added to the routing table.

Use the none option to remove the policy association.

Policy files for this command will recognize only the following policy attributes:

- Match attributes:
 - nlr *IPv6-address/mask-len*
 - route-origin [ospfv3 | ospfv3-extern1 | ospfv3-extern2 | ospfv3-inter | ospfv3-intra]
- Action (set) attributes
 - cost *cost*
 - tag *number*

Any other policy attribute will not be recognized and will be ignored.

Example

The following example applies the policy campuseast to OSPFv3 routes:

```
configure ospfv3 import-policy campuseast
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.



configure ospfv3 interface area

```
configure ospfv3 {domain domainName} [vlan vlan_name | tunnel tunnel_name] area
area_identifier
```

Description

Moves an interface from one OSPFv3 area to another.

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.

Default

Area 0.0.0.0.

Usage Guidelines

Use this command to move an already configured interface from one area to another. The instance ID associated with the interface will be unchanged.

Example

The following command moves the VLAN accounting to the OSPFv3 area 0.0.0.6:

```
configure ospfv3 vlan accounting area 0.0.0.6
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure ospfv3 interface cost

```
configure ospfv3 {domain domainName} [vlan vlan_name | tunnel tunnel_name | [vlan
| tunnel] all]] cost [automatic | cost]
```

Description

Configures the cost of one or all interface(s).

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
all	Specifies all IPv6 configured VLANs or all IPv6 tunnels.
automatic	Determine the advertised cost from the OSPFv3 metric table.
<i>cost</i>	Specifies the cost metric. Range is 1 to 65535.

Default

The default cost is automatic.

Usage Guidelines

Use this command to set the cost of an interface (a VLAN or tunnel) manually, if the default cost needs to be overwritten. The interface cost is advertised as the link cost in router-LSA.

Example

The following command configures the cost metric of the VLAN accounting:

```
configure ospfv3 vlan accounting cost 10
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.



configure ospfv3 interface priority

```
configure ospfv3 {domain domainName} [vlan vlan_name | tunnel tunnel_name | [vlan
| tunnel] all] priority priority
```

Description

Configures the priority used in the designated router and backup designated router election algorithm for one or all OSPFv3 interface(s).

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
all	Specifies all IPv6 configured VLANs or all IPv6 tunnels.
<i>priority</i>	Specifies a priority range. The range is 0 through 255.

Default

The default setting is 1.

Usage Guidelines

When two routers are attached to a network, both attempt to become the designated router. The one with the higher priority takes precedence. If there is a tie, the router with the higher router ID takes precedence. Setting the value to 0 ensures that the router is never selected as the designated router or backup designated router.

Example

The following command sets the priority of the interface VLAN corporate to 10:

```
configure ospfv3 domain ospf-internal vlan corporate priority 10
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure ospfv3 interface timer

```
configure ospfv3 {domain domainName} [vlan vlan_name | tunnel tunnel_name | [vlan
| tunnel] all] timer {retransmit-interval} retransmit_interval {transit-delay}
transit_delay {hello-interval} hello_interval {dead-interval} dead_interval
```

Description

Configures the timers for all interfaces in the same OSPFv3 area.

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
all	Specifies all IPv6 configured VLANs or all IPv6 tunnels.
<i>retransmit_interval</i>	Specifies the length of time that the router waits before retransmitting an LSA that is not acknowledged. The range is 1 to 3600 seconds.
<i>transit_delay</i>	Specifies the length of time it takes to transmit an LSA packet over the interface. The range is 1 to 3600 seconds.
<i>hello_interval</i>	Specifies the interval at which routers send hello packets. The range is 1 to 65535 seconds.
<i>dead_interval</i>	Specifies the interval after which a neighboring router is declared down due to the fact that hello packets are no longer received from the neighbor. The range is 1 to 65535 seconds.

Default

- Retransmit interval—Default: 5 seconds.
- Transit delay—Default: 1 second.
- Hello interval—Default: 10 seconds.
- Dead interval—Default: 40 seconds.

Usage Guidelines

Use this command to configure the OSPFv3 timers on a per-interface basis.

Specify the following:

- retransmit interval—If you set an interval that is too short, unnecessary retransmissions will result.
- transit delay—The transit delay must be greater than 0.
- hello interval—Smaller times allow routers to discover each other more quickly, but also increase network traffic.
- dead interval—This interval should be a multiple of the hello interval.



The value of the dead interval and the hello interval must be same for all OSPFv3 routers connected to a common link. The value of the dead interval and the hello interval are advertised by OSPFv3 in Hello packets. The shorter the hello interval, the earlier topological changes will be detected, but more routing traffic will ensue.

The retransmit interval must be greater than the expected round trip delay between any two routers on the attached network. The setting of this parameter must be conservative, or needless retransmission will result.



Note

The wait interval for the interface is not separately configurable. It is always equal to the dead interval.

Example

The following command sets the timers for the VLAN corporate:

```
configure ospfv3 domain ospf-default vlan corporate timer retransmit-interval
10 transit-delay 2 hello-interval 20 dead-interval 80
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospfv3 metric-table

```
configure ospfv3 {domain domainName} metric-table 10M cost_10m 100M cost_100m 1G
cost_1g {10G cost_10g} {40G cost_40g}
```

Description

Configures the automatic interface costs for 10 Mbps, 100 Mbps, and 1 Gbps interfaces, and optionally, the 10 Gbps and 40 Gbps interfaces.

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>cost_x</i>	Specifies the interface cost for the indicated interfaces. Range is 1 to 65535.



Default

- 10 Mbps—The default cost is 100.
- 100 Mbps—The default cost is 50.
- 1 Gbps—The default cost is 40.
- 10 Gbps—The default cost is 20.
- 40 Gbps—The default cost is 20.

Usage Guidelines

The value of the costs cannot be greater for higher speed interfaces. In other words, the following condition must be true:

```
cost_10m >= cost_100m >= cost_1g >= cost_10g >= cost_40g
```

Example

The following command configures the automatic interface costs for 10 Mbps, 100 Mbps, and 1 Gbps interfaces:

```
configure ospfv3 metric-table 10m 200 100m 100 1g 20
```

The following example displays the output of the show ospfv3 command:

```
show ospfv3
OSPF Domain Name      : OSPF-Default
OSPFv3                 : Disabled
RouterId Selection    : Automatic
RouterId               : 0.0.0.0
ABR                    : No
ASBR                   : No
ExtLSAs                : 0
ExtLSASummary         : 0
OriginatedNewLSAs     : 0
ReceivedNewLSAs       : 0
SpfHoldTime           : 10s
Num of Areas          : 1
10M Cost               : 100
100M Cost              : 50
1000M Cost (1G)       : 40
10000M Cost (10G)     : 20
40000M Cost (40G)     : 20
Router Alert          : Disabled
ASExternal LSALimit   : Disabled
Timeout (Count)       : Disabled (0)
Originate Default     : Disabled
Import Policy File    : none
Redistribute:
Protocol              Status   Cost   Type  Tag   Policy
direct                Disabled 20     2     ---  none
e-bgp                 Disabled 20     2     ---  none
i-bgp                 Disabled 20     2     ---  none
ripng                 Disabled 20     2     ---  none
static                Disabled 20     2     ---  none
isis-level-1          Disabled 20     2     ---  none
isis-level-2          Disabled 20     2     ---  none
isis-level-1-external Disabled 20     2     ---  none
isis-level-2-external Disabled 20     2     ---  none
```



History

This command was first available in ExtremeXOS 11.2.

The 40 Gbps parameter was added in ExtremeXOS 12.6.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospfv3 routerid

```
configure ospfv3 {domain domainName} routerid [automatic | router_identifier]
```

Description

Configures the OSPFv3 router ID. If automatic is specified, the switch uses the highest IPv4 interface address as the OSPFv3 router ID.

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
automatic	Specifies to use automatic addressing.
<i>router_identifier</i>	Specifies a router identifier, a four-byte, dotted decimal number.

Default

Automatic.

Usage Guidelines

Each switch that is configured to run OSPFv3 must have a unique router ID. The router ID is a four-byte, dotted decimal number, like an IPv4 address. Even though the IP address format has changed from IPv4 to IPv6, the router ID format has not. It is recommended that you manually set the router ID of the switches participating in OSPFv3, instead of having the switch automatically choose its router ID based on the highest interface IPv4 address (if it exists). Not performing this configuration in larger, dynamic environments could result in an older link-state database remaining in use.

This command is accepted only when OSPFv3 is globally disabled.



Note

Do not set the router ID to 0.0.0.0.



Example

The following command sets the router ID to 10.1.6.1:

```
configure ospfv3 routerid 10.1.6.1
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospfv3 spf-hold-time

```
configure ospfv3 {domain domainName} spf-hold-time seconds
```

Description

Configures the minimum number of seconds between Shortest Path First (SPF) recalculations.

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>seconds</i>	Specifies a time in seconds. The range is 0 to 300 seconds.

Default

3 seconds.

Usage Guidelines

Setting the interval too high will force OSPFv3 to run SPF calculations less frequently. This will reduce the CPU load, but will cause delay in routes getting updated in the IP routing table. Setting the interval too low will decrease the interval between SPF calculations, but will increase the processing load on CPU.



Example

The following command configures the minimum number of seconds between Shortest Path First (SPF) recalculations:

```
configure ospfv3 spf-hold-time 6
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

configure ospfv3 virtual-link timer

```
configure ospfv3 {domain domainName} virtual-link {routerid} router_identifier
{area} area_identifier timer {retransmit-interval} retransmit_interval {transit-
delay} transit_delay {hello-interval} hello_interval {dead-interval}
dead_interval
```

Description

Configures the timers for a virtual link.

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>router_identifier</i>	Specifies the router ID of the other end of the link.
<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
<i>retransmit_interval</i>	Specifies the length of time that the router waits before retransmitting an LSA that is not acknowledged. The range is 1 to 3600 seconds.
<i>transit_delay</i>	Specifies the length of time it takes to transmit an LSA packet over the interface. The range is 1 to 3600 seconds.
<i>hello_interval</i>	Specifies the interval at which routers send hello packets. The range is 1 to 65535 seconds.
<i>dead_interval</i>	Specifies the interval after which a neighboring router is declared down due to the fact that hello packets are no longer received from the neighbor. The range is 1 to 65535 seconds.



Default

- Retransmit interval—Default: 5 seconds.
- Transit delay—Default: 1 second.
- Hello interval—Default: 10 seconds.
- Dead interval—Default: 40 seconds.

Usage Guidelines

In OSPFv3, all areas must be connected to a backbone area. If the connection to the backbone is lost, it can be repaired by establishing a virtual link.

The smaller the hello interval, the faster topological changes will be detected, but more routing traffic will ensue.

The setting of the retransmit interval should be conservative, or needless retransmissions will result. The value should be larger for serial lines and virtual links.

The transmit delay value should take into account the transmission and propagation delays for the interface.



Note

The wait interval is not separately configurable. It is always equal to the dead interval.

Example

The following command sets the timers on the virtual link to router 6.6.6.6 transiting area 0.0.0.2:

```
configure ospfv3 virtual-link 6.6.6.6 area 0.0.0.2 timer 10 transit-delay 1
hello-interval 20 dead-interval 200
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

create ospfv3 area

```
create ospfv3 {domain domainName} area area_identifier
```

Description

Creates an OSPFv3 area.



Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.

Default

Area 0.0.0.0.

Usage Guidelines

Area 0.0.0.0 does not need to be created. It exists by default.

Example

The following command creates a non-backbone OSPFv3 area:

```
create ospfv3 area 1.2.3.4
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

delete ospfv3 area

```
delete ospfv3 {domain domainName} area [area_identifier | all]
```

Description

Deletes an OSPFv3 area or all OSPFv3 areas.

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
all	Specifies all areas.



Default

N/A.

Usage Guidelines

An OSPFv3 area cannot be deleted if it has an associated interface. Also, area 0.0.0.0 cannot be deleted.

Example

The following command deletes an OSPFv3 area:

```
delete ospfv3 area 1.2.3.4
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable ospfv3

```
disable ospfv3 {domain domainName}
```

Description

Disables OSPFv3 for the router.

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
-------------------	---

Default

N/A.

Usage Guidelines

None.



Example

The following command disables OSPFv3 for the router:

```
disable ospfv3
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

disable ospfv3 export

```
disable ospfv3 {domain domainName} export [direct | ripng | static | isis | isis-level-1 | isis-level-1-external | isis-level-2 | isis-level-2-external | bgp]
```

Description

Disables redistribution of routes to OSPFv3.

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
direct	Specifies direct routes.
ripng	Specifies RIP routes.
static	Specifies static routes.
isis	Specifies IS-IS routes.
isis-level-1	Specifies IS-IS Level 1 routes.
isis-level-1-external	Specifies IS-IS Level 1 External routes.
isis-level-2	Specifies IS-IS Level 2 routes.
isis-level-2-external	Specifies IS-IS Level 2 External routes.
bgp	Specifies BGP IPv6 routes.

Default

The default setting is disabled.



Usage Guidelines

Use this command to stop OSPFv3 from exporting routes derived from other protocols.

Example

The following command disables OSPFv3 to export RIPng routes to other OSPFv3 routers:

```
disable ospfv3 export ripng
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

enable ospfv3

```
enable ospfv3 {domain domainName}
```

Description

Enables OSPFv3 for the router.

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
-------------------	---

Default

N/A.

Usage Guidelines

When OSPFv3 is enabled, it will start exchanging Hellos on all of its active interfaces. It will also start exporting routes into OSPFv3 routing domain from other protocols, if enabled.

When OSPFv3 is disabled, it will release all the run-time allocated resources like adjacencies, link state advertisements, run-time memory, etc.

OSPFv3 can be enabled successfully if and only if:

- At least one of the VLANs in the current virtual router has one IPv4 address configured



–OR–

- You explicitly configure the OSPFv3 router ID, a four-byte, dotted decimal number

Example

The following command enables OSPFv3 for the router:

```
enable ospfv3
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable ospfv3 export

```
enable ospfv3 {domain domainName} export [direct | ripng | static | isis | isis-  
level-1 | isis-level-1-external | isis-level-2 | isis-level-2-external | bgp]  
[cost cost type [ase-type-1 | ase-type-2] | policy_map]
```

Description

Enables redistribution of routes to OSPFv3.

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
direct	Specifies direct routes.
ripng	Specifies RIPng routes.
static	Specifies static routes.
isis	Specifies IS-IS routes.
isis-level-1	Specifies IS-IS Level 1 routes.
isis-level-1-external	Specifies IS-IS Level 1 External routes.
isis-level-2	Specifies IS-IS Level 2 routes.
isis-level-2-external	Specifies IS-IS Level 2 External routes.
bgp	Specifies BGP IPv6 routes.
<i>cost</i>	Specifies a cost metric.



ase-type-1	Specifies AS-external type 1 routes.
ase-type-2	Specifies AS-external type 2 routes.
number	Specifies a tag value.
<i>policy_map</i>	Specifies a policy.

Default

The default tag number is 0. The default setting is disabled.

Usage Guidelines

The cost metric is inserted for all RIPv3-learned, static, and direct routes injected into OSPFv3. If the cost metric is set to 0, the cost is inserted from the route. The tag value is used only by special routing applications. Use 0 if you do not have specific requirements for using a tag. The tag value in this instance has no relationship with 802.1Q VLAN tagging.



Note

Setting the tag value is not supported in this release.

The same cost, type, and tag values can be inserted for all the export routes, or a policy can be used for selective insertion. When a policy is associated with the export command, the policy is applied on every exported route. The exported routes can also be filtered using a policy.

Policy files for this command will only recognize the following policy attributes:

- Match attributes
 - *nlri IPv6-address/mask-len*
- Action (set) attributes
 - *cost <cost>*
 - *tag <number>*
 - *cost-type [ase-type-1 | ase-type-2]*
 - *permit*
 - *deny*

Any other policy attribute will not be recognized and will be ignored.

The following is an example OSPFv3 export policy file:

```
entry first {
  if match any{
    nlri 2001:db8:200:300:/64;
    nlri 2001:db8:2146:23d1::/64;
    nlri 2001:db8:af31:3d0::/64;
    nlri 2001:db8:f6:2341::/64;
  } then {
    deny;
  }
}
entry second {
```



```

if match any{
    nlri 2001:db8:304::/48;
    nlri 2001:db8:ca11::/48;
    nlri 2001:db8:da36::/48;
    nlri 2001:db8:f6a6::/48;
} then {
    cost 220;
    cost-type ase-type-2;
    permit;
}
}

```

Example

The following command enables OSPFv3 to export RIPng-related routes and associates a policy redistrib:

```
enable ospfv3 export ripng redistrib
```

History

This command was first available in ExtremeXOS 11.2.

The tag keyword was removed in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

show ospfv3

```
show ospfv3 {domain domainName}
```

Description

Displays global OSPFv3 information.

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
-------------------	---

Default

N/A.



Usage Guidelines

None.

Example

The following command displays global OSPFv3 information:

```
show ospfv3
```

The following is sample output:

```

OSPF Domain Name      : OSPF-Default
OSPFv3                : Disabled
RouterId Selection    : Automatic
ABR                   : No
ExtLSAChecksum        : 0x0
ReceivedNewLSAs       : 0
Num of Areas          : 1
100M Cost              : 50
10000M Cost (10G)    : 20
(40G)
Router Alert          : Disabled
ASExternal LSALimit  : Disabled
Originate Default     : Disabled
Import Policy File    : none
Redistribute:
Protocol              Status   Cost   Type  Tag   Policy
direct                Disabled 20     2     ---   none
ripng                  Disabled 20     2     ---   none
static                Disabled 20     2     ---   none
isis-level-1          Disabled 20     2     ---   none
isis-level-2          Disabled 20     2     ---   none
isis-level-1-external Disabled 20     2     ---   none
isis-level-2-external Disabled 20     2     ---   none
bgp                    Disabled 20     2     ---
none

```

History

This command was first available in ExtremeXOS 11.2.

The 40G parameter was added in ExtremeXOS 12.6.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.



show ospfv3 area

```
show ospfv3 {domain domainName} area {area_identifier | detail}
```

Description

Displays information about OSPFv3 areas.

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
detail	Specifies to display the information in detailed format.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays summary information about the OSPFv3 areas:

```
show ospfv3 area
```

The following is sample output:

```

AREA ID      Type Summ  Def   Num  Num  Num  SPF  Num  LSA
Metric ABR  ASBR Intf Runs LSAs  Checksum
0.0.0.0      NORM ---- - 0    0    1    7    7    0x3155b
1.0.0.0      NORM ---- - 1    1    1    6    9    0x4793d
2.0.0.0      NORM ---- - 0    0    1    5   10    0x47174
3.0.0.0      NORM ---- - 1    0    1    3   12    0x420cf
5.0.0.0      NORM ---- - 1    0    1    4   10    0x3b5b1

```

The following command displays information about OSPFv3 area 1.0.0.0:

```
show ospfv3 area 1.0.0.0
```



The following is sample output:

```

Area Identifier      : 1.0.0.0                Type                : NORM
Router ID           : 20.0.0.1              Num of Interfaces   : 1
Spf Runs            : 6                    Num ABRs            : 1
Num ASBRs           : 1                    Num DC-Bit LSAs     : 1
Num Indication LSAs : 1                    Num of DoNotAge LSAs : 1
Num LSAs            : 9                    LSA Chksum          : 0x4793d
Num of Nbrs         : 1                    Num of Virtual Nbrs : 0
Interfaces:
Interface Name      Ospf State   DR ID            BDR ID
to65                E   BDR             0.0.0.65         20.0.0.1
accounts            E   DR              80.0.0.5         0.0.0.0
finance             E   BDR             90.0.0.7         66.0.0.4
engineering         E   ODR             192.168.0.1     165.0.0.3
Corporate           E   ODR             201.0.16.6      204.0.0.1
Inter-Area route Filter: ospfSummPolicy
External route Filter: ospfExtPolicy
Configured Address Ranges:
Addr: fffe:408:1449::/48 Type: 3 Advt: Yes
Addr: ffe0:930:2781::/40 Type: 7 Advt: No

```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show ospfv3 interfaces

```
show ospfv3 {domain domainName} interfaces {vlan vlan_name | tunnel tunnel_name | area area_identifier | detail}
```

Description

Displays information about one or all OSPFv3 interfaces.

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
detail	Specifies to display the information in detailed format.



Default

If no argument is specified, all OSPFv3 interfaces are displayed.

Usage Guidelines

None.

Example

The following command shows a summary of the OSPFv3 interfaces:

```
show ospfv3 interfaces
```

The following is sample output from the command:

```

VLAN/Tunnel IPv6 Address          AREA ID      Flags Cost
State  Neighbors
ixia      4:5:6:7::1/64          2.0.0.0      -rif  5/A
DR        1
to-jmpr   111:222:333::7/48      5.0.0.0      -rif  5/A
BDR       1
to-Zebra  3ffe:506::4/48         0.0.0.0      -rif  5/A
BDR       1
to5       234:567::7/48          3.0.0.0      -rif  5/A
BDR       1
to65     10:203:134:7::7/48     1.0.0.0      -rif  5/A
BDR       1
Flags : (f) Interface Forwarding Enabled, (i) Interface OSPF Enabled,
(p) Passive Interface, (r) Router OSPF Enable.

```

The following command displays information about the OSPFv3 interfaces on the VLAN to5:

```
show ospfv3 interfaces vlan to5
```

The following is sample output:

```

Interface          : to5          Enabled           : ENABLED
Router             : ENABLED       AreaID           : 3.0.0.0
RouterID           : 20.0.0.1     Link Type        : broadcast
Passive            : No           Cost             : 40A
Priority           : 1           Transit Delay    : 1s
Hello Interval     : 10s          Rtr Dead Time    : 40s
Retransmit Interval : 5s          Wait Timer       : 40s
Interface ID       : 63           Instance ID      : 0
State              : BDR          Number of state chg : 2
Hello due in       : 3s           Number of events  : 3
Total Num of Nbrs  : 1           Nbrs in FULL State : 1
Hellos Rxed        : 94          Hellos Txed      : 94
DB Description Rxed : 4           DB Description Txed : 3
LSA Request Rxed   : 1           LSA Request Txed  : 1
LSA Update Rxed    : 8           LSA Update Txed   : 7
LSA Ack Rxed       : 6           LSA Ack Txed      : 5
In Discards        : 0

```



```

DR RtId          : 10.0.0.5          BDR RtId          : 20.0.0.1
DR Interface addr : fe80::280:c8ff:feb9:1cf1
BDR Interface addr : fe80::280:c8ff:feb9:2089
Neighbors:
RtrId: 10.0.0.5  IpAddr: fe80::280:c8ff:feb9:1cf1  Pri: 1  Type: Auto
State: FULL DR: 10.0.0.5  BDR: 20.0.0.1  Dead Time: 00:00:36
Options: 0x13 (-|R|-|-|E|V6)  Opaque LSA: No

```

The following command output shows BFD protection configuration information:

```

# show ospfv3 interfaces
Interface IPV6 Address          AREA ID          Flags
Cost   State   Nbrs
HQ_10_0_4  2000::d00:202/64          0.0.0.0          -rifb 40/A
DR        1
Flags : (b) BFD protection configured, (f) Interface Forwarding Enabled,
        (i) Interface OSPFv3 Enabled, (p) Passive Interface, (r) Router
        OSPFv3 Enable.
Cost   : (A) Automatic cost, (C) Configured cost.

```

The following command output displays the BFD session state:

```

# show ospfv3 interfaces {vlan} <vlan-name>
Interface          : HQ_10_0_4          Enabled          : ENABLED
Router             : ENABLED          AreaID          : 0.0.0.0
RouterID          : 2.2.2.2          Link Type       : broadcast
Passive           : No          Cost            : 40/A
Priority          : 1          Transit Delay   : 1s
Hello Interval    : 10s          Rtr Dead Time   : 40s
Retransmit Interval : 5s          Wait Timer      : 40s
Interface ID      : 50          Instance ID     : 0
State             : DR          Number of state chg : 2
Hello due in      : 4s          Number of events  : 4
Total Num of Nbrs : 1          Nbrs in FULL State : 1
Hellos Rxed      : 1306         Hellos Txed      : 1306
DB Description Rxed : 5          DB Description Txed : 3
LSA Request Rxed  : 1          LSA Request Txed  : 1
LSA Update Rxed   : 18         LSA Update Txed   : 37
LSA Ack Rxed      : 36         LSA Ack Txed      : 17
In Discards       : 0
DR RtId           : 2.2.2.2          BDR RtId       : 1.1.1.1
DR Interface addr  : fe80::201:30ff:fe10:3b16
BDR Interface addr : fe80::201:30ff:fe10:3ae6
BFD Protection     : On

Neighbors:
RtrId: 1.1.1.1  IpAddr: fe80::201:30ff:fe10:3ae6  Pri: 1  Type: Auto
State: FULL DR: 2.2.2.2  BDR: 1.1.1.1  Dead Time: 00:00:40
Options: 0x13 (-|R|-|-|E|V6)  Opaque LSA: No
BFD Session State: Pending

```



History

This command was first available in ExtremeXOS 11.2.

BFD example output was added in 15.3.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

show ospfv3 lsdb

```
show ospfv3 {domain domainName} lsdb {detail} {area [area_identifier | all]
{lstype [router | network | inter-prefix | inter-router | intra-prefix]} | [vlan
[vlan_name | all] | tunnel [tunnel_name | all]] {lstype link} | lstype [as-
external | router | network | inter-prefix | inter-router | intra-prefix | link]}
{lsid lsid_address} {adv-router router_identifier}
```

Description

Displays a table of the current Link-State Database (LSDB).

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
detail	Specifies to display all fields of matching LSAs in a multi-line format.
<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
all	Specifies all OSPFv3 areas, IPv6 configured VLANs, or IPv6 tunnels.
link	Link LSA
as-external	AS External LSA
router	Router LSA
network	Network LSA
inter-prefix	Inter Area Prefix LSA
inter-router	Inter Area Router LSA
intra-prefix	Intra Area Prefix LSA
<i>lsid_address</i>	Specifies the link state ID of the LSA.
<i>router_identifier</i>	Specifies the router identifier of the advertising router.



Default

Display in summary format.

Usage Guidelines

ExtremeXOS provides several filtering criteria for the `show ospfv3 lsdb` command. You can specify multiple search criteria and only the results matching all of the criteria are displayed. This allows you to control the displayed entries in large routing tables.

A common use of this command is to omit all optional parameters, resulting in the following shortened form:

```
show ospfv3 lsdb
```

The shortened form displays all areas and all types in a summary format.

You can filter the display using either the area ID, the remote router ID, or the link-state ID. The default setting is all with no detail. If detail is specified, each entry includes complete LSA information.

Example

The following command displays all areas and all types in a summary format:

```
show ospfv3 lsdb
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show ospfv3 lsdb stats

```
show ospfv3 {domain domainName} lsdb stats {area [area_identifier | all] {lstype
[router | network | inter-prefix | inter-router | intra-prefix]} | [vlan
[vlan_name | all] | tunnel [tunnel_name | all]] {lstype link} | lstype [as-
external | router | network | inter-prefix | inter-router | intra-prefix | link]}
{lsid lsid_address} {adv-router router_identifier}
```

Description

Displays a table of the current Link-State Database (LSDB) statistics.



Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
all	Specifies all OSPFv3 areas, IPv6 configured VLANs, or IPv6 tunnels.
<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
link	Link LSA.
router	Router LSA.
network	Network LSA.
inter-prefix	Inter Area Prefix LSA.
inter-router	Inter Area Router LSA.
intra-prefix	Intra Area Prefix LSA.
as-external	AS External LSA.
<i>lsid_address</i>	Specifies the link state ID of the LSA.
<i>router_identifier</i>	Specifies the router identifier of the advertising router.

Default

Display in summary format.

Usage Guidelines

ExtremeXOS provides several filtering criteria for the `show ospfv3 lsdb stats` command. You can specify multiple search criteria and only the results matching all of the criteria are displayed. This allows you to control the displayed entries in large routing tables.

A common use of this command is to omit all optional parameters, resulting in the following shortened form:

```
show ospfv3 lsdb stats
```

The shortened form displays all areas and all types in a summary format.

You can filter the display using either the area ID, the remote router ID, or the link-state ID. The default setting is all.

Example

The following command displays all areas and all types in a summary format:

```
show ospfv3 lsdb stats
```



History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.

show ospfv3 memory

```
show ospfv3 memory {detail | memoryType}
```

Description

Displays OSPFv3 specific memory usage.

Syntax Description

detail	Displays detail information.
<i>memoryType</i>	Specifies the memory type usage to display.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays OSPFv3 specific memory for all types:

```
show ospfv3 memory detail
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.



show ospfv3 neighbor

```
show ospfv3 {domain domainName} neighbor {routerid ip_address} {vlan vlan_name |
tunnel tunnel_name} {detail}
```

Description

Displays information about an OSPFv3 neighbor.

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>ip_address</i>	Specifies a neighbor router ID.
<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
detail	Specifies detail information.

Default

If no argument is specified, all OSPFv3 neighbors are displayed.

Usage Guidelines

None.

Example

The following command displays information about the OSPFv3 neighbors on the VLAN accounting:

```
show ospfv3 neighbor vlan accounting
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in Feature License Requirements in the *ExtremeXOS Concepts Guide*.



show ospfv3 virtual-link

```
show ospfv3 {domain domainName} virtual-link {{routerid} router_identifier {area}
area_identifier}
```

Description

Displays virtual link(s) information.

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>router_identifier</i>	Specifies a router identifier, a four-byte, dotted decimal number.
<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.

Default

N/A.

Usage Guidelines

Router-identifier—Router ID for the other end of the link.

Area-identifier—Transit area used for connecting the two end-points. The transit area cannot have an area identifier of 0.0.0.0 and cannot be a stub or NSSA area.

Example

The following command displays information about the virtual link to a particular router:

```
show ospfv3 virtual-link 1.2.3.4 10.1.6.1
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



unconfigure ospfv3

```
unconfigure ospfv3 {domain domainName} {area area_identifier | vlan vlan_name | tunnel tunnel_name}
```

Description

Resets one or all OSPFv3 interfaces to the default settings.

Syntax Description

<i>domainName</i>	Specifies an OSPFv3 domain. OSPF-Default is the only one currently supported.
<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel_name</i>	Specifies an IPv6 tunnel.

Default

N/A.

Usage Guidelines

ExtremeXOS OSPFv3 allows you to change certain configurable OSPFv3 parameters on the fly. This command selectively resets the configurable parameters to their default values. The following is the list of parameters whose values will be reset to their defaults:

- Interface:
 - Hello Interval.
 - Dead Interval.
 - Transmit Delay.
 - Retransmit Interval.
 - Priority.
 - Cost.
- Area:
 - All the parameters of Interfaces associated with this area.
 - Inter-Area-Prefix-LSA Filter.
 - AS-External-LSA Filter.
- OSPF Global:
 - All parameters of all areas in this OSPF domain.
 - SPF Delay interval.
 - Interface Cost metric Table.
 - Route Redistribution.



Example

The following command resets the OSPFv3 interface to the default settings on the VLAN accounting:

```
unconfigure ospfv3 accounting
```

The following command unconfigures the parameters of the area 0.0.0.1 (and all its associated interfaces):

```
unconfigure ospfv3 domain ospf-default area 0.0.0.1
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



40 IS-IS Commands

```
clear isis counters
clear isis counters area
clear isis counters vlan
configure isis add vlan
configure isis area add area-address
configure isis area add summary-address
configure isis area area-password
configure isis area delete area-address
configure isis area delete summary-address
configure isis area domain-password
configure isis area interlevel-filter level 1-to-2
configure isis area interlevel-filter level 2-to-1
configure isis area is-type level
configure isis area metric-style
configure isis area overload-bit on-startup
configure isis area system-id
configure isis area timer lsp-gen-interval
configure isis area timer lsp-refresh-interval
configure isis area timer max-lsp-lifetime
configure isis area timer restart
configure isis area timer spf-interval
configure isis area topology-mode
configure isis circuit-type
configure isis delete vlan
configure isis hello-multiplier
configure isis import-policy
configure isis link-type
configure isis mesh
configure isis metric
configure isis password vlan
configure isis priority
configure isis restart
configure isis restart grace-period
configure isis timer csnp-interval
configure isis timer hello-interval
configure isis timer lsp-interval
configure isis timer restart-hello-interval
```



```
configure isis timer retransmit-interval
configure isis wide-metric
create isis area
delete isis area
disable isis
disable isis area adjacency-check
disable isis area dynamic-hostname
disable isis area export
disable isis area export ipv6
disable isis area originate-default
disable isis area overload-bit
disable isis hello-padding
disable isis restart-helper
enable isis
enable isis area adjacency-check
enable isis area dynamic-hostname
enable isis area export
enable isis area export ipv6
enable isis area originate-default
enable isis area overload-bit
enable isis hello-padding
enable isis restart-helper
show isis
show isis area
show isis area summary-addresses
show isis counters
show isis lsdb
show isis neighbors
show isis topology
show isis vlan
unconfigure isis area
unconfigure isis vlan
```

This chapter describes commands for doing the following:

- Configuring IS-IS
- Displaying IS-IS information
- Managing IS-IS

For an introduction to the IS-IS feature, see the *ExtremeXOS Concepts Guide*.

clear isis counters

```
clear isis counters
```



Description

This command clears all IS-IS-related counters in the current virtual router.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command clears all area and VLAN counters.

The following area counters are cleared: corrupted LSPs, LSPDB overloads, manual address from area count, LSP sequence number wraps, LSP sequence number skips, LSP purges, partition changes, and SPF calculations.

The following VLAN counters are cleared: adjacency changes, adjacency initialization failures, rejected adjacencies, ID field length mismatches, maximum area address mismatches, authentication type failures, authentication failures, DIS changes, hello PDU TX and RX count, LSP TX and RX count, CSNP TX and RX count, PSNP TX and RX count, unknown PDU type TX and RX count.

Example

The following command clears all IS-IS counters:

```
clear isis counters
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

clear isis counters area

```
clear isis counters area [area_name | all]
```

Description

This command clears all IS-IS counters for the specified router process or all router processes.



Syntax Description

<i>area_name</i>	Specifies the router process for which counters are cleared.
all	Clears IS-IS counters for all router processes.

Default

N/A.

Usage Guidelines

The following counters are cleared: corrupted LSPs, LSPDB overloads, manual address from area count, LSP sequence number wraps, LSP sequence number skips, LSP purges, partition changes, SPF calculations, authentication type failures, authentication failures, and ID field length mismatches.

Example

The following command clears the IS-IS counters for areax:

```
clear isis counters area areax
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

clear isis counters vlan

```
clear isis counters [vlan all | {vlan} vlan_name]
```

Description

This command clears all IS-IS counters for one or all VLANs.

Syntax Description

vlan all	Clears the counters for all VLANs.
<i>vlan_name</i>	Specifies a single VLAN for which counters are cleared.

Default

N/A.



Usage Guidelines

This command only affects VLANs that have been added to IS-IS router processes. The following counters are cleared: adjacency changes, adjacency initialization failures, rejected adjacencies, ID field length mismatches, maximum area address mismatches, authentication type failures, authentication failures, DIS changes, hello PDU TX and RX count, LSP TX and RX count, CSNP TX and RX count, PSNP TX and RX count, unknown PDU type TX and RX count.

Example

The following command clears the IS-IS counters for all VLANs:

```
clear isis counters vlan all
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis add vlan

```
configure isis add [vlan all | {vlan} vlan_name] area area_name {ipv4 | ipv6}
```

Description

This command associates the specified VLAN interface with the specified IS-IS router process.

Syntax Description

vlan all	Adds all IS-IS eligible VLANs to the router process.
<i>vlan_name</i>	Specifies a single IS-IS eligible VLAN to be added to the router process.
<i>area_name</i>	Identifies the router process to which the VLANs are added.
ipv4 ipv6	Specifies the VLAN IP address type, IPv4 or IPv6, to be added. If you do not specify an IP address type, the VLAN is added for the IPv4 address type. To support both IP address types on the same VLAN, enter the command twice, using a different IP address type each time.

Default

IPv4.



Usage Guidelines

An IS-IS-eligible interface is one that already has the appropriate IP address type (IPv4 or IPv6) address assigned to it. The VLAN must have an IPv4 address assigned to it if `ipv4` is specified or an IPv6 address assigned to it if `ipv6` is specified. In the event that a VLAN address is unconfigured, the interface is automatically removed from the IS-IS router.

VLANs are added to an IS-IS router process to form adjacencies with neighboring IS-IS routers. Hello PDUs are transmitted over these interfaces once the router process is enabled and has a system ID and area address. IP forwarding, IPv6 forwarding, or both must be enabled on the interface. If the router process operates at both L1 and L2, interfaces can be configured to form adjacencies in only a specific level.

Example

The following command adds VLAN `SJvlan` with an IPv4 address type to `areax`:

```
configure isis add SJvlan area areax
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis area add area-address

```
configure isis area area_name add area-address area_address
```

Description

This command adds an IS-IS area address to the specified routing process.

Syntax Description

<i>area_name</i>	Specifies the area name of the IS-IS process to which to add the area address.
<i>area_address</i>	Specifies an IS-IS area address to add to the IS-IS process. The area address can be from 1 to 13 bytes long and must be entered in the following format: 0101.0102.0103.0104.0105.0106.07.

Default

None.



Usage Guidelines

The IS-IS area address defines an L1 or L2 area within an AS. An IS-IS routing process must be assigned at least one area address before it can send or process PDUs. The area address must be configured appropriately. Level 1 routers only form adjacencies with other level 1 routers with at least one area address in common. Multiple area addresses may be configured, which may be desirable during a topological transition. The maximum number of area addresses that can be configured is 3.

Example

The following command assigns area address 0011.03 to areax:

```
configure isis area areax add area-address 0011.03
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis area add summary-address

```
configure isis area area_name add summary-address [ipv4_address_mask | ipv6_address_mask] {level [1 | 2]}
```

Description

This command adds an IPv4 or IPv6 summary address for the specified level on the specified router process.

Syntax Description

<i>area_name</i>	Specifies the router process to which the summary address is to be added.
<i>ipv4_address_mask</i>	Specifies an IPv4 summary address.
<i>ipv6_address_mask</i>	Specifies an IPv6 summary address.
level	Specifies the IS-IS level for the summary address. The level 1 option summarizes level 2 routes leaked to level 1. The level 2 option summarizes level 1 routes that are advertised into level 2.

Default

No summarization.



Usage Guidelines

Route summaries are useful for minimizing the number of LSPs required to describe reachability for an area. The summary address is advertised instead of the actual reachable addresses. This is particularly useful for L1/L2 routers in which the summary address is used in a single LSP instead of including a part or all of the addresses reachable in its level 1 area.

Note that a summary address is only advertised if at least one route matches the summary address. If there is no route present that matches the summary address exactly, a blackhole route is installed for the summary address. If an interlevel filter permits any route matched by the summary address, and that route is present, the summary address is advertised.

If multiple summary addresses are installed in which one or more supersede each other (10.0.0.0/8 and 10.0.0.0/16, for example), only the more specific summary addresses are advertised.

Example

The following command adds an IPv4 summary address to areax:

```
configure isis area areax add summary-address 10.0.0.0/8
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis area area-password

```
configure isis area area_name area-password [none | {encrypted} simple password
{authenticate-snp {tx-only}}]
```

Description

This command sets or clears the password for level 1 LSPs.

Syntax Description

<i>area_name</i>	Specifies the router process to which the password configuration applies.
none	Disables level 1 password authentication.



encrypted simple password	Enables password authentication and specifies that the supplied password is encrypted and must be decrypted prior to placement in a TLV.
authenticate-snp tx-only	Enables password authentication and level 1 SNP authentication. If the tx-only keyword is specified, the password is included in SNPs on transmission, but received SNPs are not authenticated.

Default

None.

Usage Guidelines

Only plain text passwords are supported. Passwords may be up to 254 alphanumeric characters in length. Although passwords are plaintext in the protocol, they are displayed and saved in an encrypted form.

When password authentication is enabled, received packets are authenticated against the configured password and are discarded if the password does not match. Authentication TLVs are included in transmitted level 1 LSPs with a configured password.

Example

The following command configures the password extreme for areax:

```
configure isis area areax area-password simple extreme
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis area delete area-address

```
configure isis area area_name delete area-address area_address
```

Description

This command deletes an area address from the specified routing process.



Syntax Description

<i>area_name</i>	Specifies the area name of the IS-IS process from which to delete the area address.
<i>area_address</i>	Specifies the area address name to delete from the IS-IS process.

Default

None.

Usage Guidelines

If this router process has only one area address configured, this command also causes the routing process to stop sending or processing IS-IS PDUs.

Example

The following command deletes the 0011.03 area address from areax:

```
configure isis area areax delete area-address 0011.03
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis area delete summary-address

```
configure isis area area_name delete summary-address [ipv4_address_mask | ipv6_address_mask] {level [1 | 2]}
```

Description

This command removes the specified IPv4 or IPv6 summary address from the specified router process at the specified level.

Syntax Description

<i>area_name</i>	Specifies the router process from which the summary address is to be deleted.
<i>ipv4_address_mask</i>	Specifies an IPv4 summary address.



<i>ipv6_address_mask</i>	Specifies an IPv6 summary address.
level	Specifies the IS-IS level for the summary address.

Default

No summarization.

Usage Guidelines

Individual reachable addresses that were superseded by the summary address are now advertised in separate LSPs.

Example

The following command deletes an IPv4 summary address from areax:

```
configure isis area areax delete summary-address 10.0.0.0/8
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis area domain-password

```
configure isis area area_name domain-password [none | {encrypted} simple password {authenticate-snp {tx-only}}]
```

Description

This command sets or clears the password for Level 2 LSPs.

Syntax Description

<i>area_name</i>	Specifies the router process for which the password is set or cleared.
none	Disables level 2 password authentication.
encrypted	Specifies that the supplied password is encrypted and must be decrypted prior to using it in a TLV.



<i>password</i>	Specifies a password. Passwords may be up to 254 alphanumeric characters in length.
authenticate-snp tx-only	If the optional <code>authenticate-snp</code> keyword is included, level 2 SNPs are also authenticated on receive and the password is included on transmission. If <code>tx-only</code> is specified, the password is included in SNPs on transmission, but received SNPs are not authenticated.

Default

None.

Usage Guidelines

Packets received are authenticated against the configured password and are discarded if the password does not match. Authentication TLVs are included in transmitted level 2 LSPs with the configured password. Only plain text passwords are supported. Although LSPs contain plain text passwords, passwords are displayed and saved in an encrypted form.

Example

The following command sets the domain password to Extreme:

```
configure isis area areax domain-password simple Extreme
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis area interlevel-filter level 1-to-2

```
configure isis area area_name interlevel-filter level 1-to-2 [policy | none]
{ipv4 | ipv6}
```

Description

This command provides a method of restricting L1 routes from being redistributed into the L2 domain on an L1/L2 router.



Syntax Description

<i>area_name</i>	Specifies the router process for which this configuration change applies.
<i>policy</i>	Specifies a policy to control how L1 routes are redistributed.
none	Removes any previously configured interlevel filters.
ipv4 ipv6	Applies the interlevel filter to IPv4 or IPv6. If neither IPv4 nor IPv6 is specified, this command applies to IPv4.

Default

None.

Usage Guidelines

This command has no effect on level 1-only and level 2-only routers. Normally all L1 routes are redistributed into L2 on an L1/L2 router. Routes are permitted unless explicitly denied in the policy. This command does not necessarily disable level 1 to level 2 redistribution unless the configured policy effectively filters out all routes. For policies, the `nlri match` attribute is supported, and the `permit` and `deny set` attributes are supported.

Example

The following command removes any previously configured interlevel filters in area `areax` for IPv4:

```
configure isis area areax interlevel-filter level 1-to-2 none
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis area interlevel-filter level 2-to-1

```
configure isis area area_name interlevel-filter level 2-to-1 [policy | block-all  
| allow-all] {ipv4 | ipv6}
```

Description

This command enables route leaking from level 2 to level 1 on an L1/L2 router.



Syntax Description

<i>area_name</i>	Specifies the router process for which this configuration change applies.
<i>policy</i>	Specifies a policy to control how L2 routes are leaked to L1.
block-all	Blocks all route leaking.
allow-all	Leaks all routes into level 1.
ipv4 ipv6	Applies the interlevel filter to IPv4 or IPv6. If neither IPv4 nor IPv6 is specified, this command applies to IPv4.

Default

block-all.

Usage Guidelines

When a policy is supplied with this command, all routes are leaked unless explicitly denied in the policy. This command has no effect on level 1-only and level 2-only routers. For policies, the nlri match attribute is supported, and the permit and deny set attributes are supported.

Example

The following command configures areax to leak all level 2 routes to level 1 for IPv4:

```
configure isis area areax interlevel-filter level 2-to-1 allow-all
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis area is-type level

```
configure isis area area_name is-type level [1 | 2 | both-1-and-2]
```

Description

This command configures the specified router process to operate as a level 1, level 2, or level 1/level 2 router.



Syntax Description

area_name	Specifies the router process you are configuring.
level	Specifies the IS-IS operation level for the router.

Default

both-1-and-2.

Usage Guidelines

Adjacencies are only formed with other routers of the same level. In addition, level 1 adjacencies are only formed with other level 1 routers with the same area address.

If there are no other L2 areas, the default is both-1-and-2. If an L2 or L1/L2 area is already present, the default is L1. This is because there can be only one L2 area in each system.

Example

The following command configures the areax router to operate at level 1:

```
configure isis area areax is-type level 1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis area metric-style

```
configure isis area area_name metric-style [[narrow | wide] {transition}] | transition] {level [1 | 2]}
```

Description

This command specifies the metric style for the specified router process and IS-IS level.

Syntax Description

area_name	Specifies the router process for which the metric style is to be configured.
narrow	Specifies the narrow metric style, which uses the 6-bit default metric. Only narrow metrics are encoded in originated TLVs; only narrow SPF calculations are performed.



narrow transition	Specifies the narrow metric style, which uses the 6-bit default metric. Only narrow metrics are encoded in originated TLVs; both narrow and wide SPF calculations are performed.
wide	Specifies the wide metric style, which uses the 24-bit metric specified in RFC 3784. Only wide metrics are encoded in originated TLVs; only wide SPF calculations are performed.
wide transition	Specifies the wide metric style, which uses the 24-bit metric specified in RFC 3784. Only wide metrics are encoded in originated TLVs; both narrow and wide SPF calculations are performed.
transition	Specifies both the narrow and wide metrics. Both narrow and wide metric types are encoded in TLVs; both narrow and wide SPF calculations are performed.
level	Specifies the IS-IS level to which the metric style applies.

Default

Narrow.

Usage Guidelines

Refer to RFC 3787, Section 5.1, for information on how to migrate a network from narrow metric-style to wide metric-style. Note that Section 5.2 is not supported. As a result, each interface's narrow and wide metric values must match while transitioning the metric style. Only when the entire network has transitioned to wide metric style should the interface metrics be configured differently than the configured narrow metric.

Example

The following command configures area *areax* for the narrow metric style:

```
configure isis area areax metric-style narrow
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis area overload-bit on-startup

```
configure isis area area_name overload-bit on-startup [ off | {suppress [external  
| interlevel | all]} seconds]
```



Description

This command enables or disables the overload bit feature while the specified IS-IS process is initializing.

Syntax Description

<i>area_name</i>	Specifies the area name of the IS-IS process for which this feature is to be enabled or disabled.
off	Disables the overload bit feature during initialization.
suppress	Specifies that one or all types of reachability information is to be suppressed or excluded from LSPs during initialization.
external	When included with the suppress option, this specifies that external reachability information is to be excluded from LSPs during initialization.
interlevel	When included with the suppress option, this specifies that interlevel reachability information is to be excluded from LSPs during initialization.
all	When included with the suppress option, this specifies that external and interlevel reachability information is to be excluded from LSPs during initialization.
<i>seconds</i>	Specifies the period (in seconds) during which this feature is enabled at initialization.

Default

Off.

Usage Guidelines

This command configures the overload bit to be set only while the configured router is initializing, and only for the period of time specified. This can be useful to minimize network churn while a new router joins and learns the topology. The suppress options are used during startup if the router process is level 1/level 2 or is running another protocol, such as BGP (in order to wait for the other protocol to converge). Note that in the latter case, there is no signaling between protocols to indicate convergence. Again, this can reduce churn while the topologies are learned during router initialization.

Note



Although `enable isis area area_name overload-bit {suppress [external | interlevel | all]}` and `disable isis area area_name overload-bit` override the overload bit behavior configured by the `configure isis area area_name overload-bit on-startup [off | {suppress [external | interlevel | all]}seconds]` command, the enable and disable commands do not modify the configured parameters.



Example

The following command enables the areax overload bit feature for 15 seconds during initialization:

```
configure isis area areax overload-bit on-startup 15
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis area system-id

```
configure isis area area_name system-id [automatic | system_id]
```

Description

This command configures the system ID for an IS-IS router process.

Syntax Description

<i>area_name</i>	Specifies the area name of the IS-IS process to which to add the system ID.
automatic	Sets the system ID to the system MAC address.
<i>system_id</i>	Specifies the 6-byte system ID using three sets of four hexadecimal digits, where each set is separated by a period. For example: 001B.1F62.1201.

Default

Automatic (system MAC address is used).

Usage Guidelines

The system ID must be a unique ID within the AS. Typically a system MAC address is used as the system ID. Sometimes a combination of one of the router's IP addresses and 2 prefix bytes are used. The assignment of the system ID may vary depending on how the AS is chosen to be administered.

Example

The following example configures an IS-IS system ID for areax:

```
configure isis area areax system-id 001B.1F62.1201
```



History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis area timer lsp-gen-interval

```
configure isis area area_name timer lsp-gen-interval seconds {level[1| 2]}
```

Description

This command configures the minimum time required to wait before regenerating the same LSP.

Syntax Description

<i>area_name</i>	Specifies the router process for which you want to configure the LSP generation interval.
<i>seconds</i>	Specifies the generation level in seconds. The range is 1 to 120 seconds.
level	Specifies the level to which you want to apply the configuration. If neither level 1 nor level 2 is specified, the configuration applies to both levels.

Default

30 seconds.

Usage Guidelines

In link flapping situations in a mesh network, this can greatly reduce the amount of network traffic generated from LSP flooding.

Example

The following command sets the LSP generation interval to a value of 40 seconds:

```
configure isis area areax timer lsp-gen-interval 40
```

History

This command was first available in ExtremeXOS 12.1.



Platform Availability

This command is available on platforms with a Core license.

configure isis area timer lsp-refresh-interval

```
configure isis area area_name timer lsp-refresh-interval seconds
```

Description

This command configures the refresh rate for locally originated LSPs.

Syntax Description

<i>area_name</i>	Specifies the router process for which you are setting the LSP refresh timer.
<i>seconds</i>	Specifies the LSP refresh interval. The range is 1 to 65535 seconds.

Default

900 seconds.

Usage Guidelines

This value should be configured to be less than the maximum LSP lifetime value, which is set with the `configure isis area area_name timer max-lsp-lifetimeseconds` command. Locally originated LSPs are purged and retransmitted at the specified interval regardless of link state.

Example

The following command sets the LSP refresh timer for areax to 1200 seconds:

```
configure isis area areax timer lsp-refresh-interval 1200
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis area timer max-lsp-lifetime

```
configure isis area area_name timer max-lsp-lifetime seconds
```



Description

This command configures the LSP lifetime timer for locally originated LSPs.

Syntax Description

<i>area_name</i>	Specifies the router process for which you want to configure the LSP lifetime timer.
<i>seconds</i>	Specifies the LSP lifetime in seconds. The range is 1 to 65535 seconds.

Default

1200 seconds.

Usage Guidelines

This value should be configured to be greater than the LSP refresh interval, which is set with the `configure isis area area_name timer lsp-refresh-intervalseconds` command. The remaining lifetime value is included in LSPs when they are flooded. Routers age out LSPs from other routers using the remaining lifetime provided in the LSP. If a refreshed version of the LSP is not received before it is aged out, an SPF recalculation occurs, possibly resulting in routing around the router from which the LSP originated.

Example

The following command configures the LSP lifetime timer for 1800 seconds:

```
configure isis area areax timer max-lsp-lifetime 1800
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis area timer restart

```
configure isis area area_name timer restart seconds {level [1 | 2]}
```

Description

This command configures the IS-IS T2 timer for the specified router process and level.



Syntax Description

<i>area_name</i>	Specifies the router process for which the T2 timer configuration applies.
<i>seconds</i>	Specifies the T2 timer value. The range is 5 to 65535 seconds.
level	Specifies the IS-IS level to which this timer configuration applies. If neither level 1 nor level 2 is specified, the configuration applies to both levels.

Default

60 seconds.

Usage Guidelines

The T2 timer is the restart timer for the LSP database for an IS-IS level. If the T2 timer for the respective level expires before the database has been resynchronized, SPF is run for that level.

Example

The following command configures the areax level 1 T2 timer for 90 seconds:

```
configure isis area areax timer restart 90 level 1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis area timer spf-interval

```
configure isis area area_name timer spf-interval seconds {level[1|2]}
```

Description

This command specifies the minimum time to wait between SPF calculations.

Syntax Description

<i>area_name</i>	Specifies the router process for which you are configuring the SPF interval.
<i>seconds</i>	Specifies the minimum time between SPF calculations. The range is 1 to 120 seconds.
level	Specifies the IS-IS level to which the timer configuration applies. If neither level 1 nor level 2 is specified, the configuration applies to both levels.



Default

10 seconds.

Usage Guidelines

This helps prevent switch CPU overloading when a link flap causes several back-to-back SPF calculations.

Example

The following command configures the SPF interval timer for 30 seconds on areax:

```
configure isis area areax timer spf-interval 30
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis area topology-mode

```
configure isis area area_name topology-mode [single | multi | transition] {level [1 | 2]}
```

Description

This command enables or disables use of multi-topology TLVs as specified in draft-ietf-isis-wg-multi-topology-11.

Syntax Description

<i>area_name</i>	Specifies the router process to be configured.
single	Specifies a single topology, where extended TLVs are used in SPF calculation and TLVs.
multi	Specifies a multi topology, where only the multi-topology TLVs are used in SPF calculation and TLVs.
transition	Specifies a transition topology, where both extended and multi-topology TLVs are used in SPF calculation and TLVs. The transition option is useful when migrating a routing domain.
level	For L1/L2 routers, this applies the configuration to IS-IS level 1 or level 2. If the level option is not specified, the configuration applies to both L1 and L2 areas. This option has no affect on L1-only and L2-only routers.



Default

Single.

Usage Guidelines

Multi-topology capability is desirable if both an IPv4 topology and an IPv6 topology exist with different routing paths.

Extreme supports MT IDs 0 and 2 (IPv4 unicast and IPv6 unicast) only.

Example

The following command configures the transition topology mode for areaX:

```
configure isis area areaX topology-mode transition
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis circuit-type

```
configure isis [vlan all | {vlan} vlan_name] circuit-type level [1 | 2 | both-1-and-2]
```

Description

This command configures the circuit type level for one or all IS-IS VLANs.

Syntax Description

vlan all	Applies the selected circuit type to all IS-IS VLANs.
<i>vlan_name</i>	Specifies a single VLAN to which the circuit type configuration applies.
level [1 2 both-1-and-2]	Sets the circuit type level to level 1, level 2, or to both level 1 and level 2.

Default

Both-1-and-2.



Usage Guidelines

Hello PDUs are only sent on the specified level for the selected VLANs. This can be useful for level 1/level 2 routers that are neighbors.

Note that for per-level VLAN configurable parameters L1 and L1/L2, point-to-point interfaces use the level 1 parameters, and L2-only point-to-point interfaces use the L2 parameters.

Example

The following command configures all IS-IS VLANs to use circuit type level 1:

```
configure isis vlan all circuit-type level 1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis delete vlan

```
configure isis delete [vlan all | {vlan} vlan_name] {area area_name} {ipv4 | ipv6}
```

Description

This command removes a VLAN interface from the specified router process.

Syntax Description

vlan all	Deletes all IS-IS VLANs.
<i>vlan_name</i>	Specifies a single VLAN to delete.
<i>area_name</i>	Specifies the router process from which the VLAN is deleted. If you do not specify an IS-IS area, the software deletes the VLAN from the configured IS-IS area.
ipv4 ipv6	Specifies the IP address type for which the VLAN is deleted. If you do not specify an IP address type, the VLAN for the IPv4 address type is deleted. If the VLAN was added as IPv6, the ipv6 option must be used to remove the VLAN. If the VLAN was added as both IPv4 and IPv6, each VLAN IP address type must be deleted with a separate command.



Default

N/A.

Usage Guidelines

The associated adjacency is removed, causing the removal of the corresponding LSP if there is one, and causing an SPF recalculation if the router process is enabled. Hello PDUs are no longer sent on the specified interface. This command applies to IS-IS-enabled VLANs only.

Example

The following command deletes the IPv4 address type for all VLANs in areax:

```
configure isis delete vlan all area areax
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis hello-multiplier

```
configure isis [vlan all | {vlan} vlan_name] hello-multiplier multiplier {level  
[1 | 2]}
```

Description

This command sets the hello multiplier for one or all IS-IS VLANs.

Syntax Description

vlan all	Applies the configuration to all IS-IS VLANs.
<i>vlan_name</i>	Specifies a single VLAN to which the configuration applies.
<i>multiplier</i>	Sets the hello multiplier. The range is 2 to 100.
level [1 2]	Limits the configuration to either level 1 or level 2. If neither level 1 nor level 2 is specified, the configuration applies to both levels.

Default

3.



Usage Guidelines

The hello multiplier is used in conjunction with the hello interval to compute the holding time. The holding time is included in hello PDUs and is calculated by multiplying the hello multiplier by the hello interval. If the hello interval is set to minimal, the holding time is set to 1 second and the hello interval is calculated by dividing 1 second by the hello multiplier. For example, a hello interval of minimal and a hello multiplier of 4 means that the hold interval is set to 250 ms (and the holding time to 1 second). The holding time tells the neighboring router how long to wait before declaring the sending router dead.

Example

The following command sets the SIPvlan hello multiplier to 4:

```
configure isis SIPvlan hello-multiplier 4
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis import-policy

```
configure isis import-policy [policy-map | none]
```

Description

This command applies a policy map for routes imported to the FIB from all IS-IS router processes on this virtual router.

Syntax Description

<i>policy-map</i>	Specifies the policy to apply.
none	Removes any policies assigned to this virtual router.

Default

None.

Usage Guidelines

IS-IS policy files support the following policy match conditions:

- `nlri IPv4-address/mask-len IPv6-address/mask-len`



- `route-origin [isis-level-1 | isis-level-2 | isis-level-1-external | isis-level-2-external]`

IS-IS policy files support the following policy action statements:

- `cost`

Example

The following command applies the IS-IS policy `policy2` to the virtual router:

```
configure isis import-policy policy2
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis link-type

```
configure isis [vlan all | {vlan} vlan_name] link-type [broadcast | point-to-point]
```

Description

This command specifies the link type for one or all IS-IS VLANs.

Syntax Description

 vlan all 	Applies the link type configuration to all IS-IS VLANs.
<i> vlan_name </i>	Specifies a single IS-IS VLAN to which the link type configuration is applied.
 broadcast 	Selects the broadcast link type for the specified VLANs.
 point-to-point 	Selects the point-to-point link type for the specified VLANs.

Default

Broadcast.

Usage Guidelines

On broadcast interfaces, a DIS is elected. There is no DIS election on point-to-point interfaces. If it is known that only two routers will be present on a physical network, it may be desirable to set their connecting interfaces to point-to-point mode. This reduces the overhead associated with DIS election.



and periodic CSNP transmissions and processing. In addition, if the adjacency is both level 1 and level 2, only one set of hello PDUs are sent on a point-to-point interface whereas hello PDUs are sent for both levels on broadcast interfaces. Interfaces in point-to-point mode must have an IP address assigned to them. Unnumbered interfaces are not supported.

For point-to-point interfaces, level 1 parameters apply to L1-only and L1/L2 interfaces. Level 2 parameters apply to L2-only point-to-point interfaces.

Example

The following command configures all IS-IS VLANs to use the broadcast link type:

```
configure isis vlan all link-type broadcast
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis mesh

```
configure isis [vlan all | {vlan} vlan_name] mesh [block-none | block-all | block-group group_id]
```

Description

This command configures LSP flooding behavior for the specified interface.

Syntax Description

vlan all	Applies the configuration to all IS-IS VLANs.
<i>vlan_name</i>	Specifies a single VLAN to which the configuration applies.
block-none	Disables LSP blocking.
block-all	Blocks all LSPs. No LSPs are flooded out of the selected interface.
block-group	Blocks LSPs that contain the specified group ID.
<i>group_id</i>	Specifies a group ID number. The range is 1 to 4294967295.

Default

Block-none.



Usage Guidelines

In a mesh environment, which is a set of fully interconnected point-to-point interfaces, LSP flooding can generate N2 PDUs because no router can tell which routers have and have not received the flooded LSP. By carefully selecting the links over which LSPs are flooded, traffic can be greatly reduced at the cost of some resiliency. Using mesh group IDs instead of a full block (the block-all option) allows a finer granularity of control.

Example

The following command configures blocking on SJvlan for group 5:

```
configure isis SJvlan mesh block-group 5
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis metric

```
configure isis [vlan all | {vlan} vlan_name] metric metric {level[1|2]}
```

Description

This command sets the narrow metric for one or all IS-IS VLANs.

Syntax Description

vlan all	Applies the configuration to all IS-IS VLANs.
<i>vlan_name</i>	Specifies a single VLAN to which the configuration applies.
metric <i>metric</i>	Sets the metric value. The range is 1 to 63.
level [1 2]	Limits the configuration to either level 1 or level 2. If neither level 1 nor level 2 is specified, the configuration applies to both levels.

Default

10.

Usage Guidelines

If narrow metrics are enabled, this value is used in the associated LSPs for the selected VLANs.



Example

The following command sets the narrow metric for all IS-IS VLANs to 15:

```
configure isis vlan all metric 15
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis password vlan

```
configure isis [vlan all | {vlan} vlan_name] password [none | {encrypted} simple  
password] level [1 | 2]
```

Description

This command sets or clears the authentication password for one or all IS-IS VLANs.

Syntax Description

vlan all	Applies the password configuration to all IS-IS VLANs.
<i>vlan_name</i>	Specifies a single VLAN to which the password configuration is applied.
none	Clears the password configuration and disables hello PDU authentication.
encrypted	Specifies that the supplied password is encrypted and must be decrypted prior to using it in a TLV.
<i>password</i>	Specifies the password. Passwords may be up to 254 alphanumeric characters in length.
level [1 2]	Limits the password configuration to level 1 or level 2. If neither level 1 or level 2 is specified, the configuration applies to both levels.

Default

None.

Usage Guidelines

If configured, the specified password is included in Hello PDUs for the specified level. In addition, received Hello PDUs on the specified interface are authenticated with the same password. Hello PDUs that are not authenticated are discarded.



Only plain text passwords are supported. Note that if the password is changed on an interface with an existing adjacency, the neighboring router needs to be configured as well. Depending on how timers are configured, the adjacency may time out while transitioning between passwords. Although passwords appear in plain text during configuration, they are displayed and saved in encrypted form.

Example

The following command assigns password Extreme to all level 1 VLANs configured for IS-IS:

```
configure isis vlan all password simple Extreme level 1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis priority

```
configure isis [vlan all | {vlan} vlan_name] priority priority {level[1 | 2]}
```

Description

This command sets the priority used for DIS election on broadcast interfaces.

Syntax Description

vlan all	Applies the configuration to all IS-IS VLANs.
<i>vlan_name</i>	Specifies a single VLAN to which the configuration applies.
priority <i>priority</i>	Sets the priority value. The range is 0 to 127.
level [1 2]	Limits the configuration to either level 1 or level 2. If neither level 1 nor level 2 is specified, the configuration applies to both levels.

Default

64.

Usage Guidelines

A higher priority value is preferred over a lower priority value. The priority is encoded in level 1 or level 2 hello PDUs. This command is not valid for point-to-point interfaces. Note that a priority of 0 has no special meaning other than the fact that it is the lowest priority. A router with a priority of 0 can still become the DIS.



Example

The following command configures priority level 32 for SIPvlan:

```
configure isis SIPvlan priority 32
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis restart

```
configure isis restart [ none | planned | unplanned | both ]
```

Description

This command configures IS-IS graceful restart behavior.

Syntax Description

none	Disables IS-IS graceful restart. When graceful restart is disabled, this router still operates as a helper to other restarting routers.
planned	Initiates IS-IS graceful restart only in response to the restart process isis or run msm-failover commands.
unplanned	Initiates graceful restart only when the IS-IS process is restarted due to a process crash or an unplanned failover.
both	Initiates graceful restart for all events supported by the planned and unplanned options.

Default

None.

Usage Guidelines

The command options specify under which circumstances graceful restart is to be performed. This command has no effect during normal switch boot up. All IS-IS routing processes in the current virtual router are affected by this command.

All neighboring routers must support IS-IS restart in order for graceful restart to work. If graceful restart is not performed after a process restart or failover, the router's adjacencies are re-initialized



causing SPF recalculation throughout the network and, if the overload bit is not configured to be set during startup, churn as adjacencies change state and LSPs are learned.



Note

The planned and unplanned command options do not affect the actual restart protocol operation of IS-IS; they only determine when the restart process occurs.

Example

The following command configures the switch to initiate a graceful restart for all events supported by the planned and unplanned options:

```
configure isis restart both
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis restart grace-period

```
configure isis restart grace-period seconds
```

Description

This command configures the T3 global restart timer for all IS-IS router processes on the current virtual router.

Syntax Description

<i>seconds</i>	Specifies the restart grace period in seconds. The range is 1 to 65535 seconds.
----------------	---

Default

65535.

Usage Guidelines

If the grace period expires before LSP resynchronization is complete, the virtual router sets the overload bit in LSPs that it originates.



Example

The following command sets the restart grace period to 5000 seconds:

```
configure isis restart grace-period 5000
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis timer csnp-interval

```
configure isis [vlan all | {vlan} vlan_name] timer csnp-interval seconds {level [1 | 2]}
```

Description

This command sets the minimum time between consecutive CSNP transmissions on the specified interface.

Syntax Description

vlan all	Applies the timer configuration to all IS-IS VLANs.
<i>vlan_name</i>	Specifies a single VLAN to which the timer configuration applies.
<i>seconds</i>	Sets the timer interval. The range is 1 to 65535 seconds.
level [1 2]	Limits the configuration to either level 1 or level 2. If neither level 1 nor level 2 is specified, the configuration applies to both levels.

Default

10 seconds.

Usage Guidelines

Periodic CSNPs are only sent on broadcast interfaces and only by the DIS.



Example

The following command sets the CSNP interval time for all IS-IS VLANs to 15 seconds:

```
configure isis vlan all timer csnp-interval 15
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis timer hello-interval

```
configure isis [vlan all | {vlan} vlan_name] timer hello-interval [seconds | minimal] {level [1 | 2]}
```

Description

This command sets the interval between two consecutive hello transmissions.

Syntax Description

vlan all	Applies the timer configuration to all IS-IS VLANs.
<i>vlan_name</i>	Specifies a single VLAN to which the timer configuration applies.
<i>seconds</i>	Sets the timer interval. The range is 1 to 65535 seconds.
minimal	Specifies that the hello interval is calculated by dividing 1 second by the hello multiplier.
level [1 2]	Limits the configuration to either level 1 or level 2. If neither level 1 nor level 2 is specified, the configuration applies to both levels.

Default

10 seconds.

Usage Guidelines

If this router is the elected DIS, hellos are sent three times more frequently than the configured interval.

When the timer configuration is set to minimal, the holding time included in the PDU is set to 1 second. Otherwise, the holding time is computed by multiplying the hello interval by the hello multiplier. The holding time tells the neighboring router how long to wait before declaring the sending router dead.



Example

The following command sets the hello interval timer for all VLANs to 15 seconds:

```
configure isis vlan all timer hello-interval 15
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis timer lsp-interval

```
configure isis [vlan all | {vlan} vlan_name] timer lsp-interval milliseconds
```

Description

This command sets the minimum time between LSP transmissions.

Syntax Description

vlan all	Applies the timer configuration to all IS-IS VLANs.
<i>vlan_name</i>	Specifies a single VLAN to which the timer configuration applies.
<i>milliseconds</i>	Specifies the timer value. The range is 1 to 4294967295 milliseconds.

Default

33 milliseconds.

Usage Guidelines

This is used to throttle LSP flooding. Higher values reduce network traffic and can help keep underpowered routers from becoming overloaded during network events. Lower values speed up convergence.

Example

The following command sets the minimal LSP interval for IS-IS VLANs to 66 milliseconds:

```
configure isis vlan all timer lsp-interval 66
```



History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis timer restart-hello-interval

```
configure isis [vlan all | {vlan} vlan_name] timer restart-hello-interval seconds
{level [1 | 2]}
```

Description

This command configures the T1 restart retransmit timer for one or all VLANs.

Syntax Description

vlan all	Specifies that the T1 restart timer configuration applies to all VLANs.
<i>vlan_name</i>	Specifies a VLAN to which the T1 restart timer configuration applies.
<i>seconds</i>	Specifies the T1 restart timer value. The range is 1 to 65535 seconds.
level [1 2]	Limits the configuration change to level 1 or level 2. If neither level 1 nor level 2 is specified, the configuration applies to both levels.

Default

3 seconds.

Usage Guidelines

If, after sending a restart request, the router process associated with this interface does not receive a restart acknowledgement and a CSNP within the period specified by this command, another restart request is sent.

Example

The following command sets the T1 restart timer to 6 seconds on all level 1 VLANs:

```
configure isis vlan all timer restart-hello-interval 6 level 1
```

History

This command was first available in ExtremeXOS 12.1.



Platform Availability

This command is available on platforms with a Core license.

configure isis timer retransmit-interval

```
configure isis [vlan all | {vlan} vlan_name] timer retransmit-interval seconds
```

Description

This command sets the time to wait for an acknowledgement of a transmitted LSP on a point-to-point interface.

Syntax Description

vlan all	Applies the timer value to all IS-IS VLANs.
<i>vlan_name</i>	Specifies a single VLAN to which the timer configuration applies.
<i>seconds</i>	Defines the timer value. The range is 0 to 65535 seconds.

Default

5 seconds.

Usage Guidelines

If an acknowledgement is not received when the timer expires, the LSP is resent and the timer is reset.

Example

The following command sets the retransmit interval for the SJvlan to 10 seconds:

```
configure isis SJvlan timer retransmit-interval 10
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

configure isis wide-metric

```
configure isis [vlan all | {vlan} vlan_name] wide-metric metric {level[1 | 2]}
```



Description

This command sets the wide metric value for one or all IS-IS VLANs.

Syntax Description

vlan all	Applies the configuration to all IS-IS VLANs.
<i>vlan_name</i>	Specifies a single VLAN to which the configuration applies.
<i>metric</i>	Sets the metric. The range is 1 to 16777214.
level [1 2]	Limits the configuration change to either level 1 or level 2. If neither level 1 nor level 2 is specified, the configuration applies to both levels.

Default

10.

Usage Guidelines

If the wide metric style is enabled on the associated IS-IS router process, the wide metric value is used in Extended IP reachability TLVs, Extended IS Reachability TLVs, and IPv6 Reachability TLVs in LSPs.

Example

The following command sets the wide metric to 15 for all IS-IS VLANs:

```
configure isis vlan all wide-metric 15
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

create isis area

```
create isis area area_name
```

Description

This command creates an IS-IS router process in the current virtual router.



Syntax Description

<code>area_name</code>	Defines a name for the new IS-IS router process.
------------------------	--

Default

N/A.

Usage Guidelines

No PDUs are sent until after the following events:

- The router process has been enabled
- The router process has been assigned a system ID and area address
- The router process has at least one interface (VLAN) that has IPv4 or IPv6 forwarding enabled.

By default, newly created IS-IS router processes are Level 1/Level 2 routers if a level 2 router process does not already exist in the current virtual router. No more than one IS-IS router process may be configured as a level 2 router. IS-IS router processes on different virtual routers may have the same name, but this is not recommended as it may cause confusion when administering the switch. The router process name supplied with this command may be optionally used as the hostname for this router process when dynamic hostname exchange support is enabled.

The area name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see [Object Names](#) in the ExtremeXOS Concepts Guide.

A maximum of one area can be created per VR in this release.

Example

The following command creates a new IS-IS router process named areax:

```
create isis area areax
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

delete isis area

```
delete isis area [all | area_name]
```



Description

This command disables and deletes the specified IS-IS router process in the current virtual router.

Syntax Description

all	Deletes all IS-IS router processes.
<i>area_name</i>	Specifies the name of the IS-IS router process to be deleted.

Default

None.

Usage Guidelines

All configuration for the specified router is lost. All routes learned from this router process are purged from the routing tables.

Example

The following command deletes the IS-IS process named areax:

```
delete isis area areax
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

disable isis

```
disable isis {area area_name}
```

Description

This command disables the specified IS-IS router process on the current virtual router.

Syntax Description

<i>area_name</i>	Specifies the name of the IS-IS router process to be disabled.
------------------	--



Default

Disabled.

Usage Guidelines

IS-IS PDUs are no longer sent or processed on this IS-IS router process. The LSP and neighbor databases are purged. IS-IS routes are purged from the routing table. This command should only be used during planned network outages. This command has no effect on router processes that are already disabled.

Example

The following command disables the IS-IS process named areax:

```
disable isis area areax
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

disable isis area adjacency-check

```
disable isis area area_name adjacency-check {ipv4| ipv6}
```

Description

This command disables the checking of the following TLVs when forming adjacencies: Protocols Supported and IP Interface Address.

Syntax Description

<i>area_name</i>	Specifies the name of the IS-IS router process that should no longer perform the adjacency check.
ipv4	Specifies that the adjacency check should no longer be performed on IPv4 interfaces.
ipv6	Specifies that the adjacency check should no longer be performed on IPv6 interfaces.

Default

IPv4: Enabled.



IPv6: Enabled.

Usage Guidelines

When adjacency checking is disabled, adjacencies may be formed on interfaces that do not reside on the same subnet or do not support IPv4 (if disabled for IPv4) or IPv6 (if disabled for IPv6). If neither `ipv4` nor `ipv6` is specified, this command applies to IPv4.

Example

The following command directs the IS-IS process named `areax` to disable adjacency checks on IPv6 interfaces:

```
disable isis area areax adjacency-check ipv6
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

disable isis area dynamic-hostname

```
disable isis area area_name dynamic-hostname
```

Description

This command disables the dynamic hostname feature.

Syntax Description

<i>area_name</i>	Specifies the name of the IS-IS process for which the dynamic-hostname feature is to be disabled.
------------------	---

Default

Disabled.

Usage Guidelines

The specified router process no longer includes code 137 TLVs in its LSPs and names are no longer displayed in `show` commands.



Example

The following command disables the display of area names or SNMP names instead of system IDs:

```
disable isis area areax dynamic-hostname
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

disable isis area export

```
disable isis area area_name export {ipv4} route-type
```

Description

This command disables IPv4 route redistribution of the specified type into IS-IS.

Syntax Description

<i>area_name</i>	Specifies the IS-IS router process for which route redistribution is disabled.
ipv4	Specifies that the configuration change is for IPv4 IS-IS routing.
<i>route-type</i>	Selects the type of export route to disable. The valid route types are: bgp, direct, e-bgp, i-bgp, ospf, ospf-extern1, ospf-extern2, ospf-inter, ospf-intra, rip, and static.

Default

All types are disabled.

Usage Guidelines

None.

Example

The following command disables RIP route distribution into areax:

```
disable isis area areax export rip
```



History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

disable isis area export ipv6

```
disable isis area area_name export ipv6 route-type
```

Description

This command disables IPv6 route redistribution of the specified type into IS-IS.

Syntax Description

<i>area_name</i>	Specifies the IS-IS router process for which route redistribution is disabled.
<i>route-type</i>	Selects the type of export route to disable. The valid route types are: direct, ospfv3, ospfv3-extern1, ospfv3-extern2, ospfv3-inter, ospfv3-intra, ripng, bgp, and static.

Default

All types are disabled.

Usage Guidelines

None.

Example

The following command disables RIPng route distribution into areax:

```
disable isis area areax export ipv6 ripng
```

History

This command was first available in ExtremeXOS 12.1.

Support for BGP was added in ExtremeXOS 12.6.0-BGP.

Platform Availability

This command is available on platforms with a Core license.



disable isis area originate-default

```
disable isis area area_name originate-default {ipv4 | ipv6}
```

Description

This command disables the generation of one or all default routes in the LSPs for the specified router process.

Syntax Description

area_name	Specifies the name of the IS-IS router process that should no longer generate the default route.
ipv4	Specifies that the router process should no longer generate the default IPv4 route.
ipv6	Specifies that the router process should no longer generate the default IPv6 route.

Default

IPv4: Disabled.

IPv6: Disabled.

Usage Guidelines

This applies to level 2 routing only. By default this command disables IPv4 default route origination. The optional ipv6 keyword disables IPv6 default route origination. This command has no effect on router processes that are already disabled for default route origination on level 1-only router processes.

Example

The following command directs the IS-IS process named areax to stop generating the default IPv4 route in it's LSPs:

```
disable isis area areax originate-default
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.



disable isis area overload-bit

```
disable isis area area_name overload-bit
```

Description

This command disables the overload-bit feature.

Syntax Description

<i>area_name</i>	Specifies the area name of the IS-IS process for which this feature is to be disabled.
------------------	--

Default

Disabled.

Usage Guidelines

Disabling the overload bit feature causes an SPF recalculation throughout the network. In addition, external and interlevel router redistribution is no longer suppressed if those options were included when the overload bit was enabled. If the overload bit is currently set as a result of the overload-bit on-startup command, this command overrides the configuration and disables this feature.

Example

The following command disables the overload bit feature for areax:

```
disable isis area areax overload-bit
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

disable isis hello-padding

```
disable isis [vlan all | {vlan} vlan_name] hello-padding
```

Description

This command disables the padding of Hello PDUs for one or all IS-IS VLANs.



Syntax Description

vlan all	Disables hello padding on all IS-IS VLANs.
<i>vlan_name</i>	Specifies a single VLAN on which to disable hello padding.

Default

Enabled.

Usage Guidelines

Implicit adjacency MTU verification is not performed when hello padding is disabled.

Example

The following command disables hello padding on all IS-IS VLANs:

```
disable isis vlan all hello-padding
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

disable isis restart-helper

```
disable isis restart-helper
```

Description

This command disables the IS-IS restart helper.

Syntax Description

This command has no arguments or variables.

Default

Enabled.



Usage Guidelines

When this feature is disabled, the router does not act as a restart helper and may time out a restarting router's adjacency per normal operation.

Example

The following command disables the IS-IS restart helper:

```
disable isis restart-helper
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

enable isis

```
enable isis {area area_name}
```

Description

This command enables the specified IS-IS router process on the current virtual router.

Syntax Description

<i>area_name</i>	Specifies the name of the IS-IS router process to be enabled.
------------------	---

Default

Disabled.

Usage Guidelines

If no area name is specified, all IS-IS router processes on the current virtual router are enabled. Once a router process is enabled, IS-IS PDUs are sent and processed provided that the following conditions are met:

- The router process has a system ID and area address configured.
- At least one associated VLAN interface has IPv4 or IPv6 forwarding enabled.

This command has no effect on router processes that are already enabled.



Example

The following command enables the IS-IS process named areax:

```
enable isis area areax
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

enable isis area adjacency-check

```
enable isis area area_name adjacency-check {ipv4 | ipv6}
```

Description

This command enables the checking of the following TLVs when forming adjacencies: Protocols Supported and IP Interface Address.

Syntax Description

<i>area_name</i>	Specifies the name of the IS-IS router process that should perform the adjacency check.
ipv4	Specifies that the adjacency check is to be performed in IPv4 interfaces.
ipv6	Specifies that the adjacency check is to be performed in IPv6 interfaces.

Default

ipv4/ipv6: Enabled.

Usage Guidelines

When enabled for IPv4, IPv4 adjacencies may only be formed with neighbors whose connected interface supports IPv4 and is on the same subnet as the receiving interface. Similarly, when enabled for IPv6, IPv6 adjacencies may only be formed with neighbors whose connected interface supports IPv6 and is on the same link local subnet as the receiving interface. For each enabled protocol, if both criteria are not met, received Hello PDUs are discarded. By default, IPv4 routing is affected by this command. The optional `ipv6` keyword enables adjacency checking for IPv6 interfaces on the specified router process. It may be necessary to disable adjacency checking in multi-topology environments where a neighbor may only form an IPv4 or an IPv6 adjacency, but not both.



Example

The following command directs the IS-IS process named `areax` to perform adjacency checks on IPv6 interfaces:

```
enable isis area areax adjacency-check ipv6
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

enable isis area dynamic-hostname

```
enable isis area area_name dynamic-hostname [area-name | snmp-name]
```

Description

This command enables the dynamic hostname feature, which displays either the area name or the SNMP name instead of a IS-IS router system ID in select show commands.

Syntax Description

<i>area_name</i>	Specifies the name of the IS-IS process for which the dynamic-hostname feature is to be enabled.
area-name	Specifies that affected show commands display the area name instead of the IS-IS system ID.
snmp-name	Specifies that affected show commands display the SNMP name instead of the IS-IS system ID.

Default

Disabled.

Usage Guidelines

This command enables support for the dynamic hostname exchange feature defined by RFC 2763.



Example

The following command enables the display of IS-IS area names:

```
enable isis area areax dynamic-hostname area-name
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

enable isis area export

```
enable isis area area_name export {ipv4} route-type [policy | metric mvalue
{metric-type [internal | external]}] {level[1 | 2 | both-1-and-2]}
```

Description

This command enables IPv4 route redistribution into IS-IS for direct, static, BGP, RIP, or OSPF routes.

Syntax Description

<i>area_name</i>	Specifies the IS-IS router process that receives the exported routes.
ipv4	Specifies that the redistributed routes are for use in IPv4 IS-IS routing.
<i>route-type</i>	Selects the type of route for export. The valid route types are: bgp, direct, e-bgp, i-bgp, ospf, ospf-extern1, ospf-extern2, ospf-inter, ospf-intra, rip, and static.
<i>policy</i>	Specifies a policy that controls how routes are redistributed into IS-IS.
<i>mvalue</i>	Specifies a metric to assign to the routes exported to IS-IS. The range is 0 to 4261412864.
metric-type [internal external]	Specifies a metric type, which is internal or external, to assign to the routes exported to IS-IS.
level [1 2 both-1-and-2]	Limits the use of redistributed routes to level 1, level 2, or both.

Default

All types are disabled.



Usage Guidelines

If wide metrics are enabled, redistributed routes are included in the Extended IP Reachability TLV in LSPs. If wide metrics are not enabled, redistributed routes are added to IP External Reachability TLV in LSPs. For policies, the nlri match attribute is supported, and the cost, cost-type internal, permit, and deny set attributes are supported.

Example

The following command exports RIP routes to IS-IS and assigns the internal metric type and metric value 5 to the redistributed routes:

```
enable isis area areax export rip metric 5 metric-type internal
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

enable isis area export ipv6

```
enable isis area area_name export ipv6 route-type [policy | metric mvalue]
{level[1 | 2 | both-1-and-2]}
```

Description

This command enables IPv6 route redistribution into IS-IS for direct, static, RIPng, or OSPFv3 routes.

Syntax Description

<i>area_name</i>	Specifies the IS-IS router process that receives the exported routes.
<i>route-type</i>	Selects the type of route for export. The valid route types are: direct, ospfv3, ospfv3-extern1, ospfv3-extern2, ospfv3-inter, ospfv3-intra, ripng, bgp, and static.
<i>policy</i>	Specifies a policy that controls how routes are redistributed into IS-IS.
<i>mvalue</i>	Specifies a metric to assign to the routes exported to IS-IS. The range is 0 to 4261412864.
level [1 2 both-1-and-2]	Limits the use of redistributed routes to level 1, level 2, or both.

Default

All types are disabled.



Usage Guidelines

If a policy is specified, the policy is used to determine what specific routes are redistributed into IS-IS. Otherwise, the specified metric and type are assigned to the redistributed routes. Redistributed routes are added to the IPv6 External Reachability TLV in LSPs. For policies, the nlri match attribute is supported, and the cost, cost-type internal, permit, and deny set attributes are supported.

Example

The following command exports RIPng routes to IS-IS and assigns the internal metric type and metric value 5 to the redistributed routes:

```
enable isis area areax export ipv6 ripng metric 5 metric-type internal
```

History

This command was first available in ExtremeXOS 12.1.

Support for BGP was added in ExtremeXOS 12.6.0-BGP.

Platform Availability

This command is available on platforms with a Core license.

enable isis area originate-default

```
enable isis area area_name originate-default {ipv4 | ipv6}
```

Description

This command causes the specified IS-IS router process to generate the default route in its LSPs.

Syntax Description

area_name	Specifies the name of the IS-IS router process that should generate the default route.
ipv4	Specifies that the router process should generate the default IPv4 route.
ipv6	Specifies that the router process should generate the default IPv6 route.

Default

IPv4: Disabled

IPv6: Disabled



Usage Guidelines

This applies to level 2 routing only. In contrast, level 1 routers compute the default route as the nearest attached L1/L2 router. When enabled, the router process generates an IPv4 default route unless the `ipv6` option is specified. Only one level 2 router in the IS-IS domain should be configured to originate a default route. This command has no effect on router processes that are already enabled for default route origination or on router processes that are level 1-only.

Example

The following command directs the IS-IS process named `areax` to generate the default IPv4 route in it's LSPs:

```
enable isis area areax originate-default
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

enable isis area overload-bit

```
enable isis area area_name overload-bit {suppress [external | interlevel | all]}
```

Description

This command enables the overload-bit feature, which signals other routers that they are no longer permitted to use this router as a transit or forwarding node.

Syntax Description

<i>area_name</i>	Specifies the area name of the IS-IS process for which this feature is to be enabled.
suppress	Specifies that one or all types of reachability information is to be suppressed or excluded from LSPs.
external	When included with the <code>suppress</code> option, this specifies that external reachability information is to be excluded from LSPs.
interlevel	When included with the <code>suppress</code> option, this specifies that interlevel reachability information is to be excluded from LSPs.
all	When included with the <code>suppress</code> option, this specifies that external and interlevel reachability information is to be excluded from LSPs.



Default

Disabled.

Usage Guidelines

When the overload bit feature is enabled, the router process still receives and processes LSPs.

Example

The following command enables the overload bit feature for area:

```
enable isis area areax overload-bit
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

enable isis hello-padding

```
enable isis [vlan all | {vlan} vlan_name] hello-padding
```

Description

This command enables the padding of hello PDUs on one or all VLANs.

Syntax Description

 vlan all 	Enables hello padding on all IS-IS VLANs.
<i> vlan_name </i>	Specifies a single VLAN on which hello padding is enabled.

Default

Enabled.

Usage Guidelines

When hello padding is enabled, IS-IS pads hello packets to the interface MTU. This is used among neighbors to verify that adjacencies have the same MTU configured on either end. The disadvantage of hello padding is the price of bandwidth consumed by larger packets.



Example

The following command enables hello padding on the SJVlan VLAN:

```
enable isis SJvlan hello-padding
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

enable isis restart-helper

enable isis restart-helper

Description

This command enables the IS-IS router to act as a restart helper according to draft-ietf-isis-restart-02—Restart signaling for IS-IS.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

None.

Example

The following command enables the IS-IS restart helper:

```
enable isis restart-helper
```

History

This command was first available in ExtremeXOS 12.1.



Platform Availability

This command is available on platforms with a Core license.

show isis

show isis

Description

This command displays the global IS-IS configuration information as well as a summarized router process listing.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The displayed global configuration information includes the restart enablement, restart grace period, and import-policy setting. The router process listing includes the area name, system ID, whether it's enabled, the IS type, and a count of associated interfaces and area addresses. This command applies only to the IS-IS router processes running in the current virtual router.

Example

The following command displays IS-IS information:

```
show isis
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

show isis area

```
show isis area [area_name | all]
```



Description

This command displays configuration information for a specific router process or for all IS-IS router processes.

Syntax Description

<i>area_name</i>	Specifies the router process for which to display IS-IS area information.
all	Displays information for all IS-IS router processes.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays IS-IS configuration for areax:

```
show isis area areax
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

show isis area summary-addresses

```
show isis area area_name summary-addresses
```

Description

This command displays the configured IPv4 and IPv6 summary addresses for the specified area.

Syntax Description

<i>area_name</i>	Specifies the router process for which you want to display summary addresses.
------------------	---



Default

N/A.

Usage Guidelines

None.

Example

The following command displays the summary addresses for areax:

```
show isis area areax summary-addresses
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

show isis counters

```
show isis counters {area [area_name | all] | vlan [vlan_name | all]}
```

Description

This command displays counters for an area or a VLAN.

Syntax Description

<i>area_name</i>	Specifies a router process for which to display the IS-IS counters. If you do not specify an IS-IS area, the software displays counters for all areas.
<i>vlan_name</i>	Specifies a VLAN for which to display IS-IS counters.
all	Displays IS-IS counters for all areas or VLANs.

Default

None.

Usage Guidelines

If you enter the show isis counters command without any additional keywords or parameters, the software displays the counters for all areas.



Example

The following command displays the IS-IS counters for the configured area:

```
show isis counters
```

The following command displays the IS-IS counters for the SJVlan VLAN:

```
show isis counters vlan SJvlan
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

show isis lsdb

```
show isis lsdb {area area_name {lsp-id lsp_id}} {level [1|2]} {detail | stats}
```

Description

Displays a summary of the IS-IS link state database for one or all IS-IS router processes running in the current virtual router.

Syntax Description

area_name	Specifies the name of a router process for which to display the IS-IS link state database. If the area name is omitted, this command displays information for all areas.
lsp_id	Limits the display to the specified LSP ID, which is specified in the form <i>system id.pseudonode ID-LSP number</i> . For example: 0102.0ff2.0023.00-01. The pseudonode ID and LSP numbers are optional. If they are not included, multiple listings might appear.
level [1 2]	Limits the display to LSPs for either level 1 or level 2.
detail	Expands the display to include the LSP TLVs and IPv4 and IPv6 reachability information. The displayed information varies depending on what is included in the LSPs.
stats	Displays counts of the number of LSPs and prefixes in the LSP database.

Default

N/A.



Usage Guidelines

None.

Example

The following command displays information for a specific LSP:

```
show isis lsdb area areax lsp-id 0102.0ff2.0023.00-01
```

The following example shows the display for the stats option:

```
(debug) Switch.6 # show isis lsdb stats
Area "a1" :
IS-IS Level-1 Link State Database:
LSPs (including fragments) : 4
Internal Prefixes (Type 128) : 7
External Prefixes (Type 130) : 0
IPv4 Prefixes (Type 135) : 0
IPv6 Prefixes (Type 236) : 0
MT IPv4 Prefixes (Type 235) : 0
MT IPv6 Prefixes (Type 237) : 0
IS-IS Level-2 Link State Database:
LSPs (including fragments) : 1
Internal Prefixes (Type 128) : 5
External Prefixes (Type 130) : 0
IPv4 Prefixes (Type 135) : 0
IPv6 Prefixes (Type 236) : 0
MT IPv4 Prefixes (Type 235) : 0
MT IPv6 Prefixes (Type 237) : 0
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

show isis neighbors

```
show isis neighbors {area area_name} {vlan vlan_name} {ipv4 | ipv6} {detail}
```

Description

This command displays information about neighbors and their adjacencies.



Syntax Description

<i>area_name</i>	Specifies a router process for which to specify neighbor information. If you do not specify an IS-IS area, the software displays the neighbors in all areas.
<i>vlan_name</i>	Specifies a VLAN for which to specify neighbor information.
ipv4	Displays only the neighbors that advertise the IPv4 protocol as supported.
ipv6	Displays only the neighbors that advertise the IPv6 protocol as supported.
detail	Displays detailed information for IS-IS neighbors.

Default

N/A.

Usage Guidelines

If you do not specify either the `ipv4` or the `ipv6` keyword, this command displays all neighbors regardless of the supported protocol.

Example

The following command displays IS-IS neighbor information for area:

```
show isis neighbors area areax
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

show isis topology

```
show isis topology {area area_name {level [1 | 2]}} {ipv4 | ipv6}
```

Description

This command displays the topology for IPv4, IPv6, or both IPv4 and IPv6 for the specified area and level.



Syntax Description

area_name	Specifies the router process for which the topology applies.
level [1 2]	Specifies the IS-IS level of the topology you want to view.
ipv4	Selects the IPv4 topology for display. If you omit the ipv4 and ipv6 options, the IPv4 topology appears.
ipv6	Selects the IPv6 topology for display. If you omit this option, the IPv4 topology appears.

Default

None.

Usage Guidelines

Each known IS in the area or domain is displayed along with the next-hop and metric information.

Example

The following command display IPv4 topology information for areax:

```
show isis topology area areax ipv4
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

show isis vlan

```
show isis vlan {enabled | { vlan_name | all} }
```

Description

This command displays configuration and status information about the specified IS-IS interface.

Syntax Description

enabled	Displays information only for IS-IS-enabled VLANs.
vlan_name	Specifies a VLAN for which to display IS-IS interface information.
all	Displays information on all IS-IS interfaces.



Default

N/A.

Usage Guidelines

None.

Example

The following command displays IS-IS interface information for the SJvlan VLAN:

```

show isis vlan SJvlan
ISIS Interfaces Summary :
-----
VLAN      Area      State   Cfg      Address
-----
bd10k     a1        u46n-   b1246    11.1.1.1/24
v2        a1        u46--   p1246    2.1.1.2/24
2001:db8:2010::1/64
v3        a1        u46-g   p1246    2001:db8:2011::2/64
State Flags :
u - Links up, d - Links down,
4 - IPv4 forwarding enabled, 6 - IPv6 forwarding enabled,
n - Multinetted (v4), g - Multiple global addresses (v6)
Cfg Flags   :
b - Broadcast interface, p - Point-To-Point interface,
1 - L1 circuit type, 2 - L2 circuit type,
4 - ISIS-enabled for IPv4, 6 - ISIS-enabled for IPv6

```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

unconfigure isis area

```

unconfigure isis area area_name {level [1|2]}

```

Description

This command resets most configurable parameters for the specified router process to the default values.



Syntax Description

<i>area_name</i>	Specifies the router process to be unconfigured.
level [1 2]	Specifies either level 1 or level 2.

Default

N/A.

Usage Guidelines

This command does not delete interfaces from the router process, but it does disable them. The system ID and IS type are not changed. Where appropriate, the default values apply to level 1, level 2, and both IPv4 and IPv6.

Example

The following command resets the area configuration parameters for areax:

```
unconfigure isis area areax
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.

unconfigure isis vlan

```
unconfigure isis [vlan all | {vlan} vlan_name] {level [1|2]}
```

Description

This command resets all configurable interface parameters to the defaults on one or all VLANs.

Syntax Description

vlan all	Unconfigures IS-IS for all VLANs.
<i>vlan_name</i>	Specifies a single VLAN for which IS-IS is unconfigured.
level [1 2]	Specifies either level 1 or level 2.



Default

N/A.

Usage Guidelines

None.

Example

The following command resets IS-IS configuration parameters for Srvlan:

```
unconfigure isis Srvlan
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Core license.



41 BGP Commands

```
clear bgp flap-statistics
clear bgp neighbor counters
configure bgp add aggregate-address
configure bgp add confederation-peer sub-AS-number
configure bgp add network
configure bgp as-display-format
configure bgp as-number
configure bgp cluster-id
configure bgp confederation-id
configure bgp delete aggregate-address
configure bgp delete confederation-peer sub-AS-number
configure bgp delete network
configure bgp export shutdown-priority
configure bgp import-policy
configure bgp local-preference
configure bgp maximum-paths
configure bgp med
configure bgp neighbor allowas-in
configure bgp neighbor dampening
configure bgp neighbor description
configure bgp neighbor dont-allowas-in
configure bgp neighbor maximum-prefix
configure bgp neighbor next-hop-self
configure bgp neighbor no-dampening
configure bgp neighbor password
configure bgp neighbor peer-group
configure bgp neighbor route-policy
configure bgp neighbor route-reflector-client
configure bgp neighbor send-community
configure bgp neighbor shutdown-priority
configure bgp neighbor soft-reset
configure bgp neighbor source-interface
configure bgp neighbor timer
configure bgp neighbor weight
configure bgp peer-group allowas-in
configure bgp peer-group dampening
configure bgp peer-group dont-allowas-in
```

```
configure bgp peer-group maximum-prefix
configure bgp peer-group next-hop-self
configure bgp peer-group no-dampening
configure bgp peer-group password
configure bgp peer-group remote-AS-number
configure bgp peer-group route-policy
configure bgp peer-group route-reflector-client
configure bgp peer-group send-community
configure bgp peer-group soft-reset
configure bgp peer-group source-interface
configure bgp peer-group timer
configure bgp peer-group weight
configure bgp restart
configure bgp restart address-family
configure bgp restart restart-time
configure bgp restart stale-route-time
configure bgp restart update-delay
configure bgp routerid
configure bgp soft-reconfiguration
create bgp neighbor peer-group
create bgp neighbor remote-AS-number
create bgp peer-group
delete bgp neighbor
delete bgp peer-group
disable bgp
disable bgp adj-rib-out
disable bgp advertise-inactive-route
disable bgp aggregation
disable bgp always-compare-med
disable bgp community format
disable bgp export
disable bgp export vr
disable bgp fast-external-fallover
disable bgp neighbor
disable bgp neighbor capability
disable bgp neighbor originate-default
disable bgp neighbor remove-private-AS-numbers
disable bgp neighbor soft-in-reset
disable bgp peer-group
disable bgp peer-group capability
disable bgp peer-group originate-default
disable bgp peer-group remove-private-AS-numbers
```



```

disable bgp peer-group soft-in-reset
enable bgp
enable bgp adj-rib-out
enable bgp advertise-inactive-route
enable bgp aggregation
enable bgp always-compare-med
enable bgp community format
enable bgp export
enable bgp export vr
enable bgp fast-external-fallover
enable bgp neighbor
enable bgp neighbor capability
enable bgp neighbor originate-default
enable bgp neighbor remove-private-AS-numbers
enable bgp neighbor soft-in-reset
enable bgp peer-group
enable bgp peer-group capability
enable bgp peer-group originate-default
enable bgp peer-group remove-private-AS-numbers
enable bgp peer-group soft-in-reset
show bgp
show bgp memory
show bgp neighbor
show bgp neighbor [flap-statistics | suppressed-routes]
show bgp peer-group
show bgp routes
show bgp routes summary

```

This chapter describes commands for doing the following:

- Configuring BGP
- Displaying BGP information
- Managing BGP



Note

Many of the BGP commands in this chapter are enhanced to support BGP configurations in a VRF (PE – CE). All of these commands can be executed in a context of a VRF, and the configuration is applied to the BGP instance running inside a VRF.

For an introduction to the BGP feature, see the *ExtremeXOS Concepts Guide*.

clear bgp flap-statistics



```
clear bgp {neighbor} remoteaddr {address-family [ipv4-unicast | ipv4-multicast |
ipv6-unicast | ipv6-multicast |vpn4]} flap-statistics [all | rd rd_value |as-
path path expression | community [no-advertise | no-export | no-export-subconfed
| number community_num | AS_Num:Num] | network [any / netMaskLen |
networkPrefixFilter] {exact}]
```

Description

Clears flap statistics for routes to specified neighbors.

Syntax Description

all	Specifies flap statistics for all routes.
<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support.
<i>rd_value</i>	Specifies the Route Distinguisher (RD) value for the Layer 3 VPN routes for which you want to clear flap statistics.
no-advertise	Specifies the no-advertise community attribute.
no-export	Specifies the no-export community attribute.
no-export-subconfed	Specifies the no-export-subconfed community attribute.
<i>community_num</i>	Specifies a community number.
<i>AS_Num</i>	Specifies an autonomous system ID (0-65535).
<i>Num</i>	Specifies a community number.
any	Specifies all routes with a given or larger mask length.
<i>netMaskLen</i>	Specifies a subnet mask length (number of bits).
<i>networkPrefixFilter</i>	Specifies an IP address and netmask.
exact	Specifies an exact match with the IP address and subnet mask.

Default

If no address family is specified, IPv4 unicast is the default.



Note

You must specify an IPv6 address family for an IPv6 peer, because an IPv6 peer does not support the default IPv4 unicast address family. Similarly, if you specify an IPv4 peer and an address family in the command, an IPv4 address family must be specified.



Usage Guidelines

Use this command to clear flap statistics for a specified BGP neighbor.

The option `network any / netMaskLen` clears the statistics for all BGP routes whose mask length is equal to or greater than `maskLength`, irrespective of their network address.

The option `network any / netMaskLen exact` clears the statistics for all BGP routes whose mask length is exactly equal to `maskLength`, irrespective of their network address.

To clear flap statistics on Layer 3 VPNs, you must configure this feature in the context of the MPLS-enabled VR; this feature is not supported for BGP routes on the CE (VRF) side of the PE router.

This command applies to the current VR or VRF context.

Example

The following command clears the flap statistics for a specified neighbor:

```
clear bgp neighbor 10.10.10.10 flap-statistics all
```

History

This command was first available in ExtremeXOS 10.1.

The `any / <netMaskLen>` options were added in ExtremeXOS 11.0.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 in BGP was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

clear bgp neighbor counters

```
clear bgp neighbor [remoteaddr | all] counters
```

Description

Resets the BGP counters for one or all BGP neighbor sessions to zero.



Syntax Description

<code>remoteaddr</code>	Specifies the IPv4 or IPv6 address of a specific BGP neighbor.
<code>all</code>	Specifies that counters for all BGP neighbors should be reset.

Default

N/A.

Usage Guidelines

This command resets the following counters:

- In-total-msgs
- Out-total-msgs
- In-updates
- Out-updates
- FsmTransitions

The command `clear counters` also resets all counter for all BGP neighbors. For BGP, the `clearcounters` command is equivalent to the following BGP command:

```
clear bgp neighbor all counters
```

This command applies to the current VR or VRF context.

Example

The following command resets the counters for the BGP neighbor at 10.20.30.55:

```
clear bgp neighbor 10.20.30.55 counters
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure bgp add aggregate-address

```
configure bgp add aggregate-address {address-family [ipv4-unicast | ipv4-
multicast | ipv6-unicast | ipv6-multicast]} ipaddress/masklength {as-match | as-
set} {summary-only} {advertise-policy policy} {attribute-policy policy}
```

Description

Configures a BGP aggregate route.

Syntax Description

address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.
<i>ipaddress/masklength</i>	Specifies an IP network address and mask length.
as-match	Generates autonomous system sequence path information (order of AS numbers in AS_PATH is preserved).
as-set	Generates autonomous system set path information (order of AS numbers in AS_PATH is not preserved).
summary-only	Specifies to send only aggregated routes to the neighbors.
advertise-policy	Specifies the policy used to select routes for this aggregated route.
attribute-policy	Specifies the policy used to set the attributes of the aggregated route.

Default

If no address family is specified, IPv4 unicast is the default.

N/A.

Usage Guidelines

Route aggregation is the process of combining the characteristics of several routes so that they are advertised as a single route. Aggregation reduces the amount of information that a BGP speaker must store and exchange with other BGP speakers. Reducing the information that is stored and exchanged also reduces the size of the routing table.

Before you can create an aggregate route, you must enable BGP aggregation using the following command:

```
enable bgp aggregation
```



If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

If the specified address is an IPv4 address, an IPv4 address family must be specified with the command. If the specified address is an IPv6 address, an IPv6 address family must be specified with the command.

BGP supports overlapping routes. For example, you can configure both of the following aggregate addresses:

- 192.0.0.0/8
- 192.168.0.0/16

After you create an aggregate route, the aggregate route remains inactive until BGP receives a route with an IP address and mask that conforms to an aggregate route. When a conforming route is received, the aggregate route becomes active and is advertised to BGP neighbors. If the summary-only option is specified, only the aggregate route becomes active and is advertised. If the summary-only option is omitted, any conforming aggregate routes and the received route are advertised to BGP neighbors.

Example

The following command configures a BGP aggregate route:

```
configure bgp add aggregate-address 192.1.1.4/30
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for overlapping aggregate addresses was added in ExtremeXOS 12.1.3.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp add confederation-peer sub-AS-number

```
configure bgp add confederation-peer sub-AS-number number
```



Description

Adds a sub-AS to a confederation.

Syntax Description

<i>number</i>	Specifies a sub-AS number of the confederation. The range is 1 to 4294967295.
---------------	---

Default

N/A.

Usage Guidelines

Before you can add a sub-AS to a confederation on the switch, you must disable any BGP neighbor sessions that are configured with the same AS number as a remote AS number. To disable BGP neighbor sessions, use the following command:

```
disable bgp neighbor [remoteaddr | all]
```

Invoke the `configure bgp add confederation-peer sub-AS-number` command multiple times to add multiple sub-ASs.

IBGP requires networks to use a fully-meshed router configuration. This requirement does not scale well, especially when BGP is used as an interior gateway protocol. One way to reduce the size of a fully-meshed AS is to divide the AS into multiple sub-autonomous systems and group them into a routing confederation. Within the confederation, all BGP speakers in each sub-AS must be fully-meshed. The confederation is advertised to other networks as a single AS.

The AS number is a 4-byte AS number in either the ASPLAIN or the ASDOT format as described in RFC 5396, Textual Representation of Autonomous System (AS) Numbers.

Example

The following command adds one sub-AS to a confederation using the ASPLAIN 4-byte AS number format:

```
configure bgp add confederation-peer sub-AS-number 65536
```

The following command adds one sub-AS to a confederation using the ASDOT 4-byte AS number format:

```
configure bgp add confederation-peer sub-AS-number 1.15
```

History

This command was first available in ExtremeXOS 10.1.



This command required a specific license in ExtremeXOS 11.1.

Support for 4-byte AS numbers was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp add network

```
configure bgp add network {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast]} ipaddress/masklength {network-policy policy}
```

Description

Adds a network to be originated from this router.

Syntax Description

address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.
<i>ipaddress/masklength</i>	Specifies an IP network address and mask length.
<i>policy</i>	Name of policy to be associated with network export. Policy can filter and/or change the route parameters.

Default

If no address family is specified, IPv4 unicast is the default.

N/A.

Usage Guidelines

The network must be present in the routing table.

Using the export command to redistribute routes complements the redistribution of routes using the `configure bgp add network` command. The `configure bgp add network` command adds the route to BGP only if the route is present in the routing table. The `enable bgp export` command redistributes an individual route from the routing table to BGP. If you use both commands to redistribute routes, the routes redistributed using the network command take precedence over routes redistributed using the export command.



If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

If the specified address is an IPv4 address, an IPv4 address family must be specified with the command. If the specified address is an IPv6 address, an IPv6 address family must be specified with the command.

Example

The following command adds a network to be originated from this router:

```
configure bgp add network 192.1.1.16/32
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp as-display-format

```
configure bgp as-display-format [asdot | asplain]
```

Description

Configures the AS number format displayed in show commands.

Syntax Description

asdot	Specifies the ASDOT format.
asplain	Specifies the ASPLAIN format.

Default

N/A.



Usage Guidelines

The ASPLAIN and ASDOT formats are described in RFC 5396, Textual Representation of Autonomous System (AS) Numbers.

Example

The following command selects the ASDOT 4-byte AS number format:

```
configure bgp as-display-format asdot
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp as-number

```
configure bgp AS-number number
```

Description

Changes the local AS number used by BGP.

Syntax Description

<i>number</i>	Specifies a local AS number. The range is 1 to 4294967295.
---------------	--

Default

N/A.

Usage Guidelines

BGP must be disabled before the AS number can be changed.

This command applies to the current VR or VRF context.

The AS number is a 4-byte AS number in either the ASPLAIN or the ASDOT format as described in RFC 5396, Textual Representation of Autonomous System (AS) Numbers.



Example

The following command specifies a local AS number using the ASPLAIN 4-byte AS number format:

```
configure bgp AS-number 65551
```

The following command specifies a local AS number using the ASDOT 4-byte AS number format:

```
configure bgp AS-number 1.15
```



Note

To remove the configured bgp as-number, assign as-number value as 0, i.e. configure bgp AS-number 0.

The following command configures the BGP router ID:

```
configure bgp routerid
```



Note

To remove the configured bgp routerid, give routerid value as 0.0.0.0 i.e. configure bgp routerid 0.0.0.0.

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for 4-byte AS numbers was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp cluster-id

```
configure bgp cluster-id cluster-id
```

Description

Configures the local cluster ID.



Syntax Description

<code>cluster-id</code>	Specifies a 4 byte field used by a route reflector to recognize updates from other route reflectors in the same cluster. The range is 0 - 4294967295.
-------------------------	---

Default

N/A.

Usage Guidelines

BGP must be disabled before the cluster ID can be changed.

Used when multiple route reflectors are used within the same cluster of clients.

Example

The following command appends a BGP route reflector cluster ID to the cluster list of a route:

```
configure bgp cluster-id 40000
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp confederation-id

```
configure bgp confederation-id number
```

Description

Specifies a BGP routing confederation ID.

Syntax Description

<code>confederation-id</code>	Specifies a routing confederation identifier, which is a 4-byte AS number in the range of 1 to 4294967295.
-------------------------------	--



Default

N/A.

Usage Guidelines

IBGP requires that networks use a fully-meshed router configuration. This requirement does not scale well, especially when BGP is used as an interior gateway protocol. One way to reduce the size of a fully-meshed AS is to divide the AS into multiple sub-autonomous systems and group them into a routing confederation. Within the confederation, each sub-AS must be fully-meshed. The confederation is advertised to other networks as a single AS.

The confederation ID is a 4-byte AS number in either the ASPLAIN or the ASDOT format as described in RFC 5396, Textual Representation of Autonomous System (AS) Numbers.

BGP must be disabled before the confederation ID can be changed.

Use a confederation ID of 0 to indicate no confederation. You cannot unconfigure the confederation ID while confederation peers are configured. You must delete the confederation peers before you unconfigure the confederation ID.

Example

The following command specifies a BGP routing confederation ID using the ASPLAIN 4-byte AS number format:

```
configure bgp confederation-id 65551
```

The following command specifies a BGP routing confederation ID using the ASDOT 4-byte AS number format:

```
configure bgp confederation-id 1.15
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for 4-byte AS numbers was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure bgp delete aggregate-address

```
configure bgp delete aggregate-address {address-family [ipv4-unicast | ipv4-
multicast | ipv6-unicast | ipv6-multicast]} [ ipaddress/masklength | all]
```

Description

Deletes one or all BGP aggregated routes.

Syntax Description

address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.
<i>ipaddress/masklength</i>	Specifies an IP network address and netmask length.
all	Specifies all aggregated routes in the specified address family. If you do not specify an address family, all aggregated routes in all address families are deleted.

Default

N/A.

Usage Guidelines

Route aggregation is the process of combining the characteristics of several routes so that they are advertised as a single route. Aggregation reduces the amount of information that a BGP speaker must store and exchange with other BGP speakers. Reducing the information that is stored and exchanged also reduces the size of the routing table.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Note



If the specified address is an IPv4 address, an IPv4 address family must be specified with the command. If the specified address is an IPv6 address, an IPv6 address family must be specified with the command.

Example

The following command deletes a BGP aggregate route:

```
configure bgp delete aggregate-address 192.1.1.4/30
```

History

This command was first available in ExtremeXOS 10.1.



This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp delete confederation-peer sub-AS-number

```
configure bgp delete confederation-peer sub-AS-number number
```

Description

Specifies a sub-AS that should be deleted from a confederation.

Syntax Description

sub-AS-number	Specifies a sub-AS.
----------------------	---------------------

Default

N/A.

Usage Guidelines

Before you can change the configuration with this command, you must disable the BGP neighbors in the confederation using the following command:

```
disable bgp neighbor [remoteaddr | all]
```

Example

The following command deletes a sub-AS from a confederation using the ASPLAIN 4-byte AS number format:

```
configure bgp delete confederation-peer sub-AS-number 65551
```

The following command deletes a sub-AS from a confederation using the ASDOT 4-byte AS number format:

```
configure bgp delete confederation-peer sub-AS-number 1.15
```



History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for 4-byte AS numbers was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp delete network

```
configure bgp delete network {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast]} [all | ipaddress/masklength]
```

Description

Deletes a network to be originated from this router.

Syntax Description

address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.
all	Specifies all networks for the specified address family. If no address family is specified, all networks for all address families are deleted.
<i>ipaddress/masklength</i>	Specifies an IP network address and netmask length.

Default

N/A.

Usage Guidelines

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Note



If the specified address is an IPv4 address, an IPv4 address family must be specified with the command. If the specified address is an IPv6 address, an IPv6 address family must be specified with the command.



Example

The following command deletes a network to be originated from this router:

```
configure bgp delete network 192.1.1.12/30
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp export shutdown-priority

```
configure bgp export route_type {{address-family} address_family} shutdown-  
priority number
```

Description

Configures the shutdown priority for IGP export.

Syntax Description

<i>route_type</i>	Specifies the BGP export route type.
address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.
<i>number</i>	Specifies the shutdown priority. The range is 0 - 65,535.

Default

The default value is 2048.

If no address family is specified, IPv4 unicast is the default.

Note



You must specify an IPv6 address family for an IPv6 peer, because an IPv6 peer does not support the default IPv4 unicast address family. Similarly, if you specify an IPv4 peer and an address family in the command, an IPv4 address family must be specified.



Usage Guidelines

To export IPv6 protocols to BGP, you must specify an IPv6 address family.



Note

This command is not currently supported, and is not recommended for use.

Higher priority values lower the chance of an IGP export to be automatically disabled in case BGP or the system goes to a low memory condition.



Note

For this command to execute, the specified protocol must support the specified address family. For example, the command fails if you specify OSPF and the IPv6 unicast address family. You can specify blackhole, direct, static, and IS-IS routes with IPv4 or IPv6 address families.

Example

The following command configures the shutdown priority of BGP exported OSPF routes to 1000:

```
configure bgp export ospf shutdown-priority 1000
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp import-policy

```
configure bgp import-policy [policy-name | none]
```

Description

Configures the import policy for BGP.



Syntax Description

<i>policy-name</i>	Specifies the policy.
none	Specifies no policy.

Default

N/A.

Usage Guidelines

Use the none keyword to remove a BGP import policy.

An import policy is used to modify route attributes while adding BGP routes to the IP route table.

Example

The following command configures a policy `imprt_plcy` for BGP:

```
configure bgp import-policy imprt_plcy
```

The following command unconfigures the import policy for BGP:

```
configure bgp import-policy none
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp local-preference

```
configure bgp local-preference number
```

Description

Changes the default local preference attribute.



Syntax Description

<i>number</i>	Specifies a value used to advertise this router's degree of preference to other routers within the AS. Range is 0 to 2147483647.
---------------	--

Default

100.

Usage Guidelines

BGP must be disabled before the local preference attribute can be changed.

BGP selects routes based on the following precedence (from highest to lowest):

- higher weight
- higher local preference
- shortest length (shortest AS path)
- lowest origin code
- lowest MED
- route from external peer
- lowest cost to Next Hop
- lowest routerID

Local preference is used to determine a preferred exit point from an AS. Local preferences are exchanged throughout the AS. A change in the local-preference can result in a change in routing and forwarding of traffic leaving the AS.

Example

The following command changes the default local preference attribute to 500:

```
configure bgp local-preference 500
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure bgp maximum-paths

configure bgp maximum-paths *max-paths*

Description

Enables or disables the BGP ECMP feature and specifies the maximum number of paths supported on the current VR.

Syntax Description

<i>max-paths</i>	Specifies the maximum number of paths. The range is 1 to 8. The value 1 disables BGP ECMP. A value greater than 1 enables BGP ECMP and specifies the maximum number of paths.
------------------	---

Default

One. BGP ECMP is disabled.

Usage Guidelines

This command triggers the BGP decision process, causing BGP to re-install the entire BGP routing table into the IP forwarding table. This activity requires a significant amount of switch processor resources, so we recommend that you enable or disable the BGP ECMP feature before enabling the BGP protocol globally on a VR. To ensure that BGP ECMP routes are programmed in the hardware, enter the `enable iproute sharing` command.



Note

BGP must be disabled before you can change the configuration with this command.

Example

The following command enables BGP ECMP and sets the maximum number of paths to 4:

```
configure bgp maximum-paths 4
```

History

This command was first available in ExtremeXOS 12.1.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure bgp med

```
configure bgp med [none | bgp_med]
```

Description

Configures the metric to be included in the Multi-Exit-Discriminator (MED) path attribute. The MED path attribute is included in route updates sent to external peers if a value is configured.

Syntax Description

none	Specifies not to use a multi-exist-discriminator number.
<i>bgp_med</i>	Specifies a multi-exit-discriminator number. The range is 0-2147483647.

Default

N/A.

Usage Guidelines

BGP selects routes based on the following precedence (from highest to lowest):

- higher weight
- higher local preference
- shortest length (shortest AS path)
- lowest origin code
- lowest MED
- route from external peer
- lowest cost to Next Hop
- lowest routerID



Note

BGP must be disabled before you can change the configuration with this command.

Example

The following command configures the metric to be included in the MED path attribute:

```
configure bgp med 3
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp neighbor allowas-in

```
configure bgp neighbor [all | remoteaddr] {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4]} allowas-in {max-as-occurrence as-count}
```

Description

Configures EBGP to receive and accept a looped EBGP route from the specified neighbor, provided the number of occurrences of local AS number in AS-Path is less than or equal to the value of *as-count*.

Syntax Description

all	Specifies that the configuration applies to all neighbors in the specified address family. If no address family is specified or if an IPv4 address is specified, the configuration applies to all IPv4 neighbors. If an IPv6 address family is specified, the configuration applies to all IPv6 neighbors.
<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor. The switch uses the IP address format to determine if the address is an IPv4 or IPv6 address.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support. This address family is applicable for PE to PE BGP neighbor sessions only. This keyword may prompt warning or error messages if executed for a regular BGP neighbor session or for a PE to CE neighbor session.
<i>as-count</i>	The maximum number of occurrences of local AS number in the received route AS-Path. If the number of occurrences of local AS number in AS-Path is more than <i>as-count</i> , the route is not accepted. The valid range is from 1-16.

Default

This feature is disabled by default.

If no *as-count* is specified, the *as-count* defaults to 3.

If no address family is specified and an IPv4 address is detected, IPv4 unicast is the default address family.



Usage Guidelines

In a hub and spoke configuration, it becomes necessary to accept an inbound EBGP route even though the route's AS-Path contains the receiver's own AS-number. In such network topologies, this feature can be enabled.



Note

A looped AS path is always allowed for IBGP, irrespective of the BGP configuration.

All EBGP routes with looped AS-Path are silently discarded by default.

This command applies to the current VR or VRF context.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default, the IPv4 unicast address family, applies and the command fails. Similarly an IPv4 peer only supports IPv4 address families and the command fails if an IPv6 address family is specified.

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the MPLS-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following example enables BGP to accept looped BGP routes that contains a maximum of 6 occurrences of receiver's AS-number in AS-Path attribute:

```
configure bgp neighbor 192.162.17.54 allowas-in max-as-occurrence 6
```

History

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure bgp neighbor dampening

```
configure bgp neighbor [all | remoteaddr] {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4]} dampening {{half-life half-life-minutes {reuse-limit reuse-limit-number suppress-limit suppress-limit-number max-suppress max-suppress-minutes} | policy-filter [policy-name | none]}}
```

Description

Configures the route flap dampening feature for a BGP neighbor.

Syntax Description

all	Specifies that the configuration applies to all neighbors in the specified address family. If no address family is specified or if an IPv4 address is specified, the configuration applies to all IPv4 neighbors. If an IPv6 address family is specified, the configuration applies to all IPv6 neighbors.
<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor. The switch uses the IP address format to determine if the address is an IPv4 or IPv6 address.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support. This address family is applicable for PE to PE BGP neighbor sessions only. Using this keyword may prompt warning or error messages if executed for a regular BGP neighbor session, or for a PE to CE neighbor session.
half-life	Specifies the dampening half life. Range is 1 to 45 minutes.
reuse-limit	Specifies the reuse limit. Range is 1 to 20000.
suppress-limit	Specifies the suppress limit. Range is 1 to 20000.
max-suppress	Specifies the maximum hold down time. Range is 1 to 255 minutes.
policy-filter	Specifies a policy.

Default

This feature is disabled by default.

If no address family is specified and an IPv4 address is detected, IPv4 unicast is the default address family.

Usage Guidelines

This command applies to the current VR or VRF context.



If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Note



For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default, the IPv4 unicast address family, applies and the command fails. Similarly an IPv4 peer only supports IPv4 address families and the command fails if an IPv6 address family is specified.

The half life is the period of time, in minutes, during which the accumulated penalty of a route is reduced by half. The range is 1 to 45 minutes, and the default is 15 minutes.

The reuse limit is the penalty value below which a route is used again. The range is 1-20,000, and the default is 750.

The suppress limit is the penalty value above which a route is suppressed. The range is 1-20,000, and the default is 2,000.

The maximum hold down time is the maximum time a route can be suppressed, no matter how unstable it has been, as long as it no longer flaps. The range is 1-255 minutes, and the default is 4 * the half life.

If you change dampening parameters when routes are in suppressed or history state, the new dampening parameters apply only to routes in the active state. Routes in the suppressed or history state continue to use the old dampening parameters until they become active, at which time they use the updated dampening parameters.

Instead of explicitly configuring the dampening parameters using the command line, you can associate a policy using the policy-filter option. Multiple sets of parameters can be supplied using a policy.

Use the following command to disable route flap dampening for BGP neighbors:

```
configure bgp neighbor [remoteaddr | all] {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpnv4]} no-dampening
```

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the MPLS-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command configures route flap dampening to the BGP neighbor at 192.168.1.22 to the default values:

```
configure bgp neighbor 192.168.1.22 dampening
```

History

This command was first available in ExtremeXOS 10.1.



This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp neighbor description

```
configure bgp neighbor [all | remoteaddr] description {description}
```

Description

Configures a description for a BGP neighbor.

Syntax Description

all	Specifies all IPv4 and IPv6 neighbors.
<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
<i>description</i>	Specifies a string used to describe the neighbor.

Default

The description is a NULL string by default.

Usage Guidelines

Use this command to attach a description to a BGP neighbor. This description is displayed in the output of the `show bgp neighbor` command when you specify the detail option, or when you specify a particular neighbor. Enclose the string in double quotes if there are any blank spaces in the string. The maximum length of the string is 56 characters.

If you do not specify the *description* parameter, the description is reset to the default.

This command applies to the current VR or VRF context.

Example

The following command configures the description for the BGP neighbor 192.168.1.22 to Toledo_5:

```
configure bgp neighbor 192.168.1.22 description Toledo_5
```



History

This command was first available in ExtremeXOS 11.0.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp neighbor dont-allowas-in

```
configure bgp neighbor [all | remoteaddr] {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4] } dont-allowas-in
```

Description

Disables EBGP from receiving and accepting a looped EBGP route from the specified neighbor, provided the number of occurrences of local AS number in AS-Path is less than or equal to the value of `as-count`.

Syntax Description

all	Specifies that the configuration change applies to all neighbors in the specified address family. If no address family is specified or if an IPv4 address is specified, the configuration change applies to all IPv4 neighbors. If an IPv6 address family is specified, the configuration change applies to all IPv6 neighbors.
<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor. The switch uses the IP address format to determine if the address is an IPv4 or IPv6 address.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support.

Default

This feature is disabled by default.

If no `as-count` is specified, the `as-count` defaults to 3.

If no address family is specified and an IPv4 address is detected, IPv4 unicast is the default address family.



Usage Guidelines

In a hub and spoke configuration, it becomes necessary to accept an inbound EBGP route even though the route's AS-Path contains the receiver's own AS-number. In such network topologies, this feature can be enabled.



Note

A looped AS path is always allowed for IBGP, irrespective of the BGP configuration.

All EBGP routes with looped AS-Path are silently discarded by default.

This command applies to the current VR or VRF context.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default, the IPv4 unicast address family, applies and the command fails. Similarly an IPv4 peer only supports IPv4 address families and the command fails if an IPv6 address family is specified.

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the MPLS-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

History

This command was first available in ExtremeXOS 12.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp neighbor maximum-prefix

```
configure bgp neighbor [remoteaddr | all] {address-family [ipv4-unicast | ipv4-
multicast | ipv6-unicast | ipv6-multicast | vpnv4]} maximum-prefix number
{{threshold percent} {teardown {holddown-interval seconds}} {send-traps}}
```

Description

Configures the maximum number of IP prefixes accepted from a BGP neighbor.



Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor. The switch uses the IP address format to determine if the address is an IPv4 or IPv6 address.
all	Specifies that the configuration applies to all neighbors in the specified address family. If no address family is specified or if an IPv4 address is specified, the configuration applies to all IPv4 neighbors. If an IPv6 address family is specified, the configuration applies to all IPv6 neighbors.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpnv4	Specifies the VPNv4 address family for Layer 3 VPN support. This address family is applicable for PE to PE BGP neighbor sessions only. This keyword may prompt warning or error messages if executed for a regular BGP neighbor session, or for a PE to CE neighbor session.
<i>number</i>	Specifies the maximum number of prefixes that can be accepted. The range is 0 to 4294967294. A value of 0 disables prefix limit feature.
<i>percent</i>	Specifies the percentage of the maximum prefix (threshold) at which a warning message is printed in the log (and console), and/or a trap is sent to the SNMP manager.
teardown	Specifies that the peer session is torn down when the maximum is exceeded.
<i>seconds</i>	Specifies the length of time before the session is re-established, if the session is torn down due to maximum prefix exceeded. If the hold-down interval is zero or not specified, it is kept down until the peer is enabled. The range is 30 to 86400 seconds.
send-traps	Specifies sending "number of prefix reached threshold" and "number of prefix exceed the max-prefix limit" SNMP traps.

Default

This feature is disabled by default.

The default threshold is 75%.

By default, teardown is not specified.

By default, send-traps is not specified.

Usage Guidelines

This command applies to the current VR or VRF context.



If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Note



For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default, the IPv4 unicast address family, applies and the command fails. Similarly an IPv4 peer only supports IPv4 address families and the command fails if an IPv6 address family is specified.

Configure the peer group before configuring the neighbors. To configure the peer group, use the following command:

```
configure bgp peer-group peer-group-name {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4]} maximum-prefix number
{{threshold percent} {teardown {holddown-interval seconds}} {send-traps}
```

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the MPLS-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command configures the maximum number of IP prefixes accepted from all neighbors to 5000, sets the threshold for warning messages to 60%, and specifies SNMP traps:

```
configure bgp neighbor all maximum-prefix 5000 threshold 60 send-traps
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp neighbor next-hop-self

```
configure bgp neighbor [remoteaddr | all] {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4]} [next-hop-self | no-next-hop-self]
```



Description

Configures the next hop address used in the outgoing updates to be the address of the BGP connection originating the update.

Syntax Description

<i>remoteaddr</i>	Specifies an IP address.
all	Specifies all neighbors.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support. This address family is applicable for PE to PE BGP neighbor sessions only. This keyword may prompt warning or error messages if executed for a regular BGP neighbor session, or for a PE to CE neighbor session.
next-hop-self	Specifies that the next hop address used in the updates be the address of the BGP connection originating it.
no-next-hop-self	Specifies that the next hop address used in the updates not be the address of the BGP connection originating it (lets BGP decide what would be the next hop).

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

This command applies to the current VR or VRF context. These settings apply to the peer group and all neighbors of the peer group.



Note

The BGP neighbor must be disabled before you can change the configuration with this command.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default, the IPv4 unicast address family, applies and the command fails. Similarly an IPv4 peer only supports IPv4 address families and the command fails if an IPv6 address family is specified.



To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the MPLS-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command configures the next hop address used in the updates to be the address of the BGP connection originating it:

```
configure bgp neighbor 172.16.5.25 next-hop-self
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp neighbor no-dampening

```
configure bgp neighbor [remoteaddr | all] {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4]} no-dampening
```

Description

Configures no route flap dampening over BGP peer sessions (disables route flap dampening).

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor. The switch uses the IP address format to determine if the address is an IPv4 or IPv6 address.
all	Specifies that the configuration applies to all neighbors in the specified address family. If no address family is specified or if an IPv4 address is specified, the configuration applies to all IPv4 neighbors. If an IPv6 address family is specified, the configuration applies to all IPv6 neighbors.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.



ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support.

Default

This feature is disabled by default.

If no address family is specified and an IPv4 address is detected, IPv4 unicast is the default address family.

Usage Guidelines

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Note



For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default, the IPv4 unicast address family, applies and the command fails. Similarly an IPv4 peer only supports IPv4 address families and the command fails if an IPv6 address family is specified.

Use the following command to enable route flap dampening for BGP neighbors:

```
configure bgp neighbor [all | remoteaddr] {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4]} dampening {{half-life half-life-minutes {reuse-limit reuse-limit-number suppress-limit suppress-limit-number max-suppress max-suppress-minutes} | policy-filter [policy-name | none]}}
```

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the MPLS-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command disables route flap dampening to the BGP neighbor at 192.168.1.22:

```
configure bgp neighbor 192.168.1.22 no-dampening
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp neighbor password

```
configure bgp neighbor [all | remoteaddr] password [none | {encrypted}
tcpPassword]
```

Description

Configures an RSA Data Security, Inc. MD5 Message-Digest Algorithm secret password for a neighbor.

Syntax Description

all	Specifies all IPv4 and IPv6 neighbors.
<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
none	Specifies not to use a password
encrypted	Specifies an encrypted string; do not use.
<i>tcpPassword</i>	Specifies a password string.

Default

N/A.

Usage Guidelines

This command applies to the current VR or VRF context.

You must disable the BGP neighbor before changing the password.

When a password is configured, TCP RSA Data Security, Inc. MD5 Message-Digest Algorithm authentication is enabled on the TCP connection that is established with the neighbor.

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

To change any one of the following parameters you must disable and re-enable the peer session:

- timer
- source-interface
- soft-in-reset
- password

Changing a route reflector client automatically disables and enables the peer session.



The encrypted option is used by the switch when generating a configuration file, and when parsing a switch-generated configuration file. Do not select the encrypted option in the CLI.

Example

The following command configures the password for a neighbor as Extreme:

```
configure bgp neighbor 192.168.1.5 password extreme
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp neighbor peer-group

```
configure bgp neighbor [all | remoteaddr] peer-group [peer-group-name | none]
{acquire-all}
```

Description

Configures an existing neighbor as the member of a peer group.

Syntax Description

all	Specifies all neighbors.
<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
<i>peer-group-name</i>	Specifies a peer group name.
none	Removes the neighbor from the peer group.
acquire-all	Specifies that all parameters should be inherited by the neighbor from the peer group.

Default

By default, remote AS (if configured for the peer group), source-interface, outbound route policy, send-community and next-hop-self settings are inherited.



Usage Guidelines

This command applies to the current VR or VRF context.

If `acquire-all` is not specified, only the default parameters are inherited by the neighbor.

When you remove a neighbor from a peer group, it retains the parameter settings of the group. The parameter values are not reset to those the neighbor had before it inherited the peer group values.

To create a new neighbor and add it to a BGP peer group, use the following command:

```
create bgp neighbor remoteaddr peer-group peer-group-name {multi-hop}
```

The new neighbor is created as part of the peer group and inherits all of the existing parameters of the peer group. The peer group must have a remote AS configured.

If you are adding an IPv4 peer to a peer group and no IPv4 address family capabilities are assigned to the specified peer group, the IPv4 unicast and multicast address families are automatically enabled for that peer group. If you adding an IPv6 peer to a peer group and no IPv6 address family capabilities are assigned to the peer group, you must explicitly enable the IPv6 address family capabilities you want to support.

Note



If the peer group or any member of the peer group has been configured with an IPv4 or IPv6 address family, the peer group only accepts peers that are configured to use that family. For example, if a peer group is configured for the IPv4 unicast address family, the switch will not allow you to add an IPv6 peer. Likewise, an IPv6 peer group cannot accept an IPv4 peer.

Example

The following command configures an existing neighbor as the member of the peer group `outer`:

```
configure bgp neighbor 192.1.1.22 peer-group outer
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure bgp neighbor route-policy

```
configure bgp neighbor [remoteaddr | all] {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4]} route-policy [in | out] [none | policy]
```

Description

Configures a route map filter for a neighbor.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor. The switch uses the IP address format to determine if the address is an IPv4 or IPv6 address.
all	Specifies that the configuration applies to all neighbors in the specified address family. If no address family is specified or if an IPv4 address is specified, the configuration applies to all IPv4 neighbors. If an IPv6 address family is specified, the configuration applies to all IPv6 neighbors.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support. This address family is applicable for PE to PE BGP neighbor sessions only. This keyword may prompt warning or error messages if executed for a regular BGP neighbor session, or for a PE to CE neighbor session.
in	Specifies to install the filter on the input side.
out	Specifies to install the filter on the output side.
none	Specifies to remove the filter.
<i>policy</i>	Specifies a policy.

Default

If no address family is specified and an IPv4 address is detected, IPv4 unicast is the default address family.

Usage Guidelines

This command applies to the current VR or VRF context.



The policy can be installed on the input or output side of the router. The policy is used to modify or filter the NLRI information and the path attributes associated with it when exchanging updates with the neighbor.



Note

A policy file applied to BGP neighbors cannot have NLRI for both IPv4 and IPv6 address families defined in the same policy file.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default, the IPv4 unicast address family, applies and the command fails. Similarly an IPv4 peer only supports IPv4 address families and the command fails if an IPv6 address family is specified.

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the MPLS-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command configures the route-policy filter for a neighbor based on the policy nosales:

```
configure bgp neighbor 192.168.1.22 route-policy in nosales
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp neighbor route-reflector-client

```
configure bgp neighbor [remoteaddr | all] [route-reflector-client | no-route-reflector-client]
```



Description

Configures a BGP neighbor to be a route reflector client.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
all	Specifies all IPv4 and IPv6 neighbors.
route-reflector-client	Specifies for the BGP neighbor to be a route reflector client.
no-route-reflector-client	Specifies for the BGP neighbor not to be a route reflector client.

Default

N/A.

Usage Guidelines

Another way to overcome the difficulties of creating a fully-meshed AS is to use route reflectors. Route reflectors allow a single router to serve as a central routing point for the AS or sub-AS.

Use this command to implicitly define the router to be a route reflector. The neighbor must be in the same AS as the router.

This command applies to the current VR or VRF context.

When changing the route reflector status of a peer, the peer is automatically disabled and re-enabled and a warning message appears on the console and in the log.

A cluster is formed by the route reflector and its client routers. Peer routers that are not part of the cluster must be fully meshed according to the rules of BGP.

Example

The following command configures a BGP neighbor to be a route reflector client:

```
configure bgp neighbor 192.168.1.5 route-reflector-client
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp neighbor send-community

```
configure bgp neighbor [remoteaddr | all] {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4]} [send-community | dont-send-community] {both | extended | standard}
```

Description

Configures whether the community path attribute associated with a BGP NLRI should be included in the route updates sent to the BGP neighbor.

Syntax Description

<i>remoteaddr</i>	Specifies an IP address of a BGP neighbor.
all	Specifies all neighbors.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support. This address family is applicable for PE to PE BGP neighbor sessions only. This keyword may prompt warning or error messages if executed for a regular BGP neighbor session, or for a PE to CE neighbor session.
send-community	Specifies to include the community path attribute.
dont-send-community	Specifies not to include the community path attribute.
both	Send both standard and extended community attributes to this BGP neighbor, or neighbors in peer group
extended	Send only extended communities to this BGP neighbor or neighbors in peer group
standard	Send only standard communities to this BGP neighbor or neighbors in peer group

Default

If no address family is specified, IPv4 unicast is the default. If no optional keyword (both, standard or extended) is specified, standard is assumed.



Usage Guidelines

A BGP community is a group of BGP destinations that require common handling. ExtremeXOS supports the following well-known BGP community attributes:

- no-export
- no-advertise
- no-export-subconfed

The command is additive; that is, if the command is executed twice with the standard or extended option, both the extended and standard communities are sent to the BGP neighbor.

This command applies to the current VR or VRF context.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Note



For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default, the IPv4 unicast address family, applies and the command fails. Similarly an IPv4 peer only supports IPv4 address families and the command fails if an IPv6 address family is specified.

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the MPLS-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command includes the community path attribute associated with a BGP NLRI in the route updates sent to all BGP neighbors:

```
configure bgp neighbor all send-community
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Options to control the advertisement of extended community attributes were added in ExtremeXOS12.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp neighbor shutdown-priority

```
configure bgp neighbor [all | remoteaddr] shutdown-priority number
```

Description

Configures the shutdown priority for a BGP neighbor.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
all	Specifies all IPv4 and IPv6 neighbors.
<i>number</i>	Specifies the shutdown priority. The range is 0 - 65,535.

Default

The default value is 1024.

Usage Guidelines



Note

This command is not currently supported, and is not recommended for use.

Higher priority values lower the chance of a BGP neighbor to be automatically disabled in case BGP or the system goes to a low memory condition.

Example

The following command configures the shutdown priority of the BGP neighbor 10.0.20.1 to 500:

```
configure bgp neighbor 10.0.20.1 shutdown-priority 1000
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp neighbor soft-reset

```
configure bgp neighbor [remoteaddr | all] {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4]} soft-reset {in | out}
```

Description

Applies the current input or output routing policy to the routing information already exchanged with the neighbor.

Syntax Description

<i>remoteaddr</i>	Specifies an IP address of a BGP neighbor.
all	Specifies all neighbors.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support. This address family is applicable for PE to PE BGP neighbor sessions only. This keyword may prompt warning or error messages if executed for a regular BGP neighbor session, or for a PE to CE neighbor session.
soft-reset	Do a soft reconfiguration for the BGP neighbor.
in	Specifies to apply the input routing policy.
out	Specifies to apply the output routing policy.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

The input/output policy is determined by the route policy configured for the neighbor on the input and/or output side of the router. This command does not affect the switch configuration.

If both the local BGP neighbor and the neighbor router support the route refresh capability (ExtremeWare does not support this feature), a dynamic soft input reset can be performed. The `configure bgp neighbor soft-reset` command triggers the generation of a Route-Refresh message to the neighbor. As a response to the Route-Refresh message, the neighbor sends the entire BGP routing table in updates and the switch applies the appropriate routing policy to the updates.



This command applies to the current VR or VRF context.

If the route-refresh capability is not supported by the neighbor (like ExtremeWare), the `configure bgp neighbor soft-reset` command reprocesses the BGP route database using the policy configured for that neighbor.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Note



For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default, the IPv4 unicast address family, applies and the command fails. Similarly an IPv4 peer only supports IPv4 address families and the command fails if an IPv6 address family is specified.

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the MPLS-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command applies the current input routing policy to the routing information already exchanged with the neighbor:

```
configure bgp neighbor 192.168.1.5 soft-reset in
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp neighbor source-interface

```
configure bgp neighbor [remoteaddr | all] source-interface [any | ipaddress  
ipAddr]
```



Description

Changes the BGP source interface for TCP connections.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
all	Specifies all IPv4 and IPv6 neighbors.
any	Specifies any source interface.
<i>ipAddr</i>	Specifies the IP address of a source interface.

Default

Any.

Usage Guidelines

The source interface IP address must be a valid IP address of any VLAN configured on the switch.

This command applies to the current VR or VRF context.

Example

The following command changes the BGP source interface to 10.43.55.10:

```
configure bgp neighbor 192.168.1.5 source-interface ipaddress 10.43.55.10
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp neighbor timer

```
configure bgp neighbor [remoteaddr | all] timer keep-alive keepalive hold-time holdtime
```



Description

Configures the BGP neighbor timers.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
all	Specifies all IPv4 and IPv6 neighbors.
<i>keepalive</i>	Specifies a BGP neighbor timer keepalive time in seconds. The range is 0 to 21,845 seconds.
<i>holdtime</i>	Specifies a BGP neighbor timer hold time in seconds. The range is 0 and 3to65,535 seconds.

Default

The default keepalive setting is 60 seconds. The default hold time is 180 seconds.

Usage Guidelines

You must disable the BGP neighbor before changing the timer values.

This command applies to the current VR or VRF context.

Example

The following command configures the BGP neighbor timers:

```
configure bgp neighbor 192.168.1.5 timer keep-alive 120 hold-time 360
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp neighbor weight

```
configure bgp neighbor [remoteaddr | all] weight weight
```



Description

Assigns a locally-used weight to a neighbor connection for the route selection algorithm.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
all	Specifies all IPv4 and IPv6 neighbors.
<i>weight</i>	Specifies a BGP neighbor weight.

Default

By default, the weight is 1.

Usage Guidelines

All routes learned from this peer are assigned the same weight. The route with the highest weight is more preferable when multiple routes are available to the same network. The range is 0 to 65,535.

BGP selects routes based on the following precedence (from highest to lowest):

- higher weight
- higher local preference
- shortest length (shortest AS path)
- lowest origin code
- lowest MED
- route from external peer
- lowest cost to Next Hop
- lowest routerID

This command applies to the current VR or VRF context.

Example

The following command assigns a locally used weight of 10 to a neighbor connection:

```
configure bgp neighbor 192.168.1.5 weight 10
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp peer-group allowas-in

```
configure bgp peer-group peer-group-name {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4]} allowas-in {max-as-occurrence as-count}
```

Description

Configures BGP to receive and accept a looped BGP route from the neighbors of the specified peer group, provided the number of occurrences of local AS number in AS-Path is less than or equal to that specified in *as-count*.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support.
<i>as-count</i>	The maximum number of occurrences of local AS number in the received route AS-Path. If the number of occurrences of local AS number in AS-Path is more than <i>as-count</i> , the route is not accepted. The valid range is from 1-16.

Default

This feature is disabled by default.

If no *as-count* is specified, the *as-count* defaults to 3.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

In a hub and spoke configuration, it becomes necessary to accept an inbound BGP route even though the route's AS-Path contains the receiver's own AS-number. In such network topologies, this feature can be enabled.

This feature can also be enabled for both IBGP and EBGp neighbors, wherever necessary.



This command applies to the current VR or VRF context.



Note

BGP neighbors do not inherit the allow-as configuration from their peer group unless you explicitly specify the `acquire-all` option when adding a neighbor to a peer-group.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

If the specified peer group contains IPv6 peers, it is an IPv6 peer group and you must specify an IPv6 address-family. When the specified peer group is an IPv6 peer group, this command fails if no address family is specified or if an IPv4 address-family is specified. This command also fails if an IPv6 address family is specified for an IPv4 peer-group.

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the MPLS-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following example enables BGP to accept looped BGP routes that contains a maximum of 8 occurrences of receiver's AS-number in AS-Path attribute:

```
configure bgp peer-group internal allow-as max-as-occurrence 8
```

History

This command was first available in ExtremeXOS 12.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp peer-group dampening

```
configure bgp peer-group peer-group-name {address-family [ipv4-unicast | ipv4-
multicast | ipv6-unicast | ipv6-multicast | vpnv4]} dampening {{half-life half-
life-minutes {reuse-limit reuse-limit-number suppress-limit suppress-limit-number
max-suppress max-suppress-minutes}} | policy-filter [policy-name | none]}
```



Description

Configures route flap dampening for a BGP peer group.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support.
<i>half-life-minutes</i>	Specifies the dampening half life.
<i>reuse-limit-number</i>	Specifies the reuse limit.
<i>suppress-limit-number</i>	Specifies the suppress limit.
<i>max-suppress-minutes</i>	Specifies the maximum hold down time.
<i>policy-name</i>	Specifies a policy
none	Removes any policy association.

Default

This feature is disabled by default.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

This command applies to the current VR or VRF context.

The half life is the period of time, in minutes, during which the accumulated penalty of a route is reduced by half. The range is 1 to 45 minutes, and the default is 15 minutes.

The reuse limit is the penalty value below which a route is used again. The range is 1-20,000, and the default is 750.

The suppress limit is the penalty value above which a route is suppressed. The range is 1-20,000, and the default is 2,000.

The maximum hold down time is the maximum time a route can be suppressed, no matter how unstable it has been, as long as it no longer flaps. The range is 1-255 minutes, and the default is 4 * the half life.

If you change dampening parameters when routes are in suppressed or history state, the new dampening parameters apply only to routes in the active state. Routes in the suppressed or history state continue to use the old dampening parameters until they become active, at which time they use the updated dampening parameters.



Instead of explicitly configuring the dampening parameters using the command line, you can associate a policy using the `policy-filter` option. Multiple sets of parameters can be supplied using a policy.

Use the following command to disable route flap dampening for a BGP peer-group:

```
configure bgp peer-group peer-group-name {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4]} no-dampening
```

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Note



If the specified peer group contains IPv6 peers, it is an IPv6 peer group and you must specify an IPv6 address-family. When the specified peer group is an IPv6 peer group, this command fails if no address family is specified or if an IPv4 address-family is specified. This command also fails if an IPv6 address family is specified for an IPv4 peer-group.

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the MPLS-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command configures route flap dampening for the BGP peer group `outer`:

```
configure bgp peer-group outer dampening
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp peer-group dont-allowas-in

```
configure bgp peer-group peer-group-name {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4]} dont-allowas-in
```



Description

Disables BGP from receiving and accepting a looped BGP route from the neighbors of the specified peer group, provided the number of occurrences of local AS number in AS-Path is less than or equal to that specified in `as-count`.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support.

Default

This feature is disabled by default.

If no `as-count` is specified, the `as-count` defaults to 3.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

This command applies to the current VR or VRF context.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Note



If the specified peer group contains IPv6 peers, it is an IPv6 peer group and you must specify an IPv6 address-family. When the specified peer group is an IPv6 peer group, this command fails if no address family is specified or if an IPv4 address-family is specified. This command also fails if an IPv6 address family is specified for an IPv4 peer-group.

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the MPLS-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Note



BGP neighbors do not inherit the `allowas-in` configuration from their peer group unless you explicitly specify the `acquire-all` option when adding a neighbor to a peer-group.

History

This command was first available in ExtremeXOS 12.1.



Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp peer-group maximum-prefix

```
configure bgp peer-group peer-group-name {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4]} maximum-prefix number
{{threshold percent} {teardown {holddown-interval seconds}} {send-traps}
```

Description

Configures the maximum number of IP prefixes accepted for all neighbors in the peer group.

Syntax Description

name	Specifies a peer group.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support.
<i>number</i>	Specifies the maximum number of prefixes that can be accepted. The range is 0 to 4294967294. A value of 0 disables prefix limit feature.
<i>percent</i>	Specifies the percentage of the maximum prefix (threshold) at which a warning message is printed in the log (and on the console). An SNMP trap can also be sent.
teardown	Specifies that the peer session is torn down when the maximum is exceeded.
<i>seconds</i>	Specifies the length of time before the session is re-established, if the session has been torn down due to exceeding the max limit. If the hold down interval is 0 or not specified, it is kept down until the peer is enabled. The range is 30 to 86400 seconds.
send-traps	Specifies sending “number of prefix reached threshold” and “number of prefix exceed the max-prefix limit” SNMP traps.

Default

This feature is disabled by default.



The default threshold is 75%.

By default, teardown is not specified.

By default, send-traps is not specified.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

This command applies to the current VR or VRF context.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

If the specified peer group contains IPv6 peers, it is an IPv6 peer group and you must specify an IPv6 address-family. When the specified peer group is an IPv6 peer group, this command fails if no address family is specified or if an IPv4 address-family is specified. This command also fails if an IPv6 address family is specified for an IPv4 peer-group.

Configure the peer group before configuring the neighbors. To configure the neighbors, use the following command:

```
configure bgp neighbor 192.168.1.1 maximum-prefix
```

After you enter this command, the switch automatically disables and enables all neighbors in the peer group before the change takes effect.

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the MPLS-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command configures the maximum number of IP prefixes accepted from the peer group outer to 5000, sets the threshold for warning messages to 60%, and specifies SNMP traps:

```
configure bgp peer-group outer maximum-prefix 5000 threshold 60 send-traps
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp peer-group next-hop-self

```
configure bgp peer-group peer-group-name {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4]} [next-hop-self | no-next-hop-self]
```

Description

Configures the next hop address used in the updates to be the address of the BGP connection originating the update.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support.
next-hop-self	Specifies that the next hop address used in the updates be the address of the BGP connection originating it.
no-next-hop-self	Specifies that the next hop address used in the updates not be the address of the BGP connection originating it (Let the BGP protocol decide the next hop).

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

This command applies to the current VR or VRF context.

These settings apply to the peer group and all neighbors of the peer group.



If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Note

If the specified peer group contains IPv6 peers, it is an IPv6 peer group and you must specify an IPv6 address-family. When the specified peer group is an IPv6 peer group, this command fails if no address family is specified or if an IPv4 address-family is specified. This command also fails if an IPv6 address family is specified for an IPv4 peer-group.

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the MPLS-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command configures the next hop address used in the updates to be the address of the BGP connection originating it:

```
configure bgp peer-group outer next-hop-self
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp peer-group no-dampening

```
configure bgp peer-group peer-group-name {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4]} no-dampening
```

Description

Configures no route flap dampening for a BGP peer group (disables route flap dampening).



Syntax Description

<i>peer-group-name</i>	Specifies a BGP peer group.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support.

Default

This feature is disabled by default.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Note



If the specified peer group contains IPv6 peers, it is an IPv6 peer group and you must specify an IPv6 address-family. When the specified peer group is an IPv6 peer group, this command fails if no address family is specified or if an IPv4 address-family is specified. This command also fails if an IPv6 address family is specified for an IPv4 peer-group.

Use the following command to enable route flap dampening for a BGP peer-group:

```
configure bgp peer-group peer-group-name {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4]} dampening {{half-life half-life-minutes {reuse-limit reuse-limit-number suppress-limit suppress-limit-number max-suppress max-suppress-minutes}} | policy-filter [policy-name | none]}
```

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the MPLS-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command disables route flap dampening to the BGP peer group outer:

```
configure bgp peer-group outer no-dampening
```

History

This command was first available in ExtremeXOS 10.1.



This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp peer-group password

```
configure bgp peer-group peer-group-name password [none | tcpPassword]
```

Description

Configures the TCP RSA Data Security, Inc. MD5 Message-Digest Algorithm secret password for a peer group and all neighbors of the peer group.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
none	Specifies no password.
<i>tcpPassword</i>	Specifies a password.

Default

N/A.

Usage Guidelines

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

This command applies to the current VR or VRF context.

After you enter this command, the switch automatically disables and enables all neighbors in the peer group before the change takes effect.

Example

The following command configures the password as Extreme for the peer group outer and its neighbors:

```
configure bgp peer-group outer password extreme
```



History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp peer-group remote-AS-number

```
configure bgp peer-group peer-group-name remote-AS-number number
```

Description

Configures the remote AS number for a peer group and all the neighbors of the peer group.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
<i>number</i>	Specifies a remote AS number. The range is 1 to 4294967295.

Default

N/A.

Usage Guidelines

The AS number is a 4-byte AS number in either the ASPLAIN or the ASDOT format as described in RFC 5396, Textual Representation of Autonomous System (AS) Numbers.

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

After you enter this command, the switch automatically disables and enables all neighbors in the peer group before the change takes effect.

Example

The following command configures the remote AS number for the peer group `outer` and its neighbors using the ASPLAIN 4-byte AS number format:

```
configure bgp peer-group outer remote-AS-number 65536
```



The following command configures the remote AS number for the peer group abc and its neighbors using the ASDOT 4-byte AS number format:

```
configure bgp peer-group abc remote-AS-number 1.10
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for 4-byte AS numbers was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp peer-group route-policy

```
configure bgp peer-group peer-group-name {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4]} route-policy [in | out] [none | policy]
```

Description

Configures the policy for a peer group and all the neighbors of the peer group.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support.
in	Specifies to install the policy on the input side.
out	Specifies to install the policy on the output side.
none	Specifies to remove the filter.
<i>policy</i>	Specifies a policy.



Default

There is no default policy configuration.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

This command applies to the current VR or VRF context.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Note



If the specified peer group contains IPv6 peers, it is an IPv6 peer group and you must specify an IPv6 address-family. When the specified peer group is an IPv6 peer group, this command fails if no address family is specified or if an IPv4 address-family is specified. This command also fails if an IPv6 address family is specified for an IPv4 peer-group.

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the MPLS-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command configures the route policy for the peer group outer and its neighbors using the policy nosales:

```
configure bgp peer-group outer route-policy in nosales
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure bgp peer-group route-reflector-client

```
configure bgp peer-group peer-group-name [route-reflector-client | no-route-reflector-client]
```

Description

Configures all the peers in a peer group to be a route reflector client.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
route-reflector-client	Specifies that all the neighbors in the peer group be a route reflector client.
no-route-reflector-client	Specifies that all the neighbors in the peer group not be a route reflector client.

Default

N/A.

Usage Guidelines

This command implicitly defines this router to be a route reflector.

This command applies to the current VR or VRF context.

The peer group must be in the same AS of this router.

Example

The following command configures the peer group outer as a route reflector client:

```
configure bgp peer-group outer route-reflector-client
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure bgp peer-group send-community

```
configure bgp peer-group peer-group-name {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4]} [send-community | dont-send-community] {both | extended | standard}
```

Description

Configures whether communities should be sent to neighbors as part of route updates.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support.
send-community	Specifies that communities are sent to neighbors as part of route updates.
dont-send-community	Specifies that communities are not sent to neighbors as part of route updates.
both	Send both standard and extended community attributes to this BGP neighbor, or neighbors in peer group.
extended	Send only extended communities to this BGP neighbor or neighbors in peer group.
standard	Send only standard communities to this BGP neighbor or neighbors in peer group.

Default

If no address family is specified, IPv4 unicast is the default. If no optional keyword (both, standard or extended) is specified, standard is assumed.

Usage Guidelines

These settings apply to the peer group and all neighbors of the peer group.

This command applies to the current VR or VRF context.



If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Note



If the specified peer group contains IPv6 peers, it is an IPv6 peer group and you must specify an IPv6 address-family. When the specified peer group is an IPv6 peer group, this command fails if no address family is specified or if an IPv4 address-family is specified. This command also fails if an IPv6 address family is specified for an IPv4 peer-group.

The command is additive; that is, if the command is executed twice with the standard or extended option, both the extended and standard communities are sent to the BGP neighbor.

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the MPLS-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command configures communities to be sent to neighbors as part of route updates:

```
configure bgp peer-group outer send-community
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Options to control the advertisement of extended community attributes were added in ExtremeXOS12.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp peer-group soft-reset

```
configure bgp peer-group peer-group-name {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpnv4]} soft-reset {in | out}
```



Description

Applies the current input/output routing policy to the neighbors in the peer group.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support.
in	Specifies to apply the input routing policy.
out	Specifies to apply the output routing policy.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

The input/output routing policy is determined by the route policy configured for the neighbors in the peer group on the input/output side of the router. This command does not affect configuration of the switch.

This command applies to the current VR or VRF context.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Note



If the specified peer group contains IPv6 peers, it is an IPv6 peer group and you must specify an IPv6 address-family. When the specified peer group is an IPv6 peer group, this command fails if no address family is specified or if an IPv4 address-family is specified. This command also fails if an IPv6 address family is specified for an IPv4 peer-group.

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

Any configuration change with this command automatically disables and enables the neighbors before the changes.

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the MPLS-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.



Example

The following command applies the current input routing policy to the neighbors in the peer group outer:

```
configure bgp peer-group outer soft-reset in
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp peer-group source-interface

```
configure bgp peer-group peer-group-name source-interface [any | ipaddress ipAddr]
```

Description

Configures the source interface for a peer group and all the neighbors of the peer group.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
any	Specifies any source interface.
<i>ipAddr</i>	Specifies an interface.

Default

N/A.

Usage Guidelines

The source interface IP address must be a valid IP address of a VLAN configured on the switch.

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.



This command applies to the current VR or VRF context.

After you enter this command, the switch automatically disables and enables the neighbors so that the changes can take effect.

Example

The following command configures the source interface for the peer group `outer` and its neighbors on 10.34.25.10:

```
configure bgp peer-group outer source-interface ipaddress 10.34.25.10
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp peer-group timer

```
configure bgp peer-group peer-group-name timer keep-alive seconds hold-time seconds
```

Description

Configures the keepalive timer and hold timer values for a peer group and all the neighbors of the peer group.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
keep-alive <i>seconds</i>	Specifies a keepalive time in seconds. Range is 0 to 21845.
hold-time <i>seconds</i>	Specifies a hold-time in seconds. Range is 0 and 3 to 65535.

Default

N/A.



Usage Guidelines

This command applies to the current VR or VRF context.

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

After you enter this command, the switch automatically disables and enables all neighbors in the peer group before the change takes effect.

Example

The following command configures the keepalive timer and hold timer values for the peer group `outer` and its neighbors:

```
configure bgp peer-group outer timer keep-alive 30 hold-time 90
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp peer-group weight

```
configure bgp peer-group peer-group-name weight number
```

Description

Configures the weight for the peer group and all the neighbors of the peer group.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
<i>number</i>	Specifies a BGP peer group weight. Range is 0 to 65535.

Default

N/A.



Usage Guidelines

BGP selects routes based on the following precedence (from highest to lowest):

- higher weight
- higher local preference
- shortest length (shortest AS path)
- lowest origin code
- lowest MED
- route from external peer
- lowest cost to Next Hop
- lowest routerID

This command applies to the current VR or VRF context.

Example

The following command configures the weight for the peer group `outer` and its neighbors:

```
configure bgp peer-group outer weight 5
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp restart

```
configure bgp restart [none | planned | unplanned | both | aware-only]
```

Description

Configures the router as a graceful BGP restart router.

Syntax Description

none	Do not act as a graceful BGP restart router.
planned	Only act as a graceful BGP restart router for planned restarts.
unplanned	Only act as a graceful BGP restart router for unplanned restarts.



both	Act as a graceful BGP restart router for both planned and unplanned restarts.
aware-only	Only act as a graceful BGP receiver (helper) router.

Default

The default is none; graceful restart is disabled.

Usage Guidelines

This command configures the router as a graceful BGP router. You can decide to configure a router to enter graceful restart for only planned restarts, for only unplanned restarts, or for both. Also, you can decide to configure a router to be a receiver only (which helps a restarting BGP router to perform the graceful restart process), and not to do graceful restarts itself.

After a graceful restart, the switch preserves the time stamps for all BGP routes in the RIB that were received before the stale timer expired. After restart, the capabilities for all BGP peers are renegotiated.



Note

End of Restart (EOR) messages are not sent to BGP peers if the graceful restart feature is disabled.

This command cannot be used while BGP is enabled globally on the switch.

Example

The following command configures a router to perform graceful BGP restarts only for planned restarts:

```
configure bgp restart planned
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp restart address-family

```
configure bgp restart [add | delete] address-family [ipv4-unicast | ipv4-  
multicast | ipv6-unicast | ipv6-multicast]
```



Description

Configures the address family used with graceful BGP restart.

Syntax Description

add	Add the address family.
delete	Remove the address family.
address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.

Default

The default is IPv4 unicast.

Usage Guidelines

Before you can enter this command, you must disable BGP services on the switch with the `disable bgp` command.

This command configures the address family participating in graceful BGP restart. An address family can be added or deleted. By adding an address family, BGP instructs the switch to preserve BGP routes of that address family during a graceful restart. The local OPEN message contains all the added address families.

Note

When graceful restart is enabled on the switch, the IPv4 unicast address family support is added by default. Graceful restart for other address families must be explicitly added using this command.



For BGP graceful restart to inter-operate with Cisco routers, any restarting routers connected to Cisco routers must be configured with the command, `enable bgp neighbor capability`, in the following form: `enable bgp neighbor remoteaddr capability ipv4-unicast`. The command must be executed before BGP is enabled globally on the switch.

Example

The following command configures a router to add IPv4 unicast addresses to graceful BGP restarts:

```
configure bgp restart add address-family ipv4-unicast
```

History

This command was first available in ExtremeXOS 11.4.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp restart restart-time

```
configure bgp restart restart-time seconds
```

Description

Configures the restart time used with graceful BGP restart. This is the maximum time a receiver router waits for a restarting router to come back up.

Syntax Description

<i>seconds</i>	Specifies the restart time. The range is 1 to 3600 seconds.
----------------	---

Default

The default is 120 seconds.

Usage Guidelines

Before you can enter this command, you must disable BGP services on the switch with the `disable bgp` command.

This command configures the restart timer. This timer is started on the receiver router when it detects the neighbor router is restarting (usually when the peer TCP session is reset). At that time, routes from the restarting router are marked as stale, but are preserved in the routing table. The timer is stopped when the restarting BGP neighbor goes to the ESTABLISHED state (it has finished restarting). If the timer expires, the stale routes are deleted.

Example

The following command configures the graceful BGP restart timer:

```
configure bgp restart restart-time 200
```

History

This command was first available in ExtremeXOS 11.4.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp restart stale-route-time

```
configure bgp restart stale-route-time seconds
```

Description

Configures the stale route timer used with graceful BGP restart. This is the maximum time to hold stale paths on receiver routers while its neighbor gracefully restarts.

Syntax Description

<i>seconds</i>	Specifies the stale route time. The range is 1 to 3600 seconds.
----------------	---

Default

The default is 360 seconds.

Usage Guidelines

Before you can enter this command, you must disable BGP services on the switch with the `disable bgp` command.

This command configures the stale route timer. This timer is started when the restarting BGP peer goes to the ESTABLISHED state after it restarts. The timer is stopped when the restarting BGP peer sends EOR messages for all address families. When the timer is stopped, or it expires, the stale routes are deleted.

Example

The following command configures the graceful BGP stale route timer:

```
configure bgp restart stale-route-time 400
```

History

This command was first available in ExtremeXOS 11.4.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp restart update-delay

```
configure bgp restart update-delay seconds
```

Description

Configures the update delay timer used with graceful BGP restart. This is the maximum time to delay updating BGP routes to the local IP route table.

Syntax Description

<i>seconds</i>	Specifies the stale route time. The range is 1 to 3600 seconds.
----------------	---

Default

The default is 600 seconds.

Usage Guidelines

Before you can enter this command, you must disable BGP services on the switch with the `disable bgp` command.

This command configures the update delay timer. Usually, a restarting router waits to receive EOR messages from all the receiving BGP neighbors before it starts the route update. Otherwise, it does the route selection when the timer expires.

Example

The following command configures the graceful BGP update delay timer:

```
configure bgp restart update-delay 800
```

History

This command was first available in ExtremeXOS 11.4.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp routerid

```
configure bgp routerid router identifier
```

Description

Changes the router identifier.

Syntax Description

<i>router identifier</i>	Specifies a router identifier in the IPv4 address format.
--------------------------	---

Default

N/A.

Usage Guidelines

BGP must be disabled before changing the router ID.

BGP selects routes based on the following precedence (from highest to lowest):

- Higher weight
- Higher local preference
- Shortest length (shortest AS path)
- Lowest origin code
- Lowest MED
- Route from external peer
- Lowest cost to Next Hop
- Lowest router ID

This command applies to the current VR or VRF context.



Example

The following command changes the router ID:

```
configure bgp routerid 192.1.1.13
```



Note

To remove the configured bgp routerid, give routerid value as 0.0.0.0 i.e. configure bgp routerid 0.0.0.0.

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure bgp soft-reconfiguration

```
configure bgp soft-reconfiguration
```

Description

Immediately applies the route policy associated with the network command, aggregation, import, and redistribution.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command does not affect the switch configuration.

This command applies to the current VR or VRF context.



Example

The following command applies the route policy associated with the network command, aggregation, import, and redistribution:

```
configure bgp soft-reconfiguration
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

create bgp neighbor peer-group

```
create bgp neighbor remoteaddr peer-group peer-group-name {multi-hop}
```

Description

Creates a new neighbor and makes it part of the peer group.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
<i>peer-group-name</i>	Specifies a peer group.
multi-hop	Specifies to allow connections to EBGP peers that are not directly connected.

Default

N/A.

Usage Guidelines

You can specify an IPv4 or IPv6 address for the BGP peer. The address can be a global unicast or a link-local address. IPv6 link-local remote addresses are supported only for EBGP single-hop peerings.

If you are adding an IPv4 peer to a peer group and no IPv4 address family capabilities are assigned to the specified peer group, the IPv4 unicast and multicast address families are automatically enabled for that peer group. If you adding an IPv6 peer to a peer group and no IPv6 address family capabilities are



assigned to the peer group, you must explicitly enable the IPv6 address family capabilities you want to support.



Note

If the peer group or any member of the peer group has been configured with an IPv4 or IPv6 address family, the peer group only accepts peers that are configured to use that family. For example, if a peer group is configured for the IPv4 unicast address family, the switch will not allow you to add an IPv6 peer. Likewise, an IPv6 peer group cannot accept an IPv4 peer.

If the multihop keyword is not specified, the IP addresses of the EBGP speaker and peer must belong to the same subnet.

All the parameters of the neighbor are inherited from the peer group. The peer group should have the remote AS configured.

To add an existing neighbor to a peer group, use the following command:

```
configure bgp neighbor [all | remoteaddr] peer-group [peer-group-name | none]
{acquire-all}
```

If you do not specify acquire-all, only the mandatory parameters are inherited from the peer group. If you specify acquire-all, all of the parameters of the peer group are inherited. This command disables the neighbor before adding it to the peer group.

Example

The following command creates a new neighbor and makes it part of the peer group outer:

```
create bgp neighbor 192.1.1.22 peer-group outer
```

The following example specifies how to create a neighbor peer group in a VRF (PE - CE neighbor session):

```
virtual-router <vr_vrf_name>
create bgp neighbor <remoteaddr> remote-AS-number <asNumber> {multi-hop}
create bgp neighbor <remoteaddr> peer-group <peer-group-name> {multi-hop}
delete bgp [{neighbor} <remoteaddr> | neighbor all ]
[create | delete] bgp peer-group <peer-group-name>
```

BGP maintains a separate RIB (RIB-In, RIB-Loc and RIB-Out) for each of the VRF it is configured to run. So routes received from a peer in VRF1 are not mixed up with routes from a peer in VRF2. Additionally, BGP routes in a VRF are regular IPv4 routes of address family ipv4. The BGP decision algorithm occurs inside a VRF and is not impacted by any BGP activity in other VRF. There can be two BGP neighbors with the same peer IP address in two different VRFs.

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.



Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for L3 VPN was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

create bgp neighbor remote-AS-number

```
create bgp neighbor remoteaddr remote-AS-number as-number {multi-hop}
```

Description

Creates a new BGP peer.

Syntax Description

<i>remoteaddr</i>	Specifies an IP address of the BGP neighbor.
<i>as-number</i>	Specifies a remote AS number. The range is 1 to 4294967295.
multi-hop	Specifies to allow connections to EBGP peers that are not directly connected.

Default

N/A.

Usage Guidelines

You can specify an IPv4 or IPv6 address for the BGP peer. The address can be a global unicast or a link-local address. IPv6 link-local remote addresses are supported only for EBGP single-hop peerings.

If the multihop keyword is not specified, the IP addresses of the EBGP speaker and peer must belong to the same subnet.

The AS number is a 4-byte AS number in either the ASPLAIN or the ASDOT format as described in RFC 5396, Textual Representation of Autonomous System (AS) Numbers.

If the AS number is the same as the AS number provided in the configure bgp as command, then the peer is consider an IBGP peer, otherwise the neighbor is an EBGP peer. The BGP session to a newly created peer is not started until the `enable bgp neighbor` command is issued.



Example

The following command specifies a BGP peer AS number using the ASPLAIN 4-byte AS number format:

```
create bgp neighbor 10.0.0.1 remote-AS-number 65540
```

The following command specifies a BGP peer AS number using the ASDOT 4-byte AS number format:

```
create bgp neighbor 10.0.0.1 remote-AS-number 1.5
```

The following command specifies a BGP peer using an IPv6 address:

```
create bgp neighbor fe80::204:96ff:fe1e:a8f1%vlan1 remote-AS-number 200
```

The following example specifies how to create a neighbor peer group in a VRF (PE - CE neighbor session):

```
virtual-router <vr_vrf_name>
create bgp neighbor <remoteaddr> remote-AS-number <asNumber> {multi-hop}
create bgp neighbor <remoteaddr> peer-group <peer-group-name> {multi-hop}
delete bgp [{neighbor} <remoteaddr> | neighbor all ]
[create | delete] bgp peer-group <peer-group-name>
```

BGP maintains a separate RIB (RIB-In, RIB-Loc and RIB-Out) for each of the VRF it is configured to run. So routes received from a peer in VRF1 are not mixed up with routes from a peer in VRF2. Additionally, BGP routes in a VRF are regular IPv4 routes of address family ipv4. The BGP decision algorithm occurs inside a VRF and is not impacted by any BGP activity in other VRF. There can be two BGP neighbors with the same peer IP address in two different VRFs.

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for 4-byte AS numbers was first available in ExtremeXOS 12.4.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for L3 VPN was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



create bgp peer-group

```
create bgp peer-group peer-group-name
```

Description

Creates a new peer group.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
------------------------	-------------------------

Default

N/A.

Usage Guidelines

You can use BGP peer groups to group together up to 512 BGP neighbors. All neighbors within the peer group inherit the parameters of the BGP peer group. The following mandatory parameters are shared by all neighbors in a peer group:

- source-interface
- out-nlri-filter
- out-aspath-filter
- out-route-policy
- send-community
- next-hop-self

The BGP peer group name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

No IPv4 or IPv6 address family capabilities are added to a new peer group. When the first IPv4 peer is added to a peer group, the IPv4 unicast and multicast families are enabled by default. No IPv6 address family capabilities are automatically added when an IPv6 peer is added to a peer group; you must explicitly add any IPv6 address family capabilities that you want for a peer group.

Example

The following command creates a new peer group named outer:

```
create bgp peer-group outer
```



The following example specifies how to create a neighbor peer group in a VRF (PE – CE neighbor session):

```
virtual-router <vr_vrf_name>
create bgp neighbor <remoteaddr> remote-AS-number <asNumber> {multi-hop}
create bgp neighbor <remoteaddr> peer-group <peer-group-name> {multi-hop}
delete bgp [{neighbor} <remoteaddr> | neighbor all ]
[create | delete] bgp peer-group <peer-group-name>
```

BGP maintains a separate RIB (RIB-In, RIB-Loc and RIB-Out) for each of the VRF it is configured to run. So routes received from a peer in VRF1 are not mixed up with routes from a peer in VRF2. Additionally, BGP routes in a VRF are regular IPv4 routes of address family ipv4. The BGP decision algorithm occurs inside a VRF and is not impacted by any BGP activity in other VRF. There can be two BGP neighbors with the same peer IP address in two different VRFs.

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for L3 VPN was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

delete bgp neighbor

```
delete bgp neighbor [remoteaddr | all]
```

Description

Deletes one or all BGP neighbors.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of the BGP neighbor to be deleted.
all	Specifies all IPv4 and IPv6 neighbors.

Default

N/A.



Usage Guidelines

You can use global unicast remote addresses to delete all BGP peer types. You can use link-local remote address to delete only EBGp single-hop peers.

Example

The following command deletes the specified IPv4 BGP neighbor:

```
delete bgp neighbor 192.168.1.17
```

The following command deletes the specified IPv6 BGP neighbor:

```
delete bgp neighbor fe80::204:96ff:fe1e:a8f1%v1an1
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

delete bgp peer-group

```
delete bgp peer-group peer-group-name
```

Description

Deletes a peer group.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
------------------------	-------------------------

Default

N/A.



Usage Guidelines

Use this command to delete a specific BGP peer group.

Example

The following command deletes the peer group named outer:

```
delete bgp peer-group outer
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable bgp

disable bgp

Description

Disables BGP.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Use this command to disable BGP on the router.



Example

The following command disables BGP:

```
disable bgp
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable bgp adj-rib-out

```
disable bgp adj-rib-out
```

Description

Disables local storage of Adj-Rib-Out (ARO) data to reduce memory usage.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

The default configuration for this feature conserves memory usage by not storing the Adj-RIB-Out for each peer. This results in reduced performance during outbound route advertisements, withdraws, outbound policy evaluations, and the display of transmitted routes in response to CLI commands.

If your configuration has a relatively low number of peers, you can enable this feature and benefit from the increased performance. If your configuration has a relatively large number of peers, you might want to disable this feature to reduce memory usage.

This command applies to the BGP instance for the current VR or VRF context.



Example

The following command disables ARO data storage:

```
disable bgp adj-rib-out
```

History

This command was first available in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on the platforms that support this Core license feature as listed in the *ExtremeXOS Concepts Guide, Feature License Requirements*

disable bgp advertise-inactive-route

```
disable bgp {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast]} advertise-inactive-route
```

Description

Disables advertisement of BGP inactive routes, which are defined as those routes that rated best by BGP and not best in the IP routing table.

Syntax Description

address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.
-----------------------	--

Default

Disabled.

If no address family is specified, IPv4 unicast is the default address family.

Usage Guidelines

This command can be successfully executed only when BGP is globally disabled. If you want to disable inactive route advertisement and BGP is enabled, you must disable BGP (`disable bgp`), disable this feature, and then enable BGP (`enable bgp`).

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Example

The following command disables inactive route advertisement for IPv4 unicast traffic:

```
disable bgp address-family ipv4-unicast advertise-inactive-route
```

History

This command was first available in ExtremeXOS 12.1.3.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable bgp aggregation

disable bgp aggregation

Description

Disables BGP route aggregation.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Route aggregation is the process of combining the characteristics of several routes so that they are advertised as a single route. Aggregation reduces the amount of information that a BGP speaker must store and exchange with other BGP speakers. Reducing the information that is stored and exchanged also reduces the size of the routing table.

Use this command to disable BGP route aggregation.



Example

The following command disables BGP route aggregation:

```
disable bgp aggregation
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable bgp always-compare-med

```
disable bgp always-compare-med
```

Description

Disables BGP from comparing Multi Exit Discriminators (MEDs) for paths from neighbors in different Autonomous Systems (AS).

Syntax Description

This command has no arguments or variables.

Default

ExtremeXOS does not compare MEDs for paths from neighbors in different AS.

Usage Guidelines

The MED is one of the parameters that is considered when selecting the best path among many alternative paths. The path with a lower MED is preferred over a path with a higher MED. By default, during the best path selection process, MED comparison is done only among paths from the same AS.

BGP must be disabled before you can change the configuration with this command.



Example

The following command disables MED from being used in comparison among paths from different AS:

```
disable bgp always-compare-med
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable bgp community format

```
disable bgp community format AS-number : number
```

Description

Disables the AS-number:number format of display for communities in the output of show commands.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Using this command, communities are displayed as a single decimal value.

Example

The following command disables the AS-number:number format of display for communities:

```
disable bgp community format AS-number : number
```



History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable bgp export

```
disable bgp export route_type {{address-family} address_family}
```

For Layer 3 VPNs:

```
disable bgp export route_type {{address-family} address_family}
```

Description

Disables BGP from exporting routes from other protocols to BGP peers.

Syntax Description

bgp	For Layer 3 VPNs, this specifies that BGP routes learned from CE routers are to be exported to remote PE routers.
<i>route_type</i>	Specifies the BGP export route type.
address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.

Default

Disabled.

If no address family is specified, IPv4 unicast is the default.



Note

You must specify an IPv6 address family for an IPv6 peer, because an IPv6 peer does not support the default IPv4 unicast address family. Similarly, if you specify an IPv4 peer and an address family in the command, an IPv4 address family must be specified.

Usage Guidelines

The exporting of routes between any two routing protocols is a discrete configuration function. For example, you must configure the switch to export routes from OSPF to BGP and, if desired, you must configure the switch to export routes from BGP to OSPF. You must first configure both protocols and



then verify the independent operation of each. Then you can configure the routes to export from OSPF to BGP, and the routes to export from BGP to OSPF.

You can use policies to associate BGP attributes including Community, NextHop, MED, Origin, and Local Preference with the routes. Policies can also be used to filter out exported routes.

Using the export command to redistribute routes complements the redistribution of routes using the `configure bgp add network` command. The `configure bgp add network` command adds the route to BGP only if the route is present in the routing table. The `enable bgp export` command redistributes an individual route from the routing table to BGP. If you use both commands to redistribute routes, the routes redistributed using the network command take precedence over routes redistributed using the export command.

Note



For this command to execute, the specified protocol must support the specified address family. For example, the command fails if you specify OSPF and the IPv6 unicast address family. You can specify blackhole, direct, static, and IS-IS routes with IPv4 or IPv6 address families.

For Layer 3 VPNs, the `disable bgp export` command must be entered in the context of the VRF that supports the Layer 3 VPN.

When the export source is the Layer 3 VPN, you can specify `direct`, or `remote-vpn` to disable route export to the VRF. The destination address family must be `ipv4-unicast`.

When the export source is the VRF, you can specify `direct`, or `bgp` to disable route export to the VPN. The destination address family must be `vpn4`.

Example

The following command disables BGP from exporting routes from the OSPF protocol to BGP peers:

```
disable bgp export ospf
```

The following command disables the export of BGP routes from a VRF to a VPN:

```
disable bgp export bgp address-family vpnv4
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

The blackhole option was added in ExtremeXOS 12.1.3.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable bgp export vr

```
disable bgp export {vr} vr_name route_type {address-family} vpnv4
```

Description

For IPv4 and IPv6 routes, this command disables the PE router to export and redistribute local VRF routes to remote PE routers through BGP.

Syntax Description

vr	Specifies the source VPN VRF of the exported routes.
<i>vr_name</i>	Specifies the name of the source VPN VRF.
<i>route_type</i>	Specifies the source or origin of the route types to be exported to remote PE routers. Valid Types: blackhole, direct, and bgp.
<i>address-family</i>	Specifies the address family for the exported routes. Valid types are ipv4-unicast, vpnv4.
vpnv4	Specifies that routes from the VRF are exported as vpnv4 routes over MPBGP.

Default

Disabled.

Usage Guidelines

This command disables a PE router to advertise learned routes from CE routers to remote PE routers in a Service Provider's backbone. Executing this command allows the PE router to convert VRF native IPv4 routes into VPN-IPv4 routes and advertise to all remote PE BGP neighbors as VPN-IPv4 routes.

- For Layer 3 VPNs, you must enter the **disable bgp export** vrcommand in the context of the VRF that supports the Layer 3 VPN.
- When the export source is the Layer 3 VPN, you can specify direct, or remote-vpn to disable route export to the VRF. The destination address family must be ipv4-unicast.
- This export command is applicable in Parent VR context only. If you execute it in a VRF context, an error message is returned.
- The source VPN VRF must be a child of the Parent VR.
- BGP need not be added to a VPN VRF to export routes from a VPN VRF.
- The direction of where the redistribution is targeted is implicit on the keywords used, For eg:- remote-vpn only applies to remote routes from PE redistributed to CE, hence we cannot use it with address family vpnv4. Similarly bgp only applies to EBGp routes from CE exported as VPN routes,



hence we use it only with address family vpnv4. Other sources such as “static” and “direct” are redistributed both ways.

Example

The following command disables BGP to advertise a vpnv4 route named “corp1_vpn_vrf”:

```
switch 19 # disable bgp export "corp1_vpn_vrf" bgp address-family vpnv4
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

The blackhole option was added in ExtremeXOS 12.1.3.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable bgp fast-external-fallover

```
disable bgp fast-external-fallover
```

Description

Disables BGP fast external fallover functionality.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables the BGP fast external fallover on the router. This command applies to all directly-connected external BGP neighbors.



When BGP fast external fallover is enabled, the directly-connected EBGP neighbor session is immediately reset when the connecting link goes down.

If BGP fast external fallover is disabled, BGP waits until the default hold timer expires (3 keepalives) to reset the neighboring session. In addition, BGP might teardown the session somewhat earlier than hold timer expiry if BGP detects that the TCP session and it's directly connected link is broken (BGP detects this while sending or receiving data from TCP socket).

Example

The following command disables BGP fast external fallover:

```
disable bgp fast-external-fallover
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable bgp neighbor

```
disable bgp neighbor [remoteaddr | all]
```

Description

Disables the BGP session.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
all	Specifies all IPv4 and IPv6 neighbors.

Default

Disabled.

Usage Guidelines

After the session has been disabled, all the information in the route information base (RIB) for the neighbor is flushed.



This command applies to the current VR or VRF context.

Example

The following command disables the BGP session:

```
disable bgp neighbor 192.168.1.17
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable bgp neighbor capability

```
disable bgp neighbor [all | remoteaddr] capability [ipv4-unicast | ipv4-multicast  
| ipv6-unicast | ipv6-multicast | vpnv4 | route-refresh]
```

Description

This command disables an address family or the route-refresh capability for one or all neighbors.

Syntax Description

all	Specifies that the configuration applies to all neighbors in the specified address family. If no address family is specified or if an IPv4 address is specified, the configuration applies to all IPv4 neighbors. If an IPv6 address family is specified, the configuration applies to all IPv6 neighbors.
<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor. The switch uses the IP address format to determine if the address is an IPv4 or IPv6 address.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpnv4	Specifies the VPNv4 address family for Layer 3 VPN support.
route-refresh	Specifies ROUTE-REFRESH message capabilities.



Default

The following capabilities are enabled by default for IPv4 peers: IPv4 unicast, IPv4 multicast, and route refresh.

The following capabilities are enabled by default for IPv6 peers: route refresh.

Usage Guidelines

This command applies to the current VR or VRF context.

Note



To inter-operate with Cisco routers for BGP graceful restart, you must enable IPv4 unicast address capability.

For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default, the IPv4 unicast address family, applies and the command fails. Similarly an IPv4 peer only supports IPv4 address families and the command fails if an IPv6 address family is specified.

Example

The following command disables the route-refresh feature for all neighbors:

```
disable bgp neighbor all capability route-refresh
```

The following command disables the VPNv4 address family for a neighbor:

```
disable bgp neighbor 192.168.96.235 capability vpnv4
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



disable bgp neighbor originate-default

```
disable bgp [{neighbor} remoteaddr | neighbor all] {address-family [ipv4-unicast
| ipv4-multicast | ipv6-unicast | ipv6-multicast]} originate-default
```

Description

Removes a default route to a single BGP neighbor or to all BGP neighbors.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
all	Specifies that the configuration applies to all neighbors in the specified address family. If no address family is specified, the configuration applies to the IP Unicast family on all IPv4 peers. If an IPv4 address family is specified, the configuration applies to all IPv4 neighbors. If an IPv6 address family is specified, the configuration applies to all IPv6 neighbors.
address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.

Default

Disabled. BGP does not automatically originate and advertise default routes to BGP neighbors.

If no address family is specified and an IPv4 address is detected, IPv4 unicast is the default address family.

Note



You must specify an IPv6 address family for an IPv6 peer, because an IPv6 peer does not support the default IPv4 unicast address family. Similarly, if you specify an IPv4 peer and an address family in the command, an IPv4 address family must be specified.

Usage Guidelines

This command can be successfully executed at any time, irrespective of whether local BGP or the remote BGP peer is enabled or disabled.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Example

The following command removes default routes for IPv4 unicast traffic for all BGP peer nodes:

```
disable bgp neighbor all originate-default
```



History

This command was first available in ExtremeXOS 12.1.3.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable bgp neighbor remove-private-AS-numbers

```
disable bgp neighbor [remoteaddr | all] remove-private-AS-numbers
```

Description

Disables the removal of private AS numbers from the AS path in route updates sent to EBGP peers.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
all	Specifies all IPv4 and IPv6 neighbors.

Default

Disabled.

Usage Guidelines

Private AS numbers are AS numbers in the range 64512 through 65534. You can remove private AS numbers from the AS path attribute in updates that are sent to external BGP (EBGP) neighbors.

Possible reasons for using private AS numbers include:

- The remote AS does not have officially allocated AS numbers.
- You want to conserve AS numbers if you are multi-homed to the local AS.

Private AS numbers should not be advertised on the Internet. Private AS numbers can only be used locally within an administrative domain. Therefore, when routes are advertised out to the Internet, the private AS number can be stripped out from the AS paths of the advertised routes using this feature.

This command applies to the current VR or VRF context.



Example

The following command disables the removal of private AS numbers from the AS path in route updates sent to the EBGP peers:

```
disable bgp neighbor 192.168.1.17 remove-private-AS-numbers
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable bgp neighbor soft-in-reset

```
disable bgp neighbor [all | remoteaddr] {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpnv4]} soft-in-reset
```

Description

Disables the soft input reset feature.

Syntax Description

all	Specifies that the configuration applies to all neighbors in the specified address family. If no address family is specified, the configuration applies to the IP Unicast family on all IPv4 peers. If an IPv4 address family is specified, the configuration applies to all IPv4 neighbors. If an IPv6 address family is specified, the configuration applies to all IPv6 neighbors.
<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpnv4	Specifies the VPNv4 address family for Layer 3 VPN support.



Default

Disabled.

If no address family is specified and an IPv4 address is detected, IPv4 unicast is the default address family.

Usage Guidelines

Disabling the soft input reset feature can potentially limit the amount of system memory consumed by the RIB-in.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Note



For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default, the IPv4 unicast address family, applies and the command fails. Similarly an IPv4 peer only supports IPv4 address families and the command fails if an IPv6 address family is specified.

Before you can change the configuration with this command, you must disable BGP, and you must disable the corresponding BGP neighbor session using the following command:

```
disable bgp neighbor [remoteaddr | all]
```

To disable this feature on Layer 3 VPNs, you must do so in the context of the MPLS-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

This command applies to the current VR or VRF context.

Example

The following command disables the soft input reset for the neighbor at 192.168.1.17:

```
disable bgp neighbor 192.168.1.17 soft-in-reset
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable bgp peer-group

```
disable bgp peer-group peer-group-name
```

Description

Disables a BGP peer group and all its BGP neighbors.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
------------------------	-------------------------

Default

Disabled.

Usage Guidelines

This command applies to the current VR or VRF context.

Example

The following command disables the BGP peer group outer:

```
disable bgp peer-group outer
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



disable bgp peer-group capability

```
disable bgp peer-group peer-group-name capability [ipv4-unicast | ipv4-multicast
| ipv6-unicast | ipv6-multicast | vpnv4 | route-refresh]
```

Description

This command disables an address family or the route-refresh capability for a peer group.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpnv4	Specifies the VPNv4 address family for Layer 3 VPN support.
route-refresh	Specifies ROUTE-REFRESH message capabilities.

Default

All capabilities are enabled for IPv4 peer groups by default.

Only the route refresh capability is enabled for peer groups by default.

Usage Guidelines

This command applies to the current VR or VRF context.



Note

To inter-operate with Cisco routers for BGP graceful restart, you must enable IPv4 unicast address capability.

Example

The following command disables the route-refresh feature for the peer group outer:

```
disable bgp peer-group outer route-refresh
```

The following command disables the VPNv4 address family for a peer group:

```
disable bgp peer-group backbone capability vpnv4
```



History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable bgp peer-group originate-default

```
disable bgp {peer-group} peer-group-name {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast]} originate-default
```

Description

Removes default routes to all BGP neighbors in the specified peer group.

Syntax Description

<i>peer-group-name</i>	Specifies the BGP peer group for which the default routes are removed.
address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.

Default

Disabled. BGP does not automatically originate and advertise default routes to BGP neighbors.

Usage Guidelines

This command can be successfully executed at any time, irrespective of whether local BGP or the remote BGP peers are enabled or disabled.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Note



If the specified peer group contains IPv6 peers, it is an IPv6 peer group and you must specify an IPv6 address-family. When the specified peer group is an IPv6 peer group, this command fails if no address family is specified or if an IPv4 address-family is specified. This command also fails if an IPv6 address family is specified for an IPv4 peer-group.



Example

The following command removes default routes for IPv4 unicast traffic for all nodes in the test BGP peer group:

```
disable bgp peer-group test originate-default
```

History

This command was first available in ExtremeXOS 12.2.2.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable bgp peer-group remove-private-AS-numbers

```
disable bgp peer-group peer-group-name remove-private-AS-numbers
```

Description

Disables the removal of private autonomous system (AS) numbers from the AS_Path attribute of outbound updates.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
------------------------	-------------------------

Default

Disabled.

Usage Guidelines

This command applies to the current VR or VRF context.

Example

The following command disables the BGP peer group outer from removing private AS numbers:

```
disable bgp peer-group outer remove-private-AS-numbers
```



History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable bgp peer-group soft-in-reset

```
disable bgp peer-group peer-group-name {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast]} soft-in-reset
```

Description

Disables the soft input reset feature.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.

Default

Disabled.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Disabling the soft input reset feature can potentially limit the amount of system memory consumed by the RIB-in.

After you enter this command, the switch automatically disables and enables all neighbors in the peer group before the change takes effect.

This command applies to the current VR or VRF context.



If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Note

If the specified peer group contains IPv6 peers, it is an IPv6 peer group and you must specify an IPv6 address-family. When the specified peer group is an IPv6 peer group, this command fails if no address family is specified or if an IPv4 address-family is specified. This command also fails if an IPv6 address family is specified for an IPv4 peer-group.

Example

The following command disables the soft input reset feature:

```
disable bgp peer-group outer soft-in-reset
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable bgp

enable bgp

Description

Enables BGP.

Syntax Description

This command has no arguments or variables.

Default

Disabled.



Usage Guidelines

This command enables the Border Gateway Protocol (BGP) on the router. Before invoking this command, the local AS number and BGP router ID must be configured.

Example

The following command enables BGP:

```
enable bgp
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable bgp adj-rib-out

```
enable bgp adj-rib-out
```

Description

Enables local storage of Adj-Rib-Out (ARO) data to support the display of transmitted routes by other CLI commands.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

The default configuration for this feature conserves memory usage by not storing the Adj-RIB-Out for each peer. This results in reduced performance during outbound route advertisements, withdraws, outbound policy evaluations, and the display of transmitted routes in response to CLI commands.



If your configuration has a relatively low number of peers, you can enable this feature and benefit from the increased performance. If your configuration has a relatively large number of peers, you might want to disable this feature to reduce memory usage.

This command applies to the BGP instance for the current VR or VRF context.

Example

The following command enables ARO data storage:

```
enable bgp adj-rib-out
```

History

This command was first available in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on the platforms that support this Core license feature as listed in the *ExtremeXOS Concepts Guide, Feature License Requirements*

enable bgp advertise-inactive-route

```
enable bgp {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast]} advertise-inactive-route
```

Description

Enables advertisement of BGP inactive routes, which are defined as those routes that are rated best by BGP and not best in the IP routing table.

Syntax Description

address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.
-----------------------	--

Default

Disabled.

If no address family is specified, IPv4 unicast is the default address family.

Usage Guidelines

This command can be successfully executed only when BGP is globally disabled. It is best to enable this feature before you enable BGP (`enable bgp`). If BGP is enabled, you must disable BGP (`disable bgp`), enable this feature, and then enable BGP.



If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Example

The following command enables inactive route advertisement for IPv4 unicast traffic:

```
enable bgp address-family ipv4-unicast advertise-inactive-route
```

History

This command was first available in ExtremeXOS 12.1.3.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable bgp aggregation

enable bgp aggregation

Description

Enables BGP route aggregation.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Route aggregation is the process of combining the characteristics of several routes so that they are advertised as a single route. Aggregation reduces the amount of information that a BGP speaker must store and exchange with other BGP speakers. Reducing the information that is stored and exchanged also reduces the size of the routing table.

To use BGP route aggregation, follow these steps:



- 1 Enable aggregation using the following command:

```
enable bgp aggregation
```

- 2 Create an aggregate route using the following command:

```
configure bgp add aggregate-address {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast]} ipaddress/masklength {as-match | as-set} {summary-only} {advertise-policy policy} {attribute-policy policy}
```

Example

The following command enables BGP route aggregation:

```
enable bgp aggregation
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable bgp always-compare-med

```
enable bgp always-compare-med
```

Description

Enables BGP to use the Multi Exit Discriminator (MED) from neighbors in different autonomous systems (ASs) in the route selection algorithm.

Syntax Description

This command has no arguments or variables.

Default

Disabled.



Usage Guidelines

MED is only used when comparing paths from the same AS, unless `always-compare-med` is enabled. When this command is issued, MEDs from different AS are used in comparing paths. A MED value of zero is treated as the lowest MED and therefore the most preferred route.

BGP must be disabled before you can change the configuration with this command.

Example

The following command enables BGP to use the Multi Exit Discriminator (MED) from neighbors in different autonomous systems in the route selection algorithm:

```
enable bgp always-compare-med
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable bgp community format

```
enable bgp community format AS-number : number
```

Description

Enables the `as-number:number` format of display for the communities in the output of `show` commands.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

If not enabled, the communities are displayed as a single decimal value.



Example

The following command enables the AS-number:number format of display for communities:

```
enable bgp community format AS-number : number
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable bgp export

For IPv4 and IPv6 routes:

```
enable bgp export route_type {{address-family} address_family} {export-policy policy-name}
```

For VPNv4 routes:

```
enable bgp export remote-vpn {{address-family} ipv4-unicast} {export-policy policy-name}
```

Description

For IPv4 and IPv6 routes, this command enables the export of routes learned from BGP peers to the specified protocol.

For VPNv4 routes, this command enables the exchange of routes between a BGP PE router and a CE router.

Syntax Description

bgp	For Layer 3 VPNs, this specifies that BGP routes learned from CE routers are to be exported to remote PE routers.
<i>route_type</i>	Specifies the BGP export route type. Valid <i>route_type</i> values are: blackhole; direct; isis; isis-level-1; isis-level-2; isis-level-1-external; isis-level-2-external; ospf; ospf-extern1; ospf-extern2; ospf-inter; ospf-intra; rip; static; ospfv3; ospfv3-extern1; ospfv3-extern2; ospfv3-inter; ospfv3-intra; ripng;
<i>address-family</i>	Valid <i>address_family</i> values are: ipv4-unicast; ipv4-multicast; ipv6-unicast; ipv6-multicast



remote-vpn	For Layer 3 VPNs, this specifies that BGP routes learned from remote PE routers are to be exported to the local VRF.
<i>policy-name</i>	Name of policy to be associated with network export. Policy can filter and/or change the route parameters.

Default

Disabled.

If no address family is specified for an IPv6 protocol, the default IPv6 unicast family applies; otherwise if no address family is specified, IPv4 unicast is the default.

Usage Guidelines

The exporting of routes between any two routing protocols is a discrete configuration function. For example, you must configure the switch to export routes from OSPF to BGP and, if desired, you must configure the switch to export routes from BGP to OSPF. You must first configure both protocols and then verify the independent operation of each. Then, you can configure the routes to export from OSPF to BGP, and the routes to export from BGP to OSPF.

You can use a policy to associate BGP attributes including Community, NextHop, MED, Origin, and Local Preference with the routes. A policy can also be used to filter out exported routes.

Using the export command to redistribute routes complements the redistribution of routes using the `configure bgp add network` command. The `configure bgp add network` command adds the route to BGP only if the route is present in the routing table. The `enable bgp export` command redistributes an individual route from the routing table to BGP. If you use both commands to redistribute routes, the routes redistributed using the network command take precedence over routes redistributed using the export command.

Note



For this command to execute, the specified protocol must support the specified address family. For example, the command fails if you specify OSPF and the IPv6 unicast address family. You can specify blackhole, direct, static, and IS-IS routes with IPv4 or IPv6 address families.

To export Layer 3 VPN routes to the CE peer in a VPN VRF, the source must be remote-vpn and destination address family must be ipv4-unicast.

Example

The following command enables BGP to export OSPF routes to other BGP routers:

```
enable bgp export ospf
```



The following command enables export of Layer 3 VPN Routes received from the PE Core in a VPN-VRF to its CE peers:

```
enable bgp export remote-vpn address-family ipv4-unicast
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

The blackhole option was added in ExtremeXOS 12.1.3.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable bgp export vr

```
enable bgp export {vr} vr_name route_type {address-family} vpnv4 {export-policy policy_name}
```

Description

For IPv4 and IPv6 routes, this command enables the PE router to export and redistribute local VRF routes to remote PE routers through BGP .

Syntax Description

vr	Specifies the source VPN VRF of the exported routes .
<i>vr_name</i>	Specifies the name of the source VPN VRF.
<i>route_type</i>	Specifies the source or origin of the route types to be exported to remote PE routers. Valid Types: blackhole, direct, and bgp, and static .
<i>address-family</i>	Specifies the address family for the exported routes. Valid types are vpnv4.
export-policy vpnv4	(Optional) The export policy can be specified when you enable bgp export. Specifies that routes from the VRF are exported as vpnv4 routes over MPBGP.
<i>policy_name</i>	Name of export policy to be associated with export of VRF routes into BGP's VPN-IPv4 domain for advertisement to other PE routers.



Default

Disabled.

Usage Guidelines

This command enables a PE router to advertise learned routes from CE routers to remote PE routers in a Service Provider's backbone. Executing this command allows the PE router to convert VRF native IPv4 routes into VPN-IPv4 routes and advertise to all remote PE BGP neighbors as VPN-IPv4 routes.

- This export command is applicable in Parent VR context only. If you execute it in a VRF context, an error message is returned.
- The source VPN VRF must be a child of the Parent VR.
- BGP need not be added to a VPN VRF to export routes from a VPN VRF.
- The direction of where the redistribution is targeted is implicit on the keywords used. Similarly `bgp` only applies to EBGp routes from CE exported as VPN routes, hence we use it only with address family `vpn4`. Other sources such as "static" and "direct" are redistributed both ways.
- Use `show vr parent_vr_name` to check routes exported from various VPN VRFs into the MBGP's VPN-IPv4 domain.
- Use `show vr vpn_vrf_name` to check routes exported from a VPN VRF into the MBGP's VPN-IPv4 domain.

Example

The following command enables BGP to advertise a `vpn4` route named "corp1_vpn_vrf":

```
switch 19 # enable bgp export "corp1_vpn_vrf" bgp address-family vpn4
```

History

This command was first added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable bgp fast-external-fallover

```
enable bgp fast-external-fallover
```

Description

Enables BGP fast external fallover functionality.



Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables the BGP fast external fallover on the router. This command applies to all directly-connected external BGP neighbors.

When BGP fast external fallover is enabled, the directly-connected EBGP neighbor session is immediately reset when the connecting link goes down.

If BGP fast external fallover is disabled, BGP waits until the default hold timer expires (3 keepalives) to reset the neighboring session. In addition, BGP might teardown the session somewhat earlier than hold timer expiry if BGP detects that the TCP session and it's directly connected link is broken (BGP detects this while sending or receiving data from TCP socket).

Example

The following command enables BGP fast external fallover:

```
enable bgp fast-external-fallover
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable bgp neighbor

```
enable bgp neighbor [remoteaddr | all]
```

Description

Enables the BGP session. The neighbor must be created before the BGP neighbor session can be enabled.



Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
<i>all</i>	Specifies all IPv4 and IPv6 neighbors.

Default

Disabled.

Usage Guidelines

To create a new neighbor and add it to a BGP peer group, use the following command:

```
create bgp neighbor remoteaddr peer-group peer-group-name {multi-hop}
```

This command applies to the current VR or VRF context.

Example

The following command enables the BGP neighbor session:

```
enable bgp neighbor 192.168.1.17
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable bgp neighbor capability

```
enable bgp neighbor [all | remoteaddr] capability [ipv4-unicast | ipv4-multicast  
| ipv6-unicast | ipv6-multicast | vpnv4 | route-refresh]
```

Description

This command enables multi protocol BGP (MBGP) and route-refresh capabilities for one or all BGP neighbors.



Syntax Description

all	Specifies that the configuration applies to all neighbors in the specified address family. If no address family is specified or if an IPv4 address is specified, the configuration applies to all IPv4 neighbors. If an IPv6 address family is specified, the configuration applies to all IPv6 neighbors.
<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor. The switch uses the IP address format to determine if the address is an IPv4 or IPv6 address.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies VPN ipv4 unicast address family for a BGP neighbor. This is a required configuration for PE to PE BGP neighbor session. You must configure it before you enable a neighbor.
route-refresh	Specifies ROUTE-REFRESH message capabilities.

Default

The following capabilities are enabled by default for IPv4 peers: IPv4 unicast, IPv4 multicast, and route refresh.

The following capabilities are enabled by default for IPv6 peers: route refresh.

Note



For IPv4 peers, the IPv4 unicast and IPv4 multicast capabilities are enabled by default to support legacy peers that do not support MBGP. All other capabilities (except route-refresh) are disabled by default.

Usage Guidelines

When you change the capability configuration, you must enable the BGP neighbor before the configuration becomes active. If the BGP neighbor was enabled before the change, you must disable and enable the BGP neighbor. After the capabilities have been enabled, the BGP neighbor announces its capabilities to neighbors in an OPEN message.

When one or more address families are enabled, routes from the specified address families are updated, accepted, and installed. If more than one address family capability is enabled, or if the VPNv4 address family is enabled, the MBGP extension is automatically enabled. To disable MBGP, you must disable all enabled address families.

To support Layer 3 VPNs, you must enable the VPNv4 address family for all MBGP peers that will distribute VPNv4 routes across the service provider backbone. The VPNv4 address family must be enabled on the MPLS-enabled VR; it is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Use the **vpn4** keyword for all PE to PE BGP neighbor sessions. This instructs BGP to negotiate the vpn4 address family in an open message with other PE routers. If this command is executed when a



BGP neighbor session is established, it will take effect only after BGP session is reset. We recommend that you execute this command when a BGP neighbor is operationally down. Do not issue this command for a neighbor that is part of a VRF (PE - CE), or a warning message will be displayed.

**Note**

To inter-operate with Cisco routers for BGP graceful restart, you must enable the IPv4 unicast address capability.

This command applies to the current VR or VRF context.

**Note**

For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default, the IPv4 unicast address family, applies and the command fails. Similarly an IPv4 peer only supports IPv4 address families and the command fails if an IPv6 address family is specified.

**Note**

You must enable a VPN IPv4 unicast address family for a BGP neighbor for a PE to PE BGP neighbor session before you enable the neighbor.

Example

The following command enables the route-refresh feature for all neighbors:

```
enable bgp neighbor all capability route-refresh
```

The following command enables the VPNv4 address family for a BGP neighbor:

```
virtual router corp1_vrf  
enable bgp neighbor 192.168.96.235 capability vpnv4
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



enable bgp neighbor originate-default

```
enable bgp [{neighbor} remoteaddr | neighbor all] {address-family [ipv4-unicast |
ipv4-multicast | ipv6-unicast | ipv6-multicast]} originate-default {policy policy-
name}
```

Description

Enables the origination and advertisement of a default route to a single BGP neighbor or to all BGP neighbors.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
all	Specifies that the configuration applies to all neighbors in the specified address family. If no address family is specified, the configuration applies to the IP Unicast family on all IPv4 peers. If an IPv4 address family is specified, the configuration applies to all IPv4 neighbors. If an IPv6 address family is specified, the configuration applies to all IPv6 neighbors.
address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.
<i>policy-name</i>	Specifies a policy to be applied to the default route origination.

Default

Disabled. BGP does not automatically originate and advertise default routes to BGP neighbors.

If no address family is specified and an IPv4 address is detected, IPv4 unicast is the default address family.

Note



You must specify an IPv6 address family for an IPv6 peer, because an IPv6 peer does not support the default IPv4 unicast address family. Similarly, if you specify an IPv4 peer and an address family in the command, an IPv4 address family must be specified.

Usage Guidelines

This command can be successfully executed at any time, irrespective of whether local BGP or the remote BGP peer is enabled or disabled. The default route or routes are created regardless of whether or not there are matching entries in the IP route table.

When a BGP neighbor is added to a peer group, it does not inherit the default route origination configuration from the peer group. Also, default route origination for a neighbor and the associated peer group can be different.

If a policy is configured and specified in the command, a default route can be originated only if there is a route in the local BGP RIB that matches the policy's match rules. The default route's attribute can be modified using the same policy file by including statements in the set block of the policy.



If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Example

The following command enables the origination and advertisement of default routes for IPv4 unicast traffic for all BGP peer nodes:

```
enable bgp neighbor all originate-default
```

History

This command was first available in ExtremeXOS 12.3.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable bgp neighbor remove-private-AS-numbers

```
enable bgp neighbor [remoteaddr | all] remove-private-AS-numbers
```

Description

Enables the removal of private AS numbers from the AS path in route updates sent to EBGP peers.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
all	Specifies all IPv4 and IPv6 neighbors.

Default

Disabled.

Usage Guidelines

Private AS numbers are AS numbers in the range 64512 through 65534. You can remove private AS numbers from the AS path attribute in updates that are sent to external BGP (EBGP) neighbors. Possible reasons for using private AS numbers include:



- The remote AS does not have officially allocated AS numbers.
- You want to conserve AS numbers if you are multi-homed to the local AS.

Private AS numbers should not be advertised on the Internet. Private AS numbers can only be used locally within an administrative domain. Therefore, when routes are advertised out to the Internet, the routes can be stripped out from the AS paths of the advertised routes using this feature.

This command applies to the current VR or VRF context.

Example

The following command enables the removal of private AS numbers from the AS path in route updates sent to the EBGP peers:

```
enable bgp neighbor 192.168.1.17 remove-private-AS-numbers
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable bgp neighbor soft-in-reset

```
enable bgp neighbor [all | remoteaddr] {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpnv4]} soft-in-reset
```

Description

Enables the soft input reset feature.

Syntax Description

all	Specifies that the configuration applies to all neighbors in the specified address family. If no address family is specified, the configuration applies to the IP Unicast family on all IPv4 peers. If an IPv4 address family is specified, the configuration applies to all IPv4 neighbors. If an IPv6 address family is specified, the configuration applies to all IPv6 neighbors.
<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.



ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support.

Default

Disabled.

If no address family is specified and an IPv4 address is detected, IPv4 unicast is the default address family.

Usage Guidelines

Before you can change the configuration with this command, you must disable BGP, and you must disable the corresponding BGP neighbor session using the following command:

```
disable bgp neighbor [remoteaddr | all]
```

To enable this feature on Layer 3 VPNs, you must do so in the context of the MPLS-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

This command applies to the current VR or VRF context.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Note



For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default, the IPv4 unicast address family, applies and the command fails. Similarly an IPv4 peer only supports IPv4 address families and the command fails if an IPv6 address family is specified.

Example

The following command enables the soft recognition feature:

```
enable bgp neighbor 192.168.1.17 soft-in-reset
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.



Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable bgp peer-group

```
enable bgp peer-group peer-group-name
```

Description

Enables a peer group and all the neighbors of a peer group.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
------------------------	-------------------------

Default

Disabled.

Usage Guidelines

You can use BGP peer groups to group together up to 200512 BGP neighbors. All neighbors within the peer group inherit the parameters of the BGP peer group. The following mandatory parameters are shared by all neighbors in a peer group:

- source-interface
- out-nlri-filter
- out-aspath-filter
- out-route-map
- send-community
- next-hop-self

This command applies to the current VR or VRF context.

Example

The following command enables the BGP peer group outer and all its neighbors:

```
enable bgp peer-group outer
```



History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable bgp peer-group capability

```
enable bgp peer-group peer-group-name capability [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4 | route-refresh ]
```

Description

This command enables BGP Multiprotocol (MP) and route-refresh capabilities for a peer-group.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
route-refresh	Specifies ROUTE-REFRESH message capabilities.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support.

Default

All capabilities are enabled for IPv4 peer groups by default.

The route refresh capability is enabled for IPv6 peer groups by default.

Usage Guidelines

This command enables BGP Multiprotocol or route-refresh capabilities for a peer group. When you change the capability configuration, you must enable the BGP peer group before the configuration becomes active. If the BGP peer group was enabled before the change, you must disable and enable the BGP peer group. After the capabilities have been enabled, the BGP peer announces its capabilities to neighbors in an OPEN message.



When one or more address families are enabled, routes from the specified address families are updated, accepted, and installed. If more than one address family capability is enabled, or if the VPNv4 address family is enabled, the MBGP extension is automatically enabled. To disable MBGP, you must disable all enabled address families.

A peer group can be configured for either IPv4 or IPv6 address families, but not both. Because a peer-group cannot support both IPv4 and IPv6 peers, the switch prevents the enabling of address families that are not compatible with peers that are already in the peer-group. Similarly if a particular address family is enabled for the peer-group, a peer that is incompatible with the existing peer-group configuration cannot be added to the group.

To support Layer 3 VPNs, you must enable the VPNv4 address family for all MBGP peers that will distribute VPNv4 routes across the service provider backbone. The VPNv4 address family must be enabled on the MPLS-enabled VR; it is not supported for BGP neighbors on the CE (VRF) side of the PE router.

This command applies to the current VR or VRF context.

**Note**

To inter-operate with Cisco routers for BGP graceful restart, you must enable IPv4 unicast address capability.

Example

The following command enables the route-refresh feature for the peer group outer:

```
enable bgp peer-group outer capability route-refresh
```

The following command enables the VPNv4 address family for a peer group:

```
enable bgp peer-group backbone capability vpnv4
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



enable bgp peer-group originate-default

```
enable bgp {peer-group} peer-group-name {address-family [ipv4-unicast | ipv4-
multicast | ipv6-unicast | ipv6-multicast]} originate-default {policy policy_name}
```

Description

Enables the origination and advertisement of default routes to all BGP neighbors in the specified peer group.

Syntax Description

peer-group <i>peer-group-name</i>	Specifies the BGP peer group for which the default routes are originated and advertised.
address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.
<i>policy_name</i>	Specifies a policy to be applied to the default routes during origination.

Default

Disabled. BGP does not automatically originate and advertise default routes to BGP neighbors.

Usage Guidelines

This command can be successfully executed at any time, irrespective of whether local BGP or the remote BGP peers are enabled or disabled. The default routes are created regardless of whether or not there are matching entries in the IGP route table.

When a BGP neighbor is added to a peer group, it does not inherit the default route origination configuration from the peer group. Also, default route origination for a neighbor and the associated peer group can be different.

If a policy is configured and specified in the command, a default route can be originated only if there is a route in the local BGP RIB that matches the policy's match rules. The default route's attribute can be modified using the same policy file by including statements in the set block of the policy.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Note



If the specified peer group contains IPv6 peers, it is an IPv6 peer group and you must specify an IPv6 address-family. When the specified peer group is an IPv6 peer group, this command fails if no address family is specified or if an IPv4 address-family is specified. This command also fails if an IPv6 address family is specified for an IPv4 peer-group.



Example

The following command enables the origination and advertisement of default routes for IPv4 unicast traffic for all nodes in the test BGP peer group:

```
enable bgp peer-group test originate-default
```

History

This command was first available in ExtremeXOS 12.2.2.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable bgp peer-group remove-private-AS-numbers

```
enable bgp peer-group peer-group-name remove-private-AS-numbers
```

Description

Enables the removal of private autonomous system (AS) numbers from the AS_Path attribute of outbound updates.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
------------------------	-------------------------

Default

Disabled.

Usage Guidelines

This command applies to the current VR or VRF context.

Example

The following command enables the BGP peer group outer from removing private AS numbers:

```
enable bgp peer-group outer remove-private-AS-numbers
```



History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable bgp peer-group soft-in-reset

```
enable bgp peer-group peer-group-name {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4]} soft-in-reset
```

Description

Enables the soft input reset feature.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support.

Default

Disabled.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Disabling the soft input reset feature can potentially limit the amount of system memory consumed by the RIB-in.

After you enter this command, the switch automatically disables and enables all neighbors in the peer group before the change takes effect.

This command applies to the current VR or VRF context.



If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Note

If the specified peer group contains IPv6 peers, it is an IPv6 peer group and you must specify an IPv6 address-family. When the specified peer group is an IPv6 peer group, this command defaults to IPv4 unicast if no address family is specified. This command fails if an IPv6 address family is specified for an IPv4 peer-group.

Example

The following command enables the soft input reset feature:

```
enable bgp peer-group outer soft-in-reset
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show bgp

show bgp

Description

Displays BGP configuration information.

Syntax Description

This command has no arguments or variables.



Default

N/A.

Usage Guidelines

None.

Example

The following command examples display various BGP configurations:

Output for **show bgp** for a VRF (PE-CE Protocol, RD and RT configured):

```
(virtual-router vrf-foo) BD-12802.15 # show bgp
Enabled                : No                OperStatus           :          Down
RouterId               : 3.3.3.3          AS                   : 200
LocalPref              : 100              MED                  : None
Always-Compare-MED    : Disabled          Aggregation          : Disabled
Route Reflector        : No                RR ClusterId         : 0
IGP Synchronization   : Disabled          New Community Format  : Disabled
Fast Ext Fallover     : Disabled          MPLS LSP as Next-Hop: No
AS Disp Format         : Asplain            Maximum ECMP Paths   : 1
ConfedId               : 0                Outbound rt. filter  : Enabled
Confed Peers           :
Networks               : 2
  ipv4-unicast 10.0.0.0/16 network-policy nwk.pol
  ipv4-multicast 11.0.0.0/16 network-policy nwk.pol
Aggregate Networks    : 2
  ipv4-unicast 21.0.0.0/8 as-match advertise-policy: agg.pol
  ipv4-multicast 22.0.0.0/8 as-set summary-only advertise-policy: agg.pol
```

Route Statistics:

Address family	EBGP	IBGP	Redist.
ipv4-unicast	0	0	0
ipv4-multicast	0	0	0

Redistribute :

Address Family

Route Type	Flags	Priority	Policy
ipv4-unicast			
Direct	EO	2048	None
ipv6-multicast			
Direct	EO	2048	None



Flags: (E) Export Enabled, (L) Export Operationally Off due to Low Memory,

(O) Export Operationally On

Advertise Inactive Routes:

ipv4-unicast : Disabled
 ipv4-multicast : Disabled

Output of **show bgp** for a VRF (PE-CE Protocol, RD and RT “not” configured):

```
BD-12802.5 # show bgp
Enabled                : No                OperStatus           : Down
RouterId              : 3.3.3.3            AS                   : 200
LocalPref             : 100                MED                  : None
Always-Compare-MED   : Disabled           Aggregation          : Disabled
Route Reflector       : No                RR ClusterId         : 0
IGP Synchronization : Disabled           New Community Format  : Disabled
Fast Ext Fallover    : Disabled           MPLS LSP as Next-Hop: No
AS Disp Format        : Asplain            Maximum ECMP Paths   : 1
ConfedId             : 0                Outbound rt. filter  : Enabled
Confed Peers         :
Networks              : 4
  ipv4-unicast 10.0.0.0/16 network-policy nwk.pol
  ipv4-multicast 11.0.0.0/16 network-policy nwk.pol
  ipv6-unicast 2001::/64 network-pol nwk6.pol
  ipv6-multicast 2001::/64 network-pol nwk6.pol
Aggregate Networks   : 4
  ipv4-unicast 21.0.0.0/8 as-match advertise-policy: agg.pol
  ipv4-multicast 22.0.0.0/8 as-set summary-only advertise-policy: agg.pol
  ipv6-unicast 2003::/64 as-match advertise-policy: agg6.pol
  ipv6-multicast 2004::/64 as-set advertise-policy: agg6.pol
```

Route Statistics:

Address family	EBGP	IBGP	Redist.
-----	-----	-----	-----
ipv4-unicast	0	0	0
ipv4-multicast	0	0	0
ipv6-unicast	0	0	0
ipv6-multicast	0	0	0

Redistribute:

```
-----
--
Address Family
  Route Type      Flags      Priority      Policy
-----
--
ipv4-unicast
  Direct          EO          2048          None
ipv6-multicast
  Direct          EO          2048          None
-----
--
```

Flags: (E) Export Enabled, (L) Export Operationally Off due to Low Memory,

(O) Export Operationally On



```

Advertise Inactive Routes:
  ipv4-unicast    : Disabled
  ipv4-multicast  : Disabled
  ipv6-unicast    : Disabled
  ipv6-multicast  : Disabled

```

If BGP is added as a protocol inside a heavy-weight VR, normal BGP peering applies with the addition of vpnv4 address family support:

```

BD-12802.5 # show bgp
Enabled          : No                OperStatus      : Down
RouterId         : 3.3.3.3           AS              : 200
LocalPref        : 100               MED             : None
Always-Compare-MED : Disabled       Aggregation     : Disabled
Route Reflector  : No                RR ClusterId    : 0
IGP Synchronization : Disabled      New Community Format: Disabled
Fast Ext Fallover  : Disabled        MPLS LSP as Next-Hop: No
AS Disp Format    : Asplain           Maximum ECMP Paths : 1
ConfedId         : 0                 Outbound rt. filter : Enabled
Confed Peers     :
Networks         : 4
  ipv4-unicast 10.0.0.0/16 network-policy nwk.pol
  ipv4-multicast 11.0.0.0/16 network-policy nwk.pol
  ipv6-unicast 2001::/64 network-pol nwk6.pol
  ipv6-multicast 2001::/64 network-pol
    nwk6.pol
Aggregate Networks : 4
  ipv4-unicast 21.0.0.0/8 as-match advertise-policy: agg.pol
  ipv4-multicast 22.0.0.0/8 as-set summary-only advertise-policy: agg.pol
  ipv6-unicast 2003::/64 as-match advertise-policy: agg6.pol
  ipv6-multicast 2004::/64 as-set advertise-policy: agg6.pol

```

Route Statistics:

Address family	EBGP	IBGP	Redist.
-----	-----	-----	-----
ipv4-unicast	0	0	0
ipv4-multicast	0	0	0
vpnv4	0	0	0
ipv6-unicast	0	0	0
ipv6-multicast	0	0	0

Redistribute:

ipv4	Admin	Operational	Shutdown	Policy
unicast	Status	Status	Priority	
-----	-----	-----	-----	-----
Direct	Disabled	Down	2048	None
Static	Disabled	Down	2048	None
RIP	Disabled	Down	2048	None
BlackHole	Disabled	Down	2048	None
OSPFIntra	Disabled	Down	2048	None
OSPFInter	Disabled	Down	2048	None
OSPFExt1	Disabled	Down	2048	None
OSPFExt2	Disabled	Down	2048	None
ISISL1	Disabled	Down	2048	None
ISISL2	Disabled	Down	2048	None
ISISL1Ext	Disabled	Down	2048	None
ISISL2Ext	Disabled	Down	2048	None



```

ipv4      Admin      Operational  Shutdown  Policy
multicast Status      Status      Priority

-----
Direct    Disabled  Down        2048      None
Static    Disabled  Down        2048      None
RIP       Disabled  Down        2048      None
BlackHole Disabled  Down        2048      None
OSPFIntra Disabled  Down        2048      None
OSPFInter Disabled  Down        2048      None
OSPFExt1  Disabled  Down        2048      None
OSPFExt2  Disabled  Down        2048      None
ISISL1    Disabled  Down        2048      None
ISISL2    Disabled  Down        2048      None
ISISL1Ext Disabled  Down        2048      None
ISISL2Ext Disabled  Down        2048      None

ipv6      Admin      Operational  Shutdown  Policy
unicast   Status      Status      Priority

-----
Direct    Disabled  Down        2048      None
Static    Disabled  Down        2048      None
Ripng     Disabled  Down        2048      None
OspfV3-intra Disabled  Down        2048      None
OspfV3-inter  Disabled  Down        2048      None
OspfV3-extern1 Disabled  Down        2048      None
OspfV3-extern2 Disabled  Down        2048      None
ISISL1    Disabled  Down        2048      None
ISISL2    Disabled  Down        2048      None
ISISL1Ext Disabled  Down        2048      None
ISISL2Ext Disabled  Down        2048      None

ipv6      Admin      Operational  Shutdown  Policy
multicast Status      Status      Priority

-----
Direct    Disabled  Down        2048      None
Static    Disabled  Down        2048      None
Ripng     Disabled  Down        2048      None
OspfV3-intra Disabled  Down        2048      None
OspfV3-inter  Disabled  Down        2048      None
OspfV3-extern1 Disabled  Down        2048      None
OspfV3-extern2 Disabled  Down        2048      None
ISISL1    Disabled  Down        2048      None
ISISL2    Disabled  Down        2048      None
ISISL1Ext Disabled  Down        2048      None
ISISL2Ext Disabled  Down        2048      None

```

```

Advertise Inactive Routes:
ipv4-unicast : Disabled
ipv4-multicast : Disabled
ipv6-unicast : Disabled
ipv6-multicast : Disabled

```

History

This command was first available in ExtremeXOS 10.1.



This command required a specific license in ExtremeXOS 11.1.

This command was modified in Extreme EXOS 15.3 to reflect its operation in VRs and VRFs.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show bgp memory

```
show bgp memory {detail | memoryType}
```

Description

Displays BGP specific memory usage.

Syntax Description

detail	Displays detail information.
<i>memoryType</i>	Specifies the memory type usage to display.

Default

N/A.

Usage Guidelines

To see the memory types that you can display, enter the show bgp memory command without any attributes.

Example

The following command displays detailed BGP output for a specific memory types:

```
Switch.16.3 # sh bgp memory
BGP Memory Information
-----
Current Memory Utilization Level:      GREEN
-----
Type                AN                AB
-----
Callbacks           1141           17039828
Buffers             19             8456
```



Memory Utilization Statistics:

```

-----
Module Type      Module Id      MemType
Name                                     Size      AN/AB/HN/HB
-----
PCT_NBASE_ROOT  0x0000000000  16777219
MEM_PROCESS_ENTRY
PCT_NBASE_ROOT  0x0000000000  16777230      212      8/1696/8/1696
MEM_NBB_DIAGS_BLOCK
PCT_NBASE_ROOT  0x0000000000  16777233      3212     8/25696/8/25696
MEM_UNFORMATTED
PCT_NBASE_ROOT  0x0000000000  50528257      1508     1/1508/1/1508
0x0003030001
PCT_NBASE_ROOT  0x0000000000  50921473      732      1/732/1/732
0x0003090001
PCT_NBASE_ROOT  0x0000000000  52232193      2004     1/2004/1/2004
0x00031d0001
PCT_NBASE_ROOT  0x0000000000  1090584577    1508     1/1508/1/1508
MEM_QBRM_LOCAL
PCT_NBASE_ROOT  0x0000000000  1090650113    9660     2/19320/2/19320
MEM_QBNM_LOCAL
PCT_NBASE_ROOT  0x0000000000  1107361793    1508     2/3016/2/3016
MEM_QVB_LOCAL
PCT_SCK         0x0001109000  16777220      3076     1/3076/1/3076
MEM_NBB_POOL_CB
PCT_QVB         0x0001104000  1107361803    108      9/972/9/972
MEM_QVB_RV_REM_CB
PCT_QVB         0x0001104000  1107361806    60       6/360/6/360
MEM_QVB_AS_PATH_CB
PCT_QVB         0x0001104000  1107361807    60       4/240/4/240
MEM_QVB_RTM_CB

```

```

Flags : AN - Number of Allocations,          AB - Total
Allocation in Bytes
        : HN - Number of High Water Marks for Allocation,  HB - Total High
Water Mark Allocations in Bytes
t16.3 # sh bgp memory 1107361807  BGP Memory Information

```

```

-----
Current Memory Utilization Level:      GREEN
-----

```

```

-----
Type      AN      AB
-----
Callbacks      1141      17039828
Buffers        19        8456

```

Memory Statistics for MEM_QVB_RTM_CB:

```

-----
MemId      Size      AN      AB
-----
001107361807      516      1      516

```



Flags : AN - Number of Allocations, AB - Total Allocation
in Bytes
: HN - Number of High Water Marks for Allocation, HB - Total High Water
Mark Allocations in Bytes

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

This command is updated to reflect L3 VPN changes in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show bgp neighbor

For IPv4 and IPv6 address families:

```
show bgp {neighbor} remoteaddr {address-family [ipv4-unicast | ipv4-multicast |
ipv6-unicast | ipv6-multicast]} [accepted-routes | received-routes | rejected-
routes | transmitted-routes] {detail} [all | as-path path-expression | community
[no-advertise | no-export | no-export-subconfed | number community_number |
autonomous-system-id :bgp-community] | network [any / netMaskLen |
networkPrefixFilter] {exact}]
```

For the VPNv4 address family:

```
show bgp {neighbor} remoteaddr address-family vpnv4 [accepted-routes | received-
routes | rejected-routes | transmitted-routes] {detail} [all | as-path path-
expression | community [no-advertise | no-export | no-export-subconfed | number
community_number | autonomous-system-id :bgp-community] | rd rd_value network
[any / netMaskLen | networkPrefixFilter] {exact}]
```

Description

Displays information about routes to a specified neighbor.

Syntax Description

<code>remoteaddr</code>	Specifies an IPv4 or IPv6 address that identifies a BGP neighbor.
<code>ipv4-unicast</code>	Specifies IPv4 unicast routes.
<code>ipv4-multicast</code>	Specifies IPv4 multicast routes.



ipv6-unicast	Specifies IPv6 unicast routes.
ipv6-multicast	Specifies IPv6 multicast routes.
vpn4	Specifies VPNv4 routes.
accepted-routes	Specifies that only accepted routes are displayed.
received-routes	Specifies that only received routes are displayed.
rejected-routes	Specifies that only rejected routes are displayed.
transmitted-routes	Specifies that only transmitted routes are displayed.
detail	Specifies to display the information in detailed format.
all	Specifies all routes.
<i>path-expression</i>	Display routes that match the specified AA path expression.
no-advertise	Specifies the no-advertise community attribute.
no-export	Specifies the no-export community attribute.
no-export-subconfed	Specifies the no-export-subconfed community attribute.
<i>community_number</i>	Specifies a community number.
<i>autonomous-system-id</i>	Specifies an autonomous system ID (0-65535).
<i>bgp-community</i>	Specifies the BGP community number.
rd	Specifies the Route Distinguisher (RD) value for the Layer 3 VPN routes for which you want to clear flap statistics.
any	Specifies all routes with a given or larger mask length.
<i>netMaskLen</i>	Specifies a IPv4 or IPv6 subnet mask length (number of bits).
<i>networkPrefixFilter</i>	Specifies an IPv4 or IPv6 address and netmask.
exact	Specifies an exact match with the IP address and subnet mask.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

show bgp neighbor now supports v6 unicast and multicast and vpn4 address families. This command applies to the current VR or VRF context.



Note

If this command displays Bad Source Address, the BGP neighbor IP address is unavailable. Possible causes for this condition include a deleted or unconfigured VLAN or IP address.

The option **network any / netMaskLen** displays all BGP routes whose mask length is equal to or greater than *maskLength*, irrespective of their network address.



The option `network any / netMaskLen exact` displays all BGP routes whose mask length is exactly equal to `maskLength`, irrespective of their network address.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Note



For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default address-family, i.e. IPv4 unicast is assumed and hence no address-family information appears. Similarly an IPv4 peer only supports IPv4 address families and no address-family information appears if an IPv6 address family is specified.

To display Layer 3 VPN information, you must enter this command in the context of on the MPLS-enabled VR; it is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command displays sample output for **show bgp neighbor summary**:

```
Switch.18 # show bgp nei ?
<cr>          Execute the command
detail        display all information available about BGP neighbor
<remoteaddr> BGP neighbor IP address  "11.0.0.2" "3001::1"

BD-12802.1 # show bgp neighbor

Peer          AS    Weight    State          InMsgs OutMsgs(InQ)  Up/
Down
-----
-
Ie-- 11.0.0.2    100    0          OPENSENT       0      9      (0  )
0:8:27:21
Ie-- 3001::1     100    0          ESTABLISEHD    4      3      (0  )
0:8:27:21

Flags:  (d) disabled, (e) enabled, (E) external peer, (I) internal peer
        (m) EBGP multihop, (r) route reflector client

BGP Peer Statistics
Total Peers      : 2
EBGP Peers       : 0
IBGP Peers       : 2
RR Client        : 0
EBGP Multihop    : 0
Enabled          : 2
Disabled         : 0
```



The following example displays show output for an IPv4 peer:

```
switch 19 # sh bgp neighbor 11.0.0.5 det

Peer Description      :
IBGP Peer            : 1.0.0.3          AS              : 1000
Enabled              : Yes              OperStatus       : Up
Weight              : 1                Shutdown-Priority : 1024
ConnectRetry        : 120              MinAsOrig        : 5
HoldTimeCfg         : 180              KeepaliveCfg     : 60
Source Interface    : 1.0.0.2          RRClient         : No
EBGP-Multihop       : No               Remove Private AS : No
Capabilities Config : ipv4-unicast,ipv4-multicast,4-Byte-As,vpnv4,route-
refresh (old & new)
Policy for NLRI Type ipv4-unicast
  In Policy          : None
  Out Policy         : None
  NextHopSelf       : Disabled          Send Communities  : No
  Soft Input Recfg  : Disabled          Allow Looped AS-Path: No
Policy for NLRI Type ipv4-multicast
  In Policy          : None
  Out Policy         : None
  NextHopSelf       : Disabled          Send Communities  : No
  Soft Input Recfg  : Disabled          Allow Looped AS-Path: No
Policy for NLRI Type vpnv4
  In Policy          : None
  Out Policy         : None
  NextHopSelf       : Disabled          Send Communities  : Standard,
Extended
  Soft Input Recfg  : Disabled          Allow Looped AS-Path: No
Policy for NLRI Type ipv6-unicast
  In Policy          : None
  Out Policy         : None
  NextHopSelf       : Disabled          Send Communities  : No
  Soft Input Recfg  : Disabled          Allow Looped AS-Path: No
Policy for NLRI Type ipv6-multicast
  In Policy          : None
  Out Policy         : None
  NextHopSelf       : Disabled          Send Communities  : No
  Soft Input Recfg  : Disabled          Allow Looped AS-Path: No
State                : ESTABLISHED
FSM Up since         : Thu Aug  8 12:09:15 2013
  (Duration:0:0:03:28)
Remote Addr          : 1.0.0.3          Local Addr        : 1.0.0.2
Remote Port          : 36879           Local Port         : 179
Remote RouterId      : 1.0.0.3          Local RouterId    : 1.0.0.2
HoldTimeNegotiated  : 180              KeepAliveNegotiated : 60
FsmTransitions       : 1
InUpdateElapsedTime : 00:00:00:05         InMsgElapsedTime  : 0:0:00:05
InUpdates            : 5                OutUpdates (in TxQ) : 5 (0)
InTotalMsgs          : 11              OutTotalMsgs      : 10
InRouteRefreshes    : 0                OutRouteRefreshes : 0
Route Statistics for NLRI Type ipv4-unicast
  Received           : 0                Accepted           : 0
  Rejected           : 0                Active             : 0
  Suppressed         : 0
Route Statistics for NLRI Type ipv4-multicast
  Received           : 0                Accepted           : 0
```



```

    Rejected          : 0                Active          : 0
    Suppressed        : 0
Route Statistics for NLRI Type vpnv4
    Received          : 4                Accepted         : 4
    Rejected          : 0                Active           : 0
    Suppressed        : 0
Capabilities Tx      : ipv4-unicast,ipv4-multicast,4-Byte-As,vpnv4,route-
refresh (old & new),graceful-restart
Capabilities Rx      : ipv4-unicast,ipv4-multicast,4-Byte-As,vpnv4,route-
refresh (old & new),graceful-restart
NLRI for the session: ipv4-unicast,ipv4-multicast,vpnv4
Last State           : ESTABLISHED      Last Event        : RX_KEEP
LastError            : 'None'

Restart Capability   : Aware-Only
Restart Bit Set     : No                Restart Time       : 120
Restart Status      : OFF

BGP Peer Statistics
    Total Peers      : 1
    EBGPeers         : 0                IBGP Peers        : 1
    RR Client        : 0                EBGPeers Multihop : 0
    Enabled           : 1                Disabled           : 0

```

The following example displays output for an IPv6 peer:

```

switch 19 # show bgp neighbor 3001::1 det
EBGP Peer          : 3001::1          AS                : 5
Enabled            : Yes              OperStatus         : Up
Weight             : 1                Shutdown-Priority  : 1024
ConnectRetry       : 120              MinAsOrig          : 15
HoldTimeCfg        : 180              KeepaliveCfg       : 60
Source Interface   : Not configured   RRClient           : No
EBGP-MultiHop     : No                Remove Private AS  : No
Capabilities Config : ipv6-unicast, ipv6-multicast, 4-Byte-As, route-
refresh
Policy for NLRI Type ipv6-unicast
    In Policy        : None
    Out Policy        : None
    NextHopSelf      : Disabled        Send Communities   : No
    Soft Input Recfg : Disabled        Allow Looped AS-Path: No
    RFD HalfLife     : 0m              RFD Reuse          : 0
    RFD Suppress     : 0                RFD Max-Suppress   : 0m
Policy for NLRI Type ipv6-multicast
    In Policy        : None
    Out Policy        : None
    NextHopSelf      : Disabled        Send Communities   : No
    Soft Input Recfg : Disabled        Allow Looped AS-Path: No
    RFD HalfLife     : 0m              RFD Reuse          : 0
    RFD Suppress     : 0                RFD Max-Suppress   : 0m
State              : ESTABLISHED
FSM Up since       : Mon Apr 19 21:20:02 2010 (Duration: 0:0:00:23)
Remote Addr        : 3001::1          Local Addr         : 3001::6
Remote Port        : 56539            Local Port         : 179
Remote RouterId    : 1.0.0.5          Local RouterId     : 1.0.0.6
HoldTimeNegotiated : 180              KeepAliveNegotiated : 60
FsmTransitions     : 5

```



```

InUpdateElapsedTime : 00:00:00:23      InMsgElapsedTime      : 0:0:00:23
InUpdates           : 1                 OutUpdates (in TxQ)  : 4 (0)
InTotalMsgs        : 1                 OutTotalMsgs        : 4
InRouteRefreshes   : 0                 OutRouteRefreshes   : 0
Route Statistics for NLRI Type ipv6-unicast
    Received        : 6
    Accepted        : 6
    Rejected        : 0
    Active          : 6
    Suppressed      : 0
Route Statistics for NLRI Type ipv6-multicast
    Received        : 0
    Accepted        : 0
    Rejected        : 0
    Active          : 0
    Suppressed      : 0
Capabilities Tx      : ipv6-unicast, ipv6-multicast, 4-Byte-AS, route-refresh
(old &
new), vpvnv4
Capabilities Rx      : ipv6-unicast, ipv6-multicast
NLRI for the session: ipv6-unicast, ipv6-multicast
Error                : 'Hold Timer Expired'                Tx: 3      Rx: 0
Last State           : ESTABLISHED                          Last Event      : RX_UPDATE
LastError            : 'Hold Timer Expired' (TX) on: Mon Apr 19 20:50:26
2010
BGP Peer Statistics
    Total Peers      : 1
    EBGP Peers       : 1                 IBGP Peers      : 0
    RR Client        : 0                 EBGP Multihop   : 0
    Enabled          : 1
    Disabled         : 0

```

The following example displays show output for transmitted routes:

```

switch.19 #show bgp 11.0.0.2 transmitted-routes all
Advertised Routes:
Destination      LPref Weight MED      Peer      Next-Hop      AS-
Path
-----
-----
>? 1.1.1.1/32    100
100
>? 11.0.0.0/24  100
100
>? 101.0.0.0/24 100
100
>? 103.0.0.0/24 100
100
>? 103.0.0.1/32 100
100
Flags: (*) Preferred BGP route, (>) Active, (d) Suppressed, (h) History

(s) Stale, (m) Multipath, (u) Unfeasible
Origin: (?) Incomplete, (e) EGP, (i) IGP
BGP Route Statistics
    Advertised Routes : 5

```



The following example displays show output for rejected routes:

```
switch.19 # show bgp 11.0.0.2 rejected-routes all
Rejected Routes:
Destination      LPref Weight MED      Peer          Next-Hop      AS-
Path
-----
-----
u ? 1.1.1.1/32   100
100

Flags: (*) Preferred BGP route, (>) Active, (d) Suppressed, (h) History

      (s) Stale, (m) Multipath, (u) Unfeasible
Origin: (?) Incomplete, (e) EGP, (i) IGP
BGP Route Statistics
  Total Rxed Routes : 5
  Rejected Routes   : 1
  Unfeasible Routes : 1
```

The following example displays show output for accepted routes:

```
switch.21 # show bgp 11.0.0.2 accepted-routes all
Rejected Routes:
Destination      LPref Weight MED      Peer          Next-Hop      AS-Path
-----
-----
>? 11.0.0.0/24   100
100
>? 101.0.0.0/24  100
100
>? 103.0.0.0/24  100
100
>? 103.0.0.1/32  100
100

Flags: (*) Preferred BGP route, (>) Active, (d) Suppressed, (h) History

      (s) Stale, (m) Multipath, (u) Unfeasible
Origin: (?) Incomplete, (e) EGP, (i) IGP
BGP Route Statistics
  Total Rxed Routes : 5
  Feasible Routes   : 4
  Active Routes     : 4
```

History

This command was first available in ExtremeXOS 10.1.

The any / *netMaskLen* options were added in ExtremeXOS 11.0.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.



Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show bgp neighbor [flap-statistics | suppressed-routes]

For IPv4 and IPv6 address families:

```
show bgp {neighbor} remoteaddr {address-family [ipv4-unicast | ipv4-multicast |
ipv6-unicast | ipv6-multicast]} [flap-statistics | suppressed-routes] {detail}
[all | as-path path-expression | community [no-advertise | no-export | no-export-
subconfed | number community-number | autonomous-system-id : bgp-community] |
network [any / netMaskLen | networkPrefixFilter] {exact}]
```

For the VPNv4 address family:

```
show bgp {neighbor} remoteaddr address-family vpnv4 [flap-statistics |
suppressed-routes] {detail} [all | as-path path-expression | community [no-
advertise | no-export | no-export-subconfed | number community-number |
autonomous-system-id : bgp-community] | rd rd_value network [any / netMaskLen |
networkPrefixFilter] {exact}]
```

Description

Displays flap statistics or suppressed-route information about a specified neighbor.

Syntax Description

<code>remoteaddr</code>	Specifies an IPv4 or IPv6 address that identifies a BGP neighbor.
<code>ipv4-unicast</code>	Specifies the IPv4 unicast address family.
<code>ipv4-multicast</code>	Specifies an IPv4 multicast address family.
<code>ipv6-unicast</code>	Specifies the IPv6 unicast address family.
<code>ipv6-multicast</code>	Specifies an IPv6 multicast address family.
<code>vpnv4</code>	Specifies the VPNv4 address family for Layer 3 VPN support.
<code>flap-statistics</code>	Specifies that only flap-statistics should be displayed (for route flap dampening enabled routes).
<code>suppressed-routes</code>	Specifies that only suppressed routes should be displayed (for route flap dampening enabled routes).
<code>detail</code>	Specifies to display the information in detailed format.
<code>all</code>	Specifies all routes.
<code>path-expression</code>	Display routes that match the specified AA path expression.



no-advertise	Specifies the no-advertise community attribute.
no-export	Specifies the no-export community attribute.
no-export-subconfed	Specifies the no-export-subconfed community attribute.
<i>community_number</i>	Specifies a community number.
<i>autonomous-system-id</i>	Specifies an autonomous system ID (0-65535).
<i>bgp-community</i>	Specifies the BGP community number.
rd	Specifies the Route Distinguisher (RD) value for the Layer 3 VPN routes for which you want to clear flap statistics.
any	Specifies all routes with a given or larger mask length.
<i>netMaskLen</i>	Specifies a IPv4 or IPv6 subnet mask length (number of bits).
<i>networkPrefixFilter</i>	Specifies an IPv4 or IPv6 address and netmask.
exact	Specifies an exact match with the IP address and subnet mask.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

This command applies to the current VR or VRF context.



Note

If this command displays Bad Source Address, the BGP neighbor IP address is unavailable. Possible causes for this condition include a deleted or unconfigured VLAN or IP address.

The option `network any / netMaskLen` displays all BGP routes whose mask length is equal to or greater than *maskLength*, irrespective of their network address.

The option `network any / netMaskLen exact` displays all BGP routes whose mask length is exactly equal to *maskLength*, irrespective of their network address.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default, the IPv4 unicast address family, applies and no address-family information appears. Similarly an IPv4 peer only supports IPv4 address families and no address-family information appears if an IPv6 address family is specified.



To display Layer 3 VPN information, you must enter this command in the context of on the MPLS-enabled VR; it is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command displays flap statistics for the specified IPv4 neighbor:

```
* Switch.18 # show bgp neighbor 10.0.0.0 flap-statistics
BGP Routes Flap Statistics
Destination          NextHop          Penalty Flaps Duration Reuse
AS-Path
-----
* ?100:1:100.0.0.0/8  11.0.0.2        100
100
Flags: (*) Preferred BGP route, (>) Active, (d) Suppressed, (h) History
(s) Stale, (m) Multipath, (u) Unfeasible
Origin: (?) Incomplete, (e) EGP, (i) IGP
Total Number of Flapped Routes: 1
```

The following command displays flap statistics for the specified IPv6 neighbor:

```
* Switch.21 # show bgp neighbor 2001::64:: address-family ipv6-unicast flap-
statistics
BGP Routes Flap Statistics
Destination          NextHop          Penalty Flaps Duration Reuse
AS-Path
-----
* ?2001::/64        3001::1         100
100
Flags: (*) Preferred BGP route, (>) Active, (d) Suppressed, (h) History
(s) Stale, (m) Multipath, (u) Unfeasible
Origin: (?) Incomplete, (e) EGP, (i) IGP
Total Number of Flapped Routes: 1
```

History

This command was first available in ExtremeXOS 10.1.

The any / *netMaskLen* options were added in ExtremeXOS 11.0.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show bgp peer-group

```
show bgp peer-group {detail | peer-group-name {detail}}
```

Description

Displays the peer groups configured in the system.

Syntax Description

detail	Specifies to display the information in detailed format.
<i>peer-group-name</i>	Specifies a peer group.

Default

N/A.

Usage Guidelines

If the detail keyword is specified then the parameters of the neighbors in the peer group, which are different from the ones that are configured in the peer group, are displayed.

If no peer group name is specified, all the peer group information is displayed.

This command applies to the current VR or VRF context.

Example

The following command displays information for the outer peer group:

```
* (debug) Summit-PC.19 # show bgp peer-group "outer"
Peer Group           : outer
Enabled              : No                AS                : 65551
Router Enabled       : Yes                Weight             : 1
ConnectRetry        : 120                 MinAsOrig          : 15
HoldTimeCfg         : 180                 KeepaliveCfg       : 60
Source Interface     : Not configured     RRClient           : No
Remove Private AS   : No                 Router-Alert       : Disabled
Capabilities Config  : ipv4-unicast ipv4-multicast route-refresh 4-Byte-AS
Policy for NLRI Type ipv4-unicast
In Policy            : None
Out Policy           : None
NextHopSelf         : Disabled            Send Communities   : No
```



```

Soft Input Recfg : Disabled          Allow Looped AS-Path: No
Policy for NLRI Type ipv4-multicast
In Policy       : None
Out Policy      : None
NextHopSelf    : Disabled          Send Communities   : No
Soft Input Recfg : Disabled          Allow Looped AS-Path: No
Peers          : 11.11.11.11
BGP Peer Group Statistics
Total Peer Groups : 1
Enabled          : 0
Disabled        : 1

```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show bgp routes

For IPv4 and IPv6 address families:

```

show bgp routes {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast |
ipv6-multicast]} {detail} [all | as-path path-expression | community [no-
advertise | no-export | no-export-subconfed | number community_number |
autonomous-system-id : bgp-community] | network [any / netMaskLen |
networkPrefixFilter] {exact}]

```

For the VPNv4 address family:

```

show bgp routes address-family vpnv4 {detail} [all | as-path path-expression |
community [no-advertise | no-export | no-export-subconfed | number community-
number | autonomous-system-id : bgp-community] | rd rd network [any / netMaskLen
| networkPrefixFilter] {exact}]

```

Description

Displays the BGP route information base (RIB).

Syntax Description

ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.



ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support.
all	Specifies all routes.
<i>path-expression</i>	Display routes that match the specified AA path expression.
no-advertise	Specifies the no-advertise community attribute.
no-export	Specifies the no-export community attribute.
no-export-subconfed	Specifies the no-export-subconfed community attribute.
<i>community_number</i>	Specifies a community number.
<i>autonomous-system-id</i>	Specifies an autonomous system ID (0-65535).
<i>bgp-community</i>	Specifies the BGP community number.
rd	Specifies the Route Distinguisher (RD) value for the Layer 3 VPN routes for which you want to clear flap statistics.
any	Specifies all routes with a given or larger mask length.
<i>netMaskLen</i>	Specifies a IPv4 or IPv6 subnet mask length (number of bits).
<i>networkPrefixFilter</i>	Specifies an IPv4 or IPv6 address and netmask.
exact	Specifies an exact match with the IP address and subnet mask.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

The option `network any / netMaskLen` displays all BGP routes whose mask length is equal to or greater than *maskLength*, irrespective of their network address.

The option `network any / netMaskLen exact` displays all BGP routes whose mask length is exactly equal to *maskLength*, irrespective of their network address.

To display Layer 3 VPN information, you must enter this command in the context of on the MPLS-enabled VR; it is not supported for BGP neighbors on the CE (VRF) side of the PE router.

You can only execute the `show for vpn4` address family in a VR context. If you execute this command in a VRF context, the “Cannot execute command in VRF context” error is displayed.

This command applies to the current VR or VRF context.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Example

The following command displays detailed information about all BGP routes:

```
* Switch.5 # show bgp routes all
Received Routes:
Destination          LPref Weight MED      Peer          Next-Hop      AS-
Path
-----
-----
*>? 1.1.1.1/32       100   0          11.0.0.1     11.0.0.1
100
* ? 11.0.0.0/24     100   0          11.0.0.1     11.0.0.1
100
*>? 101.0.0.0/24    100   0          11.0.0.1     11.0.0.1
100
u ? 103.0.0.0/24    100   0          11.0.0.1     11.0.0.1
100
Flags: (*) Preferred BGP route, (>) Active, (d) Suppressed, (h) History

        (s) Stale, (m) Multipath, (u) Unfeasible

Origin: (?) Incomplete, (e) EGP, (i) IGP

BGP Route Statistics
  Total Rxed Routes : 4
  Feasible Routes   : 3
  Active Routes     : 2
  Rejected Routes   : 0
  Unfeasible Routes : 1
Route Statistics on Session Type
  Routes from Int Peer: 4
  Routes from Ext Peer: 0
```

The following example displays a detailed show output:

```
Route: 11.0.0.0/24, Peer 11.0.0.1, Unfeasible
Origin Incomplete, Next-Hop 11.0.0.1, LPref 100, MED 0
Weight 0, RR Orig ID 0.0.0.0
AS-Path: 100
DampInfo: Penalty 0 Flapped 0 times in 00:10:47
Flags: (*) Preferred BGP route, (>) Active, (d) Suppressed, (h) History

        (s) Stale, (m) Multipath, (u) Unfeasible

Origin: (?) Incomplete, (e) EGP, (i) IGP
BGP Route Statistics
  Total Rxed Routes : 1
  Feasible Routes   : 0
  Active Routes     : 0
  Rejected Routes   : 5
  Unfeasible Routes : 1
Route Statistics on Session Type
  Routes from Int Peer : 5
  Routes from Ext Peer : 0
```



The following command displays BGP information for the IPv6 address family:

```
Switch.21 # show bgp routes address-family ipv6-unicast all
Received Routes:
  Destination                               LPref Weight MED
  Peer                                       Next-Hop                               AS-Path
-----
-----
*>? 2001::/64                               100           0
120
    3000::1                                 3001::1      100,      200

Flags: (*) Preferred BGP route, (>) Active, (d) Suppressed, (h) History
       (s) Stale, (m) Multipath, (u) Unfeasible

Origin: (?)
        Incomplete, (e) EGP, (i) IGP

BGP Route Statistics
  Total Rxed Routes : 1
  Feasible Routes   : 1
  Active Routes     : 1
  Rejected Routes   : 0
  Unfeasible Routes : 0

Route Statistics on Session Type
  Routes from Int Peer: 1
  Routes from Ext Peer: 0
```

The following example displays detailed show output for the IPv6 address family:

```
switch.21 # show bgp routes address-family ipv6-unicast all
Route: 2001::/64, Peer 3000::1,
Unfeasible, Origin Incomplete,
Next-Hop 3001::1,
LPref 100, MED 0,
Weight 0, RR Orig ID 0.0.0.0
AS-Path: 100
DampInfo: Penalty 0 Flapped 0 times in 00:10:47

Route: 2002::/64, Peer 3000::1,
Active, Origin Incomplete,
Next-Hop 3001::1,
LPref 100, MED 0,
Weight 0, RR Orig ID 0.0.0.0
AS-Path: 100
DampInfo: Penalty 0 Flapped 0 times in 00:10:47

BGP Route Statistics
  Total Rxed Routes : 2
  Feasible Routes   : 1
  Active Routes     : 1
  Rejected Routes   : 0
  Unfeasible Routes : 1

Route Statistics on Session Type
```



```

Routes from Int Peer: 2
Routes from Ext Peer: 0

```

The following examples display detailed show output for the IPv4 address family:

```

switch.21 # show bgp routes address-family vpnv4 all
Received Routes:
Destination                                     LPref Weight  MED

Peer                Next-Hop                AS-Path
-----
*?> 100:1:10.0.0.0/8                100    0    120

    11.0.0.2                11.0.0.2
    100, 200
Flags: (*) Preferred BGP route, (>) Active, (d) Suppressed, (h) History

        (s) Stale, (m) Multipath, (u) Unfeasible

Origin: (?) Incomplete, (e) EGP, (i) IGP

BGP Route Statistics
Total Rxed Routes : 1
Feasible Routes   : 1
Active Routes     : 1
Rejected Routes   : 0
Unfeasible Routes : 0
Route Statistics on Session Type
Routes from Int Peer: 1
Routes from Ext Peer: 0
-----
-----

switch.21 # show bgp routes address-family ipv6-unicast all
Route: 100:1:10.0.0.0/8, Peer 11.0.0.2,
Unfeasible, Origin Incomplete,
Next-Hop 11.0.0.2,
LPref 100, MED 0,
Weight 0, RR Orig ID 0.0.0.0
AS-Path: 100
DampInfo: Penalty 0 Flapped 0 times in 00:10:47

BGP Route Statistics
Total Rxed Routes : 1
Feasible Routes   : 0
Active Routes     : 0
Rejected Routes   : 0
Unfeasible Routes : 1
Route Statistics on Session Type
Routes from Int Peer: 1
Routes from Ext Peer: 0

```

History

This command was first available in ExtremeXOS 10.1.



The any / *netMaskLen* options were added in ExtremeXOS 11.0.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show bgp routes summary

```
show bgp routes {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast |
ipv6-multicast | vpnv4]} summary
```

Description

Displays a summary of the BGP route information base (RIB).

Syntax Description

ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpnv4	Specifies the VPNv4 address family for Layer 3 VPN support.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

This command applies to the current VR or VRF context.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

To display Layer 3 VPN information, you must enter this command in the context of on the MPLS-enabled VR; it is not supported for BGP neighbors on the CE (VRF) side of the PE router.



When the `show bgp routes summary` command is issued with **address-family vpnv4**, the command will impact the behavior of PE to PE neighbor sessions and display/clear the VPN-IPv4 RIB of BGP.

Example

The following command displays a summary of the BGP route information base (RIB) for IPv4 multicast:

```
show bgp routes address-family ipv4-multicast summary
```

The following example displays VPN routes with RD 100:1:

```
virtual-router corpl_vrf
  show bgp routes address-family vpnv4 rd 100:1
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



42 L3 VPN Commands

disable snmp traps l3vpn
enable snmp traps l3vpn

Layer 3 Virtual Private Networks (L3 VPN) is a specific implementation of a Provider Provisioned VPN (PPVPN). L3VPN is a way to create a tunnel between customer sites through a Provider Backbone, and that tunnel is established and maintained by the Service Provider.

This chapter describes commands for configuring L3 VPN.

disable snmp traps l3vpn

```
disable snmp traps l3vpn {vr vr_name
```

Description

Use this command to turn off SNMP trap support for L3 VPN.

Syntax Description

<i>vr_name</i>	Specifies the name of the parent VR where this RFC 4382 scalar is applied. If VR <i>name</i> is not provided, then this command is applied to the VR in the current context.
----------------	--

Default

Enabled.

Usage Guidelines

Use this command to disable L3VPN SNMP traps.

Example

The following command disables L3 VPN SNMP traps support on the switch:

```
disable snmp traps l3vpn vr vr-default
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

enable snmp traps l3vpn

```
enable snmp traps l3vpn {vr vr_name}
```

Description

Use this command to turn on SNMP trap support for L3 VPN.

Syntax Description

<i>vr_name</i>	Specifies the name of the parent VR where this RFC 4382 scalar is applied. If VR name is not provided, then this command is applied to the VR in the current context.
----------------	---

Default

Enabled.

Usage Guidelines

Use this command to enable generation of L3VPN SNMP traps—mplsL3VpnVrfUp and mplsL3VpnVrfDown. These trap notifications are sent under the following conditions:

- mplsL3VpnVrfUp – first IP VLAN becomes active and administrative state is enabled.
- mplsL3VpnVrfDown – last active IP VLAN becomes inactive OR administrative state is disabled.

Example

The following command enables L3 VPN SNMP traps support on the switch:

```
enable snmp traps l3vpn vr vr-default
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.



43 OpenFlow Commands

```
clear openflow counters
configure openflow controller
debug openflow
debug openflow show flows
disable openflow
disable openflow vlan
enable openflow
enable openflow vlan
show openflow
show openflow controller
show openflow flows
show openflow vlan
unconfigure openflow controller
```

OpenFlow provides a standardized, flexible operational tool for building virtualized networks. OpenFlow enables switching control plane features to be implemented and evolved in a hardware independent manner.

This chapter describes commands for configuring OpenFlow on the switch. There are additional EXOS commands that you can use to configure OpenFlow. Please refer to the *ExtremeXOS Concepts Guide*, 15.3 for more configuration details.

clear openflow counters

```
clear openflow counters {flow | controller {primary | secondary}}
```

Description

Globally clears the flow error count, packets sent and received. **controller** {**primary** | **secondary**} clears the connection counters of the primary, secondary, or both controllers.

Syntax Description

This command has no keywords or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command clears Openflow counters on the switch:

```
clear openflow counters
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure openflow controller

```
configure openflow controller [primary | secondary] [in-band [port port-number | discovery] | out-of-band [active [ipaddress ipaddress | hostname host_name] {port} | passive port]] {tls} {vr vr_name} {rate-limit rate_limit {burst-size burst-size}}
```

Description

Configures the OpenFlow controller(s) that the switch will communicate with.

Syntax Description

primary	Specifies the primary openflow controller.
secondary	Specifies the secondary openflow controller.
port	Specifies the port number for in-band mode .
<i>port-number</i>	Specifies the physical port number.
out-of-band	Specifies the out-of-band connection to the controller.
active	Specifies that you actively connect to the controller .
ipaddress	Specifies that you use an IP address for active out-of-band mode; it might be followed by tcp port.
<i>ipaddress</i>	Configures an IP address, for example: 192.168.32.25.
hostname	Specifies the hostname.
<i>port</i>	Specifies the TCP port. for example: 6643.



passive	Configures the passive mode for out-of-band; you must specify a tcp port.
tls	Specifies that you use the Transport Layer Security (TLS) option.
vr	Specifies that you use the virtual router option.
<i>vr_name</i>	Specifies the name of the virtual router.
rate-limit	Specifies the rate-limit Packet-In packets sent to the controller.
<i>rate_limit</i>	Specifies packets per second. Default is 1000, the range is 100-2147483647.
burst-size	Specifies that you use the burst-size with rate-limit.
<i>burst-size</i>	Specifies the burst size in bytes; the range is 1500-65536.

Default

If *burst-size* is not specified, the default is 1500 bytes. If *rate_limit* is not specified, the default value is 1000.

Usage Guidelines

Use this command to configure the OpenFlow controller(s) that the switch will communicate with.

If only a secondary controller is configured, it will be treated as a primary controller until a primary controller is configured.

OpenFlow attempts to communicate with the primary controller until connectivity fails, in which case it automatically fails over to the secondary controller, if configured. 'out-of-band' control enables controller(s) to connect to the switch using a non-OpenFlow vlan. 'vr' specifies the virtual router used by the switch to communicate with the controller(s).

The **rate-limit** *rate* and **burst-size** *burstSize* options limit the rate and burst-size of messages sent from the switch to the controller.

Example

The following example illustrates how to use the `configure openflow controller` command:

```
configure openflow controller primary out-of-band active ipaddress
10.1.1.1 6633 vr vr-mgmt
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Openflow feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



debug openflow

Short reference description.

```
debug openflow {on | off | on} {{{verbosity verbosity} {output output_file}} | {{output
output_file} {verbosity verbosity}}}}
```

Description

Captures OpenFlow protocol packets for analysis.

Syntax Description

on	Turn debug on.
off	Turn debug off.
verbosity	Verbosity of output.
output	Output packet capture information to a file.
<i>output_file</i>	Output filename.
<i>verbosity</i>	0 (default) is the least detailed, 5 is the most detailed.

Default

0 is the default value for *verbosity*.

Usage Guidelines

Use this command to decode OpenFlow protocol packets sent to and from the connected OpenFlow controllers for analysis.

Example

The following example turns debugging off:

```
debug openflow off
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

debug openflow show flows



```
debug openflow show [tables | controller stats | flows [vendor-table | exos-tree]
| flow flow_no]
```

Description

Displays the flows currently configured by the active controller. The command is used to show the contents of the reference code flow table datastructure, or an EXOS-specific flow table datastructure.

Syntax Description

tables	Displays internal VLAN tables.
controller stats	Displays controller connection counters.
vendor-table	Displays the flows in vendor datastructure .
exos-tree	Displays flows in binary tree maintained by EXOS .
flow flow_no	Displays the match conditions and actions of <i>flow_no</i>

Default

None.

Usage Guidelines

Used to view internal tables, counters, and datastructures for debugging purposes.

Example

The following example displays openflow flow statistics:

```
debug openflow show flows exos-tree
=====
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable openflow

```
disable openflow
```



Description

Globally disables the Openflow application on the switch.

Syntax Description

disable	Disables openflow.
----------------	--------------------

Default

The default is disabled.

Usage Guidelines

None.

Example

The following command disables Openflow on the switch:

```
disable openflow
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Openflow feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable openflow vlan

```
disable openflow {vlan} vlan_name
```

Description

Disables OpenFlow on a specific VLAN.

Syntax Description

vlan	Specifies that OpenFlow is disabled on a VLAN.
<i>vlan_name</i>	Specifies the name of the VLAN to disable.



Default

None.

Usage Guidelines

You must specify a VLAN name to disable.

Example

The following command disables OpenFlow on VLAN 1:

```
disable openflow vlan 1
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable openflow

enable openflow

Description

Globally enables the Openflow application on the switch.

Syntax Description

enable	Enables openflow.
---------------	-------------------

Default

The default is disabled.

Usage Guidelines

You do not have to issue this command before you issue other OpenFlow commands.



Example

The following command enables Openflow on the switch:

```
enable openflow
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable openflow vlan

```
enable openflow {vlan} vlan_name
```

Description

Enables OpenFlow on specific VLANs .

Syntax Description

vlan	Specifies the VLAN on which to enable Openflow.
<i>vlan_name</i>	Specifies the VLAN name.

Default

No VLANs are enabled for OpenFlow by default.

Usage Guidelines

Only one VLAN can be enabled for OpenFlow on the switch.

Example

The following command specifies the ports to enable Openflow on the switch:

```
enable openflow
```



History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the OpenFlow feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show openflow

show openflow

Description

Shows whether OpenFlow is enabled or disabled globally on the switch.

Syntax Description

This command has no keywords or variables.

Default

None.

Usage Guidelines

None.

Example

The following example displays the current configuration for the primary controller:

```
show openflow
openflow is enabled!
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



show openflow controller

```
show openflow controller {primary | secondary}
```

Description

Shows the OpenFlow controller configuration and status on the switch.

Syntax Description

primary	Specifies the primary openflow controller.
secondary	Specifies the secondary openflow controller.

Default

None.

Usage Guidelines

None.

Example

The following example displays the current configuration for the primary controller:

```
show openflow controller
Controller      : Primary
  Configured    : Yes
  Datapath ID   : abcdef0123456789
  Target        : tcp:10.1.1.1:6633
  VR            : VR-Default
  Mode          : out-of-band Active
  Status        : ACTIVE          SSL           : Disabled
  Probe(secs)  : 30              Uptime(secs) : 130
  Rate Limit   : 1000           Burst Size   : 250
  Packets Sent : 9              Packets Received : 8
Controller     : Secondary
  Configured    : No
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



show openflow flows

```
show openflow flows {flow_name}
```

Description

Display the match conditions and actions of installed OpenFlow flows.

Syntax Description

<i>flow_name</i>	Specifies the flow name.
------------------	--------------------------

Default

None.

Usage Guidelines

Use this command to determine the number and details of OpenFlow flows installed on the switch by an OpenFlow controller.

Example

The following example displays the current OpenFlow flows:

```
show openflow flows
Total number of flows: 1
Flow name Type Duration (secs) Prio Packets
-----
of_12345 FDB 9223372036854775807 65535 18446744073709551615
Match:  Input port: 2
      Src MAC: 00:11:22:33:44:55
      Dst MAC: 00:11:22:33:44:55
      VLAN ID: 1234
      VLAN priority: 255
      Ethernet type: 0x8888
      IP TOS: 0x1234
      IP protocol: 0x1234
      IP src address: 255.255.255.255
      IP dst address: 255.255.255.255
      Transport src port: 65535
      Transport dst port: 65535
      Actions: Output port: 3, Output port: 4, Drop
```

History

This command was first available in ExtremeXOS 15.3.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show openflow vlan

```
show openflow {vlan} vlan_name
```

Description

Shows the OpenFlow configuration state for the specified ports.

Syntax Description

 vlan 	Specifies that show output is restricted to a specified VLAN.
<i> vlan_name </i>	Specifies a named VLAN.

Default

None.

Usage Guidelines

If the VLAN name is specified, the output is restricted to that VLAN.

Example

The following command displays show output for all configured OpenFlow ports:

```
show openflow
OpenFlow is enabled.
  Controller      : Primary
  Status          : ACTIVE
  Datapath ID    : 000011112222
  VR              : VR_Default
  Mode           : out-of-band Active
  Target         : tcp:10.1.1.2:6633
  Uptime(secs)  : 200s
Secondary controller: Not configured.
```

VLAN	VID	Ports	Flows Active	Error
of1	20	5	999999	999999

```
Total number of VLAN(s): 1

show openflow vlan of1
Primary controller: tcp:10.1.1.2:6633, out-of-band, Active, Uptime: 200s
```



Secondary controller: Disabled.

VLAN	VID	Flows	
		Ports	Active Error
of1	20	5	999999 999999
Total number of VLAN(s): 1			

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure openflow controller

```
unconfigure openflow controller [primary | secondary]
```

Description

Unconfigures the OpenFlow controller(s).

Syntax Description

primary	Specifies the primary openflow controller.
secondary	Specifies the secondary openflow controller.

Default

N/A.

Usage Guidelines

None.

Example

The following example unconfigures the primary controller:

```
unconfigure openflow controller primary
```



History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Openflow feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



44 IP Multicast Commands

```
clear igmp group
clear igmp snooping
clear pim cache
clear pim snooping
configure forwarding ipmc compression
configure forwarding ipmc lookup-key
configure igmp
configure igmp router-alert receive-required
configure igmp router-alert transmit
configure igmp snooping filters
configure igmp snooping flood-list
configure igmp snooping leave-timeout
configure igmp snooping timer
configure igmp snooping vlan ports add dynamic group
configure igmp snooping vlan ports add static group
configure igmp snooping vlan ports add static router
configure igmp snooping vlan ports delete static group
configure igmp snooping vlan ports delete static router
configure igmp snooping vlan ports filter
configure igmp snooping vlan ports set join-limit
configure igmp ssm-map add
configure igmp ssm-map delete
configure ipmcforwarding
configure ipmroute add
configure ipmroute delete
configure iproute add (Multicast)
configure iproute delete
configure mcast ipv4 cache timeout
configure mvr add receiver
configure mvr add vlan
configure mvr delete receiver
configure mvr delete vlan
configure mvr mvr-address
configure mvr static group
configure pim add vlan
configure pim border
configure pim cbsr
```

```
configure pim crp static
configure pim crp timer
configure pim crp vlan
configure pim delete vlan
configure pim dr-priority
configure pim iproute sharing hash
configure pim register-policy
configure pim register-policy rp
configure pim register-rate-limit-interval
configure pim register-suppress-interval register-probe-interval
configure pim register-checksum-to
configure pim shutdown-priority
configure pim spt-threshold
configure pim ssm range
configure pim state-refresh
configure pim state-refresh timer origination-interval
configure pim state-refresh timer source-active-timer
configure pim state-refresh ttl
configure pim timer vlan
configure pim vlan trusted-gateway
disable igmp
disable igmp snooping
disable igmp snooping vlan fast-leave
disable igmp ssm-map
disable ipmcforwarding
disable mvr
disable pim
disable pim iproute sharing
disable pim snooping
disable pim ssm vlan
enable igmp
enable igmp snooping
enable igmp snooping vlan fast-leave
enable igmp snooping with-proxy
enable igmp ssm-map
enable ipmcforwarding
enable mvr
enable pim
enable pim iproute sharing
enable pim snooping
enable pim ssm vlan
mrinfo
```



```

mtrace
rtlookup
rtlookup rpf
show igmp
show igmp group
show igmp snooping
show igmp snooping cache
show igmp snooping vlan
show igmp snooping vlan filter
show igmp snooping vlan static
show igmp ssm-map
show ipmroute
show iproute multicast
show L2stats
show mcast cache
show mvr
show mvr cache
show pim
show pim cache
show pim snooping
unconfigure igmp
unconfigure igmp snooping vlan ports set join-limit
unconfigure igmp ssm-map
unconfigure pim
unconfigure pim ssm range

```

This chapter describes commands for configuring and managing the following IPv4 multicast features:

- Multicast routing
- Protocol Independent Multicast (PIM)
- Internet Group Management Protocol (IGMP)
- Multicast VLAN Registration (MVR)

For an introduction to these features, see the ExtremeXOS Concepts Guide.

New Commands in ExtremeXOS 15.2:

- `configure mcast ipv4 cache timeout`

clear igmp group

```
clear igmp group {grpipaddress} {{vlan} name}
```

Description

Removes one or all IGMP groups.



Syntax Description

<i>grpipaddress</i>	Specifies the group IP address.
<i>name</i>	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

This command can be used by network operations to manually remove learned IGMP group entries instantly. Traffic is impacted until the IGMP groups are relearned. Use this command for diagnostic purposes only.

Example

The following command clears all IGMP groups from VLAN accounting:

```
clear igmp group accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

clear igmp snooping

```
clear igmp snooping {{vlan} name}
```

Description

Removes one or all IGMP snooping entries.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------



Default

N/A.

Usage Guidelines

This command can be used by network operations to manually remove IGMP snooping entries instantly. However, removing an IGMP snooping entry can disrupt the normal forwarding of multicast traffic, until the snooping entries are learned again.

The dynamic IGMP snooping entries are removed, then recreated upon the next general query. The static router entry and static group entries are removed and recreated immediately.

This command clears both the IGMPv2 and IGMPv3 snooping entries.

Example

The following command clears IGMP snooping from VLAN accounting:

```
clear igmp snooping accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

clear pim cache

```
clear pim cache {ipv4 | ipv6} {group_addr {source_addr}}
```

Description

Resets the IP multicast cache table.

Syntax Description

ipv4	Specifies an IPv4 address.
ipv6	Specifies an IPv6 address.
<i>group_addr</i>	Specifies a group address.
<i>source_addr</i>	Specifies a source IP address.



Default

If no options are specified, all IP multicast cache entries are flushed.

Usage Guidelines

This command can be used by network operators to manually remove IPMC software and hardware forwarding cache entries instantly. If the stream is available, caches are re-created, otherwise caches are removed permanently. This command can disrupt the normal forwarding of multicast traffic.

Example

The following command resets the IP multicast table for group 224.1.2.3:

```
clear pim cache 224.1.2.3
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

clear pim snooping

```
clear pim snooping {vlan} name
```

Description

Clears all PIM snooping neighbors, joins received on the VLAN, and the VLAN forwarding entries.

Syntax Description

<i>name</i>	Specifies the VLAN to which this command applies.
-------------	---

Default

N/A.

Usage Guidelines

None.



Example

The following command clears the PIM snooping database for the Default VLAN:

```
clear pim snooping "Default"
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure forwarding ipmc compression

```
configure forwarding ipmc compression {group-table | off}
```

Description

Enables or disables compression of entries in the IP multicast group table to facilitate improved IP multicast scaling.

Syntax Description

group-table	Enables compression.
off	Disables compression.

Default

group-table.

Usage Guidelines

Compression of IP multicast group table entries allows the switch to process more multicast traffic using the faster switch hardware instead of the relatively slower switch software. Compression requires additional processing. Disable this feature if you suspect a problem exposed by IP multicast compression.



When you enable or disable this feature, all IP multicast entries are flushed, and this can result in a temporary loss of multicast traffic while the IP multicast entries are relearned.



Note

On BlackDiamond X8 series switches and BlackDiamond 8800 series switches, all IP multicast forwarding entries utilizing the same IP multicast group table entry share a single backplane link, limiting the total throughput to 12Gbps.

To display the compression feature configuration, enter the command:

```
show forwarding configuration
```

Example

The following command disables compression:

```
configure forwarding ipmc compression off
```

History

This command was first available in ExtremeXOS 12.2.

Platform Availability

This command is available on all Summit family switches, BlackDiamond X8 series switches and BlackDiamond 8000 series modules.

configure forwarding ipmc lookup-key

```
configure forwarding ipmc lookup-key [group-vlan | source-group-vlan | mac-vlan | mixed-mode]
```

Description

Enables you to choose the lookup-key for multicast forwarding.

Syntax Description

group-vlan	Specifies that IP multicast forwarding database entries are programmed as (*,GroupIP,VlanId).
source-group-vlan	Specifies that IP multicast forwarding database entries are programmed as (SourceIP, GroupIP, VlanId). (Default).



mac-vlan	Specifies that IP multicast forwarding database entries are programmed as (Mac, VlanId).
mixed-mode	Specifies that IP multicast forwarding database entries are programmed as follows: L3 cache entries (PIM/MVR/PVLAN) use source-group-vlan; L2 cache entries (IGMP/MLD/PIM snooping) use mac-vlan.

Default

source-group-vlan.

Usage Guidelines

Use this command to choose the lookup-key for multicast forwarding.

The following restrictions apply to this command:

The `configure forwarding ipmc lookup-key mac-vlan` command is disallowed under the following conditions.

- If IPMC forwarding is enabled on at least on one VLAN
- If MVR is enabled either globally or on a VLAN

Similarly, enabling the above two features are disallowed, when the ipmc lookup-key is **mac-vlan**.

The following warning message is displayed when the **mac-vlan** option is specified:

Warning: Usage of multicast IP addresses that could result in overlapping MAC addresses should be avoided. Example: Using 225.1.1.1, 226.1.1.1 and 225.129.1.1 should be avoided. Either one of the addresses could be used. Using multicast with PVLAN should be avoided with this forwarding option.

- Mixed-mode `configure forwarding ipmc lookup-key mixed-mode`

If the chassis or stack has a member node with Felix/Helix/Firebolt*, then the command is disallowed. After enabling this mode, if a new member with unsupported chipset joins, then that card will be failed. The following warning message is displayed when the “mixed-mode” option is specified: **Warning: Usage of multicast IP addresses that could result in overlapping MAC addresses should be avoided for snooping (IGMP/MLD/PIM snooping) controlled traffic.**

Example: Using 225.1.1.1, 226.1.1.1 and 225.129.1.1 should be avoided. Either one of the addresses could be used.

- The `configure igmp snooping forwarding-mode [group-vlan | source-group-vlan]` command was introduced to support (*, G, V) forwarding before the IPMC compression feature was introduced. Because we are introduced IPv6 multicast support in EXOS 15.2, this command is deprecated, and the new `configure forwarding ipmc lookup-key` command now covers both IPv4 and IPv6.



Example

The following command specifies that IP multicast forwarding database entries are programmed as (*,GroupIP,VlanId):

```
configure forwarding ipmc lookup-key group-vlan
```

To display the ipmc lookup-key configuration, enter the command:

```
show forwarding configuration
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all Summit family switches, BlackDiamond X8 series switches and BlackDiamond 8000 series modules.

configure igmp

```
configure igmp query_interval query_response_interval last_member_query_interval  
{robustness}
```

Description

Configures the Internet Group Management Protocol (IGMP) timers.

Syntax Description

<i>query_interval</i>	Specifies the interval (in seconds) between general queries.
<i>query_response_interval</i>	Specifies the maximum query response time (in seconds).
<i>last_member_query_interval</i>	Specifies the maximum group-specific query response time (in seconds).
<i>robustness</i>	Specifies the degree of robustness for the network.

Default

- query interval—125 seconds
- query response interval—10 seconds
- last member query interval—1 second
- robustness—2



Usage Guidelines

Timers are based on RFC2236. Specify the following:

- query interval—The amount of time, in seconds, the system waits between sending out general queries. The range is 1 to 429,496,729 seconds.
- query response interval—The maximum response time inserted into the periodic general queries. The range is 1 to 25 seconds.
- last member query interval—The maximum response time inserted into a group-specific query sent in response to a leave group message. The range is 1 to 25 seconds.
- robustness—The degree of robustness of the network. The range is 2 to 7. This parameter allows tuning for the expected packet loss on a link. If a link is expected to have packet loss, this parameter can be increased.
- The group timeout is defined by the formula: $\text{group_timeout} = (\text{query_interval} \times \text{robustness}) + \text{query_response_interval}$, according to RFC 2236. You can explicitly define the host timeout using the `configure igmp snooping timer router_timeout host_timeout {vr vrname}` command. The effective `host_timeout` is the lesser value of the `group_timeout` and the configured `host_timeout`.

Example

The following command configures the IGMP timers:

```
configure igmp 100 5 1 3
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure igmp router-alert receive-required

```
configure igmp router-alert receive-required [on | off] {{vlan} vlan_name}
```

Description

Controls when the router-alert option is required for IGMPv2 and IGMPv3 packet reception and processing.



Syntax Description

vlan	Applies the configuration only to the specified VLAN. If no VLAN is specified, the configuration applies to all VLANs.
-------------	--

Default

Off—All IGMP packets are received and processed.

Usage Guidelines

By default, the ExtremeXOS software receives and processes all IGMP packets, regardless of the setting of the router-alert option within a packet. The default configuration works with all switches that support the ExtremeXOS software.

IETF standards require that a router accept and process IGMPv2 and IGMPv3 packets only when the router-alert option is set. The on setting for this command sets the ExtremeXOS software to comply with the IETF standards and should be used when the switch will be used with third-party switches that expect IETF compliant behavior.

Example

The following command configures the switch for IETF compliant IGMP packet processing:

```
configure igmp router-alert receive-required on
```

History

This command was first available in ExtremeXOS 12.5.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure igmp router-alert transmit

```
configure igmp router-alert transmit [on | off] {{vlan} vlan_name}
```

Description

Controls whether the router-alert option is set when forwarding IGMPv2 and IGMPv3 packets.



Syntax Description

vlan	Applies the configuration only to the specified VLAN. If no VLAN is specified, the configuration applies to all VLANs.
-------------	--

Default

On—The router-alert option is set when forwarding IGMPv2 and IGMPv3 packets.

Usage Guidelines

IETF standards require that a router set the router-alert option in forwarded IGMPv2 and IGMPv3 packets. The ExtremeXOS software has been updated to comply with this requirement using the default settings.

Earlier versions of the ExtremeXOS software forwarded all IGMP packets without setting the router-alert option. If compatibility issues arise, you can configure the software to use the legacy behavior by using this command with the off option.

Example

The following command configures the switch for IETF compliant IGMP packet processing:

```
configure igmp router-alert transmit on
```

History

This command was first available in ExtremeXOS 12.5.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure igmp snooping filters

```
configure igmp snooping filters [per-port | per-vlan]
```

Description

Selects the type of IGMP snooping filters that are installed.



Syntax Description

per-port	Installs the per-port IGMP snooping filters.
per-vlan	Installs the per-VLAN IGMP snooping filters.

Default

per-port.

Usage Guidelines

This command applies only to Summit family switches and BlackDiamond 8800 series switches.

Use the per-vlan option when the number of VLANs configured on the switch is lower than the maximum numbers listed in the following table. This option conserves usage of the hardware Layer 3 multicast forwarding table.

When the number of configured VLANs is larger than the maximum values listed here, select the per-port option. Each VLAN requires additional interface hardware ACL resources. The per-port option conserves usage of the interface hardware ACL resources.

Table 55: Maximum Number of VLANs Supported by per-VLAN IGMP Snooping Filters

Summit Switch and BlackDiamond 8000 Series Module Type	Maximum Number of VLANs When per-VLAN Snooping Filters are Installed ¹³
a Series	1000
c Series	2000
e Series	448
xl Series	2000

To display the IGMP snooping filters configuration, use the show igmp snooping command.

Note



For MLD Snooping, the maximum number of VLANs is half of the numbers provided in this table.

The maximum number specified here is individual limit for IGMP snooping filters. If both IGMP and MLD snooping filters are used, the maximum numbers are lower than the ones specified.

Example

The following command configures the switch to install the per-VLAN IGMP snooping filters:

```
configure igmp snooping filters per-vlan
```

¹³ The actual maximum value is smaller if other processes require entries in the interface ACL table.

¹³ The actual maximum value is smaller if other processes require entries in the interface ACL table.



History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure igmp snooping flood-list

```
configure igmp snooping flood-list [policy | none] {vr vrname}
```

Description

Configures certain multicast addresses to be slow path flooded within the VLAN.

Syntax Description

<i>policy</i>	Specifies a policy file with a list of multicast addresses to be handled.
none	Specifies no policy file is to be used.
<i>vrname</i>	Specifies a virtual router.

Default

None.

Usage Guidelines

With this command, a user can configure certain multicast addresses to be slow path flooded within the VLAN, which otherwise are fast path forwarded according to IGMP and/or Layer 3 multicast protocol.

A policy file is a text file with the extension, .pol. It can be created or edited with any text editor. The specified policy file <policy file> should contain a list of addresses which determine if certain multicast streams are to be treated specially. Typically, if the switch receives a stream with a destination address which is in the <policy file> in 'permit' mode, that stream is software flooded and no hardware entry is installed.

When adding an IP address into the policy file, a 32-bit host address is recommended.

This feature is meant to solve the multicast connectivity problem for unknown destination addresses within system reserved ranges. Specifically this feature was introduced to solve the problem of recognizing certain streams as control packets.

To create a policy file for the snooping flood-list, use the following template:



```
# This is a template for IGMP Snooping Flood-list Policy File # Add your group
addresses between "Start" and "End" # Do not touch the rest of the file!!!! entry
igmpFlood { if match any { #----- Start of group addresses
----- nlri 234.1.1.1/32; nlri 239.1.1.1/32; #----- end
of group addresses ----- } then { permit; } } entry catch_all { if
{ } then { deny; } }
```

Note



The switch does not validate any IP address in the policy file used in this command. Therefore, slow-path flooding should be used only for streams which are very infrequent, such as control packets. It should not be used for multicast data packets. This option overrides any default mechanism of hardware forwarding (with respect to IGMP, PIM, or DVMRP) so it should be used with caution.

Slow path flooding is done within the L2 VLAN only.

Use the none option to effectively disable slow path flooding.

You can use the `show igmp` command to see the configuration of slow path flooding.

Example

The following command configures the multicast data stream specified in access1 for slow path flooding:

```
configure igmp snooping flood-list access1
```

The following command specifies that no policy file is to be used, this effectively disabling slow path flooding:

```
configure igmp snooping flood-list none
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure igmp snooping leave-timeout

```
configure igmp snooping leave-timeout leave_timeout_ms {vr vrname}
```



Description

Configures the IGMP snooping leave timeout.

Syntax Description

<i>leave_timeout_ms</i>	Specifies an IGMP leave timeout value in milliseconds.
<i>vrname</i>	Specifies a virtual router.

Default

1000 ms.

Usage Guidelines

The leave-timeout is the IGMP leave override interval. If no other hosts override the IGMP leave by the end of this interval, the receiver port is removed.

The range is 0 - 175000 ms (175 seconds). For timeout values of one second or less, you must set the leave-timeout to a multiple of 100ms. For values of more than one second, you must set the leave-timeout to a multiple of 1000ms (one second).

Example

The following command configures the IGMP snooping leave timeout:

```
configure igmp snooping leave-timeout 10000
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure igmp snooping timer

```
configure igmp snooping timer router_timeout host_timeout {vr vrname}
```

Description

Configures the IGMP snooping timers.



Syntax Description

<i>router_timeout</i>	Specifies the time in seconds before removing a router snooping entry.
<i>host_timeout</i>	Specifies the time in seconds before removing a host's group snooping entry.
<i>vrname</i>	Specifies a virtual router.

Default

The router timeout default setting is 260 seconds. The host timeout setting is 260 seconds.

Usage Guidelines

Timers should be set to approximately 2.5 times the router query interval in use on the network. Specify the following:

- router timeout—The maximum time, in seconds, that a router snooping entry can remain in the IGMP snooping table without receiving a router report. If a report is not received, the entry is deleted. The range is 10 to 214,748,364 seconds (6.8 years). The default setting is 260 seconds.
- host timeout—The maximum time, in seconds, that a group snooping entry can remain in the IGMP snooping table without receiving a group report. If a report is not received, the entry is deleted. The range is 10 to 214,748,364 seconds. The default setting is 260 seconds.

Note



The *host_timeout* value should be less than or equal to the query timeout value, which is defined by the following configure igmp command timers as follows: $\text{query_interval} \times \text{robustness} + \text{query_response_interval}$.

IGMP snooping expects at least one device on every VLAN to periodically generate IGMP query messages. Without an IGMP querier, the switch eventually stops forwarding IP multicast packets to any port, because the IGMP snooping entries time out, based on the value specified in *host_timeout* or *router_timeout*.

Example

The following command configures the IGMP snooping timers:

```
configure igmp snooping timer 600 600
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure igmp snooping vlan ports add dynamic group

```
configure igmp snooping {vlan} vlan_name {ports portlist} add dynamic group
[ v4group ]
```

Description

Configures an IGMP dynamic group.

Syntax Description

<i>vlan_name</i>	Specifies a vlan name.
<i>portlist</i>	Specifies a port list.
v4group	Specifies a version 4 group.

Default

N/A.

Usage Guidelines

This command is not saved in the configuration. The following message is displayed on execution of this command:

```
INFO: This command is not saved in the configuration.
```

Example

```
show igmp group
Group Address      Ver Vlan      Port      Age
224.0.0.2          2   v1          1         40
224.0.0.6          2   v1          1         44
225.1.1.1(s)      2   v3          Lpbk      0

show igmp snooping
Igmp Snooping Flag      : forward-all-router
Igmp Snooping Flood-list : none
Igmp Snooping Proxy     : Enable
Igmp Snooping Filters   : per-port

Vlan      Vid  Port  #Senders #Receivers Router Enable
-----
Default   1    0     0         0         Yes  Yes
v1        10   3     3         0         Yes  Yes

v3        4088 1     0         0         Yes  Yes
          Lpbk 1     1         1         No   No
```



History

This command was first available in ExtremeXOS 15.3.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure igmp snooping vlan ports add static group

```
configure igmp snooping {vlan} vlanname {ports portlist }add static group ip address
```

Description

Configures VLAN ports to receive the traffic from a multicast group, even if no IGMP joins have been received on the port.

Syntax Description

<i>vlanname</i>	Specifies a VLAN name.
<i>portlist</i>	Specifies one or more ports or slots and ports.
<i>ip address</i>	Specifies the multicast group IP address.

Default

None.

Usage Guidelines

Use this command to forward a particular multicast group to VLAN ports. In effect, this command emulates a host on the port that has joined the multicast group. As long as the port is configured with the static entry, multicast traffic for that multicast group is forwarded to that port.

This command is for IGMPv2 only.

The switch sends proxy IGMP messages in place of those generated by a real host. The proxy messages use the VLAN IP address for source address of the messages. If the VLAN has no IP address assigned, the proxy IGMP message uses 0.0.0.0 as the source IP address.

The multicast group should be in the class-D multicast address space, but should not be in the multicast control subnet range (224.0.0.x/24).

If the ports also have an IGMP filter configured, the filter entries take precedence. IGMP filters are configured using the command:

```
configure igmp snooping vlan vlanname ports portlist filterpolicy file
```



Example

The following command configures a static IGMP entry so the multicast group 224.34.15.37 is forwarded to VLAN marketing on ports 2:1-2:4:

```
configure igmp snooping marketing ports 2:1-2:4 add static group 224.34.15.37
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure igmp snooping vlan ports add static router

```
configure igmp snooping {vlan} vlanname ports portlist add static router
```

Description

Configures VLAN ports to forward the traffic from all multicast groups, even if no IGMP joins have been received on the port.

Syntax Description

<i>vlanname</i>	Specifies a VLAN name.
<i>portlist</i>	Specifies one or more ports or slots and ports.

Default

None.

Usage Guidelines

Use this command to forward all multicast groups to the specified VLAN ports. In effect, this command emulates a multicast router attached to those ports. As long as the ports are configured with the static entry, all available multicast traffic is forwarded to those ports.



Example

The following command configures a static IGMP entry so all multicast groups are forwarded to VLAN marketing on ports 2:1-2:4:

```
configure igmp snooping marketing ports 2:1-2:4 add static router
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure igmp snooping vlan ports delete static group

```
configure igmp snooping {vlan} vlanname ports portlist delete static  
group[ip_address | all]
```

Description

Removes the port configuration that causes multicast group traffic to be forwarded, even if no IGMP leaves have been received on the port.

Syntax Description

<i>vlanname</i>	Specifies a VLAN name.
<i>portlist</i>	Specifies one or more ports or slots and ports.
<i>ip_address</i>	Specifies the multicast group IP address.
all	Delete all the static groups.

Default

None.

Usage Guidelines

Use this command to remove a static group entry created by the following command:

```
configure igmp snooping vlan vlanname portsportlist add static group ipaddress
```



Example

The following command removes a static IGMP entry that forwards the multicast group 224.34.15.37 to the VLAN marketing on ports 2:1-2:4:

```
configure igmp snooping marketing ports 2:1-2:4 delete static group
224.34.15.37
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure igmp snooping vlan ports delete static router

```
configure igmp snooping vlan vlanname ports portlist delete static router
```

Description

Removes the configuration that causes VLAN ports to forward the traffic from all multicast groups, even if no IGMP joins have been received on the port.

Syntax Description

<i>vlanname</i>	Specifies a VLAN name.
<i>portlist</i>	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

None.

Usage Guidelines

Use this command to remove an entry created by the following command:

```
configure igmp snooping vlan vlanname ports portlist add static router
```



Example

The following command removes the static IGMP entry that caused all multicast groups to be forwarded to VLAN marketing on ports 2:1-2:4:

```
configure igmp snooping marketing ports 2:1-2:4 delete static router
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure igmp snooping vlan ports filter

```
configure igmp snooping vlan vlanname ports portlist filter [policy | none]
```

Description

Configures an IGMP snooping policy file filter on VLAN ports.

Syntax Description

<i>vlanname</i>	Specifies a VLAN name.
<i>portlist</i>	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
<i>policy</i>	Specifies the policy file for the filter.

Default

None.

Usage Guidelines

Use this command to filter multicast groups to the specified VLAN ports.

The policy file used by this command is a text file that contains the class-D addresses of the multicast groups that you wish to block.

To remove IGMP snooping filtering from a port, use the none keyword version of the command.

Use the following template to create a snooping filter policy file:




```
# # Add your group addresses between "Start" and "end" # Do not touch the rest of
the file!!!! entry igmpFilter { if match any { #----- Start of group
addresses ----- nlri 239.11.0.0/16; nlri 239.10.10.4/32;
#----- end of group addresses ----- } then
{ deny; } } entry catch_all { if { } then { permit; } }
```

Example

The following command configures the policy file `ap_multicast` to filter multicast packets forwarded to VLAN marketing on ports 2:1-2:4:

```
configure igmp snooping marketing ports 2:1-2:4 filter ap_multicast
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure igmp snooping vlan ports set join-limit

```
configure igmp snooping {vlan} vlanname ports portlist set join-limit {num}
```

Description

Configures VLAN ports to support a maximum number of IGMP joins.

Syntax Description

<i>vlanname</i>	Specifies a VLAN name.
<i>portlist</i>	Specifies one or more ports or slots and ports.
<i>num</i>	Specifies the maximum number of joins permitted on the ports. The range is 1 to 500.

Default

No limit.

Usage Guidelines

None.



Example

The following command configures port 2:1 in the Default VLAN to support a maximum of 100 IGMP joins:

```
configure igmp snooping "Default" ports 2:1 set join-limit 100
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure igmp ssm-map add

```
configure igmp ssm-map add group_ip [prefix | mask] [source_ip | src_domain_name]  
{vr vr-name}
```

Description

Configures an IGMP SSM mapping.

Syntax Description

<i>group_ip</i>	Specifies the multicast IP address for the group mapping.
<i>prefix</i>	Specifies a prefix length for the multicast group IP address. The range is 4 to 32.
<i>mask</i>	Specifies the network mask for the group multicast IP address.
<i>source_ip</i>	The IP address for a multicast group source.
<i>src_domain_name</i>	The source domain name for the multicast group source.
<i>vr-name</i>	Specifies a virtual router name. If the VR name is omitted, the switch uses the VR specified by the current CLI VR context.

Default

N/A.

Usage Guidelines

IGMP SSM mapping operates only with IPv4.



Example

The following command configures an IGMP-SSM mapping for the range of multicast IP addresses at 232.1.1.0/24 to IP host 172.16.8.1:

```
configure igmp ssm-map add 232.1.1.0/24 172.16.8.1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure igmp ssm-map delete

```
configure igmp ssm-map delete group_ip [prefix] | mask] [source_ip | all] v vr-name}
```

Description

Unconfigures an SSM mapping.

Syntax Description

<i>group_ip</i>	Specifies the multicast IP address for the group mapping.
<i>prefix</i>	Specifies a prefix length for the multicast group IP address. The range is 4 to 32.
<i>mask</i>	Specifies the network mask for the group multicast IP address.
<i>source_ip</i>	The IP address for a multicast group source.
all	Specifies that all sources for the specified group or mask are deleted.
<i>vr-name</i>	Specifies a virtual router name. If the VR name is omitted, the switch uses the VR specified by the current CLI VR context.

Default

N/A.

Usage Guidelines

None.



Example

The following command deletes an IGMP-SSM mapping for the range of multicast IP addresses at 232.1.1.0/24 to IP host 172.16.8.1:

```
configure igmp ssm-map delete 232.1.1.0/24 172.16.8.1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ipmcforwarding

```
configure ipmcforwarding to-cpu [auto | off] ports port_list
```

Description

Configure whether IP multicast CPU filters are installed automatically.

Syntax Description

auto	The software will automatically program IP multicast processing based on configuration.
off	IP multicast packets received on this port are always flooded with no CPU processing.
<i>port_list</i>	Specifies one or more ports.

Default

N/A.

Usage Guidelines

IP forwarding and IPMC forwarding must be enabled for the configuration to operate.



Example

The following example configures automatic operation for port 2.1:

```
configure ipmcforwarding to-cpu auto ports 2.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ipmroute add

```
configure ipmroute add [source-net mask-len | source-net mask] {{protocol}}
protocol rpf-address {metric} {vr vr-name}
```

Description

Adds a static multicast route to the multicast routing table.

Syntax Description

<i>source-net</i>	Specifies an IP address/mask length.
<i>mask-len</i>	Mask length for the IP multicast source's subnet. Range is [1-32].
<i>mask</i>	Specifies a subnet mask.
<i>protocol</i>	Unicast routing protocol that is to be used for route learning.
<i>rpf-address</i>	Next hop through which the multicast source can be reached.
<i>metric</i>	Specifies a cost metric.
<i>vr-name</i>	Specifies the virtual router to which the route is added.

Default

The following defaults apply:

- *metric*—1
- *vr-name*—VR of the current CLI context
- *protocol*—none



Usage Guidelines

This command allows you to statically configure where multicast sources are located (even though the unicast routing table has different entries). It allows you to configure a multicast static route in such a way as to have non-congruent topology for Unicast and Multicast topology and traffic.

Example

The following command configures a multicast static route for all multicast sources within network subnet 192.168.0.0/16. Those sources are reachable through the gateway 192.75.0.91.

```
configure ipmroute add 192.168.0.0/16 192.75.0.91
```

The following example configures multicast static route for all sources via a single gateway with a metric of 100:

```
configure ipmroute add 0.0.0.0/0 192.75.0.91 100
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure ipmroute delete

```
configure ipmroute delete [source-net/mask-len | source-net mask] {{protocol}  
protocol} rpf-address {vr vr-name}
```

Description

Deletes a static multicast address from the multicast routing table.

Syntax Description

<i>source-net</i>	Specifies an IP address/mask length.
<i>mask-len</i>	Mask length for the IP multicast source's subnet. Range is [1-32].
<i>mask</i>	Specifies a subnet mask.
<i>protocol</i>	Unicast routing protocol that is to be used for route learning.



<i>rpf-address</i>	Next hop through which the multicast source can be reached.
<i>vr-name</i>	Specifies the virtual router to which the route is added.

Default

vr-name is the VR of the current CLI context.

Usage Guidelines

This command allows you to delete an existing multicast static route. It allows you to configure congruent topology for unicast and multicast packets and traffic.

Example

The following command deletes a multicast static route:

```
configure ipmroute delete 192.168.0.0/16 192.75.0.91
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure iproute add (Multicast)

```
configure iproute add [ipNetmask | ip_addr mask] gateway {metric} {multicast | multicast-only | unicast | unicast-only} {vr vrname}
```

Description

Adds a static route to the routing table.

Syntax Description

<i>ipNetmask</i>	Specifies an IP address/mask length.
<i>ip_addr</i>	Specifies an IP address.
<i>mask</i>	Specifies a subnet mask.
<i>gateway</i>	Specifies a VLAN gateway.



metric	Specifies a cost metric.
vrname	Specifies the virtual router to which the route is added.
multicast	Adds the specified route to the multicast routing table.
multicast-only	Adds the specified route to the multicast routing table. This option is provided for backward compatibility with releases prior to ExtremeXOS Release 12.1.
unicast	Adds the specified route to the unicast routing table.
unicast-only	Adds the specified route to the unicast routing table. This option is provided for backward compatibility with releases prior to ExtremeXOS Release 12.1.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

Use a mask value of 255.255.255.255 to indicate a host entry.



Note

Although dynamic unicast routes can be captured in the multicast routing table, unicast static routes cannot be captured in the multicast routing table. To create a static route for the multicast routing table, you must specify the multicast option.

Example

The following command adds a static address to the multicast routing table:

```
configure iproute add 10.1.1.0/24 123.45.67.1 5 multicast
```

History

This command was first available in ExtremeXOS 10.1.

The multicast and unicast keywords were first available in ExtremeXOS 12.1. These keywords replace the multicast-only and unicast-only keywords, which remain in the software for backward compatibility.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure iproute delete

```
configure iproute delete [ipNetmask | ipaddress mask] gateway {multicast | multicast-only | unicast | unicast-only} {vr vrname}
```

Description

Deletes a static address from the routing table.

Syntax Description

<i>ipNetmask</i>	Specifies an IP address/mask length.
<i>ipaddress</i>	Specifies an IP address.
<i>mask</i>	Specifies a subnet mask.
<i>gateway</i>	Specifies a VLAN gateway.
multicast	Specifies a multicast route to delete.
multicast-only	Specifies a multicast route to delete.
unicast	Specifies a unicast route to delete.
unicast-only	Specifies a unicast route to delete.
<i>vrname</i>	Specifies the virtual router to which the route is deleted.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

Use a value of 255.255.255.255 or /32 for mask to indicate a host entry.

Example

The following command deletes an address from the multicast routing table:

```
configure iproute delete 10.101.0.0/24 10.101.0.1 multicast
```

History

This command was first available in ExtremeXOS 10.1.

The multicast and unicast keywords were first available in ExtremeXOS 12.1. These keywords replace the multicast-only and unicast-only keywords, which remain in the software for backward compatibility.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mcast ipv4 cache timeout

```
configure mcast ipv4 cache timeout {seconds | none}
```

Description

Configures the IPv4 multicast cache timeout.

Syntax Description

<i>seconds</i>	Idle time after which cache entries are deleted.
none	Cache entries are not timed out.

Default

300 seconds.

Usage Guidelines

Cache timeout is the time after which the cache entries are deleted, if traffic is not received for that duration. The applies only for snooping and MVR caches and does not apply for PIM caches.

The range is 90 to 100000 seconds. You can use the option none if you do not want the cache entry to be deleted. If none is configured, the cache entries could be deleted only using the following command:

```
clear igmp snooping
```

Example

```
configure mcast ipv4 cache timeout 400
configure mcast ipv4 cache timeout none
```

History

This command was first available in ExtremeXOS 15.2.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mvr add receiver

```
configure mvr vlan vlan-name add receiver port port-list
```

Description

Configures a port to receive MVR multicast streams.

Syntax Description

<i>vlan-name</i>	Specifies a VLAN name.
<i>port-list</i>	A list of ports or slots and ports.

Default

N/A.

Usage Guidelines

This command is used to add a group of virtual ports for multicast forwarding through MVR. By default, some ports on non-MVR VLANs (router ports, primary and secondary EAPS ports), are excluded from the MVR cache egress list. This command is used to override these rules, so that if valid IGMP memberships are received, or a router is detected, streams are forwarded out on the ports.

Example

The following command adds the ports 1:1 and 1:2 of VLAN v1 to MVR for forwarding:

```
configure mvr vlan v1 add receiver port 1:1-1:2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MVR feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure mvr add vlan

```
configure mvr add vlan vlan-name
```

Description

Configures a VLAN as an MVR VLAN.

Syntax Description

<i>vlan-name</i>	Specifies a VLAN name.
------------------	------------------------

Default

N/A.

Usage Guidelines

Configures MVR on the specified VLAN. When a multicast stream in the specified MVR address range is received on the VLAN, it is leaked to all other VLAN ports where the corresponding IGMP join message is received. By default, the entire multicast address range 224.0.0.0/4, except for the multicast control range 224.0.0.0/24 is used for MVR. To change the MVR address range, use the following command:

```
configure mvr vlan vlan-name mvr-address {policy-name | none}
```

Example

The following command configures VLAN v1 as an MVR VLAN:

```
configure mvr add vlan v1
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MVR feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mvr delete receiver

```
configure mvr vlan vlan-name delete receiver port port-list
```



Description

Configures a port not to receive MVR multicast streams.

Syntax Description

<i>vlan-name</i>	Specifies a VLAN name.
<i>port-list</i>	A list of ports or slots and ports.

Default

N/A.

Usage Guidelines

This command is used to delete a group of virtual ports for multicast forwarding through MVR. After using this command, the ports revert to the default forwarding rules.

Example

The following command deletes the ports 1:1 and 1:2 of VLAN v1 to MVR for forwarding:

```
configure mvr vlan v1 delete receiver port 1:1-1:2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MVR feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mvr delete vlan

```
configure mvr delete vlan vlan-name
```

Description

Deletes a VLAN from MVR.

Syntax Description

<i>vlan-name</i>	Specifies a VLAN name.
------------------	------------------------



Default

N/A.

Usage Guidelines

Removes MVR from the specified VLAN.

Example

The following command configures VLAN v1 as a non-MVR VLAN:

```
configure mvr delete vlan v1
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MVR feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mvr mvr-address

```
configure mvr vlan vlan-name mvr-address {policy-name | none}
```

Description

Configures the MVR address range on a VLAN.

Syntax Description

<i>vlan-name</i>	Specifies a VLAN name.
<i>policy-name</i>	Specifies a policy file.

Default

The default address range is 224.0.0.0/4 (all multicast addresses), but excluding 224.0.0.0/24 (the multicast control range).

Usage Guidelines

If no policy file is specified (the none option), the entire multicast address range 224.0.0.0/4, except for the multicast control range 224.0.0.0/24 is used for MVR.



MVR must first be configured on the VLAN before using this command.

If the policy is later refreshed, groups denied and newly allowed groups in the policy are flushed from fast path forwarding. This allows synching existing channels with the new policy, without disturbing existing channels.

The following is a sample policy file `mvrpol.pol`. This policy configures 236.1.1.0/24 as the MVR address range. Any address outside this range has the standard switching behavior on an MVR VLAN.

```
Entry extreme1 {
  if match any {
    nlri 236.1.1.0/24 ;
  }
  then {
    permit ;
  }
}
```

Example

The following command configures the MVR address range specified in the policy file `mvrpol.pol` for the VLAN `v1`:

```
configure mvr vlan v1 mvr-address mvrpol
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MVR feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mvr static group

```
configure mvr vlan vlan-name static group {policy-name | none}
```

Description

Configures the MVR static group address range on a VLAN.



Syntax Description

<i>vlan-name</i>	Specifies a VLAN name.
<i>policy-name</i>	Specifies a policy file.

Default

By default, all the MVR group addresses work in static mode.

Usage Guidelines

If no policy file is specified (the none option), the entire multicast address range 224.0.0.0/4, except for the multicast control range 224.0.0.0/24, is used for static groups in MVR.

MVR must first be configured on the VLAN before using this command.

The following is a sample policy file mvrpol.pol. This policy configures 236.1.1.0/24 as the MVR static group address range. Any MVR addresses outside this range are dynamically registered through IGMP. An MVR VLAN will proxy join only for addresses that are not in the static group. If you want all the multicast groups to be dynamic, use a policy file with this command that denies all multicast addresses.

```
Entry extremel {
    if match any {
        nlri 236.1.1.0/24 ;
    }
    then {
        permit ;
    }
}
```

Example

The following command configures the MVR static group address range specified in the policy file mvrpol.pol for the VLAN v1:

```
configure mvr vlan v1 static group mvrpol
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MVR feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure pim add vlan

```
configure pim {ipv4 | ipv6} add vlan [vlan-name | all] {dense | sparse} {passive}
```

Description

Configures an IP interface for PIM.

Syntax Description

IPv4	Specifies the IPv4 address family.
IPv6	Specifies the IPv6 address family.
<i> vlan-name </i>	Specifies a VLAN name.
all	Specifies all VLANs.
dense	Specifies PIM dense mode (PIM-DM).
sparse	Specifies PIM sparse mode (PIM-SM).
passive	Specifies a passive interface.

Default

Dense.

Usage Guidelines

When an IP interface is created, per-interface PIM configuration is disabled by default.

The switch supports both dense mode and sparse mode operation. You can configure dense mode or sparse mode on a per-interface basis. After they are enabled, some interfaces can run dense mode, while others run sparse mode.

Passive interfaces are host only interfaces that allow a multicast stream from other VLANs to be forwarded to edge hosts. Since they do not peer with other PIM routers, do not connect a multicast router to a passive interface.

In order for the interface to participate in PIM, PIM must be enabled on the switch using the following command:

```
enable pim
```

Example

The following command enables PIM-DM multicast routing on VLAN accounting:

```
configure pim add vlan accounting dense
```



History

This command was first available in ExtremeXOS 10.1.

The passive parameter was added in ExtremeXOS 11.1.

The IPv4 and IPv6 options were added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure pim border

```
configure pim {ipv4 | ipv6} [{vlan} vlan_name | vlan all] border
```

Description

Configures a PIM VLAN as a border VLAN, which is used to demarcate a PIM domain when using MSDP.

Syntax Description

ipv4	Configures a PIM timer on IPv4 router interfaces.
ipv6	Configures a PIM timer on IPv6 router interfaces.
<i>vlan_name</i>	Specifies a VLAN name.
vlan all	Specifies all VLANs.
border	Interface is domain border.

Default

None.

Usage Guidelines

MSDP is used to connect multiple multicast routing domains. A PIM-SM domain is created by limiting the reach of PIM BSR advertisements. When a border VLAN is configured, PIM BSR advertisements are not forwarded out of the PIM VLAN.

Example

The following command configures a PIM border on a VLAN called “vlan_border”:

```
configure pim vlan_border border
```



History

This command was first available in ExtremeXOS 10.1.

The **ipv4** and **ipv6** keywords were added giving an option to support this functionality in IPv6 as well, in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the ExtremeXOS Concepts Guide.

configure pim cbsr

```
configure pim cbsr {ipv4 | ipv6} [{vlan} vlan_name {priority [0-254]} | none]
```

Description

Configures a candidate bootstrap router for PIM sparse-mode operation.

Syntax Description

IPv4	Specifies the IPv4 address family.
IPv6	Specifies the IPv6 address family.
<i>vlan_name</i>	Specifies a VLAN name.
<i>priority</i>	Specifies a priority setting. The range is 0 - 254.
none	Specifies to delete a CBSR.

Default

The default setting for priority is 0, and indicates the lowest priority.

Usage Guidelines

The VLAN specified for CBSR must have PIM enabled for it to take effect. After PIM is enabled, CBSRs advertise themselves in the PIM domain. A bootstrap router (BSR) is elected among all the candidates based on CBSR priority. To break the tie among routers with the same priority setting, the router with the numerically higher IP address is chosen.

An Extreme XOS switch can support up to 145 RPs per group when it is configured as a PIM BSR (bootstrap router). If more than 145 RPs are configured for a single group, the BSR ignores the group and does not advertise the RPs. Non-BSR switches can process more than 145 RPs in the BSR message.



Example

The following command configures a candidate bootstrap router on the VLAN accounting:

```
configure pim cbsr vlan accounting 30
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure pim crp static

```
configure pim {ipv4 | ipv6} crp static ip_address [none | policy] {priority [0-254]}
```

Description

Configures a rendezvous point and its associated groups statically, for PIM sparse mode operation.

Syntax Description

ipv4	Specifies an IPv4 address.
ipv6	Specifies an IPv6 address.
<i>ip_address</i>	Specifies a static CRP address.
none	Deletes the static rendezvous point.
<i>policy</i>	Specifies a policy file name.
<i>priority</i>	Specifies a priority setting. The range is 0 - 254.

Default

The default setting for priority is 0, which indicates highest priority.

Usage Guidelines

In PIM-SM, the router sends a join message to the rendezvous point (RP). The RP is a central multicast router that is responsible for receiving and distributing multicast packets. If you use a static RP, all switches in your network must be configured with the same RP address for the same group (range).



ExtremeXOS switches support up to 50 RPs in a switch, and up to 180 groups (group/mask entries) in a single RP policy file. If you configure more than 180 group entries in a single RP policy file, the switch will not process entries added after the first 180.

The policy file contains a list of multicast group addresses served by this RP.

This policy file is not used for filtering purposes. As used with this command, the policy file is just a container for a list of addresses. So a typical policy file used for RP configuration looks a little different from a policy used for other purposes.

If routers have different group-to-RP mappings, due to misconfiguration of the static RP (or any other reason), traffic is disrupted.

Example

The following command statically configures an RP and its associated groups defined in policy file rp-list:

```
configure pim crp static 10.0.3.1 rp-list
```

The following is a sample policy file:

```
entry extremel {  
  if match any { }  
  then { nlri 224.0.0.0/4 ;  
        nlri 239.255.0.0/24 ;  
        nlri 232.0.0.0/8 ;  
        nlri 238.1.0.0/16 ;  
        nlri 232.232.0.0/20 ;  
  }  
}
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure pim crp timer

```
configure pim {ipv4 | ipv6} crp timer crp_adv_interval
```



Description

Configures the candidate rendezvous point advertising interval in PIM sparse mode operation.

Syntax Description

ipv4	Specifies an IPv4 address.
ipv6	Specifies an IPv6 address.
<i>crp_adv_interval</i>	Specifies a candidate rendezvous point advertising interval in seconds. The range is 1 to 1,717,986,918.

Default

The default is 60 seconds.

Usage Guidelines

Increasing this time results in increased convergence time for CRP information to the PIM routers.

Example

The following command configures the candidate rendezvous point advertising interval to 120 seconds:

```
configure pim crp timer 120
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure pim crp vlan

```
configure pim {ipv4 | ipv6} crp vlan vlan_name [none | policy] {priority}
```

Description

Configures the dynamic candidate rendezvous point (CRP) for PIM sparse-mode operation.



Syntax Description

IPv4	Specifies the IPv4 address family.
IPv6	Specifies the IPv6 address family.
<i>vlan_name</i>	Specifies a VLAN name.
none	Specifies to delete a CRP.
<i>policy</i>	Specifies a policy file name.
<i>priority</i>	Specifies a priority setting. The range is 0 - 254.

Default

The default setting for priority is 0 and indicates the highest priority.

Usage Guidelines

ExtremeXOS switches support up to 50 RPs in a switch, and up to 180 groups (group/mask entries) in a single RP policy file. If you configure more than 180 group entries in single RP policy file, then switch will not process entries added after first 180.

The policy file contains the list of multicast group addresses serviced by this RP. This set of group addresses are advertised as candidate RPs. Each router then elects the common RP for a group address based on a common algorithm. This group to RP mapping should be consistent on all routers.

This policy file is not used for filtering purposes. As used with this command, the policy file is just a container for a list of addresses. So a typical policy file used for RP configuration looks a little different from a policy used for other purposes. The following is a sample policy file which configures the CRP for the address ranges 239.0.0.0/24 and 232.144.27.0/24:

```
entry extremel {
    if match any {
    }
    then {
        nlri 239.0.0.0/24 ;
        nlri 232.144.27.0/24 ;
    }
}
```

The VLAN specified for a CRP must have PIM configured.

To delete a CRP, use the keyword none as the access policy.

Example

The following command configures the candidate rendezvous point for PIM sparse-mode operation on the VLAN HQ_10_0_3 with the policy rp-list and priority set to 30:

```
configure pim crp HQ_10_0_3 rp-list 30
```



History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure pim delete vlan

```
configure pim delete vlan [vlanname | all]
```

Description

Disables PIM on a router interface.

Syntax Description

<i>vlanname</i>	Specifies a VLAN name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables PIM on VLAN accounting:

```
configure pim delete vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure pim dr-priority

```
configure pim {ipv4 | ipv6} [ {vlan} vlan_name | vlan all ] dr-priority priority
```

Description

Configures the designated router (DR) priority that is advertised in PIM hello messages.

Syntax Description

ipv4	IPv4 address family (default).
ipv6	IPv6 address family.
vlan all	Apply to all VLANs.
dr-priority	Designated Router Priority for VLAN.
<i>priority</i>	Priority value for VLAN (default 1). The range is 1 to 4294967295.

Default

The default setting for **dr-priority** is 1.

Usage Guidelines

The DR_Priority option allows a network administrator to give preference to a particular router in the DR election process by giving it a numerically larger DR priority. The DR_Priority option is included in every hello message, even if no DR priority is explicitly configured on that interface. This is necessary because priority-based DR election is only enabled when all neighbors on an interface advertise that they are capable of using the DR_Priority Option. The default priority is 1.

The DR priority is a 32-bit unsigned number, and the numerically larger priority is always preferred. A router's idea of the current DR on an interface can change when a PIM hello message is received, when a neighbor times out, or when a router's own DR priority changes. If the router becomes the DR or ceases to be the DR, this will normally cause the DR register state machine to change state. Subsequent actions are determined by that state machine. DR election process on interface is:

- If any one of the neighbor on the interface is not advertised the DR priority (not DR capable) then DR priority will not considered for the all the neighbors in the circuit, and the primary IP address will be considered for all the neighbors.
- Higher DR priority or higher primary address will be elected as DR.

Example

```
configure pim ipv4 vlan accounting dr-priority 10
```



History

This command was first available in ExtremeXOS 15.3.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure pim iproute sharing hash

```
configure pim {ipv4 | ipv6} iproute sharing hash [source | group | source-group | source-group-nexthop]
```

Description

This command is used to configure the PIM ECMP hash algorithm.

Syntax Description

IPv4	Specifies the IPv4 address family.
IPv6	Specifies the IPv6 address family.
hash	Configure Hash Algorithm for Equal Cost Multipath Routing.
source	Hash for route sharing is based on source address only.
group	Hash for route sharing is based on group address only.
source-group	Hash for route sharing is based on source and group addresses.
source-group-nexthop	Hash for route sharing is based on source, group, and next hop addresses (default).

Default

Source-group-nexthop.

Usage Guidelines

Use this command to configure the PIM ECMP hash algorithm.

Example

The following command configures the PIM ECMP hash algorithm based on source-group-nexthop:

```
configure pim ipv6 iproute sharing hash source-group-nexthop
```



History

This command was first available in ExtremeXOS 15.3.2.

Platform Availability

This command is available on platforms that support the appropriate license. All platforms except Summit X440 support IP route sharing in the ExtremeXOS 15.3.2 release. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure pim register-policy

```
configure pim {ipv4 | ipv6} register-policy [rp_policy_name | none]
```

Description

Configures the register filter at the First Hop Router (FHR).

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
<i>rp_policy_name</i>	Specifies the Policy File for Register filter.
none	Unconfigures the configured RP Register filter.

Default

IPv4 is the default value.

Usage Guidelines

None.

Example

The following command configures an IPv4 register policy named "entry_policy" at the FHR:

```
configure pim ipv4 register-policy entry_policy
```

History

This command was first available in ExtremeXOS 15.3.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure pim register-policy rp

```
configure pim {ipv4 | ipv6} register-policy rp [rp_policy_name | none]
```

Description

Configures the register filter at the Rendezvous Point .

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
<i>rp_policy_name</i>	Specifies the Policy File for RP Register filter.
none	Unconfigures the configured RP Register filter.

Default

N/A.

Usage Guidelines

None.

Example

The following command configures IPv4 register policy named "entry_policy":

```
configure pim ipv4 register-policy rp entry_policy
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure pim register-rate-limit-interval

```
configure pim {ipv4 | ipv6} register-rate-limit-interval interval
```

Description

Configures the initial PIM-SM periodic register rate.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
<i>interval</i>	Specifies an interval time in seconds. Range is 0 - 60. Default is 0.

Default

Default is 0.

Usage Guidelines

Configuring a non-zero interval time can reduce the CPU load on the first hop switch, in case register stop messages are not received normally.

When a non-zero value is configured, the first hop switch sends a few register messages and then waits for a corresponding register stop from the RP for *time* seconds. The process is repeated until the register stop is received. This command should be used when the (S,G) tree between the first hop router and the RP is not converging quickly.

The default value is zero in default mode, the switch sends continuous register messages until the register stop is received.

Example

The following command configures the initial PIM register rate limit interval:

```
configure pim register-rate-limit-interval 2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure pim register-suppress-interval register-probe-interval

```
configure pim {ipv4 | ipv6} register-suppress-interval reg-interval register-probe-interval probe_interval
```

Description

Configures an interval for periodically sending null-registers.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
<i>reg-interval</i>	Specifies an interval time in seconds. Range is 30 - 200 seconds. Default is 60.
<i>probe-interval</i>	Specifies an interval time in seconds. Default is 5.

Default

The following defaults apply:

- register-suppress-interval—60
- register-probe-interval—5

Usage Guidelines

The register-probe-interval time should be set less than the register-suppress-interval time. By default, a null register is sent every 55 seconds (register-suppress-interval - register-probe-interval). A response to the null register is expected within register probe interval. By specifying a larger interval, a CPU peak load can be avoided because the null-registers are generated less frequently. The register probe time should be less than half of the register suppress time, for best results.

Example

The following command configures the register suppress interval and register probe time:

```
configure pim register-suppress-interval 90 register-probe time 10
```

History

This command was first available in ExtremeXOS 10.1.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure pim register-checksum-to

```
configure pim {ipv4 | ipv6} register-checksum-to [include-data | exclude-data]
```

Description

Configures the checksum computation to either include data (for compatibility with Cisco Systems products) or to exclude data (for RFC-compliant operation), in the register message.

Syntax Description

IPv4	Specifies the IPv4 address family.
IPv6	Specifies the IPv6 address family.
include-data	Specifies to include data.
exclude-data	Specifies to exclude data.

Default

Exclude data.

Usage Guidelines

None.

Example

The following command configures the checksum mode to include data for compatibility with Cisco Systems products:

```
configure pim register-checksum-to include-data
```

History

This command was first available in ExtremeXOS 10.1.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure pim shutdown-priority

```
configure pim {ipv4 | ipv6} [ {vlan} vlan_name | vlan all ] shutdown-priority
number
```

Description

Configures the priority for out of memory shutdown.

Syntax Description

ipv4	Configures a PIM timer on IPv4 router interfaces.
ipv6	Configures a PIM timer on IPv6 router interfaces.
<i>vlan_name</i>	Specifies a VLAN name.
all	Specifies all VLANs.
<i>number</i>	Priority for VLAN range is [0 - 65535].

Default

IPv4.

Usage Guidelines

None.

Example

The following command configures the shutdown priority for VLAN 36:

```
config pim vlan v36 shutdown-priority 22
```

History

This command was first available in ExtremeXOS 12.4.

The **ipv4** and **ipv6** keywords were added giving an option to support this functionality in IPv6 as well, in ExtremeXOS 15.3.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure pim spt-threshold

```
configure pim {ipv4 | ipv6} spt-threshold leaf-threshold {rp_threshold}
```

Description

Configures the threshold, in kbps, for switching to SPT. On leaf routers, this setting is based on data packets. On the RP, this setting is based on register packets.

Syntax Description

IPv4	Specifies the IPv4 address family.
IPv6	Specifies the IPv6 address family.
<i>leaf-threshold</i>	Specifies the rate of traffic per (s,g,v) group in kbps for the last hop. Range is 0 - 419403.
<i>rp_threshold</i>	Specifies an RP threshold. Range is 0 - 419403.

Default

The default setting is 0 for both parameters.

Usage Guidelines

For the best performance, use default value of 0.

Example

The following command sets the threshold for switching to SPT:

```
configure pim spt-threshold 4 16
```

History

This command was first available in ExtremeXOS 10.1.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure pim ssm range

```
configure pim {ipv4 | ipv6} ssm range [default | policy policy-name]
```

Description

Configures the range of multicast addresses for PIM SSM.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
default	Specifies the default address range, 232.0.0.0/8.
<i>policy-name</i>	Specifies a policy that defines the SSM address range.

Default

By default, no SSM range is configured. Using this command with the default keyword sets the range to 232.0.0.0/8. To reset the switch to the initial state, use the `unconfigure pim ssm range` command.

Usage Guidelines

You must disable PIM before configuring or unconfiguring a PIM-SSM range. Use the `disable pim` command.

Initially, no range is configured for SSM. After a range is configured, you can remove the range with the `unconfigure pim ssm range` command. If you wish to change the PIM SSM range, you must first unconfigure the existing range, and then configure the new range.

SSM requires that hosts use IGMPv3 messages to register to receive multicast group packets. When a range is configured for SSM, any IGMPv2 messages for an address in the range are ignored. Also, any IGMPv3 Exclude messages are ignored.

Note



If a PIM-SSM range is configured, IGMPv2 messages and IGMPv3 exclude messages within the PIM-SSM range are ignored on all IP interfaces, whether or not PIM-SSM is configured on the interfaces.



To specify a range different from the default PIM SSM range, create a policy file. The match statement of the policy file contains the group addresses to be treated as PIM SSM addresses. For example, to specify the PIM SSM address range as 232.0.0.0/8 and 233.0.0.0/8, use the following policy file:

```
Entry extremel {
  if match any {
    nlri 232.0.0.0/8 ;
    nlri 233.0.0.0/8 ;
  }
  then {
    permit ;
  }
}
```

Example

The following command sets the PIM SSM range to 232.0.0.0/8 and 233.0.0.0/8, if the policy file `ssmrange.pol` contains the policy example used above:

```
configure pim ssm range policy ssmrange.pol
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure pim state-refresh

```
configure pim {ipv4 | ipv6} state-refresh {vlan} [vlan_name | all] [on | off]
```

Description

Enables or disables the PIM-DM state refresh feature on one or all VLANs.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
<i>vlan_name</i>	Specifies a VLAN on which to enable or disable the PIM-DM state refresh feature.



on	Enables the PIM-DM state refresh feature on the specified VLANs.
off	Disables the PIM-DM state refresh feature on the specified VLANs.

Default

Disabled.

Usage Guidelines

When this feature is disabled on an interface, the interface behaves as follows:

- State refresh messages are not originated.
- State refresh messages received on the interface are dropped without processing.
- State refresh messages received on other interfaces are not forwarded to the disabled interface.

Example

The following command enables the PIM-DM state refresh feature on VLAN blue:

```
configure pim state-refresh blue on
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure pim state-refresh timer origination-interval

```
configure pim state-refresh timer origination-interval interval
```

Description

Configures the interval at which state refresh messages are originated.

Syntax Description

<i>interval</i>	Specifies a refresh interval in seconds. The range is 30 to 90 seconds.
-----------------	---



Default

60 seconds.

Usage Guidelines

None.

Example

The following command configures the interval to 45 seconds:

```
configure pim state-refresh timer origination-interval 45
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure pim state-refresh timer source-active-timer

```
configure pim {ipv4 | ipv6} state-refresh timer source-active-timer interval
```

Description

Defines how long a multicast source (S,G) is considered active after a packet is received from the source.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
<i>interval</i>	Specifies a source-active timer interval in seconds. The range is 90 to 300 seconds.

Default

210 seconds.



Usage Guidelines

None.

Example

The following command configures the interval to 45 seconds:

```
configure pim state-refresh timer source-active-timer 180
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure pim state-refresh ttl

```
configure pim {ipv4 | ipv6} state-refresh ttl ttlvalue
```

Description

Configures a time-to-live (TTL) value for PIM-DM state refresh messages.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
<i>ttl_value</i>	Specifies a TTL value. The range is 1 to 64.

Default

16.

Usage Guidelines

None.



Example

The following command the TTL value for 24:

```
configure pim state-refresh ttl 24
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure pim timer vlan

```
configure pim {ipv4 | ipv6} timer hello_interval jp_interval [{vlan} vlan_name |
vlan all]
```

Description

Configures the global PIM timers on the specified router interfaces.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
<i>hello_interval</i>	Specifies the amount of time before a hello message is sent out by the PIM router. The range is 1 to 65,535 seconds.
<i>jp_interval</i>	Specifies the join/prune interval. The range is 1 to 65,535 seconds.
<i>vlan_name</i>	Specifies a VLAN name.
all	Specifies all VLANs.

Default

- hello_interval—30 seconds.
- jp_interval—60 seconds.

Usage Guidelines

These default timers should only be adjusted when excess PIM control packets are observed on the interface.



Example

The following command configures the PIM timers on the VLAN accounting:

```
configure pim timer 150 300 vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure pim vlan trusted-gateway

```
configure pim {ipv4 | ipv6} [{vlan} vlan_name | vlan all ] trusted-gateway
[policy | none]
```

Description

Configures a trusted neighbor policy.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
<i>vlan_name</i>	Specifies a VLAN name.
vlan all	Specifies all VLANs.
<i>policy</i>	Specifies an policy file name.
none	Specifies no policy file, so all gateways are trusted.

Default

No policy file, so all gateways are trusted.

Usage Guidelines

Because PIM leverages the unicast routing capability that is already present in the switch, the access policy capabilities are, by nature, different. When the PIM protocol is used for routing IP multicast traffic, the switch can be configured to use a policy file to determine trusted PIM router neighbors for the VLAN on the switch running PIM. This is a security feature for the PIM interface.



Example

The following command configures a trusted neighbor policy on the VLAN backbone:

```
configure pim vlan backbone trusted-gateway nointernet
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable igmp

```
disable igmp {vlan name}
```

Description

Disables IGMP on a router interface. If no VLAN is specified, IGMP is disabled on all router interfaces.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

Enabled.

Usage Guidelines

IGMP is a protocol used by an IP host to register its IP multicast group membership with a router. Periodically, the router queries the multicast group to see if the group is still in use. If the group is still active, hosts respond to the query, and group registration is maintained.

IGMP is enabled by default on the switch. However, the switch can be configured to disable the generation and processing of IGMP packets. IGMP should be enabled when the switch is configured to perform IP multicast routing.

This command disables IGMPv2 and IGMPv3.



Example

The following command disables IGMP on VLAN accounting:

```
disable igmp vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable igmp snooping

```
disable igmp snooping {forward-mcrouter-only | with-proxy | vlan name}
```

Description

Disables IGMP snooping.

Syntax Description

forward-mcrouter-only	Specifies that the switch forwards all multicast traffic to the multicast router only.
with-proxy	Disables the IGMP snooping proxy.
<i>name</i>	Specifies a VLAN.

Default

IGMP snooping and the with-proxy option are enabled by default, but forward-mcrouter-only option is disabled by default.

Usage Guidelines

If a VLAN is specified, IGMP snooping is disabled only on that VLAN, otherwise IGMP snooping is disabled on all VLANs.

This command applies to both IGMPv2 and IGMPv3.

If the switch is in the forward-mcrouter-only mode, then the command `disable igmp snooping forward-mcrouter-only` changes the mode so that all multicast traffic is forwarded to any IP router. If not in the forward-mcrouter-mode, the command `disable igmp snooping forward-mcrouter-only` has no effect.



To change the snooping mode you must disable IP multicast forwarding. Use the command:

```
disable ipmcforwarding
```

The with-proxy option can be used for troubleshooting purpose. It should be enabled for normal network operation.

Enabling the proxy allows the switch to suppress the duplicate join requests on a group to forward to the connected Layer 3 switch. The proxy also suppresses unnecessary IGMP leave messages so that they are forwarded only when the last member leaves the group.

Example

The following command disables IGMP snooping on the VLAN accounting:

```
disable igmp snooping accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable igmp snooping vlan fast-leave

```
disable igmp snooping {vlan} name fast-leave
```

Description

Disables the IGMP snooping fast leave feature on the specified VLAN.

Syntax Description

<i>name</i>	Specifies a VLAN.
-------------	-------------------

Default

Disabled.

Usage Guidelines

None.



Example

The following command disables the IGMP snooping fast leave feature on the default VLAN:

```
disable igmp snooping "Default" fast-leave
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable igmp ssm-map

```
disable igmp ssm-map {vr vr-name}
```

Description

Disables IGMP SSM mapping.

Syntax Description

<i>vr-name</i>	Specifies a virtual router name. If the VR name is omitted, the switch disables mapping on the VR specified by the current CLI VR context.
----------------	--

Default

Disabled on all interfaces.

Usage Guidelines

None.

Example

The following command disables IGMP-SSM mapping on the VR in the current CLI VR context:

```
disable igmp ssm-map
```



History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable ipmcforwarding

```
disable ipmcforwarding {vlan name}
```

Description

Disables IP multicast forwarding on a router interface.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

Disabled.

Usage Guidelines

If no options are specified, all configured IP interfaces are affected. When new IP interfaces are added, IP multicast forwarding is disabled by default.

IP forwarding must be enabled before enabling IP multicast forwarding.

Disabling IP multicast forwarding disables any Layer 3 multicast routing for the streams coming to the interface.

Example

The following command disables IP multicast forwarding on the VLAN accounting:

```
disable ipmcforwarding vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable mvr

disable mvr

Description

Disables MVR on the system.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables MVR on the system:

```
disable mvr
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MVR feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable pim

disable pim {ipv4 | ipv6}



Description

Disables PIM on the system.

Syntax Description

IPv4	Specifies the IPv4 address family.
IPv6	Specifies the IPv6 address family.
priority	Specifies a priority setting. The range is 0 - 254.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables PIM on the system:

```
disable pim
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable pim iproute sharing

```
disable pim {ipv4 | ipv6} iproute sharing
```

Description

Disables the PIM Equal Cost Multi Path (ECMP) feature.



Syntax Description

iproute	IP Route
sharing	Equal Cost Multipath Routing

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables the PIM ECMP feature:

```
disable pim ipv4 iproute sharing
```

History

This command was first available in ExtremeXOS 15.3.2.

Platform Availability

This command is available on platforms that support the appropriate license. All platforms except Summit X440 support IP route sharing in the ExtremeXOS 15.3.2 release. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable pim snooping

```
disable pim snooping {{vlan} name}
```

Description

Disables PIM snooping and clears all the snooping PIM neighbors, joins received on the VLAN, and the forwarding entries belonging to one or all VLANs.

Syntax Description

<i>name</i>	Specifies a VLAN.
-------------	-------------------



Default

Disabled.

Usage Guidelines

None.

Example

The following command disables PIM snooping for all VLANs on the switch:

```
disable pim snooping
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable pim ssm vlan

```
disable pim {ipv4 | ipv6} ssm vlan [vlan_name | all]
```

Description

Disables PIM SSM on a router interface.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
<i>vlan_name</i>	Specifies a VLAN name.
all	Specifies all VLANs.

Default

Disabled on all interfaces.



Usage Guidelines

This command disables PIM-SSM on the specified Layer 3 VLAN.

IGMPv3 include messages for multicast addresses in the SSM range is only processed by PIM if PIM-SSM is enabled on the interface. Any non-IGMPv3 messages in the SSM range are not processed by PIM on any switch interface, whether SSM is enabled or not.

Example

The following command disables PIM-SSM multicast routing on VLAN accounting:

```
disable pim ssm vlan accounting
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable igmp

```
enable igmp {vlan vlan name } {IGMPv1 | IGMPv2 | IGMPv3}
```

Description

Enables IGMP on a router interface. If no VLAN is specified, IGMP is enabled on all router interfaces.

Syntax Description

<i>vlan name</i>	Specifies a VLAN name.
IGMPv1	Specifies the compatibility mode as IGMPv1.
IGMPv2	Specifies the compatibility mode as IGMPv2.
IGMPv3	Specifies the compatibility mode as IGMPv3.

Default

Enabled, set to IGMPv2 compatibility mode.



Usage Guidelines

IGMP is a protocol used by an IP host to register its IP multicast group membership with a router. Periodically, the router queries the multicast group to see if the group is still in use. If the group is still active, IP hosts respond to the query, and group registration is maintained.

IGMPv2 is enabled by default on the switch. However, the switch can be configured to disable the generation and processing of IGMP packets. IGMP should be enabled when the switch is configured to perform IP multicast routing.

Example

The following command enables IGMPv2 on the VLAN accounting:

```
enable igmp vlan accounting
```

The following command enables IGMPv3 on the VLAN finance:

```
enable igmp vlan finance igmpv3
```

History

This command was first available in ExtremeXOS 10.1.

The IGMPv1, IGMPv2, and IGMPv3 options were added in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable igmp snooping

```
enable igmp snooping {forward-mcrouter-only | {vlan} name | with-proxy vr vrname}
```

Description

Enables IGMP snooping on one or all VLANs.



Syntax Description

forward-mcrouter-only	Specifies that the switch forward all multicast traffic to the multicast router only.
<i>name</i>	Specifies a VLAN or VMAN on which to enable IGMP snooping.
with-proxy vr <i>vrname</i>	Controls how join and leave messages are forwarded from the specified virtual router. If this option is specified, one join message per query is forwarded, and a leave message is forwarded only if it is from the last receiver on the VLAN.

Default

Enabled.

Usage Guidelines

This command applies to both IGMPv2 and IGMPv3.

IGMP snooping is enabled by default on the switch. If you are using multicast routing, IGMP snooping can be enabled or disabled. If IGMP snooping is disabled, all IGMP and IP multicast traffic floods within a given VLAN or VMAN.

The forward-mcrouter-only, vlan, and with-proxy options control three separate and independent features. You can manage one feature at a time with the enable igmp snooping command, and you can enter the command multiple times as needed to control each feature. For example, you can enter the command twice to enable both the forward-mcrouter-only and with-proxy options.

If a VLAN or VMAN is specified with the enable igmp snooping command, IGMP snooping is enabled only on that VLAN or VMAN. If no options are specified, IGMP snooping is enabled on all VLANs.



Note

IGMP snooping is not supported on SVLANs on any platform.

The with-proxy option enables the IGMP snooping proxy feature, which reduces the number of join and leave messages forwarded on the virtual router as described in the table above. This feature is enabled by default.

An optional optimization for IGMP snooping is the strict recognition of routers only if the remote devices are running a multicast protocol. Two IGMP snooping modes are supported:

- The forward-mcrouter-only mode forwards all multicast traffic to the multicast router (that is, the router running PIM, DVMRP or CBT).



- When not in the forward-mcrouter-only mode, the switch forwards all multicast traffic to any IP router (multicast or not), and any active member port to the local network that has one or more subscribers.

Note



The forward-mcrouter-only mode for IGMP snooping is enabled/disabled on a switch-wide basis, not on a per-VLAN basis. In other words, all the interfaces enabled for IGMP snooping are either in the forward-mcrouter-only mode or in the non-forward-mcrouter-only mode, and not a mixture of the two modes.

To change the snooping mode you must disable IP multicast forwarding. To disable IP multicast forwarding, use the command:

```
disable ipmcforwarding {vlan name}
```

To change the IGMP snooping mode from the non-forward-mcrouter-only mode to the forward-mcrouter-only mode, use the commands:

```
disable ipmcforwarding enable igmp snooping forward-mcrouter-only enable ipmcforwarding {vlan name}
```

To change the IGMP snooping mode from the forward-mcrouter-only mode to the non-forward-mcrouter-only mode, use the commands:

```
disable ipmcforwarding disable igmp snooping forward-mcrouter-only enable ipmcforwarding {vlan name}
```

Example

The following command enables IGMP snooping on the switch:

```
enable igmp snooping
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable igmp snooping vlan fast-leave

```
enable igmp snooping {vlan} name fast-leave
```



Description

Enables the IGMP snooping fast leave feature on the specified VLAN.

Syntax Description

<i>name</i>	Specifies a VLAN.
-------------	-------------------

Default

Disabled.

Usage Guidelines

The fast leave feature operates only with IGMPv2.

To view the fast leave feature configuration, use the `show configuration msmgr` command. This `show` command displays the fast leave configuration only when the feature is enabled.

Example

The following command enables the IGMP snooping fast leave feature on the default VLAN:

```
enable igmp snooping "Default" fast-leave
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable igmp snooping with-proxy

```
enable igmp snooping with-proxy
```

Description

Enables the IGMP snooping proxy. The default setting is enabled.

Syntax Description

This command has no arguments or variables.



Default

Enabled.

Usage Guidelines

Enabling the proxy allows the switch to suppress the duplicate join requests on a group to forward to the connected Layer 3 switch. The proxy also suppresses unnecessary IGMP leave messages so that they are forwarded only when the last member leaves the group.

This command can be used for troubleshooting purpose. It should be enabled for normal network operation. The command does not alter the snooping setting.

This feature can be enabled when IGMPv3 is enabled; however, it is not effective for IGMPv3.

Example

The following command enables the IGMP snooping proxy:

```
enable igmp snooping with-proxy
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable igmp ssm-map

```
enable igmp ssm-map {vr vr-name}
```

Description

Enables IGMP SSM mapping on a VR.

Syntax Description

<i>vr-name</i>	Specifies a virtual router name. If the VR name is omitted, the switch uses the VR specified by the current CLI VR context.
----------------	---

Default

Disabled on all interfaces.



Usage Guidelines

Configure the range of multicast addresses for PIM SSM before you enable IGMP SSM mapping. IGMP SSM mapping operates only with IPv4.

Example

The following command enables IGMP-SSM mapping on the VR in the current CLI VR context:

```
enable igmp ssm-map
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable ipmcforwarding

```
enable ipmcforwarding {vlan name}
```

Description

Enables IP multicast forwarding on an IP interface.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

Disabled.

Usage Guidelines

If no options are specified, all configured IP interfaces are affected. When new IP interfaces are added, IPMC forwarding is disabled by default.

IP forwarding must be enabled before enabling IPMC forwarding.



Example

The following command enables IPMC forwarding on the VLAN accounting:

```
enable ipmcforwarding vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable mvr

enable mvr

Description

Enables MVR on the system.

Syntax Descripton

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables MVR on the system:

```
enable mvr
```

History

This command was first available in ExtremeXOS 11.4.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MVR feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable pim

```
enable pim {ipv4 | ipv6}
```

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables PIM on the system:

```
enable pim
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



enable pim iproute sharing

```
enable pim {ipv4 | ipv6} iproute sharing
```

Description

Enables the PIM Equal Cost Multi Path (ECMP) feature.

Syntax Description

iproute	IP Route.
sharing	Equal Cost Multipath Routing.

Default

Disabled.

Usage Guidelines

Use this feature to allow downstream PIM router to choose multiple ECMP path to source via hash from one of the following selections without affecting the existing unicast routing algorithm:

- Source
- Group
- Source-Group
- Source-Group-Next-Hop

This feature does load splitting, not load balancing and operates on a per (S, G) and (*;G) basis splitting the load onto the available equal cost paths by hashing according to the selection criteria defined by the user.

Make sure that IP route sharing is also enabled using `enable iproute {ipv4| ipv6} sharing`.

Example

The following command enables the PIM ECMP feature:

```
enable pim ipv4 iproute sharing
```

History

This command was first available in ExtremeXOS 15.3.2.

Platform Availability

This command is available on platforms that support the appropriate license. All platforms except Summit X440 support IP route sharing in the ExtremeXOS 15.3.2 release. For complete information



about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable pim snooping

```
enable pim snooping {{vlan} name}
```

Description

Enables PIM snooping on one or all VLANs.

Syntax Description

<i>name</i>	Specifies a VLAN.
-------------	-------------------

Default

Disabled.

Usage Guidelines

PIM snooping does not require PIM to be enabled. However, IGMP snooping must be disabled on VLANs that use PIM snooping. PIM snooping and MVR cannot be enabled simultaneously on a switch. PIM snooping should not be enabled on a VLAN that supports PIM-DM neighbors.

Example

The following command enables PIM snooping on the default VLAN:

```
enable pim snooping default
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable pim ssm vlan

```
enable pim {ipv4 | ipv6} ssm vlan [vlan_name | all]
```



Description

Enables PIM SSM on an IP interface.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
<i>vlan_name</i>	Specifies a VLAN name.
all	Specifies all VLANs.

Default

Disabled on all interfaces.

Usage Guidelines

This command enables PIM-SSM on the specified Layer 3 VLAN.

PIM-SM must also be configured on the interface for PIM to begin operating (which includes enabling IP multicast forwarding).

IGMPv3 include messages for multicast addresses in the SSM range are only processed by PIM if PIM-SSM is enabled on the interface. Any non-IGMPv3 include messages in the SSM range are not processed by PIM on any switch interface, whether SSM is enabled or not.

Example

The following command enables PIM-SSM multicast routing on VLAN accounting:

```
enable pim ssm vlan accounting
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

mrinfo

```
mrinfo {router_address} {from from_address} {timeout seconds} {multiple-response-timeout multi_resp_timeout} {vr vname}
```



Description

Requests information from a multicast router.

Syntax Description

<i>router_address</i>	Specifies the unicast IP address of the router for which you want information.
<i>from_address</i>	Specifies the unicast IP address of the interface where the mrimfo request is generated.
<i>seconds</i>	Specifies a maximum time to wait for a response. The range is 1 to 30 seconds.
<i>multi_resp_timeout</i>	Specifies a maximum time to wait for additional responses after the first response is received. The range is 0 to 3 seconds.
<i>vrname</i>	Specifies a VR name.

Default

router_address: One of the local interface addresses.

from: IP address of interface from which the mrimfo query is generated.

timeout: 3 seconds

multiple-response-timeout: 1 second

vr: DefaultVR

Usage Guidelines

The last column of the mrimfo command output displays information in the following format:

[Metric/threshold/type/flags]

This information is described in [detail here](#).

Table 56: mrimfo Command Display Data

Data	Description
Metric	This should always be 1 because mrimfo queries the directly connected interfaces of a device.
Threshold	This should always be 0 because the threshold feature is not supported in ExtremeXOS software.
Type	The type specifies the multicast protocol type. Because the ExtremeXOS software only supports PIM, this value is always pim.
querier	The querier flag indicates that the queried router is the IGMP querier.
leaf	The leaf flag indicates that the IP interface has no neighbor router.
down	The down flag indicates that the interface link status is down.



Example

The following command requests information from multicast router 1.1.1.1:

```
Switch.1 # mrimfo 1.1.1.1
1.1.1.1 [Flags:PGM]
2.2.2.1 -> 2.2.2.2 [1/0/pim/querier]
1.1.1.1 -> 0.0.0.0 [1/0/pim/querier/leaf]
8.8.8.1 -> 8.8.8.4 [1/0/pim/querier]
3.3.3.1 -> 0.0.0.0 [1/0/pim/down]
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

mtrace

```
mtrace source src_address {destination dest_address} {group grp_address} {from
from_address} {gateway gw_address} {timeout seconds} {maximum-hops number}
{router-alert [include | exclude]} {vr vrname}
```

Description

Traces multicast traffic from the receiver back to the source.

Syntax Description

<i>src_address</i>	Specifies the unicast IP address of the multicast source.
<i>dest_address</i>	Specifies the unicast IP address of the multicast group receiver.
<i>grp_address</i>	Specifies the multicast IP address of the group.
<i>from_address</i>	Specifies the unicast IP address of the interface where the mtrace request originates. This is used as the IP destination address of the mtrace response packet.
<i>gw_address</i>	Specifies the gateway router IP address of the multicast router to which the unicast mtrace query is sent.
<i>seconds</i>	Specifies a maximum time to wait for the mtrace response before making the next attempt. The range is 1 to 30 seconds.
<i>number</i>	Specifies the maximum number of hops for the trace. The range is 1 to 255.



router-alert	Specifies whether the router-alert option is included or excluded in mtrace packets.
<i>vrname</i>	Specifies a VR name.

Default

destination: IP address of interface from which mtrace query is generated.

group: 0.0.0.0

from: IP address of interface from which mtrace query is generated.

gateway: 224.0.0.2 when the destination is in the same subnet as one of the IP interfaces. For a non-local destination address, it is mandatory to provide a valid multicast router address.

timeout: 3 seconds

maximum-hops: 32

router-alert: include

vr: DefaultVR

Usage Guidelines

The multicast traceroute initiator node generates a multicast query and waits for timeout period to expire. If there is no response for the timeout period, the initiator node makes 2 more attempts. If no response is received after 3 attempts, the initiator node moves to a hop-by-hop trace by manipulating the maximum hop fields to perform a linear search.

The multicast trace response data contains the following fields:

- Incoming interface address—Interface on which traffic is expected from the specific source and group
- Outgoing interface address—Interface on which traffic is forwarded from the specified source and group towards the destination
- Previous hop router address
- Input packet count on incoming interface
- Output packet count on outgoing interface
- Total number of packets for this source-group pair
- Multicast routing protocol
- Forwarding code

Extreme Networks switches set the packet count statistics field to 0xffffffff to indicate that this field is not supported.

The last column of the mtrace command output displays forwarding codes, which are described in the following table.



Table 57: mtrace Command Forwarding Codes

Forwarding Code	Description
Wrong interface	mtrace request arrived on an interface to which this router would not forward for this source and group.
Prune sent upstream	This router has sent a prune request upstream for the source and group in the mtrace request.
Output pruned	This router has stopped forwarding for this source and group in response to a prune request from the next hop router.
Hit scope boundary ¹⁴	The group is subject to administrative scoping at this hop.
No route	This router has no route for the source or group and no way to determine a potential route.
Wrong Last Hop	This router is not the proper last-hop router.
Not forwarding	This router is not forwarding for this source and group on the outgoing interface for an unspecified reason.
Reached RP/Core	Reached rendezvous point or core.
RPF Interface	mtrace request arrived on the expected RPF interface (upstream interface) for this source and group.
Multicast disabled	mtrace request arrived on an interface which is not enabled for multicast.
Info. Hidden	One or more hops have been hidden from this trace.
No space in packet	There was not enough room to insert another response data block in the packet.
Next router no mtrace	The previous hop router does not understand mtrace requests.
Admin. Prohibited	mtrace is administratively prohibited.

Example

The following command initiates an mtrace for group 225.1.1.1 at IP address 1.1.1.100:

```
Switch.6 # mtrace source 1.1.1.100 group 225.1.1.1
Mtrace from 1.1.1.100 to Self via 225.1.1.1
0          34.2.2.4
-1         34.2.2.4 PIM thresh^ 0          1.1.1.100/32 RPF Interface
-2         34.2.2.3 PIM thresh^ 0          1.1.1.100/32
-3         23.1.1.2 PIM thresh^ 0          1.1.1.100/32
-4         2.2.2.1 PIM thresh^ 0          1.1.1.100/32
Round trip time 9 ms; total ttl of 4 required.
```

History

This command was first available in ExtremeXOS 12.4.

The router-alert option was added in ExtremeXOS 12.5.3.

¹⁴ ExtremeXOS switches along the mtrace path do not provide this forwarding code.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

rtlookup

```
rtlookup [ipv4_address | ipv6_address] { unicast | multicast | rpf } {vr vr_name}
```

Description

Displays the available routes to the specified IP address.

Syntax Description

<i>ipv4_address</i>	Specifies an IPv4 address.
<i>ipv6_address</i>	Specifies an IPv6 address.
unicast	Displays the routes from the unicast routing table in the current router context.
multicast	Displays the routes from the multicast routing table in the current router context.
rpf	Displays the RPF route to the specified destination.
<i>vr_name</i>	Specifies the virtual router for which to display the route.

Default

vr_name is the VR of the current CLI context.

When no option (unicast or multicast) is provided, this command displays the route in the unicast routing table.

Usage Guidelines

None.

Example

The following example displays the route lookup for 12.1.20.12 in the multicast routing table for the default VR:

```
BD30K # rtlookup 12.1.20.12 multicast vr vr-default
@mbe 12.1.0.0/16      50.1.10.21      1      UG---S--m--- toronto      0d:0h:
41m:1s
Origin(Ori): (b) BlackHole, (be) EBGp, (bg) BGP, (bi) IBGP, (bo) BOOTP
(ct) CBT, (d) Direct, (df) DownIF, (dv) DVMRP, (e1) ISISL1Ext
(e2) ISISL2Ext, (h) Hardcoded, (i) ICMP, (i1) ISISL1 (i2) ISISL2
```



```
(mb) MBGP, (mbe) MBGPExt, (mbi) MBGPInter, (mp) MPLS Lsp
(mo) MOSPF (o) OSPF, (o1) OSPFExt1, (o2) OSPFExt2
(oa) OSPFIntra, (oe) OSPFAsExt, (or) OSPFInter, (pd) PIM-DM, (ps) PIM-SM
(r) RIP, (ra) RtAdvrt, (s) Static, (sv) SLB_VIP, (un) UnKnown
(*) Preferred unicast route (@) Preferred multicast route
(#) Preferred unicast and multicast route
Flags: (B) BlackHole, (D) Dynamic, (G) Gateway, (H) Host Route
(L) Matching LDP LSP, (l) Calculated LDP LSP, (m) Multicast
(P) LPM-routing, (R) Modified, (S) Static, (s) Static LSP
(T) Matching RSVP-TE LSP, (t) Calculated RSVP-TE LSP, (u) Unicast, (U) Up
(c) Compressed Route
Mariner # rtlookup 12.1.20.12 multicast vr vr-default
No route to 12.1.10.12
```

History

This command was first available in ExtremeXOS 10.1.

The xhostname option was removed in ExtremeXOS 11.0.

Support for IPv6 was added in ExtremeXOS 11.2.

The unicast and multicast options were added in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

rtlookup rpf

```
rtlookup [ipv4_address | ipv6_address] rpf {vr vr_name}
```

Description

Displays the RPF for a specified multicast source.

Syntax Description

<i>ipv4_address</i>	Specifies an IPv4 address.
<i>ipv6_address</i>	Specifies an IPv6 address.
rpf	Selects the RPF for the specified multicast source.
<i>vr_name</i>	Specifies the virtual router for which to display the route.

Default

vr_name is the VR of the current CLI context.



Usage Guidelines

None.

Example

The following example displays the RPF lookup for multicast source 12.1.20.12 in the default VR:

```
BD10K # rtlookup 12.1.20.12 rpf vr vr-default
Ori Prefix          Route          Gateway          VLAN
@d 12.1.10.22      12.1.10.0/24    12.1.10.10      v1
Origin(Ori): (b) BlackHole, (be) EBGP, (bg) BGP, (bi) IBGP, (bo) BOOTP
(ct) CBT, (d) Direct, (df) DownIF, (dv) DVMRP, (e1) ISISL1Ext
(e2) ISISL2Ext, (h) Hardcoded, (i) ICMP, (i1) ISISL1 (i2) ISISL2
(mb) MBGP, (mbe) MBGPExt, (mbi) MBGPInter, (mp) MPLS Lsp
(mo) MOSPF (o) OSPF, (o1) OSPFExt1, (o2) OSPFExt2
(oa) OSPFIntra, (oe) OSPFAsExt, (or)OSPFIInter, (pd) PIM-DM,(ps) PIM-SM
(r) RIP, (ra) RtAdvrt, (s) Static, (sv) SLB_VIP, (un) UnKnown
(*) unicast route (@) multicast route
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show igmp

```
show igmp {vlan} {vlanname}
```

Description

This command can be used to display an IGMP-related configuration and group information, per VLAN.

Syntax Description

<i>vlanname</i>	Specifies a VLAN name.
-----------------	------------------------

Default

N/A.



Usage Guidelines

The output of this command shows:

- The VLAN name.
- The router interface IP address and subnet mask.
- If the interface is active (up), by the letter U.
- If IP forwarding is enabled for the interface, by the letter f.
- If multicast forwarding is enabled, by the letter M.
- If IGMP is enabled, by the letter i.
- If IGMP snooping is enabled, by the letter z.

Example

The following command displays the IGMP configuration:

```
show igmp
VLAN                IP Address          Flags      nLRMA  nLeMA  IGMPver
Default             0.0.0.0            / 0      ---izpt-    0      0      3
isc                 50.50.50.1        /24      ---izpt-    0      0      3
v1                  0.0.0.0            / 0      U--izpt-    0      2      3
v3000               1.1.1.1            /24      ---izpt-    0      0      3
v666                6.0.0.1            /16      ---izpt-    0      0      3
Flags: (f) Forwarding Enabled, i) IGMP Enabled
(m) Multicast Forwarding Enabled, (p) IGMP Proxy Query Enabled
(r) Receive Router Alert Required (t) Transmit Router Alert
(U) Interface Up, (z) IGMP Snooping Enabled
(nLeMA) Number of Learned Multicast Addresses
(nLRMA) Number of Locally Registered Multicast Addresses
```

The following command displays the IGMP configuration for VLAN vlan1:

```
show igmp vlan1
Query Interval      :    125 sec
Max Response Time  :     10 sec
Last Member Query   :      1 sec
Robustness          :      2
Interface on VLAN vlan1 is enabled and up.
inet 0.0.0.0/0
Locally registered multicast addresses:
Learned multicast addresses (Last Querier=118.1.1.100):
224.0.0.2           224.0.0.22
s = static igmp member
Flags:
IP Fwding   NO      IPmc Fwding  NO      IGMP YES
IGMP Ver    V3      Snooping    YES     Proxy Query YES
XmitRtrAlrt YES     RcvRtrAlrtReq NO
```

History

This command was first available in ExtremeXOS 10.1.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show igmp group

```
show igmp group {{vlan} {name} | {grpipaddress}} {IGMPv3}
```

Description

Lists the IGMP group membership for the specified VLAN.

Syntax Description

<i>name</i>	Specifies a VLAN name.
<i>grpipaddress</i>	Specifies a group IP address.
IGMPv3	Display the IGMP group in IGMPv3 format (if group record is IGMPv3 compatible, otherwise display in earlier format).

Default

IGMPv2.

Usage Guidelines

If no VLAN is specified all VLANs are displayed. You can also filter the display by group address and by multicast stream sender address.

The output of this command shows:

- The multicast group address received.
- The version of the IGMP group.
- The name of the VLAN where the group address is being received.
- The physical port where the group address is being received. If multiple ports within the VLAN have subscribers for the group, all the ports are listed.
- The age since the last IGMP report for this group was received.



Note

The show igmp group command output is populated on the router that is the PIM Rendezvous Point.



Example

The following command lists the IGMP group membership:

```
show igmp group
```

The following is sample output from this command:

```
Group Address      Ver  Vlan          Port    Age
239.2.4.70        2   banana        7       101
224.0.1.24        2   banana        7       107
239.255.255.254   2   banana        7       103
Total: 3
```

History

This command was first available in ExtremeXOS 10.1.

The IGMPv3 option was added in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show igmp snooping

```
show igmp snooping {detail {IGMPv3}}
```

Description

Displays IGMP snooping registration information for all VLANs.

Syntax Description

detail	Displays the information in detailed format.
IGMPv3	Display the IGMP group in IGMPv3 format (if group record is IGMPv3 compatible, otherwise display in earlier format).

Default

IGMPv2.



Usage Guidelines

None.

Example

The following command displays IGMP snooping registration information for all VLANs:

```
show igmp snooping
Igmp Snooping Flag           : forward-all-router
Igmp Snooping Flood-list    : none
Igmp Snooping Proxy         : Disable
Igmp Snooping Filters       : per-port
Vlan          Vid  Port  #Senders #Receivers Router Enable
-----
Default      1    0    0         0         Yes
v1           4090 0    0         0         Yes
```

History

This command was first available in ExtremeXOS 10.1.

The IGMP Forwarding Lookup mode output was removed from this command in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show igmp snooping cache

This command is provided for backward compatibility. The recommended command is:

```
show mcast cache {{vlan} name} {{[group grpaddressMask | grpaddressMask] {source sourceIP | sourceIP}} {type [snooping | pim | mvr]}| {summary}}
```

The syntax for the original form of this command is:

```
show igmp snooping cache {{vlan} name} {{group}
grpaddressMask}
```

Description

Displays multicast cache entries added by IGMP snooping for all VLANs and groups. The display can be limited to specific VLANs or groups.



Syntax Description

<i>name</i>	Specifies a VLAN name.
<i>grpaddressMask</i>	Specifies a multicast group address and mask.

Default

Displays information for all VLANs and groups.

Usage Guidelines

None.

Example

The following command displays IGMP snooping cache information for all VLANs and groups:

```
BD-8808.2 # show igmp snooping cache
```

This command display is the same as for the following preferred command:

```
show mcast cache {{vlan} name} {[group grpaddressMask |grpaddressMask] {source
sourceIP |sourceIP}} {type [snooping | pim | mvr]}| {summary}}
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show igmp snooping vlan

```
show igmp snooping {vlan} name {port port} {IGMPv3}
```

Description

Displays IGMP snooping registration information for a specific VLAN. The display can be further limited to a specific port or to only IGMPv3 entries.



Syntax Description

<i>name</i>	Specifies a VLAN name.
<i>port</i>	Specifies a single port for which information is displayed.
IGMPv3	Display the IGMP group in IGMPv3 format (if group record is IGMPv3 compatible, otherwise display in earlier format).

Default

IGMPv2.

Usage Guidelines

The two types of IGMP snooping entry are sender entry and subscribed entry.

The following information is displayed in a per-interface format:

- Group membership information
- Router entry
- Timeout information
- Sender entry

Example

The following command displays IGMP snooping registration information on VLAN v1:

```
show igmp snooping vlan v1
Router Timeout           :    260 sec
Host Timeout            :    260 sec
Igmp Snooping Fast Leave Time :   1000 ms
VLAN v1 d                (4084) Snooping=Enabled
Port  Host              Subscribed      Age      Group-Limit
25    118.1.1.100         All Groups      3        0
```

The following command displays IGMP snooping registration information for port 2:1 on VLAN test:

```
show igmp snooping test port 2:1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



show igmp snooping vlan filter

```
show igmp snooping {vlan} name filter
```

Description

Displays IGMP snooping filters.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

None.

Usage Guidelines

Use this command to display IGMP snooping filters configured on the specified VLAN. When no VLAN is specified, all the filters are displayed.

Example

The following command displays the IGMP snooping filter configured on VLAN vlan101:

```
show igmp snooping vlan101 filter
Filter          Port Flags
igmppermit0    5:10 a
Flags: (a) Active
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show igmp snooping vlan static

```
show igmp snooping {vlan} name static [group | router]
```



Description

Displays static IGMP snooping entries.

Syntax Description

<i>name</i>	Specifies a VLAN name.
group	Displays static multicast groups.
router	Displays static router entries.

Default

None.

Usage Guidelines

Use this command to display the IGMP snooping static groups or router ports configured on the specified VLAN. When no VLAN is specified, all the static groups or router ports are displayed.

Example

The following command displays the IGMP snooping static groups configured on VLAN vlan101:

```
show igmp snooping vlan101 static group
VLAN vlan101 (4094)
Group          Port   Flags
239.1.1.2      29     s-
239.1.1.2      30     s-
239.1.1.2      31     sa
239.1.1.2      32     s-
239.1.1.2      34     s-
Total number of configured static IGMP groups = 5
Flags: (s) Static, (a) Active
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show igmp ssm-map

```
show igmp ssm-map {group_ip} {vr vr-name}
```



Description

Displays the IGMP SSM feature status (enabled or disabled), the mappings for the specified multicast group IP address, and the total count of maps.

Syntax Description

<i>group_ip</i>	Specifies an IP multicast group, for which all mappings in the PIM SSM range are to be displayed. If no group address is specified, the switch displays all IGMP-SSM mappings.
<i>vr-name</i>	Specifies a virtual router name. If the VR name is omitted, the switch displays the mappings on the VR specified by the current CLI VR context.

Default

N/A.

Usage Guidelines

When a target group is specified, this command displays all mapping entries for the configured range in which the group IP address resides.

Example

The following command displays the mappings for the multicast group IP address 232.1.1.2:

```
show igmp ssm-map 232.1.1.2
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show ipmroute

```
show ipmroute {source-net mask-len | source-net mask | summary} {vr vr-name}
```

Description

Displays the contents of the IP multicast routing table or the route origin priority.



Syntax Description

<i>source-net</i>	Specifies an IP address/mask length.
<i>mask-len</i>	Mask length for the IP multicast source's subnet. Range is [1-32].
<i>mask</i>	Specifies a subnet mask.
summary	Displays the statistics of multicast static routes.
<i>vr-name</i>	Specifies the virtual router to which the route is added.

Default

vr-name is the VR of the current CLI context.

Usage Guidelines

This command allows you to view the configured multicast static routes. You can specify the filtering criteria on this CLI to view only the desired route. The multicast static routes are displayed in ascending order of their prefix (same order as show iproute displays).

Example

The following example displays a multicast static route from a default virtual router:

```
* (debug) Summit-PC.19 # show ipmroute
Destination      Gateway          Mtr  Flags Protocol      VLAN
Default Route   20.20.20.1      255  UG   None          pc4-1
*1.1.0.0/16     20.20.20.1      10   UG   bgp           pc4-1
*11.0.0.0/8     30.30.30.1      12   U-   None          pc5-3
11.22.0.0/16    20.20.20.1      10   UG   None          pc4-1
*11.22.33.0/24  30.30.30.1      8    U-   None          pc5-3
11.22.33.44/32  20.20.20.1      4    UG   None          pc4-1
*12.0.0.0/8     20.20.20.1      0    UG   None          pc4-1
12.24.0.0/16    30.30.30.1      0    U-   None          pc5-3
*12.24.48.96/32 30.30.30.1      2    U-   ospf-extern1 pc5-3
44.66.0.0/16    30.30.30.1      0    U-   None          pc5-3
Flags: (*) Active, (G) Gateway, (U) Up
Mask distribution:
1 default routes          2 routes at length 8
4 routes at length 16     1 routes at length 24
2 routes at length 32
Total number of multicast static routes = 10
```

History

This command was first available in ExtremeXOS 11.6.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show iproute multicast

```
show iproute {ipv4} {{vlan} name | [ipaddress netmask | ipNetmask] | origin
[direct | static | mbgp | imbgp | embgp]} multicast {vr vr_name}
```

Description

Displays all or a filtered set of multicast routes in the IP multicast routing table.

Syntax Description

ipv4	Selects only IPv4 multicast routes.
name	Specifies a VLAN for which to display multicast routes.
ipaddress netmask	Specifies an IP address and network mask (in dotted decimal notation) for which to display multicast routes.
ipNetmask	Specifies the IP address and network mask in classless inter domain routing (CIDR) notation.
origin	Limits the displayed multicast routes to those generated by the specified origin. Origin options select direct routes, static routes, and routes created by the MBGP, IMBGP, and EMBGP protocols.
vr_name	Specifies the virtual router for which to display multicast routes.

Default

vr_name is the VR of the current CLI context.

Usage Guidelines

This command does not display unicast routes, which can be used for multicast traffic.

Example

The following example displays all the routes in multicast routing table:

```
BD10K # show iproute multicast
Ori  Destination      Gateway      Mtr  Flags          VLAN          Duration
@d   3.3.3.3/32         3.3.3.3      1    U-----m---  lpbk          12d:1h:
30m:36s
@d   28.0.0.0/24       28.0.0.15   1    U-----m---  trunk28       12d:1h:
30m:36s
@mbe 77.0.0.0/24        50.1.10.21  1    UG---S--m---  toronto       0d:0h:
```



```

41m:1s
@mbe 77.0.1.0/24          50.1.10.21      1      UG---S--m--- toronto  0d:0h:
41m:1s
@mbe 77.0.2.0/24          50.1.10.21      1      UG---S--m--- toronto  0d:0h:
41m:1s
@mbe 77.0.3.0/24          50.1.10.21      1      UG---S--m--- toronto  0d:0h:
41m:1s
@mbe 77.0.4.0/24          50.1.10.21      1      UG---S--m--- toronto  0d:0h:
41m:1s
@mbe 77.0.5.0/24          50.1.10.21      1      UG---S--m--- toronto  0d:0h:
41m:1s
@mbe 77.0.6.0/24          50.1.10.21      1      UG---S--m--- toronto  0d:0h:
41m:1s
@mbe 77.0.10.0/24         50.1.10.21      1      UG---S--m--- toronto  0d:0h:
41m:1s
@mbe 77.0.11.0/24         50.1.10.21      1      UG---S--m--- toronto  0d:0h:
41m:1s
@mbe 77.0.12.0/24         50.1.10.21      1      UG---S--m--- toronto  0d:0h:
41m:1s
@mbe 77.0.13.0/24         50.1.10.21      1      UG---S--m--- toronto  0d:0h:
41m:1s
@mbe 77.0.14.0/24         50.1.10.21      1      UG---S--m--- toronto  0d:0h:
41m:1s
@d 82.0.0.0/24            82.0.0.15       1      U-----m--- trunk28-2 12d:1h:
30m:36s
Origin(Ori): (b) BlackHole, (be) EBGp, (bg) BGP, (bi) IBGP, (bo) BOOTP
(ct) CBT, (d) Direct, (df) DownIF, (dv) DVMRP, (e1) ISISL1Ext
(e2) ISISL2Ext, (h) Hardcoded, (i) ICMP, (i1) ISISL1 (i2) ISISL2
(is) ISIS, (mb) MBGP, (mbe) MBGPExt, (mbi) MBGPInter, (mp) MPLS Lsp
(mo) MOSPF (o) OSPF, (o1) OSPFExt1, (o2) OSPFExt2
(oa) OSPFIntra, (oe) OSPFAsExt, (or) OSPFInter, (pd) PIM-DM, (ps) PIM-SM
(r) RIP, (ra) RtAdvrt, (s) Static, (sv) SLB_VIP, (un) UnKnown
(*) Preferred unicast route (@) Preferred multicast route
(#) Preferred unicast and multicast route
Flags: (B) BlackHole, (D) Dynamic, (G) Gateway, (H) Host Route
(L) Matching LDP LSP, (l) Calculated LDP LSP, (m) Multicast
(P) LPM-routing, (R) Modified, (S) Static, (s) Static LSP
(T) Matching RSVP-TE LSP, (t) Calculated RSVP-TE LSP, (u) Unicast, (U) Up
(f) Provided to FIB (c) Compressed Route
Mask distribution:
14 routes at length 24          1 routes at length 32
Route Origin distribution:
3 routes from Direct
Total number of routes = 15
Total number of compressed routes = 0

```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



show L2stats

```
show L2stats {vlan vlan_name}
```

Description

Displays the counters for the number of packets bridged, switched, and snooped (Layer 2 statistics).

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
------------------	------------------------

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the counters for the number of packets bridged, switched, and snooped (Layer 2 statistics) for the VLAN accounting:

```
show L2stats accounting
```



Note

You can also enter the command as show l2stats. We use the uppercase letter here to avoid confusion with the numeral 1.

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

show mcast cache

```
show mcast {ipv4 | ipv6} cache {{vlan} name} {[group grpaddressMask |  
grpaddressMask] {source sourceIP | sourceIP}} {type [snooping | pim | mvr]}|  
{summary}}
```



Description

Displays multicast cache information.

The display can be limited to entries for specific VLANs or groups, and it can be limited to specific types of entries, such as those created by snooping protocols, PIM, or MVR.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
<i>name</i>	Specifies a VLAN name.
<i>grpaddressMask</i>	Specifies a multicast group address and mask.
<i>sourceIP</i>	Specifies the source IP address for a multicast group.
snooping	Limits the display to cache entries created by PIM or IGMP snooping.
pim	Limits the display to cache entries created by PIM.
mvr	Limits the display to cache entries created by MVR.
summary	Specifies the summary display format.

Default

Displays information for all entries in the multicast cache.

Usage Guidelines

If the `configure forwarding ipmc lookup-key mac-vlan` command is configured, the following message is displayed:

NOTE: Traffic is forwarded based on MAC address. Actual traffic forwarded based on the installed MAC address need not be the same displayed in this command, if overlapping IP multicast addresses are used in the network.

If the mode is **mixed-mode**, the following message is displayed:

NOTE: Traffic could be forwarded based on MAC address. Actual traffic forwarded based on the installed MAC address need not be the same displayed in this command, if overlapping IP multicast addresses are used in the network.

Example

The following command displays all multicast cache information:

```
show mcast cache
Snooping/MVR Cache Timeout: 300 sec
Type Group          Sender          Age  InVlan
snoop 225.1.1.1      222.222.222.222  17  snvlan
```



```

Vlan          Port      Vid
snvlan        2          400
23            400
snoop 224.0.0.5          100.1.2.2      2    pmvlan2
Vlan          Port      Vid
pmvlan2       4          402
snoop 224.0.0.5          100.1.3.3      17   pmvlan3
Vlan          Port      Vid
pmvlan3       23         403
snoop 224.0.0.13        100.1.2.2      11   pmvlan2
Vlan          Port      Vid
pmvlan2       4          402
snoop 224.0.0.13        100.1.3.3      14   pmvlan3
Vlan          Port      Vid
pmvlan3       23         403
pim 226.1.1.1          100.1.1.12     0    pmvlan1
Vlan          Port      Vid
pmvlan2       4          402
pmvlan3       23         403
Multicast cache distribution:
5 entries from Snooping          0 entries from MVR          1 entries from
PIM
Total Cache Entries: 6

```

The following command displays summary cache information for VLAN pmvlan1:

```

show mcast cache vlan pmvlan1 summary
Snooping/MVR Cache Timeout: 300 sec
=====MULTICAST CACHE SUMMARY=====
Multicast cache distribution:
5 entries from Snooping          0 entries from MVR          1 entries from
PIM
pmvlan1: Multicast cache distribution:
0 entries from Snooping          0 entries from MVR          1 entries from
PIM
Total Cache Entries: 6
Total Cache Entries for VLAN pmvlan1: 1

```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

show mvr

```
show mvr {vlan vlan_name}
```



Description

Displays the MVR configuration on the switch.

Syntax Description

<code>vlan_name</code>	Specifies a VLAN name.
------------------------	------------------------

Default

N/A.

Usage Guidelines

If a VLAN is specified, information for the VLAN is displayed.

Example

The following command displays the MVR configuration for the VLAN accounting:

```
show mvr accounting
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MVR feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show mvr cache

This command is provided for backward compatibility. The recommended command is:

```
show mcast cache {{vlan} vlan_name} {[group grpaddressMask | grpaddressMask]
{source sourceIP | sourceIP}} {type [snooping | pim | mvr]}| {summary}}
```

The syntax for the original form of this command is:

```
show mvr cache {vlan vlan_name}
```

Description

Displays the multicast cache entries added by MVR.



Syntax Description

<code>vlan_name</code>	Specifies a VLAN name.
------------------------	------------------------

Default

N/A.

Usage Guidelines

If no VLAN is specified, information for all the VLANs is displayed.

Example

The following command displays the multicast cache in the MVR range for the VLAN vlan110:

```
Switch.78 # show mvr cache vlan110
```

This command display is the same as for the following preferred command:

```
show mcast cache {{vlan} name} {[group grpaddressMask | grpaddressMask] {source sourceIP | sourceIP}} {type [snooping | pim | mvr]}| {summary}}
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MVR feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show pim

```
show pim {ipv4 | ipv6 | detail | rp-set {group_addr} | vlan vlan_name}
```

Description

Displays the PIM configuration and statistics.

Syntax Description

<code>ipv4</code>	Displays PIM IPv4 configuration information.
<code>ipv6</code>	Displays PIM IPv6 configuration information.
<code>detail</code>	Displays show output in the detailed format.



<i>group_addr</i>	Specifies an IP multicast group, for which the RP is to be displayed.
<i>vlan_name</i>	Specifies a VLAN name.

Default

IPv4 is the default for the `show pim {ipv4 | ipv6}` command.

If no VLAN is specified, the configuration is displayed for all PIM interfaces.

If no multicast group is specified for the `rp-set` option (Rendezvous Point set), all RPs are displayed.

Usage Guidelines

The detail version of this command displays the global statistics for PIM, as well as the details of each PIM enabled VLAN.

Example

The following command displays the global PIM configuration and statistics:

```
Switch. 14 # show pim
PIM Enabled, Version 2
PIM CRP Disabled
BSR state           : ACCEPT_ANY ; BSR Hash Mask Length: 255.255.255.252
Current BSR Info    : 0.0.0.0 (Priority 0)
Configured BSR Info : 0.0.0.0 (Priority 0)
CRP Adv Interval    : 60 sec ; CRP Holdtime: 150
BSR Interval        : 60 sec ; BSR Timeout : 130
Cache Timer         : 210 sec ; Prune Timer : 210
Assert Timeout      : 210 sec ; Register Suppression Timeout,Probe: 60, 5
Generation Id       : 0x4a01a1a6
PIM-DM State Refresh TTL           : 16
PIM-DM State Refresh Source Active Timer : 210 sec
PIM-DM State Refresh Origination Interval: 60 sec
Threshold for Last Hop Routers: 0 kbps
Threshold for RP                   : 0 kbps
Register-Rate-Limit-Interval       : Always active
PIM SSM address range              : None
PIM IP Route Sharing               : Disabled
PIM IP Route Sharing Hash          : Source-Group-Next Hop
Register Checksum to include data
Active Sparse Ckts 0 Dense Ckts 2
Global Packet Statistics (In/Out)
C-RP-Advs           0           0
Registers           0           0
RegisterStops       0           0
VLAN                Cid IP Address      Designated      Flags          Hello  J/P
Nbrs
Router              Int    Int
v36                  2 36.36.36.3  /16 36.36.36.6   rifmd-----R  30   60   1
vixia                2 64.1.1.1    /16 64.1.1.1     rifmd-----R  30   60   0
Legend: J/P Int: Join/Prune Interval
Flags : r - Router PIM Enabled, i - Interface PIM Enabled, f - Interface,
Forwarding Enabled, m - Interface Multicast Forwarding Enabled,
```



s - Sparse mode, d - Dense mode, c - CRP enabled,
 t - Trusted Gateway configured, n - Multinetted VLAN,
 p - Passive Mode, S - Source Specific Multicast, b - Border,
 R - State Refresh Enabled.

The following command displays the detailed PIM configuration and statistics:

```
Switch.22 # show pim detail
PIM Enabled, Version 2
PIM CRP Disabled
BSR state          : ACCEPT_ANY ; BSR Hash Mask Length: 255.255.255.252
Current BSR Info   : 0.0.0.0 (Priority 0)
Configured BSR Info : 0.0.0.0 (Priority 0)
CRP Adv Interval   : 60 sec ; CRP Holdtime: 150
BSR Interval       : 60 sec ; BSR Timeout : 130
Cache Timer        : 210 sec ; Prune Timer : 210
Assert Timeout     : 210 sec ; Register Suppression Timeout,Probe: 60, 5
Generation Id      : 0x4a01a1a6
PIM-DM State Refresh Source Active Timer : 210 sec
PIM-DM State Refresh TTL                  : 16
PIM-DM State Refresh Origination Interval: 60 sec
Threshold for Last Hop Routers: 0 kbps
Threshold for RP                          : 0 kbps
Register-Rate-Limit-Interval : Always active
PIM SSM address range           : None
PIM IP Route Sharing            : Disabled
PIM IP Route Sharing Hash       : Source-Group-Next Hop
Register Checksum to include data
Active Sparse Ckts 0 Dense Ckts 1
Global Packet Statistics (In/Out)
C-RP-Advs          0          0
Registers          0          0
RegisterStops      0          0
PIM DENSE Interface[2] on VLAN v36 is enabled and up
IP adr: 36.36.36.3   mask: 255.255.0.0   DR of the net: 36.36.36.6
Passive             : No
Hello Interval      : 30 sec
Neighbor Time out   : 105 sec
Join/Prune Interval : 60 sec
Join/Prune holdtime : 210 sec
Trusted Gateway     : none
CRP group List      : none with priority 0
Shutdown priority   : 1024
Source Specific Multicast : Disabled
State Refresh       : On
State Refresh Capable : Yes
Border              : No
Neighbor IP address  Generation Id   Expires State Refresh
36.36.36.6          0x4a01a39d    90      On
Packet Statistics (In/Out)
Hellos              20          20  Bootstraps          0          0
Join/Prunes         0          0  Asserts             0          0
Grafts              0          0  GraftAcks           0          0
State Refresh       0          0
```



The following command displays the elected, active RP for the group 239.255.255.1:

```
show pim rp-set 239.255.255.1
Group          Mask          C-RP          Origin      Priority
224.0.0.0     240.0.0.0    10.10.10.2   Bootstrap  0
224.0.0.0     240.0.0.0    124.124.124.124 Bootstrap  0
224.0.0.0     240.0.0.0    124.124.124.124 static     0
239.255.255.0 255.255.255.0 124.124.124.124 Bootstrap  0
Elected RP is 124.124.124.124
```

The following command displays the PIM configuration for VLAN v3:

```
# show pim v3
PIM SPARSE Interface[2] on VLAN v3 is enabled and up
IP adr: 30.30.30.1   mask: 255.255.255.0   DR of the net: 30.30.30.2

DR Priority          : 1
Passive              : No
Hello Interval      : 30 sec
Neighbor Time out   : 105 sec
Join/Prune Interval : 60 sec
Join/Prune holdtime : 210 sec
Trusted Gateway     : none
CRP group List      : pimPolicy with priority 0
Shutdown priority   : 1024
Source Specific Multicast : Disabled
State Refresh       : Off
State Refresh Capable : No
Border              : No

Neighbor IP address      Generation      State      DR
                        Id          Expires    Refresh    Priority
30.30.30.2              0x5199b2db    105       No         1

Packet Statistics (In/Out)
Hellos                  41            40         Bootstraps      0
Join/Prunes             0             0         Asserts          0
Grafts                  0             0         GraftAcks       0
State Refresh           0             0
```

The following is PIM IPv4 show output for the show register policy configuration, including drop counters:

```
Switch. 14 # show pim ipv4
PIM Enabled, Version 2
PIM CRP Disabled
BSR state          : ACCEPT_ANY ; BSR Hash Mask Length: 255.255.255.252
Current BSR Info   : 0.0.0.0 (Priority 0)
Configured BSR Info : 0.0.0.0 (Priority 0)
CRP Adv Interval   : 60 sec ; CRP Holdtime: 150
BSR Interval       : 60 sec ; BSR Timeout : 130
Cache Timer        : 210 sec ; Prune Timer : 210
Assert Timeout     : 210 sec ; Register Suppression Timeout,Probe: 60, 5
Generation Id      : 0x4fd9a7f3
PIM-DM State Refresh TTL          : 16
PIM-DM State Refresh Source Active Timer : 210
PIM-DM State Refresh Origination Interval : 60
```




```

Threshold for Last Hop Routers: 0 kbps
Threshold for RP                : 0 kbps
Register-Rate-Limit-Interval  : Always active
PIM SSM address range         : None
PIM Register Policy            : swl_rp_filter
PIM Register Policy RP        : None
Register Checksum to include data
PIM IP Route Sharing           : Disabled
PIM IP Route Sharing Hash      : Source-Group-Next Hop

```

```

Active Sparse Ckts 0 Dense Ckts 0 State Refresh Ckts 0
  Global Packet Statistics (In          Out          Drop)
  C-RP-Advts          0            0            0
  Registers            1            0            1
  RegisterStops       0            0            0

```

The following is PIM show output with IP Route Sharing information:

```

PEER2_460.17 # show pim
PIM Enabled, Version 2
PIM CRP Disabled
BSR state           : ACCEPT_ANY ; BSR Hash Mask : 255.255.255.252
Current BSR Info    : 0.0.0.0 (Priority 0)
Configured BSR Info : 0.0.0.0 (Priority 0)
CRP Adv Interval    : 60 sec ; CRP Holdtime: 150
BSR Interval        : 60 sec ; BSR Timeout : 130
Cache Timer         : 210 sec ; Prune Timer : 210
Assert Timeout      : 210 sec ; Register Suppression Timeout,Probe: 60, 5
Generation Id       : 0x50c60413
PIM-DM State Refresh TTL           : 16
PIM-DM State Refresh Source Active Timer : 210
PIM-DM State Refresh Origination Interval : 60
Threshold for Last Hop Routers: 0 kbps
Threshold for RP                : 0 kbps
Register-Rate-Limit-Interval  : Always active
PIM SSM address range         : None
PIM Register Policy            : None
PIM Register Policy RP        : None
PIM IP Route Sharing           : Disabled

```

```

Register Checksum to include data
Active Sparse Ckts 0 Dense Ckts 0 State Refresh Ckts 0

```

```

  Global Packet Statistics ( In          Out          Drop )
  C-RP-Advts          0            0            0
  Registers            0            0
  RegisterStops       0            0

```

```

VLAN      Cid IP Address          Designated      Flags          Hello J/P
Nbrs
Router
Int
v1          1 1.1.1.9          / 24 1.1.1.9    rifms----- 30    60
0

```



```
v2          2 2.2.2.9      / 24 2.2.2.9      rifms----- 30    60
0
```

The following command shows the output for the `show pim ipv6 v3` command:

```
# show pim ipv6 v3
  PIM SPARSE Interface[1] on VLAN v3 is enabled and up
  Global IP adr           : 2010::2/64
  Local IP adr            : fe80::204:96ff:fe27:f2c6/64
  DR of the net           : fe80::204:96ff:fe27:f2c6
  DR Priority              : 1

  Passive                  : No
  Hello Interval           : 30 sec
  Neighbor Time out       : 105 sec
  Join/Prune Interval     : 60 sec
  Join/Prune holdtime    : 210 sec
  Trusted Gateway         : none
  CRP group List          : none with priority 0
  Shutdown priority       : 1024
  Source Specific Multicast : Disabled
  State Refresh           : Off
  State Refresh Capable   : No
  Border                   : No
  Secondary Interfaces: 2003::2/ 64
```

Neighbor IP address	Generation Id	Expires	State Refresh	DR Priority
fe80::204:96ff:fe26:6c89	0x5192f6f5	101		
No	1			

```
Packet Statistics (In/Out)
  Hellos                5          6
  Bootstraps            0          0
  Join/Prunes           0          0
  Asserts               0          0
  Grafts                0          0
  GraftAcks             0          0
  State Refresh         0          0
```

History

This command was first available in ExtremeXOS 10.1.

The PIM-SSM information was added in ExtremeXOS 11.4.

Border VLAN information was added in ExtremeXOS 12.0.

The `ipv6` keyword was added to PIM Register Policy Filter feature in ExtremeXOS 15.3.

DR Priority output was added in ExtremeXOS 15.3.2.

IP Route Sharing output was added in ExtremeXOS 15.3.2.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show pim cache

```
show pim { ipv4 | ipv6 } cache {{detail} | {state-refresh} {mlag-peer-info}
{group_addr {source_addr}}
```

Description

Displays the multicast cache entries created by PIM.

Syntax Description

ipv4	Displays IPv4 PIM information
ipv6	Displays IPv6 PIM information.
detail	Specifies to display the information in detailed format.
<i>group_addr</i>	Specifies an IP group address.
<i>source_addr</i>	Specifies an IP source address.
state-refresh	Specifies to display the PIM cache entries with state refresh parameters.
mlag-peer-info	Shows MLAG peer related information.

Default

N/A.

Usage Guidelines

Displays the following information:

- IP group address
- IP source address / source mask
- Upstream neighbor (RPF neighbor)
- Interface (VLAN) to upstream neighbor
- Cache expire time
- Egress and prune interface list

When the detail option is specified, the switch displays the egress VLAN list and the pruned VLAN list.



Example

The following command displays the PIM cache entry for group 239.255.255.1:

```
Switch.33 # show pim cache 239.255.255.1
Index Dest Group      Source          InVlan  Origin
[0000] 239.255.255.1  124.124.124.124 (WR) v4      Sparse
Entry timer is not run; UpstNbr: 200.124.124.24
EgressIfList = vbs15(0)(FW)(SM)(I)
[0001] 239.255.255.1  118.5.1.1 (S)          vbs5    Sparse
Expires after 186 secs UpstNbr: 0.0.0.0
RP: 124.124.124.124 via 200.124.124.24 in v4
EgressIfList = vbs15(0)(FW)(SM)(I) , vpim5(170)(FW)(SM)(S)
PrunedIfList = v4(0)(SM)
Number of multicast cache = 20
Entry flags :-
R: RP tree. S: Source tree. W: Any source.
Egress/Pruned interface flags :-
SM: Sparse Mode           DM: Dense Mode
Fw: Forwarding           PP: Prune pending
AL: Assert Loser         N: Neighbor present
I: IGMP member present   S: (s,g) join received
Z: (*,g) join received   Y: (*,*,rp) join received
```

The following command displays the PIM-DM cache entry with state-refresh information for group 225.0.0.1:

```
Switch.5 # show pim cache state-refresh 225.0.0.1
Index Dest Group      Source          InVlan  Origin
[0001] 225.0.0.1      64.1.1.100 (S)          vixia   Dense   Not Pruned
Expires after 204 secs UpstNbr: 0.0.0.0
Refresh State: Originator(20), TTL: 16
EgressIfList = v36(0)(FW)(DM)(N)
[0001] 225.0.0.1      65.1.1.100 (S)          vixia   Dense   Not Pruned
Expires after 195 secs UpstNbr: 65.1.1.200
Refresh State: Not-Originator(25), TTL: 8
EgressIfList = v36(0)(FW)(DM)(N)
```

The following command displays the ingress VLAN information of all MLAG peers:

```
Switch.5 # show pim cache mlag-peer-info
Index Dest Group      Source          InVlan  Origin
[0001] 225.0.0.1      64.1.1.100 (S)          vixia   Sparse
Expires after 210 secs UpstNbr: 0.0.0.0
[S1] Peer Ingress VLAN: 1.1.1.1/24 (Same)
EgressIfList = v36(0)(FW)(DM)(I)
```

History

This command was first available in ExtremeXOS 10.1.

This command was modified in ExtremeXOS 15.2 to display MLAG peer information.



The **IPv4** and **IPv6** keyword options were added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show pim snooping

```
show pim snooping {vlan} vlan_name
```

Description

Displays the PIM snooping configuration for a VLAN.

Syntax Description

<code>vlan_name</code>	Specifies a VLAN.
------------------------	-------------------

Default

Disabled.

Usage Guidelines

None.

Example

The following command displays the PIM snooping configuration for the default VLAN:

```
BD-8810Rack3.8 # show pim snooping default
Global PIM Snooping DISABLED
Default          Snooping DISABLED
```

History

This command was first available in ExtremeXOS 12.1.

The **ipv4** and **ipv6** keyword options were added in ExtremeXOS 15.3.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure igmp

unconfigure igmp

Description

Resets all IGMP settings to their default values and clears the IGMP group table.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command resets all IGMP settings to their default values and clears the IGMP group table:

```
unconfigure igmp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure igmp snooping vlan ports set join-limit

```
unconfigure igmp snooping {vlan} vlan_name ports port_list set join-limit
```



Description

Removes the join limit set on VLAN ports.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

No limit.

Usage Guidelines

None.

Example

The following command removes the join limit for port 2:1 in the Default VLAN:

```
unconfigure igmp snooping "Default" ports 2:1 set join-limit
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure igmp ssm-map

```
unconfigure igmp ssm-map {v vr-name}
```

Description

Unconfigures all SSM mappings on the virtual router.

Syntax Description

<i>vr-name</i>	Specifies a virtual router name. If the VR name is omitted, the switch uses the VR specified by the current CLI VR context.
----------------	---



Default

N/A.

Usage Guidelines

None.

Example

The following command deletes all IGMP-SSM mappings on the virtual router xyz:

```
unconfigure igmp ssm-map vr xyz
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure pim

```
unconfigure pim {ipv4 | ipv6} {vlan vlan_name} | {tunnel} tunnel_name] border
```

Description

Resets all PIM settings on an IPv4 or IPv6 module, or on one or all VLANs, to their default values.

Syntax Description

ipv4	Specifies the ipv4 module from which PIM is to be unconfigured.
ipv6	Specifies the ipv6 module from which PIM is to be unconfigured.
<i>vlan_name</i>	Specifies the VLAN from which PIM is to be unconfigured.
tunnel	Specifies the tunnel which PIM is to be unconfigured.
<i>tunnel_name</i>	Specifies the tunnel name.
border	Specifies the border.

Default

If no VLAN is specified, the configuration is reset for all PIM interfaces.



Usage Guidelines

If you unconfigure PIM, you also unconfigure PIM-SSM, removing the PIM-SSM range.

Example

The following command resets all PIM settings on the VLAN accounting:

```
unconfigure pim vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

The **ipv4** and **ipv6** keywords were added, giving an option to support this functionality in IPv6 as well, in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure pim ssm range

```
unconfigure pim {ipv4 | ipv6} ssm range
```

Description

Unconfigures the range of multicast addresses for PIM SSM.

Syntax Description

ipv4	Configures PIM functionality on IPv4 router interfaces.
ipv6	Configures PIM functionality on IPv6 router interfaces.

Default

By default, no SSM range is configured.

Usage Guidelines

You must disable PIM before configuring or unconfiguring a PIM-SSM range. Use the `disable pim` command.



Initially, no range is configured for SSM. After a range is configured, you can remove the range with the `unconfigure pim ssm range` command.

When no range is configured for PIM SSM, the switch does not use PIM SSM for any multicast groups.

Example

The following command removes the PIM SSM range:

```
unconfigure pim ssm range
```

History

This command was first available in ExtremeXOS 11.4.

The **ipv4** and **ipv6** keywords were added, giving an option to support this functionality in IPv6 as well, in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



45 IPv6 Multicast Commands

```
clear mld counters
clear mld group
clear mld snooping
configure mld
configure mcast ipv6 cache timeout
configure mld snooping fast-learning
configure mld snooping filters
configure mld snooping vlan ports add dynamic group
configure mld snooping vlan ports add static group
configure mld snooping vlan ports delete static group
configure mld snooping vlan ports add static router
configure mld snooping vlan ports delete static router
configure mld snooping vlan ports filter
configure mld snooping vlan ports join-limit
configure mld snooping flood-list
configure mld snooping leave-timeout
configure mld snooping timer
disable ipmcforwarding ipv6
disable mld
disable mld snooping
enable ipmcforwarding ipv6
enable mld
enable mld snooping
enable mld snooping with-proxy
show mcast ipv6 cache
show mld
show mld counters
show mld group
show mld snooping
show mld snooping vlan filter
show mld snooping vlan static
unconfigure mld
```

This chapter describes commands for doing the following:

- Configuring IPv6 multicast routing
- Displaying IPv6 multicast information

For an introduction to the IPv6 multicast feature, see the *ExtremeXOS Concepts Guide*.

This chapter contains information about the following commands:

- `clear mld counters`
- `clear mld group`
- `clear mld snooping`
- `configure mld`
- `configure mcast ipv6 cache timeout`
- `configure mld snooping fast-learning`
- `configure mld snooping filters`
- `configure mld snooping vlan ports add static group`
- `configure mld snooping vlan ports delete static group`
- `configure mld snooping vlan ports add static router`
- `configure mld snooping vlan ports delete static router`
- `configure mld snooping vlan ports delete static group`
- `configure mld snooping vlan ports join-limit`
- `configure mld snooping flood-list`
- `configure mld snooping leave-timeout`
- `configure mld snooping timer`
- `disable mld`
- `disable mld snooping`
- `enable mld`
- `enable mld snooping`
- `enable mld snooping with-proxy`
- `show mcast ipv6 cache`
- `show mld`
- `show mld counters`
- `show mld group`
- `show mld snooping`
- `show mld snooping vlan filter`
- `show mld snooping vlan static`
- `unconfigure mld`

For an introduction to the IPv6 multicast feature, see the Multicast Routing Overview in the *ExtremeXOS Concepts Guide*.

clear mld counters

```
clear mld counters {{vlan} vlan_name}
```

Description

Clears MLD statistics counters.



Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
------------------	------------------------

Default

N/A.

Usage Guidelines

Use this command to manually clear MLD statistics counters.

Example

The following example clears all MLD counters for all VLANs:

```
clear mld counters
```

If a VLAN is specified, only the counters on the specific VLAN is cleared.

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

clear mld group

```
clear mld group {v6grpipaddress} {{vlan} name}
```

Description

Removes one or all MLD groups.

Syntax Description

<i>v6grpipaddress</i>	Specifies the group IP address.
<i>name</i>	Specifies a VLAN name.

Default

N/A.



Usage Guidelines

This command is used to manually remove learned MLD group entries instantly.

Example

The following command clears all MLD groups from VLAN accounting:

```
clear mld group accounting
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

clear mld snooping

```
clear mld snooping [{vlan} name]
```

Description

Removes one or all MLD snooping entries.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

N/A.

Usage Guidelines

This command can be used by network operations to manually remove MLD snooping entries instantly. However, removing an MLD snooping entry can disrupt the normal forwarding of multicast traffic, until the snooping entries are learned again.

The static and dynamic MLD snooping entries are removed, then recreated upon the next general query. The static router entry is removed and recreated immediately.



Example

The following command clears MLD snooping from VLAN accounting:

```
clear mld snooping accounting
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mld

```
configure mld query_interval query_response_interval last_member_query_interval  
{robustness}
```

Description

Configures the Multicast Listener Discovery (MLD) timers.

Syntax Description

<i>query_interval</i>	Specifies the interval (in seconds) between general queries.
<i>query_response_interval</i>	Specifies the maximum query response time (in seconds).
<i>last_member_query_interval</i>	Specifies the maximum group-specific query response time (in seconds).
<i>robustness</i>	Specifies the degree of robustness for the network.

Default

- query interval—125 seconds
- query response interval—10 seconds
- last member query interval—1 second
- robustness—2

Usage Guidelines

Timers are based on RFC2710. Specify the following:



- query interval—The amount of time, in seconds, the system waits between sending out general queries. The range is 1 to 429,496,729 seconds.
- query response interval—The maximum response time inserted into the periodic general queries. The range is 1 to 25 seconds.
- last member query interval—The maximum response time inserted into a group-specific query sent in response to a leave group message. The range is 1 to 25 seconds.
- robustness—The degree of robustness of the network. The range is 2 to 7.

Example

The following command configures the MLD timers:

```
configure mld 100 5 1 3
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mcast ipv6 cache timeout

```
configure mcast ipv6 cache timeout {seconds | none}
```

Description

Configures the IPv6 multicast cache timeout.

Syntax Description

<i>seconds</i>	Idle time after which cache entries are deleted.
none	Cache entries are not timed out.

Default

300 seconds.

Usage Guidelines

Cache timeout is the time after which the cache entries are deleted, if traffic is not received for that duration. The applies only for snooping and MVR caches and does not apply for PIM caches.



The range is 90 to 100000 seconds. You can use the option none if you do not want the cache entry to be deleted. If none is configured, the cache entries could be deleted only using the following command:

```
clear igmp snooping
```

Example

```
configure mcast ipv6 cache timeout 400
configure mcast ipv6 cache timeout none
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MLD snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mld snooping fast-learning

```
configure mld snooping fast-learning [on | off] [vlan vlan_name]
```

Description

Configures fast-learning mode.

Syntax Description

<i>vlan_name</i>	Specifies a vlan name
------------------	-----------------------

Default

off.

Usage Guidelines

When MLD snooping is enabled on a VLAN, learning of group entries will happen only when the next periodic query is sent by the querier in the network. When fast-learning is turned on using this command, a general is sent under the following conditions:

- When MLD snooping is enabled.
- When MLD snooping VLAN is operationally up.
- Group join limit changed through configuration.



Query generated for faster learning uses unspecified address as the source address (both L2 and L3), unless the switch generating the triggered query is the querier for the network.

Example

```
configure mld snooping fast-learning on
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MLD snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mld snooping filters

```
configure mld snooping filters [per-port | per-vlan]
```

Description

Selects the type of MLD snooping filters that are installed.

Syntax Description

per-port	Installs the per-port MLD snooping filters.
per-vlan	Installs the per-VLAN MLD snooping filters.

Default

per-port.

Usage Guidelines

This command applies only to Summit family switches and BlackDiamond 8800 series switches.

Use the per-vlan option when the number of VLANs configured on the switch is lower than half of the maximum numbers listed in [Table 60](#). This option conserves usage of the hardware Layer 3 multicast forwarding table.

When the number of configured VLANs is larger than half of the maximum values listed in [Table 60](#), select the per-port option. Each VLAN requires additional interface hardware ACL resources. The per-port option conserves usage of the interface hardware ACL resources.



To display the MLD snooping filters configuration, use the `show mld snooping` command.

Example

The following command configures the switch to install the per-VLAN MLD snooping filters:

```
configure mld snooping filters per-vlan
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MLD snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mld snooping vlan ports add dynamic group

```
configure mld snooping {vlan} vlan_name {ports portlist} add dynamic group
[ v6group]
```

Description

Configures an MLD dynamic group.

Syntax Description

<code>vlan_name</code>	Specifies a vlan name.
<code>portlist</code>	Specifies a port list.
<code>v6group</code>	Specifies a version 6 group.

Default

N/A.

Usage Guidelines

This command is not saved in the configuration. The following message is displayed on execution of this command:

```
INFO: This command is not saved in the configuration.
```



Example

```

show mcast cache
Snooping/MVR Cache Timeout: 300 sec

Type Group                Sender                Age  InVlan
pim  225.1.1.1            20.20.20.50          0   v1
     Vlan                Port      Vid
     v1                  1         10
     v3                  Lpbk      4088

snoop 224.0.0.5          10.10.10.1           21  v1
     Vlan                Port      Vid
     v1                  1         10      snoop 224.0.0.13
10.10.10.1              21  v1          Vlan
     Port                Vid      v1
     1                   10      Multicast cache
distribution:           2 entries from Snooping      0 entries from
MVR                    1 entries from PIM          Total Cache Entries: 3

```

History

This command was first available in ExtremeXOS 15.3.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see [Feature License Requirements](#) in the the *ExtremeXOS Concepts Guide*.

configure mld snooping vlan ports add static group

```

configure mld snooping {vlan} vlan_name {ports port_list }add static group
v6grpipaddress

```

Description

Configures VLAN ports to receive the traffic from a multicast group, even if no MLD joins have been received on the port.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>port_list</i>	Specifies one or more ports or slots and ports. On a modular switch, it can be a list of slots and ports. On a stand-alone switch, it can be one or more port numbers. In the form 1, 2, 3-5, 2:5, 2:6-2:8.
<i>v6grpipaddress</i>	Specifies the multicast group IPv6 address.



Default

None.

Usage Guidelines

Use this command to forward a particular multicast group to VLAN ports. In effect, this command emulates a host on the port that has joined the multicast group. As long as the port is configured with the static entry, multicast traffic for that multicast group is forwarded to that port.

The switch sends proxy MLD messages in place of those generated by a real host. The proxy messages use the VLAN IPv6 address for source address of the messages. If the VLAN has no IPv6 address assigned, the proxy MLD message uses 0::0 as the source IP address.

Example

The following command configures a static MLD entry so the multicast group ff02::1:1 is forwarded to VLAN marketing on ports 2:1-2:4:

```
configure mld snooping marketing ports 2:1-2:4 add static group ff02::1:1
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mld snooping vlan ports delete static group

```
configure mld snooping {vlan} vlan_name ports port_list delete static group[all | v6grpipaddress]
```

Description

Removes the configuration that causes VLAN ports to receive the traffic from a multicast group, even if no MLD joins have been received on the port.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>port_list</i>	Specifies one or more ports or slots and ports. On a modular switch, it can be a list of slots and ports. On a stand-alone switch, it can be one or more port numbers. In the form 1, 2, 3-5, 2:5, 2:6-2:8.



all	Specifies all multicast groups.
v6grpipaddress	Specifies the multicast group IPv6 address.

Default

None.

Usage Guidelines

Use this command to delete a static group from a particular VLAN port.

To add a static group, use the following command:

```
configure mld snooping {vlan} vlan_name portsport_list add static
groupv6grpipaddress configure mld snooping {vlan}vlanname portsport_list add
static groupv6grpipaddress
```

Example

The following command removes a static MLD entry so the multicast group ff02::a:b is not forwarded to VLAN marketing on ports 2:1-2:4, unless an MLD join message is received on the port:

```
configure mld snooping marketing ports 2:1-2:4 delete static group ff02::a:b
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mld snooping vlan ports add static router

```
configure mld snooping {vlan} vlan_name ports port_list add static router
```

Description

Configures VLAN ports to forward the traffic from all multicast groups, even if no MLD joins have been received on the port.



Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>port_list</i>	Specifies one or more ports or slots and ports. On a modular switch, it can be a list of slots and ports. On a stand-alone switch, it can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

None.

Usage Guidelines

Use this command to forward all multicast groups to the specified VLAN ports. In effect, this command emulates a multicast router attached to those ports. As long as the ports are configured with the static entry, all available multicast traffic is forwarded to those ports.

Example

The following command configures a static MLD entry so all multicast groups are forwarded to VLAN marketing on ports 2:1-2:4:

```
configure mld snooping marketing ports 2:1-2:4 add static router
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mld snooping vlan ports delete static router

```
configure mld snooping {vlan} vlan_name ports port_list delete static router
```

Description

Configures VLAN ports to stop forwarding the traffic from all multicast groups, unless MLD joins have been received on the port.



Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>port_list</i>	Specifies one or more ports or slots and ports. On a modular switch, it can be a list of slots and ports. On a stand-alone switch, it can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

None.

Usage Guidelines

Use this command to remove the configuration that forwards all multicast groups to the specified VLAN ports.

Example

The following command removes a static MLD entry so all multicast groups are not forwarded to VLAN marketing on ports 2:1-2:4, unless an MLD join is received on the port:

```
configure mld snooping marketing ports 2:1-2:4 delete static router
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mld snooping vlan ports filter

```
configure mld snooping vlan vlan_name ports port_list filter [policy]
```

Description

Configures a MLD snooping policy file filter on VLAN ports.



Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>port_list</i>	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a standalone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
<i>policy</i>	Specifies the policy file for the filter.

Default

None.

Usage Guidelines

Use this command to filter multicast groups to the specified VLAN ports.

The policy file used by this command is a text file that contains the IPv6 multicast addresses of the multicast groups that you wish to block.

To remove MLD snooping filtering from a port, use the none keyword version of the command.

Use the following template to create a snooping filter policy file:

```
#
# Add your group addresses between "Start" and "end"
# Do not touch the rest of the file!!!!
entry mldFilter {
  if match any {
    #----- Start of group addresses -----
    nlri FF03::1/128;
    nlri FF05::1/112;
    #----- end of group addresses -----
  } then {
    deny;
  }
}
entry catch_all {
  if {
  } then {
    permit;
  }
}
```

Example

The following command configures the policy file ap_multicast to filter multicast packets forwarded to VLAN marketing on ports 2:1-2:4:

```
configure mld snooping marketing ports 2:1-2:4 filter ap_multicast
```



History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MLD snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mld snooping vlan ports join-limit

```
configure mld snooping {vlan} vlan_name ports port_list join-limit [num_joins | no-limit]
```

Description

Configures VLAN ports to support a maximum number of MLD joins.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name
<i>port_list</i>	Specifies one or more ports or slots and ports.
num	Specifies the maximum number of joins permitted on the ports. The range is 1 to 5000.

Default

No limit.

Usage Guidelines

None.

Example

The following command configures port 2:1 in the Default VLAN to support a maximum of 100 MLD joins:

```
configure mld snooping "Default" ports 2:1 join-limit 100
```

History

This command was first available in ExtremeXOS 15.2.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MLD snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mld snooping flood-list

```
configure mld snooping flood-list [policy | none]
```

Description

Configures certain multicast addresses to be slow path flooded within the VLAN.

Syntax Description

<i>policy</i>	Specifies a policy file with a list of multicast addresses to be handled.
none	Specifies no policy file is to be used.

Default

None.

Usage Guidelines

With this command, a user can configure certain multicast addresses to be slow path flooded within the VLAN, instead of fast path forwarded according to MLD and/or Layer 3 multicast protocol.

A policy file is a text file with the extension .pol. It can be created or edited with any text editor. The specified policy file *policy file* should contain a list of addresses that determine if certain multicast streams are to be treated specially. Typically, if the switch receives a stream with a destination address which is in the *policy file* in 'permit' mode, that stream is software flooded and no hardware entry is installed.

When adding an IPv6 address into the policy file, a 128-bit host address is recommended.

This feature is meant to solve the multicast connectivity problem for unknown destination addresses within system reserved ranges. Specifically this feature was introduced to solve the problem of recognizing a certain stream as control packets.

To create a policy file for the snooping flood-list, use the following template:

```
# This is a template for MLD Snooping Flood-list Policy File
# Add your group addresses between "Start" and "End"
# Do not touch rest of file!!!!
entry mldFlood {
  if match any {
    #----- Start of group addresses -----
    nlri ff05::100:1/128;
```



```

nlri ff05::100:15/128;
#----- end of group addresses -----
} then {
permit;
}
}
entry catch_all {
if {
} then {
deny;
}
}
}

```

Note



The switch does not validate any IP address in the policy file used in this command. Therefore, slow-path flooding should be used only for streams which are very infrequent, such as control packets. It should not be used for multicast data packets. This option overrides any default mechanism of hardware forwarding (with respect to MLD or PIM) so it should be used with caution.

Slow path flooding occurs within the L2 VLAN only.

Use the `none` option to effectively disable slow path flooding.

You can use the `show mld` command to see the configuration of slow path flooding.

Note



This command has no effect in the current release, since IPv6 multicast traffic floods on all platforms.

Example

The following command configures the multicast data stream specified in `access1` for slow path flooding:

```
configure mld snooping flood-list access1
```

The following command specifies that no policy file is to be used, this effectively disabling slow path flooding:

```
configure mld snooping flood-list none
```

History

This command was first available in ExtremeXOS 11.2.



Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure mld snooping leave-timeout

```
configure mld snooping leave-timeout leave_timeout_ms
```

Description

Configures the MLD snooping leave timeout.

Syntax Description

<i>leave_timeout_ms</i>	Specifies an MLD leave timeout value in milliseconds, upon receiving an MLD done message.
-------------------------	---

Default

1000 ms.

Usage Guidelines

The range is 0 - 175000 ms (175 seconds). For timeout values of one second or less, you must set the leave-timeout to a multiple of 100ms. For values of more than one second, you must set the leave-timeout to a multiple of 1000ms (one second).

The specified time is the maximum leave timeout value. The switch could leave sooner if an MLD done message is received before the timeout occurs.

Example

The following command configures the MLD snooping leave timeout:

```
configure mld snooping leave-timeout 10000
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure mld snooping timer

```
configure mld snooping timer router_timeout host_timeout
```

Description

Configures the MLD snooping timers.

Syntax Description

<i>router_timeout</i>	Specifies the time in seconds before removing a router snooping entry.
<i>host_timeout</i>	Specifies the time in seconds before removing a host's group snooping entry.

Default

The router timeout default setting is 260 seconds. The host timeout setting is 260 seconds.

Usage Guidelines

Timers should be set to approximately 2.5 times the router query interval in use on the network. Specify the following:

- *router_timeout*—The maximum time, in seconds, that a router snooping entry can stay without receiving a router report. The range is 10 to 214,748,364 seconds (6.8 years). The default setting is 260 seconds.
- *host_timeout*—The maximum time, in seconds, that a group snooping entry can stay without receiving a group report. The range is 10 to 214,748,364 seconds (6.8 years). The default setting is 260 seconds.

MLD snooping is a Layer 2 function of the switch. It does not require multicast routing to be enabled. The feature reduces the flooding of IPv6 multicast traffic. On the VLAN, MLD snooping optimizes the usage of network bandwidth and prevents multicast traffic from being flooded to parts of the network that do not need it. The switch does not reduce any IP multicast traffic in the local multicast domain (FF02::x).

MLD snooping is enabled by default on the switch. MLD snooping expects at least one device on every VLAN to periodically generate MLD query messages. Without an MLD querier, the switch eventually stops forwarding IPv6 multicast packets to any port, because the MLD snooping entries times out, based on the value specified in host timeout.

Example

The following command configures the MLD snooping timers:

```
configure mld snooping timer 600 600
```



History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable ipmcforwarding ipv6

```
disable ipmcforwarding ipv6 {{vlan} name }
```

Description

Disables IPv6 multicast forwarding on a router interface.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

Disabled.

Usage Guidelines

If no options are specified, all configured IPv6 interfaces are affected. When new IPv6 interfaces are created, IPv6 multicast forwarding is disabled by default.

Disabling IPv6 multicast forwarding disables any Layer 3 IPv6 multicast routing for the streams coming to the interface.

Example

The following command disables IPv6 multicast forwarding on VLAN accounting:

```
disable ipmcforwarding ipv6 vlan accounting
```

History

This command was first available in ExtremeXOS 15.3.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv6 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable mld

```
disable mld {vlan name}
```

Description

Disables MLD on a router interface. If no VLAN is specified, MLD is disabled on all router interfaces.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

Disabled.

Usage Guidelines

MLD is a protocol used by an IPv6 host to register its IPv6 multicast group membership with a router. Periodically, the router queries the multicast group to see if the group is still in use. If the group is still active, hosts respond to the query, and group registration is maintained.

MLD is disabled by default on the switch. However, the switch can be configured to enable the generation and processing of MLD packets. MLD should be enabled when the switch is configured to perform IPv6 unicast or IPv6 multicast routing.

This command disables all MLD versions. When MLD is disabled, the MLDv2 compatibility mode setting is lost. If compatibility mode is not specified in the command when MLD is enabled again, MLDv1 compatibility mode is set.

Example

The following command disables MLD on VLAN accounting:

```
disable mld vlan accounting
```

History

This command was first available in ExtremeXOS 11.2.



Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable mld snooping

```
disable mld snooping {with-proxy | vlan name}
```

Description

Disables MLD snooping.

Syntax Description

with-proxy	Disables the MLD snooping proxy.
<i>name</i>	Specifies a VLAN.

Default

The with-proxy option is enabled by default, but MLD snooping and forward-mcrouter-only option is disabled by default.

Usage Guidelines

If a VLAN is specified, MLD snooping is disabled only on that VLAN, otherwise MLD snooping is disabled on all VLANs.

The with-proxy option can be used for troubleshooting purpose. It should be enabled for normal network operation.

Enabling the proxy allows the switch to suppress the duplicate join requests on a group to forward to the connected Layer 3 switch. The proxy also suppresses unnecessary MLD done messages so that they are forwarded only when the last member leaves the group.

Example

The following command disables MLD snooping on the VLAN accounting:

```
disable mld snooping accounting
```

History

This command was first available in ExtremeXOS 11.2.



Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable ipmcforwarding ipv6

```
enable ipmcforwarding ipv6 {{vlan} name }
```

Description

Enables IPv6 multicast forwarding on a router interface.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

Disabled.

Usage Guidelines

If no options are specified, all configured IPv6 interfaces are affected. When new IPv6 interfaces are created, IPv6 multicast forwarding is disabled by default.

IPv6 forwarding must be enabled before enabling IPv6 multicast forwarding.

Example

The following command enables IPv6 multicast forwarding on VLAN accounting:

```
enable ipmcforwarding ipv6 vlan accounting
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv6 multicast feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable mld



```
enable mld {vlan vla name} {MLDv1 | MLDv2}
```

Description

Enables MLD on a router interface. If no VLAN is specified, MLD is enabled on all router interfaces.

Syntax Description

<i>vlan name</i>	Specifies a VLAN name.
MLDv1	Sets the compatibility mode to MLDv1.
MLDv2	Sets the compatibility mode to MLDv2.

Default

Disabled.

Usage Guidelines

MLD is a protocol used by an IPv6 host to register its IPv6 multicast group membership with a router. Periodically, the router queries the multicast group to see if the group is still in use. If the group is still active, IPv6 hosts respond to the query, and group registration is maintained.

MLD is disabled by default on the switch. However, the switch can be configured to enable the generation and processing of MLD packets. If compatibility mode is not specified in the command, MLDv1 compatibility mode is set.

A VLAN must have an IPv6 address to support MLD.

Example

The following command enables MLDv1 on the VLAN accounting:

```
enable mld vlan accounting
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Feature License Requirements](#) *ExtremeXOS Concepts Guide*.



enable mld snooping

```
enable mld snooping vlan name
```

Description

Enables MLD snooping on the switch.

Syntax Description

<i>name</i>	Specifies a VLAN.
-------------	-------------------

Default

Disabled.

Usage Guidelines

If a VLAN is specified, MLD snooping is enabled only on that VLAN, otherwise MLD snooping is enabled on all VLANs.

A VLAN must have an IPv6 address to support MLD.

Example

The following command enables MLD snooping on the switch:

```
enable mld snooping
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable mld snooping with-proxy

```
enable mld snooping with-proxy
```

Description

Enables the MLD snooping proxy.



Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Enabling the proxy allows the switch to suppress the duplicate join requests on a group to forward to the connected Layer 3 switch. The proxy also suppresses unnecessary MLD leave messages so that they are forwarded only when the last member leaves the group.

This command can be used for troubleshooting purpose. It should be enabled for normal network operation. The command does not alter the snooping setting.

Example

The following command enables the MLD snooping proxy:

```
enable mld snooping with-proxy
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show mcast ipv6 cache

```
show mcast ipv6 cache {{vlan} name} {[group v6GrpAddressMask | v6GrpAddressMask]  
{source v6SourceIP | v6SourceIP}} {type [snooping | pim]} {with-in-port} |  
{summary}}
```

Description

Displays multicast cache information. The display can be limited to entries for specific VLANs or groups, and it can be limited to specific types of entries, such as those created by snooping protocols, or PIM.



Syntax Description

<i>name</i>	Specifies a VLAN name.
<i>v6GrpAddressMask</i>	Specifies a multicast group address and mask.
<i>v6SourceIP</i>	Specifies the source IP address for a multicast group.
snooping	Limits the display to cache entries created by MLD snooping.
pim	Limits the display to cache entries created by PIM.
summary	Specifies the summary display format.

Default

Displays information for all entries in the multicast cache.

Usage Guidelines

None.

Example

The following command displays all multicast cache information:

```
show mcast ipv6 cache
Snooping Cache Timeout: 300 sec
(ff03::1 3001::1)
Type: snoop Age: 9 Ingress Vlan: v1
Vlan      Port      Vid
v1        25        4084
(ff03::1 3001::2)
Type: snoop Age: 9 Ingress Vlan: v1
Vlan      Port      Vid
v1        25        4084
Multicast cache distribution:
2 entries from Snooping          0 entries from PIM
Total Cache Entries: 2
```

The following command displays summary cache information for VLAN v1:

```
show mcast ipv6 cache vlan v1 summary
Snooping Cache Timeout: 300 sec
=====MULTICAST CACHE SUMMARY=====
Multicast cache distribution:
2 entries from Snooping          0 entries from PIM
v1: Multicast cache distribution:
2 entries from Snooping          0 entries from PIM
Total Cache Entries: 2
Total Cache Entries for VLAN v1: 2
*X480-48t.22 #
```



History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MLD snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show mld

```
show mld {vlan} {name}
```

Description

This command can be used to display an MLD-related configuration and group information, per VLAN or for the switch as a whole.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

N/A.

Usage Guidelines

If you do not specify a VLAN, the command displays the switch configuration.

Example

The following command displays the MLD configuration:

```
show mld
```

The following is sample output from this command:

```
show mld
VLAN          IP Address          Flags  nLRMA  nLeMA
MLDver
Default      ::/0                ---iz-    0      0
0
v1           ::/0                U--iz-    0      5
0
Flags:      (f) Forwarding Enabled, (g) Fast-learning on, (i) MLD Enabled,
```



(m) Multicast Forwarding Enabled, (U) Interface Up,
 (z) MLD Snooping Enabled.
 (nLeMA) Number of Learned Multicast Addresses
 (nLRMA) Number of Locally Registered Multicast Addresses

The following command displays the MLD configuration for VLAN v1:

```
show mld v1
Query Interval      : 125 sec
Max Response Time  : 10 sec
Last Member Query   : 1 sec
Robustness         : 2
Interface on VLAN v1 is enabled and up.
inet6 ::/0
Locally registered multicast addresses:
Learned multicast addresses (Last Querier=fe80::204:96ff:fe3a:ce50):
ff02::2             ff02::1:ff56:5c2b
ff02::1:ff00:2     ff02::1:ff3a:ce50
ff02::1:ff55:5c27
s = static MLD member
Flags:
IP Fwding  NO      IPmc Fwding  NO      MLD YES
MLD Ver   v0      Snooping    YES
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Feature License Requirements](#) *ExtremeXOS Concepts Guide*.

show mld counters

```
show mld counters {{vlan} name}
```

Description

Use this command to display an MLD packet statistics.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

N/A.



Usage Guidelines

The following command displays the MLD configuration:

```
* topleft.74 # show mld counters
```

MLD Message type	Received	Originated	Forwarded
MLD Query (v1/v2)	0	20	0
MLDv1 Report	499	0	157
MLDv1 Done	101	0	91
MLDv2 Report	0	0	0
Global Statistics:			
MLD Packet unknown	0		
MLD Packet Error	617		

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Feature License Requirements](#) *ExtremeXOS Concepts Guide*.

show mld group

```
show mld group {{vlan} {name} | {v6grpipaddress}} {MLDv2}
```

Description

Lists the MLD group membership for the specified VLAN or group.

Syntax Description

<i>name</i>	Specifies a VLAN name.
<i>v6grpipaddress</i>	Specifies a group IPv6 address.
MLDv2	Display the MLD group in MLDv2 format (if group record is MLDv2 compatible, otherwise display in earlier format). This option is not supported in this release.

Default

MLDv1.



Usage Guidelines

If no VLAN is specified all VLANs are displayed. You can also filter the display by group address and by multicast stream sender address.

Example

The following command lists the MLD group membership for the VLAN accounting:

```
show mld group vtest3
```

Output from this command looks similar to the following:

Group Address	Ver	Vlan	Port	Age
ff03::1:1	1	vtest3	4:5	25
ff03::1:2	1	vtest3	4:5	25
ff02::1:ff22:124	1	vtest3	4:45	26
ff05::a:abcd	1	vtest3	4:15	23
ff05::a:abce	1	vtest3	4:15	23
ff02::1:ff22:112	1	vtest3	4:45	26
ff02::1:ff1f:a418	1	vtest3	4:45	26

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Feature License Requirements](#) *ExtremeXOS Concepts Guide*.

show mld snooping

```
show mld snooping {vlan name | detail} {MLDv2}
```



Note

MLD snooping is not supported in this software release.

Description

Displays MLD snooping registration information and a summary of all MLD timers and states.



Syntax Description

<i>name</i>	Specifies a VLAN name.
detail	Displays the information in detailed format.
MLDv2	Display the MLD group in MLDv2 format (if group record is MLDv2 compatible, otherwise display in earlier format). This option is not supported in this release.

Default

MLDv1.

Usage Guidelines

The two types of MLD snooping entries are sender entry and subscribed entry.

The following information is displayed in a per-interface format:

- Group membership information
- Router entry
- Timeout information
- Sender entry

Example

Here is an example of the show output:

```
show mld snooping
MLD Snooping Flood-list : none
MLD Snooping Proxy      : Enable
MLD Snooping Filters    : per-port
Vlan                    Vid  Port  #Senders #Receivers Router Enable
-----
Default                 1    0      0         0         No      Yes
v1                      4084 0      0         0         No      Yes
25                      1    1      Yes        0         No
41                      2    2      No         0         No
42                      2    2      No         0         No
```

The following command displays MLD snooping registration information for the VLAN V1:

```
show mld snooping v1
Router Timeout          : 260 sec
Host Timeout           : 260 sec
MLD Snooping Fast Leave Time : 1000 ms
VLAN v1                (4084) Snooping=Enabled
Port  Host                                     Age
Subscribed                                     Join Limit
25    fe80::204:96ff:fe3a:ce50                   13
ff02::1:ff3a:ce50                               No Limit
25    fe80::204:96ff:fe3a:ce50                   14
```



```

All Groups                               No Limit
41    fe80::200:8ff:fe55:5c27             13
ff02::1:ff00:2                           No Limit
41    fe80::200:8ff:fe55:5c27             13
ff02::1:ff55:5c27                         No Limit
42    fe80::200:8ff:fe56:5c2b             14
ff02::1:ff00:2                           No Limit
42    fe80::200:8ff:fe56:5c2b             13
ff02::1:ff56:5c2b                         No Limit
s = static MLD member
* X480-48t.27 #

```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Feature License Requirements](#) in the ExtremeXOS Concepts Guide.

show mld snooping vlan filter

```
show mld snooping {vlan} name filter
```

Description

Displays MLD snooping filters..

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

None.

Usage Guidelines

Use this command to display MLD snooping filters configured on the specified VLAN. When no VLAN is specified, all the filters are displayed.

Example

The following command displays the MLD snooping filter configured on VLAN vlan101:

```
show mld snooping vlan101 filter
```



```
Filter Port Flags
mldpermit0 5:10 a
Flags: (a) Active
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MLD snooping feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show mld snooping vlan static

```
show mld snooping vlan name static [group | router]
```

Description

Displays static MLD snooping entries.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

None.

Usage Guidelines

Use this command to display the MLD snooping static groups or router ports configured on the specified VLAN. When no VLAN is specified, all the static groups or router ports are displayed.

Example

The following command displays the MLD snooping static groups configured on VLAN vlan101:

```
show mld snooping vlan101 static group
```

The following is sample output for this command:

```
Group                Port    Flags
ff03::1:1:1         7       sa
```



```
ff03::1:1:1          15          sa
Flags: (s) Static, (a) Active
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Feature License Requirements](#) *ExtremeXOS Concepts Guide*.

unconfigure mld

unconfigure mld

Description

Resets all MLD settings to their default values and clears the MLD group table.

Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command resets all MLD settings to their default values and clears the MLD group table:

```
unconfigure mld
```

History

This command was first available in ExtremeXOS 11.2.



Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Feature License Requirements](#) *ExtremeXOS Concepts Guide*.



46 MSDP Commands

```
clear msdp counters
clear msdp sa-cache
configure msdp as-display-format
configure msdp max-rejected-cache
configure msdp originator-id
configure msdp peer default-peer
configure msdp peer description
configure msdp peer mesh-group
configure msdp peer no-default-peer
configure msdp peer password
configure msdp peer sa-filter
configure msdp peer sa-limit
configure msdp peer source-interface
configure msdp peer timer
configure msdp peer ttl-threshold
configure msdp sa-cache-server
configure pim border
create msdp mesh-group
create msdp peer
delete msdp mesh-group
delete msdp peer
disable msdp
disable msdp data-encapsulation
disable msdp export local-sa
disable msdp peer
disable msdp process-sa-request
enable msdp
enable msdp data-encapsulation
enable msdp export local-sa
enable msdp peer
enable msdp process-sa-request
show msdp
show msdp memory
show msdp mesh-group
show msdp peer
show msdp sa-cache
unconfigure msdp sa-cache-server
```


unconfigure pim border

This chapter describes commands for doing the following:

- Configuring MSDP.
- Displaying MSDP information.

For an introduction to the MSDP feature, see the *ExtremeXOS Concepts Guide*.

clear msdp counters

```
clear msdp counters {peer remoteaddr | peer all | system} {vr vrname}
```

Description

This command resets the MSDP counters to zero.

Syntax Description

peer all	Specifies all MSDP peers.
<i>remoteaddr</i> <i>r</i>	Specifies the IP address of the MSDP peer.
system	Clears the global MSDP counters.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

The `clear msdp counters` command clears the following MSDP counters:

- Per peer counters:
 - Number of SA messages received.
 - Number of SA messages transmitted.
 - Number of SA request messages received.
 - Number of SA request messages transmitted.
 - Number of SA response messages received.
 - Number of SA response messages transmitted.
 - Number of SA messages received without encapsulated data.
 - Number of SA messages transmitted without encapsulated data.
 - Number of SA messages received with encapsulated data.
 - Number of SA messages transmitted with encapsulated data.
 - Number of times the MSDP peer attained an “ESTABLISHED” state.



- Number of times the peer-RPF check failed.
- Number of times the TCP connection attempt failed.
- Total number of received messages.
- Total number of transmitted messages.
- Global counters:
 - None defined.

The clear counters command will also clear all MSDP counters, but it clears the counters for all other applications too.

Example

The following command clears the counters for an MSDP peer with the IP address 192.168.45.43:

```
clear msdp counters peer 192.168.45.43
```

The following command clears the all peer and global counters:

```
clear msdp counters
```

The following command clears all counters for a particular peer:

```
clear msdp counters peer 192.168.32.45
```

The following command clears the counters of all MSDP peers:

```
clear msdp counters peer all
```

The following command clears the global counters:

```
clear msdp counters system
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



clear msdp sa-cache

```
clear msdp sa-cache [{peer} remoteaddr | peer all] {group-address grp-addr} {vr
vrname}
```

Description

This command purges all SA cache entries and notifies the PIM that the SA cache is empty.

Syntax Description

peer all	Specifies all MSDP peers. All matching SA cache entries from all peers are removed from the database.
<i>grp-addr</i>	Specifies the IP address and subnet mask of the multicast group you want to clear. All SA cache entries that match the specified group address are removed from the database.
<i>remoteaddr</i>	Specifies the IP address of the MSDP peer. All matching SA cache entries learned from the specified peer are removed from the database.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

MSDP receives SA messages periodically. So, after clearing SA cache entries from the local database, MSDP relearns those entries during the next advertisement from its peer.

Example

The following command clears SA cache records for an MSDP peer with the IP address 192.168.45.43:

```
clear msdp sa-cache peer 192.168.45.43
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure msdp as-display-format

```
configure msdp as-display-format [asdot | asplain]
```

Description

Configures the AS number format displayed in show commands.

Syntax Description

asdot	Specifies the ASDOT format.
asplain	Specifies the ASPLAIN format.

Default

N/A.

Usage Guidelines

The ASPLAIN and ASDOT formats are described in RFC 5396, Textual Representation of Autonomous System (AS) Numbers.

Example

The following command selects the ASDOT 4-byte AS number format:

```
configure msdp as-display-format asdot
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure msdp max-rejected-cache

```
configure msdp max-rejected-cache max-cache {vr vrname}
```



Description

Configures the maximum limit on rejected SA cache entries that an MSDP router will store in its database.

Syntax Description

<i>max-cache</i>	Specifies the maximum number of rejected SA cache entries that the MSDP router will store in its database. To remove the limit, enter 0 (zero) for the <i>max-cache</i> value.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, the maximum cache entries stored is zero. That is, rejected SA cache entries are not stored. Any SA cache entries that are stored and not refreshed for six minutes are removed.

Usage Guidelines

SA cache are rejected because of:

- Peer-RPF failure.
- Policy denied.

When a previously rejected SA cache entry is accepted because of an RP reachability change or policy rule change, the rejected SA cache entry is moved to the accepted SA cache list.

By default, rejected SA cache entries are discarded. You can configure a limit for rejected cache entries to store them, which will help debug/diagnose some issues; however, it consumes extra memory.

Example

The following command sets the maximum rejected cache limit to 100 for an MSDP router:

```
configure msdp max-rejected-cache 100
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure msdp originator-id

```
configure msdp originator-id ip-address {vr vrname}
```

Description

Configures the originator ID for an MSDP router. The originator ID is the RP address you want to use (instead of the default) in locally originated SA messages.

Syntax Description

<i>ip-address</i>	Specifies the RP address to use in locally originated SA messages. To unconfigure an originator ID (that is, to use the default RP address), enter the IP address 0.0.0.0.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, the RP address is used as the originator ID in locally originated SA messages.

Usage Guidelines

Use this command to override the default RP address used in SA messages. Because only RPs and MSDP border routers originate SAs, there are times when it is necessary to change the ID used for this purpose. The originator ID address must be one of the interface addresses on the MSDP router.

You can configure the MSDP originator ID only when MSDP is disabled globally.

To remove an originator ID, enter the IP address 0.0.0.0.

Example

The following command configures the originator ID for an MSDP router:

```
configure msdp originator-id 10.203.134.1
```

The following command unconfigures the originator ID for an MSDP router:

```
configure msdp originator-id 0.0.0.0
```

History

This command was first available in ExtremeXOS 12.0.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure msdp peer default-peer

```
configure msdp peer [remoteaddr | all] default-peer {default-peer-policy filter-name} {vr vrname}
```

Description

This command configures a default or static RPF peer from which all MSDP SA messages are accepted. To remove the default peer, enter the `configure msdp peer no-default-peer` command.

Syntax Description

<i>filter-name</i>	Specifies the name of the policy filter associated with the default peer. The peer will be the default peer for all SA entries that are permitted by the policy filter. If an SA message is allowed by the policy filter, it will be accepted. Otherwise, the SA message has to go through the regular RPF-check. The static peer RPF check is the last step in peer RPF algorithm. So, if an SA message is denied by the default peer policy, ultimately the SA message will be rejected by MSDP.
peer all	Specifies all MSDP peers.
<i>remoteaddr</i> <i>r</i>	Specifies the IP address of the MSDP peer.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, no static RPF peer is configured.

The **default-peer-policy** keyword specifies the name of the policy filter associated with the default peer. You can configure multiple default peers with different policies. If no policy is specified, then the current peer is the default RPF peer for all SA messages.

Usage Guidelines

Configuring a default peer simplifies peer-RPF checking of SA messages. If the peer-RPF check fails, the default peer rule is applied to see if the SA messages should be accepted or rejected.

If a default peer policy is specified, the peer is the default peer only for the (Source, Group), or (S, G), that satisfies the policy. If the policy is not specified, then the default peer is used for all (S, G, RP).

You can configure multiple default peers on an MSDP router; however all default peers must either have a default policy or not. A mix of default peers, with a policy and without a policy, is not allowed.

When configuring multiple default peer rules, follow these guidelines:



- When you enter multiple default-peer commands with the default-peer-policy keyword, you can use all the default peers at the same time for different RP prefixes.
- When you enter multiple default-peer commands without the default-peer-policy keyword, you can use a single active peer to accept all SA messages. If that peer goes down, then the next configured default peer accepts all SA messages. This configuration is typically used at a stub site.

You can use the following policy attributes in a default peer policy. All other attributes are ignored.

- Match:
 - multicast-group.
 - multicast-source.
 - pim-rp.
- Set:
 - permit.
 - deny.

Example

The following command configures an MSDP peer with the IP address 192.168.45.43 as the default peer policy for “sales”:

```
configure msdp peer 192.168.45.43 default-peer default-peer-policy sales
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure msdp peer description

```
configure msdp peer remoteaddr description {peer-description} {vr vrname}
```

Description

Configures a name or description for an MSDP peer. This text is for display purposes only.



Syntax Description

<i>remoteaddr</i>	Specifies the IP address of the MSDP peer.
<i>peer-description</i>	Specifies the name or description of the MSDP peer. The maximum is 63 characters.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, no name or description is specified.

Usage Guidelines

Use this command to configure a name or description to make an MSDP peer easier to identify. The description is visible in the output of the `show msdp peer` command.

To remove the description, use this command without a description string.

Example

The following command configures the name “internal_peer” to an MSDP peer:

```
configure msdp peer 192.168.45.43 description internal_peer
```

The following command removes the description from an MSDP peer:

```
configure msdp peer 192.168.45.43 description
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure msdp peer mesh-group

```
configure msdp peer [remoteaddr | all] mesh-group [mesh-group-name | none] {vr  
vrname}
```



Description

This command configures an MSDP peer to become a member of a mesh-group. To remove a peer from a mesh-group, enter the none CLI keyword for the mesh-group.

Syntax Description

<i>mesh-group-name</i>	Specifies the name of the MSDP mesh-group.
none	Removes a peer from a mesh-group.
peer all	Specifies all MSDP peers.
<i>remoteaddr</i>	Specifies the IP address of the MSDP peer.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

A mesh-group is a group of MSDP peers with fully meshed MSDP connectivity. Any SA messages received from a peer in a mesh-group are not forwarded to other peers in the same mesh-group.

Mesh-groups achieve two goals:

- Reduce SA message flooding.
- Simplify peer-RPF flooding.

Example

The following command configures an MSDP peer with the IP address 192.168.45.43 to become a member of a mesh-group called "intra":

```
configure msdp peer 192.168.45.43 mesh-group intra
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure msdp peer no-default-peer

```
configure msdp peer [remoteaddr | all] no-default-peer {vr vrname}
```

Description

This command removes a default peer.

Syntax Description

peer all	Specifies all MSDP peers.
<i>remoteaddr</i>	Specifies the IP address of the MSDP peer.
no-default-peer	Removes a default peer.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

N/A.

Example

The following command removes all MSDP peers:

```
configure msdp peer all no-default-peer
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure msdp peer password

```
configure msdp peer [remoteaddr | all] password [none | {encrypted} tcpPassword]
{vr vrname}
```

Description

This command configures a TCP RSA Data Security, Inc. MD5 Message-Digest Algorithm password for an MSDP peer. This command enables TCP RSA Data Security, Inc. MD5 Message-Digest Algorithm authentication for a MSDP peer. When a password is configured, MSDP receives only authenticated MSDP messages from its peers. All MSDP messages that fail TCP RSA Data Security, Inc. MD5 Message-Digest Algorithm authentication are dropped.

Syntax Description

peer all	Specifies all MSDP peers.
<i>remoteaddr</i>	Specifies the IP address of the MSDP peer.
none	Removes the previously configured password.
encrypted	Encrypts the password for RSA Data Security, Inc. MD5 Message-Digest Algorithm authentication. To improve security, the password displays in encrypted format and cannot be seen as simple text. Additionally, the password is saved in encrypted format.
<i>tcpPassword</i>	Specifies the password to use for RSA Data Security, Inc. MD5 Message-Digest Algorithm authentication at the TCP level. The password must be an ASCII string with a maximum of 31 characters.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Defaults

By default, TCP RSA Data Security, Inc. MD5 Message-Digest Algorithm authentication is disabled for the MSDP peer.

Usage Guidelines

We recommend that you enable TCP RSA Data Security, Inc. MD5 Message-Digest Algorithm authentication for all MSDP peers to protect MSDP sessions from attacks. You can execute this command only when the MSDP peer is disabled or when MSDP is globally disabled on that VR.

Example

The following command configures a password for the MSDP peer with the IP address 192.168.45.43, which automatically enables TCP MD5 authentication:

```
configure msdp peer 192.168.45.43 password test123
```



The following command removes the password:

```
configure msdp peer 192.168.45.43 password none
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure msdp peer sa-filter

```
configure msdp peer [remoteaddr | all] sa-filter [in | out] [filter-name | none]
{vr vr_name}
```

Description

This command configures an incoming or outgoing policy filter for SA messages.

Syntax Description

peer all	Specifies all MSDP peers.
<i>remoteaddr</i>	Specifies the IP address of the MSDP peer.
in	Associates the SA filter with inbound SA messages.
out	Associates the SA filter with outbound SA messages.
<i>filter-name</i>	Specifies the name of the policy associated with an SA filter. To remove an SA filter, enter the none CLI keyword for <i>filter-name</i> .
<i>vr_name</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, no SA filter is configured for an MSDP peer. That is, incoming and outgoing SA messages are not filtered.

Usage Guidelines

This command configures an SA filter such that only a specified set of SA messages are accepted or sent to a peer. Note that an SA filter does not adversely impact the flow of SA request and response messages.



To remove an SA filter, enter the **none** CLI keyword for *filter-name*.

You can use the following policy attributes in an SA filter policy. All other attributes are ignored.

- Match:
 - multicast-group.
 - multicast-source.
 - pim-rp.
- Set:
 - permit.
 - deny.

Example

The following command configures an incoming SA messages filter on an MSDP peer with the IP address 192.168.45.43:

```
configure msdp peer 192.168.45.43 sa-filter in allow_229
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure msdp peer sa-limit

```
configure msdp peer [remoteaddr | all] sa-limit max-sa {vr vr_name}
```

Description

This command allows you to limit the number of SA entries from an MSDP peer that the router will allow in the SA cache. To allow an unlimited number of SA entries, use 0 (zero) as the value for *max-sa*.

Syntax Description

peer all	Specifies all MSDP peers.
<i>remoteaddr</i>	Specifies the IP address of the MSDP peer.



<i>max-sa</i>	Specifies the maximum number of SA entries from an MSDP peer allowed in the SA cache. To specify an unlimited number of SA entries, use 0 (zero) as the value for <i>max-sa</i> .
<i>vr_name</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, no SA entry limit is set. The router can receive an unlimited number of SA entries from an MSDP peer.

Usage Guidelines

You can use this command to prevent a distributed denial of service (DOS) attack. We recommend that you configure an MSDP SA limit on all MSDP peer sessions. Note that a rejected SA cache entry is not included in the number of SA cache entries received from a peer.

Example

The following command configures the SA entry limit of 500 for the MSDP peer with the IP address 192.168.45.43:

```
configure msdp peer 192.168.45.43 sa-limit 500
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure msdp peer source-interface

```
configure msdp peer [remoteaddr | all] source-interface [ipaddress | any] {vr  
vrname}
```

Description

This command configures the source interface for the MSDP peer TCP connection.



Syntax Description

<code>peer all</code>	Specifies all MSDP peers.
<code>remoteaddress</code>	Specifies the IP address of the MSDP peer.
<code>ipaddress</code>	Specifies the IP address of the MSDP router interface to use on one end of a TCP connection. The <code>ipaddress</code> must be one of the MSDP router interface addresses; otherwise, the command fails and an error message displays.
<code>any</code>	Specifies to use any interface as one end of the TCP connection. The source interface is selected based on the IP route entry used to reach the MSDP peer. The egress interface that reaches the MSDP peer is used as the source interface for the TCP connection. Basically, this command removes the previously configured source interface of the MSDP peer.
<code>vrname</code>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Defaults

By default, the source interface is selected based on the IP route entry used to reach the MSDP peer. The egress interface that reaches the MSDP peer is used as the source interface for the TCP connection.

Usage Guidelines

You must first disable MSDP or the MSDP peer before using this command. Extreme Networks recommends that you configure a source interface for MSDP peers that are not directly connected. We also recommend using the loopback address as the MSDP peer connection endpoint.

Example

The following command configures a source interface for an MSDP peer with the IP address 192.168.45.43:

```
configure msdp peer 192.168.45.43 source-interface 60.0.0.5
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



configure msdp peer timer

```
configure msdp peer [remoteaddr | all] timer keep-alive keep-alive-sec hold-time
hold-time-sec {vr vrname}
```

Description

The command configures the keep-alive and hold timer intervals of the MSDP peers.

Syntax Description

<i>peer all</i>	Specifies all MSDP peers.
<i>remoteaddr</i>	Specifies the IP address of the MSDP peer.
<i>keep-alive-sec</i>	Specifies the keep-alive timer interval in seconds, in the range of 1 through 60 seconds.
<i>hold-time-sec</i>	Specifies the hold timer interval in seconds, in the range of 3 through 75 seconds.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, the:

- Keep-alive timer interval is 60 seconds.
- Hold timer interval is 75 seconds.
- SA timer interval is 60 seconds.

Usage Guidelines

You can use this command only when either MSDP or the MSDP peer is disabled. The hold timer interval must be greater than the keep-alive timer interval.

Example

The following command configures the keep-alive and hold timer intervals for the MSDP peer 55.0.0.83:

```
configure msdp peer 55.0.0.83 timer keep-alive 30 hold-time 60
```

History

This command was first available in ExtremeXOS 12.0.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure msdp peer ttl-threshold

```
configure msdp peer [remoteaddr | all] ttl-threshold t1 {vr vrname}
```

Description

Configures the limit to which multicast data packets are sent in SA messages to an MSDP peer. If the time-to-live (TTL) in the IP header of an encapsulated data packet exceeds the TTL threshold configured, encapsulated data is not forwarded to MSDP peers.

Syntax Description

<i>remoteaddr</i> <i>r</i>	Specifies the IP address of the MSDP peer on which to configure a TTL threshold.
all	Specifies all MSDP peers.
<i>t1</i>	Specifies the TTL value. The range is 0 through 255. To restore the default value, enter a TTL value of 0 (zero).
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

The default value is zero, meaning all multicast data packets are forwarded to the peer regardless of the TTL value in the IP header of the encapsulated data packet.

Usage Guidelines

This command allows you to configure a TTL value to limit multicast data traffic.

Example

The following command configures a TTL threshold of 5:

```
configure msdp peer 192.168.45.43 ttl-threshold 5
```

History

This command was first available in ExtremeXOS 12.0.



Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure msdp sa-cache-server

```
configure msdp sa-cache-server remoteaddr {vr vr_name}
```

Description

Configures the MSDP router to send SA request messages to the MSDP peer when a new member becomes active in a group.

Syntax Description

<i>remoteaddr</i>	Specifies the IP address of the MSDP peer from which the local router requests SA messages when a new member becomes active in a group, and MSDP has no cache entry for the group in the local database.
<i>vr_name</i>	Specifies the name of the virtual router on which the MSDP cache server is configured. If a virtual router name is not specified, it is extracted from the current CLI context.

Default

By default, the router does not send SA request messages to its MSDP peers when a new member joins a group and wants to receive multicast traffic. The new member simply waits to receive SA messages, which eventually arrive.

Usage Guidelines

You can use this command to force a new member of a group to learn the current active multicast sources in a connected PIM-SM domain that are sending to a group. The router will send SA request messages to the specified MSDP peer when a new member joins a group and MSDP doesn't have a cache entry for that group in the local database. The peer replies with the information in an SA cache response message.

Note



An MSDP peer must exist before it can be configured as an SA cache server. The `configure msdp sa-cache-server` command accepts the value for *remoteaddr* only if it is an existing peer's IP address.



Example

The following command configures an MSDP cache server:

```
configure msdp sa-cache-server 172.19.34.5
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

configure pim border

```
configure pim vlan_name border
```

Description

Configures a PIM VLAN as a border VLAN, which is used to demarcate a PIM domain when using MSDP.

Syntax Description

<i>vlan_name</i> Specifies a VLAN name.

Default

N/A.

Usage Guidelines

MSDP is used to connect multiple multicast routing domains. A PIM-SM domain is created by limiting the reach of PIM BSR advertisements. When a border VLAN is configured, PIM BSR advertisements are not forwarded out of the PIM VLAN.

Example

The following command configures a PIM border on a VLAN called “vlan_border”:

```
configure pim vlan_border border
```



History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

create msdp mesh-group

```
create msdp mesh-group mesh-group-name {vr vrname}
```

Description

Creates an MSDP mesh-group.

Syntax Description

<i>mesh-group-name</i>	Specifies the name for the MSDP mesh-group.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

A mesh-group is a group of MSDP peers with fully meshed MSDP connectivity. Create a mesh-group to:

- Reduce SA message flooding.
- Simplify peer-RPF flooding.

SA messages received from a peer in a mesh-group are not forwarded to other peers in the same mesh-group, which reduces SA message flooding.

A mesh group name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see [Object Names](#) in the *ExtremeXOS Concepts Guide*.



Example

The following command creates a mesh-group called “verizon”:

```
create msdp mesh-group verizon
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

create msdp peer

```
create msdp peer remoteaddr {remote-as remote-AS} {vr vrname}
```

Description

Creates an MSDP peer.

Syntax Description

remoteaddr <i>r</i>	Specifies the IP address of the MSDP router to configure as an MSDP peer.
remote-AS	Specifies the autonomous system (AS) number of the MSDP peer. This optional parameter is deprecated in ExtremeXOS 12.1, though the option is still available in the CLI for backward compatibility. The software ignores this parameter.
vrname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

The BGP route database is used by MSDP to determine the AS number for the peer. You can display the AS number (which can be a 2-byte for 4-byte AS number) using the command: `show msdp [peer {detail} | {peer} remoteaddr] {vrvrname}`.



Example

The following command creates an MSDP peer:

```
create msdp peer 192.168.45.43 remote-as 65001
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

delete msdp mesh-group

```
delete msdp mesh-group mesh-group-name {vr vrname}
```

Description

Removes an MSDP mesh-group.

Syntax Description

<i>mesh-group-name</i>	Specifies the name of the MSDP mesh-group. The character string can be a maximum of 31 characters.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

A mesh-group is a group of MSDP peers with fully meshed MSDP connectivity. Mesh-groups are used to achieve two goals:

- Reduce SA message flooding.
- Simplify peer-RPF flooding.

SA messages received from a peer in a mesh-group are not forwarded to other peers in the same mesh-group.



Use the `delete msdp mesh-group` command only if you created a mesh-group that you want to remove. By default, there is no MSDP mesh-group.

Example

The following command removes a mesh-group called “verizon”:

```
delete msdp mesh-group verizon
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

delete msdp peer

```
delete msdp peer [all | remoteaddr] {vr vr_name}
```

Description

Deletes an MSDP peer.

Syntax Description

all	Deletes all MSDP peers.
<i>remoteaddr</i>	Specifies the IP address of the MSDP router to configure as an MSDP peer.
<i>vr_name</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

None.



Example

The following command deletes an MSDP peer:

```
delete msdp peer 192.168.45.43
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable msdp

```
disable msdp {vr vrname}
```

Description

Disables MSDP on a virtual router.

Syntax Description

<i>vrname</i>	Specifies the name of the virtual router on which MSDP is being enabled or disabled. If a name is not specified, it is extracted from the current CLI context.
----------------------	--

Default

MSDP is disabled by default.

Usage Guidelines

Use this command to disable MSDP on a virtual router.

Example

The following command disables MSDP on a virtual router:

```
disable msdp
```



History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable msdp data-encapsulation

```
disable msdp data-encapsulation {vr vrname}
```

Description

Disables the encapsulation of locally originated SA messages with multicast data (if available).

Syntax Description

<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.
----------------------	---

Default

By default, multicast data packet encapsulation is enabled for locally originated SA messages.

Usage Guidelines

N/A.

Example

The following command disables multicast data packet encapsulation:

```
disable msdp data-encapsulation
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



disable msdp export local-sa

```
disable msdp export local-sa {vr vrname}
```

Description

Disables the advertisement of local sources to groups for which the router is an RP.

Syntax Description

vrname	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.
---------------	---

Default

By default, the export of local sources is enabled. All sources are advertised if the router is an RP for the groups. Use this command to disable it.

Usage Guidelines

You can create a policy to filter out some of the local sources so that they are not advertised to MSDP peers and exposed to the external multicast domain. To configure an export filter, you must first disable the export of local sources (with the `disable msdp export local-sa` command), and then re-enable it with an export filter (with the `enable msdp export local-sa export-filter` command).

Example

The following command disables the advertisement of local sources:

```
disable msdp export local-sa
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable msdp peer

```
disable msdp [{peer} remoteaddr | peer all] {vr vr_name}
```



Description

Configures the administrative state of an MSDP peer.

Syntax Description

<i>remoteaddr</i>	Specifies the IP address of the MSDP peer to disable.
all	Disables all MSDP peers.
<i>vr_name</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, MSDP peers are disabled.

Usage Guidelines

Use this command to administratively disable MSDP peers to stop exchanging SA messages.

Example

The following command disables an MSDP peer:

```
disable mstp peer 192.168.45.43
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

disable mstp process-sa-request

```
disable mstp [{peer} remoteaddr | peer all] process-sa-request {vr vrname}
```

Description

This command configures a router to reject SA request messages from a specified peer or all peers.



Syntax Description

peer all	Specifies all MSDP peers.
<i>remoteaddr</i> <i>r</i>	Specifies the IP address of the MSDP peer.
<i>vrrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, all SA request messages are accepted from all peers.

Usage Guidelines

Use this command to configure the router to reject SA request messages from a specified peer or all peers.

You cannot change an SA request filter while SA request processing is enabled for an MSDP peer. You must first disable SA request processing for a peer and then re-enable it with an SA request filter.

You can use the following policy attributes in an SA request policy. All other attributes are ignored.

- Match:
 - multicast-group.
 - multicast-source.
 - pim-rp.
- Set:
 - permit.
 - deny.

Example

The following command disables processing of SA request messages received from a peer with the IP address 192.168.45.43:

```
disable msdp peer 192.168.45.43 process-sa-request
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



enable msdp

```
enable msdp {vr vrname}
```

Description

Enables MSDP on a virtual router.

Syntax Description

vrname	Specifies the name of the virtual router on which MSDP is being enabled or disabled. If a name is not specified, it is extracted from the current CLI context.
---------------	--

Default

MSDP is disabled by default.

Usage Guidelines

Use this command to enable MSDP on a virtual router.

Example

The following command enables MSDP on a virtual router:

```
enable msdp
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable msdp data-encapsulation

```
enable msdp data-encapsulation {vr vrname}
```

Description

Enables the encapsulation of locally originated SA messages with multicast data (if available).



Syntax Description

<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.
----------------------	---

Default

By default, multicast data packet encapsulation is enabled for locally originated SA messages. Multicast data packets with a packet size of up to 8 KB are encapsulated in SA messages.

Usage Guidelines

Enable data encapsulation to handle bursty sources.

Example

The following command enables multicast data packet encapsulation:

```
enable msdp data-encapsulation
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable msdp export local-sa

```
enable msdp export local-sa {export-filter filter-name} {vr vrname}
```

Description

Enables the advertisement of local sources to groups for which the router is an RP.

Syntax Description

<i>filter-name</i>	Specifies the policy to associate with the export of local sources. No policy is specified by default.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.



Default

By default, the export of local sources is enabled. All sources are advertised if the router is an RP for the groups.

Usage Guidelines

You can create a policy to filter out some of the local sources so that they are not advertised to MSDP peers and exposed to the external multicast domain. To configure an export filter, you must first disable the export of local sources (with the `disable msdp export local-sa` command), and then re-enable it with an export filter (with the `enable msdp export local-sa export-filter` command).

You can use the following policy attributes in an export policy. All other attributes are ignored.

- Match:
 - multicast-group.
 - multicast-source.
 - pim-rp.
- Set:
 - permit.
 - deny.

Please note that the syntax for “multicast-group”, “multicast-source,” and “pim-rp” are the same as for the “nlri” policy attribute.

```
[multicast-group | multicast-source | pim-rp] [<ipaddress> | any]/<mask-length> {exact}
[multicast-group | multicast-source | pim-rp] [<ipaddress> | any] mask <mask> {exact}
```

An example of an MSDP policy file follows:

```
entry allow_internal_rp {
  if match any {
    multicast-group 234.67.89.0/24;
    multicast-source 23.123.45.0/24;
    pim-rp 10.203.134.5/32;
  } then {
    permit;
  }
}
entry deny_local_group239 {
  if match any {
    multicast-group 239.0.0.0/8;
    multicast-source 23.123.45.0/24;
  } then {
    deny;
  }
}
entry allow_external_rp_172 {
  if {
```




```

multicast-group 234.172.0.0/16;
} then {
permit
}
}
# deny remaining entries

```

Example

The following command enables the advertisement of local sources:

```
enable msdp export local-sa
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable msdp peer

```
enable msdp [{peer} remoteaddr | peer all] {vr vr_name}
```

Description

Configures the administrative state of an MSDP peer.

Syntax Description

all	Enables all MSDP peers.
<i>remoteaddr</i>	Specifies the IP address of the MSDP peer to configure.
<i>vr_name</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, MSDP peers are disabled.



Usage Guidelines

You must use this command to administratively enable the MSDP peers before they can establish peering sessions and start exchanging SA messages.

Example

The following command enables an MSDP peer:

```
enable msdp peer 192.168.45.43
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

enable msdp process-sa-request

```
enable msdp [{peer} remoteaddr | peer all] process-sa-request {sa-request-filter
filter-name } {vr vr_name}
```

Description

This command configures MSDP to receive and process SA request messages from a specified peer or all peers. If an SA request filter is specified, only SA request messages from those groups permitted are accepted. All others are ignored.

Syntax Description

peer all	Specifies all MSDP peers.
<i>filter-name</i>	Specifies the name of the policy filter associated with SA request processing.
<i>remoteaddr</i>	Specifies the IP address of the MSDP peer.
<i>vr_name</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, all SA request messages are accepted from peers.



Usage Guidelines

Use this command to configure the router to accept all or just some SA request messages from peers. If no policy is specified, all SA request messages are accepted. If a policy is specified, only SA request messages from those groups permitted are accepted, and all others are ignored.

You cannot change an SA request filter while SA request processing is enabled for an MSDP peer. You must first disable SA request processing for a peer and then re-enable it with an SA request filter.

You can use the following policy attributes in an SA request policy. All other attributes are ignored.

- Match:
 - multicast-group.
 - multicast-source.
 - pim-rp.
- Set:
 - permit.
 - deny.

Example

The following command enables processing of SA request messages received from a peer with the IP address 192.168.45.43:

```
enable msdp peer 192.168.45.43 process-sa-request sa-request-filter
intra_domain
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show msdp

```
show msdp {vr vrname}
```

Description

This command displays global configuration and run-time parameters for MSDP.



Syntax Description

<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.
----------------------	---

Default

N/A.

Usage Guidelines

Use this command to verify the global configuration parameters of MSDP.

Example

The following command displays global configuration and run-time parameters for MSDP:

```
Switch.2 # show msdp
MSDP Enabled           : No           VR-Name                : VR-Default
Originator RP Addr    : not configured SA Cache ageout time  : 360
Store SA Cache        : Yes           SA Cache Server        : not
configured
Export Local SAs      : Yes           Export SA filter       : not
configured
Max Rejected Cache    : not configured Encapsulate data       : Yes
Num of Rejected SAs   : 0             Total Num of SAs      : 0
Num of Local SAs      : 0             AS Disp Format         : Asdot
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show msdp memory

```
show msdp memory {detail | memoryType}
```

Description

This command displays current memory utilization of the MSDP process, including all virtual router instances of the MSDP process.



Syntax Description

detail	Displays detailed statistics for all memory types.
<i>memoryType</i>	Displays statistics for a particular memory type.

Default

N/A.

Usage Guidelines

Use this command to view and diagnose the memory utilization of the MSDP process.

Example

The following displays current memory utilization of the MSDP process, including all virtual router instances of the MSDP process:

```
show msdp memory
```

The following is sample output from this command:

```
MSDP Memory Information
-----
Bytes Allocated: 79792  AllocFailed: 0  OversizeAlloc: 0
Current Memory Utilization Level: GREEN
Memory Utilization Statistics
-----
Size      16      32      48      64      80      96      128     256     1024    4096
8192  12288
-----
-----  Used Blocks  0      0      256     263      3
0       2       0       1       0       0       4       0       0       0
peer    0       0       0       0       0       0       0       0       0
0       0       4
mesh-group  0       0       0       3       0       0       0       0       0
0       0       0
sa-node  0       0       0       255     0       0       0       0       0
0       0       0
sa-entry 0       0       255     0       0       0       0       0       0
0       0       0
vr-node  0       0       0       0       0       0       0       0       0       1
0       0       0
rt-cache 0       0       0       5       0       0       0       0       0
0       0       0
rp-node  0       0       1       0       0       0       0       0       0
0       0       0
client  0       0       0       0       0       0       2       0       0
0       0       0
misc    0       0       0       0       3       0       0       0       0
0       0       0
```



History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show msdp mesh-group

```
show msdp [mesh-group {detail} | {mesh-group} mesh-group-name] {vr vrname}
```

Description

This command displays configuration information about MSDP mesh-groups.

Syntax Description

detail	Displays detailed information about MSDP mesh-groups.
<i>mesh-group-name</i>	Specifies the name of the MSDP mesh-group. The character string can be a maximum of 31 characters.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

Use this command to display configuration information about MSDP mesh-groups, as follows:

- For summary information, enter the show msdp mesh-group command.
- For detailed information, enter the show msdp mesh-group detail command.
- For detailed information about a specific mesh-group, enter the show msdp mesh-group *name* command.

Example

The following command displays the peer count for a mesh-group:

```
show msdp mesh-group
```



The following is sample output from this command:

```

MeshGroupName                PeerCount
-----
external                      0
internal                      0
msdp_mesh                    4

```

The following command displays detailed information about a mesh-group called "msdp_mesh":

```
show msdp mesh-group "msdp_mesh"
```

The following is sample output from this command:

```

Mesh Group Name      : msdp_mesh Num of Peers : 4
Peers                : 54.172.168.97  55.0.0.83      124.56.78.90
221.160.90.228

```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show msdp peer

```
show msdp [peer {detail} | {peer} remoteaddr] {vr vr_name}
```

Description

This command displays configuration and run-time parameters about MSDP peers.

Syntax Description

detail	Displays detailed information about MSDP peers.
<i>remoteaddr</i>	Specifies the IP address of the MSDP peer.
<i>vr_name</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.



Default

N/A.

Usage Guidelines

Use this command to verify the configuration and run-time parameters for MSDP peers, as follows:

- For summary information, enter the `show msdp peer` command.
- For detailed information for all peers, enter the `show msdp peer detail` command.
- For detailed information for a specific peer, enter the `show msdp peer remoteaddr` command.

Example

The following command displays configuration and run-time parameters for MSDP peers:

```
show msdp peer
```

The following is sample output from this command:

```
Peer Address      AS      State      Up/Down      Resets      SA_Cnt      Name
-----
-d 54.172.168.97  14490  DISABLED   00:31:36    0           0           test
*e 55.0.0.83      100    ESTABLISHED 00:21:04    1           0           to-Hawaii
-d 124.56.78.90   2345   DISABLED   00:31:36    0           0
-d 221.160.90.228 23456  DISABLED   00:31:36    0           0
Flags: (*) default peer, (d) disabled, (e) enabled
```

The following command displays output from an MSDP peer with the IP address 16.0.0.2:

```
* Switch.8 # show msdp peer 16.0.0.2
MSDP Peer           : 16.0.0.2
Enabled             : No
Keepalive Interval  : 60
Source Address      : not known
Default Peer        : No
configured
Process In Request  : Yes
configured
Maximum SA Limit    : not configured
Input SA Filter     : not configured
configured
State               : DISABLED
Local Port          : 0
In Total Msgs       : 0
In SA Msgs          : 0
In SA Req Msgs     : 0
In SA Resp Msgs    : 0
Time since Last Msg : 00:00:02
Connection Attempts : 0
RPF Fails           : 0
AS Number           : 100.100
Holdtimer Interval  : 75
TTL Threshold       : 0
Default Peer Filter : not
In Request filter   : not
Mesh Group          : not
Output SA Filter    : not
Uptime/Downtime    : 00:00:02
Remote Port         : 0
Out Total Msgs     : 0
Out SA Msgs        : 0
Out SA Req Msgs    : 0
Out SA Resp Msgs   : 0
Hold Tmr Exp in    : 00:00:00
Entered Established : 0
Output Queue Size   : 0
```



History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

show msdp sa-cache

```
show msdp [sa-cache | rejected-sa-cache] {group-address grp-addr} {source-address
src-addr} {as-number as-num} {originator-rp originator-rp-addr} {local} {peer
remoteaddr} {vr vrname}
```

Description

This command displays the SA cache database. The following quadruplet per SA cache entry displays: {Group, Source, originating RP, and peer}. In addition, information about the following displays: the cache uptime, aging, whether sources are local or remote, etc.

Syntax Description

<i>grp-addr</i>	Displays the SA cache within the specified group address range.
<i>src-addr</i>	Displays the SA cache within the specified source address range.
<i>as-num</i>	Displays all SA cache that originated from the specified Autonomous System (AS) number.
<i>originator-rp-addr</i>	Displays all SA cache entries that were originated by the specified rendezvous point.
local	Displays locally originated SA cache entries only.
<i>remoteaddr</i>	Displays the SA cache entries received from the MSDP peer with the specified IP address.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

Use this command to view and troubleshoot the SA cache database. There are various filtering criteria you can use to display just a subset of the SA cache database. The following are some of the criteria, which you can use together or separately, to display information about the SA cache:

- Filtering on the group address range.
- Filtering on the source address range.



- Filtering on the originator rendezvous point address.
- Filtering of the advertising MSDP peer.
- Locally originated SA cache.
- Rejected SA cache.

Example

The following command displays the SA cache database:

```
show msdp sa-cache
Group Address      Source Address  Originator      Peer Address    Age/Ageout In
-----
--
235.100.200.1     10.20.30.1     60.0.0.5       10.0.0.1       00:44:24/05:10
235.100.200.2     10.20.30.2     60.0.0.5       192.0.0.16     00:44:24/05:16
235.100.200.3     10.20.30.3     60.0.0.5       10.0.0.1       00:44:24/05:10
235.100.200.4     10.20.30.4     60.0.0.5       10.0.0.1       00:44:24/05:10
235.100.200.5     10.20.30.5     60.0.0.5       55.0.0.5       00:44:24/05:01
235.100.200.6     10.20.30.6     60.0.0.5       178.54.67.23   00:44:24/05:17
235.100.200.7     10.20.30.7     60.0.0.5       112.234.213.12 00:44:24/05:43
235.100.200.8     10.20.30.8     60.0.0.5       10.0.0.1       00:44:24/05:10
235.100.200.9     10.20.30.9     60.0.0.5       10.0.0.1       00:44:24/05:10
235.100.200.10    10.20.30.10    60.0.0.5       0.0.0.0        00:44:24/00:00
235.100.200.11    10.20.30.11    60.0.0.5       0.0.0.0        00:44:24/00:00
235.100.200.12    10.20.30.12    60.0.0.5       0.0.0.0        00:44:24/00:00
235.100.200.13    10.20.30.13    60.0.0.5       0.0.0.0        00:44:24/00:00
235.100.200.14    10.20.30.14    60.0.0.5       0.0.0.0        00:44:24/00:00
235.100.200.15    10.20.30.15    60.0.0.5       0.0.0.0        00:44:24/00:00
235.100.200.16    10.20.30.16    60.0.0.5       0.0.0.0        00:44:24/00:00
235.100.200.17    10.20.30.17    60.0.0.5       0.0.0.0        00:44:24/00:00
235.100.200.18    10.20.30.18    60.0.0.5       0.0.0.0        00:44:24/00:00
235.100.200.19    10.20.30.19    60.0.0.5       0.0.0.0        00:44:25/00:00
Number of accepted SAs      : 255
Number of rejected SAs     : 0
Flags: (a) Accepted, (f) Filtered by policy, (r) RPF check failed
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure msdp sa-cache-server

```
unconfigure msdp sa-cache-server {vr vrname}
```



Description

Removes the MSDP SA cache server.

Syntax Description

<i>vrname</i>	Specifies the name of the virtual router on which the MSDP cache server is configured. If a virtual router name is not specified, it is extracted from the current CLI context.
----------------------	---

Default

By default, the router does not send SA request messages to its MSDP peers when a new member joins a group and wants to receive multicast traffic. The new member simply waits to receive SA messages, which eventually arrive.

Usage Guidelines

Use this command to remove the MSDP SA cache server you specified with the `configure msdp sa-cache-server` command.

Example

The following command removes the MSDP SA cache server:

```
unconfigure msdp sa-cache-server
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

unconfigure pim border

```
unconfigure pim vlan_name border
```

Description

Unconfigures a PIM VLAN that has been configured as a border VLAN, which is used to demarcate a PIM domain when using MSDP.



Syntax Description

<code>vlan_name</code> Specifies a VLAN name.

Default

By default, no PIM VLANs are configured as border VLANs.

Usage Guidelines

A PIM-SM domain is created by limiting the reach of PIM BSR advertisements. When a border VLAN is configured, PIM BSR advertisements are not forwarded out of the PIM VLAN. Use the `unconfigure pim border` command to remove the border functionality of the specified PIM VLAN.

Example

The following command unconfigures a PIM border on a VLAN called “vlan_border”:

```
unconfigure pim vlan_border border
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.



47 Configuration and Image Commands

```
clear license-info
configure firmware
download image
enable license
enable license file
install bootrom
install firmware
install image
load script
run update
save configuration
save configuration as-script
show configuration
show licenses
show memorycard
show script output autoexec
show script output default
synchronize
unconfigure switch
uninstall image
upload configuration
use configuration
use image
```

This appendix describes commands for:

- Downloading and using a new switch software image
- Saving, uploading, and downloading switch configuration information
- Downloading and installing a new BootROM image and switch rebooting

The switch software image contains the executable code that runs on the switch. An image comes preinstalled from the factory. The image can be upgraded by downloading a new version from a Trivial File Transfer Protocol (TFTP) server on the network. If you have a switch with a compact flash card or a USB 2.0 storage device, you can also download a new version from the storage device.

A switch can store up to two images; a primary and a secondary image. You can download a new image into either one of these, and you can select which image will load on the next switch reboot.

The configuration is the customized set of parameters that you have selected to run on the switch. As you make configuration changes, the new settings are stored in run-time memory. To retain the settings, and have them load when you reboot the switch, you must save the configuration to nonvolatile storage.

The switch can store multiple user-defined configuration files, each with its own file name. By default, the switch has two pre-named configurations: a primary and a secondary configuration. You can select to which configuration you want the changes saved, or you can save the changes to a new configuration file. You can also select which configuration will be used on the next switch reboot.

The BootROM initializes certain important switch variables during the switch boot process. In specific situations, you can upgrade the BootROM on the Summit family switches and SummitStack by download from a TFTP server on the network. On the BlackDiamond X8 and BlackDiamond 8800 series switches, you can upgrade the firmware, including the BootROM, when you upgrade the software image.

clear license-info

```
clear license-info {software}
```

Description

This command, which should be used only in conjunction with a representative from Extreme Networks, clears the licensing information from the switch.

Syntax Description

software	Specifies ExtremeXOS base software license.
-----------------	---

Default

N/A.

Usage Guidelines



Note

Use this command only under the guidance of an Extreme Networks representative.

When you issue the command, the following message is displayed:

```
This will clear the license information stored in EEPROM and also delete
the license file (license.xlic). Are you sure you want to continue? (y/N)
```

When you reply “yes”, the license information is removed from the EEPROM and the switch deletes the license.xlic file permanently.



Example

The following command removes licensing information from the switch:

```
clear license-info
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

configure firmware

```
configure firmware [auto-install | install-on-demand]
```

Description

Configures the way a BlackDiamond 8800 series switch performs a system firmware upgrade.

Syntax Description

auto-install	Specifies ExtremeXOS to automatically upgrade the firmware if the software detects a newer firmware image is available. The switch does not prompt you to confirm the firmware upgrade.
install-on-demand	Specifies the switch to prompt you to upgrade the firmware when ExtremeXOS determines that a newer firmware image is available. This is the default behavior.

Default

The default is install-on-demand.

Usage Guidelines

Use the `configure firmware` [**auto-install** | **install-on-demand**] and `install firmware {force}` commands to upgrade the BootROM images on the MSM and I/O modules and the firmware on the PSU controllers installed in BlackDiamond X8 and BlackDiamond 8800 series switches.

Firmware images are bundled with ExtremeXOS software images. ExtremeXOS automatically compares the existing firmware image flashed into the hardware with the firmware image bundled with the ExtremeXOS image when you:

- Download a new version of ExtremeXOS to the alternate (inactive) partition.
- Install a new module into an active chassis.



After a firmware image upgrade, messages are sent to the log.

In addition, if you insert an I/O module that has older firmware present into a chassis that is configured for auto-install then the following behaviors occur:

On the BlackDiamond 8800, newer firmware is installed automatically upon insertion and the following log message is generated:

```
<Info:HAL.Card.Info> MSM-A: Slot 2 bootrom updated to 1.0.1.3
```

There is also a log message that states that the newer firmware will not be used until after the next reboot of the I/O module.

On the BlackDiamond X8, newer firmware will not be installed automatically. A log message is generated informing the user that newer firmware is available and the command to use to install it. This behavior applies to MM, IO, FM, and FanTray only.

If you select the auto-install parameter, you are not prompted to confirm the firmware upgrade. Whenever ExtremeXOS determines a newer firmware image is available, the firmware is automatically upgraded.

If you use the default configuration install-on-demand, you have the opportunity to cancel the firmware upgrade. If you install a new software image, and a new firmware image is available, the switch prompts you to upgrade the firmware. Enter y to upgrade the firmware image. Enter n to cancel the firmware upgrade for the specified hardware and continue scanning for other hardware that needs to be upgraded.

The following command downloads the switch software image. The secondary partition must be the alternate partition.

```
download image bd8800-11.3.0.10.xos secondary
```

If you download a new image and new firmware images are available, you see messages similar to the following:

```
Do you want to install image after downloading? (y - yes, n - no, <cr> - cancel) Yes Downloading to
MSM-A Saving configuration on secondary MSM ..... done! Installing to secondary partition! Installing
to MSM-A... Image installed successfully Installing version 1.0.0.16 of the MSM bootrom(s). Do you want
to continue? (y/n) Yes Installing version 1.0.0.24 of the IO module bootrom(s). Do you want to
continue? (y/n) Yes Installing version 2.4 of the PSU control module firmware. Do you want to
continue? (y/n) Yes Installing bootrom... MSM bootrom(s) installed successfully Installing bootrom... IO
module bootrom(s) installed successfully Installing firmware... PSU controller firmware installed
successfully ...
```

Displaying BootROM and Firmware Versions

To display the BootROM (firmware) version for all modules and PSU controllers installed in the switch, use the `show version` command.



Recovering From a Corrupted BootROM

If your default BootROM image becomes corrupted, you can force the MSM to boot from an alternate BootROM image by inserting a pen into the Alternate (A) and Reset (R) holes on the BlackDiamond 8800 MSM and applying pressure. For more information, please refer to the hardware documentation.

Example

The following command automatically upgrades the firmware when a newer firmware image is present without prompting you to confirm the upgrade:

```
configure firmware auto-install
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available only on the BlackDiamond X8 and BlackDiamond 8800 series switches.

download image

Using TFTP:

```
download image [[hostname | ipaddress] filename {{vr} vrname} | memorycard  
filename] {partition} {msm slotid}
```

To download an image to a stack:

```
download image [[hostname | ipaddress] filename {{vr} vrname} | memorycard  
filename] {partition} {slot slot number}
```

Description

Downloads a new version of the ExtremeXOS software image.

The image file can be downloaded using TFTP which is not a secure method or SFTP and SCP2 which are secure methods. The procedure using TFTP begins above and using SFTP/SCP2.

Note



Beginning with ExtremeXOS 12.1, an ExtremeXOS core image must be downloaded and installed on the alternate (non-active) partition. If a user tries to download to an active partition, the error message **Error: Image can only be installed to the non-active partition.** is displayed.



Syntax Description

<i>hostname</i>	Specifies the hostname of the TFTP server from which the image should be obtained.
<i>ipaddress</i>	Specifies the IP address of TFTP server from which the image should be obtained.
memorycard	Specifies that the image should be obtained from a removable storage device, which can be a compact flash card or a USB 2.0 storage device. Note: This parameter is available for compact flash cards on BlackDiamond 8800 series switches and for USB 2.0 storage devices on BlackDiamond X8, Summit X460, X480, X670, X670V, and X770 switches.
<i>filename</i>	Specifies the filename of the new image.
<i>vrname</i>	Specifies the name of the virtual router. Note: User-created VRs are supported only on the platforms listed for this feature in the Feature License Requirements in the <i>ExtremeXOS Concepts Guide</i> .
<i>partition</i>	Specifies which partition the image should be saved to: primary or secondary. Select primary to save the image to the primary partition and secondary to save the image to the secondary partition.
<i>slotid</i>	Specifies the MSM/MM where the software image should be downloaded. A specifies the MSM/MM installed in slot A.B specifies the MSM/MM installed in slot B. Note: This parameter is available only on modular switches.
<i>slot number</i>	Specifies the slot where the software image should be downloaded. The value may be from 1 to 8. Note: This parameter is available only on stackable switches in a stack.

Default

Stores the downloaded image in the alternate (inactive) partition.

SFTP and SCP2 provide secure methods of downloading the ExtremeXOS software image files, *.xos or *.xmod. You can use one of three procedures:

- From the switch, running the command SCP2. connect to and “get” from a remote server. This is similar to the download image command.
- From outside the switch, connect to the switch which is acting as the server and “put” from the remote server. There is no TFTP equivalent for this method.
 - Using SFTP, or
 - Using SCP2

Example of these procedures are included [the example below](#).



Usage Guidelines

Prior to downloading an image on the switch, you must download the image you received from Extreme Networks to a TFTP server on your network. If your switch has a removable storage device, you can also download the image to that device.



Note

Unlike ExtremeWare, the `download image` command in ExtremeXOS causes the switch to use the newly downloaded software image during the next switch reboot. To modify or reset the software image used during a switch reboot, use the `use image` command. Use the `use image` command after downloading and installing the image for it to be effective.

Specify the `ipaddress` or `hostname` parameters to download an image from a TFTP server on the network. Use of the `hostname` parameter requires that DNS be enabled.

Specify `memorycard` to download a an image from a removable storage device. Use a PC with appropriate hardware such as a compact flash reader/writer and follow the manufacturer's instructions to access the compact flash card and place the image onto the card. For more information about installing a removable storage device, see the hardware documentation.

Core Software Images

The switch can store up to two core images: a primary and a secondary. When downloading a new image, you select which partition (primary or secondary) to install the new image. The ExtremeXOS core image must be downloaded and installed to the alternate partition.

Image Filenames

The software image file can be an `.xos` file, which contains an ExtremeXOS core image, or an `.xmod` file, which contains an ExtremeXOS modular software package. Modular software packages have additional functionality that supplement a core image.

You can identify the appropriate image or module for your platform based on the filename of the image. The following table lists the filename prefixes for each platform:

Table 58: Filename Prefixes

Platform	Filename Prefixes
BlackDiamond X8	bdX-
BlackDiamond 8810	bd8800-
BlackDiamond 8806	bd8800-
Summit family	summitX-

For example, if you have a BlackDiamond 8806 switch, download image filenames with the prefix `bd8800-`.



Displaying the Software Image Versions

To display the software image version running on the switch, use the `show version` or `show switch` commands.

Host Name and Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-) Permitted only for host names
- Underscore (_) Permitted only for host names
- Colon (:)

When naming or configuring an IP address for your network server, remember the requirements listed above.

Local and Remote Filename Character Restrictions

This section provides information about the characters supported by the switch for local and remote filenames.

When specifying a local or remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)
- Slash (/) Permitted only for remote files

When naming a local or remote file, remember the requirements listed above.

Messages Displayed by the Switch

When you download a new image, you will see the following message:

Do you want to install image after downloading? (y - yes, n - no, <cr> - cancel)

Do one of the following:

- Enter y if you want to install the image after download.
- Enter n if you want to install the image at a later time.
- Press [Enter] if you want to cancel the download.



Core Dump Messages

If you configure the switch to write core dump (debug) files to the internal memory card and attempt to download a new software image, you might have insufficient space to complete the image download. If this occurs, you must decide whether to continue the software download or move or delete the core dump files from the internal memory. For example, if you have a switch with removable storage device and available space, transfer the files to the device. On switches without removable storage devices, transfer the files from the internal memory card to a TFTP server. This frees up space on the internal memory card while keeping the core dump files.

The switch displays a message similar to the following and prompts you to take action:

Core dumps are present in internal-memory and must be removed before this download can continue. (Please refer to documentation for the “configure debug core-dumps” command for additional information) Do you want to continue with download and remove existing core dumps? (y/n)

Enter y to remove the core dump files and download the new software image. Enter n to cancel this action and transfer the files before downloading the image.

For information about configuring and sending core dump information, see the `configure debug core-dumps` and `save debug tracefiles memorycard` commands.

SummitStack Only

You can issue this command only from the Master node.

If a slot is not specified, the image is downloaded to every node in the Active Topology. If a slot is specified, the image is downloaded to that slot only.

If all nodes to be downloaded are not running the same partition, the command is not executed and following message is displayed:

Error: all nodes do not have the same image partition selected.

If all nodes to be downloaded have the same partition selected but the EXOS is currently running from the selected partition, the command is not executed and the following message is displayed:

Error: the image partition selected must not be the active partition.

Downloading a New Image

The following assumes you have already downloaded the image to a network TFTP server or removable storage device. The information in this section provides more detailed information for downloading a new image to your switch.



Note

Always refer to the most recent version of the release notes for the most current download instructions.

Step 1--Verifying the Virtual Router

If you loaded the image onto a removable storage device, proceed to step 2.



If you loaded the image onto a TFTP server, use one of the following ping commands to confirm which virtual router reaches your TFTP server:

```
ping vr vr-Mgmt host
```

```
ping vr vr-Default host
```

At least one of these commands must successfully reach your TFTP server for you to download the image. After verifying the virtual router that reaches your TFTP server, specify that virtual router when you download the image.

Step 2--Viewing the Partition

To view your selected and booted partition, use the following command: `show switch`

Output from this command includes the selected and booted images and if they are in the primary or the secondary partition.

Step 3--Selecting the Partition

The image must be downloaded and installed to the alternate (inactive) partition. To specify the partition when downloading and installing the image, use one of the following commands:

- (SummitStack) `download image [[hostname | ipaddress] filename {{vr} vrname} | memorycard filename] {partition} {msm slotid}`
- (Modular switches) `download image [[hostname | ipaddress] filename {{vr}vrname} | memorycard filename] {partition} {slot slot number}`

Step 4--Downloading and Installing the Image

To download the image, use the appropriate, previously described, `download image` command.

Downloading an ExtremeXOS core image

An ExtremeXOS core image uses the file extension `.xos`.

- Before the download begins, the switch asks if you want to install the image immediately after the download is finished.

Enter `y` to install the image after download. Enter `n` to install the image at a later time.

When you install the image to the alternate (inactive) partition; you do not need to reboot the switch until you are ready to use the image.

If you install the image at a later time, the image is still downloaded and saved to the switch, but you must use the following command to install the software and reboot the switch:

```
install image fname {partition} {msm slotid} {reboot}
```

Where `fname` specifies the filename of the new, downloaded image.



Downloading an ExtremeXOS module image.

An ExtremeXOS module image has functionality that supplements a core image. You download and install a module onto an already installed core image. The version number of the core image and the module must match. For example, the module `bd10K-11.0.0.25-ssh.xmod` can only be installed onto the core image `bd10K-11.0.0.25.xos`.

To install a module to the inactive partition, use the `download image` command to download the module to the inactive partition, and use the `install firmware` command to install it, if you did not choose to install when the image was downloaded. Remember, the core image on the inactive partition must be of the same version as the module. When you make the inactive partition active, by issuing the `use image` command and rebooting the switch, the module is also activated at boot time.

To install a module to the active partition (except on SummitStack), use the `download image` command to download the module to the active partition, and use the `install firmware` command to install it, if you did not choose to install when the image was downloaded. Remember, the core image on the active partition must be of the same version as the module. If you reboot the switch, the module will also be activated, but you can activate the module without rebooting the switch by issuing the `run update` command. After issuing that command, all the functionality, and command line interface (CLI) commands, of the module will be available.

Performing a Hitless Upgrade--Modular Switches Only

Hitless upgrade is a mechanism that allows you to upgrade the ExtremeXOS software running on the switch without taking the switch out of service. Some additional benefits of using hitless upgrade include:

- Minimizing network downtime.
- Reducing the amount of traffic lost.

Although any method of upgrading software can have an impact on network operation, including interrupting Layer 2 network operation, performing a hitless upgrade can decrease that impact.

You must have two MSMs installed in your switch to perform a hitless upgrade. With two MSMs installed in the switch, one assumes the role of primary and the other assumes the role of backup. The primary MSM provides all of the switch management functions including bringing up and programming the I/O modules, running the bridging and routing protocols, and configuring the switch. The primary MSM also synchronizes its configurations with the backup MSM which allows the backup to take over the management functions of the primary.



Note

If you download an image to the backup MSM, the image passes through the primary MSM before the image is downloaded to the backup MSM.

Before performing a hitless upgrade, review the following list to confirm that your system supports hitless upgrade:

- BlackDiamond 8800 series switch with a mix of BlackDiamond 8000 c-, e-, and xl-series modules installed—Both MSMs are running ExtremeXOS 11.5 or later. r.

To perform a hitless upgrade, do the following:

View current switch information using the following command: `show switch`



Determine your selected and booted partition, verify which MSM is the primary and which is the backup, and confirm that the MSMs are synchronized.

Output from this command indicates, for each MSM, the selected and booted images and if they are in the primary or the secondary partition. The selected image partition indicates which image will be used at the next reboot. The booted image partition indicates the image used at the last reboot. It is the active partition.

The current state indicates which MSM is the primary (displayed as MASTER), which MSM is the backup (displayed as BACKUP), and if the backup MSM is synchronized with the primary MSM (displayed as In Sync).

Select the partition to download the image to and download and install the new ExtremeXOS software on the backup MSM using the following command:

```
download image [[hostname | ipaddress] filename {{vr} vrname} | memorycard
filename] {partition} {msm slotid}
```



Note

If the backup MSM is installed in slot B, specify msm B. If the backup MSM is installed in slot A, specify msm A.

- Before the download begins, the switch asks if you want to install the image immediately after the download is finished.

Enter `y` to install the image after download. Enter `n` to install the image at a later time.

When you install the image after download to the alternate partition, you need to reboot only the Backup MSM that the newer code was downloaded and installed on. Use the `reboot msm b` command (if "b" is the Backup MSM).

If you install the image at a later time, use the following command to install the software:

```
install image fname {partition} {msm slotid} {reboot}
```

Verify that the backup MSM comes up correctly and that the MSMs are synchronized using the following command: `show switch`

The current state indicates which MSM is the primary (displayed as MASTER), which MSM is the backup (displayed as BACKUP), and if the backup MSM is synchronized with the primary MSM (displayed as In Sync).

Initiate failover from the primary MSM to the backup MSM using the following command: `run msm-failover`

When you failover from the primary MSM to the backup MSM, the backup becomes the new primary, runs the newly downloaded software, and provides all of the switch management functions.

If you have a BlackDiamond 8800 series switch and the new ExtremeXOS image supports hitless upgrade but is not compatible with the current running I/O module image (the I/O version numbers do not match), you cannot perform a hitless upgrade.

The switch displays a warning message similar to the following:



WARNING: Failover will not be hitless due to incompatible images. Traffic will be interrupted. Are you sure you want to failover? (y/n)

You can either continue the upgrade or cancel the action. If you continue the upgrade, the primary MSM downloads the new image to the I/O module and reboots.

Verify that the backup MSM comes up correctly and that the MSMs are synchronized using the following command: `show switch`

The current state indicates which MSM is the primary (displayed as MASTER), which MSM is the backup (displayed as BACKUP), and if the backup MSM is synchronized with the primary MSM (displayed as In Sync).

Select the partition to download the image to and download and install the new ExtremeXOS software on the new backup MSM (this was the original primary MSM) using the following command:

```
download image [hostname | ipaddress] filename {vr vrname} msm slotid
```



Note

If the new backup MSM is installed in slot A, specify `msm A`. If the new backup MSM/MM is installed in slot B, specify `msm B`.

Before the download begins, the switch asks if you want to install the image immediately after the download is finished.

- When you download and install the software image on the alternate partition, you need to reboot only the Backup MSM that the newer code was downloaded and installed on. This can be done using the `reboot msm a` command (if "a" is the New Backup)
- If you install the image at a later time, use the following command to install the software:

```
install image fname {partition} {msm slotid} {reboot}
```

Verify that the new backup MSM comes up correctly and that the MSMs are synchronized using the following command: `show switch`

The current state indicates which MSM is the primary (displayed as MASTER), which MSM is the backup (displayed as BACKUP), and if the backup MSM is synchronized with the primary MSM (displayed as In Sync).

Optionally, initiate failover from the new primary MSM to the new backup MSM using the following command: `run msm-failover`

When you failover from the new primary MSM to the new backup MSM, this optional step restores the switch to the original primary and backup MSM.

Optionally, confirm that the failover is successful by checking the current state of the MSMs using the following command: `show switch`

You can also perform a hitless upgrade on ExtremeXOS modular software packages (.xmod files). To perform a hitless upgrade of a software package, you must install the core software image first, and the version number of the modular software package must match the version number of the core image that it will be running with. For more information about hitless upgrade, see the *ExtremeXOS Concepts Guide*.



Hitless Upgrade Caveats for the BlackDiamond 8800 Series Switches Only

The following is a summary of hitless upgrade caveats for only the BlackDiamond 8800 series switches:

- If you are running ExtremeXOS 11.4 or earlier, **do not** attempt to perform a hitless upgrade to ExtremeXOS 11.5 or later. If attempted, the backup MSM enters the non-operational state.

To recover from the non-operational state, do the following:

- 1 From the primary MSM, use the **synchronize** command to return the MSMs to the same version of software. To confirm the MSMs are synchronized, use the **show switch** command.

The current state indicates which MSM is the primary (displayed as MASTER), which MSM is the backup (displayed as BACKUP), and if the backup MSM is synchronized with the primary MSM (displayed as In Sync).

- 2 After you recover from the non-operational state and confirm the MSMs are synchronized, perform a normal code upgrade to install and upgrade the image on the switch. For more information, see the sections starting with [Step 1—Verifying the Virtual Router through Downloading an ExtremeXOS core image](#).

Note



ExtremeXOS 11.5 introduced support for the BlackDiamond 8800 e-series modules. If your switch is running ExtremeXOS 11.4 or earlier, you must upgrade to ExtremeXOS 11.5 to operate the modules. Hitless upgrade is not supported between major releases. Do not attempt to perform a hitless upgrade. To upgrade the switch from ExtremeXOS 11.4 or earlier to ExtremeXOS 11.5 or later, reboot the switch after downloading and installing the new image to both installed MSMs.

Example—Using TFTP on Modular Switches

The following command downloads the switch software image from the TFTP server at 10.10.15.4, from the file named bd12K-12.4.1.1.xos without specifying the desired partition:

```
download image 10.10.15.4 bd12K-12.4.1.1.xos
Note: The inactive partition (secondary) will be used for installation.
Do you want to install image after downloading? (y - yes, n - no, <cr> -
cancel) Yes
Downloading to MSM-
A.....
Installing to secondary partition!
Installing to MSM-
A.....
.....
.
.....
.
.....
```



The following command downloads the switch software image from the TFTP server at 10.10.15.4, from the file named bd10K-11.0.0.25.xos specifying the desired partition. The secondary partition is the alternate partition in this example.

```
download image 10.10.15.4 bd10K-11.0.0.25.xos secondary
```

On a modular switch, when you download an image into the alternate partition, you see output similar to the following:

```
Do you want to install image after downloading? (y - yes, n - no, <cr> -
cancel) Yes
```

```
Downloading to MSM-
```

```
A.....
```

```
Downloading to MSM-B.....
```

```
Installing to secondary partition!
```

```
Installing to MSM-
```

```
B.....
```

```
.....
```

```
.
```

```
.....
```

```
.
```

```
.....
```

```
.
```

```
Installing to MSM-
```

```
A.....
```

```
.....
```

```
.
```

```
.....
```

```
.
```

```
.....
```

```
.
```

If you answer yes to installing the image, the switch reboots upon completion of the installation.

Example—Using TFTP on Summit Switches

The following command downloads the switch software image from the TFTP server at 10.10.15.4, from the file named summitX440-11.2.0.15.xos:

```
download image 10.10.15.4 summitX440-12.4.1.1.xos
```

```
Note: The inactive partition (secondary) will be used for installation.
```

When you download an image into the alternate partition, you see output similar to the following:

```
Do you want to install image after downloading? (y - yes, n - no, <cr> -
cancel) Yes
```

```
Downloading to
```

```
Switch.....
```




```

.....
.....
.....
.....
.....

```

Using SFTP and SCP2

The following commands show the three procedures:

- From the Switch using SCP2:

```

(debug) SummitX440-24x.5 # scp2 root@10.203.133.13:/tmp/summitX-12.0.0.4-
ssh.xmod summitX-12.0.0.4-ssh.xmod Download to /scratch/summitX-12.0.0.4-
ssh.xmod on switch Connecting to 10.203.133.13... root@10.203.133.13's
password: Fetching /tmp/summitX-12.0.0.4-ssh.xmod to /scratch/
summitX-12.0.0.4-ssh.xmod /tmp/summitX-12.0.0.4-ssh.xmod 100% 997KB 996.9KB/s
00:00 Please wait validating image summitX-12.0.0.4-ssh.xmod ,this will take
approximately 30 seconds ... * (debug) SummitX440-24x.6 #

```

- From Outside the Switch using SFTP:

```

sftp admin@10.203.3.231 Connecting to 10.203.3.231... admin@10.203.3.231's
password: sftp> put summitX-12.0.0.4-ssh.xmod Examplesftp.xmod Uploading
summitX-12.0.0.4-ssh.xmod to /config/Examplesftp.xmod sftp>

```

- From Outside the Switch using SCP2

- scp2 summitX-12.0.0.4-ssh.xmod admin@10.203.3.231:Examplescp.xmod
- Keyboard-interactive:
- Keyboard-interactive authentication
- Enter password for admin:
- summitX-12.0.0.4-ssh.xmod | 997kB | 997kB/s | TOC: 00:00:01 | 100%
- [-/cvs/branch-fixes_v1200b4.c4/exos/cougar/release]\$

Following download, the output from a show log command, shown below, is similar to the following for the last two procedures—the two from outside the switch.

```
show log
```

```

04/03/2007 14:32:29.31 <Info:AAA.LogSsh> Got Image file Example.xmod 04/03/2007 14:32:29.31
<Info:AAA.LogSsh> Validating Image file, this could take approximately 30 seconds.. Example.xmod
04/03/2007 14:32:30.89 <Info:AAA.LogSsh> Image file Example.xmod successfully validated can now
install image

```

History

This command was first available in ExtremeXOS 10.1.

The memorycard option was added in ExtremeXOS 11.0.

The msm parameter was added in ExtremeXOS 11.1.



Support for USB 2.0 storage devices was added in ExtremeXOS 12.5.3.

Platform Availability

This command is available on all platforms.

enable license

```
enable license {software} [key | capacity_key ]
```

Description

Enables software license or feature pack that allows you to use advanced features.

Syntax Description

<i>key</i>	Specifies your hexadecimal license key in format xxxx-xxxx-xxxx-xxxx-xxxx.
<i>capacity-key</i>	Specifies your Hexadecimal license key in format xxxx-xxxx-xxxx-xxxx-xxxx-xx.

Default

N/A.

Usage Guidelines

The software license levels that apply to ExtremeXOS software are described in the [Feature License Requirements](#) in the *ExtremeXOS Concepts Guide*.

To obtain a software license, specify the key in the format xxxx-xxxx-xxxx-xxxx-xxxx.

You obtain the software license key (or feature pack key) either by ordering it from the factory or by obtaining a license voucher from your Extreme Networks supplier. You can obtain a regular software license or a trial software license, which allows you use of the license for either 30, 60 or 90 days; you cannot downgrade software licenses.

The voucher contains all the necessary information on the software license, whether regular or trial, and number of days for trial software license.

After you enable the software license or feature pack by entering the software key, the system returns a message that you either successfully or unsuccessfully set the license.

Once you enable the software license (or if you do not use the correct key, attempt to downgrade the license, or already installed the software license) you see one of the following messages:

```
Enabled license successfully. Error: Unable to set license using supplied key. Error: Unable to set license - downgrade of licenses is not supported. Error: Unable to set license - license is already enabled. Error: Unable to set license - trial license already enabled.
```



If you enable a trial license, the system generates a daily message showing the number of days until expiry.

Once installed (or enabled), the software license goes with the switch chassis itself (not with the MSM/MM module on modular switches). The software license information is stored in EEPROM in the modular switches and NVRAM on the stand-alone switches and SummitStack; the information persists through reboots, software upgrades, power outages, and reconfigurations.

If you attempt to execute a command and you do not either have the required software license or have reached the limits defined by the current software license level, the system returns one of the following messages:

```
Error: This command cannot be executed at the current license level.
Error: You have reached the maximum limit for this feature at this
license level.
```

If you attempt to execute a command and you do not have the required feature pack, the system also returns a message.

To protect against attacks to install maliciously created license keys, the system has an exponential delay of each failed attempt to install a license.

To view the type of software license you are currently running on the switch, use the `show licenses` command. The license key number is not displayed, but the type of software license is displayed in the `show licenses` output. This command can be run on any node in a SummitStack, regardless of its node role (Master, Standby, or Backup).

Capacity license can be for either a switch or a blade, for example, TCAM license is per blade. The key is a hexadecimal value of 11 bytes, which contains all the necessary information regarding the type and nature of the license. The key length can be used to differentiate capacity license from existing software license whose length is 10 bytes.

For TCAM license change, the following warning message is displayed if the license is successfully enabled: `Warning: A reboot switch or disable and enable slot n is required before the new license takes effect.`

Example

The following command enables a software license on the switch:

```
enable license 2d5e-0e84-e87d-c3fe-bfff
```

```
enable license 2d5e-0e84-e87d-c3fe-bfff-ae
```

```
Half External Tables Usage license successfully enabled on the
BDXC-40G12X-XL module in Slot-3.
```

```
Warning: A reboot switch or disable and enable slot 3 is required
before the new license takes effect.
```

Messages to user for different scenarios:



- No module in the switch match the part # and serial # in the license key.

```
"Error: No module in the switch matches the supplied key."
```
- Key format is incorrect.

```
"Error: Incorrect key format."
```
- License key is re-applied.

```
"Full External Tables Usage license is already enabled on the  
BDXC-40G12X-XL module in Slot-3".
```
- EEPROM Read/Write failure.

```
"Error: Failed to enable External Tables Usage license on the  
BDXC-40G12X-XL module in Slot-3. This may be due to damage or  
corruption of the EEPROM on the module and should be reported."
```

History

This command was first available in ExtremeXOS 11.1.

The software parameter was added in ExtremeXOS 11.6.

The *capacity-key* variable was added in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

enable license file

```
enable license file filename
```

Description

Enables the text file that applies software licenses and feature packs licenses to more than one switch at a time.

Syntax Description

filename	Specifies the filename that you download onto the switch using TFTP; the file extension is .xlic.
----------	---

Default

N/A.



Usage Guidelines

You download the license file to the switch using TFTP or SCP. The file name extension for this file is `.xlic`; for example, you may see a file named `systemlic.xlic`.

Using this file, you enable the software and feature pack licenses for more than one switch simultaneously. The file can contain licenses for some or all of the Extreme Networks switches that the customer owns. During upload, only those license keys destined for the specific switch are used to attempt enabling the licenses. The license file is a text file that has the switch serial number, software license type, and license key; it is removed from the switch after the licenses are enabled.

After you enable the license file, the system returns one or more of the following messages:

```
Enabled license successfully. Error: Unable to set license <license_name>
using supplied key. Error: Unable to set license <license_name> -
downgrade of licenses is not supported. Error: Unable to set license
<license_name> - license is already enabled. Error: Unable to set license
<license_name> - trial license already enabled.
```

To protect against attacks to install maliciously created license keys, the system has an exponential delay of each failed attempt to install a license.

Example

The following command enables a license file on the specified Extreme Networks switches:

```
enable license file santaclara.xlic
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.

install bootrom

```
install bootrom fname {reboot} {msm slotid}
```

On a SummitStack use:

```
install bootrom fname {reboot} {slot slotid}
```

Description

Installs a new version of the ExtremeXOS BootROM image.



Syntax Description

<i>fname</i>	Specifies the BootROM image file.
reboot	Reboots the switch after the image is installed.
<i>slotid</i>	Specifies the MSM/MM where the BootROM image should be installed. A specifies the MSM/MM installed in slot A. B specifies the MSM/MM installed in slot B. Note: This parameter is available only on the SummitStack. On a SummitStack, the slotid specifies the node on which the BootROM image should be installed.

Default

N/A.

Usage Guidelines

When you download a BootROM image, the system asks if you want to install the image immediately after the download is finished. If you choose to install the image at a later time, use this command to install the software on the switch.

The BootROM image file is an .xbr file, and this file contains the executable code.

Displaying the BootROM Versions

To display the BootROM version for the switch and all of the modules and PSU controllers installed in a modular switch, use the `show version` command.

Local Filename Character Restrictions

This section provides information about the characters supported by the switch for local and remote filenames.

When specifying a local or remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)

When naming a local file, remember the requirements listed above.

SummitStack Only

You can issue this command only from the Master node.



Example

The following command installs the bootrom image file bd10K-1.0.1.5-bootrom.xbr:

```
install bootrom bd10K-1.0.1.5-bootrom.xbr
```

On the Summit series switch, you see output similar to the following:

```
Installing bootrom...
Writing bootrom
.....
.....
.....
Verifying Flash contents...
.....
.....
.....
.....
.....
bootrom written.
Bootrom installed successfully
```

History

This command was first available in ExtremeXOS 11.0.

The msm parameter was added in ExtremeXOS 11.1.

From ExtremeXOS 12.0, this command is supported on a stack. The slot parameter is added. The slot parameter is applicable only when the switch is in a stack.

Platform Availability

This command is available only on Summit family switches and SummitStack.

install firmware

```
install firmware {force}
```

Description

Installs the firmware bundled with the ExtremeXOS image on the BlackDiamond X8 and BlackDiamond 8800 series switches.

Syntax Description

force	Specifies that a new image is installed without a version check.
--------------	--



Default

N/A.

Usage Guidelines

On BlackDiamond X8 switches, use the `install firmware` command to upgrade the BootROM images on the MM, I/O, and Fabric modules and the firmware on the fan bars installed on the switch. On BlackDiamond 8800 series switches, use the `install firmware` command to upgrade the BootROM images on the MSM and I/O modules and the firmware on the PSU controllers installed in the switch.

Firmware images are bundled with ExtremeXOS software images.

On BlackDiamond X8 and BlackDiamond 8800 series switches, the ExtremeXOS software automatically compares the existing firmware image flashed into the hardware with the firmware image bundled with the ExtremeXOS image. You can also use the `install firmware` command to compare the firmware images.

Before using the `install firmware` command, wait until the `show slot` command indicates the MSMs and I/O modules are operational. When the modules are operational, use the `install firmware` command.

BlackDiamond X8 Switches.

The switch scans MM, I/O, Fabric modules and the Fanbar controllers for a possible firmware upgrade. If the bundled firmware image is newer than the existing firmware image, the switch prompts you to confirm the upgrade.

- Enter `y` to upgrade the firmware.
- Enter `n` to cancel the firmware upgrade for the specified hardware and continue scanning for other hardware that needs to be upgraded.
- Enter `<cr>` to cancel the upgrade. After a firmware image upgrade, messages are sent to the log.

The Fanbar controller firmware is used immediately after it is installed without rebooting the switch. The new BootROM and firmware overwrite the older versions flashed into the hardware. Use the `reboot` command to reboot the switch and activate the new BootROM and firmware.

During the firmware upgrade, do not cycle down or disrupt the power to the switch. If a power interruption occurs, the firmware may be corrupted and need to be recovered. ExtremeXOS automatically recovers corrupted firmware; however, the time it takes for the switch to boot-up may increase.

The switch displays status messages after you use the `install firmware` command. The output varies depending upon your platform and the software version running on your system.

BlackDiamond 8800 Series Switches

The switch scans the I/O and MSM modules and the PSU controllers for a possible firmware upgrade. If the bundled firmware image is newer than the existing firmware image, the switch prompts you to confirm the upgrade.



- Enter `y` to upgrade the firmware.
- Enter `n` to cancel the firmware upgrade for the specified hardware and continue scanning for other hardware that needs to be upgraded.
- Enter `cr` to cancel the upgrade. After a firmware image upgrade, messages are sent to the log.

The PSU controller firmware is used immediately after it is installed without rebooting the switch. The new BootROM and firmware overwrite the older versions flashed into the hardware. Use the `reboot` command to reboot the switch and activate the new BootROM and firmware.

During the firmware upgrade, do not cycle down or disrupt the power to the switch. If a power interruption occurs, the firmware may be corrupted and need to be recovered. ExtremeXOS automatically recovers corrupted firmware; however, the time it takes for the switch to boot-up may increase.

The switch displays status messages after you use the `install firmware` command. The output varies depending upon your platform and the software version running on your system.



Note

If the information in the most current version of the ExtremeXOS Installation and Release Notes differs from the information in this section, follow the release notes.

Sample Output--BlackDiamond X8 Series Switch

The following is sample output from a BlackDiamond X8 switch:

```
Installing version 1.0.0.4 of the BIOS for MMs. Do you want to continue?
(y - yes, n - no, <cr> - cancel) Yes
Installing version 1.0.0.8 of the bootrom for BD-X series
I/O and Fabric modules. Do you want to continue?
(y - yes, n - no, <cr> - cancel) Yes
Installing version 1.0.2.9 of fanbar firmware on all fanbars. Do you want to
continue?
(y - yes, n - no, <cr> - cancel) Yes
Installing MM FPGA image version 0.1.22. Do you want to continue?
(y - yes, n - no, <cr> - cancel) Yes
Installing IO/FM FPGA image version 0.0.36. Do you want to continue?
(y - yes, n - no, <cr> - cancel) Yes
Installing firmware...
Updating fanbar firmware
Fan Bar(5)/SlotId(21) - The secondary partition is active, will now boot to
opposite.
Fan Bar(5)/SlotId(21) - Programming secondary partition with /tmp/
fanbar_secondary.bin...
Fan Bar(5)/SlotId(21) - Validating secondary partition with /tmp/
fanbar_secondary.bin...
Fan Bar(5)/SlotId(21) - Images are identical.
Fan Bar(5)/SlotId(21) - Booting updated secondary partition...
Fan Bar(5)/SlotId(21) - Unconditional success!
Fan Bar(4)/SlotId(20) - The secondary partition is active, will now boot to
opposite.
Fan Bar(4)/SlotId(20) - Programming secondary partition with /tmp/
fanbar_secondary.bin...
Fan Bar(4)/SlotId(20) - Validating secondary partition with /tmp/
```



```
fanbar_secondary.bin...
Fan Bar(4)/SlotId(20) - Images are identical.
Fan Bar(4)/SlotId(20) - Booting updated secondary partition...
Fan Bar(4)/SlotId(20) - Unconditional success!
Fan Bar(1)/SlotId(17) - The secondary partition is active, will now boot to
opposite.
Fan Bar(1)/SlotId(17) - Programming secondary partition with /tmp/
fanbar_secondary.bin...
Fan Bar(1)/SlotId(17) - Validating secondary partition with /tmp/
fanbar_secondary.bin...
Fan Bar(1)/SlotId(17) - Images are identical.
Fan Bar(1)/SlotId(17) - Booting updated secondary partition...
Fan Bar(1)/SlotId(17) - Unconditional success!
Fan Bar(3)/SlotId(19) - The secondary partition is active, will now boot to
opposite.
Fan Bar(3)/SlotId(19) - Programming secondary partition with /tmp/
fanbar_secondary.bin...
Fan Bar(3)/SlotId(19) - Validating secondary partition with /tmp/
fanbar_secondary.bin...
Fan Bar(3)/SlotId(19) - Images are identical.
Fan Bar(3)/SlotId(19) - Booting updated secondary partition...
Fan Bar(3)/SlotId(19) - Unconditional success!
Fan Bar(2)/SlotId(18) - The secondary partition is active, will now boot to
opposite.
Fan Bar(2)/SlotId(18) - Programming secondary partition with /tmp/
fanbar_secondary.bin...
Fan Bar(2)/SlotId(18) - Validating secondary partition with /tmp/
fanbar_secondary.bin...
Fan Bar(2)/SlotId(18) - Images are identical.
Fan Bar(2)/SlotId(18) - Booting updated secondary partition...
Fan Bar(2)/SlotId(18) - Unconditional success!
Fan Bar(1)/SlotId(17) - Successfully programmed.
Fan Bar(2)/SlotId(18) - Successfully programmed.
Fan Bar(3)/SlotId(19) - Successfully programmed.
Fan Bar(4)/SlotId(20) - Successfully programmed.
Fan Bar(5)/SlotId(21) - Successfully programmed.
Completed updating fanbar firmware
Firmware for fanbars have been updated successfully.
Installing firmware...
Updating MM FPGA image
MM-A - Programming.
Reading input file image - 255061 bytes.
Writing file image to flash - 255061 bytes.
New FPGA image has been programmed
MM-B - Programming.
Reading input file image - 255061 bytes.
Writing file image to flash - 255061 bytes.
9301 bytes remaining at 12288.0 bytes per second leaves      0 seconds left.
New FPGA image has been programmed
MM FPGA image has been updated successfully.
Installing firmware...
Updating IO/FM FPGA image
Slot-1 - No Module Present.
Slot-2 - Programming.
Reading input file image - 158453 bytes.
Writing file image to flash - 158453 bytes.
35573 bytes remaining at 12288.0 bytes per second leaves      2 seconds left.
New FPGA image has been programmed
```



```

Slot-3 - No Module Present.
Slot-4 - No Module Present.
Slot-5 - Programming.
Reading input file image - 158453 bytes.
Writing file image to flash - 158453 bytes.
39669 bytes remaining at 12288.0 bytes per second leaves      3 seconds left.
New FPGA image has been programmed
Slot-6 - No Module Present.
Slot-7 - Programming.
Reading input file image - 158453 bytes.
Writing file image to flash - 158453 bytes.
43765 bytes remaining at 11468.8 bytes per second leaves      3 seconds left.
New FPGA image has been programmed
Slot-8 - No Module Present.
FM-1 - Programming.
Reading input file image - 158453 bytes.
Writing file image to flash - 158453 bytes.
31477 bytes remaining at 13107.2 bytes per second leaves      2 seconds left.
New FPGA image has been programmed
FM-2 - Programming.
Reading input file image - 158453 bytes.
Writing file image to flash - 158453 bytes.
43765 bytes remaining at 12288.0 bytes per second leaves      3 seconds left.
New FPGA image has been programmed
FM-3 - Programming.
Reading input file image - 158453 bytes.
Writing file image to flash - 158453 bytes.
39669 bytes remaining at 11468.8 bytes per second leaves      3 seconds left.
New FPGA image has been programmed
FM-4 - No Module Present.
IO/FM FPGA image has been updated successfully.
Installing Bootrom for MM(s) may take several minutes, please do not reboot
system during this time
Installing bootrom...
Bootrom for MMs installed successfully and will be activated upon next MM
reboot
Installing bootrom...
Bootrom for BD-X I/O and Fabric modules installed successfully and will be
activated upon next module reboot
This image will be used only after rebooting the switch!

```

Sample Output--BlackDiamond 8800 Series Switch

The following is sample output from a BlackDiamond 8800 series switch:

```

Installing version 1.0.4.2 of the bootrom for MSMs. Do you want to continue?
(y - yes, n - no, <cr> - cancel) Yes
Do you want to save configuration changes to primary.cfg? (y or n) Yes
Saving configuration on primary MSM ..... done!
Installing version 1.0.4.0 of the bootrom for I/O modules.
Do you want to continue?
(y - yes, n - no, <cr> - cancel) Yes
Installing version 1.0.1.0 of the bootrom for newer (e.g. 8900-series)
I/O modules. Do you want to continue?
(y - yes, n - no, <cr> - cancel) Yes
Installing version 2.13 of the firmware for PSU control modules.

```



```

Do you want to continue?
(y - yes, n - no, <cr> - cancel) Yes
Installing bootrom...
MSM bootrom(s) installed successfully
MSM bootrom(s) will be activated upon next MSM reboot
Installing bootrom...
IO module bootrom(s) installed successfully
IO module bootrom(s) will be activated upon next IO module reboot
Installing firmware...
PSU controller firmware installed successfully

```

Additional Behavior--BlackDiamond X8 and BlackDiamond 8800 Series Switches Only

During a firmware upgrade, the switch prompts you to save your configuration changes to the current, active configuration. Enter y to save your configuration changes to the current, active configuration. Enter n if you do not want to save your changes.

In earlier versions of ExtremeXOS, you are required to immediately reboot the system after a firmware upgrade. In ExtremeXOS 11.3.3 and later, the system displays a message informing you that the new firmware image will be activated the next time you reboot the system.

Use the `configure firmware [auto-install | install-on-demand]` command to configure how the switch performs a system firmware upgrade. If you select the auto-install parameter, you are not prompted to confirm the firmware upgrade. If you use the default configuration install-on-demand, you can cancel the firmware upgrade.

Power over Ethernet (PoE) firmware is always automatically upgraded or downgraded to match the operational code image. This configuration is not applicable to PoE firmware.

Recovering From a Corrupted BootROM--BlackDiamond 8800 Series Switches Only

If your default BootROM image becomes corrupted, you can force the MSM to boot from an alternate BootROM image by inserting a pen into the Alternate (A) and Reset (R) holes on the BlackDiamond 8800 MSM and applying pressure. For more information, please refer to the hardware documentation.

Displaying BootROM and Firmware Versions

To display the BootROM (firmware) version for all modules and PSU controllers installed in the switch, use the `show version` command.

Example

The following command installs the newer firmware image(s):

```
install firmware
```



History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available only on BlackDiamond X8 and BlackDiamond 8800 series switches.

install image

```
install image fname {partition} {msm slotid} {reboot}
```

On a SummitStack use:

```
install image fname {partition} {slot slot number} {reboot}
```

Description

Installs a new version of the ExtremeXOS software image.

Note



Beginning with ExtremeXOS 12.1, an ExtremeXOS core image must be installed on the alternate (non-active) partition. If a user tries to install on an active partition, the following error message is displayed:

```
Error: Image can only be installed to the non-active partition.
```

Syntax Description

<i>fname</i>	Specifies the software image file.
<i>partition</i>	Specifies which partition the image should be saved to: primary or secondary. Select primary to save the image to the primary partition and secondary to save the image to the secondary partition.
<i>slotid</i>	Specifies the MSM/MM where the software image file should be installed. A specifies the MSM/MM installed in slot A.B specifies the MSM/MM installed in slot B. Note: This parameter is available only on modular switches and SummitStack. On a SummitStack, the slotid specifies the node to which the BootROM image should be installed.
reboot	Reboots the switch after the image is installed.

Default

N/A.



Usage Guidelines

When you download a software image, you are asked if you want to install the image immediately after the download is finished. If you choose to install the image at a later time, use this command to install the software on the switch.

The software image file can be an .xos file, which contains an ExtremeXOS core image, or an .xmod file, which contains additional functionality to supplement a core image.

SummitStack Only

You can issue this command only from a Master node. The slot parameter is available only on a stack.

Displaying the Software Image Version

To display the software image version running on the switch, use the `show version` or `show switch` commands.

Displaying the Downloaded Software Image Version.

To display a software image version that has been downloaded but not installed, use the `install image ?` command.

Local Filename Character Restrictions

This section provides information about the characters supported by the switch for local filenames.

When specifying a local filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)

When naming a local, remember the requirements listed above.

Installing an ExtremeXOS core image

Install the software image on the alternate partition. You can continue to run the currently booted image, but to run the newly installed image, you will need to set the boot partition with the `use image {partition} partition {msm slotid}` command and reboot the switch.

Installing an ExtremeXOS module image

An ExtremeXOS module image has functionality that supplements a core image. You will install a module onto an already installed core image. The version number of the core image and the module must match. For example, the module `bd10K-11.0.0.25-ssh.xmod` can only be installed onto the core image `bd10K-11.0.0.25.xos`.



To install a module to the alternate partition, use the `install firmware` command to install the module. Remember, the core image on the alternate partition must be of the same version as the module. When you make the alternate partition active, by issuing the `use image` command and rebooting the switch, the module is also activated at boot time.

To install a module to the active partition, use the `install firmware` command to install the module. Remember, the core image on the active partition must be of the same version as the module. If you reboot the switch, the module will also be activated, but you can activate the module without rebooting the switch by issuing the `run update` command. After issuing that command, all the functionality, and CLI commands, of the module will be available.

Performing a Hitless Upgrade--Modular Switches Only

If you specify the `msm` parameter on a BlackDiamond 8800 series switch, you can initiate hitless upgrade between the primary and backup MSMs installed in the switch.

Hitless upgrade is a mechanism that allows you to upgrade the ExtremeXOS software running on the switch without taking the switch out of service. Some additional benefits of using hitless upgrade include:

- Minimizing network downtime
- Reducing the amount of traffic lost

Although any method of upgrading software can have an impact on network operation, including interrupting Layer 2 network operation, performing a hitless upgrade can decrease that impact.

Regardless of how you upgrade the software, you must:

- Review the following list to confirm that your system supports hitless upgrade:
 - BlackDiamond 8800 series switch with a mix of BlackDiamond 8000 c-, e-, and xl-series modules installed—Both MSMs are running ExtremeXOS 11.5 or later.
- View the current switch information to determine your selected and booted image partitions, verify which MSM is the primary and which is the backup, and confirm that the MSMs are synchronized using the `show switch` command
- Select the partition to use when downloading the image using the `download image [[hostname | ipaddress] filename {{vr} vrname} | memorycard filename] {partition} {msm slotid}` command.

When performing a hitless upgrade, you must:

- Download the software to the backup MSM. Download the image to the alternate partition.
- Use the `install image fname {partition} {msm slotid} {reboot}` command to install the software image at a later time.
- Reboot the Backup MSM using the `reboot msm b` command.
- Use the `run msm-failover` command to initiate failover from the primary MSM to the backup MSM. The original primary MSM becomes the new backup MSM.
- Download the software to the new backup MSM. Again, download the image to the alternate partition, and use the `install image fname {partition} {msm slotid} {reboot}` command to install the software image at a later time.

For more detailed information about hitless upgrade, see the `download image` command.



Example

The following command installs the software image file `bd8800-11.3.0.10.xos` on a BlackDiamond 8810 switch:

```
install image bd8800-11.3.0.10.xos
```

The following command installs the software image file `summitX440-11.5.1.2.xos` on a Summit family switch:

```
install image summitX440-11.5.1.2.xos
```

The following command displays a software image version that has been downloaded but not displayed:

```
install image ?
```

A message similar to the following is displayed:

```
# install image ?  
<fname>      Image file name  
"summitX-12.1.0.52.xos"
```

History

This command was first available in ExtremeXOS 10.1.

The `msm` parameter was added in ExtremeXOS 11.1.

The `slot` parameter was added to support SummitStack in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.

load script

```
load script filename {arg1} {arg2} ... {arg9}
```

Description

Loads (plays back) an ASCII-formatted configuration file or a user-written script file on the switch.



Syntax Description

<i>filename</i>	Specifies the user-defined name of the ASCII-formatted configuration file or a user-written script file. The script file is known as the XOS script file and uses the .xsf file extension.
arg	Specifies up to nine variable values that can be specified by the user. The variables are created with the names CLI.ARGV1, CLI.ARGV2, ... CLI.ARGV9.

Default

N/A.

Usage Guidelines

Use this command to load an ASCII-formatted configuration file or a user-written script file.

Configuration File: After downloading the configuration file from the TFTP server, this command loads and restores the ASCII-formatted configuration file to the switch.

An ASCII-formatted configuration file uses the .xsf file extension, not the .cfg file extension. The .xsf file extension (known as the XOS script file) saves the XML-based configuration in an ASCII format readable by a text editor.

For more detailed information about the ASCII configuration file, including the steps involved to upload, download, and save the configuration, see the `upload configuration [hostname | ipaddress] filename {vr vr-name}` command.

User-Written Script File: After writing a script, this command executes the script and passes arguments to it. As with the configuration files, these files use the .xsf file extension that is automatically added.

The command allows up to nine optional variable values to be passed to the script. These are created with the names CLI.ARGV1, CLI.ARGV2, CLI.ARGV3, ... CLI.ARGV9.

In addition, two other variables are always created. CLI.ARGV0 gives the count of the number of parameters passed, and CLI.ARGV1 contains the name of the script that is being executed.

To check the variable values use the command, `show var`.



Note

Only the .xsf extension is used. The load script command assumes an .xsf extension and retries opening the file if the file cannot be found with the original specified name or no extension is provided.

Example

The following command loads the ASCII-formatted configuration named configbackup.xsf:

```
load script configbackup.xsf
```



After issuing this command, the ASCII configuration quickly scrolls across the screen. The following is an example of the type of information displayed when loading the ASCII configuration file:

```
script.meg_upload_config1.xsf.389 # enable snmp access
script.meg_upload_config1.xsf.390 # enable snmp traps
script.meg_upload_config1.xsf.391 # configure mstp region purple
script.meg_upload_config1.xsf.392 # configure mstp revision 3
script.meg_upload_config1.xsf.393 # configure mstp format 0
script.meg_upload_config1.xsf.394 # create stpd s0
```

History

This command was first available in ExtremeXOS 11.4.

Multiple arguments for user-written scripts were added in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

run update

run update

Description

Activates a newly installed modular software package.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

After you install a modular software package to the active partition, use this command to make the update active. This command causes the ExtremeXOS system to start the newly installed processes contained in the package, without rebooting the switch.

If you installed the package to the inactive partition, you need to reboot the switch to activate the package.



Example

The following command activates any newly installed modular software packages installed on the active partition:

```
run update
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

save configuration

```
save configuration {primary | secondary | existing-config | new-config}
```

Description

Saves the current configuration from the switch's runtime memory to non-volatile memory.

Syntax Description

primary	Specifies the primary saved configuration.
secondary	Specifies the secondary saved configuration.
<i>existing-config</i>	Specifies an existing user-defined configuration.
<i>new-config</i>	Specifies a new user-defined configuration.

Default

Saves the current configuration to the location used on the last reboot.

Usage Guidelines

The configuration takes effect on the next reboot.

Each file name must be unique and can be up to 32 characters long but cannot include any spaces, commas, or special characters.

Configuration files have a .cfg file extension. When you enter the name of the file in the CLI, the system automatically adds the .cfg file extension. Do not use this command with ASCII-formatted configuration files. Those configuration files have an .xsf file extension. For more information about using ASCII-formatted configuration files see the `upload configuration [hostname | ipaddress] filename {vr vr-name}` and the `load script filename {arg1} {arg2} ... {arg9}` commands.



This command also displays in alphabetical order a list of available configurations. The following is sample output that displays the primary, secondary, and user-created and defined configurations (“test” and “XOS1” are the names of the user-created and defined configurations):

```
exsh.9 # save configuration
<cr>           Execute the command
primary        Primary configuration file
secondary      Secondary configuration file
<existing-config> Existing configuration file name
"test" "XOS1"
<new-config>   New configuration file name
```

The switch prompts you to save your configuration changes. Enter y to save the changes or n to cancel the process.

If you enter n, the switch displays a message similar to the following:

```
Save configuration cancelled.
```

If you enter y, the switch saves the configuration and displays a series of messages. The following sections provide information about the messages displayed when you save a configuration on your switch.

Note



Configuration files are forward compatible only and not backward compatible. That is, configuration files created in a newer release, such as ExtremeXOS 12.4, might contain commands that do not work properly in an older release, such as ExtremeXOS 12.1.

Local Filename Character Restrictions

This section provides information about the characters supported by the switch for local filenames.

When specifying a local filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)

When naming a local file, remember the requirements listed above.

Saving a New Configuration

If you create and save a configuration with a new file name, the switch saves the new configuration and then prompts you to select the newly created configuration as the switch’s default configuration.



The following sample output is similar to the message displayed:

```
Do you want to save configuration to test1.cfg? (y/n) Yes
Saving configuration on primary MSM ..... done!
Configuration saved to test1.cfg successfully.
```

The switch then prompts you to select which configuration to use to bootup the system. The following sample output is similar to the message displayed:

```
The current selected default configuration database to boot up the system
(primary.cfg) is different than the one just saved (test.cfg).
Do you want to make test.cfg the default database? (y/n)
```

Enter y to use the new configuration as the default configuration. Enter n to cancel the operation and keep using the current default, active configuration.

Saving an Existing Configuration

If you make and save changes to an existing configuration, the switch prompts you to save and override the existing configuration.

The following sample output is similar to the message displayed:

```
The configuration file test.cfg already exists.
Do you want to save configuration to test.cfg and overwrite it? (y/n) Yes
Saving configuration on primary MSM ..... done!
Configuration saved to test.cfg successfully.
```

The following sample output on a SummitStack is similar to the output displayed:

```
The configuration file primary.cfg already exists.
Do you want to save configuration to primary.cfg and overwrite it? (y/N) Yes
Saving configuration on primary ..... done!
Synchronizing configuration to backup .... done!
Saving config on Standbys (Slots: 1).
...
Configuration saved on Standby (Slot 1): done!
```

If you override an existing configuration that is not the current default, active configuration, the switch prompts you to select which configuration to use to bootup the system. The following sample output is similar to the message displayed:

```
The current selected default configuration database to boot up the system
(primary.cfg) is different than the one just saved (test.cfg).
Do you want to make test.cfg the default database? (y/n) No
Default configuration database selection cancelled.
```



Enter y to use the updated configuration as the default configuration. Enter n to cancel the operation and keep using the current default, active configuration.

Example

The following command saves the current switch configuration to the configuration file named XOS1:

```
save configuration XOS1
```

The following command save the current switch configuration to the secondary configuration file:

```
save configuration secondary
```

History

This command was first available in ExtremeXOS 10.1.

The status messages displayed by the switch were updated in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

save configuration as-script

```
save configuration as-script script-name
```

Description

Saves the running configuration as a script.

Syntax Description

<i>script-name</i>	Specifies the name of the file to save the configuration to. The script file is known as the XOS script file and uses the .xsf file extension.
--------------------	--

Default

N/A.

Usage Guidelines

The save configuration as-script command allows the user to save the current configuration as a script and export it out of the box for later use.



For SummitStack only

The script is saved on all the nodes in a SummitStack when the save configuration as-script command is executed.

Example

The following command saves a running ASCII-formatted configuration named primary.xsf.

```
save configuration as-script primary.xsf
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

show configuration

```
show configuration {module-name} {detail}
```

Description

Displays the current configuration for the system or the specified module.

Syntax Description

<i>module-name</i>	Specifies the name of configuration module. The term configuration module refers to feature in ExtremeXOS. By displaying a module, you can view the commands used to configure that feature. For example, to display all of the configurations that you made for only STP, specify the stp as the module-name.
detail	Displays configuration data including default. If the detail option is not specified, only the configuration changes you made to the factory defaults are shown.

Default

N/A.

Usage Guidelines

If the output scrolls off the top of the screen, you can use the `enable clipaging` command to pause the display when the output fills the screen. The default for clipaging is enabled.



We recommend using the `show configuration` command to view on the CLI your currently running switch configuration. These files have the .cfg file extension. Do not use a text editor to view or modify your XML-based switch configuration files.

To save the configuration file as an ASCII-formatted file, and to view it with a text editor, see the `upload configuration [hostname | ipaddress] filename {vr vr-name}` and the `load script filename {arg1} {arg2} ... {arg9}` commands.

Beginning with ExtremeXOS 12.1, when you specify `show configuration` only, the switch displays configuration information for each of the switch modules excluding the default data.

You can display only the configuration of a module of interest by using the module-name keyword. For example, some of the modules are AAA, ACL, BGP, EDP, FDB, SNMP, and VLAN. Use [TAB]-completion to see a list.

You must have administrator access to view the output of the `show configuration` command.

Depending on the software version running on your switch, the configurations on your switch, and the type of switch you have, additional or different configuration information may be displayed.

Example

This command shows the current configuration of the OSPF module in the switch:

```
show configuration ospf
```

The following is sample output from this command:

```
# Module ospf configuration.
#
configure ospf routerid automatic
configure ospf spf-hold-time 3
configure ospf metric-table 10M 10 100M 5 1G 4 10G 2
configure ospf lsa-batch-interval 30
configure ospf import-policy none
configure ospf ase-limit 0
disable ospf originate-default
disable ospf use-ip-router-alert
disable ospf
configure ospf restart none
configure ospf restart grace-period 120
disable ospf export direct
disable ospf export static
disable ospf export rip
disable ospf export e-bgp
disable ospf export i-bgp
configure ospf area 0.0.0.0 external-filter none
configure ospf area 0.0.0.0 interarea-filter none
SummitX440-24x #
```



History

This command was first available in ExtremeXOS 11.0.

The detail variable was added in ExtremeXOS 12.1.

The display of blackhole output was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

show licenses

show licenses

Description

Displays current software license level and feature packs enabled on your switches.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The command displays information on the software license level and feature packs enabled on the switch, including the trial license and days left to expiry.

To see the license information on a SummitStack node, run the command while logged into that node.

Note



Refer to the specific chapter that discusses each feature of the *ExtremeXOS Concepts Guide* to determine if a license is required for some functionality. If not noted, all functionality is available, and license is not required.

Example

The following command displays the license level configuration:

```
show licenses
```



On a SummitStack, the output from this command looks similar to the following:

```
Slot-3 Stack.12 > show licenses
Enabled License Level:
Advanced Edge
Enabled Feature Packs:
None
Effective License Level:
Edge
```

The show output also includes the current status of the non-Extreme optical device license:

```
# show license
Enabled License Level:
    Advanced Core
Enabled Feature Packs:
    MPLS ServiceProviderEdge NonExtremeOptics
```

History

This command was first available in ExtremeXOS 11.1.

The information on enabled feature packs was added in ExtremeXOS 11.4.

The information on the trial licenses was added in ExtremeXOS 11.6.

The show output was updated to include the current status of the non-Extreme optical device license in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

show memorycard

show memorycard

Description

Displays whether a compact flash card or USB 2.0 storage device is present in the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.



Usage Guidelines

None.

Example

This command shows whether a removable storage device is present on the switch:

```
show memorycard
```

If you do not have a removable storage device installed, the output is similar to the following:

```
Memorycard is not present.
```

On the BlackDiamond X8 switch, there are two USB ports, and so the command indicates which port is being used:

```
BD-X8.1 # show memorycard  
Memorycard is present in USB-1.
```

History

This command was first available in ExtremeXOS 11.0.

Support for USB 2.0 storage devices was added in ExtremeXOS 12.5.3.

Platform Availability

This command is available on BlackDiamond X8 and 8800 series switches and Summit X460, X480, X670, X670V, and X770 switches.

show script output autoexec

```
show script output autoexec
```

Description

Shows the results of executing the autoexec script.

Syntax Description

This command has no arguments or variables.

Default

N/A



Usage Guidelines

Use this command to show results when a `autoexec.xsf` file is executed. The file is not executed when a `default.xsf` file has been executed.

The CLI script file `autoexec.xsf` is executed after the configuration has been loaded. Its purpose is to run some commands after every reboot. It can also be used to revert to the original configuration following changes made by UPM executed persistent commands.

Example

This command shows the results of executing the `autoexec` script:

```
show script output autoexec
```

When there is no `autoexec.xsf` file, there is no response.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

show script output default

```
show script output default
```

Description

Shows the results of executing `default.xsf` on bootup.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to show results when a `default.xsf` file is loaded.

An existing `default.xsf` file is executed if the switch comes up in an unconfigured state because the configuration file is missing, or the configuration file cannot be determined due to a corrupt NVRAM or



other problems. This returns the switch to some basic configuration. When default.xsf is executed, the show switch command shows default.xsf as the booted configuration file.

Example

This command shows the results of executing the autoexec script:

```
show script output default
```

When there is no default.xsf file, there is no response.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms.

synchronize

```
synchronize {slot slotid}
```

Description

The synchronize command replicates all saved images and configurations from the primary MSM/MM or node to the backup MSM/MM or target node on a switch or SummitStack.

Syntax Description

<i>slotid</i>	On a SummitStack, the slotid specifies the target node that has to be synchronized with the Master node. If the slotid is omitted, the target is the backup node. NOTE: This parameter is available only with SummitStack.
---------------	---

Default

N/A.

Usage Guidelines

This command:

- Reboots the backup MSM/MM or target node to prepare it for synchronizing with the primary MSM/MM or node.



- Performs a binary copy of the primary MSM/MM or node to the backup MSM/MM or target node, including the primary and secondary software images, all configurations and policies, and temporary files.
- Reboots the backup MSM/MM or target node after replication is complete.

During a synchronization, half of the switch fabric is lost. When the primary MSM/MM or node finishes replicating its configurations and images to the backup MSM/MM or target node, the full switch fabric is restored.

To use the `synchronize` command make sure your switch or SummitStack is running the following software:

- BlackDiamond 8800 series switch with a mix of BlackDiamond 8000 c-, e-, and xl-series modules installed—Both MSMs are running ExtremeXOS 11.5 or later.
- SummitStack—All nodes are running ExtremeXOS 12.0 or later.

When you install a backup MSM/MM, you are not prompted to synchronize the images and the configurations from the primary. If not synchronized, the backup uses its image and the primary's configuration. This image/configuration mismatch will likely cause the switch to operate differently after failover. Use the `synchronize` command to replicate all saved images and configurations from the primary to the backup.

If you have not saved your runtime configuration, you are prompted to save it when you use the `synchronize` command. A message similar to the following appears:

```
Do you want to save configuration changes to primary.cfg? (y or n)
```

Enter y to save the configuration and continue with synchronizing the MSMs/MMs. Enter n to cancel the operation. If you enter y, messages similar to the following appear:

```
Saving configuration on primary MSM ..... done! Synchronizing
configuration to backup MSM .. done!
```

After the configuration has been saved and replicated to the backup MSM/MM, synchronization begins.

After the initial reboot, if the backup MSM/MM is not available or does not respond within 120 seconds, the synchronize operation fails.

Use the `show switch {detail}` command to verify that the backup MSM/MM is in sync with the primary MSM/MM.

BlackDiamond 8800 Series Switch Only

On a BlackDiamond 8800 series switch, the I/O ports on the backup MSM go down when you synchronize the MSMs. When the primary MSM finishes replicating its configurations and images to the backup MSM, the I/O ports on the backup MSM come back up.

SummitStack Only.

While using the command `synchronize {slot slot-number}` if the slot number is provided, that slot is the target of the synchronize operation. If the slot number is not provided, the Backup node is synchronized. This command can be executed only on the Master node.



The synchronize command preserves the following stacking configurations on the target node:

- slot number
- master-capable configuration
- alternate IP address, subnetwork mask, and default gateway
- priority
- license restriction

Thus a synchronized node comes up in the same place in the active topology that it occupied before the synchronize command was issued.

If the synchronizing switch cannot determine the stacking support configuration on the target switch, the following message is displayed:

```
Error: Information for the target switch is temporarily unavailable.
Please retry the command.
```

ExtremeXOS software does not allow a synchronize operation on a SummitStack between a Summit X460 or X670 switch and a Summit X480 switch. If one is attempted, the following message is displayed:

```
Error: the target slot's partitions are not compatible with the Master's
for synchronize.
```

Example

The following example assumes you have already saved your runtime configuration.

The following command replicates all saved images and configurations from the master MSM to the backup MSM:

```
synchronize
```

After you enter synchronize, status messages similar to the following appear:

```
Synchronize will reboot the backup MSM, then overwrite all code images
and configs with a copy from the master MSM.
Synchronize currently requires ExtremeXOS version 11 or greater on
the backup MSM
DO NOT interrupt synchronize, the backup MSM may become unbootable!
OK to continue? (y/n) Yes
Rebooting Backup MSM...
NOTE: The command line is locked during synchronize
synchronizing...
synchronizing nvram...
synchronizing nvram...
synchronizing nvram...
synchronizing nvram...
synchronizing nvram...
synchronizing nvram...
synchronizing XOS...
[=====] 100% XOS
```



```
Synchronize complete - rebooting backup MSM...
BD-8808.4 #
```

History

This command was first available in ExtremeXOS 11.0.

The slot parameter was added to support SummitStack in ExtremeXOS 12.0.

Platform Availability

This command is available only on modular switches and SummitStack.

unconfigure switch

```
unconfigure switch {all}
```

Description

Returns the switch configuration to its factory default settings and reboots the switch.

Syntax Description

all	Specifies that the entire configuration should be changed to the default values, including the management IP address, failsafe account, and SummitStack-specific parameters, and the switch rebooted.
------------	---

Default

N/A.

Usage Guidelines

Use `unconfigure switch` to reset the configuration to factory defaults, but without erasing the configuration. This preserves users account information, date and time settings, SummitStack configuration, and so on.

Include the parameter `all` to clear the entire current configuration, including all switch and SummitStack parameters, and reboot using the last used image and factory default configuration.

The command `unconfigure switch all` does not clear licensing information. The license cannot be disabled once it is enabled on the switch.

For SummitStack only.

The `all` option also resets all stacking-specific parameters to defaults. To reset only the stacking-specific parameters to defaults, enter the `unconfigure stacking` command.



Beginning with ExtremeXOS 12.5, stacking support and stacking port selection are reset only on the local node. When stacking support of any kind is supported on the platform, the following message is added to the output that is shown on the console after this command has been confirmed:

```
Stacking-support will be unconfigured on this node only.
```

Example

The following command preserves the entire current configuration (but does not reload the current configuration after the switch reboots) and reboots the switch or SummitStack using the last specified saved image and factory default configuration:

```
unconfigure switch all
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

uninstall image

```
uninstall image fname partition {msm slotid} {reboot}
```

On a SummitStack, use:

```
uninstall image fname partition {slot slot number } {reboot}
```

Description

Uninstalls an ExtremeXOS software package.

Syntax Description

<i>fname</i>	Specifies the software package to uninstall.
<i>partition</i>	Specifies which partition the package was installed to: primary or secondary. Select primary to remove it from the primary partition and secondary to remove it from the secondary partition.
<i>slotid</i>	Specifies the MSM where the package should be uninstalled. A specifies the MSM/MM installed in slot A.B specifies the MSM/MM installed in slot B. Note: This parameter is available only on modular switches and SummitStack. On a SummitStack, the slot number specifies the node on which the BootROM image should be uninstalled.
reboot	Reboots the switch after the package is uninstalled.



Default

N/A.

Usage Guidelines

Use this command to uninstall a software package previously installed on the switch.

When you uninstall a software package, the switch prompts you to save your changes to your currently active configuration file:

Uninstallation of the EXOS module Do you want to save configuration changes to primary.cfg? (y or n)

Enter **y** to save the changes to your configuration file. Enter **n** to not save the changes to your configuration file.

Local Filename Character Restrictions

This section provides information about the characters supported by the switch for local and remote filenames.

When specifying a local filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)

When naming a local file, remember the requirements previously described.

SummitStack Only

You can issue this command only from a Master node.

Example

The following command uninstalls the modular software package for Secure Shell, bd10K-11.2.0.18-ssh.xmod, from the secondary partition:

```
uninstall image bd10K-11.2.0.18-ssh.xmod secondary
```

History

This command was first available in ExtremeXOS 11.0.

The msm parameter was added in ExtremeXOS 11.1.

The slot parameter was added to support SummitStack in ExtremeXOS 12.0.



Platform Availability

This command is available on all platforms.

upload configuration

```
upload configuration [hostname | ipaddress] filename {vr vr-name}
```

Description

Uploads the current configuration in ASCII format to a TFTP server on your network.

Syntax Description

<i>hostname</i>	Specifies the hostname of the TFTP server where you want to download the configuration file. You must have DNS enabled.
<i>ipaddress</i>	Specifies the IP address of the TFTP server where you want to download the configuration file.
<i>filename</i>	Specifies a user-defined name for the configuration file. You must use the .xsf file extension when naming an ASCII-formatted configuration file.
<i>vr-name</i>	Specifies the name of the virtual router. By default the switch uses VR-Mgmt for this command. NOTE: User-created VRs are supported only on the platforms listed for this feature in the Feature License Requirements in the ExtremeXOS Concepts Guide.

Default

Uploads the current configuration in ASCII format immediately to a TFTP server.

Usage Guidelines

Specify the *ipaddress* or *hostname* parameters to upload the current, active configuration file from the switch to a TFTP server on the network. Use of the *hostname* parameter requires that DNS be enabled.

The uploaded ASCII file retains the CLI format. This allows you to do the following:

- Modify the configuration using a text editor, and later download a copy of the file to the same switch or to one or more different switches.
- Send a copy of the configuration file to Extreme Networks Technical Support for problem-solving purposes.

This command is not applicable to XML-based configurations. Those files use the .cfg file extension.

If you want to view your configuration in ASCII format, use the .xsf file extension (known as the XOS script file) when you save the configuration file on the switch. This saves the XML-based configuration in an ASCII format readable by a text editor.

If you successfully upload the active configuration to the network TFTP server, the switch displays a message similar to the following:



```
Uploading meg_upload_config1.xsf to 10.10.10.10 ... done!
```

If the switch displays a timeout error message similar to the following:

```
failed! Error: timeout
```

Make sure you entered the correct host name or IP address of the TFTP server

If the switch displays an unreachable network error similar to the following:

```
failed! Error: Network is unreachable
```

Make sure you entered the correct virtual router. By default the switch uses VR-Mgmt for this command.

Summary of Steps

The following summary only describes the CLI involved to transfer the configuration and load it on the switch; it is assumed that you know how to modify the configuration file with a text editor. As previously described, to use these commands, use the .xsf file extension. These steps are not applicable to configurations that use the .cfg file extension.

To work with an ASCII-formatted configuration file, complete the following tasks:

- Upload the configuration to a network TFTP server using the following command:

```
upload configuration [hostname | ipaddress] filename {vr vr-name}
```

After the configuration file is on the TFTP server, use a text editor to the desired changes.

- Download the configuration from the TFTP server to the switch using one of the following commands:

```
tftp [host-name | ip-address] -g -r remote-file
```

```
tftp get [host-name | ip-address]remote-file
```

- Verify the configuration file is on the switch using the following command:

```
ls
```

- Load and restore the new configuration file on the switch using the following command:

```
load script filename {arg1} {arg2} ... {arg9}
```

- Save the configuration to the configuration database so the switch can reapply the configuration after switch reboot using the following command:

```
save configuration {primary | secondary | existing-config | new-config}
```

When you save the configuration file, the switch automatically adds the .cfg file extension to the filename. This saves the ASCII configuration as an XML-based configuration file.

The following describes the steps in more detail.



Uploading the ASCII Configuration File To a TFTP Server

To upload the current switch configuration as an ASCII-based file to the TFTP server, use the `upload configuration` command and save the configuration with the `.xsf` file extension.

For example, to transfer the current switch configuration as an ASCII-based file named `meg_upload_config1.xsf` to the TFTP server with an IP address of 10.10.10.10, do the following:

```
upload configuration 10.10.10.10 meg_upload_config1.xsf
```

If you successfully upload the configuration to the TFTP server, the switch displays a message similar to the following:

```
Uploading meg_upload_config1.xsf to 10.10.10.10 ... done!
```

Downloading the ASCII Configuration File to the Switch

To download the configuration from the TFTP server to the switch, use the `tftp` command. For example, to retrieve the configuration file named `meg-upload_config1.xsf` from a TFTP server with an IP address of 10.10.10.10, you can use one of the following commands:

```
tftp 10.10.10.10 -g -r meg_upload_config1.xsf
tftp get 10.10.10.10 meg_upload_config1.xsf
```

If you successfully download the configuration to the switch, the switch displays a message similar to the following:

```
Downloading meg_upload_config1.xsf to switch... done!
```

Verifying that the ASCII Configuration File is on the Switch

To confirm that the ASCII configuration file is on the switch, use the `ls` command. The file with an `.xsf` extension is the ASCII configuration.

The following sample output contains an ASCII configuration file:

```
-rw-r--r-- 1 root 0 98362 Nov 2 13:53 Nov022005.cfg
-rw-r--r-- 1 root 0 117136 Dec 12 12:56 epicenter.cfg
-rw-r--r-- 1 root 0 68 Oct 26 11:17 mcastgroup.pol
-rw-r--r-- 1 root 0 21203 Dec 13 15:40 meg_upload_config1.xsf
-rw-r--r-- 1 root 0 119521 Dec 6 14:35 primary.cfg
-rw-r--r-- 1 root 0 96931 Nov 11 11:01 primary_11_11_05.cfg
-rw-r--r-- 1 root 0 92692 Jul 19 16:42 secondary.cfg
```

Loading the ASCII Configuration File

After downloading the configuration file, you must load the new configuration on the switch. To load and restore the ASCII configuration file, use the `load script filename {arg1} {arg2} ...`



{arg9} command. After issuing this command, the ASCII configuration quickly scrolls across the screen.

The following is an example of the type of information displayed when loading the ASCII configuration file:

```
script.meg_upload_config1.xsf.389 # enable snmp access
script.meg_upload_config1.xsf.390 # enable snmp traps
script.meg_upload_config1.xsf.391 # configure mstp region purple
script.meg_upload_config1.xsf.392 # configure mstp revision 3
script.meg_upload_config1.xsf.393 # configure mstp format 0
script.meg_upload_config1.xsf.394 # create stpd s0
```

Instead of entering each command individually, the script runs and loads the CLI on the switch.

Saving the Configuration

After you load the configuration, save it to the configuration database for use by the switch. This allows the switch to reapply the configuration after a switch reboot. To save the configuration, use the `save configuration {primary | secondary | existing-config | new-config }` command.

When you save the configuration file, the switch automatically adds the .cfg file extension to the filename. This saves the ASCII configuration as an XML-based configuration file.

You can use any name for the configuration. For example, after loading the file `meg_upload_config1.xsf`, you need to save it to the switch. To save the configuration as `configuration1.cfg`, do the following:

```
save configuration configuration1
```

Host Name and Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-) Permitted only for host names
- Underscore (_) Permitted only for host names
- Colon (:)

When naming or configuring an IP address for your network server, remember the requirements listed above.

Remote Filename Character Restrictions

This section provides information about the characters supported by the switch for remote filenames.



When specifying a remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)
- Slash (/) Permitted only for remote files

When naming a remote file, remember the requirements previously described.

Example

The following command uploads the current switch configuration as an ASCII-based file named `configbackup.xsf` to the TFTP server with an IP address of 10.10.10.10:

```
upload configuration 10.10.10.10 configbackup.xsf
```

If you successfully upload the configuration to the TFTP server, the switch displays a message similar to the following:

```
Uploading configbackup.xsf to 10.10.10.10 ... done!
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on all platforms.

use configuration

```
use configuration [primary | secondary | file_name]
```

Description

Configures the switch to use a previously saved configuration on the next reboot.

Syntax Description

primary	Specifies the configuration file named <code>primary.cfg</code> .
secondary	Specifies the configuration file named <code>secondary.cfg</code> .
<i>file_name</i>	Specifies an existing user-defined configuration file name (displays a list of available user-defined configuration files).



Default

N/A.

Usage Guidelines

XML-based configuration files have a .cfg file extension. When you enter the name of the file in the CLI, the system automatically adds the .cfg file extension.

Do not use this command with ASCII-formatted configuration files. Those configuration files have an .xsf file extension. For more information about using and saving ASCII-formatted configuration files see the `upload configuration [hostname | ipaddress] filename {vr vr-name }` and the `load script filename {arg1} {arg2} ... {arg9}` commands.

There is no special significance to the primary and secondary configurations. They are just conveniences to specify the files primary.cfg and secondary.cfg.

When you configure the switch to use a previously saved configuration, the switch displays the following message:

```
The selected configuration will take effect after the next switch reboot.
```

You can create a new configuration file by saving your current switch configurations and using that file on the next reboot. For example, to create a new configuration named test1 based on your current CLI session and switch configurations, use the following command:

```
save configuration test1
```

Tracking and Displaying Switch Configuration Files

To keep track of your configuration file names, use the `ls` command to display the files saved on your switch. Files with the .cfg extension are configuration files. In addition, you can see a list of available configuration files when you use the `use configuration` command.

The following is sample output from this command (“test” and “XOS1” are the names of the user-created and defined configurations):

```
exsh.1 # use configuration
primary          Primary configuration file
secondary        Secondary configuration file
<file-name>     Configuration file name
"test" "XOS1"
```

You can also use the `ls` command to display a list of the current configuration and policy files in the system.

Displaying the Active Configuration

To view the currently active, running configuration, use the `show switch` command.



Local Filename Character Restrictions

This section provides information about the characters supported by the switch for local filenames.

When specifying a local filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)

When naming a local file, remember the requirements listed above.

Example

The following command specifies that the next reboot should use the saved configuration file named XOS1.cfg:

```
use configuration XOS1
```

The following command specifies that the next reboot should use the configuration saved in the primary partition:

```
use configuration primary
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

use image

```
use image {partition} partition {msm slotid}
```

On a SummitStack, use:

```
use image {partition} partition {slot slotid}
```

Description

Configures the switch to use a saved image on the next reboot.



Syntax Description

<i>partition</i>	Specifies which image to use on the next reboot, the one stored on the primary partition, or the one stored on the secondary partition.
<i>slotid</i>	A specifies the MSM/MM installed in slot A.B specifies the MSM/MM installed in slot B. NOTE: This parameter is available only on modular switches and SummitStack. On a SummitStack, the slotid specifies the node on which the BootROM image is selected.

Default

The currently booted image.

Usage Guidelines

This command specifies which image to use on the next reboot. Two images can be stored, one on the primary partition, one on the secondary partition. To view your current (active) partition and the selected partition for the next reboot or installation, use the following command:

```
show switch
```

Output from this command includes the selected and booted images and if they are in the primary or the secondary partition. Primary indicates the saved image in the primary partition; secondary indicates the saved image in the secondary partition.

SummitStack Only

You can issue this command only from a Master node. The image to use is stored in NVRAM on all target nodes.

If a slot number is not provided, the partition is selected on all nodes in the Active Topology.

Example

Using TFTP

The following command configures the switch to use the image stored in the primary partition on the next reboot:

```
use image partion primary
```

A message similar to the following is displayed:

```
To take effect of partition change please reboot the switch!
```



History

This command was first available in ExtremeXOS 10.1.

The msm parameter was added in ExtremeXOS 11.1.

The slot parameter was added to support SummitStack in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms.



48 CNA Agent Commands

```
clear cna-testplug counters
configure cna-testplug scheduler ipaddress
configure cna-testplug vlan
disable cna-testplug
enable cna-testplug
show cna-testplug
```

The Converged Network Analyzer (CNA) Agent is part of the CNA software from Avaya Inc.[®] that is used to test network conditions. Use the CNA Agent only if your network includes an Avaya solution that uses CNA.

The entire CNA software package consists of multiple parts. The user interface is a combination of a Java applet hosted from the CNA Server and a Command Line Interface (CLI). You obtain all parts (except the CNA Agent) from your Avaya representative, along with the accompanying documentation. You configure and manage the CNA Agent using the CLI.

The CNA Agent is a downloadable software module that is used when running the Avaya CNA software.



Note

You must download and install the Secure Shell (SSH2) software module prior to downloading and installing the CNA Agent software module.

Using the CNA software, the CNA Agent runs the requested tests and returns the test results. The CNA Agent runs the following tests as directed by the CNA Server:

- Traceroute
- RTP
- Ping
- TCPconnect
- Merge

You enable the software and configure the CNA Agent to communicate with the CNA Server, to clear the test counters, and to display connection status and test results.

clear cna-testplug counters

```
clear cna-testplug counters
```

Description

Clears all counters maintained by the CNA Agent and resets the counters to zero.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

You can also use the `clear counters` command to reset the internal counters for the CNA Agent and return them to 0.

Example

The following command clears all the counters on the CNA Agent and returns the values to zero:

```
clear cna-testplug counters
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure cna-testplug scheduler ipaddress

```
configure cna-testplug scheduler ipaddress ip_address
```

Description

Configures the IP address of the CNA Server using open Secure Socket Layer (SSL); you enter the IP address of the CNA server.

Syntax

<i>ip_address</i>	Specifies the IP address of the CNA Server that communicates with the CNA Agent to schedule tests and receive the results.
-------------------	--

Default

N/A.



Usage Guidelines

Use this command to configure the CNA Agent with the IP address of the CNA Server. The CNA Server requests the timing and type of networking testing, and the CNA Agent runs the tests.



Note

You use the CNA Agent only if you are running the Avaya CNA solution; you must have other pieces of the CNA (available from Avaya) to run these tests.

You enter the IP address of the CNA server.

This command sets up the encryption key that is subsequently used for all communication between the CNA Agent and the CNA Server.



Note

You must have previously installed the Secure Shell (SSH2) downloadable software module, which contains SSL, to use the CNA Agent software.

Example

The following command enters the CNA Server's IP address to 10.6.13.116; the CNA Agent uses this IP address to communicate with the CNA Server:

```
configure cna-testplug scheduler ipaddress 10.6.13.116
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

configure cna-testplug vlan

```
configure cna-testplug vlan vlan_name
```

Description

Configures the CNA Agent (test plug) to an interface. By default, the CNA Agent is bound to the Default VLAN.

Syntax

<i>vlan_name</i>	Specifies which interface IP address the CNA Server uses to communicate with the CNA Agent.
------------------	---



Default

Default VLAN.

Usage Guidelines

The interface IP address is specified when setting up the SSL connection with the CNA Server (when you issue the `configure cna-testplug scheduler ipaddress` command. The CNA Server attempts to establish the socket connection on the interface specified in the `configure cna-testplug vlan` interface to conduct the actual tests.

The system uses the primary IP address if the VLAN has more than one IP address. By default, the CNA Server uses the Default VLAN.



Note

Extreme Networks recommends that you put IP telephones on the same virtual router.

Example

The following command instructs the CNA Server to use the interface associated with VLAN gateway to conduct tests:

```
configure cna-testplug vlan gateway
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

disable cna-testplug

```
disable cna-testplug
```

Description

Disables the CNA Agent.

Syntax

This command has no arguments or variables.



Default

Disabled.

Usage Guidelines

You must enable the CNA Agent before it can run the network tests requested by the CNA Server.

Example

The following command disables the CNA Agent:

```
disable cna-testplug
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

enable cna-testplug

```
enable cna-testplug
```

Description

Once enabled, the CNA Agent coordinates with the CNA Server to test the network for throughput.

Syntax

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

If you previously configured the IP address for the CNA Server and the VLAN interface, the CNA Agent immediately registers with the CNA Server upon being enabled and begins running the requested tests.



Example

The following command enables the CNA Agent:

```
enable cna-testplug
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.

show cna-testplug

show cna-testplug

Description

Displays the statistics and connection status with the CNA software from Avaya. The display includes configured CNA Agent (test plug) and CNA Server (scheduler) connections and the number of tests conducted on each connection.

Syntax

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the following information:

- Hardware Name—The name of the Extreme Networks device running the CNA Agent (test plug)
- Firmware version—The version of ExtremeXOS firmware running on the device
- Interface VLAN—VLAN (and virtual router) interface the CNA Server uses to schedule and run the tests
- IP address on the interface—IP address of the interface the CNA Server uses; the connection on which the CNA Agent tests and sends results
- CNA Test plug version—Version of CNA Agent (test plug) software, in the following format: MajorRev.MinorRev.Build
- Interface version with Scheduler—Shows compatibility between the CNA Agent (test plug) and the CNA Server (scheduler) software
- Enabled or Disabled



- Status—Connection status to the CNA Server:
 - Registered
 - Unregistered
- Errors—Number of errors in tests of connectivity
- Total tests received—Total test requests received by the CNA Agent
- Scheduler (SBC)—IP address and port number on which the CNA Server communicates with the CNA Agent; the connection on which the CNA Agent listens
- Listening ports
 - Test requests(from ANS)—Number of tests that the CNA Agent was requested to run by the Adaptive Networking Software (ANS) on the CNA Server
 - RTP test requests(from test plugs)—Number of RTP streams that the CNA Agent initiates
 - RTP and Traceroute responses(from test plugs)—Number of responses by the CNA Agent to requests for RTP and Traceroute tests
- Last Test—Last test that CNA Agent performed
- Result Last Test—Results of the last test, which the CNA Agent sends to the CNA Server
- Test—Tests run by the CNA Agent
- Count—Number of tests successfully run

Example

The following command displays CNA Agent statistics and connection status on the Summit series switch:

```
show cna
HW Name:                SummitX
Firmware version:       11.3.0.11
Interface VLAN:         "Default" on Virtual router "VR-Default"
IP address of the Interface: 10.203.128.126
CNA Test plug version:  3.0.2
CNA Interface version:  17
Admin:                  Enabled
Status:                 Registered
Errors:                 22
Total tests received:   23529
Scheduler (SBC):        10.203.128.127
Listening ports
--Test requests(from ANS):           50000
--RTP test requests(from test plugs): 50001
--RTP and Traceroute responses(from test plugs): 50016
Last Test: RTP to 10.203.128.124
Result Last Test: delay: 0.136375, jitter: 0.007525, loss: 0.000000(in milli
secs)
Test Results:
Name          COUNT          FAILED
-----
Traceroute    74898            0
RTP           145146           22
Ping          0                0
```



Tcpconnect	0	0
Merge	93	0

**Note**

Adaptive Networking Software (ANS) runs on the CNA Server.

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms.



A Troubleshooting Commands

Extreme Loop Recovery Protocol
disable log debug-mode
clear elrp counters
clear esvt traffic-test
configure debug core-dumps
configure elrp-client disable ports
configure elrp-client one-shot
configure elrp-client periodic
configure forwarding fabric hash
configure forwarding hash-algorithm
configure forwarding hash-recursion-level
configure tech-support add collector
configure tech-support collector
configure tech-support collector data-set
configure tech-support collector frequency error-detected
configure tech-support collector report
configure tech-support delete collector
disable elrp-client
disable led locator
disable log debug-mode
disable tech-support collector
eject memorycard
enable elrp-client
enable led locator
enable log debug-mode
enable tech-support collector
nslookup
run diagnostics
run elrp
run esvt traffic-test
run tech-support report
save debug tracefiles memorycard
show debug
show diagnostics
show elrp
show elrp disabled-ports
show esvt traffic-test


```
show forwarding configuration
show tech support
show tech-support collector
stop esvt traffic-test
top
unconfigure elrp-client
unconfigure elrp-client disable ports
upload debug
```

If you encounter problems when using your switch, ExtremeXOS provides troubleshooting commands. Use these commands only under the guidance of Extreme Networks technical personnel.

This appendix describes commands for troubleshooting your switch, including:

- Running diagnostics and displaying diagnostic test results.
- Enabling and disabling the standalone Extreme Loop Recovery Protocol (ELRP) client.
- Enabling and disabling debug mode for Event Management System (EMS) components.

You can contact Extreme Networks Technical Support at (800) 998-2408 or (408) 579-2826.

Extreme Loop Recovery Protocol

Extreme Loop Recovery Protocol (ELRP) is a feature of ExtremeXOS that allows you to detect Layer 2 loops in the network.

You can use ELRP with other protocols such as the Extreme Standby Router Protocol (ESRP) and the Ethernet Automatic Protection Switching Protocol (EAPS).

A switch running ELRP transmits multicast packets with special MAC destination address out of some or all of the ports belonging to a Virtual LAN (VLAN). All the other switches in the network treat this packet as a regular, multicast packet and flood it to all the ports belonging to the VLAN. If the packets transmitted by a switch are received back by that switch, this indicates a loop in the Layer 2 network.

Standalone ELRP provides the ability to send ELRP packets, either periodically or on an ad hoc “one-shot” basis on a specified subset of VLAN ports. If any of these transmitted packets is received back then standalone ELRP can perform configured actions such as sending a log message to the system log file and, in the case of periodic transmission, sending a trap to the SNMP manager and disabling the port where the packet ingresssed.

ELRP is discussed in more detail in the *ExtremeXOS Concepts Guide* chapter, [Troubleshooting](#).

Details of using ELRP with ESRP are discussed in the *ExtremeXOS Concepts Guide* chapter, [ESRP](#), and the commands used to configure ELRP with ESRP are described in [ESRP Commands](#).

disable log debug-mode

```
disable log debug-mode
```



Description

Disables debug mode. The switch stops generating debug events.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables debug mode. Debug mode must be enabled prior to configuring advanced debugging capabilities. These include allowing debug messages, which can severely degrade performance. For typical network device monitoring, debug mode should remain disabled, the default setting. Debug mode should only be enabled when advised by technical support, or when advanced diagnosis is required. The debug mode setting is saved to FLASH.

The following configuration options require that debug mode be enabled:

- Including a severity of debug-summary, debug-verbose, or debug-data when configuring filters.
- Target format options process-name, process-id, source-function, and source-line.

Example

The following command disables debug mode:

```
disable log debug-mode
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

clear elrp counters

```
clear elrp counters
```

Description

Clears and resets the ELRP counters.



Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

You should view the switch statistics before you delete the ELRP counters. Use the `show log counters` command to display event statistics.

Viewing and maintaining statistics on a regular basis allows you to see how well your network is performing. If you keep simple daily records, you will see trends emerging and notice problems arising before they cause major network faults. By clearing the counters, you can see fresh statistics for the time period that you are monitoring.

Example

The following command clears all switch statistics related to ELRP:

```
clear elrp counters
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

clear esvt traffic-test

```
clear esvt traffic-test {{vlan} vlan_name}
```

Description

Clears all measurements of the Service verification test.

Syntax Description

vlan	Specifies the VLAN for the traffic test.
<i>vlan_name</i>	Specifies a VLAN name.



Default

N/A.

Usage Guidelines

Use this command to clear all measurements of the Service verification test .

Example

The following example clears all measurements of the Extreme Service Verification tool on vlan1:

```
clear esvt traffic-test vl
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available in X460, X480, E4G-200, E4G-400, X670, X770, Stacking, BDx8 and BD8800.

configure debug core-dumps

```
configure debug core-dumps [internal-memory | memorycard | off]
```

Description

Enables or disables the sending of core dump files to the internal memory card, a compact flash card, or a USB 2.0 storage device.

Syntax Description

internal-memory	Specifies that saving debug information to the internal memory card is enabled. This is the default behavior. Use this parameter only under the guidance of Extreme Networks Technical Support personnel.
memorycard	Enables saving debug information to a removable storage device, which can be a compact flash card or a USB 2.0 storage device. Use this parameter only under the guidance of Extreme Networks Technical Support personnel. NOTE: This parameter is available for compact flash cards on BlackDiamond 8800 and for USB 2.0 storage devices on Summit X460, X480, X670, and X670V switches.
off	Specifies that the switch does not save core dump files to memory or to removable storage devices.



Default

Beginning with ExtremeXOS 11.6, configure debug core-dumps internal-memory is enabled by default.

Usage Guidelines



Note

Use this command only under the guidance of Extreme Networks Technical Support personnel to troubleshoot the switch.

The switch only generates core dump files and writes them to the specified device in the following situations:

- If an ExtremeXOS process fails.
- When forced under the guidance of Extreme Networks Technical Support.

If you configure the switch to write core dump files to the internal memory and attempt to download a new software image, you might have insufficient space to complete the image download. If this occurs, move or delete the core dump files from the internal memory. For example, if the switch supports a removable storage device that has space available, transfer the files to the device. On switches without removable storage devices, transfer the files from the internal memory card to a TFTP server. This frees up space on the internal memory card while keeping the core dump files.

Before you can enable and save debug information to a removable storage device, you must install the device. For more information about installing a removable storage device, refer to the hardware documentation.

After you use the `eject memorycard` command and manually remove a removable storage device, this setting is automatically changed to off.

Stackables in Stack Mode

This command works only from the master node. If you enable it on stack master, it is applicable for all nodes.

Example

The following example enables a switch to save debug information to a removable storage device:

```
configure debug core-dumps memorycard
```

The following example enables the switch to save debug information to the internal memory card:

```
configure debug core-dumps internal-memory
```

History

This command was first available in ExtremeXOS 11.1.



The internal-memory parameter was added in ExtremeXOS 11.2.

Support for USB 2.0 storage devices was added in ExtremeXOS 12.5.3.

Platform Availability

This command is available on all platforms.

configure elrp-client disable ports

```
configure elrp-client disable-ports [exclude | include] [ ports | eaps-ring-ports ]
```

Description

Creates an ELRP exclude port list.

Syntax Description

exclude	Specifies that selected ports are to be excluded from ELRP disabling.
include	Specifies that selected ports are to be included in ELRP disabling.
<i>ports</i>	Specifies one or more ports to be excluded or included.
eaps-ring-ports	Specifies whether EAPS ring ports are to be excluded or included.

Default

All ports, together with EAPS ring ports, are included by default; that is, they will be disabled if a loop is detected on that port.

Usage Guidelines

Use this command to specify ports or EAPS ring ports that are to be part of an ELRP exclude port list. Use the `exclude` option to add ports to the exclude port list. Use the `include` option to remove them from the list.

When ELRP detects a loop and has been configured to automatically disable the port where a looped ELRP PDU is received and an exclude port list has been configured, it will check to determine if that port is on the exclude port list. If that port is on the list, ELRP will not disable it; if it is not on the list, it will be disabled.

Any port on the switch can be added to or removed from the list. EAPS ring ports can be part of the list.

To display the ports that are include in the exclude port list, use the `show elrp disabled-ports` command.

To remove the exclude port list, use the `unconfigure elrp-client disable ports` command.



Example

The following example adds port 2:1 to an ELRP exclude port list:

```
configure elrp-client disable-ports exclude 2:1,2:3
```

History

This command was first available in ExtremeXOS 12.5.3.

Platform Availability

This command is available on all platforms.

configure elrp-client one-shot

```
configure elrp-client one-shot vlan_name ports [ports | all] interval sec retry
count [log | print | print-and-log]
```

Description

Starts one-time, non-periodic ELRP packet transmission on the specified ports of the VLAN using the specified count and interval.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>ports</i>	Specifies the set of VLAN ports for packet transmission.
all	Specifies all ports of this VLAN for packet transmission.
<i>sec</i>	Specifies the interval (in seconds) between consecutive packet transmissions. The range is 1 to 600 seconds. The default is 1 second.
<i>count</i>	Specifies the number of times ELRP packets must be transmitted. The range is 1 to 255 times. The default is 3 times.
log	Specifies that a message should be logged in the system log file when ELRP packets are received back indicating detection of network loop, or no packets are received within the specified duration.
print	Specifies that a message should be printed to the console when ELRP packets are received back indicating detection of network loop, or no packets are received within the specified duration.
print-and-log	Specifies that a message should be logged in the system log file and printed to the console when ELRP packets are received back indicating detection of network loop, or no packets are received within the specified duration.



Default

`sec`—The interval between consecutive packet transmissions is 1 second.

`count`—The number of times ELRP packets must be transmitted is 3.

Usage Guidelines

This command starts one-time, non-periodic ELRP packet transmission on the specified ports of the VLAN using the specified count and interval. If any of these transmitted packets is returned, indicating loopback detection, the ELRP client can perform a configured action such as logging a message in the system log file or printing a log message to the console. There is no need to send a trap to the SNMP manager for non-periodic requests.

Note



This command is backward compatible with Extreme Networks switches running ExtremeWare. If your network contains switches running only ExtremeXOS, you can also use the `run elrp vlan_name {portsports} {intervalsec} {retrycount}` to perform one-time ELRP packet transmission.

Use the `configure elrp-client periodic` command to configure periodic transmission of ELRP packets.

The ELRP client must be enabled globally in order for it to work on any VLANs. Use the `enable elrp-client` command to globally enable the ELRP client.

The ELRP client can be disabled globally so that none of the ELRP VLAN configurations take effect. Use the `disable elrp-client` command to globally disable the ELRP client.

Example

The following example starts one-time, non-periodic ELRP packet transmission on all ports of the VLAN sales, uses the default interval and transmission times, and sends messages to the console:

```
configure elrp-client one-shot sales ports all interval 1 retry 3 print
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

configure elrp-client periodic

```
configure elrp-client periodic vlan_name ports [ports | all] interval sec [log |
log-and-trap | trap] {disable-port {{duration seconds} | permanent}
```



Description

Starts periodic ELRP packet transmission on the specified ports of the VLAN using the specified interval.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>ports</i>	Specifies the set of VLAN ports for packet transmission.
all	Specifies all ports of this VLAN for packet transmission.
<i>sec</i>	Specifies the interval (in seconds) between consecutive packet transmissions. The range is 1 to 600 seconds. The default is 1 second.
log	Specifies that a message should be logged in the system log file when ELRP packets are received back indicating detection of network loop, or no packets are received within the specified duration.
log-and-trap	Specifies that a message should be logged in the system log file and a trap message should be sent to the SNMP manager when ELRP packets are received back indicating detection of network loop, or no packets are received within the specified duration.
trap	Specifies that a trap message should be sent to the SNMP manager when ELRP packets are received back indicating detection of network loop, or no packets are received within the specified duration.
disable-port	Specifies that the port should be disabled where the looped PDU is received.
duration	Specifies a hold time that the port is kept disabled before re-enabling.
<i>seconds</i>	The number of seconds the port is kept disabled.
permanent	Specifies that the port is disabled permanently. User intervention is required to enable.

Default

The default interval between consecutive packet transmissions is 1 second.

If a duration in seconds is not specified, the default is permanent.

Usage Guidelines

This command starts periodic ELRP packet transmission on the specified ports of the VLAN using the specified interval. If any of these transmitted packets is returned, indicating loopback detection, the ELRP client performs a configured action of logging a message in the system log file and/or sending a trap to the SNMP manager.

Beginning with ExtremeXOS 12.4, you have the option to automatically disable the port where the looped packet arrives and to specify the time interval for which the port remains disabled. When that specified time expires, the port is automatically enabled.

Should a loop occur on multiple ports, only the first port in the VLAN on which the PDU is received is disabled. The second port is ignored for 1 or 2 seconds and then if another PDU is received, that port is disabled until the loop is gone. This prevents shutting down all ports in the VLAN.



Use either the `configure elrp-client one-shot` or the `run elrp` command to configure non-periodic, one-time transmission of ELRP packets.

The ELRP client must be enabled globally in order for it to work on any VLANs. Use the `enable elrp-client` command to globally enable the ELRP client.

The ELRP client can be disabled globally so that none of the ELRP VLAN configurations take effect. Use the `disable elrp-client` command to globally disable the ELRP client.

The ExtremeXOS software does not support ELRP and Network Login on the same port.

Use the `show elrp` command to check the ELRP status and the `show elrp disabled-ports` command to view details of ELRP disabled ports.

Example

The following example starts periodic ELRP packet transmission on slot 3, port 2 of VLAN marketing, sends packet transmissions every 2 seconds, sends messages to the log, and should a loop be detected, disables the port for 5 seconds:

```
configure elrp-client periodic marketing ports 3:2 interval 2 log disable-
port duration 5
```

History

This command was first available in ExtremeXOS 11.1.

The disable port feature was added in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

configure forwarding fabric hash

```
configure forwarding fabric hash [default | source-port | packet {algorithm [crc
|xor] | dynamic-mode [spray | eligibility | none}}] {slot slot-number}
```

Description

A 40Gb port must use multiple 20Gb HGd links to carry its traffic. In the degenerate case, a particular HGd link is always the target of the hash result. With a 40Gb link hashing into a 20Gb link, half the traffic is lost. Extreme will attempt to configure the packet hash algorithm to accommodate most common situations; however, some customers may have uncommon situations and may wish to adjust the hash. For this reason, the user will be able to specify values for some of the hash parameters in normal packet hash mode:

- Source port hash (cannot be used when 40Gb ports are in use).
- Packet field hash calculation algorithm (CRC or XOR).



The packet hash is configurable. A Dynamic Load Balance (DLB) algorithm allows distribution of flows to the links in the switch fabric trunk that are currently carrying the least load. DLB uses the packet field hash, but does not use the hash code that is reduced to a number modulo the number of links in the group. Instead the number used is 15 bits for a total of 32K possible hash codes. This 15-bit hash code indexes into a 32K-entry “flow table”. Each time an unused entry is allocated to one or more “micro-flows” (i.e., to flows that generate the same 15-bit hash value), a load calculation is performed and the link with the most available bandwidth is assigned to the flow table entry. DLB offers two modes:

- Spray mode causes the link assignment to occur on every packet transmission. This is similar to a “round-robin” hash. Every packet goes to the link with the least load but ordering within flows is not guaranteed. This mode is useful for Bandwidth Management Testing (BMT).
- Eligibility mode keeps the link fixed to the flow entry until a 32ms inactivity timeout occurs. Ordering within flows is guaranteed.

For blades that provide 10Gb ports only (and for 40Gb blades that are entirely configured for 10Gb operation), source port hashing will be used by default. While source port hashing can be used for non-blocking operation, such operation depends on the switch fabric distribution. . For example, suppose three 10Gb ports on one BDXA-10G48X card hash to the same switch fabric “channel”. Also suppose that the three 10Gb ports’ traffic is aggregated into the same 40Gb port on a different I/O blade. Since each switch fabric channel can only provide 20Gb of bandwidth to a single 40Gb port, then the switch fabric channel will try to send 30Gb/s of Ethernet packets to the same 20Gb switch fabric link, causing congestion in the channel on the FM blade.

To avoid this situation, the user can configure DLB Eligibility mode which will cause per 10Gb Ethernet port traffic to be distributed to all switch fabric channels instead of being sent to the same channel. Use the following command:

```
configure forwarding fabric hash [default |source-port |
packet {algorithm [crc |xor] |
dynamic-mode [spray | eligibility | none]]] {slot slot-number}
```

Syntax Description

default	Resets the hashing algorithm to the default value on the specified slot or slots.
source-port	Indicates that the system should distribute traffic based on the ingress port (the system controls the mapping between the ingress port and the switch fabric ingress port).
packet	Indicates that packet fields are to be used to calculate a hash value that will be used to select the switch fabric port that will carry the packet. Use algorithm to select from crc or xor. Each of these is a 16-bit result. If algorithm is not specified, packet hashing occurs and is set to a default value.
dynamic-mode	Dynamic Load Balancing (DLB) allows flow distribution to links within the switch fabric trunk based on the current load on the individual trunk links. DLB is configurable in one of three modes: spray, eligibility, and none (direct packet hash without considering the load).
<i>slot-number</i>	Selects the I/O slot to which the command is applied.



Default

The default option resets the hashing algorithm to the default value on the specified slot or slots.

Usage Guidelines



Note

Modify the hash algorithm only with the guidance of Extreme Networks technical personnel.

For the packet hash, the following fields are used depending on packet type:

- IPv4 packets – IP source, IP destination, VLAN ID, source and destination L4 ports, and L3 protocol ID.
- IPv6 packets – Collapsed IP source and destination addresses, VLAN ID, source and destination L4 ports, and L3 protocol ID. (Collapsed source and destination addresses are the 32-bit result of the exclusive OR of the bits IP[127-96], IP[95-64], IP[63-32], and IP[31-0].)
- L2 packets – MAC source, MAC destination, Ethertype, and VLAN ID.
- L2 MPLS – Payload MAC source, payload MAC destination, payload Ethertype, and payload outer VLAN ID.
- L3 MPLS – Payload source and destination IPv4 or collapsed IPv6 addresses, tunnel VLAN ID, source and destination L4 ports, and L3 protocol field.
- Non-terminated MPLS (label swapped or L2 switched on outer-most header) – RPID, labels 1, 2, and 3, and payload source and destination IPv4 or collapsed IPv6 addresses.
- MAC-in-MAC with tunnel termination – L2 payload MAC DA, L2 payload MAC SA, L2 payload Ethertype, and L2 payload outer VLAN ID.
- MAC-in-MAC without termination – RPID, tunnel MAC DA, tunnel MAC SA, and ISID.
- FiberChannel over Ethernet – RPID, source ID, destination ID, originator exchange ID, responder exchange ID, fabric ID, and VLAN ID.
- TRILL – Egress RBridge nickname, Ingress RBridge nickname, payload MAC destination (low 32-bits), payload MAC source (low 32-bits), payload Ethertype, and payload VLAN ID.

Example

Example output not yet available and will be provided in a future release.

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command variant is available only on the BlackDiamond X8 switch.

configure forwarding hash-algorithm

```
configure forwarding hash-algorithm [crc16 | crc32] {dual-hash [on | off]}
```



Description

Modifies hardware table utilization by configuring the hash algorithm or dual-hash settings.

Syntax Description

crc16	Specifies the CRC16 hash algorithm.
crc32	Specifies the CRC32 hash algorithm. This is the default setting.
on	Specifies that the dual-hash feature be turned on for the L3 Hash Table for hardware with dual-hash capability. With dual-hash on, each hash bucket is divided into two half-buckets with independent hash algorithms. One half-bucket uses the configured hash algorithm (CRC16 or CRC32), and the other half-bucket uses an alternate hash algorithm. This is the default setting. The “dual-hash” function applies only to the BlackDiamond 8800 series switches.
off	Specifies that the dual-hash feature be turned off, even for hardware with dual-hash capability. The “dual-hash” function applies only to the BlackDiamond 8800 series switches.

Default

In ExtremeXOS 11.5, the default hash algorithm is `crc32`.

In ExtremeXOS 11.4 and earlier, the default hash algorithm is `crc16`.

The dual-hash default is “on.”

Usage Guidelines



Note

Modify the hardware table hash algorithm only with the guidance of Extreme Networks technical personnel.

The switch uses a hash algorithm to decide where to store the addresses in the hardware table. The standard, default hash algorithm works well for most systems; however, for some addresses with certain patterns, the hardware may attempt to store address information in the same section of the hardware.

If you are running ExtremeXOS 11.4 or earlier and experience a full hardware table that affects Layer 2, IP local host, and IP multicast forwarding, you see messages similar to the following in the log:

```
<Info:HAL.IPv4Adj.Info> : adj 136.159.188.109: IP add error is Table full
for new or newly resolved ARP, egress valid <Info:HAL.IPv4Adj.Info> : adj
136.159.188.109: returned -17 for L3 table bucket 181
<Warn:HAL.IPv4Mc.Warning> : Could not allocate a hardware S,G,V entry
(889f4648,effffffa,70) - hardware table resource exceeded (rv=-17).
```

If you are running ExtremeXOS 11.5 or later and experience a full hardware table that affects Layer 2, IP local host, and IP multicast forwarding, you see messages similar to the following in the log:

```
<HAL.IPv4Adj.L3TblFull> MSM-A: IPv4 unicast entry not added. Hardware L3
Table full.
```



(ExtremeXOS 12.1 and later have the Extended IPv4 Host Cache feature and do not display this HAL.IPv4Adj.L3TblFull message on the MSM for a full hash table condition.)

```
<Card.IPv4Adj.Warning> Slot 4: IPv4 unicast entry not added. Hardware L3
Table full.
<HAL.IPv4Mc.GrpTblFullEnt> MSM-A: IPv4 multicast entry
(10.0.0.1,224.1.1.1,vlan 1) not added. Hardware Group Table full.
<Card.IPv4Mc.Warning> Slot-4: IPv4 multicast entry not added. Hardware L3
Table full.
```

In the previously described situations, you can configure a different hash algorithm to select a different section of the hardware to store addresses. You must save your configuration and reboot the switch to modify the hash algorithm used by the hardware table. Typically, the dual-hash feature improves hash utilization. You must save your configuration and reboot the switch to turn dual-hash on or off.

Upgrading to ExtremeXOS 11.5

When you upgrade to ExtremeXOS 11.5, the hash algorithm automatically becomes `crc32`. For example, if you saved a configuration using an image from ExtremeXOS 11.4 or earlier with the hash algorithm set to `crc16`, when ExtremeXOS 11.5 loads, the hash algorithm becomes `crc32`. To change the hash algorithm to `crc16`, use the `configure forwarding hash-algorithm crc16` and save your switch configuration.

Example

The following command modifies the hardware table hash algorithm to `crc16`:

```
configure forwarding hash-algorithm crc16
```

The switch displays the following message to describe the change and to prompt you to save your configuration and reboot the switch:

```
Configured hash alorithm has been changed to 'crc16' with L3 dual-hash
support 'on' for applicable HW.
Warning: This command will only take effect after a save and reboot
```

The following command disables dual-hashing on BlackDiamond 8000 c- and xl-series I/O modules.

```
configure forwarding hash-algorithm crc32 dual-hash off
```

The switch displays the following message:

```
Configured hash algorithm has been changed to 'crc32' with L3 dual-hash
support 'off' for applicable HW.
Warning: This command will only take effect after a save and reboot.
```

To display the results, use the `show forwarding configuration` command.



History

This command was first available in ExtremeXOS 11.3.2.

The default hash algorithm was changed to crc32 in ExtremeXOS 11.5.

Platform Availability

This command is available only on the BlackDiamond X8, BlackDiamond 8800 series switches, SummitStack, and the Summit family switches.

configure forwarding hash-recursion-level

configure forwarding hash-recursion-level 0-3

Description

Modifies hardware table utilization by configuring the dual hashing recursion level.

Syntax Description

0-3	Sets the maximum number of L3 hash buckets to modify to make room for a new entry.
-----	--

Default

The default is "1."

Usage Guidelines

This command allows you to select the dual hashing "recursion level" for hardware with the dual-hash feature. The setting applies only if dual-hash is configured or defaulted to "on" using the configure forwarding hash-algorithm command.

The configured recursion level is the maximum number of existing hash entries to move in an attempt to add a new hash entry. A higher recursion level may provide better hash utilization at the expense of additional CPU processing. This command does not require a system reboot. However, the new recursion level takes effect only for addresses added after the command is issued.

Example

The following command modifies the dual-hash recursion level to modify up to two L3 hash buckets in an attempt to add a new entry:

```
configure forwarding hash-recursion-level 2
```



History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the BlackDiamond X8 and BlackDiamond 8800 series switches.

configure tech-support add collector

```
configure tech-support add collector [hostname | ip_address] tcp-port port {vr
vr_name} {from source_ip_address} {ssl [on | off]}
```

Description

This command adds collectors that the switch attempts to connect to for the purpose of forwarding status reports. The collector is identified by its hostname or IP address.

This command also configures the initial value of the TCP port that the collector is listening to, the VR name and source IP address that the switch uses to attempt to connect to the collector, and the SSL mode whether the switch needs to turn SSL on or off when it connects to the collector.

Syntax Description

<i>hostname</i>	Host name of the collector.
<i>ip_address</i>	IPv4 address of the collector.
tcp-port	TCP port number that the collector is listening.
<i>port</i>	Port number. The range is 1-65535.
vr <i>vr_name</i>	Specifies the Virtual router and virtual router name. The default name is VR-Mgmt.
from <i>source_ip_address</i>	Specifies the source and the source IPv4 address. The default source is the IP address on VLAN Mgmt.
ssl	Specifies the Secure Sockets Layer. The default is on if the ssh.xmod is installed, otherwise it is off.
on	Specifies that SSL is on. This is the default setting if the ssh.xmod is installed.
off	Specifies that SSL is off, or that the ssh.xmod is not installed.

Default

Disabled.

Usage Guidelines

This command adds collectors that the switch attempts to connect to for the purpose of forwarding status reports. The collector is identified by its hostname or IP address. Each added collector needs to have a unique hostname or IP address. If the specified hostname or IP address has already existed, an



error message *'ERROR: The collector 1.1.1.1 already exists'* is displayed. Other commands use hostname or IP address to specify the collector that the command reconfigures, deletes, runs reports for, or shows configuration and status.

This command also configures the initial value of the TCP port that the collector is listening to, the VR name and source IP address that the switch uses to connect to the collector, and the SSL mode that determines if the switch needs to turn SSL on/off when connecting to the collector. The SSL feature is provided by the `ssh.xmod`. If SSL is configured on, and `ssh.xmod` is not installed, an error message displays:

```
ERROR: SSL module is not installed
```

When the `techSupport.xmod` is started, a default collector is added with the IP address set to 65.222.234.14. Additionally, if the `ssh.xmod` is installed, the TCP port set to 443, otherwise it is set to 80. the VR set to VR-Mgmt, the Source IP Address set to the primary IP Address on VLAN Mgmt, and SSL set to on if `ssh.xmod` is installed or set to Off otherwise. The purpose of having a default collector configured is to minimize the configuration required for a customer to enable `techSupport`.

Example

The following command :

```
disable tech-support collector
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

configure tech-support collector

```
configure tech-support collector [hostname | ip_address] tcp-port port {vr  
vr_name} {from source_ip_address} {ssl [on | off]}
```

Description

This command reconfigures the TCP port, the VR, the Source IP Address, and SSL mode of an existing collector.

Syntax Description

<i>hostname</i>	Host name of the collector.
<i>ip_address</i>	IPv4 address of the collector.



tcp-port	TCP port number that the collector is listening.
port	Port number. The range is 1-65535.
vr vr_name	Specifies the Virtual router and virtual router name. The default name is VR-Mgmt.
from source_ip_addresses	Specifies the source and the source IPv4 address. The default source is the IP address on VLAN Mgmt.
ssl	Specifies the Secure Sockets Layer. The default is on if the ssh.xmod is installed, otherwise it is off.
on	Specifies that SSL is on. This is the default setting if the ssh.xmod is installed.
off	Specifies that SSL is off, or that the ssh.xmod is not installed.

Default

Disabled.

Usage Guidelines

This command reconfigures the TCP port, the VR, the Source IP Address, and SSL mode of an existing collector. The collector to be reconfigured is specified by its hostname or IP address. If the specified collector does not exist, an error message **ERROR: The collector 1.1.1.1 does not exists** is displayed.

Example

The following command reconfigures the tech support collector:

```
configure tech-support collector
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

configure tech-support collector data-set

```
configure tech-support collector [ all hostname | ip_address ] data-set [ summary  
| detail ]
```



Description

This command configures the amount and type of data that is included in the status report for a collector.

Syntax Description

all	Configures report data set for all existing collectors.
<i>hostname</i>	Specifies the host name of the collector.
<i>ip_address</i>	Specifies the IPv4 address of the collector.
data-set	Specifies the report data set. The default is detail if ssh.xmod is installed, otherwise it is summary .
detail	Specifies the output of <code>show tech-support area</code> for area general, config, log, VLAN, and EPM.
summary	Specifies the output of <code>show tech-support all</code> command.

Default

The default is **detail** if ssh.xmod is installed.

Usage Guidelines

This command configures the amount and type of data that is included in the status report for a collector. When you specify **all**, it configures a report data set for all existing collectors; otherwise report data is set for a particular collector specified by the *hostname* or *IP address*. When the data set is set to **summary**, the status report sent by the switch includes installed EXOS and Bootrom image versions, the active partition, serial number, equipment type, installed hardware options, stored SRAM contents, basic switch configuration, and log messages. The output of the summary option is collected from the `show tech area` command for the area general, configuration, log, VLAN, and EPM. Changing the report data set to **detail** will send the full output of the `show tech` command. When a collector is added, the data set is set to **detail** if the ssh.xmod is installed, otherwise, it is set to **summary**.

Example

The following command example configures a specific collector to display a detailed output set:

```
configure tech-support collector 65.222.234.14 data-set detail
```

History

This command was first available in ExtremeXOS 15.4.



Platform Availability

This command is available on all platforms.

configure tech-support collector frequency error-detected

```
configure tech-support collector [ all | hostname | ip_address] frequency [bootup
[on | off] | error-detected [on | off]| daily [on {time hour} | off]]
```

Description

This command configures how often the switch sends status reports for a collector.

Syntax Description

all	Configures report mode for all report collectors.
<i>hostname</i>	Configures report mode based on the host name of the collector.
<i>ip_address</i>	Configures report mode based on the IPv4 address of the collector.
bootup	Send status report when the switch boots up. The default value is on.
on	Specifies that the status reporting is on at bootup.
off	Specifies that the status reporting is off at bootup.
error-detected	Specifies that a status report is sent when a critical severity event is logged. The default value is off.
on	Specifies that error-detected reporting is on.
off	Specifies that error-detected reporting is off. This is the default value.
daily	Specifies that status reports are sent once a day. The default value is off.
on <i>time hour</i>	Specifies the time to send the report. Specifies the hour 0-23. The default value is 0 (12:00AM).
off	Specifies that the daily status reports are off. This is the default value.

Default

Disabled.

Usage Guidelines

This command configures the frequency that the switch sends status reports for a collector. By specifying **all**, it configures report frequency for all existing collectors; otherwise it configures report frequency for a particular collector specified by the hostname or IP address. If the **bootup** option is set to **on**, the switch sends a status report when the switch boots up. If the **error-detected** option is set to **on**, the switch sends a status report when a critical severity event is logged. If the **daily** option is set to **on**, the switch sends a status report once a day regardless of the switch's operational status during the last 24 hour period.



Optionally, you can specify the hour that the report is sent. The default hour is 0, and the valid range is 0 to 23, where 0 is 12:00 AM local time and 23 is 11:00 PM local time. You can enable or disable each option (**bootup**, **error-detected** or **daily**) independently. When all three options of a collector are turned off, the switch does not send any status report to that collector even if the report mode of the collector is set to automatic. When a collector is added, the **bootup** option is set to on, and the **error-detected** and **daily** option is set to off.

Example

The following command example configures the report mode on all existing collectors:

```
configure tech-support collector all report
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

configure tech-support collector report

```
configure tech-support collector [hostname | ip_address] report [ automatic | manual ]
```

Description

This command configures the report mode for a collector.

Syntax Description

all	Configures report mode for all report collectors.
<i>hostname</i>	Configures report mode based on the host name of the collector.
<i>ip_address</i>	Configures report mode based on the IPv4 address of the collector.
automatic	Automatically reports switch status to the configured collector (Default).
manual	Manually reports switch status to the configured collector through the <code>run tech-support report</code> .

Default

Disabled.

If enabled, the automatic collector is the default report setting.



Usage Guidelines

This command configures the report mode for a collector. When you specify **all**, it configures report mode for all existing collectors, otherwise it configures report mode for a particular collector specified by the hostname, or IP address. When the report mode is set to **automatic**, the switch automatically attempts to connect to the cloud-hosted collector, and reports the switch status information based on the frequency and data set setting of the collector. Changing the configuration to **manual** restricts reporting to user initiated mode using the *run tech-support* command for that collector. When a collector is added, the report mode is set to automatic by default.

Example

The following command example configures the report mode on all existing collectors:

```
configure tech-support collector all report
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

configure tech-support delete collector

```
configure tech-support delete collector [ all | hostname | ip_address]
```

Description

This command deletes existing collectors.

Syntax Description

all	Specifies that you delete all report collectors.
<i>hostname</i>	Specifies the host name of the collector you want to delete.
<i>ip_address</i>	Specifies the IPv4 address of the collector you want to delete.

Default

Disabled.



Usage Guidelines

This command deletes existing collectors. If you specify **all**, it deletes all existing collectors; otherwise it deletes the collector specified by the hostname or IP address. If the specified collector does not exist, an error message `ERROR: The collector 1.1.1.1 does not exist` is displayed.

Example

The following example deletes all collectors :

```
configure tech-support delete collector all
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

disable elrp-client

disable elrp-client

Description

Disables the ELRP client (standalone ELRP) globally.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables the ELRP globally so that none of the ELRP VLAN configurations take effect.

The ELRP client must be enabled globally in order for it to work on any VLANs. Use the `enable elrp-client` command to globally enable the ELRP client.



Example

The following command globally disables the ELRP client:

```
disable elrp-client
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

disable led locator

```
disable led locator { slot [slot | all ]}
```

Description

Disables the front panel LEDs from flashing on a switch or chassis.

Syntax Description

slot <i>slot</i>	Slot number.
all	All slots.

Default

N/A.

Usage Guidelines

None.

Example

The following example disables the front panel LEDs on all slots:

```
disable led locator all
```

History

This command was first available in ExtremeXOS 15.3.



Platform Availability

This command is available on all platforms.

disable log debug-mode

disable log debug-mode

Description

Disables debug mode. The switch stops generating debug events.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables debug mode. Debug mode must be enabled prior to configuring advanced debugging capabilities. These include allowing debug messages, which can severely degrade performance. For typical network device monitoring, debug mode should remain disabled, the default setting. Debug mode should only be enabled when advised by technical support, or when advanced diagnosis is required. The debug mode setting is saved to FLASH.

The following configuration options require that debug mode be enabled:

- Including a severity of debug-summary, debug-verbose, or debug-data when configuring filters.
- Target format options process-name, process-id, source-function, and source-line.

Example

The following command disables debug mode:

```
disable log debug-mode
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.



disable tech-support collector

disable tech-support collector

Description

Disables the tech support feature.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables the tech-support feature. In the ExtremeXOS 15.4 release, the feature is disabled by default. When the feature is disabled, the previous scheduled reports are canceled, and the bootup event and critical severity events are ignored.

Example

The following command disables the tech-support feature:

```
disable tech-support collector
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

eject memorycard

eject memorycard

Description

Ensures that the compact flash card or USB 2.0 storage device can be safely removed from the switch.



Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

After the switch writes to a compact flash card or USB 2.0 storage device, and before you can view the contents on the device, you must ensure it is safe to remove the device from the switch. Use the `eject memorycard` command to prepare the device for removal. After you issue the `eject memorycard` command, you can manually remove the device.

If the `configure debug coredumps memorycard` command is in effect when you issue the `eject memorycard` command, the behavior is similar to issuing the `configure debug coredumps off` command.

For more information about removing a compact flash card or USB 2.0 storage device, refer to the hardware documentation.

To access and read the data on the card, use a PC with appropriate hardware such as a compact flash reader/writer and follow the manufacturer's instructions to access the compact flash card and read the data.

Example

The following command prepares a compact flash card or USB 2.0 storage device to be removed from the switch:

```
eject memorycard
```

On the BlackDiamond X8 switch, there are two USB ports, and so the command indicates which port is being ejected.

If only USB slot 1 has a USB storage device:

```
BD-X8.3 # eject memorycard
sync filesystem...
unmount filesystem...
memorycard unmounted from USB-1
```

If both USB slots have a USB storage device:

```
BD-X8.3 # eject memorycard
sync filesystem...
unmount filesystem...
```



```
memorycard unmounted from USB-1  
memorycard auto-mounted on USB-2
```



Note

When both USB slots have a USB storage device, the first one inserted is the one currently visible to the user. If both devices were left in the slots since the last reboot, the system may make either one visible to the user, that is, it is not deterministic.

History

This command was first available in ExtremeXOS 11.1.

Support for USB 2.0 storage devices was added in ExtremeXOS 12.5.3.

Platform Availability

This command is available on BlackDiamond X8 and 8800 series switches and Summit X460, X480, X670, X670V, and X770 switches.

enable elrp-client

```
enable elrp-client
```

Description

Enables the Extreme Loop Recovery Protocol (ELRP) client (standalone ELRP) globally.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

The ELRP client must be enabled globally in order for it to work on any VLANs.

The ELRP client can be disabled globally so that none of the ELRP VLAN configurations take effect. Use the `disable elrp-client` command to globally disable the ELRP client.

The ExtremeXOS software does not support ELRP and Network Login on the same port.



Example

The following command globally enables the ELRP client:

```
enable elrp-client
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

enable led locator

```
enable led locator {timeout [seconds | none]} {pattern [alternating | flash-all | high-to-low | scanner]} {slot [slot | all ]}
```

Description

Configures the front panel LEDs to flash so a switch/chassis can be easily located in a crowded lab/data center.

Syntax Description

timeout	Limit the LED display time to <i>seconds</i> before returning to normal operation.
<i>seconds</i>	The length of time to display the flashing LEDs. The default is 300 seconds. The maximum value is 1 week (604800 seconds).
none	Display LED pattern until disabled.
pattern	Configures the LED display pattern.
alternating	Groups of LEDs are lit in alternating patterns (Default).
flash-all	All LEDs flash on and off.
high-to-low	LED's are lit in descending port order.
scanner	A group of 4 LED's is lit back and forth.
slot <i>slot</i>	Slot number.
all	All slots.

Default

The default **timeout** length is 300 seconds.

The default pattern is alternating.



Usage Guidelines

Use this command to enable the front panel LEDs to flash so that a switch/chassis can be easily located in a crowded lab, or data center.

Example

The following example enables the front panel LEDs to flash in an alternating pattern for one hour on all slots:

```
enable led locator timeout 3600 pattern alternating all
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms.

enable log debug-mode

enable log debug-mode

Description

Enables debug mode. The switch generates debug events.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables debug mode. Debug mode must be enabled prior to configuring advanced debugging capabilities. These include allowing debug messages, which can severely degrade performance. For typical network device monitoring, debug mode should remain disabled, the default setting. Debug mode should only be enabled when advised by technical support, or when advanced diagnosis is required. The debug mode setting is saved to FLASH.

The following configuration options require that debug mode be enabled:

- Including a severity of debug-summary, debug-verbose, or debug-data when configuring filters.
- Target format options process-name, process-id, source-function, and source-line.



Example

The following command enables debug mode:

```
enable log debug-mode
```

When you enable debug mode, the following message appears:

```
WARNING: Debug mode should only be enabled when advised by technical support, or when advanced diagnosis is required. Performance degradation is possible. Debug mode now enabled.
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

enable tech-support collector

```
enable tech-support collector
```

Description

Enables the tech support feature.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command turns on the tech-support feature. In the ExtremeXOS 15.4 release, the feature is disabled by default. When the feature is disabled, the previous scheduled reports are canceled, and the bootup event and critical severity events are ignored.

When the feature is enabled, if any configured collector has the report mode set to automatic, the switch automatically attempts to send switch status reports to those collectors based on the configuration setting for each individual collector.



You can always use the `run tech-support report` command to trigger a one-time report to a particular collector, or all collectors, regardless if the feature is enabled or disabled or if the collector's report mode is set to automatic or manual.

Example

The following command enables the tech-support feature:

```
enable tech-support collector
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

nslookup

```
nslookup {IPv4 | IPv6} hostname
```

Description

Displays the IP address of the requested host.

Syntax Description

IPv4	Lookup only IPv4 address(es).
IPv6	Lookup only IPv6 address(es).
<i>hostname</i>	Specifies the hostname.

Default

Lookup both IPv4 and IPv6 addresses.

Usage Guidelines

For nslookup to work, you must configure the DNS client, and the switch must be able to reach the DNS server.

By default, the command looks for both IPv4 and IPv6 addresses and reports an error only when neither an IPv4 address nor an IPv6 address is found for the host.



If the IPv4 or IPv6 option is specified, DNS lookup happens only for that address type, and an error is reported when no address of that type is found.

Host Name and Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z).
- Numerals (0-9).
- Period (.).
- Dash (-) Permitted only for host names.
- Underscore (_) Permitted only for host names.
- Colon (:).

When naming or configuring an IP address for your network server, remember the requirements listed above.

Example

The following command looks up the IP addresses of a computer with the name myhost.mydomain that has 2 IPv4 addresses and 1 IPv6 address:

```
nslookup myhost.mydomain
```

The following is sample output from the command on a Summit switch:

```
Host "myhost.mydomain" has the IPv4 address 192.168.1.1
Host "myhost.mydomain" has the IPv4 address 192.168.1.2
Host "myhost.mydomain" has the IPv6 address 2000::1
```

History

This command was first available in ExtremeXOS 10.1.

Support for using an IP address to obtain the name of the host was added in ExtremeXOS 11.0. Support for looking up IPv6 addresses was added in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

run diagnostics

```
run diagnostics [extended | normal | stack-port] {slot [slot | A | B]}
```



Description

Runs normal or extended diagnostics on the switch, slot, node, or management module. On the Summit family switches, this command also runs diagnostics on the stacking ports.

This command is not supported in stacking mode. But if you issue the show diagnostics command from the master node, it will show the diagnostic results for all the nodes.

Syntax Description

extended	Runs an extended diagnostic routine. Takes the ports offline, and performs extensive ASIC and packet loopback tests on all of the ports. BlackDiamond 8800 series switches only—If you have a Power over Ethernet (PoE) module installed, the switch also performs an extended PoE test, which tests the functionality of the inline power adapter.
normal	Runs a normal diagnostic routine. Takes the ports offline, and performs a simple ASIC and packet loopback test on all of the ports.
stack-port	Runs the diagnostic routine on the stack ports. NOTE: This parameter is available only on Summit family switches that are not operating in stacking mode.
<i>slot</i>	Specifies the slot number of an I/O module. BlackDiamond X8 switches only - The slot argument is used to refer to both I/O and fabric modules. NOTE: This parameter is available only on modular switches.
A B	Specifies which MSM/MM to run diagnostics on. A specifies the MSM/MM installed in slot A.B specifies the MSM/MM installed in slot B. NOTE: This parameter is available only on modular switches.

Default

N/A.

Usage Guidelines

Depending on your platform, use this command to run diagnostics on the switch, slot, management module, or stack port.

Running Diagnostics--BlackDiamond X8 Switches Only

If you run the diagnostic routine on an I/O or Fabric module, that module is taken offline while the diagnostic test is performed. The module does not forward traffic. Once the diagnostic test is completed, the I/O or Fabric module is automatically reset and becomes operational again. On a management module, the module is taken offline while the diagnostics test is performed. Once the diagnostic test is completed, the management module (MM) reboots, and becomes operational again.

After the switch runs the diagnostic routine, test results are saved in the module's EEPROM and messages are logged to the syslog. On an I/O or Fabric or MM module, the extended diagnostic routine can require significantly more time to complete.



Running Diagnostics--BlackDiamond X8 Series Switches and BlackDiamond 8800 Series Switches Only

If you run the diagnostic routine on an I/O module, that module is taken offline while the diagnostic test is performed. The ports on the module do not forward traffic. Once the diagnostic test is completed, the I/O module is automatically reset and becomes operational again.

After the switch runs the diagnostic routine, test results are saved in the module's EEPROM and messages are logged to the syslog.

On an I/O module, the extended diagnostic routine can require significantly more time to complete, depending on the number of ports on the module.

On a management module, the module is taken offline while the diagnostics test is performed. Once the diagnostic test is completed, the MSM reboots, and becomes operational again.

Note



BlackDiamond 8810 switch—If you run diagnostics on slots 5 and 6 with an MSM installed in those slots, the diagnostic routine tests the I/O subsystem of the MSM. BlackDiamond 8806 switch—if you run diagnostics on slots 3 and 4 with an MSM installed in those slots, the diagnostic routine tests the I/O subsystem of the MSM. BlackDiamond 8800 series switches—To run diagnostics on the management portion of the master MSM, specify slot A or B.

Running Diagnostics--Summit Family Switches Only

If you run the diagnostic routine on the switch, it reboots and then performs the diagnostic test. During the test, traffic to and from the ports on the switch is temporarily unavailable. When the diagnostic test is complete, the switch reboots and becomes operational again.

To run the diagnostic routine on the stack ports, you need a dedicated stacking cable that connects stack port 1 to stack port 2, which are located at the rear of the switch. The stacking cable is available from Extreme Networks. The switch performs a hardware test to confirm that the stack ports are operational; traffic to and from the ports on the switch is temporarily unavailable. This Bit Error Rate Test (BERT) provides an analysis of the number of bits transmitted in error.

After the switch runs the diagnostic routine, test results are saved to the switch's EEPROM and messages are logged to the syslog.

To run diagnostics on a Summit switch that is in a SummitStack, first disable stacking on that switch, then restart the switch. Once restarted, log into the switch via its console port, and run diagnostics. The switch will perform the diagnostic tests and then restart. Once restarted, log into the switch via its console port and enable stacking, then reboot the switch. Once restarted, the switch will rejoin the stack.

Viewing Diagnostics

To view results of the last diagnostics test run, use the following command:



```
show diagnostics {[cr] | slot [slot | A | B]}
```

**Note**

The slot, A, and B parameters are available only on modular switches.

If the results indicate that the diagnostic failed on a module, replace the module with another module of the same type.

If the results indicate that the diagnostic failed on the switch, contact Extreme Networks Technical Support.

BlackDiamond X8 switch example.

The following commands runs normal diagnostics on the I/O module installed in slot 2 :

```
run diagnostics normal slot 2
```

The switch displays a warning similar to the following about the impact of this test. You also have the opportunity to continue or cancel the test:

```
Running Diagnostics will disrupt network traffic.  
Are you sure you want to continue? (y/n)
```

Enter y to continue and run the diagnostics. Enter n to cancel the operation.

The following commands runs normal diagnostics on the first Fabric module :

```
run diagnostics extended slot fm-1
```

The switch displays a warning similar to the following about the impact of this test. You also have the opportunity to continue or cancel the test:

```
Running Diagnostics will disrupt network traffic.  
Are you sure you want to continue? (y/n)
```

Enter y to continue and run the diagnostics. Enter n to cancel the operation.

The following commands runs normal diagnostics on the MM-A :

```
run diagnostics extended slot a
```



The switch displays a warning similar to the following about the impact of this test. You also have the opportunity to continue or cancel the test:

```
Running Diagnostics will disrupt network traffic.
Are you sure you want to continue? (y/n)
```

Enter y to continue and run the diagnostics. Enter n to cancel the operation.

Sample Console Output while running Diagnostics on BlackDiamond X8 MM.

```
Initializing operational diagnostics...
Version 1.8 Image Release built by project
BDX Management Module detected in slot A
Board Rev      : 0x0
FPGA Build     : 0x14
FPGA Minor Rev : 0x 1
Running Diags in normal mode (0x 1)
PCI unit 0: Dev 0xb842, Rev 0x01, Chip BCM56842_A0, Driver BCM56840_B0
SOC unit 0 attached to PCI device BCM56842_A0
Running Power On Self Test...(Normal)
Test temperature start      - PASS
Test sdram                  - PASS
Test boot flash             - PASS
Test internal cf card       - PASS
Test usb                    - PASS
Test south bridge           - PASS
Test mgmt loopback          - PASS
Test nvram                  - PASS
Test dump flash             - PASS
Test voltage check          - PASS
Test fpga reg               - PASS
Test fpga mem               - PASS
Test port loopback          - PASS
Test port snake             - PASS
Test asic0 mem              - PASS
Test asic0 reg              - PASS
Test temperature report     - PASS
Diagnostics passes!
Current Time: Tue Jan 24 15:42:59 2012
```

BlackDiamond 8800 series switch example

The following command runs normal diagnostics on the I/O module installed in slot 2 of a BlackDiamond 8800 series switch:

```
run diagnostics normal slot 2
```



The switch displays a warning similar to the following about the impact of this test. You also have the opportunity to continue or cancel the test:

```
Running Diagnostics will disrupt network traffic.
Are you sure you want to continue? (y/n)
```

Enter y to continue and run the diagnostics. Enter n to cancel the operation.

Summit family switch example

The following command runs normal diagnostics on the Summit family switch:

```
run diagnostics normal
```

The switch displays a warning similar to the following about the impact of this test. You also have the opportunity to continue or cancel the test:

```
Running Diagnostics will disrupt network traffic.
Are you sure you want to continue? (y/n)
```

Enter y to continue and run the diagnostics. Enter n to cancel the operation.

The following command runs diagnostics on the stack ports on a Summit family switch:

```
run diagnostics stack-port
```

If you issue this command with a console connection, the switch displays the following information. You also have the opportunity to continue or cancel the test:

```
Summit Diagnostics Mode Enabled, Starting Diagnostics...
Found X440a-24T in Motherboard
Motherboard CPLD Revision: 2
Starting stacking port diagnostics
*****
*                                                                 *
* Please connect a cable between Stack Port 1 and Stack Port 2. *
*                                                                 *
*       Press S to skip test, ENTER key to continue.             *
*                                                                 *
*****
```

Press [Enter] to continue and run the diagnostics. Enter S to cancel the operation.

If you continue with diagnostics, the switch displays messages similar to the following:

```
Stack Port 1 and Stack Port 2
BERT .....
.....
```



```

.....
.....
Stacking ports
Port 1 (Device 0 - Device port 26)
Lane 0 PASSED.
Lane 1 PASSED.
Lane 2 PASSED.
Lane 3 PASSED.
Port 2 (Device 0 - Device port 27)
Lane 0 PASSED.
Lane 1 PASSED.
Lane 2 PASSED.
Lane 3 PASSED.
DIAGNOSTIC PASS: run test bert stacking
Summit Diagnostics completed, rebooting system...

```

If you issue this command with a Telnet connection, the switch displays a warning similar to the following about the impact of this test. You also have the opportunity to continue or cancel the test:

```

Running Diagnostics will disrupt network traffic.
Are you sure you want to continue? (y/n)

```

Enter y to continue and run the diagnostics. Enter n to cancel the operation.

History

This command was first available in ExtremeXOS 10.1.

This command was modified in ExtremeXOS 11.0 to run diagnostics on management modules from the command line interface (CLI).

The stack-port parameter for the Summit family switches was added in ExtremeXOS 11.5.

Platform Availability

This command is available on all platforms.

run elrp

```

run elrp vlan_name {ports ports} {interval sec} {retry count}

```

Description

Starts one-time, non-periodic ELRP packet transmission on the specified ports of the VLAN using the specified count and interval.



Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>ports</i>	Specifies the set of VLAN ports for packet transmission.
<i>sec</i>	Specifies the interval (in seconds) between consecutive packet transmissions. The range is 1 to 64 seconds. The default is 1 second.
<i>count</i>	Specifies the number of times ELRP packets must be transmitted. The range is 3 to 255 times. The default is 10 times.

Default

sec—The interval between consecutive packet transmissions is 1 second.

count—The number of time ELRP packets must be transmitted is 10.

Usage Guidelines

This command starts one-time, non-periodic ELRP packet transmission on the specified ports of the VLAN using the specified count and interval. If any of these transmitted packets is returned, indicating loopback detection, the ELRP client prints a log message to the console. There is no need to send a trap to the SNMP manager for non-periodic requests.

Note



This command is compatible with Extreme Networks switches running only the ExtremeXOS software. If your network contains switches running ExtremeXOS and switches running ExtremeWare, use the `configure elrp-client one-shot vlan_name ports [ports | all] interval sec retry count [log | print | print-and-log]` command to perform one-time ELRP packet transmission.

If you do not specify the optional interval or retry parameters, the default values are used.

Use the `configure elrp-client periodic` command to configure periodic transmission of ELRP packets.

The ELRP client must be enabled globally in order for it to work on any VLANs. Use the `enable elrp-client` command to globally enable the ELRP client.

The ELRP client can be disabled globally so that none of the ELRP VLAN configurations take effect. Use the `disable elrp-client` command to globally disable the ELRP client.

Example

The following command starts one-time, non-periodic ELRP packet transmission on the VLAN green using the default interval and packet transmission:

```
run elrp green
```



History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

run esvt traffic-test

```
run esvt traffic-test {vlan} vlan_name loopback-port loopback-port peer-switch-ip
ipaddress packet-size packet_size rate rate [Kbps|Mbps |Gbps] duration time
[seconds | minutes | hours]
```

Description

Runs the Extreme Service Verification tool on the specified ports of the VLAN using the specified count and interval.

Syntax Description

vlan	Specifies the VLAN for the traffic test.
<i>vlan_name</i>	Specifies a VLAN name.
<i>loopback-port</i>	Specifies the loopback port for the traffic test.
<i>loopback-port</i>	Specifies the <i>loopback-port</i> .
peer-switch-ip	Specifies the peer switch for the traffic test.
<i>ipaddress</i>	Specifies the IP address of the peer switch.
packet-size	Specifies the size of packet to use for service verification.
<i>packet_size</i>	Specifies the size; the range is 64 - 9216.
<i>rate</i>	Specifies the maximum service utilization rate and rate value.
Kbps	Specifies K bits per second.
Mbps	Specifies M bits per second.
Gbps	Specifies G bits per second.
duration	Specifies the duration of time to run service verification.
<i>time</i>	Specifies the duration value; the range is 1-65534.
seconds	Specifies the time in seconds.
minutes	Specifies the time in minutes..
hours	Specifies the time in hours.

Default

N/A.



Usage Guidelines

Use this command to run the Extreme Service Verification tool on the specified ports of the VLAN using the specified count and interval.

Example

The following example starts the Extreme Service Verification tool on vlan1, loopback port 2:

```
run esvt traffic-test v1 loopback-port 2 peer-switch-ip 2.2.2.2 packet-size
64 rate 200 Mbps duration 5 minutes
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available in X460, X480, E4G-200, E4G-400, X670, X770, Stacking, BDx8 and BD8800.

run tech-support report

```
run tech-support report {now | in hours | cancel} {collector [all | hostname |
ip_address]}
```

Description

This command instructs the switch to generate a report and upload it to a collector.

Syntax Description

now	Specifies that you run a report immediately. This is the default setting.
in	Specifies that you run report in a specified number of hours.
<i>hours</i>	Specifies the hours from now to run report. The range is 1-168 hours (one week).
cancel	Cancels the scheduled report.
collector	Specifies the report collector. The default value is all collectors.
all	Specifies all report collectors.
<i>hostname</i>	Specifies the host name of the collector.
<i>ip_address</i>	Specifies the IPv4 address of the collector.

Default

The default time for running reports is **now**.



The default for number of collectors is **all**.

Usage Guidelines

This command instructs the switch to generate a report and upload it to a collector. The default operation is to perform this operation immediately for all existing collectors. Optionally, you can configure a one-time trigger to perform the operation in "hours from now." The valid range is one to 168 hours (one week). If you specify the hostname or IP address, the switch runs a report for that particular collector.

Only a single one-time report per collector can be scheduled at any time. When run `tech-support report in hours` is issued before the previous scheduled one-time report completes, the previous report is cancelled, and a new one-time report is scheduled.

This command also provides a way to cancel a scheduled report for a particular collector.

Example

The following command example configures a specific collector to display a detailed output set:

```
* X460-24t.9 # run tech-support report
Connecting to 10.5.2.107:800 with SSL disabled...
Collector connected successfully.
Generating summary report for
10.5.2.107:9998.....
Report generated successfully.
Sending report to 10.5.2.107:800...
Report sent successfully

Connecting to 10.5.2.107:800 with SSL enabled...
Error: Failed to connect to the collector - Socket time out *
```

X460-24t.9 # run tech-support report in 1
Run tech-support report is scheduled on Thu Feb 21 05:06:32 2013 for the collector 10.5.2.108:800.

Run tech-support report is scheduled on Thu Feb 21 05:06:32 2013 for the collector 10.5.2.107:9998.

To cancel a scheduled report, use `'run tech-support report cancel'` command.

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.



save debug tracefiles memorycard

save debug tracefiles memorycard

Description

Copies debug information to a compact flash card or USB 2.0 storage device.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines



Note

Use this command only under the guidance of Extreme Networks Technical Support to troubleshoot the switch.

Use this command to copy debug information to an installed removable storage device. The debug information includes log files and trace files.

Progress messages are displayed that indicate the file being copied and when the copying is finished.

Beginning with ExtremeXOS 11.6, you can use the `upload debug [hostname | ipaddress] {{vr}vrname}` command to copy debug information to a network TFTP server.

Example

The following command copies debug information to a removable storage device:

```
save debug tracefiles memorycard
```

History

This command was first available in ExtremeXOS 11.0.

The syntax for this command was modified in ExtremeXOS 11.1 from `upload debug-info memorycard` to `save debug tracefiles memorycard`.

Support for USB 2.0 storage devices was added in ExtremeXOS 12.5.3.



Platform Availability

This command is available on BlackDiamond 8800 series switches and Summit X460, X480, X670, X670V, and X770 switches.

show debug

show debug

Description

This command displays whether the writing of core dump files is enabled or disabled and whether the files are written to internal memory, a compact flash card, or a USB 2.0 storage device.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

By default, the switch writes core dump files to the internal memory card. To change the default configuration, use the `configure debug core-dumps` command.

Example

The following example shows that the switch is configured to send core dump files to the internal memory card:

```
Switch.2 # show debug
Debug Settings:
  Core dumps: Enabled (internal-memory)
```

The following example shows that the sending core dump files is disabled:

```
Switch.99 # show debug
Debug Settings:
  Core dumps: Disabled
```

The following example shows that the switch is configured to send core dump files to a compact flash card or USB 2.0 storage device:

```
Switch.76 # show debug
```



```

Debug Settings:
  Core dumps: Enabled (memorycard)

```

History

This command was first available in ExtremeXOS 11.1.

Support for the internal memory card was added in ExtremeXOS 11.2.

Support for USB 2.0 storage devices was added in ExtremeXOS 12.5.3.

Platform Availability

This command is available on all platforms.

show diagnostics

```
show diagnostics {[cr] | slot [slot | A | B]}
```

On a stack:

```
show diagnostics {slot [slot_number]}
```

Description

Displays the status of the last diagnostic test run on the switch.

Syntax Description

<i>slot</i>	Specifies which I/O or Fabric module to display diagnostic status information on. NOTE: This parameter is available only on modular switches.
A B	Specifies which MSM/MM to display diagnostic status information on:A specifies the MSM installed in slot A.B specifies the MSM installed in slot B. NOTE: This parameter is available only on modular switches.
<i>slot_number</i>	Specifies the slot number of the node on which you need to run the diagnostics. Values can be from 1 to 8 or FM-1 to FM-4 (BDX8 only). Note: This parameter is available only on a stackable switch in stack mode.

Default

N/A.

Usage Guidelines

Use this command to display information from the last diagnostic test run on the switch.



To run diagnostics on a switch in a stack, you need to remove the node from the stacking mode, restart the node, and then run the diagnostics locally on the node. You can see the latest diagnostics results from the Master node of a stack after the node has rejoined the stack.

Output on BlackDiamond X8 Switches

If you have run diagnostics, a brief summary of the overall diagnostic test is displayed.

The switch displays the following diagnostic information:

- Date the test was run—The month, date, and year.
- Last Test Version—The ExtremeXOS version associated with the results.
- Summary—A brief summary of the overall diagnostic test. Options are:
 - Diagnostics Pass—The diagnostic test has passed.
 - Diagnostics Fail—One or more diagnostic test has failed.
 - Diagnostics Interrupted—The diagnostic test was interrupted on the I/O or Fabric module or MM module.

If you have never run diagnostics on a specific slot, the switch displays a message similar to the following:

```
Slot-6: G8X  
No Diagnostics Data
```

If you attempt to view diagnostics information for a slot that does not have a module installed, the switch displays a message similar to the following:

```
Slot-7: No Module Present
```

Output on the BlackDiamond 8800 Series Switches

The switch displays the following diagnostic information:

If you have run diagnostics, a brief summary of the overall diagnostic test is displayed. Options are:

- Date the test was run—The month, date, and year.
- Last Test Version—The ExtremeXOS version associated with the results.
- Summary—A brief summary of the overall diagnostic test. Options are:
 - Diagnostics Pass—The diagnostic test has passed.
 - Diagnostics Fail—One or more diagnostic test has failed.
 - Diagnostics Interrupted—The diagnostic test was interrupted on the I/O module due to initiating an MSM failover.



If you have never run diagnostics on a specific slot, the switch displays a message similar to the following:

```
Slot-6: G8X
No Diagnostics Data
```

If you attempt to view diagnostics information for a slot that does not have a module installed, the switch displays a message similar to the following:

```
Slot-7: No Module Present
```

Output on the Summit Family Switches

The switch displays the following diagnostic information:

If you have run diagnostics, information includes the:

- Date the test was run—The month, date, and year.
- Last Test Version—The ExtremeXOS version associated with the results.
- Test data—If the diagnostic test failed, the switch displays the name of the failed test. The switch displays a message similar to the following:

```
MAC memory test failed
```

- Summary—A brief summary of the overall diagnostic test. Options are:
 - Diagnostics Pass—The diagnostic test has passed.
 - Diagnostics Fail—One or more diagnostic tests have failed.

If you have never run diagnostics on the switch or stack ports, the switch displays a message similar to the following:

```
Result: FAIL
Test date run is invalid. Please run Diagnostics.
Error reading diagnostics information.
```

This message is normal and expected if you have never run diagnostics on the switch. After you run diagnostics, you should see information about the executed test.

Additional Guidelines Applicable to Modular Switches Only

If you use the `show diagnostics {[cr] | slot [slot | A | B]}` command on a slot where diagnostics have not been run, the switch displays messages similar to the following:

```
No Diagnostics Data
```

or

```
Diagnostics never run
```

If you try to display diagnostic test information on a slot where no module is installed, the switch displays messages similar to the following:



No Module Present

or

No card in slot

Running Diagnostics

To run diagnostics on an I/O module or an MSM installed in the BlackDiamond 8800 series switch, use the following command:

```
run diagnostics [extended | normal] {slot [slot | A | B]}
```

To run diagnostics on Summit family switches, use the following command:

```
run diagnostics [extended | normal | stack-port]
```

Depending on the software version running on your switch or your switch model, additional or different diagnostic information might be displayed. For more information, see the command `run diagnostics`.

Example

The following command displays the results of module diagnostics for the I/O module in slot 2:

The following is sample output from a BlackDiamond X8 switch:

```
BD-X8.5 # show diagnostics slot 2
Slot-2: BDXA-40G24X
Last Test Date: Tue Jan 24 15:09:54 2012
Last Test Version: 1.7
Summary: Diagnostics Pass
```

The following is sample output from a BlackDiamond 8800 series switch:

```
Slot-2: G24X
Last Test Date: Wed May 6 23:57:17 2009
Last Test Version: 12.3.0.1
Summary: Diagnostics Pass
```

When the version is unknown, Last Test Version reads "Unknown."

The following is sample output from the BlackDiamond 12804 switch:

```
Slotnum 2
DiagResult timestamp Fri Dec 2 11:15:29 2005
Test Temperature = 41 degrees C
Diag Test Version=1.0.1.2. S/N=05364-00058
-----
CPU System | Passed
-----
Register Test | Passed
```



```
-----  
Memory Test | Passed  
-----
```

```
System Test | Passed  
-----
```

```
Summary: Diagnostics Pass
```

The following command displays the results of the switch diagnostics for the Summit family switch:

```
show diagnostics
```

The following is sample output from this command:

```
Last Test Date: May-04-2006  
Last Test Version: 12.3.0.1  
Summary: Diagnostics Pass
```

When the version is unknown, Last Test Version reads “Unknown.”

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms.

show elrp

```
show elrp
```

Description

Displays ELRP information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the following:



- State of ELRP (enabled/disabled).
- Total number of ELRP sessions.
- ELRP packets transmitted.
- ELRP packets received.

In addition to the summary information at the top of the display, the `show elrp` command also displays the following information:

Client	Displays the name of the ELRP client.
VLAN	Displays the name of the VLAN with ELRP enabled.
Ports	Displays the set of VLAN ports used for packet transmission.
Interval	Displays the configured interval. An interval of 3 indicates that ELRP PDUs are transmitted every 3 seconds.
Count	Lists the configured number of ELRP PDUs that are transmitted. The PDUs are transmitted at the configured interval. This method of ELRP PDU transmission is used by ESRP in the pre-master state. A count of 0 indicates continuous PDU transmission. If the Cyclic value is Yes, the count is always 0.
Cyclic	Indicates whether ELRP PDUs are being continuously sent. The column shows Yes for the master VLAN because that VLAN is continuously sending ELRP PDUs for loop detection. When a VLAN is in the pre-master state, it only sends three ELRP PDUs before changing to master or slave. During this time the column shows No for that VLAN.
Pkts-Xmit	Displays the number of ELRP PDUs transmitted.
Pkts-Rcvd	Displays the number of ELRP PDUs received.
Action	Displays the configured action the switch takes when ELRP messages are received back indicating a detection of a network loop or no packets are received within the specified duration. The following list describes the actions: Print (P)—Specifies that the switch prints a message to the console. Log (L)—Specifies that the switch sends a message to the system log file. Trap (T)—Specifies that the switch sends a message to the SNMP manager. Callback (C)—Specifies a callback action. If you use ELRP with another protocol (for example ESRP), ELRP uses a callback action to notify the protocol of a loop detection.
Disable Port	Displays the configured hold time (number of seconds or permanent) for a port that was disabled with the <code>configure elrp-client periodic</code> command. When the time in seconds expires, the port is automatically enabled.

Example

The following command displays summary ELRP status information on the switch:

```
show elrp
```

The following sample output is displayed:

```
ELRP Standalone Client:      Enabled
Number of ELRP sessions:    1
Number of ELRP pkts transmitted: 14
```



```

Number of ELRP pkts received:      0
Pkt      Pkts      Disable
Client  Vlan      Ports  Int.   Count  Cyclic Xmit      Rcvd      Action
Port
-----
----
CLI     vlanp1      All   1      0      Yes   14      0        L      -
CLI     vlanp2      All   1      0      Yes   14      0        T
Perm
CLI     vlanp3      All   1      0      Yes   14      0        LT     15
-----
----
Action : (P) Print , (L) Log , (T) Trap , (C) Callback

```

History

This command was first available in ExtremeXOS 11.1.

The disable port feature was added in ExtremeXOS 12.4.

Platform Availability

This command is available on all platforms.

show elrp disabled-ports

```
show elrp disabled-ports
```

Description

Displays information about ELRP disabled ports.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the results of disabling ports using the `configure elrp-client periodic` command.

This command displays the following:

- Excluded Ports—User defined ports that will not be disabled.
- Exclude EAPS ring ports--Whether EAPS ring ports can be excluded.
- Disabled Port—The port that ELRP disabled.



- Detected VLAN—The VLAN with looping ELRP PDU(s).
- Duration—The configured time to keep the port disabled.
- Time Disabled—The time when ELRP disabled the port.

Example

The following command displays summary ELRP status information on the switch:

```
show elrp disabled-ports
```

The following sample output is displayed:

```
Exclude EAPS ring ports: Yes
Excluded Ports
-----
1:1 1:9 1:14 1:18
-----
Disabled Detected          Duration Time
Port      Vlan              (sec)   Disabled
-----
1:20     vlan1              15      Fri Jul 16 17:53:16 2010
-----
```

History

This command was first available in ExtremeXOS 12.4.

The excluded port list was added in ExtremeXOS 12.5.3.

Platform Availability

This command is available on all platforms.

show esvt traffic-test

```
show esvt traffic-test {{vlan} vlan_name}
```

Description

Displays statistics output for the Extreme Service Verification tool on the specified VLAN.

Syntax Description

vlan	\Specifies the VLAN for the traffic test.
<i>vlan_name</i>	Specifies a VLAN name.



Default

N/A.

Usage Guidelines

Use this command to display test output for the Extreme Service Verification tool on the specified VLAN.

Example

The following example displays show output for the Extreme Service Verification tool on vlan1 and vlan3:

```

VLAN  Status      Loopback  Duration  Tx rate  Tx Frame  Tx Total  Rx Total
      Port        (hh:mm:ss) (Mbps)   Size  Frames  Frames
=====
V1  Running      2    00:09:00   50   64  2000   2000
V1  Stopped     4    00:20:00  150   128  3000   3000
V3  Completed   3    00:60:00  250   256  4000   4000
V1234567> Error      7    00:10:00  100   9100   0
0
=====

```

> Indicates vlan name string truncated past 8 characters

X480-48x(10G4X).1 # show esvt traffic-test detail

```

VLAN Name          : V1
Port               : 9
IP Address         : 2.2.2.2
Status            : Running
Duration (hh:mm:ss) : 00:10:00
Time Remaining(hh:mm:ss) : 00:09:00
Loopback Port     : 2
Tx rate (Kpbs)    : 200000
Tx Frame Size (bytes) : 64
Tx Frames         : 0
Rx Frames         : 0
Frames Lost       : 0

```

```

VLAN Name          : V1
Port               : 3
IP Address         : 6.6.6.6
Status            : Stopped
Duration (hh:mm:ss) : 00:20:00
Loopback Port     : 4
Tx rate (Kpbs)    : 150
Tx Frame Size (bytes) : 128
Tx Frames         : 1000
Rx Frames         : 1000
Frames Lost       : 0

```

```

VLAN Name          : V3
Port               : 5
IP Address         : 7.7.7.7

```



```

Status : Completed
Duration (hh:mm:ss) : 00:10:00
Loopback Port : 3
Tx rate (Kpbs) : 250
Tx Frame Size (bytes) : 256
Tx Frames : 4000
Rx Frames : 4000
Frames Lost : 0

VLAN Name : V12345678901234567890123456789
Port : invalid port
IP Address : 8.8.8.8
Status : Error
Duration (hh:mm:ss) : 00:02:00
Loopback Port : 7
Tx rate (Kpbs) : 100
Tx Frame Size (bytes) : 9100
Error : Failed to communicate with peer
switch

```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available in X460, X480, E4G-200, E4G-400, X670, Stacking, BDx8 and BD8800.

show forwarding configuration

show forwarding configuration

Description

Displays the configured selection criteria for ECMP routes and load-sharing group ports and the hardware table settings, including the configured and current hash algorithm and dual-hash settings.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The output of this command displays the following information:



- Configured hash algorithm—The hash algorithm configured on the switch. After the configuration is saved and the switch is rebooted, the switch uses this hash algorithm.
- Current hash algorithm—The hash algorithm currently used by the switch.
- Configured dual-hash setting—Whether the dual-hash feature is configured 'on' or 'off' on the switch. After the configuration is saved and the switch is rebooted, the switch uses this setting.
- Current dual-hash setting—Whether the dual-hash feature is currently 'on' or 'off' on the switch.
- Dual-Hash recursion level—The current dual-hash recursion level; default is '1.'
- Sharing criteria—Current selection criterion used for ECMP route sharing as well as for load-sharing groups. Specifies which Layer 3 and Layer4 information is used in the sharing hash algorithm. For more information, see the description for the `configure forwarding sharing [L3 | L3_L4]` command.
- Group Table Compression—Whether the group table compression is currently 'on' or 'off' on the switch.
- Switching mode—Whether the switching mode is currently set to 'cut-through' or 'store-and-forward.'
- Fabric flow control—Whether flow control fabric configuration is set to 'auto' or 'off.'

It is possible for the values of the configured and the current hash, or the configured and current dual-hash settings to be different. For example, if you modified the hash algorithm and have not saved the configuration and rebooted the switch, the values might be different. In this situation, the switch also displays the following message:

NOTE: A save and reboot are required before the configured hash will take effect

Example

The following command displays the hardware forwarding algorithm configured on the switch:

```
show forwarding configuration
```

The following is sample output from this command on a BlackDiamond 8800 series switch:

```
BD-8810.1 # show forwarding configuration
L2 and L3 Forwarding table hash algorithm:
Configured hash algorithm:          crc32
Current hash algorithm:             crc32
L3 Dual-Hash configuration: (Applies only to "c" and "xl"-series HW)
Configured setting:                 on
Current setting:                    on
Dual-Hash Recursion Level:         1
Hash criteria for IP unicast traffic for L2 load sharing and ECMP route
sharing
Sharing criteria:                   L3_L4
IP multicast:
Group Table Compression:            on
IP multicast:
  Group Table Compression:          on
  Lookup-Key                        SourceIP, GroupIP, VlanId
Switch Settings:
```




```

Switching mode:                store-and-forward
Fabric Flow Control:
Fabric Flow Control:          auto

```

In addition to the data shown above, on BlackDiamond X8 series switches the following lines appear at the end of the output:

```

Fabric Hash Slot-1:
Packet Algorithm:
Dynamic mode
Eligibility
or
Fabric Hash Slot-2:
Packet Algorithm:
crcl6-ccitt
seed: 0x12345678
or
Fabric Hash Slot-5:
Source-port

```

For BlackDiamond X8 series switches, when the switch fabric hash is set to Default, there is no printed output for the related slot.

Dual-hash information appears only on BlackDiamond X8 series switches and BlackDiamond 8800 series switches.

Fabric flow control information appears only on BlackDiamond X8 series switches, BlackDiamond 8800 series switches, SummitStack, and Summit X460, X670, and X770 switches.

History

This command was first available in ExtremeXOS 11.3.2.

The flow control feature was added in ExtremeXOS 12.5.

The Forwarding Lookup-Key feature show output is added in ExtremeXOS 15.3.

Platform Availability

This command is available only on BlackDiamond X8 series switches, BlackDiamond 8800 series switches, SummitStack, and Summit family switches.

show tech support

```
show tech support {all | area} {detail} {logto [file]}
```

Description

Displays the output of various show commands to assist in monitoring and troubleshooting the switch; use only in conjunction with Extreme Networks Technical support.



Syntax Description

all	Indicates all available show command output to be displayed.
area	Specifies one tech support area. For example, if you want to view STP information, enter stp.
detail	Specifies more detailed information.
logto [file]	Instructs the switch to log the show tech output into a file located in the switch's internal memory. The default file name is show_tech.log.tgz.

Default

N/A.

Usage Guidelines



Note

Use this command only under the guidance of Extreme Networks Technical Support personnel to view your switch configurations and to troubleshoot the switch.

The `show tech` command displays the output of the following commands, among others:

- `ls internal-memory.`
- `show bootprelay.`
- `show configuration.`
- `show dhcp-client state.`
- `show diagnostics.`
- `show management.`
- `show memory.`
- `show odometers.`
- `show policy.`
- `show port rxerror.`
- `show port txerror.`
- `show power.`
- `show power budget.`
- `show power controller.`
- `show process.`
- `show radius.`
- `show session.`
- `show switch.`
- `show tacacs.`
- `show version.`
- `show vlan.`

Information about the following areas is also displayed, among others:



- `aaa`.
- `bootp`.
- `cli`.
- `stp`.

If you enter the detail keyword, the following show output is displayed, among others:

- `show log`.
- `show log configuration`.
- `show log counters all`.
- `show process detail`.

This information can be useful for your technical support representative if you experience a problem.

Depending on the software version running on your switch, the configurations running on your switch, and the type of switch you have, additional or different show command and configuration output may be displayed.

Note

When the `show tech` command is executed from the Backup MSM, the following commands will not generate proper output:



- `show debug system-dump msm/slot`
- `debug hal show version msm Master MSM a or b`
- `show iparp distributed-mode statistics`

Example

The following command displays the show command output on the switch:

```
show tech
```

History

This command was first available in ExtremeXOS 10.1.

The command name and command syntax was modified in ExtremeXOS Release 15.4 from `show tech` to `show tech support`.

Platform Availability

This command is available on all platforms.

show tech-support collector

```
show tech-support collector { hostname | ip_address }
```



Description

This command displays tech-support configuration and report status.

Syntax Description

<i>hostname</i>	Specifies the host name of the collector.
<i>ip_address</i>	Specifies the IPv4 address of the collector.

Default

NA.

Usage Guidelines

This commands displays tech-support configuration and report status, such as last report stats, when the next cyclic report is sent if report frequency daily is configured, and when the next one time report is sent if one time report is scheduled with `run tech-support in hours` command. If you do not specify a hostname or IP address, this command displays configuration and report status for all existing collectors; otherwise it displays configuration and report status for the specified collector.

Example

The following is an example of the show output:

```
Tech Support Collector:      Enabled

Report Collector:          10.5.2.107
  TCP Port:                9998
  Virtual Router:          VR-Mgmt
  Source IP Addr:          10.66.24.234
  SSL:                     Off
  Report Mode:             Automatic
  Report Data Set:         Summary
  Report Frequency:        Error-detected AND Daily at 0:00
  Last Report:            Successful
    Sent On:               Thu Feb 21 05:06:32 2013
    From:                  10.66.24.234:3546
    To:                    10.5.2.107:9998
    SSL:                   Off
    Error:                 None
  Next Cyclic Report:      Scheduled on Fri Feb 22 00:00:00 2013
  Next One Time Report:    Not Scheduled

Report Collector:          10.5.2.108
  TCP Port:                800
  Virtual Router:          VR-Mgmt
  Source IP Addr:          10.66.24.234
  SSL:                     On
  Report Mode:             Automatic
  Report Data Set:         Detail
  Report Frequency:        Bootup AND Daily at 0:00
```



```

Last Report:           Failed
Sent On:              Thu Feb 21 05:06:32 2013
From:                 10.66.24.234:3598
To:                   10.5.2.107:800
SSL:                  On
Error:                Socket time out
Next Cyclic Report:   Scheduled on Fri Feb 22 00:00:00 2013
Next One Time Report: Not Scheduled

```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms.

stop esvt traffic-test

```
stop esvt traffic-test {{vlan} vlan_name}
```

Description

Stops the Extreme Service Verification tool on the specified VLAN.

Syntax Description

vlan	Specifies the VLAN for the traffic test.
<i>vlan_name</i>	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

Use this command to stop the Extreme Service Verification tool on the specified VLAN.

Example

The following example stops the Extreme Service Verification tool on vlan1:

```
stop esvt traffic-test v1
```

History

This command was first available in ExtremeXOS 15.3.



Platform Availability

This command is available in X460, X480, E4G-200, E4G-400, X670, X770, Stacking, BDx8 and BD8800.

top

top

Description

Displays real-time CPU utilization information by process.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to show the percentage of CPU processing devoted to each process, sampled every 5 seconds.

You can change the display by typing a character while the display is active. The following table displays the supported commands.

Table 59: TOP Interactive Command Display Options

Key	Action
P	Sort process list by CPU utilization
T	Sort process list by time usage
N	Sort process list by number (process ID)
M	Sort process list by memory usage
q [Ctrl] + c	Exit the top program

For more detailed information about the top command including display options, command fields, and command usage, please refer to your UNIX documentation.

Example

The following command displays the real-time CPU utilization information by process:

```
top
```



History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms.

unconfigure elrp-client

```
unconfigure elrp-client vlan_name
```

Description

Disables a pending one-shot or periodic ELRP request for the specified VLAN.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
------------------	------------------------

Default

N/A.

Usage Guidelines

This command disables a pending one-shot or periodic ELRP request for the specified VLAN.

To start one-time, non-periodic ELRP packet transmission on specified ports of a VLAN using a particular count and interval, use one of the following commands:

- `configure elrp-client one-shot vlan_name ports [ports | all] intervalsec retrycount [log | print | print-and-log]`—(This command is backward compatible with Extreme Networks switches running the ExtremeWare software.)
- `run elrp vlan_name {portsports} {intervalsec} {retrycount}`

To configure periodic transmission of ELRP packets, use the `configure elrp-client periodic` command.

The ELRP client must be enabled globally in order for it to work on any VLANs. Use the `enable elrp-client` command to globally enable the ELRP client.

The ELRP client can be disabled globally so that none of the ELRP VLAN configurations take effect. Use the `disable elrp-client` command to globally disable the ELRP client.



Example

The following command disables a pending ELRP request on VLAN elrp1:

```
unconfigure elrp-client elrp1
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms.

unconfigure elrp-client disable ports

```
unconfigure elrp-client disable-ports
```

Description

Deletes an ELRP exclude port list.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to remove an ELRP exclude port list.

Example

The following example removes the existing ELRP exclude port list:

```
unconfigure elrp-client disable-port
```

History

This command was first available in ExtremeXOS 12.5.3.



Platform Availability

This command is available on all platforms.

upload debug

```
upload debug [hostname | ipaddress] [{vr} vr_name]
```

Description

Uploads debug information files to a tftp server. On a platform that has both primary and backup MSMs/MMs, debug information files are uploaded from both the backup and primary MSMs/MMs.

Syntax Description

<i>hostname</i>	Specifies the host name of the TFTP server to which the debug files will be uploaded to.
<i>ipaddress</i>	Specifies the IP address of the TFTP server to which the debug files will be uploaded to.
<i>vr_name</i>	Specifies the name of the virtual router.

Default

By default, the virtual router VR-Mgmt will be used.

Usage Guidelines



Note

Use this command only under the guidance of Extreme Networks Technical Support personnel to troubleshoot the switch.

Use this command to copy, upload debug information (for example, core, trace, show tech, configuration, and policy files) to the specified TFTP server.

Progress messages are displayed that indicate the file being copied and when the copying is finished. Depending on your platform, the switch displays a message similar to the following:

```
The following files on have been uploaded: Tarball Name:
TechPubsLab_C_09271428.tgz ./primary.cfg
```

You can also use this command in conjunction with the `show tech` command. Prior to uploading debug information files, the switch prompts you with the following message to run the `show tech` command with the logto file option:

```
Do you want to run show tech logto file first? (y/n)
```

Enter y to run the `show tech` command before uploading debug information. If you enter y, the `show_tech.log.tgz` file is included during the upload. Enter n to upload debug information without running the `show tech` command.



After you upload the debug information, you should see a compressed TAR file, which contains the debug information.

The TAR file naming convention is:

```
<SysName>_<{<slot#>|A|B}I|C>_<Current Time>.tgz  
- Current Time = mmddhhmm ( month(01-12), date(01-31), hour(0-24),  
minute(00-59) ).
```

Example

The following command uploads debug files to a network TFTP server:

```
upload debug 10.10.10.10
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms.



Software Licensing

Extreme Networks software may contain software from third party sources that must be licensed under the specific license terms applicable to such software. Applicable copyright information is provided below.

Copyright (c) 1995-1998 by Cisco Systems, Inc.

Permission to use, copy, modify, and distribute this software for any purpose and without fee is hereby granted, provided that this copyright and permission notice appear on all copies of the software and supporting documentation, the name of Cisco Systems, Inc. not be used in advertising or publicity pertaining to distribution of the program without specific prior notice be given in supporting documentation that modification, permission, and copying and distribution is by permission of Cisco Systems, Inc.

Cisco Systems, Inc. makes no representations about the suitability of this software for any purpose. THIS SOFTWARE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

MD5C.C - RSA Data Security, Inc., MD5 Message-Digest Algorithm

Copyright (C) 1991-2, RSA Data Security, Inc. Created 1991. All rights reserved.

License to copy and use this software is granted provided that it is identified as the "RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing this software or this function.

License is also granted to make and use derivative works provided that such works are identified as "derived from the RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing the derived work.

RSA Data Security, Inc. makes no representations concerning either the merchantability of this software or the suitability of this software for any particular purpose. It is provided "as is" without express or implied warranty of any kind.

These notices must be retained in any copies of any part of this documentation and/or software.

\$Id: md5c.c,v 1.2.4880.1 2005/06/24 01:47:07 lindak Exp \$ This code is the same as the code published by RSA Inc. It has been edited for clarity and style only.

Index

A

ACLs

- refreshing 1390
- smart refresh 1390

ARP

- and VLAN aggregation 2369

C

CLI

- ! prompt 934
- named components 66
- conventions, guide
- notice icons 58
- text 58

D

- debug erps 1912
- documentation, using 57

E

EAPS

- names 66
- enable | disable ip-fix 558

I

I/O module

- power management 132

inherit ports 1942, 1975, 1985

IPv6

- displaying 1128

ITU channels

- TDWDM XFP 485

M

modular switch

- port number 67

module recovery

- actions 976
- clearing the shutdown state 933
- displaying 977

MSTP

- identifiers 1974
- inherit ports 1942, 1975, 1985

N

names

- character types 66, 703, 704, 1114, 1367, 1470, 1859, 2045, 2171-2177, 2495, 2497, 2754, 2865, 3135
- conventions 66, 703, 704, 1114, 1367, 1470, 1859, 2045, 2171-2177, 2495, 2497, 2754, 2865, 3135

- maximum length of 66, 703, 704, 1114, 1367, 1470, 1859, 2045, 2171-2177, 2495, 2497, 2754, 2865, 3135
- VLAN, STP, EAPS 66

P

PoE

- devices 884

PoE features 884

port

- wildcard combinations 67

port-mirroring

- guidelines 449

power management

- consumption 132
- overriding 147
- re-enabling 147

prompt

- shutdown ports 934

R

refresh

- ACLs 1390

related publications 59

S

SCP2 3164

SFTP

- Using SFTP and SCP2 3164

smart refresh, ACLs 1390

software requirements for switches 62

STP

- inherit ports 1942, 1975, 1985
- names 66

switch

- software requirement 62

system recovery

- displaying 1018

T

TCP MD5 3126

troubleshooting

- AppleTalk 1108

shutdown state

- modular switches 933

tunable DWDM XFP

- ITU channels 485

V

virtual routers

- commands 1383

VLANs



- names 66
 - protocol-based 1108
- vMANs
 - names 66

W

- wildcard combinations, port 67
- WRED (weighted random early detection) 1450

