Commands: t through z

COMMAND DESCRIPTION

Copyright

Disclaimer

The contents of this document are subject to revision without notice due to continued progress in methodology, design and manufacturing. Ericsson shall have no liability for any error or damage of any kind resulting from the use of this document.

Trademark List

SmartEdge	is a registered trademark of Telefonaktiebolaget LM Ericsson.
NetOp	is a trademark of Telefonaktiebolaget LM Ericsson.

Contents

1 Command Descriptions

Commands starting with “t” through commands starting with “z” are included.

This document applies to both the Ericsson SmartEdge® and SM family routers. However, the software that applies to the SM family of systems is a subset of the SmartEdge OS; some of the functionality described in this document may not apply to SM family routers.

For information specific to the SM family chassis, including line cards, refer to the SM family chassis documentation.

For specific information about the differences between the SmartEdge and SM family routers, refer to the Technical Product Description *SM Family of Systems* (part number 5/221 02-CRA 119 1170/1) in the `Product Overview` folder of this Customer Product Information library.

1.1 table-map

```
table-map map-name
```

```
no table-map map-name
```

1.1.1 Purpose

Assigns a traffic index to routes installed for a Border Gateway Protocol (BGP) address family.

1.1.2 Command Mode

BGP address family configuration

1.1.3 Syntax Description

map-name	Name of the route map.

1.1.4 Default

A table map is not applied to a BGP address family.

1.1.5 Usage Guidelines

Use the **`table-map`** command to assign a traffic index to routes installed for a BGP address family.

Traffic index counters are maintained on interfaces with traffic index accounting enabled. Traffic indices are associated with BGP routes based on route-maps matching on BGP attributes. When IP packets are received on an interface with traffic index accounting enabled, and the route lookup for the packet's destination IP address corresponds to a BGP route with a traffic index assigned, the corresponding byte and packet counters are incremented. For more information, see the **`set traffic-index`** and **`traffic-index accounting`** commands.

Use the **`route-map`** command in BGP neighbor address family configuration mode and BGP peer group address family configuration mode to determine the attribute modifications and filtering conditions of the applied route map.

Use the **`no`** form of this command to remove the table map.

1.1.6 Examples

The following example assigns a traffic index to routes installed for a BGP address family using the `bgp-accounting` route map:

```
[local]Redback(config-ctx)#router bgp 64001

[local]Redback(config-bgp)#address-family ipv4 unicast

[local]Redback(config-bgp-af)#table-map bgp-accounting
```

1.2 tacacs+ deadtime

`tacacs+ deadtime` ***`interval`***

{**`no`** | **`default`**} **`tacacs+ deadtime`**

1.2.1 Purpose

Modifies the interval during which the SmartEdge router is to treat a nonresponsive Terminal Access Controller Access Control System Plus (TACACS+) server as "dead," and instead, try to reach another server if one is configured.

1.2.2 Command Mode

Context configuration

1.2.3 Syntax Description

`interval`	Deadtime interval in minutes. The range of values is 0 to 65,535; the default value is 5.

1.2.4 Default

The SmartEdge router waits 5 minutes after a timeout occurs before considering the affected server to be eligible to accept TACACS+ requests again.

1.2.5 Usage Guidelines

Use the **`tacacs+ deadtime`** command to modify the interval during which the SmartEdge router is to treat a nonresponsive TACACS+ server as “dead,” and try, instead, to reach another configured server.

If a server fails to respond to a TACACS+ request within the configured TACACS+ timeout window, which configured with the **`tacacs+ timeout`** command (in context configuration mode), it is declared dead. No TACACS+ requests are sent to a dead server until the server deadtime (the value of the *`interval`* argument) expires, at which time the server is again considered eligible for new TACACS+ requests and resumes its original priority. However, if all servers are unresponsive, the SmartEdge router uses local authentication, if enabled. If local authentication is disabled and all servers are unresponsive, authentication fails. Use the **`aaa authentication administrator`** command in context configuration mode to enable local authentication; for more information, see *Configuring Bridging*

Use the **`no`** form of this command or specify a value of 0 for the *`interval`* argument to disable the deadtime feature, which means that the server is always eligible for TACACS+ requests.

Use the **`default`** form of this command to reset the number of retransmission attempts to 5 minutes.

1.2.6 Examples

The following example specifies a deadtime interval of `10` minutes:

```
[local]Redback(config-ctx)#tacacs+ deadtime 10
```

1.3 tacacs+ max-retries

tacacs+ max-retries ***retries***

{**no** | **default**} **tacacs+ max-retries**

1.3.1 Purpose

Modifies the number of retransmission attempts the SmartEdge router will make to open a Transmission Control Protocol (TCP) connection to the Terminal Access Controller Access Control System Plus (TACACS+) server in the event that no response is received from the server within the timeout period.

1.3.2 Command Mode

Context configuration

1.3.3 Syntax Description

retries	Number of retransmission attempts. The range of values is 0 to 255; the default value is 3.

1.3.4 Default

The SmartEdge router makes three attempts to open a TCP connection to the TACACS+ server.

1.3.5 Usage Guidelines

Use the **tacacs+ max-retries** command to modify the number of retransmission attempts the SmartEdge router makes to open a TCP connection to the TACACS+ server in the event that no response is received from the server within a timeout period.

The timeout period is configured through the **tacacs+ timeout** command (in context configuration mode).

If no acknowledgment is received, all configured TACACS+ servers in the context are tried (moving from the last server back to the first, if necessary) until the maximum number of retransmission attempts have been made for each configured server.

Use the **no** form of this command or specify a value of 0 for the ***retries*** argument to disable the retransmission completely.

Use the `default` form of this command to reset the number of retransmission attempts to 3.

1.3.6 Examples

The following example modifies the retry count to allow the SmartEdge router to make up to `5` attempts to open a TCP connection to the TACACS+ server in the event that no response is received from the server within the timeout period:

```
[local]Redback(config-ctx)#tacacs+ max-retries 5
```

1.4 tacacs+ server

`tacacs+ server` {*`ip-addr`* | *`hostname`*} {**`key`** *`key`* | **`encrypted-key`** *`key`*} [**`port`** *`tcp-port`*]

`no tacacs+ server` {*`ip-addr`* | *`hostname`*} **`key`** *`key`* [**`port`** *`tcp-port`*]

1.4.1 Purpose

Configures the IP address or hostname for a Terminal Access Controller Access Control System Plus (TACACS+) server.

1.4.2 Command Mode

Context configuration

1.4.3 Syntax Description

`ip-addr`	IP address of the TACACS+ server.
`hostname`	Hostname of the TACACS+ server.
`key` *`key`*	Alphanumeric string indicating the authentication key that is used when communicating with the TACACS+ server.
`encrypted-key` *`key`*	Alphanumeric string representing the encrypted authentication key that is used when communicating with the TACACS+ server.
`port` *`tcp-port`*	Optional. TACACS+ server Transmission Control Protocol (TCP) port. The range of values is 1 to 65,536. If no port is specified, TCP port number 49 is used as the default.

1.4.4 Default

None

1.4.5 Usage Guidelines

Use the **tacacs+ server** command to configure the IP address or hostname for a TACACS+ server. The SmartEdge router can support up to five TACACS+ servers in each context. The servers are assigned priority based on the order configured. The first configured server is used first. If the first server becomes unavailable or unreachable, the second server is used, and so on.

For the ***hostname*** argument to take effect, Domain Name System (DNS) resolution must be enabled.

The **key*key*** construct allows you to specify the authentication key that the SmartEdge router uses to communicate with the TACACS+ server. After the configuration is saved, this key is stored in the encrypted form. The output of the **show configuration** command displays the **encrypted-key** keyword with the encrypted key indicating that the key is encrypted.

The **encrypted-key*key*** construct allows you to specify and enter an encrypted authentication key, which was previously configured in clear text using the **key*key*** construct. Use the **encrypted-key*key*** construct, if you want the **tacacs+ server** command configured with an encrypted authentication key as part of a configuration file you are loading onto a SmartEdge router. Copy the encrypted key from either the output of the **show configuration** command or a configuration file that contains it, and then paste it as the ***key*** argument within the **encrypted-key *key*** construct of the **tacacs+ server** command configuration of the configuration file you are loading. The SmartEdge router does not encrypt the encrypted key before storing it because the key is already encrypted.

Use the **no** form of this command to delete a previously configured TACACS+ server.

1.4.6 Examples

The following example shows how define a TACACS+ server with an IP address, `10.43.32.56`, and a key, `Secretkey`, for authentication:

```
[local]Redback(config-ctx)#tacacs+ server 10.43.32.56 key Secretkey port 53
```

The following example shows how to define a TACACS+ server with an IP address, `12.33.56.78`, and an encrypted key, `A04915CDD716F0CC3BC20910847B1834`, for authentication:

```
[local]Redback(config-ctx)#tacacs+ server 12.33.56.78 encrypted-key
A04915CDD716F0CC3BC20910847B1834 port 53
```

1.5 tacacs+ strip-domain

tacacs+ strip-domain

{**no** | **default**} **tacacs+ strip-domain**

1.5.1 Purpose

Specifies that the domain portion of a structured username be removed before relaying an authentication, authorization, or accounting request to a Terminal Access Controller Access Control System Plus (TACACS+) server.

1.5.2 Command Mode

Context configuration

1.5.3 Syntax Description

This command has no keywords or arguments.

1.5.4 Default

The SmartEdge router sends entire structured username, including the domain name, to the TACACS+ server.

1.5.5 Usage Guidelines

Use the **tacacs+ strip-domain** command to specify that the domain portion of a structured username be removed before relaying an authentication, authorization, or accounting request to a TACACS+ server. For example, subscriber name joe is sent rather than joe@local. The domain portion can be stripped, even if custom structured username formats have been defined using the **aaa username-format** command (in global configuration mode).

The decision to strip the domain name depends on whether subscriber and administrator records are defined with or without the domain name in the TACACS+ server configuration.

Use the **no** or **default** form of this command to disable the stripping of the domain portion of the structured username.

1.5.6 Examples

The following example prevents the domain portion of the structured username from being sent to the TACACS+ server:

```
[local]Redback(config-ctx)#tacacs+ strip-domain
```

1.6 tacacs+ timeout

tacacs+ timeout ***seconds***

default fault tacacs+ timeout

1.6.1 Purpose

Modifies the maximum amount of time the SmartEdge router waits for a response from a Terminal Access Controller Access Control System Plus (TACACS+) server before assuming that a packet is lost or that the TACACS+ server is unreachable.

1.6.2 Command Mode

Context configuration

1.6.3 Syntax Description

seconds	Timeout period in seconds. The range of values is 1 to 65,535; the default value is 10.

1.6.4 Default

The timeout interval is 10 seconds.

1.6.5 Usage Guidelines

Use the **tacacs+ timeout** command to modify the maximum amount of time that the SmartEdge router waits for a response from a TACACS+ server before assuming that a packet is lost or that the TACACS+ server is unreachable.

The timeout value is displayed in the output of the **show tacacs+ server** command.

Use the **default** form of this command to return the timeout to the default value of 10 seconds.

1.6.6 Examples

The following example sets the TACACS+ timeout to 60 seconds:

```
[local]Redback(config-ctx)#tacacs+ timeout 60
```

1.7 tag

tag ***value***

{**no** | **default**} **tag** ***value***

1.7.1 Purpose

Configures the route tag value for a dynamically verified static routing (DVSR) profile.

1.7.2 Command Mode

DVSR profile configuration

1.7.3 Syntax Description

value	Route tag value. An unsigned 32-bit integer, the range of values is 1 to 4,294,967,295; the default value is 0.

1.7.4 Default

The default route tag value is 0.

1.7.5 Usage Guidelines

Use the **tag** command to configure the route tag value for a DVSR profile. For route redistribution, the route tag can be used for route map matches.

Note: You can also define the route tag value when configuring a DVSR route. In that case, the specified DVSR route tag value overwrites the value in the DVSR profile.

Use the **no** form of this command to delete the route tag value from a DVSR profile.

Use the **default** form of this command to configure the route tag for a DVSR profile to be **0** (the default value).

1.7.6 Examples

The following example defines a DVSR profile using a route tag of `123`; however, it is not used by the DVSR route `10.1.0.0/16`, because it defines its own route tag value of `456`:

```
[local]Redback(config-ctx)#dvsr-profile abc-webfarm
[local]Redback(config-dvsr)#tag 123
[local]Redback(config-dvsr)#exit
[local]Redback(config-ctx)#ip route 10.0.0.0/8 10.10.10.10 dvsr abc-webfarm
[local]Redback(config-ctx)#ip route 10.1.0.0/16 10.10.10.10 dvsr abc-webfarm tag 456
```

1.8 talk

talk ***admin-name***[***@ctx-name***] [***tty-name***]

1.8.1 Purpose

Enables you to establish communications with another administrator during active Telnet or Secure Shell (SSH) sessions on the same SmartEdge router.

1.8.2 Command Mode

Exec

1.8.3 Syntax Description

admin-name	Name of the administrator with whom you want to establish communications.
@ctx-name	Optional. Name of the context in which the administrator account is configured. Required only if the administrator you want to talk to is in a context that is different from the one in which your administrator account is configured.
tty-name	Optional. Name of the teletypewriter (TTY) for a particular administrator session. Use this option to talk to an administrator who is logged on more than once.

1.8.4 Default

Communications with other active administrators are disabled. If the ***ctx-name*** argument is not entered, the system assumes the local context.

1.8.5 Usage Guidelines

Use the **talk** command to establish communications with other active administrators during active Telnet or SSH sessions on the same SmartEdge router. This visual communication program copies lines from one administrator's terminal to that of another administrator.

When communication is established, two administrators can type simultaneously, with their output appearing in separate windows. To exit, press **Ctrl+x+c**. The system restores the terminal to its previous state.

1.8.6 Examples

The following example displays a sample message that indicates the `admin1` administrator is contacting you through the talk program:

```
Message from Redback Talk Daemon@Redback at 5:50 ...

"admin1" wants to talk to you, respond with: talk admin1@local ttyp0
```

1.9 targeted-hello holdtime

targeted-hello holdtime ***seconds***

default fault targeted-hello holdtime

1.9.1 Purpose

Configures the time for which Label Distribution Protocol (LDP) targeted Hello adjacency is maintained in the absence of targeted Hello messages from an LDP neighbor.

1.9.2 Command Mode

LDP router configuration

1.9.3 Syntax Description

seconds	Number of seconds before LDP adjacency is deleted if LDP targeted Hello messages from an LDP neighbor are not received. The range of values is 15 to 3,600.

1.9.4 Default

The default LDP targeted Hello adjacency holdtime is 45 seconds.

1.9.5 Usage Guidelines

Use the **targeted-hello holdtime** command to configure the time for which LDP targeted Hello adjacency is maintained in the absence of targeted Hello messages from an LDP neighbor.

If LDP targeted Hello messages from an LDP neighbor are not received after the specified Hello holdtime, the LDP adjacency is deleted. If this is the last adjacency between the local LDP instance and an LDP neighbor, the LDP session to that LDP neighbor is torn down.

The locally configured targeted Hello holdtime as specified by the **targeted-hello holdtime** command is included in the targeted Hello messages sent to remote LDP neighbors. The negotiated holdtime used to timeout a targeted Hello adjacency is the minimum of the time value specified by the **targeted-hello holdtime** command and the Hello holdtime received in targeted Hello messages from the LDP neighbor of the adjacency.

Use the **hello holdtime** command in LDP router configuration mode to change the locally configured LDP link hello holdtime.

Use the **targeted-hello interval** command in LDP router configuration mode to change the locally configured LDP targeted hello interval.

Use the **default** form of this command to return to the default value of 45 seconds.

1.9.6 Examples

The following example configures a Hello holdtime of `60` seconds:

```
[local]Redback(config-ctx)#router ldp
[local]Redback(config-ldp)#targeted-hello holdtime 60
```

1.10 targeted-hello interval

targeted-hello interval ***seconds***

no targeted-hello interval ***seconds***

default targeted-hello interval ***seconds***

1.10.1 Purpose

Configures the interval between consecutive LDP targeted Hello messages used in extended LDP discovery.

1.10.2 Command Mode

LDP router configuration

1.10.3 Syntax Description

seconds	Number of seconds between consecutive LDP targeted Hello messages. The range of values is 5 to 3,600.

1.10.4 Default

The default LDP targeted Hello interval is 15 seconds.

1.10.5 Usage Guidelines

Use the **targeted-hello interval** command to configure the interval between consecutive LDP targeted Hello messages used in extended LDP discovery.

If the targeted Hello interval is explicitly configured, then the specified value is used to control targeted Hello interval regardless of the targeted Hello holdtime; however, if the targeted Hello interval is not explicitly configured, the targeted Hello interval used is the negotiated LDP targeted Hello holdtime divided by three. The negotiated LDP targeted Hello holdtime is the lesser of the received LDP targeted Hello holdtime and the locally configured LDP targeted Hello holdtime.

Use the **targeted-hello holdtime** command in LDP router configuration mode to change the locally configured LDP targeted Hello holdtime.

Use the **hello holdtime** command in LDP router configuration mode to change the locally configured LDP link Hello holdtime.

Use the **no** form of this command to use the negotiated LDP targeted Hello holdtime divided by three as the targeted-hello interval.

Use the **default** form of this command to return to the default value of 15 seconds.

1.10.6 Examples

The following example configures a targeted Hello interval of `10` seconds:

```
[local]Redback(config-ctx)#router ldp
[local]Redback(config-ldp)#targeted-hello interval 10
```

1.11 tcp keepalive

tcp keepalive [**count***count-num* | **idle***idle-time* | **interval** ***interval-time***]

default tcp keepalive {**count** | **idle** | **interval**}

1.11.1 Purpose

Modifies the Transmission Control Protocol (TCP) keepalive parameters.

1.11.2 Command Mode

Global configuration

1.11.3 Syntax Description

count ***count-num***	Optional. Maximum number of times that the SmartEdge router tries to reestablish a dropped connection. The range of values is 1 to 32; the default value is 8.
idle ***idle-time***	Optional. Maximum amount of time, in half-seconds, that the SmartEdge router allows a TCP connection to remain open. The range of values is 1 to 14,400; the default value is 14,400.
interval ***interval-time***	Optional. Amount of time, in half-seconds, that the SmartEdge router keeps an idle connection open before disconnecting it. The range of values is 1 to 300; the default value is 150.

1.11.4 Default

The values for the TCP keepalive parameters are described in the Syntax Description section.

1.11.5 Usage Guidelines

Use the **tcp keepalive**command to modify the TCP keepalive parameters.

To display the current TCP keepalive settings and TCP status, use the **show tcp** command (in any mode); for details about this command, see the *Command List*.

Use the **default** form of this command to return the TCP keepalive parameters to their default settings.

1.11.6 Examples

The following example shows how to change the count to 4 tries:

```
[local]Redback#configure
[local]Redback(config)#tcp keepalive count 4
```

1.12 tcp path-mtu-discovery

tcp path-mtu-discovery

{**no** | **default**} **tcp path-mtu-discovery**

1.12.1 Purpose

Enables the negotiation of the maximum transmission unit (MTU) for Transmission Control Protocol (TCP) sessions.

1.12.2 Command Mode

Global configuration

1.12.3 Syntax Description

This command has no keywords or arguments.

1.12.4 Default

MTU negotiation is disabled.

1.12.5 Usage Guidelines

Use the **`tcp path-mtu-discovery`** command to enable the negotiation of the MTU for TCP sessions. Enabling MTU negotiation has no effect on existing TCP sessions.

TCP has the ability to dynamically discover the largest MTU that can be used on the session pipe that minimizes fragmentation and maximizes efficiency. As described in RFC 1191, `Path MTU Discovery`, the default size of an IP packet is 576 bytes. The IP and TCP portions of the frame occupy 40 bytes leaving 536 bytes for the data payload. This payload is referred to as the maximum segment size (MSS).

This command allows the MSS (and hence the MTU) to be negotiated. When you enter this command and start a TCP session, the SYN packet sent by the SmartEdge router contains a TCP option specifying a larger MSS. This larger MSS is the MTU of the outbound interface minus 40 bytes. If the MTU of the outbound interface is 1500 bytes, the advertised MSS is 1460.

Both the SmartEdge router and the remote router must be configured for MTU negotiation to work properly. If both routers have MTU negotiation enabled, the SYN from one router to the other contains the optional TCP value advertising the higher MSS. The returning SYN then advertises the higher MSS value. If one router has MTU negotiation enabled and the second router never advertises the larger MSS, the first router is locked into sending the default values.

Use the **`no`** or **`default`** form of this command to disable the negotiation of the MTU for TCP sessions.

1.12.6 Examples

The following example enables the negotiation of the MTU for TCP sessions:

```
[local]Redback(config)#tcp path-mtu-discovery
```

1.13 tcp persist-state

`tcp persist-state {min-system-memory` ***`memory`*** **`| timeout`** ***`seconds`*** **`}`**

`default tcp persist-state {min-system-memory | timeout}`

`no tcp persist-state`

1.13.1 Purpose

Indicates when to drop a TCP session that has been in a persist state for too long.

1.13.2 Command Mode

Global configuration

1.13.3 Syntax Description

`min system memory` *`memory`*	The amount of system memory in megabytes that the system must be less than in order to drop the session. If you do not specify this amount, the system will default to 50 MB.
`timeout` *`seconds`*	The number of seconds that the session must be in the persist state for in order to drop the session. If you do not specify an amount, the system will default to 600 seconds.

1.13.4 Default

The system drops TCP sessions that use less than 50 MB of memory and remain in the persist state for more than 600 seconds.

1.13.5 Usage Guidelines

Use the **`tcp persist-state`** command to identify the parameters around which to drop a TCP session that remains in the persist state for longer than a period of time that you specify. It is possible for a TCP peer to open many TCP connections and put them into the persist state. If these connections contain a large amount of data pending to send, the pending data can consume a significant portion of the system memory. This can deny service by using up so much system memory that other functions cannot work.

You can use this command with the **`show tcp brief`** command. This command will display your settings for the **`tcp persist-state`** command. You can use the **`show tcp stat`** command to display the number of TCP sessions that were dropped as a result of a TCP session staying too long in the persist state.

Use the **`default`** form of this command to set the system to the default settings for dropping TCP sessions. The default setting is that the system drops TCP sessions that use less than 50 MB of memory and remain in the persist state for more than 600 seconds. You can use the default form of this command with the min-system-memory keyword to set the system to the default setting for memory or with the timeout keyword to set the system to the default for time in the persist state.

Use the **no** form of this command to identify that the system should never drop a TCP session in the persist state.

1.13.6 Examples

The following example shows how to set the system to use the default for system memory (50 MB) while setting the time for the system to wait to drop TCP sessions to `250` seconds:

```
[local]Redback(config)#default tcp persist-state min-system-memory
[local]Redback(config)#tcp persist-state timeout 250
```

1.14 tcp-port local

tcp-port local *loc-port*

{**no** | **default**} **tcp-port local**

1.14.1 Purpose

Assign a Transmission Control Protocol (TCP) port on which the SmartEdge router listens for Access Node Control Protocol (ANCP) sessions.

1.14.2 Command Mode

ANCP configuration

1.14.3 Syntax Description

loc-port	TCP port number. The range of values is 6,068 to 10,000; the default value is 6,068.

1.14.4 Default

The default TCP port, 6,068, is assigned as the local port.

1.14.5 Usage Guidelines

Use the **tcp-port local** command to specify the TCP port on which the SmartEdge router listens for ANCP sessions.

Use the **no** or **default** form of this command to specify the default condition.

1.14.6 Examples

The following example specifies `6070` as the port number for the local TCP port:

```
[local]Redback(config-ancp)#tcp-port local 6070
```

1.15 tcp-port remote

`tcp-port remote` ***`remote-port`***

`no tcp-port remote`

1.15.1 Purpose

Filter incoming new neighbor connections using the Transmission Control Protocol (TCP) port on which the SmartEdge router receives the General Switch Management Protocol (GSMP) messages from an Access Node Control Protocol (ANCP) neighbor peer.

1.15.2 Command Mode

ANCP neighbor configuration

1.15.3 Syntax Description

`remote-port`	TCP port number. The range of values is 1,024 to 5,000.

1.15.4 Default

If a TCP remote port number is not specified for this profile, there is no restriction on the TCP remote port number in a received GSMP adjacency protocol message from an ANCP neighbor.

1.15.5 Usage Guidelines

Use the **`tcp-port remote`** command to filter incoming new neighbor connections using the TCP port number on which the SmartEdge router receives the GSMP messages from an ANCP neighbor peer.

Use the **`no`** form of this command to specify the default condition.

1.15.6 Examples

The following example specifies `7070` as the port number for a remote TCP port:

```
[local]Redback(config-ancp-neighbor)#tcp-port remote 7070
```

1.16 telnet

telnet {***ip-addr*** | ***hostname***} [***port***]

1.16.1 Purpose

Establishes a remote Telnet session from the SmartEdge router to a host.

1.16.2 Command Mode

Exec

1.16.3 Syntax Description

ip-addr	IP address of the host with which to establish the Telnet session.
hostname	Name of the host with which to establish the Telnet session. The Domain Name System (DNS) must be enabled to use the ***hostname*** argument. The ***hostname*** argument can be specified in the format name@ctx-name, where ctx-name is either the name of an existing context or the domain alias of an existing context name.
port	Optional. Transmission Control Protocol (TCP) port used to communicate with the host. The range of values is 1 to 65,536; the default value is 23.

1.16.4 Default

None

1.16.5 Usage Guidelines

Use the **telnet** command to establish a Telnet session from the SmartEdge router to a host.

Note: By default, Telnet service is enabled in all local contexts.

You can only use the ***hostname*** argument if DNS is enabled with the **ip domain-lookup**, **ip domain-name**, and **ip name-servers** commands (in context configuration mode). For more information about these commands, see *Configuring DNS*.

Use the ***port*** argument to specify a port other than the default TCP port. Ensure that the port on the remote host is activated for Telnet.

1.16.6 Examples

The following example establishes a Telnet session with a host at IP address, 192.168.190.32:

```
[local]Redback>telnet 192.168.190.32
```

The following example establishes a Telnet session to a host at IP address, 192.168.190.32, using port 2222:

```
[local]Redback>telnet 192.168.190.32 2222
```

1.17 te-metric

te-metric ***value***

no te-metric

1.17.1 Purpose

Specifies the traffic engineer (TE) metric for an interface that is used during Constrained Shortest Path First (CSPF) calculation.

1.17.2 Command Mode

RSVP interface configuration

1.17.3 Syntax Description

value	TE metric value. Valid values are 1 through 128.

1.17.4 Default

When you do not specify a TE metric, the value is the same as the Interior Gateway Protocol (IGP) metric.

1.17.5 Usage Guidelines

Use the **`te-metric`** command to specify the TE metric for an interface that is used during CSPF calculation.

Use the **`no`** form of this command to remove the TE metric value specified on an interface.

1.17.6 Examples

The following example shows how to configure the TE metric on an interface to `15`:

```
[local]Redback#configure

[local]Redback(config)#context local

[local]Redback(config-ctx)#router rsvp

[local]Redback(config-rsvp)#interface rsvp05

[local]Redback(config-rsvp-if)#te-metric 15
```

1.18 terminal length

`terminal length` ***`length`***

`default terminal length`

1.18.1 Purpose

Sets the terminal length to be used for the administrator’s terminal for the duration of the current exec session.

1.18.2 Command Mode

Exec

1.18.3 Syntax Description

length	Number of lines to be used for the terminal length. The range of values is 0 to 512; the default value is 24.

1.18.4 Default

The default terminal length is 24 lines.

1.18.5 Usage Guidelines

Use the **terminal length** command to set the length in terminal lines for an exec session. Upon exit of the exec session, the value is reset to the default length of 24 lines. Setting the terminal length to 0 disables auto-more processing.

Use the **default** form of this command to return the terminal length to the default value.

1.18.6 Examples

The following command sets the session terminal length to 30 lines:

```
[local]Redback>terminal length 30
```

1.19 terminal monitor

terminal monitor

no terminal monitor

1.19.1 Purpose

Enables the display of system events on a remote (Telnet or Secure Shell [SSH]) session continuously as they are logged.

1.19.2 Command Mode

Exec

1.19.3 Syntax Description

This command has no keywords or arguments.

1.19.4 Default

Events are not logged to administrator terminals.

1.19.5 Usage Guidelines

Use the **`terminal monitor`** command to enable the display of events on the current terminal. This command can be useful for viewing the Event Log output while connected to a system by Telnet or SSH, rather than while working on the console.

Use the **`no`** form of this command to disable terminal monitoring.

1.19.6 Examples

The following example enables the display of logged events on the current terminal while connected using Telnet:

```
[local]Redback>terminal monitor
```

1.20 terminal width

`terminal width` ***`width`***

`default terminal width`

1.20.1 Purpose

Sets the terminal width in characters to be used for the administrator’s terminal for the duration of the current exec session.

1.20.2 Command Mode

Exec

1.20.3 Syntax Description

width	Preferred terminal width setting in characters. The range of values is 5 to 65,536; the default value is 80.

1.20.4 Default

The default terminal width is 80 characters.

1.20.5 Usage Guidelines

Use the **terminal width** command to set the width in characters of the terminal for an exec session. Upon exit from this session, the value is reset to the default width of 80 characters.

Use the **default** form of this command to change the terminal width to the default value.

1.20.6 Examples

The following command changes the session terminal width to `70` characters:

```
[local]Redback>terminal width 70
```

1.21 test aaa

test aaa {**authentication** | **accounting**} **username** ***name*** **protocol radius** [**server-ip** ***ip-addr port***]

1.21.1 Purpose

Tests the communications link to a Remote Authentication Dial-In User Service (RADIUS) server.

1.21.2 Command Mode

Exec (10)

1.21.3 Syntax Description

`authentication`	Tests the communications link with an Authentication-Request message.
`accounting`	Tests the communications link with an Accounting-Request message.
`username` ***`name`***	Structured subscriber or administrator name, including the domain name.
`protocol radius`	Tests communications to a RADIUS server.
`server-ip` ***`ip-addr port`***	Optional. IP address and User Datagram Protocol (UDP) port for the server to which the message is to be sent. The range of values for the ***`port`*** argument is 1 to 65,535. If not specified, the SmartEdge router sends the test message to all servers configured in the context.

1.21.4 Default

None

1.21.5 Usage Guidelines

Use the **`test aaa`** command to test the communications link to a RADIUS server. The SmartEdge router creates an authentication or accounting message that includes the specified username and transmits it to the server as specified by the **`server-ip`** ***`ip-addr port`*** construct.

Note: You must access the context in which the servers are configured before you can run this test.

The response from the specified server occurs asynchronously; that is, the command-line interface (CLI) prompt appears immediately after you enter the command, but the system might not display the results from the server immediately. (The length of time depends on the timeout value configured for the context.)

Table 1 lists the fields that the system displays for the **`test aaa`** command.

Table 1 Field Descriptions for the test aaa Command

Field	Description
Server	IP address and UDP port number.
Server response	• Accepted—Server accepted test message. • Timeout—No response received from the server.
Attributes list	Values for attributes configured for the server.
Send count	Number of times the test message was sent.

Table 1 Field Descriptions for the test aaa Command

Field	Description
Send time	Time at which the first test message was sent.
Response time	Time at which a response from the server was received.

1.21.6 Examples

The following example tests the communications link with a RADIUS authentication message:

```
[local]Redback#test aaa authentication username b-32@local protocol
radius server-ip 10.1.1.10 1645
```

```
Apr 29 20:22:03: %AAA-7-AUTHEN: aaa_idx 92: Sending RADIUS_SERVER_TEST
to radius.
  Pid 146: Testing radius server ...
[local]Redback#
        Radius authentication test response:
           Server:           10.1.1.10/1645
           Server response:  Accepted.
           ------------------------------
           Attributes list:
           User-Name = b-32@local
           NAS_Identifier = Redback
           NAS-Port = 16842817
           NAS-Port-Type = 0
           Platform_Type = 2
           OS_Version = 4.0.5
           Service-Type = 2
           Framed-IP-Address = 255.255.255.254
           ------------------------------
           Send count:       1
           Send time:        Apr 29 20:22:03 2005
           Response time:    Apr 29 20:22:03 2005
```

The following example tests the communications link with a RADIUS accounting message:

```
[local]Redback#test aaa accounting username b-32@local protocol
radius server-ip 10.1.1.10 1646
```

```
Apr 29 20:30:24: %AAA-7-AUTHEN: aaa_idx 92: Sending RADIUS_SERVER_TEST
to radius.
  Pid 146: Testing radius server ...
[local]Redback#
        Radius accounting test response:
           Server:                 10.1.1.10/1646
           Server response:  Accepted.
           Send count:             1
           Send time:              Apr 29 20:30:24 2005
           Response time:          Apr 29 20:30:24 2005
```

The following example tests the communications link with a RADIUS authentication message when the server is shut down:

```
[local]Redback#test aaa authentication username b-32@local protocol
radius server-ip 10.1.1.10 1645

Apr 29 20:31:09: %AAA-7-AUTHEN: aaa_idx 92: Sending RADIUS_SERVER_TEST to radius.
  Pid 146: Testing radius server ...
        Radius authentication test response:
           Server:           10.1.1.10/1645
           Server response:  Timeout.
           Send count:       4
           Send time:        Apr 29 20:31:09 2005
```

1.22 threshold

threshold {**sd-ber** ***sd-ber-exp*** | **sf-ber** ***sf-ber-exp***}

{**no** | **default**} **threshold** {**sd-ber** | **sf-ber**}

1.22.1 Purpose

Specifies the Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/SDH) signal degrade bit error rate (SD-BER) or signal fail BER (SF-BER) threshold for a SONET/SDH or WAN-PHY port.

1.22.2 Command Mode

- ATM OC configuration

- Port configuration

1.22.3 Syntax Description

sd-ber *sd-ber-exp*	Value of the exponent for the threshold. The range of values is 5 to 9; the default value is 7.
sf-ber *sf-ber-exp*	Value of the exponent for the threshold. The range of values is 3 to 5; the default value is 4.

1.22.4 Default

The default thresholds for SD-BER and SF-BER are 10E-7 and 10E-4, respectively.

1.22.5 Usage Guidelines

Use the **threshold** command to specify the SONET/SDH SD-SER or SF-BER threshold for a SONET/SDH or WAN-PHY port.

Note: This command does not apply to channelized OC-12 ports.

Use the **no** or **default** form of this command to specify the default values for the SF-BER and SD-BER thresholds.

1.22.6 Examples

The following example shows how to specify the SD-BER and SF-BER thresholds as 10E-8 and 10E-6 for port 1 on the ATM OC-12c/STM-4c traffic card in slot 3:

```
[local]Redback(config)#port atm 3/1
[local]Redback(config-atm-oc)#threshold sd-ber 8
[local]Redback(config-atm-oc)#threshold sf-ber 6
```

The following example shows how to specify the SD-BER and SF-BER thresholds as 10E-8 and 10E-6 for port 1 on a 10GE traffic card in slot 3:

```
[local]Redback(config)#port ethernet 3/1 wan-phy
Note: Creating a port may cause the card to reload. Commit to continue; abort to exit without change
[local]Redback(config-port)#commit
[local]Redback(config-port)#threshold sd-ber 8
[local]Redback(config-port)#threshold sf-ber 6
```

1.23 threshold (malicious traffic)

threshold high ***high-threshold*** **low** ***low-threshold***

```
no threshold
```

1.23.1 Purpose

Specifies the high and low threshold values for malicious traffic alarms.

1.23.2 Command Mode

Malicious-traffic alarms configuration mode

1.23.3 Syntax Description

`high` *`high-threshold`*	Specifies the high threshold value, in packets. The range is 1000 to 100,000,000,000.
`low` *`low-threshold`*	Specifies the low threshold value, in packets. The range is 1000 to 100,000,000,000.

1.23.4 Default

The high threshold has a default value of 100,000, and the low threshold has a default value of 10,000.

1.23.5 Usage Guidelines

Use the **`threshold`** command to specify the high and low threshold values, in packets, for malicious traffic alarms.

Use the **`no`** form of this command to remove the alarm threshold settings.

For information about detecting and monitoring malicious traffic, see *Configuring Malicious Traffic Detection and Monitoring*.

1.23.6 Examples

The following example shows how to specify a high threshold of 300,000 packets and a low threshold of 5000 packets:

```
[local]Redback(config-malicious-traffic-alarms)#threshold high 300000 low 5000
```

1.24 time-out

```
time-out seconds
```

{**no** | **default**} **timeout**

1.24.1 Purpose

Specifies the number of seconds for the router to wait for a dynamic tunnel to be established before bringing the current subscriber or visitor down.

1.24.2 Command Mode

Dynamic Tunnel Profile configuration

1.24.3 Syntax Description

seconds	Number of seconds for the router to wait for a dynamic tunnel to be established before bringing the current subscriber or visitor down. The range of values is 2 through 10 seconds.

1.24.4 Default

3 seconds

1.24.5 Usage Guidelines

Use the **time-out** command to specify the number of seconds for the router to wait for a dynamic tunnel to be established before bringing the current subscriber or visitor down.

Use the **no** or **default** form of this command to restore the setting to its default value of 3 seconds.

1.24.6 Examples

The following example shows how to set the timeout for `prof1` to `10` seconds:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#router mobile-ip
[local]Redback(config-mip)#dynamic-tunnel-profile prof1
[local]Redback(config-mip-dyn-tun1-profile)#time-out 10
[local]Redback(config-mip-dyn-tun1-profile)#end
```

1.25 timeout abandoned

timeout abandoned *seconds*

```
no timeout abandoned
```

1.25.1 Purpose

Configures the timeout value for P2MP TCP sessions that have no active parent session.

1.25.2 Command Mode

NAT policy configuration

1.25.3 Syntax Description

seconds	Number of seconds for the router to wait for P2MP TCP sessions that have no active parent session. The timeout is in seconds. The range of values is 4 to 262143.

1.25.4 Default

The default value is 2 hours 4 minutes.

1.25.5 Usage Guidelines

1.25.6 Example

1.26 timeout (HTTP redirect)

```
timeout timeout
```

```
no timeout
```

1.26.1 Purpose

Sets the maximum time the SmartEdge router displays the customized http-redirect message to the subscriber before the subscriber HTTP session is redirected to the preconfigured URL.

1.26.2 Command Mode

HTTP redirect profile configuration

1.26.3 Syntax Description

timeout	Timeout in seconds. The range of values is 1 to 255; the default value is 1 second.

1.26.4 Default

The default timeout is 1 second.

1.26.5 Usage Guidelines

Use the **timeout** command to set the maximum time the SmartEdge router displays the customized http-redirect message to the subscriber before the subscriber HTTP session is redirected to the preconfigured URL.

Use the **no** form of this command to specify the default condition.

1.26.6 Examples

The following example shows how to configure an HTTP redirect message and a timeout for the http-redirect profile ericsson:

```
[local]Redback(config)#http-redirect profile ericsson

[local]Redback(config-hr-profile)#message "Please wait while you are
redirected to the portal server. Thank you."

[local]Redback(config-hr-profile)#timeout 20
```

1.27 timeout (L2TP Peers)

timeout *seconds*

{**no** | **default**} **timeout**

1.27.1 Purpose

Specifies the amount of time to wait for an acknowledgment before a control message is retransmitted to a Layer 2 Tunneling Protocol (L2TP) peer.

1.27.2 Command Mode

L2TP peer configuration

1.27.3 Syntax Description

seconds	Number of seconds to wait for an acknowledgment. The range of values is 1 to 30; the default value is 3.

1.27.4 Default

3 seconds.

1.27.5 Usage Guidelines

Use the **timeout** command to specify the amount of time to wait for an acknowledgment before a control message is retransmitted to an L2TP peer. You need only increase the value if many sessions are established or if the media is slow.

Use the **no** or **default** form of this command to reset the timeout to the default.

1.27.6 Examples

The following example shows how to configure the peer so that retransmission of a control message occurs after 5 seconds without an acknowledgment:

```
[local]Redback(config-ctx)#l2tp-peer name peer1

[local]Redback(config-l2tp)#timeout 5
```

1.28 timeout login

timeout login response *minutes*

{**no** | **default**} **timeout login response**

1.28.1 Purpose

Sets the amount of time the system waits before timing out during a logon attempt after a Telnet or SSH (Secure Shell) session starts.

1.28.2 Command Mode

Global configuration

1.28.3 Syntax Description

`response` *`minutes`*	Time, in minutes, that the system waits before timing out during a logon attempt after a Telnet or SSH session starts. The range of values is 1 to 99999; the default value is 10.

1.28.4 Default

The system waits 10 minutes for a response during a logon attempt after a Telnet or SSH session starts.

1.28.5 Usage Guidelines

Use the **`timeout login`** command to set the amount of time the system waits before timing out during log on attempt after a Telnet or SSH session starts.

Use the **`no`** form of this command to disable the logon timeout value.

Use the **`default`** form of this command to configure the default logon timeout value.

1.28.6 Examples

The following example configures the system to time out if a user does not enter logon information for `5` minutes:

```
[local]Redback(config)#timeout login response 5
```

1.29 timeout (NAT)

`timeout` {**`abandoned`** *`seconds`* | **`basic`** *`seconds`* | **`fin-reset`** *`seconds`* | **`icmp`** *`seconds`* | **`syn`** *`seconds`* | **`tcp`** *`seconds`* | **`udp`** *`seconds`*}

`no timeout` {**`abandoned`** | **`basic`** | **`fin-reset`** | **`icmp`** | **`syn`** | **`tcp`** | **`udp`**}

1.29.1 Command Mode

- NAT policy configuration
- Policy group class configuration

1.29.2 Syntax Description

`abandonedse` *`conds`*	Period, in seconds, after which NAT P2MP sessions with no parent sessions time out. The range of values is 4 to 131, 071; the default value is 3600 (1 hour). This construct is supported only by policies using Network Access Port Translation (NAPT).
`basic` *`second`* *`s`*	Period, in seconds, after which basic NAT time out. The range of values is 4 to 262,143; the default value is 3,600 (1 hour). This construct is supported only for basic NAT, not NAPT.
`fin-reset` *`seconds`*	Period, in seconds, after which NAT for Transmission Control Protocol (TCP) FINISH and RESET packets time out. The range of values is 4 to 65,535; the default value is 240. This construct is supported only by policies using NAPT.
`icmp` *`seconds`*	Period, in seconds, after which NAT for Internet Control Message Protocol (ICMP) packets time out. The range of values is 4 to 65,535; the default value is 60. This construct is supported only by policies using NAPT.
`syn` *`seconds`*	Period, in seconds, after which NAT for TCP SYN packets time out. The range of values is 4 to 65,535; the default value is 128. This construct is supported only by policies using NAPT.
`tcp` *`seconds`*	Period, in seconds, after which NAT for established TCP connections time out. The range of values is 4 to 262,143. The default value is 86,400 (24 hours). This construct is supported only by policies using NAPT.
`udp` *`seconds`*	Period, in seconds, after which NAT for User Datagram Protocol (UDP) packets time out. The range of values is 4 to 65,535; the default value is 120. This construct is supported only by policies using NAPT.

1.29.3 Default

For default values, see the Syntax Description section. For the ignore action in a NAT policy, all default timeouts are 20 seconds.

1.29.4 Usage Guidelines

Use the `timeout` command to modify the period after which NAT times out if no activity occurs. Timeout applies only if there is relevant translation.

Use the `no` form of this command to reset the timeout to its default value.

1.29.5 Examples

The following example configures basic NAT to time out after no activity has occurred for 7200 seconds (2 hours):

```
[local]Redback(config-ctx)#ip nat pool NAT-POOL
[local]Redback(config-nat-pool)#address 171.71.71.0/24
[local]Redback(config-nat-pool)#exit
[local]Redback(config-ctx)#nat policy NAT-1
[local]Redback(config-policy-nat)#pool NAT-POOL local
[local]Redback(config-policy-nat)#timeout basic 7200
```

The following example configures an enhanced NAT policy with a class that times out P2MP TCP sessions with no parent sessions after one hour (3600 seconds):

```
[local]Redback(config)#context nat-context
[local]Redback(config-ctx)#nat policy nat-policy-1 enhanced
[local]Redback(config-policy-nat)#pool nat-pool-1 nat-context
[local]Redback(config-policy-nat)#access-group nat-acl
[local]Redback(config-policy-acl)#class nat-class
[local]Redback(config-policy-acl-class)#timeout abandoned 3600
```

Typically, more attributes would be added to the policy.

1.30 timeout session

timeout session idle ***minutes***

{**no** | **default**} **timeout session idle**

1.30.1 Purpose

Specifies the maximum idle timeout for any administrator account or administrator sessions on the console port (in global configuration mode), or disables the global or administrator session idle timer.

1.30.2 Command Mode

- Administrator configuration
- Global configuration

1.30.3 Syntax Description

idle ***minutes***	Time, in minutes, that the session remains connected without input before timing out. The range of values is 1 to 99,999; the default value is 10.

1.30.4 Default

The maximum session idle time for the administrator account is governed by the global session idle timer. The maximum session idle time is 10 minutes.

1.30.5 Usage Guidelines

Use the **timeout session idle** command to specify the maximum idle time for any administrator account session, or disable the global or administrator session idle timer. When specified, the system disconnects any session with no input for the specified time. The value that you specify in the administrator session overrides the value specified for the global session idle timer; if disabled, there is no timeout value.

To specify a different timeout session for a specific administrator, use this command (in administrator configuration mode); the value you specify for a specific administrator overrides the value specified for the global session idle timer.

Use the **no** form of this command to disable the global session idle timer; the global form of the command does not affect the session idle timer for a specific administrator.

Use the **default** form of this command to specify the default value for the global session idle timer.

1.30.6 Examples

The following example configures the system to disconnect any administrator session after remaining idle for `30` minutes:

```
[local]Redback(config)#timeout session idle 30
```

The following example specifies the session idle timer for this administrator to `60` minutes. This value overrides the value specified for the global session idle timer:

```
[local]Redback(config-administrator)#timeout session idle 60
```

1.31 timeout (subscriber)

For absolute timeouts, to set the minutes allowed before session termination:

timeout absolute ***minutes***

no timeout absolute

For idle timeouts, to set the direction, activity threshold, and allowed minutes of inactivity before session termination:

timeout idle {***minutes*** | **direction** {**in** | **out**} | **threshold** ***bps***}

no timeout idle {***minutes*** | **direction** {**in** | **out**} | **threshold** ***bps***}

For idle timeouts, to set the direction and allowed minutes of inactivity before session termination:

timeout idle {***minutes*** | **direction** {**in** | **out**}}

no timeout idle {***minutes*** | **direction** {**in** | **out**}}

1.31.1 Purpose

Configures the absolute or idle session timeout criteria for a subscriber session.

1.31.2 Command Mode

Subscriber configuration

1.31.3 Syntax Description

absolute	Specifies an absolute session timeout. After the time defined by the ***minutes*** argument, the subscriber is disconnected regardless of activity.
minutes	Time, in minutes, that elapses before a session times out. The range of values is 1 to 596523.
idle	Specifies an idle session timeout. If no activity above the minimum level takes place for the amount of time defined by the ***minutes*** argument, the subscriber is disconnected. Activity is measured in the direction specified by the optional **direction** {**in** | **out**} construct. Normally, the **idle** ***minutes*** option provides a session timeout value to the subscriber currently being configured; that is, the subscriber-specific session timeout value. However, when the **subscriber dhcp-lease idle-timeout** command is entered for the context, the **idle** ***minutes*** option instead sets the subscriber-specific DHCP lease times.
direction {**in** | **out**}	Optional. Specifies the direction on which the idle session timeout minutes are measured. The keyword in specifies the incoming (receive) direction, while the keyword out specifies the outgoing (transmit) direction.
threshold ***bps***	Optional. Specifies the minimum level of activity below which a subscriber session is considered inactive. Enter the argument ***bps*** in bytes per second.

1.31.4 Default

Subscriber sessions do not time out. Idle timeouts, if configured, apply to traffic in both the send and receive directions, and the bytes-per-second threshold is zero.

1.31.5 Usage Guidelines

Use the **`timeout`** command to set the absolute or idle session timeout criteria for a subscriber session. The system terminates subscriber sessions when they reach timeout.

Note: This command applies to either locally terminated or Layer 2 Tunneling Protocol (L2TP) access concentrator (LAC) subscriber sessions.

Note: Keepalive messages are not considered traffic for purposes of measuring idle time.

Use the **`no`** forms of this command to restore the default behaviors.

1.31.6 Examples

The following example sets an absolute timeout value of `20` minutes:

```
[local]Redback(config-sub)#timeout absolute 20
```

1.32 timer retry

`timer retry` seconds [**`limit`** *`count`*]

`no timer retry`

1.32.1 Purpose

Specifies the interval that the headend (ingress) router waits between attempts to establish the primary path during Constrained Shortest Path First (CSPF) calculation.

1.32.2 Command Mode

- RSVP router configuration

1.32.3 Syntax Description

seconds	Interval in seconds that the ingress router waits between attempts to establish the primary path. The range of values is 1 to 4294967295.
limit *count*	Number of times that the ingress router attempts to establish the primary path. No maximum value exists for the count.

1.32.4 Default

The ingress router waits for 30 seconds between attempts to establish the primary path.

1.32.5 Usage Guidelines

Use the **timer retry** command to specify the interval that the ingress router waits between attempts to establish the primary path during CSPF calculation.

Use the **no** form of this command to restore the default condition.

1.32.6 Examples

The following example shows how to specify that the ingress router wait 5 seconds between attempts to establish the primary path:

```
[local]Redback#configure
[local]Redback(config)#context local
[local]Redback(config-ctx)#router rsvp
[local]Redback(config-rsvp)#timer retry 5
```

1.33 timers active-open

timers active-open *interval*

{**no** | **default**} **timers active-open** *interval*

1.33.1 Purpose

Specifies the interval the router waits for a BGP peer to come up after a graceful restart before starting the best path computation.

1.33.2 Command Mode

- BGP neighbor configuration
- BGP peer group configuration

1.33.3 Syntax Description

interval	Interval in seconds that the router waits for a BGP peer to come up after a graceful restart before starting the best path computation. Range is from 1 through 600.

1.33.4 Default

The router waits 20 seconds for a BGP peer to come up before starting the best path computation.

1.33.5 Usage Guidelines

Use the **timers active-open** command to specify the interval the router waits for a BGP peer to come up after a graceful restart before starting the best path computation.

Note: For BGP sessions over port pseudowires, we recommend setting the **timers active-open** interval to be 40 seconds or more. This is because the port pseudowire typically takes 30 seconds to come up after a switchover.

The **timers active-open** interval configuration for a BGP neighbor (in BGP neighbor configuration mode) takes precedence over an interval configuration for a peer group (in BGP peer group configuration mode). For example, if you specify an interval of 40 seconds for a BGP neighbor that is part of a peer group that has a **timers active-open** interval of 20 seconds, the BGP neighbor interval (40 seconds) overrides the peer group interval of 20 seconds.

Use the **no** or **default** form of this command to return the router to the default configuration, where the router waits 20 seconds for a BGP peer to come up before starting the best path computation.

1.33.6 Examples

1.33.6.1 Configuring the timers active-open Interval for a BGP Neighbor (Peer)

The following example configures the router to wait 40 seconds for the BGP peer (whose IP address is `123.45.34.2`) to come up after a graceful restart before starting the best path computation:

```
[local]Redback(config-ctx)#router bgp 100
[local]Redback(config-bgp)#neighbor 123.45.34.2 external
[local]Redback(config-bgp-neighbor)#timers active-open 40
```

1.33.6.2 Configuring the timers active-open Interval for a BGP Peer-Group and a BGP Neighbor (Peer)

The following example configures the router to:

- Wait 50 seconds for any BGP peer in peer group `sanjose` to come up after gracefully restarting before starting the best path computation.

- Wait 100 seconds for a particular peer (whose IP address is `123.45.34.2`) to come up after a graceful restart before starting the best path computation

If the peer on IP address `123.45.34.2` is part of the peer group `sanjose`, the router waits 100 seconds for that peer to come up before starting best path computation. This is because the **`timers active-open`** interval configuration in BGP neighbor configuration mode takes precedence over the **`timers active-open`** interval configuration in BGP peer group configuration mode.

```
[local]Redback(config-ctx)#router bgp 65

[local]Redback(config-bgp)#peer-group sanjose external

[local]Redback(config-bgp-peer-group)#timers active-open 50

[local]Redback(config-bgp-peer-group)#exit

[local]Redback(config-bgp)#neighbor 123.45.34.2 external

[local]Redback(config-bgp-neighbor)#peer-group sanjose

[local]Redback(config-bgp-neighbor)#timers active-open 100
```

1.34 timers basic

`timers basic` ***`update-interval invalid-interval holddown-interval flush-interval`***

{**`no`** | **`default`**} **`timers basic`**

1.34.1 Purpose

Modifies the Routing Information Protocol (RIP) or RIP next generation (RIPng) timers for the specified RIP or RIPng instance or interface.

1.34.2 Command Mode

- RIP interface configuration
- RIPng interface configuration
- RIPng router configuration
- RIP router configuration

1.34.3 Syntax Description

update-interval	Interval, in seconds, at which RIP or RIPng updates are sent. The range of values is 1 to 32,767; the default value is 30.
invalid-interval	Interval, in seconds, before a route is declared invalid after no updates are received. This value should be at least three times the value for the ***update-interval*** argument. The range of values is 1 to 32,767; the default value is 180.
holddown-interval	Interval, in seconds, before better routes are released. The range of values is 1 to 32,767; the default value is 180.
flush-interval	Interval, in seconds, before a route is flushed from the routing table. This value must be larger than the value for the ***invalid-interval*** argument. The range of values is 1 to 32,767; the default value is 240.

1.34.4 Default

RIP and RIPng updates are sent every 30 seconds, a route is declared invalid if an update is not received after 180 seconds, better routes are released after 180 seconds, and a route is flushed from the routing table after 240 seconds.

1.34.5 Usage Guidelines

Use the **timers basic** command in RIP or RIPng interface configuration mode to modify the RIP or RIPng timers for the specified interface.

Use the **timers basic** command in RIP or RIPng router configuration mode to modify the RIP or RIPng timers for the specified instance.

Use the **no** or **default** form of this command to restore the default RIP or RIPng timer settings.

1.34.6 Examples

The following example sets the RIP timers for the RIP instance `rip001`. The update interval is set to `45` seconds, the invalid interval to `200` seconds, the holddown interval to `200` seconds, and the flush interval to `260` seconds:

```
[local]Redback(config-ctx)#rip001
[local]Redback(config-rip)#timers basic 45 200 200 260
```

The following example modifies the RIP timers for the `fe01` interface. The update interval is set to `45` seconds, the invalid interval to `200` seconds, the holddown interval to `200` seconds, and the flush interval to `260` seconds:

```
[local]Redback(config-ctx)#router rip rip002

[local]Redback(config-rip)#interface fe01

[local]Redback(config-rip-if)#timers basic 45 200 200 260
```

1.35 timers keepalive

timers keepalive ***interval*** **holdtime** ***interval***

no timers keepalive ***interval*** **holdtime** ***interval***

1.35.1 Purpose

Modifies Border Gateway Protocol (BGP) timers for the routing instance, neighbor, or peer group.

1.35.2 Command Mode

- BGP neighbor configuration
- BGP peer group configuration
- BGP router configuration

1.35.3 Syntax Description

keepalive *interval*	Interval, in seconds, at which the BGP routing process sends keepalive messages. The range of values is 1 to 65,535; the default value is 60.
holdtime *interval*	Interval, in seconds, after which, if the BGP routing process has not received a keepalive message, it considers the neighbor to be unavailable. The range of values is 3 to 65,535; the default value is 180.

1.35.4 Default

The keepalive time is 60 seconds. The hold time is 180 seconds.

1.35.5 Usage Guidelines

Use the **timers keepalive** command in BGP router configuration mode to modify keepalive and hold time timers for all BGP neighbors.

Use the **timers keepalive** command in BGP neighbor configuration mode to modify keepalive and hold time timers for a specific neighbor. Values set for a BGP neighbor override the values set for the BGP routing instance.

Use the **timers keepalive** command in BGP peer group configuration mode to modify keepalive and hold time timers for a peer group.

Note: If a neighbor is part of a peer group, and you try to apply this command in BGP neighbor configuration mode, the timer conditions are not applied to the neighbor. Use the **timers keepalive** command in BGP peer group configuration mode instead.

Use the **no** form of this command to restore timer settings to their default values.

1.35.6 Examples

The following example sets the keepalive period to `45` seconds and the holdtime to `135` seconds for only the neighbor at IP address `123.45.34.2`:

```
[local]Redback(config-ctx)#router bgp 100
[local]Redback(config-bgp)#neighbor 123.45.34.2 external
[local]Redback(config-bgp-neighbor)#timers keepalive 45 holdtime 135
```

1.36 timeslot (DS0)

timeslot ***start1* - *end1*, [*start2* - *end2*,...]**

default timeslot

1.36.1 Purpose

Defines a one or more groups of timeslots in a channelized DS1 or E1 channel (fractional T1/E1).

1.36.2 Command Mode

DS-0 group configuration

1.36.3 Syntax Description

start1 - end1	Specifies the first group of timeslots: *start1* through *end1*.
start2 - end2	Specifies the second group of timeslots: *start2* through *end2*.

1.36.4 Default

Only the starting timeslot is provisioned.

1.36.5 Usage Guidelines

The **timeslot** command defines one or more DS0 subchannels added to the *starting* DS0 subchannel (also known as a timeslot) within the parent DS1 or E1 (fractional T1/E1).

Note: The *starting* timeslot is defined as the timeslot provisioned by the *port ds0s* command prior to entering the **timeslot** command.

- You can only specify DS0 timeslots that are greater than or equal to the starting timeslot within the parent DS1 or E1 (fractional T1/E1). The *port ds0s* command identifies the lowest numbered timeslot for the subchannel within its parent DS1 or E1.

- The starting DS0 timeslot is the subchannel set by the *port ds0s* command in its *nxds0-channel-id* argument.

- The **timeslot** command adds additional timeslots to the current DS0 subchannel.

- Each channelized DS1 channel is comprised of 24 NxDS0 (DS0s, DS0-Group) or up 24 timeslots for single NxDS0 starting from the timeslot number 1.

- Each channelized E1 channel is comprised of 31 NxDS0 (DS0s, DS0-Group) or up 31 timeslots for single NxDS0 starting from the timeslot number 1.

1.36.6 Examples

In the following example, the `port ds0s 2/1:2:3:1` command establishes timeslot `1` as the default. The `timeslot 1-12` command adds timeslots 2 through 12 to the NxDS0 group (fractional T1).

```
config
!
service multiple-contexts
!
software license
 all-ports password <plain text passcode> card ch-oc3oc12-8or2-port slot 2
!
!
card ch-oc3oc12-8or2-port 2
 no shutdown
 clock-source global-reference
!
!
! Example for POS, using SDH AUG mapping au3-no-tug to channelize:
! chSTM1 -> DS3
! chSTM1 -> chDS3
! chSTM1 -> chDS3 -> DS1
! chSTM1 -> chDS3 -> chDS1 -> DS0s
! chSTM1 -> chDS3
! chSTM1 -> chDS3 -> E1
! chSTM1 -> chDS3 -> chE1 -> DS0s
!
port channelized-stm1 2/1 pos
 no shutdown
 aug-mapping au3-no-tug
 clock-source card-reference
 !
 port ds3 2/1:1
  no shutdown
  clock-source card-reference
  encapsulation ppp
  bind interface pos_chstm1->ds3_1 redkite1
 !
 port channelized-ds3 2/1:2
  no shutdown
  clock-source card-reference
 !
  port ds1 2/1:2:1
   no shutdown
   clock-source card-reference
   encapsulation ppp
   bind interface pos_chstm1->chds3->ds1_1 redkite1
  !
  port ds1 2/1:2:2
   no shutdown
   clock-source card-reference
   bind interface pos_chstm1->chds3->ds1_2 redkite1
  !
  port channelized-ds1 2/1:2:3
   no shutdown
   clock-source card-reference
  !
   port ds0s 2/1:2:3:1
    no shutdown
    timeslot 1-12
    encapsulation ppp
    bind interface pos_chstm1->chds3->chds1->ds0s_1 redkite1
```

1.37 timeslot (ces)

`timeslot` ***`ds0-group`*** | ***`ds0-group-start`*** **`-`** ***`ds0-group-end`***

1.37.1 Purpose

Configures the timeslots to be added to a CESoPSN interworking function (IWF).

1.37.2 Command Mode

DS0 Channel Config Mode.

1.37.3 Syntax Description

ds0-group	Single timeslot to add. The range of values is 1 to 31 for E1, and 1 to 24 for T1.
ds0-group-start	First in a range of timeslots to add. The range of values is 1 to 31 for E1, and 1 to 24 for T1.
ds0-group-end	Last in a range of timeslots to add. The range of values is 1 to 31 for E1, and 1 to 24 for T1.

1.37.4 Default

No timeslots are added..

1.37.5 Usage Guidelines

Operation is identical to existing STM1 channelized POS cards.

1.37.6 Examples

The following example shows how to select timeslot 16 on DS0 slot 1, port 1::

```
[local]Redback(config)#port ds0s 1/1:1:1:1
Redback(config-ds0-ces)#timeslot 16
```

The following example shows how to select the range of timeslots from 1 to 10 on DS0 slot 1, port 1::

```
[local]Redback(config)#port ds0s 1/1:1:1:1
Redback(config-ds0-ces)#timeslot 1 - 10
```

1.38 timers password

timers password *interval*

no timers password *interval*

1.38.1 Purpose

Configures the time interval, in seconds, during which an old Message Digest 5 (MD5) password can coexist with a new MD5 password for authentication.

1.38.2 Command Mode

BGP router configuration

1.38.3 Syntax Description

interval	Interval, in seconds, during which the new and old MD5 passwords coexist. The range of values is 1 to 3,600.

1.38.4 Default

The timer interval is set to 1,800 seconds.

1.38.5 Usage Guidelines

Use the **timers password** command to configure the time interval, in seconds, during which an old MD5 password can coexist with a new MD5 password for authentication. Configuring the password timer interval affects only the Border Gateway Protocol (BGP) peers which have existing MD5 passwords replaced after this configuration is committed.

1.38.6 Examples

The following example allows new MD5 passwords for BGP peers to coexist with the password being replaced for `300` seconds (5 minutes):

```
[local]Redback(config-ctx)#router bgp 1000
[local]Redback(config-bgp)#timers password 300
```

1.39 traceroute

traceroute {*ip-addr* | *hostname*} [**count** *number*] [**df**] [**dispTTL** *ttl*] [**icmp**] [**initialttl** *ttl*] [**maxttl** *ttl*] [**nh**] [**nr**] [**port** *port*] [**size** *bytes*] [**source** *ip-addr*] [**timeout** *seconds*] [**tos**] [**verbose**]

1.39.1 Purpose

Traces the IP route that packets take when traveling to the specified destination.

1.39.2 Command Mode

- Exec

1.39.3 Syntax Description

ip-addr	IP address to be traced.
hostname	Hostname to be traced. Domain Naming System (DNS) must be enabled.
count *number*	Optional. Number of probes to send. The range of values is 1 to 1,000; the default value is 3.
df	Optional. Sets the Don't Fragment (DF) bit on outbound traceroute packets. With this bit set, the traceroute packet is dropped whenever it would normally be fragmented. An Internet Control Message Protocol (ICMP) Unreachable, Needs Fragmentation packet is sent to the sender.
dispTTL *ttl*	Optional. Display time-to-live (TTL) bit is set.
icmp	Optional. Uses Internet Control Message Protocol (ICMP) echo instead of a User Datagram Protocol (UDP) datagram.
initialttl *ttl*	Optional. Initial time to live. The range of values is 1 to 255.
maxttl *ttl*	Optional. Maximum time to live. The range of values is 1 to 255.
nh	Optional. Uses the next-hop maximum transmission unit (MTU) if there is a need fragmentation error.
nr	Optional. Prints hop addresses numerically rather than symbolically.
port *port*	Optional. Destination UDP port number. The range of values is 1 to 65,535; the default value is 33,434.
size *bytes*	Optional. Datagram size in octets. The range of values is 0 to 32,768.
source *ip-addr*	Optional. IP source address of the ping packets. An interface with this IP address must exist.
timeout *seconds*	Optional. Amount of time, in seconds, for each probe sent. The range of values is 2 to 86,400; the default value is 2.
tos	Optional. Specifies the type of service (ToS) in probe packets. The range of values, in hex, is 0x0 to 0xff.
verbose	Optional. Provides verbose output.

1.39.4 Default

The **traceroute** command sends three 140-byte packets on UDP port 33434, using a timeout of 2 seconds and a time to live value of 30.

1.39.5 Usage Guidelines

Use the `traceroute` command to trace the routes that packets take when traveling to the specified destination. Each line in the display shows the next hop in the path between the system and the destination address.

You can only use the *`hostname`* argument if DNS is enabled via the `ip domain-lookup`, `ip domain-name`, and `ip name-servers` commands (in context configuration mode). For more information about these commands, see *Configuring DNS*.

If the destination IP address of the traced route results in the packet going through a Multiprotocol Label Switching (MPLS) label-switched path (LSP), the output displays the label stack along each hop of the LSP.

The `ping` and `traceroute` commands (in exec mode) can have vastly different outcomes, depending on the context in which the commands are issued. In particular, a destination (as denoted by an IP address) that can be reached by the `ping` or `traceroute` command in one context might not be reachable from another context.

Press `Ctrl+C` to stop a traceroute.

1.39.6 Examples

The following command discovers the route from the local context to the IP address `206.124.29.1`, using `100`-byte packets, UDP port `73`, ttl `20`, timeout `3`, and count `3`:

```
[local]Redback>traceroute 206.124.29.1 timeout 3 count 3 maxttl 20
port 73 size 100
```

```
traceroute to (206.124.29.1), 20 hops max, 140 byte packets
 1  155.53.145.254 (155.53.145.254)     0 ms  0 ms  0 ms
 2  155.53.200.254 (155.53.200.254)     0 ms  0 ms 16 ms
 3  206.83.66.193 (206.83.66.193)      16 ms 16 ms 16 ms
 4  206.83.90.66 (206.83.90.66)        16 ms 16 ms 16 ms
 5  157.130.193.197 (157.130.193.197)  16 ms 33 ms 16 ms
 6  157.130.194.18 (157.130.194.18)    16 ms 33 ms 16 ms
 7  209.104.192.49 (209.104.192.49)    50 ms 66 ms 50 ms
 8  209.104.198.38 (209.104.198.38)    50 ms 66 ms 66 ms
 9  206.124.1.22 (206.124.1.22)        66 ms 66 ms 66 ms
10  206.124.29.1 (206.124.29.1)        83 ms 66 ms 83 ms
```

The following example displays a destination IP address of a traced route resulting in the packet going through an MPLS LSP, and includes the label stack along each hop of the LSP:

```
[local]Redback>traceroute 5.5.5.5
```

```
se_traceroute to 5.5.5.5 (5.5.5.5), 30 hops max, 40 byte packets
 1  100.1.1.1 (100.1.1.1)   4.749 ms  4.111 ms  3.986 ms
 2  40.1.1.2 (40.1.1.2)   6.321 ms  6.457 ms  5.289 ms
         MplsLabel: 19 MplsExpBits: 0 TTL: 1
         MplsLabel: 786434 MplsExpBits: 0 TTL: 1
 3  60.1.1.1 (60.1.1.1)   3.815 ms  4.159 ms  4.120 ms
         MplsLabel: 786434 MplsExpBits: 0 TTL: 1
 4  5.5.5.5 (5.5.5.5)   5.870 ms  9.108 ms  6.639 ms
```

1.40 traceroute ipv6

traceroute ipv6 {*ipv6-addr* | *hostname*} [**first-hop** *num-hops-to-skip*] [**hop-limit** *num-hops*] [**long**] [**numeric**] [**port** *port*] [**size** *bytes*] [**source** *ip-addr*] [**timeout** *seconds*] [**verbose**]

1.40.1 Purpose

Traces the IP route that IPv6 packets take when traveling to the specified destination.

1.40.2 Command Mode

Exec

1.40.3 Syntax Description

ipv6-addr	IPv6 address to be traced.
hostname	Hostname to be traced. Domain Naming System (DNS) must be enabled.
first-hop *num-hops-to-skip*	Optional. The number of hops to skip before starting the trace. The range of values is 1 to 255.
hop-limit *num-hops*	Optional. The maximum number of hops to trace. The range of values is 1 to 255.
long	Optional. Specify long output, including numeric and hostname information.
numeric	Optional. Prints hop addresses numerically rather than symbolically.
port *port*	Optional. Destination UDP port number. The range of values is 1 to 65,535; the default value is 33,434.
size *bytes*	Optional. Datagram size in octets. The range of values is 0 to 32,768.
source *ip-addr*	Optional. IPv6 source address of the ping packets. An interface with this IPv6 address must exist.
timeout *seconds*	Optional. Amount of time, in seconds, for each probe sent. The range of values is 2 to 86,400; the default value is 2.
verbose	Optional. Provides verbose output.

1.40.4 Default

The **traceroute ipv6** command sends three 140-byte packets on UDP port 33434, using a timeout of 2 seconds and a time to live value of 30.

1.40.5 Usage Guidelines

Use the **traceroute ipv6** command to trace the routes that packets take when traveling to the specified destination. Each line in the display shows the next hop in the path between the system and the destination address.

You can only use the ***hostname*** argument if DNS is enabled via the **ip domain-lookup**, **ip domain-name**, and **ipv6 name-servers** commands (in context configuration mode). For more information about these commands, see *Configuring DNS*.

If the destination IPv6 address of the traced route results in the packet going through a Multiprotocol Label Switching (MPLS) label-switched path (LSP), the output displays the label stack along each hop of the LSP.

The **ping ipv6** and **traceroute ipv6** commands (in exec mode) can have vastly different outcomes, depending on the context in which the commands are issued. In particular, a destination (as denoted by an IPv6 address) that can be reached by the **ping ipv6** or **traceroute ipv6** command in one context might not be reachable from another context.

Press **Ctrl+C** to stop a traceroute.

1.40.6 Examples

1.40.6.1 IPV6 Trace Route

The following command discovers the route from the vpn1 context to the IPv6 address `77::103`. This example displays output for a route with three labels:

```
[vpn1]PE>traceroute ipv6 77::103


traceroute6 to 77::103 from 11::100, 30 hops max, 12 byte packets
 1  ::ffff:19.20.21.2  3.504 ms  3.263 ms  3.616 ms
        MplsLabel: 589826 MplsExpBits: 0 TTL: 1
        MplsLabel: 589825 MplsExpBits: 0 TTL: 1
 2  ::ffff:20.20.20.102  2.784 ms  3.118 ms  3.094 ms
        MplsLabel: 589826 MplsExpBits: 0 TTL: 1
        MplsLabel: 589825 MplsExpBits: 0 TTL: 2
 3  * * *
 4  fe80::230:88ff:fe04:3009  3.146 ms  2.695 ms  3.016

[local]PE#show bgp route ipv4 vpn labels
```

```
VPN RD: 100.100.100.100:1
Network            Next Hop            Rcv Label      Alloc Label
11.11.11.100/32    0.0.0.0                nolabel        589826
15.16.17.0/24      0.0.0.0                nolabel        589826
77.77.77.103/32    100.100.100.103        589826         nolabel

VPN RD: 100.100.100.103:1
Network            Next Hop            Rcv Label      Alloc Label
77.77.77.103/32    100.100.100.103        589826         nolabel
[local]PE1#

[vpn1]PE#show ipv6 route
Codes: C - connected, S - static, S dv - dvsr, R - RIP, e B - EBGP, i B - IBGP
       O   - OSPF, O3  - OSPFv3, IA - OSPF(v3) inter-area,
       N1  - OSPF(v3) NSSA external type 1, N2  - OSPF(v3) NSSA external type 2
       E1  - OSPF(v3) external type 1, E2  - OSPF(v3) external type 2
       i   - IS-IS, L1 - IS-IS level-1,  L2  - IS-IS level-2, N - NAT
       IPH - IP Host, SUB A - Subscriber address, SUB S - Subscriber static
       SUB P - AAA downloaded aggregate subscriber routes
       SUB N - Subscriber ND, SUB D - Subscriber DHCP-PD
       M F - Mobile Sub Foreign Agent, M H - Mobile Sub Home Agent
       M G - Mobile Sub GTP
       A - Derived Default, MeH - Media Nexthop
       >   - Active Route, * - LSP

Type    Network             Next Hop        Dist  Metric    UpTime  Interface

> C     11::100/128         11::100            0       0  06:09:06  net1
> C     17::/64                                0       0  06:08:58  to-CE1
> e B   18::100/128         17::2             20       0  06:08:36  to-CE1
> e B   77::103/128         100.100.100.103   20       0  06:06:18
> e B   88::/64             100.100.100.103   20       0  06:06:06
> e B   99::103/128         100.100.100.103   20       0  06:06:06
> e B   200::/64            17::2             20       0  05:30:00  to-CE1
> e B   300::/64            17::2             20       0  05:30:00  to-CE1
> e B   400::/64            17::2             20       0  05:30:00  to-CE1
> e B   500::/64            17::2             20       0  05:30:00  to-CE1
> e B   600::/64            17::2             20       0  05:30:00  to-CE1
> e B   800::/64            100.100.100.103   20       0  05:30:30
> e B   3dde::1/128         17::2             20       0  05:28:10  to-CE1
> e B   3ffe::1/128         17::2             20       0  05:28:10  to-CE1
> e B   4000:0:2:1::/64     17::2             20       0  05:28:40  to-CE1
> e B   8bbc::1/128         17::2             20       0  05:28:40  to-CE1
> e B   8f82::/64           17::2             20       0  05:28:10  to-CE1
> e B   f2cb::/64           100.100.100.103   20       0  05:27:40
[

[vpn1]PE#show nd neighbor
IPv6 Address                            Age    Link-layer Addr   State  Circuit
17::1                                   0      00:30:88:15:b0:c7 intf   1/3 vla
```

```
17::2                                   60     00:30:88:15:b0:c8 probe  1/3 vlan-i
fe80::230:88ff:fe15:b0c7                0      00:30:88:15:b0:c7 intf   1/3 vlan-i
fe80::230:88ff:fe15:b0c8                76     00:30:88:15:b0:c8 probe  1/3 vlan-i
```

```
! PE2 is the node on the remote end of the route.

[CE2]PE2#show ipv6 route
Codes: C - connected, S - static, S dv - dvsr, R - RIP, e B - EBGP, i B - IBGP
       O   - OSPF, O3  - OSPFv3, IA - OSPF(v3) inter-area,
       N1  - OSPF(v3) NSSA external type 1, N2  - OSPF(v3) NSSA external type 2
       E1  - OSPF(v3) external type 1, E2  - OSPF(v3) external type 2
       i   - IS-IS, L1 - IS-IS level-1,  L2  - IS-IS level-2, N - NAT
       IPH - IP Host, SUB A - Subscriber address, SUB S - Subscriber static
       SUB P - AAA downloaded aggregate subscriber routes
       SUB N - Subscriber ND, SUB D - Subscriber DHCP-PD
       M F - Mobile Sub Foreign Agent, M H - Mobile Sub Home Agent
       M G - Mobile Sub GTP
       A - Derived Default, MeH - Media Nexthop
       >   - Active Route, * - LSP

Type    Network             Next Hop        Dist  Metric    UpTime  Interface

> e B   11::100/128         88::2             20       0  06:12:29  to-vpn-smith
> e B   17::/64             88::2             20       0  06:12:29  to-vpn-smith
> e B   18::100/128         88::2             20       0  06:12:29  to-vpn-smith
> e B   77::103/128         88::2             20       0  06:14:30  to-vpn-smith
> C     88::/64                                0       0  06:14:50  to-vpn-smith
> C     99::103/128         99::103            0       0  06:15:00  loCE2
> e B   200::/64            88::2             20       0  05:36:12  to-vpn-smith
> e B   300::/64            88::2             20       0  05:36:12  to-vpn-smith
> e B   400::/64            88::2             20       0  05:36:12  to-vpn-smith
> e B   500::/64            88::2             20       0  05:36:12  to-vpn-smith
> e B   600::/64            88::2             20       0  05:36:12  to-vpn-smith
> C     800::/64                               0       0  05:36:42  bgp1
> e B   3dde::1/128         88::2             20       0  05:33:52  to-vpn-smith
> e B   3ffe::1/128         88::2             20       0  05:33:52  to-vpn-smith
> e B   4000:0:2:1::/64     88::2             20       0  05:34:52  to-vpn-smith
> e B   8bbc::1/128         88::2             20       0  05:34:52  to-vpn-smith
> e B   8f82::/64           88::2             20       0  05:34:22  to-vpn-smith
> i B   f2cb::/64           800::2           200       0  05:34:40
```

1.40.6.2 MPLS LSP Trace Route

The following example displays a destination IP address of a traced route resulting in the packet going through an MPLS LSP, and includes the label stack along each hop of the LSP:

```
[local]Redback>traceroute 5.5.5.5

se_traceroute to 5.5.5.5 (5.5.5.5), 30 hops max, 40 byte packets
 1  100.1.1.1 (100.1.1.1)   4.749 ms  4.111 ms  3.986 ms
 2  40.1.1.2 (40.1.1.2)   6.321 ms  6.457 ms  5.289 ms
         MplsLabel: 19 MplsExpBits: 0 TTL: 1
         MplsLabel: 786434 MplsExpBits: 0 TTL: 1
 3  60.1.1.1 (60.1.1.1)   3.815 ms  4.159 ms  4.120 ms
         MplsLabel: 786434 MplsExpBits: 0 TTL: 1
 4  5.5.5.5 (5.5.5.5)   5.870 ms  9.108 ms  6.639 ms
```

1.41 traceroute mpls

traceroute mpls {**rsvp** ***lsp-name*** | **ldp** ***ip-address***}[***count***] [**exp** ***exp-bits***] [**range** ***begin-address end- address*** increment] [**reply-mode** {**router-alert** | **udp**}] [**size** ***packet-size***] [**source** ***ip-addr***] [**timeout** ***interval***] [**verbose**]

1.41.1 Purpose

Initiates a Multiprotocol Label Switching (MPLS) trace across a Resource Reservation Protocol (RSVP) label-switched path (LSP) or a Label Distribution Protocol (LDP) LSP.

1.41.2 Command Mode

All modes

1.41.3 Syntax Description

count	Number of MPLS trace routes to send. The range of values is 1 through 4294967295.
rsvp ***lsp-name***	Name of RSVP LSP to be traced.
exp ***exp-bits***	Optional. Amount of experimental (EXP) bits in the MPLS header. The range of values is 0 through 7 bits. The default value is 0 bits.

ldp *ip-address*	IP address of LDP to be traced.
range	Optional. Range of destination addresses to be traced. The range of addresses is 127.0.0.1 to 127.255.255.255.
begin-address	Starting address of range. The default value is 127.0.0.1
end-address	Ending address of range. The default value is 127.0.0.1
increment	Allows the address range specified to increment by the increment value, instead of by 1, when cycling from beginning address to end address.
reply-mode	Optional. The reply mode for the echo request packet. The default is to forward echo replies using IPv4 through UDP.
router-alert	Sends the echo reply as an IP User Datagram Protocol (UDP) packet, with the router alert option preceding the IP header.
udp	Sends the echo request packet as an IPv4 UDP packet, with no router alert option preceding the IP header.
size *packet-size*	Optional. Packet size (in octets) for the traceroute request. The range of values is 100 to 18020.
source *ip-address*	Optional. Source IP address to use for the traceroute. If no IP address is specified, an IP interface address is selected.
any	Uses an unused UDP source port, selected from the reserved range of ports, as the source UDP port.
timeout *interval*	Optional. Interval, in seconds, to wait for each trace sent. The range of values is 1 to 3600 seconds; the default value is 2 seconds.
verbose	Optional. Displays more detailed output.

1.41.4 Default

None

1.41.5 Usage Guidelines

Use the **traceroute mpls** command to initiate a MPLS trace across a RSVP LSP or a LDP LSP. This command is useful for troubleshooting MPLS networks.

Use the *lsp-name* or *ip-address* argument to display information only for the specified LSP, or use any of the available keywords to display LSP information only for the specified keyword.

Table 2 Output Fields for the traceroute mpls Command

Field	Description
Header	Displays the following header information for an MPLS trace: • Number of traces (echoes) sent • Packet size (number of bytes) for each trace • Type of LSP being traced • IP address of the LSP being traced • IP address of the source of the trace. • Interval, in seconds, to wait for each trace sent. • Interval, in seconds, at which traces are sent
Hop	Hop number of the current node, in the path of the LSP, that is replying to the traceroute request.
Pkt	Current number of traceroute requests that have been generated. One traceroute request represents the query of the each node in the path of the LSP. By default, the number of requests transmitted is 1.
Recv Intf	IP address of the receiving interface of the node that has replied to the traceroute request.
Down Intf	IP address of the downstream interface of the node that has replied to the traceroute request.
MTU	Maximum transmission unit of the downstream interface of the node that has replied to the traceroute request.
RC	Errors associated with the query of the hop (each hop receives the traceroute request). Possible errors include Bad Protocol, Bad Request, Bad TLV, No FEC, No Label, No Response, Path Loop, and Path Too Long, If no error exists, the RC column displays the type of node in the LSP: ingress, transit, or egress.
RTT	Round-trip time (RTT) in milliseconds.

Note: By default, most **`show`** commands (in any mode) display information for the current context only or, depending on the command syntax, for all contexts. If you are an administrator for the local context, you can insert the optional **`context`** ***`ctx-name`*** construct, preceding the **`show`** command, to view output for the specified context without entering that context. For more information about using the **`context`** ***`ctx-name`*** construct, see the **`context`** command.

Note: By appending a space followed by the pipe (|) character at the end of a **`show`** command, you can filter the output using a set of modifier keywords and arguments. For more information, see *Modifying Output of show Commands* in *Using the CLI*.

1.41.6 Examples

The following example shows how to initiate an MPLS trace across a LDP LSP. This LSP contains an equal-cost multi-path (ECMP) next-hop entry at the transit node:

```
[local]Redback>traceroute mpls ldp 80.1.1.0/24 range 127.0.0.1 127.0.0.10 1
Sending 5 100-byte MPLS echos to LDP 80.1.1.0/24, source 180.180.180.180,
timeout is 1 second, send interval is 0 msec:
Pkt Hop       Recv Intf       Down Intf        Down Nbr  RTT  MTU  RC
  1   0       N/A             20.1.1.4         20.1.1.4    0    0  Ingress
  1   1       20.1.1.4         60.1.1.1        60.1.1.1    8    0  Transit
                               60.1.2.1        60.1.2.1
  1   2        60.1.1.1             N/A             N/A    8    0  Egress
  2   0             N/A        20.1.1.4        20.1.1.4    8    0  Ingress
  2   1        20.1.1.4        60.1.1.1        60.1.1.1    5    0  Transit
                               60.1.2.1        60.1.2.1
  2   2        60.1.2.1             N/A             N/A    5    0  Egress
  3   0             N/A        20.1.1.4        20.1.1.4    5    0  Ingress
  3   1        20.1.1.4        60.1.1.1        60.1.1.1    5    0  Transit
                               660.1.2.1        60.1.2.1
  3   2        60.1.1.1             N/A             N/A    5    0  Egress
  4   0             N/A        20.1.1.4        20.1.1.4    5    0  Ingress
  4   1        20.1.1.4        60.1.1.1        60.1.1.1    7    0  Transit
                               660.1.2.1        60.1.2.1
  4   2        60.1.2.1             N/A             N/A    6    0  Egress
  5   0             N/A        20.1.1.4        20.1.1.4    6    0  Ingress
  5   1        20.1.1.4        60.1.1.1        60.1.1.1    6    0  Transit
                               660.1.2.1        60.1.2.1
  5   2        60.1.1.1             N/A             N/A    5    0  Egress
  6   0             N/A        20.1.1.4        20.1.1.4    5    0  Ingress
  6   1        20.1.1.4        60.1.1.1        60.1.1.1    4    0  Transit
                               660.1.2.1       60.1.2.1
  6   2        60.1.2.1             N/A             N/A    5    0  Egress
  7   0             N/A        20.1.1.4        20.1.1.4    5    0  Ingress
  7   1        20.1.1.4        60.1.1.1        60.1.1.1    34   0  Transit
                               60.1.2.1        60.1.2.1
  7   2        60.1.1.1             N/A             N/A    7    0  Egress
  8   0             N/A        20.1.1.4        20.1.1.4    7    0  Ingress
  8   1        20.1.1.4        60.1.1.1        60.1.1.1    6    0  Transit
                               60.1.2.1        60.1.2.1
  8   2        60.1.2.1             N/A             N/A    5    0  Egress
  9   0             N/A        20.1.1.4        20.1.1.4    5    0  Ingress
  9   1        20.1.1.4        60.1.1.1        60.1.1.1    4    0  Transit
                               60.1.2.1        60.1.2.1
  9   2        60.1.1.1             N/A             N/A    5    0  Egress
  10  0             N/A        20.1.1.4        20.1.1.4    5    0  Ingress
  10  1        20.1.1.4        60.1.1.1        60.1.1.1    5    0  Transit
                               60.1.2.1        60.1.2.1
  10  2        60.1.2.1             N/A             N/A    6    0  Egress
```

1.42 track

track ***object-name***

no track ***object-name***

1.42.1 Purpose

When used in RSVP router configuration mode, creates a tracking object on an access router.

When used in RSVP LSP configuration mode, enables the LSP to track the interfaces contained in the specified object.

1.42.2 Command Mode

RSVP router configuration

RSVP LSP configuration

1.42.3 Syntax Description

object-name	Name that identifies a tracking object.

1.42.4 Default

None

1.42.5 Usage Guidelines

Use the **track** command in RSVP router configuration mode to create a tracking object on an access router. Use the **track** command in RSVP LSP configuration mode to enable an LSP to track the interfaces contained in a specific object.

Use the **no** form of this command to delete the tracking object from an access router.

1.42.6 Examples

The following example shows how to create a tracking object called `san-jose-1`. In this example, `san-jose-1` contains the interfaces called `red` and `blue`:

```
[local]Redback#configure
[local]Redback(config)#context local
[local]Redback(config-ctx)#router rsvp
[local]Redback(config-rsvp)#track san-jose-1
[local]Redback(config-rsvp-track_obj)#interface red
[local]Redback(config-rsvp-track_obj)#interface blue
```

The following example shows how to enable the LSP called `verde` to track the interfaces in the object called `san-jose-1`:

```
[local]Redback#config
[local]Redback(config)#context local
[local]Redback(config-ctx)#router rsvp
[local]Redback(config-rsvp)#lsp verde
[local]Redback(config-rsvp-lsp)#track san-jose-1
[local]Redback(config-rsvp-lsp)#exit
```

1.43 track-igp-metric

track-igp-metric

no track-igp-metric

1.43.1 Purpose

Enables Label Distribution Protocol (LDP) label-switched paths (LSPs) to inherit the Intermediate System-to-Intermediate System (IS-IS) routing metric for Border Gateway Protocol (BGP) to use when selecting a path.

1.43.2 Command Mode

LDP router configuration

1.43.3 Syntax Description

This command has no keywords or arguments.

1.43.4 Default

By default, inheriting the IS-IS routing metric is disabled.

1.43.5 Usage Guidelines

Use the **track-igp-metric** command to enable LDP LSPs to inherit the IS-IS routing metric for BGP to use when selecting a path.

Use the **no** form of this command to disable LDP LSPs from inheriting the IS-IS metric.

1.43.6 Examples

The following example enables LDP LSPs to inherit the IS-IS routing metric for BGP to use when selecting a path:

```
[local]Redback(config-ctx)#router ldp
[local]Redback(config-ldp)#track-igp-metric
```

1.44 track interface

track interface *if-name* [*context*] [**decrement** *priority*]

no track interface *if-name* [*context*] [**decrement** *priority*]

1.44.1 Purpose

Enables tracking of an interface by the specified VRRP instance.

1.44.2 Command Mode

Interface configuration

1.44.3 Syntax Description

if-name	Interface to be tracked by the VRRP instance specified with the context argument.
context	Optional. VRRP instance for tracking the interface specified with the *if-name* argument.
decrement *priority*	Optional. Specifies the amount by which the priority of the VRRP instance is decremented if the interface goes down.

1.44.4 Default

Interface tracking is disabled.

1.44.5 Usage Guidelines

Use the **track interface** command to enable the tracking of an interface by the specified VRRP instance.

Note: Only one interface can be configured for each command line. To track multiple interfaces, you must enter a separate command line for each interface you want to track.

Use the **no** form of this command to disable interface tracking by a VRRP instance.

1.44.6 Examples

The following example shows how use the local VRRP instance to track the `eth0` interface. In this case, if `eth0` goes down, the local VRRP instance priority is decremented by `100:`

```
[local]Redback#configure
[local]Redback(config)#context local
[local]Redback(config-ctx)#interface eth0
[local]Redback(config-if)#ip address 10.1.1.1/24
[local]Redback(config-if)#vrrp 1 backup
[local]Redback(config-vrrp)#virtual address 10.1.1.2
[local]Redback(config-vrrp)#priority 200
[local]Redback(config-vrrp)#track interface uplink local decrement 100
```

1.45 track network

track network ***ip-addr*** [**decrement** *priority*]

{**no** | **default**} **track network** ***ip-addr***]

1.45.1 Purpose

Enables the backup VRRP router to track a specified network through the IP network prefix of that network.

1.45.2 Command Mode

VRRP configuration mode.

1.45.3 Syntax Description

ip-addr	IP address of the network to be tracked by the VRRP instance.
decrement ***priority***	Optional. Specifies the amount by which the priority of the backup VRRP router is decremented if there is a failure in the IP network (for example, if the backup VRRP router loses connectivity to a neighbor network). The default priority value is 10.

1.45.4 Default

The default decrement priority value is 10.

1.45.5 Usage Guidelines

Use the **`track network`** command to enable the backup VRRP router to track a specified network through the IP network prefix of that network.

Note: When a network regains reachability after a switchover, the VRRP router reverts to its original priority.

A change in router priority can cause a master VRRP router to become a backup VRRP router, and a backup router to become a master router.

Note: A single backup VRRP router can track up to 10 IP network prefixes.

Use the **`no`** or **`default`** form of this command to disable IP network prefix tracking for all IP prefixes or for a specific IP address.

1.45.6 Examples

The following example configures the backup VRRP router `1` to track the network with the IP network prefix `192.168.100.1/24`. In this example, the priority of VRRP router`1` is decremented by `100` if the network with the IP network prefix `192.168.100.1/24` cannot be reached:

```
[local]SE1(config)#context local
[local]SE1(config-ctx)#interface vlan1
[local]SE1(config-if)#vrrp 1 backup
[local]SE1(config-vrrp)#track network 192.168.100.1/24 decrement 100
```

1.46 track spanning-tree

`track spanning-tree` ***`master-bridge-name context-name`***

`no track spanning-tree`

1.46.1 Purpose

Configures the current bridge to track the Rapid Spanning Tree Protocol (RSTP) master bridge specified in this command.

1.46.2 Command Mode

Bridge configuration

1.46.3 Syntax Description

master-bridge	Name of the RSTP master bridge.
context	Name of the context that contains the master bridge.

1.46.4 Default

RSTP tracking is disabled.

1.46.5 Usage Guidelines

Use the **track spanning-tree** command to configure the current bridge to track the RSTP master bridge specified.

Bridges that are not running RSTP and are enabled for tracking by the `track spanning-tree` command are called *client* bridges. The state of all client bridge circuits on the same port as a master bridge circuit follow the state of the RSTP master. When the state of the circuit controlled by the master bridge changes to blocking, forwarding, or flushing, all circuits on the same port of the tracking client bridges change to the same state.

See the RSTP tracking requirements and restrictions listed in *Configuring Bridging*.

Use the **no** form of this command to disable tracking of the RSTP master bridge.

1.46.6 Examples

The following example shows the configuration of `blue` which is an RSTP master bridge and `green` which is one of its non-RSTP client bridges:

```
[local]Redback#configure
[local]Redback(config)#context local
[local]Redback(config-ctx)#bridge blue
[local]Redback(config-bridge)#spanning-tree
[local]Redback(config-bridge-stp)#master
[local]Redback(config-bridge-stp)#end
!
[local]Redback#configure
[local]Redback(config)#context local
[local]Redback(config-ctx)#bridge green
[local]Redback(config-bridge)#track spanning-tree blue local
[local]Redback(config-bridge)#end
```

1.47 track (vrrp)

track {**vrrp** ***vrrp-id vrrp-if-name vrrp-if-context***} | **none**

no track {**vrrp** | **none**}

1.47.1 Purpose

Changes a static dot1q PVC into a track PVC, or disables following a VRRP session on the current circuit.

1.47.2 Command Mode

dot1q PVC configuration

1.47.3 Syntax Description

vrrp ***vrrp-id vrrp-if-name vrrp-if-context***	Specifies the VRRP session and interface that the track PVC follows: • ***vrrp-id*** - Virtual router ID. The range of values is 1 to 255. • ***vrrp-if-name*** - Virtual router interface name. • ***vrrp-if-context*** - Name of the context containing the virtual router interface.
none	Disables following a VRRP session on the current circuit.

1.47.4 Default

The current circuit is not a track PVC.

1.47.5 Usage Guidelines

Use the **track vrrp** command to change a static dot1q PVC into a track PVC, or to disable following a VRRP session on the current circuit.

A track PVC is a PVC that tracks the state of a VRRP interface. The track PVC carries traffic when it is active and no traffic when it is in backup state. Only one track option can be defined on a PVC.

When a static 802.1Q PVC is configured to track the state of a VRRP interface, all of its child circuits inherit this functionality and also track the state of that same VRRP interface.

Use the **no track vrrp** form of this command to change the current PVC into a static PVC. Use the **no track none** form of this command to enable following a VRRP session on the current circuit.

1.47.6 Examples

The following example shows how use the **track vrrp** command:

```
[local]Redback#configure
[local]Redback#port ethernet 5/1
[local]Redback(config)#dot1q pvc 108 encapsulation pppoe
[local]Redback(config-dot1q-pvc)#track vrrp 7 one pctx
[local]Redback(config-dot1q-pvc)#bind authentication chap context
pctx maximum 2
```

1.48 traffic-engineering

traffic-engineering [*level-1* | *level-2* | *level-1-2*]

no traffic-engineering

1.48.1 Purpose

Enables Multiprotocol Label Switching (MPLS) traffic engineering within Intermediate System-to-Intermediate System (IS-IS) routing.

1.48.2 Command Mode

IS-IS router configuration

1.48.3 Syntax Description

level-1	Optional. Traffic engineering for IS-IS level 1 routing only.
level-2	Optional. Traffic engineering for IS-IS level 2 routing only.
level-1-2	Optional. Traffic engineering for IS-IS both routing levels.

1.48.4 Default

MPLS traffic engineering is disabled.

1.48.5 Usage Guidelines

Use the **`traffic-engineering`** command to enable MPLS traffic engineering within IS-IS routing. Enabling traffic engineering allows IS-IS link-state protocol data units (LSPs) to carry traffic engineering information on IS-IS interfaces. Traffic engineering information includes link IP addresses, link bandwidth and link administrative colors.

Traffic engineering can be enabled on either IS-IS level 1, level 2, or both level 1 and level 2 routing.

Note:

- Resource Reservation Protocol (RSVP) must be configured on the interface for IS-IS traffic engineering information to be included in its LSP for the link.
- An IS-IS metric style of wide or transition must be used for traffic engineering to take effect.
- The global **`router-id`** command in context configuration mode must be configured for the IS-IS LSP to carry the specified IP address of the router ID interface.

Use the **`show isis database extensive`** command to see the traffic engineering information for the IS-IS link in the LSPs, and the **`show isis interface detail`** command to see if the interface has traffic engineering information for the routing level.

Use the **`no`** form of this command to disable MPLS traffic engineering within IS-IS routing.

1.48.6 Examples

The following example displays that IS-IS traffic engineering is enabled for IS-IS `level-2` routing:

```
[local]Redback(config-ctx)#router isis ip-backbone
[local]Redback(config-isis)#traffic-engineering level-2
```

1.49 traffic-index accounting

`traffic-index accounting`

{**`no`** | **`default`**} **`traffic-index accounting`**

1.49.1 Purpose

Enables Border Gateway Protocol (BGP) attribute-based accounting on an interface.

1.49.2 Command Mode

Interface configuration

1.49.3 Syntax Description

This command has no keywords or arguments.

1.49.4 Default

BGP attribute-based accounting is disabled.

1.49.5 Usage Guidelines

Use the **`traffic-index accounting`** command to enable BGP attribute-based accounting on an interface.

Per index counters for interfaces with BGP attribute-based accounting enabled are maintained for BGP routes assigned a traffic index. The byte and packet counters for a traffic index are incremented based on the route traversed by IP traffic received on the ingress interface. For more information, see the **`set traffic-index`** and **`table-map`** commands.

Use the **`no`** or **`default`** form of this command to disable BGP attribute-based accounting on an interface.

1.49.6 Examples

The following example enables BGP policy accounting on the interface, `value-added`:

```
[local]Redback(config)#interface value-added
[local]Redback(config-if)#ip address 10.200.1.1/30
[local]Redback(config-if)#traffic-index accounting
```

1.50 transfer-interval

`transfer-interval` ***`minutes`***

`default transfer-interval`

1.50.1 Purpose

Specifies the interval after which bulkstats data for this policy is uploaded to a remote file server.

1.50.2 Command Mode

Bulkstats configuration

1.50.3 Syntax Description

minutes	Transfer interval in minutes. The range of values is 1 to 1,440 minutes (24 hours); the default value is 60 minutes.

1.50.4 Default

The bulkstats transfer interval is 60 minutes.

1.50.5 Usage Guidelines

Use the **transfer-interval** command to specify the interval after which bulkstats data for this policy is uploaded to a remote file server. Use the **bulkstats force transfer** command in exec mode to force an immediate transfer for this policy. For information on the **bulkstats force transfer** command, see the *Command List*.

Use the **default** form of this command to return the transfer interval to 60 minutes.

1.50.6 Examples

The following example specifies that bulkstats data is transferred to a remote file server every 180 minutes:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#bulkstats policy bulk
[local]Redback(config-bulkstats)#transfer-interval 180
```

1.51 transmit-delay

transmit-delay *delay*

{**no** | **default**} **transmit-delay**

1.51.1 Purpose

Sets a delay value, increasing the age of link-state advertisements (LSAs) sent over the specified interface, sham link, or virtual link.

1.51.2 Command Mode

- OSPF interface configuration
- OSPF sham link configuration
- OSPF virtual link configuration
- OSPF3 interface configuration

1.51.3 Syntax Description

delay

Delay, in seconds. The range of values is 1 to 65,535; the default value is 1 second.

1.51.4 Default

No delay value is set. When set, the delay value is one second.

1.51.5 Usage Guidelines

Use the **transmit-delay** command to set a delay value, increasing the age of LSAs sent over the specified interface, sham link, or virtual link.

Before a link-state update packet is advertised, the age of the LSAs in the packet must be increased by a value proportionate to the speed of the interface, sham link, or virtual link; for example, on a very slow interface, sham link, or virtual link, you might set the transmit delay to two seconds to ensure that you do not receive an LSA that is less recent than the copy in the router's link-state database.

Use the **no** or **default** form of this command return the delay value to its default setting.

1.51.6 Examples

The following example sets an Open Shortest Path First (OSPF) interface transmit delay to 3 seconds:

```
[local]Redback(config-ospf-if)#transmit-delay 3
```

1.52 transmit-hold count

transmit-hold count *sec*

{**no** | **default**} **transmit-hold count**

1.52.1 Purpose

Sets the transmit hold count.

1.52.2 Command Mode

Spanning-tree configuration

1.52.3 Syntax Description

sec	Transmit hold count in seconds. The range of values is 1 to 10 seconds.

1.52.4 Default

6 seconds

1.52.5 Usage Guidelines

Use the **transmit-hold count** command to set the transmit hold count; that is, the minimum amount of time between the transmission of BPDUs. This command applies when the current bridge is the root bridge.

1.52.6 Examples

An example of setting the transmit hold count time follows:

```
[local]Redback(config-bridge-stp)#transmit-hold count 4
```

1.53 transport address

transport address *ip-addr*

1.53.1 Purpose

Configures the transport address advertised in Label Distribution Protocol (LDP) Hello messages.

1.53.2 Command Mode

LDP router configuration

1.53.3 Syntax Description

ip-addr	IP address to be advertised as the transport address. The IP address must be reachable.

1.53.4 Default

The label-switched router (LSR) router ID is used as the transport address.

1.53.5 Usage Guidelines

Use the **transport address** command to configure the transport address advertised in LDP Hello messages. Transport addresses are advertised in LDP Hello messages and are exchanged among LDP neighbors. LDP uses the local transport address as the source, and the received transport address as the destination when trying to establish a Transmission Control Protocol (TCP) connection to a neighbor. Therefore, transport addresses must be reachable. LDP also uses transport addresses to determine which of the two LSRs should perform active open.

If a transport address is not explicitly configured, the LSR router ID is used as the transport address. In this case, the router ID must be reachable; however, if a transport address is explicitly configured, then the specified value is used. In this case, the router ID is not required to be reachable.

1.53.6 Examples

The following example configures a transport address of `20.1.1.1`:

```
[local]Redback(config-ctx)#router ldp
[local]Redback(config-ldp)#transport address 20.1.1.1
```

1.54 transport protocol

transport protocol [**udp**]

```
no transport protocol [udp]
```

1.54.1 Purpose

Configures the transport protocol used to export the flow records.

1.54.2 Command Mode

Flow collector configuration

1.54.3 Syntax Description

udp	Configures the transport protocol for the flow records to be UDP.

1.54.4 Default

The default transport protocol for the flow records is UDP.

1.54.5 Usage Guidelines

Use the **transport protocol udp** command to configure the transport protocol used to export the flow records.

Note: In this release, UDP is the only supported transport protocol for the flow records.

Use the **no** form of this command to use the default transport protocol for the flow records (UDP).

1.54.6 Examples

The following example shows how to configure the SmartEdge router to use UDP for transporting flow records to an external collector called `c1`:

```
[local]Redback#configure
[local]Redback(config)#context corp2
[local]Redback(config-ctx)#flow collector c1
[local]Redback(config-flow-collector)#transport protocol udp
```

1.55 transport protocol (NAT)

1.55.1 Command Mode

NAT logging profile configuration

1.55.2 Syntax Description

`udp`	Configures the transport protocol for the flow records to be UDP.

1.55.3 Default

The default transport protocol for the flow records is UDP.

1.55.4 Usage Guidelines

Use the **`transport protocol udp`** command to configure the transport protocol used to export the flow records.

Note: UDP is the only supported transport protocol for the flow records.

Use the **`no`** form of this command to use the default transport protocol for the flow records (UDP).

For information about how to configure a NAT logging profile, see the*nat logging-profile* command and *Configure an Enhanced NAT Policy with Logging and Paired Mode*.

1.55.5 Examples

The following example show how to configure the SmartEdge router to use UDP for transporting flow records for a NAT logging profile.

```
[local]Redback#configure
Enter configuration commands, one per line, 'end' to exit
[local]Redback(config)#context nat-context
[[local]Redback(config-ctx)#nat logging-profile nat-log-profile
[local]Redback(config-nat-profile)#transport protocol udp
```

1.56 trap cesmib

`trap cesmib outage | excessive-packet-loss-rate`

```
no trap cesmib outage | excessive-packet-loss-rate
```

1.56.1 Purpose

Enables and disables the CES outage or excessive-packet-loss-rate trap.

1.56.2 Command Mode

SNMP Server Config Mode.

1.56.3 Syntax Description

`outage` Enables the CES outage trap.

`excessive-packet-loss-rate` Enables the CES excessive-packet-loss-rate trap.

1.56.4 Default

CES MIB traps are disabled.

1.56.5 Usage Guidelines

None..

1.56.6 Examples

The following example shows how to enable the outage trap and disable the excessive-packet-loss-rate trap::

```
[local]Redback(config-snmp-server)#trap cesmib outage
[local]Redback(config-snmp-server)#no trap cesmib excessive-packet-loss-rate
```

1.57 traps (SNMP server configuration)

```
traps {ifmib [ encaps | ip] | ds1mib | ds3mib | l2tpmib |
l2vpnmib | mpls13vpn interval seconds| nemib [ exclusive |
non-exclusive] | ssemib | vrrpmib}
```

```
no traps {ifmib [ encaps | ip] | ds1mib | ds3mib | l2tpmib
| l2vpnmib | mpls13vpn | nemib [ exclusive | non-exclusive]
| ssemib | vrrpmib}
```

1.57.1 Purpose

Enables Simple Network Management Protocol (SNMP) notifications for events described in the selected Management Information Bases (MIBs).

1.57.2 Command Mode

SNMP server configuration

1.57.3 Syntax Description

`ifmib encaps`	Enables linkUp and linkDown notifications as described in the Interfaces MIB (IF-MIB) for the Cisco High-Level Data Link Control (HDLC), Point-to-Point Protocol (PPP), and Frame Relay interface encapsulation layers.
`ifmib ip`	Enables linkUp and linkDown notifications as described in the IF-MIB for the IP layer.
`ds1mib`	Enables event notifications as described in the DS1-MIB for DS-1 ports.
`l2tpmib`	Enables event notifications as described in the Layer 2 Tunneling Protocol MIB (L2TP-MIB) for L2TP tunnels.
`l2vpnmib`	Enables event notification for pseudowire operational state changes as described in RBN-L2VPN-MIB.
`mplsl3vpn interval` *seconds*	Enables event notifications as described in MPLS-L3VPN-STD-MIB. When the value of mplsL3VpnVrfConfRteMxThrshTime is 0, the SNMP agent does not send the mplsL3VpnVrfNumVrfRouteMaxThreshExceeded notification again unless the number of routes drops below the mplsL3VpnVrfConfHighRteThresh threshold value. The agent will then attempt to exceed the threshold. The value range is 0 to 42929672 seconds.
`nemib exclusive`	Enables event notifications as described in the RBN-NOTIFY-ENHANCE-MIB while disabling the corresponding notifications in all other MIBs.
`nemib non-exclusive`	Enables event notifications as described in the RBN-NOTIFY-ENHANCE-MIB in addition to the corresponding notifications in all other MIBs.
`ssemib`	Enables SSE notifications for disk operation errors (I/O errors) on any of the disks on any of the SSE cards.
`vrrpmib`	Enables event notifications as described in the VRRP-MIB in addition to the corresponding notifications in all other MIBs.

1.57.4 Default

Notification of all conditions is disabled globally for all encapsulation layers, IP layers, and L2TP tunnels.

1.57.5 Usage Guidelines

Use the **traps** command to enable SNMP notifications for events described in the selected MIBs.

You can enter this command multiple times to enable notifications for encapsulation layers, IP layers, or L2TP tunnels.

The settings for the **traps** command are global; however, with the **ifmib encaps** construct, the global **traps** setting is overridden locally by setting the **traps** command (in DS-0 group configuration mode) for that specific DS-0 channel group. For more information on the **traps** command in DS-0 group configuration mode.

Note: By default, only IF-MIB physical ports generate linkUp and linkDown notifications.

When entered with the **nemib non-exclusive** and **nemib exclusive** constructs, the notifications in the RBN-NOTIFY-ENHANCE-MIB provide more information than the corresponding notifications in standard MIBs, such as IF-MIB, DS1-MIB, DS3-MIB, ENTITY-MIB, and some earlier Enterprise MIBs, such as RBN-CARDMON-MIB.

The **nemib exclusive** construct disables the following standard traps:

- linkDown and linkUp
- dsx1LineStatusChange and dsx3LineStatusChange
- entConfigChange and rbnCardAlarm

If you specify neither the **nemib non-exclusive** nor **nemib exclusive** construct, the notifications in the RBN-NOTIFY-ENHANCE-MIB are disabled by default.

Use the **no** form of this command to disable notifications of up and down conditions for encapsulation layers, IP layers, or L2TP tunnels and to use the unenhanced versions of the traps.

1.57.6 Examples

The following example enables notifications for Cisco HDLC, PPP, and Frame Relay encapsulation layers, IP layers, and L2TP tunnels:

```
[local]Redback(config)#snmp server enhance ifmib
[local]Redback(config-snmp-server)#traps ifmib encaps
[local]Redback(config-snmp-server)#traps ifmib ip
[local]Redback(config-snmp-server)#traps l2tpmib
[local]Redback(config-snmp-server)#traps nemib exclusive
```

1.58 traps (DS-0 group configuration)

traps ifmib {**enabled** | **disabled**}

{**no** | **default**} **traps ifmib**

1.58.1 Purpose

Enables linkUp and linkDown notifications for Cisco HDLC, Point-to-Point Protocol (PPP), and Frame Relay encapsulation layers (IF-MIB encapsulation layers) on the DS-0 channel group.

1.58.2 Command Mode

DS-0 group configuration

1.58.3 Syntax Description

ifmib enabled	Enables notifications for encapsulation layers on the DS-0 channel group.
ifmib disabled	Disables notifications for encapsulation layers on the DS-0 channel group.

1.58.4 Default

If this command is not entered, notification of up and down conditions is enabled or disabled by the **traps** command in (SNMP server configuration mode).

1.58.5 Usage Guidelines

Use the **traps** command to enable linkUp and linkDown notifications locally for Cisco HDLC, PPP, and Frame Relay encapsulation layers on the DS-0 channel group. This command overrides, for this DS-0 channel group, any global specification for encapsulation layers you have specified with the **traps** command (in SNMP server configuration mode).

Table 3 lists the combinations of global and local settings and the resulting notifications for encapsulation layers.

Table 3 Command Settings and Encapsulation Layer Notifications

Global	Local	Encapsulation Layer Notifications
None	None or default	None

Table 3 Command Settings and Encapsulation Layer Notifications

Global	Local	Encapsulation Layer Notifications
	`enabled`	Locally enabled for this DS-0 channel group
	`disabled`	Locally disabled for this DS-0 channel group
`ifmib encaps`	None or default	Globally enabled for all clear-channel or channelized ports and channels, including this DS-0 channel group
	`enabled`	Both globally and locally enabled for this DS-0 channel group
	`disabled`	Locally disabled for this DS-0 channel group

Use the **no** or **default** form of this command to disable encapsulation layer notifications locally; as a result, encapsulation layer notifications are enabled or disabled globally as specified with the **traps** command in (SNMP server configuration mode).

1.58.6 Examples

The following example shows how to enable encapsulation notifications globally and disable them locally for the DS-0 channel group on port `1` of a channelized E1 traffic card:

```
[local]Redback(config)#snmp server enhance ifmib
[local]Redback(config-snmp-server)#traps ifmib encaps
[local]Redback(config-snmp-server)#exit
[local]Redback(config)#port ds0s 5/1:7
[local]Redback(config-ds0-group)#traps ifmib disabled
```

1.59 trunk

trunk

{**no** | **default**} **trunk**

1.59.1 Purpose

Specifies that any circuit to which this profile is assigned is a trunk circuit.

1.59.2 Command Mode

Bridge profile configuration

1.59.3 Syntax Description

This command has no keywords or arguments.

1.59.4 Default

Any circuit to which this profile is assigned is a tributary circuit.

1.59.5 Usage Guidelines

Use the **`trunk`** command to specify that any circuit to which this profile is assigned is a trunk circuit.

Use the **`no`** or **`default`** form of this command to specify the default condition.

Tributary circuits face subscribers. Trunk circuits face service providers. In other words, tributary circuits are access facing, while trunk circuits are network facing.

1.59.6 Examples

The following example shows how to specify that the profile be a trunk profile:

```
[local]Redback(config)#bridge profile prof-isp1
[local]Redback(config-bridge-profile)#trunk
```

1.60 trunk-control

`trunk-control`

`no trunk-control`

1.60.1 Purpose

Enables control of the state of the T1/E1 trunk on a CESoPSN interworking function (IWF).

1.60.2 Command Mode

CESoPSN Config Mode.

1.60.3 Syntax Description

None.

1.60.4 Default

Trunk control is disabled.

1.60.5 Usage Guidelines

CES IWF will enforce the RFC 5086 requirement that at most one IWF on the outgoing trunk will be set to “true”.

The AIS or RDI conditions (as described in the *Egress and Ingress IWF Behavior* tables in *Configuring CESoPSN Pseudowires*) are enforced.

AIS – Imposes AIS condition on T1 or E1 trunk.

RDI – Imposes RDI condition on T1 or E1 trunk

System rejects the user configuration if trunk-control is already enabled on another channel group under same T1/E1 channel.

1.60.6 Examples

The following example shows how to enable trunk control:

```
[local]Redback(config)#port ds0s 1/1:1:1:1
Redback(config-ds0-ces)#timeslot 16
Redback(config-ds0-ces)#l2vpn local
Redback(config-ds0-ces)#cesopsn
Redback(config-ds0-cesopsn)#end-to-end-delay latency 4 jitter 160 outage-criteria 1 10
Redback(config-ds0-cesopsn)#idle-pattern 0x3f
Redback(config-ds0-cesopsn)#trunk control
```

1.61 ttl

ttl ***value***

{**no** | **default**} **ttl** ***value***

1.61.1 Purpose

Configures the time-to-live (TTL) value for a dynamically verified static routing (DVSR) profile.

1.61.2 Command Mode

DVSR profile configuration

1.61.3 Syntax Description

value	TTL value. The range of values is 1 to 255; the default value is 5.

1.61.4 Default

The default TTL value is 5.

1.61.5 Usage Guidelines

Use the **ttl** command to configure the TTL value for a DVSR profile. The TTL value controls the maximum number of hops the verification packet can traverse; for example, if there are multiple paths to reach the verify host address, you must restrict the verification packet to the shorter paths to be considered a successful verification.

Use the **no** form of this command to delete the TTL value from a DVSR profile.

Use the **default** form of this command to configure a DVSR profile to use the default TTL value of 5.

1.61.6 Examples

The following example defines a DVSR profile using a TTL value of `2`:

```
[local]Redback(config-ctx)#dvsr-profile abc-webfarm
[local]Redback(config-dvsr)#ttl 2
```

1.62 tunnel

tunnel *tunl-type tunl-name*

no tunnel *tunl-type tunl-name*

1.62.1 Purpose

Creates or selects a tunnel and accesses tunnel configuration mode.

1.62.2 Command Mode

Global configuration

1.62.3 Syntax Description

tunl-type	Type of tunnel to be created or selected, according to one of the following keywords: • **gre**—Creates or selects a Generic Routing Encapsulation (GRE) tunnel. • **ipip**—Creates or selects an IP-in-IP tunnel. • **ipv6v4-auto**—Creates or selects an overlay tunnel for which the system assigns the remote IP address. • **ipv6v4-manual**—Creates or selects an overlay tunnel for which you must assign the remote IP address.
tunl-name	Text string of up to 39 characters. This name must be unique from all other tunnels.

1.62.4 Default

No tunnels are created.

1.62.5 Usage Guidelines

Use the **tunnel** command to create or select a tunnel and enter tunnel configuration mode.

If the remote end point of a GRE tunnel is a SmartEdge router, you must also create the GRE tunnel configuration on the remote SmartEdge router.

Note: If the traffic receiving end of a GRE tunnel is a SmartEdge router, you must configure the router with a GRE tunnel whose local endpoint is the traffic receiving endpoint of the tunnel. This configuration is required in order to respond to GRE keepalives sent by the traffic sending end of the tunnel.

You can create multiple tunnels, but usually only one tunnel between sites.

Use the **no** form of this command to delete the specified tunnel and any associated parameters that you have specified in tunnel configuration mode. The keywords are not available for the **no** form of this command.

1.62.6 Examples

The following example shows how to create a GRE tunnel, `HartfordTnl`:

```
[local]Redback(config)#tunnel gre HartfordTnl
[local]Redback(config-tunnel)#
```

1.63 tunnel-auth key

tunnel-auth key ***key***

{**no** | **default**} **tunnel-auth key**

1.63.1 Purpose

Specifies a Layer 2 Tunneling Protocol (L2TP) key to be used by a peer to encrypt and decrypt information sent on the control channel.

1.63.2 Command Mode

L2TP peer configuration

1.63.3 Syntax Description

key	Key to be used by a peer to encrypt and decrypt information sent on the control channel. The key can be any alphanumeric text string of any length. Optional with the **no** form of this command.

1.63.4 Default

No password is created.

1.63.5 Usage Guidelines

Use the **tunnel-auth key** command to specify the key to be used by a peer to encrypt and decrypt information sent on the control channel.

The ***key*** argument is an alphanumeric string used for the peer password.

Use the **no** or **default** form of this command to delete any previously established primary password.

1.63.6 Examples

The following example establishes `6dkq7pv` as the password for peer `peer1`:

```
[local]Redback(config-ctx)#l2tp-peer name peer1 media udp remote
dns yellow
[local]Redback(config-l2tp)#tunnel-auth key 6dkq7pv
```

1.64 tunnel domain

tunnel domain

no tunnel domain

1.64.1 Purpose

Enables the dynamic assignment of a subscriber's Point-to-Point Protocol (PPP) session to a Layer 2 Tunneling Protocol (L2TP) peer that has the same domain alias as the subscriber's domain alias.

1.64.2 Command Mode

Subscriber configuration

1.64.3 Syntax Description

This command has no keywords or arguments.

1.64.4 Default

Dynamic assignment is disabled; subscriber PPP sessions are terminated and routed rather than tunneled.

1.64.5 Usage Guidelines

Use the **tunnel domain** command to enable the dynamic assignment of a subscriber's PPP session to an L2TP peer that has the same domain alias as the subscriber's domain alias (the **@*ctx-name*** portion of the structured subscriber name). This domain alias is also a domain alias for the context in which both are configured. You create domain aliases for a context using the **domain** command in context configuration mode.

To allow the subscriber PPP sessions to be tunneled, you must have configured the PPP for the subscriber circuit.

This command and the **tunnel name** command in subscriber configuration mode are mutually exclusive.

You can configure multiple subscribers with dynamic peer assignment if you enter this command for the default or named subscriber profile instead of individual subscriber records.

PPP sessions are tunneled in the upstream direction to the remote peer.

Use the **no** form of this command to disable dynamic assignment for a subscriber.

1.64.6 Examples

The following example shows how to configure the default subscriber profile to cause PPP sessions to be mapped to the tunnel that has the same name as the user's domain name:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#subscriber default
[local]Redback(config-sub)#tunnel domain
```

1.65 tunnel ldp-path

tunnel ldp-path

no tunnel ldp-path

1.65.1 Purpose

Enables L2VPN-to-LDP LSP mapping in an L2VPN profile.

1.65.2 Command Mode

- L2VPN profile peer configuration
- L2VPN profile backup peer configuration

1.65.3 Syntax Description

This command has no keywords or arguments.

1.65.4 Default

L2VPN-to-LDP LSP mapping is not configured in an L2VPN profile.

1.65.5 Usage Guidelines

Use the **tunnel ldp-path** command to configure L2VPN-to-LDP LSP mapping in an L2VPN profile. All L2VPN circuits attached to the specified profile are mapped to an LDP LSP.

Use the **no** form of this command to disable L2VPN-to-LDP LSP mapping in an L2VPN profile.

1.65.6 Examples

The following example shows how to enable L2VPN-to-LDP LSP mapping in an L2VPN profile:

```
[local]Redback(config)#l2vpn profile pr1
[local]Redback(config-l2vpn-xc-profile)#peer 111.111.111.111
[local]Redback(config-l2vpn-xc-profile-peer)#tunnel ldp-path
```

The following example shows how to configure a backup tunnel in an L2VPN profile. Any backup L2VPN cross-connects that have the profile called `pro` attached are mapped to an LDP LSP:

```
[local]Redback(config)#l2vpn profile prof2
[local]Redback(config-l2vpn-xc-profile)#backup peer 100.100.100.2
[local]Redback(config-l2vpn-xc-profile-peer)#tunnel ldp-path
```

1.66 tunnel lsp

tunnel lsp ***lsp-name***

no tunnel lsp ***lsp-name***

1.66.1 Purpose

Specifies a particular Resource Reservation Protocol (RSVP) tunnel for carrying traffic exiting an label-switched path (LSP).

1.66.2 Command Mode

- L2VPN profile peer configuration
- L2VPN profile backup peer configuration
- VPLS neighbor profile configuration

1.66.3 Syntax Description

lsp-name	Name of the RSVP tunnel on which to carry traffic exiting an L2VPN XC.

1.66.4 Default

L2VPN-to-RSVP tunnel mapping is not configured in an L2VPN profile.

1.66.5 Usage Guidelines

Use the **tunnel lsp** command to specify a particular RSVP tunnel for carrying traffic exiting an L2VPN XC. Any L2VPN circuits that are associated with this L2VPN profile inherit this configuration.

Use the **no** form of this command to remove L2VPN-to-RSVP tunnel mapping configuration from an L2VPN profile.

1.66.6 Examples

The following example shows how to specify an RSVP tunnel called `lsp1A` for carrying traffic exiting an L2VPN XC. All L2VPN circuits that have the L2VPN profile called `pr1` attached inherit this configuration:

```
[local]Redback(config)#l2vpn profile pr1
[local]Redback(config-l2vpn-xc-profile)peer 111.111.111.111
[local]Redback(config-l2vpn-xc-profile-peer) tunnel lsp lsp1A
```

The following example shows how to configure a backup tunnel in an L2VPN profile. Any backup L2VPN XCs that have the profile called `prof1` attached are mapped to the RSVP tunnel called `lsp-to-target2:`

```
[local]Redback(config)#l2vpn profile prof1
[local]Redback(config-l2vpn-xc-profile)#backup peer 100.100.100.2
[local]Redback(config-l2vpn-xc-profile-peer)#tunnel lsp lsp-to-target2
```

1.67 tunnel name

tunnel name ***tunl-name***

no tunnel name ***tunl-name***

1.67.1 Purpose

Statically assigns the subscriber's Point-to-Point Protocol (PPP) session to a specified Layer 2 Tunneling Protocol (L2TP) peer or group of L2TP peers.

1.67.2 Command Mode

Subscriber configuration

1.67.3 Syntax Description

tunl-name	Name of the peer or L2TP group of peers to which the subscriber is mapped.

1.67.4 Default

A PPP session is terminated rather than tunneled.

1.67.5 Usage Guidelines

Use the **tunnel name** command to statically assign the subscriber's PPP session to a specific L2TP peer or group of peers. You can use a peer name or the domain alias for the peer name, a group name, or a domain alias for the group name as the ***tunl-name*** argument, which is included in the subscriber record.

Note: This command and the **tunnel domain** command in subscriber configuration mode are mutually exclusive.

Use the **no** form of this command to remove the peer or peer group name or alias from the subscriber record.

1.67.6 Examples

The following example shows how to ensure that the subscriber uses the tunnel `freds-corp.com:`

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#subscriber name fred
[local]Redback(config-sub)#tunnel name freds-corp.com
```

1.68 tunnel-shortcut

tunnel-shortcut

no tunnel-shortcut

1.68.1 Purpose

Enables Resource Reservation Protocol (RSVP) label switched paths (LSPs) to carry label distribution protocol (LDP) control and data traffic.

1.68.2 Command Mode

- LDP router configuration
- RSVP LSP configuration
- RSVP router configuration

1.68.3 Syntax Description

This command has no keywords or arguments.

1.68.4 Default

RSVP over LDP is disabled.

1.68.5 Usage Guidelines

Use the **tunnel-shortcut** command to enable RSVP LSPs to carry LDP control and data traffic.

An RSVP LSP is enabled as a tunnel shortcut when the **tunnel-shortcut** command is:

- Enabled in RSVP router configuration mode and also in RSVP LSP configuration mode
- Enabled in RSVP router configuration mode, but not configured in RSVP LSP configuration mode
- Disabled or not configured in RSVP router configuration mode, but enabled in RSVP LSP configuration mode

Use the **no** form of this command in RSVP router configuration or RSVP LSP configuration mode to disable RSVP LSPs to carry LDP control and data traffic.

An LDP targeted neighbor is enabled as a tunnel shortcut when the **tunnel-shortcut** command is:

- Enabled in LDP router configuration mode and also for the specific targeted neighbor (**neighbor** ***address*** **targeted tunnel-shortcut**)
- Enabled in LDP router configuration mode, but not configured for the targeted neighbor
- Disabled or not configured in LDP router configuration mode, but enabled for the targeted neighbor

Use the **default** form of this command in LDP configuration mode to disable RSVP LSPs to carry LDP control and data traffic.

To enable LDP over RSVP, the following configurations must be enabled:

- The remote LDP neighbors must be configured and extended LDP discovery of the specified neighbor must be enabled. Use the **neighbor targeted** command (in LDP router configuration mode) to configure LDP peers as targeted neighbors.

- The **tunnel-shortcut** command must be configured in LDP router configuration mode for all targeted LDP neighbors or for a specific targeted LDP neighbor (**neighbor** ***address*** **targeted tunnel-shortcut**).

- The **tunnel-shortcut** command must all be configured in either RSVP router configuration mode for all RSVP LSPs or in RSVP LSP configuration mode for a specific RSVP LSP.

- You must use Open Shortest Path First (OSPF) as the Interior Gateway Protocol (IGP) and configure the **mpls tunnel-shortcut** command in OSPF router configuration mode. LDP over RSVP is not supported when Intermediate System-to-Intermediate System (IS-IS) is the IGP.

- The IP address specified for the targeted LDP neighbor and RSVP LSP egress label-switched router (LSR) must be the same. These addresses must be loopback addresses (/32). Use the **egress** command (in RSVP LSP configuration mode) to specify the IP address of the egress LSR; for more information about the **egress** command, see *Configuring Basic IP Routing*.

Note: If the RSVP LSP configured as a tunnel shortcut is also configured with next-hop fast reroute (NFRR) for link and node protection, or end-to-end backup(s), LDP traffic will also be protected. For more information, see *Configuring MPLS*.

1.68.6 Examples

The following example enables RSVP LSPs to carry LDP control and data traffic:

```
[local]Redback(config-ctx)#router ospf
[local]Redback(config-ospf)#mpls tunnel-shortcut
[local]Redback(config-ospf)#exit
[local]Redback(config-ctx)#router rsvp
[local]Redback(config-rsvp)#tunnel-shortcut
[local]Redback(config-rsvp)#exit
[local]Redback(config-ctx)#router ldp
[local]Redback(config-ldp)#neighbor 10.1.1.1 targeted tunnel-shortcut
```

1.69 tunnel-type

tunnel-type gre

no tunnel-type gre

1.69.1 Purpose

Enables use of Generic Routing Encapsulation (GRE) tunnel types by mobile nodes (MN).

1.69.2 Command Mode

HA configuration

1.69.3 Syntax Description

`gre`	Specifies Generic Routing Encapsulation tunnels.

1.69.4 Default

IP-in-IP tunnels are enabled implicitly; no optional tunnel types are enabled.

1.69.5 Usage Guidelines

Use the **tunnel-type** command to use of GRE tunnel types by MNs.

Use the **no** form of this command to specify the default condition.

1.69.6 Examples

The following example enables the GRE tunnel type:

```
[local]Redback(config)#context ha
[local]Redback(config-ctx)#router mobile-ip
[local]Redback(config-mip)#home-agent
[local]Redback(config-mip-ha)#tunnel-type gre
```

1.70 tunnel-window

tunnel-window ***messages***

{**no** | **default**} **tunnel-window**

1.70.1 Purpose

Specifies the size of the control message window that is advertised to a Layer 2 Tunneling Protocol (L2TP) peer in Start-Control-Connection-Request (SCCRQ) or Start-Control-Connection-Reply (SCCRP) messages.

1.70.2 Command Mode

L2TP peer configuration

1.70.3 Syntax Description

`messages`	Number of messages the peer can send before acknowledgment from the SmartEdge router. The range of values is 1 to 2,000; the default value is 8.

1.70.4 Default

Up to eight control messages can be sent by an L2TP peer before acknowledgment from the SmartEdge router.

1.70.5 Usage Guidelines

Use the **`tunnel-window`** command to specify the size of control message window that is advertised to an L2TP peer in SCCRQ or SCCRP messages. The size of the window controls how many messages can be sent by a peer before it must wait for acknowledgement from the SmartEdge router.

You might need to change the number of messages, depending on the number of control messages a peer can generate at one time. For example, if a peer brings up many sessions all at once, you might need to increase the number of messages. However, changing the size of the control message window does not take effect until a new tunnel to the peer is established.

We recommend that you configure the control message window size to match the size configured on the L2TP peer, unless instructed to do otherwise by technical support.

Use the **`no`** or **`default`** form of this command to set the size of the control message window to the default.

1.70.6 Examples

The following example shows how to configure the peer to be able to send up to `15` control messages before acknowledgment from the SmartEdge router:

```
[local]Redback(config-ctx)#l2tp-peer name peer1
[local]Redback(config-l2tp)#tunnel-window 15
```

1.71 turbo

`turbo`

`no turbo`

1.71.1 Purpose

Enables fast IGMP message processing in the PPA.

1.71.2 Command Mode

IGMP snooping bridge configuration

1.71.3 Syntax Description

This command has no keywords or arguments.

1.71.4 Default

Fast IGMP message processing is disabled.

1.71.5 Usage Guidelines

Use the **`turbo`** command to enable fast IGMP message processing in the PPA.

Use the **`no`** form of this command to remove a configuration for IGMP message processing in the PPA.

1.71.6 Examples

The following example enables fast IGMP message processing in the PPA:

```
[local]Router(config)#context sj1
[local]Router(config-ctx)#bridge br-sj-1
[local]Router(config-bridge)#igmp snooping
[local]Router(config-igmp-snooping)#turbo
```

1.72 type (DSL)

type {**adsl1** | **adsl2** | **adsl2+** | **vdsl1** | **vdsl2** | **sdsl**}

no type {**adsl1** | **adsl2** | **adsl2+** | **vdsl1** | **vdsl2** | **sdsl**}

1.72.1 Purpose

Specifies the digital subscriber line (DSL) data type for the overhead profile and accesses overhead type configuration mode.

1.72.2 Command Mode

Overhead profile configuration

1.72.3 Syntax Description

adsl1	Specifies the asymmetric DSL1 data type.
adsl2	Specifies the asymmetric DSL2 data type.
adsl2+	Specifies the asymmetric DSL2+ data type.
vdsl1	Specifies the very-high-data rate DSL1 data type.
vdsl2	Specifies the very-high-data rate DSL2 data type.
sdsl	Specifies the symmetric DSL data type.

1.72.4 Default

None

1.72.5 Usage Guidelines

Use the **type** command to specify the DSL data type for the overhead profile and access overhead type configuration mode.

Use the **no** form of this command to remove the specified data type from the overhead profile.

1.72.6 Examples

The following example configures an overhead profile named `prof1` with the **encapsulation** type set to `pppoa-llc` and the ADSL type set to `ADSL1`:

```
[local]Redback(config)#qos profile prof1 overhead
[local]Redback(config-profile-overhead)#encaps-access-line pppoa-llc
[local]Redback(config-profile-overhead)#type adsl1
```

1.73 unknown-dest rate-limit

unknown-dest rate-limit ***kbps*** **burst-size** ***bytes***

no unknown-dest rate-limit

1.73.1 Purpose

Sets the rate and burst tolerance for traffic to unknown destinations on any port, circuit, or Virtual Private LAN Services (VPLS) pseudowire circuit to which you assign this bridge profile.

1.73.2 Command Mode

Bridge profile configuration

1.73.3 Syntax Description

kbps	Rate in kilobits per second. The range of values is 5 to 1,000000.
burst-size *bytes*	Burst tolerance in bytes. The range of values is 1 to 12,000,000.

1.73.4 Default

No rate limiting is imposed on traffic to unknown destinations on any port, circuit, or VPLS pseudowire circuit to which this profile is assigned.

1.73.5 Usage Guidelines

Use the **unknown-dest rate-limit** command to set the rate and burst tolerance for traffic to unknown destinations on any port, circuit, or VPLS pseudowire circuit to which you assign this bridge profile. For more information about VPLS pseudowire circuits, see *Configuring VPLS*.

Note: To protect against DOS attacks, you should always configure the rate limit.

1.73.6 Examples

The following example shows how to create the `prof-isp1` bridge profile and rate limits the destination traffic to `6000000` kbps and the burst size to `10000:`

```
[local]Redback(config)#bridge profile prof-isp1
[local]Redback(config-bridge-profile)#unknown-dest rate-limit
600000 burst-size 10000
```

1.74 unmount

`unmount` {**`/md`** | ***`file-sys`***}

1.74.1 Purpose

Prepares the mass-storage device before it is physically removed from the external slot of a controller card.

1.74.2 Command Mode

Exec (10)

1.74.3 Syntax Description

`/md`	Unmounts the mass-storage device installed in the external slot of the controller card to which you are connected.
`file-sys`	Name of the file system to be unmounted; for this release, the value is always **`/md`**.

1.74.4 Default

None

1.74.5 Usage Guidelines

Use the **`unmount`** command to prepare the mass-storage device before it is physically removed from the external slot of a controller card. You must enter this command from the command-line interface (CLI) that is running on the controller card with the device to be unmounted.

Note: This command is usually not needed when removing a compact-flash (CF) card from the external slot in a SmartEdge chassis; opening the CF door causes the SmartEdge router to unmount the CF card.

Caution!

Risk of equipment failure. Removing the mass-storage device without first entering this command can permanently damage the device and cause the kernel to crash. To reduce the risk, always enter the unmount command (in exec mode) before removing the device.

Caution!

Risk of data loss. You can lose data that is being transferred to the mass-storage device if you enter the **unmount** command (in exec mode) before the data transfer operation is complete. To reduce the risk, do not enter the **unmount** command while the CF ACTIVE LED is blinking. When the operation is complete, the LED is turned off.

Caution!

Risk of data loss. You can lose data being transferred to the mass storage device if you enter the **unmount /md** command (in exec mode) for the active controller card when the CLIPS or DHCP process is still running. To reduce the risk, stop the process using the process stop command with the **clips** or **dhcp** keyword (in exec mode) before unmounting a mass storage device from the active controller card.

Note: The contents of the mass-storage devices installed in the external slots of the active and standby controller cards are not synchronized. After you have unmounted a mass-storage device, it is inaccessible; any files written to /md from the CLI running on that controller card are stored in the /md directory on its internal compact-flash card. When the mass-storage device is mounted, either by insertion or by using the **mount /md** command (in exec mode), those files are not automatically written to the device; you must copy them manually to the device.

1.74.6 Examples

The following example unmounts the mass-storage device from the active controller card:

```
[local]Redback#unmount /md
```

1.75 update-source

update-source ***if-name***

no update-source

1.75.1 Purpose

Specifies the IP address of the interface used for Border Gateway Protocol (BGP) peering.

1.75.2 Command Mode

- BGP neighbor configuration
- BGP peer group configuration

1.75.3 Syntax Description

if-name	Name of the interface used to bring up the BGP session.

1.75.4 Default

The SmartEdge router brings up BGP sessions using any interface.

1.75.5 Usage Guidelines

Use the **update-source** command to assign the interface used to bring up a BGP session with the specified neighbor or peer group.

Use the **no** form of this command to bring up BGP sessions using any interface.

1.75.6 Examples

The following example configures loopback0 as the interface used to bring up BGP sessions with the neighbor at IP address 123.45.34.2:

```
[local]Redback(config-ctx)#router bgp 100
[local]Redback(config-bgp)#neighbor 123.45.34.2 external
[local]Redback(config-bgp-neighbor)#remote-as 200
[local]Redback(config-bgp-neighbor)#update-source loopback0
```

1.76 upgrade bootrom

upgrade bootrom {**ftp:** | **scp:** | **/md**} *url* [**no-reload**]

1.76.1 Purpose

Upgrades the boot ROM image in the EEPROM on the active controller card in a working system and reloads it.

1.76.2 Command Mode

Exec

1.76.3 Syntax Description

ftp:	Specifies the File Transfer Protocol (FTP) as the protocol to use when transferring the file from a remote server.
scp:	Specifies the Secured Copy Protocol (SCP) as the protocol to use when transferring the file from a remote server.
/md	Specifies the /md directory on the mass-storage device on the active controller card as the location for the file.
url	URL for the file that contains the boot ROM image.
no-reload	Optional. Cancels the reload when upgrading the system. The upgrade does not take place until the reload occurs at a later time.

1.76.4 Default

None

1.76.5 Usage Guidelines

Use the **upgrade bootrom** command to upgrade the boot ROM image in the EEPROM on the active controller card in a working system and reload it.

Use this command on a working system; if the system is not working, contact your customer support representative for assistance with alternative methods to upgrade the boot ROM image.

Use the **show version** command (in exec mode) to display the version of the boot ROM that is currently installed. In the display, the boot ROM is referred to as the system bootstrap. Contact your local technical support representative to determine if the boot ROM that is currently installed needs to be upgraded.

Your representative can also help you access the URL to use to download the boot ROM image file.

Note: The boot ROM image is also referred to as the boot loader.

When referring to a file on a remote server, the syntax for the ***url*** argument is:

//***username***[:***passwd***]@{***ip-addr*** | ***hostname***}[//***directory***]/***filename.ext***

Note: Use double slashes (`//directory`) if the pathname to the directory on the remote server is an absolute pathname; use a single slash (`/directory`) if it is a relative pathname (under the hierarchy of username account home directory).

When referring to a file on the `/md` directory, the URL takes the following form:

[/***directory***]/***filename.ext***

Directories can be nested. The value for the ***filename*** argument can be up to 256 characters in length.

Note: If you find a file in the `/flash` directory that appears to be a boot ROM image file, it might be the result of a manual method of upgrading the boot ROM. Do not use that file to upgrade the boot ROM using the **upgrade bootrom** command. For the correct procedure, see *Installing the SmartEdge OS*.

When you specify the **ftp:** or **scp:** keywords, this command downloads the requested file to the /flash directory, upgrades the boot ROM, and then deletes the file.

When you specify the **/md** keyword, this command copies the file to the `/flash` directory from the `/md` directory, upgrades the boot ROM, and then deletes the file on the `/flash` directory; the file on the /md directory is not altered. You must have first downloaded the file to the `/md` directory using the **copy** command (in exec mode).

Note: If you inadvertently copied the file to the `/flash` directory, you can move it to the /md directory, using the **copy** and **delete** commands (in exec mode).

If there is a standby controller card installed in the system, this command either synchronizes it with the active controller card automatically or, if you are upgrading a boot ROM with version 1205 or earlier, asks you to perform the synchronization manually. If manual synchronization is needed, the following message is printed:

```
% It appears that this is an upgrade from an old bootrom,

% in order to complete the upgrade it will be needed to run

% the command "reload standby" after this card finishes

% the upgrade and returns to the CLI
```

In this case, enter the **reload standby** command (in exec mode) after the upgrade of the boot ROM is complete, and the command-line interface (CLI) prompt displays.

To temporarily defer the automatic reload of the system to process the upgrade, use the optional **no-reload** keyword. The optional **no-reload** keyword can perform more than one upgrade during one reload operation to minimize impact on network traffic. For example, you can issue the **upgrade bootrom** and **upgrade minikernel** command (in exec mode), with the optional **no-reload** keyword, followed by a **release upgrade** command (in exec mode), to upgrade to an image in the alternate partition. Thus, with one reload of the system, three components are upgraded.

Note: When using the **upgrade** command with the optional **no-reload** keyword, the actual upgrade of the boot ROM or minikernel does not happen until the system reloads at a later time.

Note: If after issuing the **upgrade bootrom** command, there is a synchronizing operation between controller cards in the system, the temporary upgrade files may be overwritten and the upgrade operation may not have the intended effects when the system reloads. To minimize this risk, use the optional **no-reload** keyword on a synchronized system, then reload the system after issuing the **upgrade bootrom** command (in exec mode), as soon as possible.

You can also use the **show redundancy** command (in any mode), to view the status of the synchronization state before you use the **upgrade bootrom** command. For more information on the **show redundancy** command (in any mode), see *Managing Hardware*.

1.76.6 Examples

The following example displays the current version of the boot ROM and then upgrades the boot ROM image in the EEPROM on the active controller card. In this example, the current version is `0b1267`. FTP is used to download the file from a remote server; the controller card is then reloaded:

```
[local]Redback#show version

Redback Networks SmartEdge OS Version SE800-5.0.5-Release
Built by sysbuild@@lx-lsf27 Mon Jan 02 10:00:01 PDT 2006
Copyright (C) 1998-2006, Redback Networks Inc. All rights reserved.
System Bootstrap version is PowerPC,1.0b1187


[local]Redback#upgrade bootrom ftp://admin@10.10.1.2//bootroms/of1267.bin

This operation will cause the box to reload, do you want to continue?y

copying from ftp://admin@10.10.1.2//bootroms/1.0b1267.bin to
local:/flash/of.bin...

Connected to 10.10.1.2.

..

*************************************

684 KB  724.97 KB/s    00:00 ETA

Nov 9 12:52:16: %DLM-6-INFO: Standby xcrp's /flash may not be in sync

226 File send OK.

700437 bytes received in 00:01 (682.20 KB/s)

221 Goodbye.

.

Nov 15 11:49:34: %ALAPI-6-INFO: XCRP in slot 7, will now reload



.

.rebooting

.

.

.
```

The following example displays the output of the **upgrade bootrom** command (in exec mode), using the optional **no-reload** keyword:

```
[local]Redback#upgrade bootrom /md/of1267.bin no-reload


copying from md:/md/of1267.bin to local:/flash/of.bin...
```

Files are copied over, but the boot ROM image is not reloaded.

1.77 upgrade minikernel

upgrade minikernel {**ftp:** | **scp:** | **/md**} *url* [**no-reload**]

1.77.1 Purpose

Upgrades the minikernel image in the EEPROM on the active controller card in a working system and reloads it.

1.77.2 Command Mode

Exec

1.77.3 Syntax Description

`ftp:`	Specifies the File Transfer Protocol (FTP) as the protocol to use when transferring the file from a remote server.
`scp:`	Specifies the Secured Copy Protocol (SCP) as the protocol to use when transferring the file from a remote server.
`/md`	Specifies the /md directory on the mass-storage device on the active controller card as the location for the file.
`url`	URL for the file that contains the minikernel image.
`no-reload`	Optional. Cancels the reload when upgrading the minikernel image. The upgrade does not take place until the reload occurs at a later time.

1.77.4 Default

None

1.77.5 Usage Guidelines

Use the **`upgrade minikernel`** command to upgrade the minikernel image in the EEPROM on the active controller card in a working system and reload it.

Use this command on a working system; if the system is not working, contact your customer support representative for assistance with alternative methods to upgrade the minikernel image.

When upgrading the minikernel image on a system that has a boot ROM image file installed with version 1205 or earlier, you must first upgrade the boot ROM image, using the **`upgrade bootrom`** command (in exec mode), before you can upgrade the minikernel image.

Contact your local technical support representative to determine if the minikernel image that is currently installed needs to be upgraded. Your representative can also help you access the URL to use to download the minikernel image file.

When referring to a file on a remote server, the syntax for the ***`url`*** argument is:

//username*[:*passwd*]@{*ip-addr* | *hostname*}[*//directory*]*/filename.ext

Note: Use double slashes (//) if the pathname to the directory on the remote server is an absolute pathname; use a single slash (/) if it is a relative pathname (under the hierarchy of username account home directory).

When referring to a file on the `/md` directory, the syntax for the ***url*** argument is:

[***/directory***]***/filename.ext***

Directories can be nested. The value for the ***filename*** argument can be up to 256 characters in length.

Note: If you find a file in the `/flash` directory that appears to be a minikernel image file, it might be the result of upgrading the minikernel image using a manual procedure.. Do not use this file to upgrade the minikernel image using the **upgrade minikernel** command. For the correct procedure, see *Installing the SmartEdge OS*.

When you specify the **ftp:** or **scp:** keyword, this command downloads the requested file to the `/flash` directory, upgrades the minikernel image, and then deletes the file.

When you specify the **/md** keyword, this command copies the file to the `/flash` directory from the `/md` directory, upgrades the minikernel image, and then deletes the file on the `/flash directory`; the file on the `/md` directory is not altered. You must have first downloaded the file to the `/md` directory using the **copy** command (in exec mode).

Note: If you inadvertently copied the file to the `/flash` directory, you can move it to the `/md` directory, using the **copy** and **delete** commands (in exec mode).

If a standby controller card is installed on the system, it is synchronized automatically with the active controller card with this command, unless you are running a boot ROM version 1205 or earlier. In that case, you must synchronize the standby controller manually, using the **reload standby** command (in exec mode), after the upgrade of the minikernel image is completed and the command-line interface (CLI) prompt displays.

To temporarily defer the automatic reload of the system to process the upgrade, use the optional **no-reload** keyword. The optional **no-reload** keyword can perform more than one upgrade during one reload operation to minimize impact on network traffic. For example, you can issue the **upgrade bootrom** and **upgrade minikernel** command (in exec mode), with the optional **no-reload** keyword, followed by a **release upgrade** command (in exec mode), to upgrade to an image in the alternate partition. Thus, with one reload of the system, three components are upgraded.

Note: When using the **upgrade** command with the optional **no-reload** keyword, the actual upgrade of the boot ROM or minikernel image does not happen until the system reloads at a later time.

Note: If after issuing the **upgrade minikernel** command, there is a synchronizing operation between controller cards in the system, the temporary upgrade files may be overwritten and the upgrade operation may not have the intended effects when the system reloads. To minimize this risk, use the optional **no-reload** keyword on a synchronized system, then reload the system after issuing the **upgrade minikernel** command (in exec mode), as soon as possible.

You can also use the **show redundancy** command (in any mode), to view the status of the synchronization state before you use the **upgrade minikernel** command. For more information on the **show redundancy** command (in any mode), see *Managing Hardware*.

1.77.6 Examples

The following example upgrades the minikernel image in the EEPROM on the active controller card using the FTP to download the file from a remote server; the controller card is then reloaded. Interactive responses are shown in bold; the `save the current configuration` message displays only if the current configuration appears to have been modified and not saved:

```
[local]Redback#upgrade minikernel ftp://userid@10.10.2.1//minikernels/netbsd.min.v20.bz2.bin

This operation will cause the box to reload, do you want to continue?y
Do you want to save the current configuration? (y/n) y
copying from ftp://userid@10.0.2.1//minikernels/netbsd.min.v20.bz2.bin
to
local:/flash/netbsd.min.bz2...
Connected to 10.0.2.1.
.
.
.
*************************************| 951 KB 1.10 MB/s
00:00 ETA
226 Transfer complete.
974310 bytes received in 00:00 (1.10 MB/s)
221 Goodbye.
Nov 15 15:37:07: %DLM-6-INFO: Standby xcrp's /flash may not be in sync
[local]Redback#Sep 25 15:37:10: %ALAPI-6-INFO: XCRP in slot 7, will now reload
.
.
.
Updating minikernel...
Erasing flash...done
Writing minikernel [974310 bytes]...done
Verifying minikernel...done
Õ
[0]Booting(2)....
Enabling L1/L2 Caches...
Reset Cause:
SCC Watchdog Reset (PPC Subsystem ONLY)
ISA reset (PPC Subsystem ONLY) .
.
.
Welcome to SmartFirmware(tm) for Ericsson PowerPC Copyright (c) 1999-2006
by Ericsson AB version of1267
SmartFirmware(tm) Copyright 1996-2006 by CodeGen, Inc.
All Rights Reserved.
Auto-boot in 0 seconds - press se* to abort, ENTER to boot:
```

The following example displays the output of the **upgrade minikernel** command using the optional **no-reload** keyword:

```
[local]Redback#upgrade minikernel /md/netbsd.min.v21.bz2 no-reload
copying from md:/md/netbsd.min.v21.bz2 to local:/flash/netbsd.min.bz2...
```

Files are copied over, but the boot ROM image is not reloaded.

1.78 url

url ***url***

no url

1.78.1 Purpose

Configures the URL to which the current subscriber HTTP session is to be redirected.

1.78.2 Command Mode

HTTP redirect profile configuration

1.78.3 Syntax Description

url

URL to which the subscriber HTTP session is to be redirected. You can add a backslash at the end of the URL followed by any of these variables to personalize the URL:

- %c—Calling-station-ID of the subscriber session.
- %d—Domain portion of the subscriber name.
- %i—IP address of the subscriber session.
- %n—NAS-port-ID of the subscriber session.
- %t—Time stamp (in seconds) indicating when the HTTP redirection is applied to the subscriber.
- %u—Username portion of the subscriber name.
- %U—Entire subscriber name used in Point-to-Point Protocol (PPP) authentication.

1.78.4 Default

An HTTP redirect URL is not configured.

1.78.5 Usage Guidelines

Use the **url** command to configure the URL to which the current subscriber session is to be redirected.

For configuration that uses **%c** (for the Calling-station-ID) or **%n** (for the NAS-port-ID), the radius attributes configuration needs to be present in the context where subscriber is terminated. For example, you can perform this configuration using the `radius attribute calling-station-id format agent-circuit-id` or `radius attribute nas-port-id format physical` command, respectively.

Caution!

Risk of redirect loop. Risk of redirect loop. Redirect can recur until an IP ACL that permits access to the new web page is applied to the subscriber record or profile. To reduce the risk, before modifying an existing URL, ensure that the subscriber record includes an IP ACL that permits access to the new URL.

Note: If the URL contains a question mark (?), press the **Escape** (**Esc**) key before you enter the **?** character. Otherwise, the SmartEdge router command-line interface (CLI) interprets the ? character as a request for help and does not allow you to complete the URL.

Use the **no** form of this command to delete the URL from the HTTP redirect profile.

1.78.6 Examples

The following example configures the URL, `www.Redirect.com:`

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#http-redirect profile Redirect
[local]Redback(config-hr-profile)#url http://www.Redirect.com
```

1.79 user-class-id

user-class-id *user-class-id* [**offset** *position*] **giaddr** *ip-addr*

no user-class-id *user-class-id*

1.79.1 Purpose

Specifies an IP address for the giaddr field in the header of Dynamic Host Configuration Protocol (DHCP) packets for the specified user class ID (option 77) field.

1.79.2 Command Mode

DHCP giaddr configuration

1.79.3 Syntax Description

user-class-id	Identifier to be matched against the contents of the DHCP option 77 ID field in DHCP discover packets, in one of the formats given in the Usage Guidelines section, for which this IP address is intended.
offset *position*	Optional. Position of the starting octet in the option 77 field which is to be matched with the specified *user-class-id* argument, according to one of the following formats: • +n or n—Starting octet is the nth octet in the received ID. The matching operation is performed on the nth and succeeding octets for the length of the string specified by the value of the *user-class-id* argument. • –n—Starting octet is the last octet in the received ID minus the previous (n–1) octets. The matching operation is performed on the succeeding octets for the length of the string specified by the value of the *user-class-id* argument. The default value is 1 (the first octet). You can also specify the first octet with a value of 0.
giaddr *ip-addr*	IP address to be inserted in the giaddr field in the header of DHCP packets for the specified user class ID.

1.79.4 Default

The giaddr field is set to the primary IP address of the interface.

1.79.5 Usage Guidelines

Use the **user-class-id** command to specify the IP address for the giaddr field in the header of DHCP packets for the specified user class ID (option 77) field. Option 77 is described in RFC 3004, *The User Class Option for DHCP.*

When the SmartEdge router receives a DHCP discover packet, the SmartEdge router performs a matching operation, comparing the contents of the option 77 field, starting at the octet within the field, as specified by the value of the ***position*** argument, with the string specified by the value of the ***user-class-id*** argument.

If more than one user class ID field is present in the option 77 field in the DHCP discover packet, the system uses only the first user class ID field to make the comparison for setting the giaddr field. The remaining user class ID fields are ignored.

If there is a match, the system inserts the specified IP address in the giaddr field in the header of DHCP packets to this client. If there is no match, the system inserts the primary IP address that you have configured for this interface.

Possible formats for the ***user-class-id*** argument are:

- Alphanumeric string, enclosed in quotation marks (“ ”); for example, “ABCD1234”
- Alphanumeric string, not enclosed in quotation marks; for example, Ericsson1
- Hex numeric string, not enclosed in quotation marks and prefaced with 0x or 0X; for example, 0Xabcd1234

Use the **giaddr** ***ip-addr*** construct to specify an IP address for the specified ***user-class-id*** argument. This IP address must be one of the secondary IP addresses that you have configured for the interface. You can specify the same IP address or different IP addresses for multiple values of the ***user-class-id*** argument.

Use the **no** form of this command to delete the giaddr IP address for the specified ***user-class-id*** argument.

Note: If you delete this DHCP proxy or relay from the configuration, using the **no** form of the **dhcp proxy** or **dhcp relay** command (in interface configuration mode), you also delete all **user-class-id** commands for that DHCP proxy or relay.

1.79.6 Examples

The following example specifies secondary IP addresses for the interface in which the DHCP proxy server is configured, and then specifies one of them as the IP address for the giaddr field for the network user class ID:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#interface voip multibind
[local]Redback(config-if)#ip address 200.1.1.1/24
[local]Redback(config-if)#ip address 200.1.2.1/24 secondary
[local]Redback(config-if)#ip address 200.1.10.1/24 secondary
[local]Redback(config-if)#dhcp proxy 16000
[local]Redback(config-dhcp-giaddr)#user-class-id network giaddr 200.1.2.1
```

1.80 valid-lifetime

In ND profile configuration mode, the syntax is:

`valid-lifetime` {*`lifetime`* | **`infinite`**}

{**`no`** | **`default`**} **`valid-lifetime`**

In ND router configuration mode or ND router interface configuration mode, the syntax is:

`valid-lifetime` *`lifetime`*

{**`no`** | **`default`**} **`valid-lifetime`**

1.80.1 Purpose

Specifies the valid lifetime in the Prefix Information option for IPv6 prefixes included in a Route Advertisement (RA) message.

1.80.2 Command Mode

- ND profile configuration
- ND router configuration
- ND router interface configuration

1.80.3 Syntax Description

`lifetime`	Optional in the ND profile configuration mode. Valid lifetime, in seconds. The range of values is 0 to 4,294,967,295.
`infinite`	Optional. Only applicable to the ND profile configuration mode. Sets the valid lifetime to the lifetime of the subscriber circuit.

1.80.4 Default

The valid lifetime default is infinite.

1.80.5 Usage Guidelines

Use the **`valid-lifetime`** command to specify the valid lifetime in the Prefix Information option for IPv6 prefixes included in an RA message. This value equals the length of time that an IPv6 prefix is valid for on-link determination. In

ND profile configuration mode, this command specifies the value for this ND profile. In ND router configuration mode, this command specifies the global value for all interfaces. In ND router interface mode, it specifies the value for this ND router interface. If specified, the setting for the interface overrides the global setting.

Use the **no** or **default** form of this command to specify the default condition.

1.80.6 Examples

The following example specifies a valid lifetime of `infinite` seconds for the ND profile **ndprofile7**:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#nd profile ndprofile7
[local]Redback(config-nd-if)#valid-lifetime infinite
```

The following example specifies a valid lifetime of `43200` seconds (12 hours) for all interfaces for this ND router:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#router nd
[local]Redback(config-nd-if)#valid-lifetime 43200
```

The following example specifies a valid lifetime of `2880` seconds (48 minutes) for the `int1` ND router interface, which overrides the global setting:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#router nd
[local]Redback(config-nd)#interface int1
[local]Redback(config-nd-if)#valid-lifetime 2880
```

1.81 vb-index

vb-index ***vb-index*** **vb-value** ***vb-value***

no vb-index

1.81.1 Purpose

Identify the object in the SNMP PDU varbind list and the object's value which corresponds to the alarm model state.

1.81.2 Command Mode

SNMP alarm model configuration

1.81.3 Syntax Description

vb-index	An integer between 1 and 512 representing the index value in an SNMPV2-TRAP-PDU or InformRequest-PDU varbind list.
vb-value ***vb-value***	The value of the object indicated by vb-index which corresponds to the specified alarm state.

1.81.4 Default

None

1.81.5 Usage Guidelines

Use the **vb-index** command to configure the index in the varbind list of the notification for the alarm model you are configuring. The index must correspond to the index in SNMPV2-TRAP-PDU or InformRequest-PDU packet for the notification. An index from a SNMPv1-Trap-PDU is not valid. The vb-value keyword identifies the placement in the varbind list of the event that you want to configure in the alarm model

Use the **no** form of this command to remove the varbind index and notification event varbind list placement configuration from the alarm model.

1.81.6 Examples

In this example, the clear state occurs when the fourth varbind in the varbind list has a value of 2.

```
[local]jazz#config
[local]jazz(config)#snmp alarm model 1 state clear
[local]jazz(config-snmp-alarmmodel)#vb-index 4 vb-value 2
[local]jazz(config-snmp-alarmmodel)#exit
```

1.82 vb-subtree

vb-subtree ***vb-subtree***

no vb-subtree

1.82.1 Purpose

Identifies the name of the VarBind in the notification.

1.82.2 Command Mode

SNMP alarm model configuration

1.82.3 Syntax Description

vb-subtree	An OID in the notification varbind list which, when matched, is used for resource identification. This object may be used in conjunction with the **res-prefix** command. If the value is 0.0 then the first object after notification OID is used.

1.82.4 Default

None

1.82.5 Usage Guidelines

Use the **vb-subtree** command to index value for alarmActiveResourceID. Each varbind in the notification is compared to the vb-subtree OID value. If the varbind OID is equal to or a subtree of this value, the search stops. The matching part is called the prefix and the remainder is called the indices. The alarmActiveResourceID is constructed by concatenating the indices to the value of the **res-prefix** keyword. If the **vb-subtree** value is 0.0 or is not set then the system will match the first varbind after the timestamp and trap OID.

Use the **no** form of this command to remove the matched VarBind name.

1.82.6 Examples

The following example shows how to identify a VarBind match with the name **AlarmID**.

```
[local]Redback#config
[local]Redback(config)#snmp alarm model 1 state clear
[local]Redback(config-snmp-alarmmodel)#vb-subtree alarmid
[local]Redback(config-snmp-alarmmodel)#exit
```

1.83 vc-id (L2VPN XC)

`vc-id` *`vc-id`* `peer` *`peer-address`*

`no vc-id`

1.83.1 Purpose

Specifies the peer that hosts the remote end of a backup Layer 2 VPN (L2VPN) cross-connect (XC).

1.83.2 Command Mode

Primary L2VPN XC configuration

1.83.3 Syntax Description

`vc-id`	Virtual circuit (VC) that host the remote end of the backup XC.
`peer` *`peer-address`*	IP address of the peer that hosts the remote end of the backup XC.

1.83.4 Default

No peer is configured to host the remote end of the backup L2VPN XC.

1.83.5 Usage Guidelines

Use the `vc-id` command to specify the peer that hosts the remote end of a backup L2VPN XC.

Use the `no` form of this command to remove a configured peer as the remote end of the backup L2VPN XC.

1.83.6 Examples

The following example shows how to specify VC 200 on the IP address 100.100.100.2 as the peer that hosts the remote end of an XC that backs up the XC 10/3 on VLAN 100 and VC 100:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#l2vpn
[local]Redback(config-l2vpn)#xc-group default
[local]Redback(config-l2vpn-xc-group)#xc 10/3 vlan-id 100 vc-id 100
profile prof2 backup
[local]Redback(config-l2vpn-xc-prime]#vc-id 200 peer 100.100.100.2
```

1.84 vc-id (Port PW)

vc-id ***vc-id*** **profile** ***prof-name***

{no | default} vc-id ***vc-id*** **profile** ***prof-name***

1.84.1 Command Mode

Port configuration mode

1.84.2 Syntax Description

vc-id	LDP VC ID for the PW; must be unique on the router; valid values can be from 0 to 4294967295..
profile *prof-name*	Previously configured L2VPN profile to be assigned to the port PW; identifies the peer at the remote end of the PW.

1.84.3 Default

None

1.84.4 Usage Guidelines

Use the **vc-id** command to assign an LDP vc-id to a port pseudowire (PW) and associate a profile with it.

1.84.5 Examples

The following example assigns VC ID `1001` and the `l2_prof1001` profile to port PW `l2-net`.

```
[local]Redback(config)#port pseudowire l2-net
[local]Redback(config-port)#vc-id 1001 profile l2_prof1001
```

1.85 vendor-class

```
vendor-class vendor-class-id [offset position] subnet-name
subnet-name

no vendor-class vendor-class-id
```

1.85.1 Purpose

Creates a static mapping between a subnet and the specified vendor class ID.

1.85.2 Command Mode

DHCP server configuration

1.85.3 Syntax Description

`vendor-class-id`	Vendor class ID for which a static mapping is to be created.
`offset` *`position`*	Optional. Position of the starting octet in the option 60 field which is to be matched with the specified *`vendor-class-id`* argument, according to one of the following formats: • +n or n—Starting octet is the nth octet in the received ID. The matching operation is performed on the nth and succeeding octets for the length of the string specified by the value of the *`vendor-class-id`* argument. • –n—Starting octet is the last octet in the received ID minus the previous (n–1) octets. The matching operation is performed on the succeeding octets for the length of the string specified by the value of the *`vendor-class-id`* argument. The default value is 1 (the first octet). You can also specify the first octet with a value of 0.
`subnet-name` *`subnet-name`*	Subnet name for the specified vendor class ID.

1.85.4 Default

No static mapping is created between a subnet and any vendor class ID.

1.85.5 Usage Guidelines

Use the **`vendor-class`** command to create a static mapping between a subnet and the specified vendor class ID.

Use the **no** form of this command to delete the static mapping between the vendor class ID and the subnet.

1.85.6 Examples

The following example specifies the `for-subs` subnet as the subnet for the `123456` vendor class ID:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#dhcp server policy
[local]Redback(config-dhcp-server)#vendor-class 123456 offset 1 subnet-name for-subs
```

1.86 vendor-class-id

vendor-class-id ***vendor-class-id*** [**offset** ***position***] **giaddr** ***ip-addr***

no vendor-class-id ***vendor-class-id***

1.86.1 Purpose

Specifies an IP address for the giaddr field in the header in Dynamic Host Configuration Protocol (DHCP) packets for the specified vendor class ID (option 60) field.

1.86.2 Command Mode

DHCP giaddr configuration

1.86.3 Syntax Description

`vendor-class-id`	Identifier to be matched against the contents of the DHCP option 60 ID field in DHCP discover packets, in one of the formats given in the Usage Guidelines section, for which this IP address is intended.
`offset` *`position`*	Optional. Position of the starting octet in the option 60 field which is to be matched with the specified *`vendor-class-id`* argument, according to one of the following formats: • +n or n—Starting octet is the nth octet in the received ID. The matching operation is performed on the nth and succeeding octets for the length of the string specified by the value of the *`vendor-class-id`* argument. • –n—Starting octet is the last octet in the received ID minus the previous (n–1) octets. The matching operation is performed on the succeeding octets for the length of the string specified by the value of the *`vendor-class-id`* argument. The default value is 1 (the first octet). You can also specify the first octet with a value of 0.
`giaddr` *`ip-addr`*	IP address to be inserted in the giaddr field in the header of DHCP packets for the specified vendor class ID.

1.86.4 Default

The giaddr field is set to the primary IP address of the interface.

1.86.5 Usage Guidelines

Use the **`vendor-class-id`** command to specify the IP address for the giaddr field in DHCP packets for the specified vendor class ID (option 60) field. option 60 is described in RFC 2131, *DHCP Options and BootP Vendor Extensions.*

When the SmartEdge router receives a DHCP discover packet, the SmartEdge router performs a matching operation, comparing the contents of the option 60 field, starting at the octet within the field, as specified by the value of the *`position`* argument, with the string specified by the value of the *`vendor-class-id`* argument.

If there is a match, the system inserts the specified IP address in the giaddr field in the header of DHCP packets to this client. If there is no match, the system inserts the primary IP address that you have configured for this interface.

Possible formats for the *`vendor-class-id`* argument are:

- Alphanumeric string, enclosed in quotation marks (“ ”); for example, “ABCD1234”

- Alphanumeric string, not enclosed in quotation marks; for example, Ericsson1

- Hex numeric string, not enclosed in quotation marks and prefaced with 0x or 0X; for example, 0Xabcd1234

Use the **giaddr** ***ip-addr*** construct to specify an IP address for the specified ***vendor-class-id*** argument. This IP address must be one of the secondary IP addresses that you have configured for the interface. You can specify the same IP address or different IP addresses for multiple values of the ***vendor-class-id*** argument.

Use the **no** form of this command to delete the giaddr IP address for the specified ***vendor-class-id*** argument.

Note: If you delete this DHCP proxy or relay from the configuration, using the **no** form of the **dhcp proxy** or **dhcp relay** command (in interface configuration mode), you also delete all **vendor-class-id** commands for that DHCP proxy or relay.

1.86.6 Examples

The following example specifies secondary IP addresses for the interface in which the DHCP proxy server is configured, and then specifies one of them as the IP address for the giaddr field for the regional vendor class ID:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#interface voip multibind
[local]Redback(config-if)#ip address 200.1.1.1/24
[local]Redback(config-if)#ip address 200.1.2.1/24 secondary
[local]Redback(config-if)#ip address 200.1.10.1/24 secondary
[local]Redback(config-if)#dhcp proxy 16000
[local]Redback(config-dhcp-giaddr)#vendor-class-id regional offset
-17 giaddr 200.1.2.1
```

1.87 verify-set

verify-set ***interval*** [**timeout-multiplier** ***count***] [**min-success** ***count***]

{**no** | **default**} **verify-set** ***interval*** [**timeout-multiplier** ***count***] [**min-success** ***count***]

1.87.1 Purpose

Configures the verify-set values for a dynamically verified static routing (DVSR) profile.

1.87.2 Command Mode

DVSR profile configuration

1.87.3 Syntax Description

interval	Interval value that defines how often DVSR route verification occurs. The interval range, in seconds, is 10 to 7,200; the default value is 20. It can only be set in 5-second increments.
timeout-multiplier *count*	Optional. Timeout multiplier. The *count* argument defines the number of verification failures that a DVSR route must have before being considered in the down state; the default value is 3.
min-success *count*	Optional. Minimum success. The *count* argument defines the number of verification successes that a DVSR route must have before being considered in the up state. Range is from 1 to 10; the default value is 3.

1.87.4 Default

For a DVSR profile, the default interval value is 20 seconds, the default timeout multiplier value is 3, and the default minimum success value is 3.

1.87.5 Usage Guidelines

Use the **verify-set** command to configure the verify-set values for a DVSR profile. The verify set values control the frequency of the verification of DVSR routes, and change the measurement of verification. The smaller the number is, the more responsive the DVSR route becomes; however, fast response may cause network instability, especially in the case of packet loss in the network.

Use the **no** form of this command to delete the verify-set value from a DVSR profile.

Use the **default** form of this command to return the verify-set values for a DVSR profile to their default values.

1.87.6 Examples

The following example defines a DVSR profile using a verification interval of 25 seconds, a timeout multiplier of 6, and a minimum success of 4:

```
[local]Redback(config-ctx)#dvsr-profile abc-webfarm
[local]Redback(config-dvsr)#verify-set 25 timeout-multiplier 6
min-success 4
```

1.88 version

version {2 | 3}

default version

1.88.1 Purpose

Configures the IGMP snooping version used by a bridge.

1.88.2 Command Mode

IGMP snooping bridge configuration

1.88.3 Syntax Description

2	Configures the bridge to use IGMP version 2. This is the default setting.
3	Configures the bridge to use IGMP version 3.

1.88.4 Default

IGMP version 2

1.88.5 Usage Guidelines

Use the **version** command to configure the IGMP snooping version used by a bridge.

For information about the differences between IGMP version 2 and IGMP version 3, see the following RFCs:

- RFC 2236, *Internet Group Management Protocol, Version 2*

- RFC 3376, *Internet Group Management Protocol, Version 3*

Use the **default** form of this command to return the bridge to the default setting in which IGMP version 2 is used.

1.88.6 Examples

The following example shows how to configure an Ethernet bridge called `br-sj-1` to use IGMP version 3:

```
[local]Redback(config)#context sj1
[local]Redback(config-ctx)#bridge br-sj-1
[local]Redback(config-bridge)#igmp snooping
[local]Redback(config-igmp-snooping)#version 3
```

1.89 violate drop

violate drop

{**no** | **default**} **violate drop**

1.89.1 Purpose

Drops packets that exceed the configured excess burst tolerance.

1.89.2 Command Mode

- Policy class rate configuration
- Policy rate configuration

1.89.3 Syntax Description

This command has no keywords or arguments.

1.89.4 Default

Packets exceeding the configured excess burst tolerance are dropped.

1.89.5 Usage Guidelines

Use the **violate drop** command to drop packets that exceed the configured excess burst tolerance. Use this command as part of a policing policy for incoming packets and as part of a metering policy for outgoing packets.

To configure the excess burst tolerance, enter the **rate** command (in policy ACL class, metering policy, or policing policy configuration mode). The following conditions determine how packets are dropped:

- If the excess burst tolerance is not configured, all packets exceeding the configured burst tolerance are dropped.
- If the excess burst tolerance is configured, all packets that exceed the excess burst tolerance are dropped.

Caution!

Risk of overriding configurations. The SmartEdge router checks for and applies marking in a specific order. To reduce the risk, remember the following guidelines: Circuit-based marking overrides class-based marking and Border Gateway Protocol (BGP) destination-based marking, through route maps, overrides both circuit-based and class-based marking.

Note: Use the **`exceed drop`** commands (in policy class rate and policy rate configuration modes) to specify how packets are dropped when the traffic rate does not exceed the configured excess burst tolerance.

Use the **`no`** or **`default`** form of this command to drop packets that exceed the configured excess-burst tolerance.

1.89.6 Examples

The following example drops packets that exceed the excess burst tolerance:

```
[local]Redback(config)#qos policy protection1 policing
[local]Redback(config-policy-policing)#rate 10000 burst 100000 excess-burst 120000
[local]Redback(config-policy-rate)#violate drop
```

1.90 violate mark dscp

`violate mark dscp` *`dscp-class`*

{**`no`** | **`default`**} **`violate mark dscp`**

1.90.1 Purpose

Assigns a quality of service (QoS) Differentiated Services Code Point (DSCP) priority to IP packets that exceed the configured QoS rate. For IPv4 packets, the DSCP marking is the upper 6 bits of the IPv4 header Type of Service (ToS) field. For IPv6 packets, the DSCP marking is the upper 6 bits of the IPv6 header Traffic Class field. In either case, the specific bits affected are those denoted by *dd* in the octet field with the format *pppddxxx*.

1.90.2 Command Mode

- Policy class rate configuration
- Policy rate configuration

1.90.3 Syntax Description

dscp-class

Priority with which packets exceeding the rate are marked. Values can be:

- An integer from 0 to 63.
- One of the keywords listed in Table 4.

1.90.4 Default

Packets exceeding the configured excess burst tolerance are dropped.

1.90.5 Usage Guidelines

Use the **violate mark dscp** command to mark packets that exceed the configured excess burst tolerance with a DSCP value.

To configure the excess burst tolerance, enter the **rate** command (in policy ACL class, metering policy, or policing policy configuration mode). Only one mark instruction can be in effect at a time. To change the mark instruction, enter the **violate mark dscp** command, specifying a new value for the *dscp-class* argument, which supersedes the one previously configured.

Table 4 lists the keywords for the *dscp-class* argument.

Table 4 DSCP Class Keywords

DSCP Class	Keyword	DSCP Class	Keyword
Assured Forwarding (AF) Class 1 /Drop precedence 1	**af11**	Class Selector 0 (same as default forwarding)	**cs0** (same as **df**)
AF Class 1/Drop precedence 2	**af12**	Class Selector 1	**cs1**
AF Class 1/Drop precedence 3	**af13**	Class Selector 2	**cs2**
AF Class 2/Drop precedence 1	**af21**	Class Selector 3	**cs3**
AF Class 2/Drop precedence 2	**af22**	Class Selector 4	**cs4**
AF Class 3/Drop precedence 3	**af23**	Class Selector 5	**cs5**
AF Class 3/Drop precedence 1	**af31**	Class Selector 6	**cs6**

Table 4 DSCP Class Keywords

DSCP Class	Keyword	DSCP Class	Keyword
AF Class 3/Drop precedence 2	`af32`	Class Selector 7	`cs7`
AF Class 3/Drop precedence 3	`af33`	Default Forwarding (same as Class Selector 0)	`df` (same as `cs0`)
AF Class 4/Drop precedence 1	`af41`	Expedited Forwarding	`ef`
AF Class 4/Drop precedence 2	`af42`		
AF Class 4/Drop precedence 3	`af43`		

Note: RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*, defines the Class Selector code points.

Caution!

Risk of packet reordering. To reduce the risk, ensure that the marking of conforming packets and exceeding packets differ only within a major DSCP class. Major DSCP classes are identified by the Class Selector code, and include CS0=DF, CS1=AF11, AF12, AF13, CS2=AF21, AF22, AF23, CS3=AF31, AF32, AF33, CS4=AF41, AF42, AF43, and CS5=EF. For example, if you mark conforming packets with AF11 and you want to avoid reordering, mark exceeding packets with AF11, AF12, or AF13 only.

Caution!

Risk of overriding configurations. The SmartEdge router checks for and applies marking in a specific order. To reduce the risk, remember the following guidelines: Circuit-based marking overrides class-based marking and Border Gateway Protocol (BGP) destination-based marking, through route maps, overrides both circuit-based and class-based marking.

Use the `no` or `default` form of this command to return to the default behavior of dropping packets that exceed the excess burst tolerance.

1.90.6 Examples

The following example configures the policy to mark all packets that exceed the excess burst tolerance with a DSCP value representing a high priority:

```
[local]Redback(config)#qos policy protection1 policing
[local]Redback(config-policy-policing)#rate 10000 burst 100000
excess-burst 120000
[local]Redback(config-policy-rate)#violate mark dscp ef
```

1.91 violate mark precedence

violate mark precedence ***prec-value***

{**no** | **default**} **violate mark precedence**

1.91.1 Purpose

Assigns a quality of service (QoS) Differentiated Services Code Point (DSCP) drop-precedence value to IP packets that exceed the configured QoS rate. For IPv4 packets, the DSCP marking is the upper six bits of the IPv4 header Type of Service (ToS) field. For IPv6 packets, the DSCP marking is upper six bits of the IPv6 header Traffic Class field. In either case, the specific bits affected are those denoted by *dd* in the octet field with the format *pppddxxx*.

1.91.2 Command Mode

- Policy class rate configuration
- Policy rate configuration

1.91.3 Syntax Description

prec-value	Drop precedence bits value. The range of values is 1 to 3.

1.91.4 Default

Packets exceeding the excess burst tolerance are dropped.

1.91.5 Usage Guidelines

Use the **violate mark precedence** command to mark packets that exceed the configured excess burst tolerance with a drop precedence value corresponding to the AF class of the packet.

To configure the excess burst tolerance, enter the **rate** command (in policy ACL class, metering policy, or policing policy configuration mode).

In general, the level of forwarding assurance of an IP packet is based on: (1) the resources allocated to the AF class to which the packet belongs, (2) the current load of the AF class, and, in case of congestion within the class, (3) the drop precedence of the packet. In case of congestion, the drop precedence of a packet determines the relative importance of the packet within the AF class. Packets with a lower drop precedence value are preferred and protected from being lost, while packets with a higher drop precedence value are discarded.

With AF classes AF1 (AF11, AF12, AF13), AF2 (AF21, AF22, AF23), AF3 (AF31, AF32, AF33), and AF4 (AF41, AF42, AF43), the second integer represents a drop precedence value. Table 5 shows how the AF drop precedence value of an incoming packet is changed when it exits the SmartEdge router after being tagged with a new drop precedence. (See also RFC 2597, ***Assured Forwarding PHB Group***.)

Table 5 Drop Precedence Values

DSCP Value of an Incoming Packet	Packet is Tagged with a Drop Precedence Value	DSCP Value of the Outgoing Packet
AF11, AF12, AF13 AF21, AF22, AF23 AF31, AF32, AF33 AF41, AF42, AF43	1	AF11 AF21 AF31 AF41
AF11, AF12, AF13 AF21, AF22, AF23 AF31, AF32, AF33 AF41, AF42, AF43	2	AF12 AF22 AF32 AF42
AF11, AF12, AF13 AF21, AF22, AF23 AF31, AF32, AF33 AF41, AF42, AF43	3	AF13 AF23 AF33 AF43

Only one mark instruction can be in effect at a time. To change the mark instruction, enter the **violate mark precedence** command, specifying a new value for the ***prec-value*** argument. This supersedes the one previously configured.

Caution!

Risk of overriding configurations. The SmartEdge router checks for and applies marking in a specific order. To reduce the risk, remember the following guidelines: Circuit-based marking overrides class-based marking and Border Gateway Protocol (BGP) destination-based marking, through route maps, overrides both circuit-based and class-based marking.

Use the **no** or **default** form of this command to return to the default behavior of dropping packets that exceed the excess burst tolerance.

1.91.6 Examples

The following example configures the policy to mark all packets that exceed the configured burst tolerance with an IP precedence value of `3`:

```
[local]Redback(config)#qos policy protection1 policing
[local]Redback(config-policy-policing)#rate 10000 burst 100000
excess-burst 120000
[local]Redback(config-policy-rate)#violate mark precedence 3
```

1.92 violate mark priority

violate mark priority {*group-num* | **ignore**} [{**drop-precedence** {*group-num* | **ignore**} | **af-drop** *drop-value*}]

{**no** | **default**} **violate mark priority**

1.92.1 Purpose

Marks packets that exceed the excess burst tolerance with a packet descriptor (PD) priority number, a drop-precedence value, or both, while leaving the packet's IP header Differentiated Services Code Point (DSCP) value unmodified.

1.92.2 Command Mode

- Policy class rate configuration
- Policy rate configuration

1.92.3 Syntax Description

group-num	Priority group number. The range of values is 0 to 7.
ignore	Specifies that the internal packet descriptor (PD) priority or drop-precedence value is not modified.
drop-precedence	Optional. Enables you to specify a setting for either the drop-precedence portion of the PD quality of service (QoS) field or the PD QoS priority number, or both.
af-drop *drop-value*	Optional. Target internal drop-precedence value in two-bit format; leaves the least significant bit unmodified. The range of values is 1 to 3.

1.92.4 Default

Packets exceeding the excess burst tolerance are dropped.

1.92.5 Usage Guidelines

Use the **violate mark priority** command to mark packets that exceed the excess burst tolerance with a PD QoS PD QoS priority number, a drop-precedence value, or both, while preserving the packet's IP header. To configure the excess burst tolerance, enter the **rate** command (in policy ACL class, metering policy, or policing policy configuration mode).

A PD QoS priority number is an internal value used by the SmartEdge router to determine into which egress queue the inbound packet should be placed. The type of service (ToS) value, Differentiated Services Code Point (DSCP) value, and Multiprotocol Label Switching (MPLS) experimental (EXP) bits are unchanged by this command. The actual queue number depends on the number of queues configured on the circuit. For more information, see the **num-queues** command in *Commands: mp through n*.

The SmartEdge router uses the factory preset, or default, mapping of a PD QoS priority number to queue, according to the number of queues configured on a circuit; see Table 6.

Table 6 Default Mapping of Priority Groups

Priority Group	8 Queues	4 Queues	2 Queues	1 Queue
0	Queue 0	Queue 0	Queue 0	Queue 0
1	Queue 1	Queue 1	Queue 1	Queue 0
2	Queue 2	Queue 1	Queue 1	Queue 0
3	Queue 3	Queue 2	Queue 1	Queue 0
4	Queue 4	Queue 2	Queue 1	Queue 0

Table 6 Default Mapping of Priority Groups

Priority Group	8 Queues	4 Queues	2 Queues	1 Queue
5	Queue 5	Queue 2	Queue 1	Queue 0
6	Queue 6	Queue 2	Queue 1	Queue 0
7	Queue 7	Queue 3	Queue 1	Queue 0

Only one mark instruction can be in effect at a time. To change the mark instruction, enter the **violate mark priority** command, specifying a new value for the ***group-num*** argument. This supersedes the value previously configured.

Caution!

Risk of overriding configurations. The SmartEdge router checks for and applies marking in a specific order. To reduce the risk, remember the following guidelines: Circuit-based marking overrides class-based marking and Border Gateway Protocol (BGP) destination-based marking, through route maps, overrides both circuit-based and class-based marking.

Note: By default, the SmartEdge router assigns a PD QoS priority number to each egress queue according to the number of queues configured on a circuit. You can override the default mapping of packets into egress queues by creating a customized queue priority map using the **qos queue-map** command (in global configuration mode).

Use the **no** or **default** form of this command to return to the default behavior of dropping packets that exceed the excess burst tolerance.

1.92.6 Examples

The following example configures the policy to mark all packets that exceed the configured burst tolerance with a PD QoS priority number of 3:

```
[local]Redback(config)#qos policy protection1 policing
[local]Redback(config-policy-policing)#rate 10000 burst 100000
excess-burst 120000
[local]Redback(config-policy-rate)#violate mark priority 3
```

1.93 violate no-action

violate no-action

```
{no | default} violate no-action
```

1.93.1 Purpose

Specifies that no action is taken on packets that exceed the configured excess burst tolerance.

1.93.2 Command Mode

- Policy class rate configuration
- Policy rate configuration

1.93.3 Syntax Description

This command has no keywords or arguments.

1.93.4 Default

Packets exceeding the excess burst tolerance are dropped.

1.93.5 Usage Guidelines

Use the **`violate no-action`** command to specify that no action is taken on packets that exceed the excess burst tolerance.

To configure the excess burst tolerance, enter the **`rate`** command (in policy ACL class, metering policy, or policing policy configuration mode).

Caution!

Risk of overriding configurations. The SmartEdge router checks for and applies marking in a specific order. To reduce the risk, remember the following guidelines: Circuit-based marking overrides class-based marking and Border Gateway Protocol (BGP) destination-based marking, through route maps, overrides both circuit-based and class-based marking.

Use the **`no`** or **`default`** form of this command to return to the default behavior of dropping packets that exceed the excess burst tolerance.

1.93.6 Examples

The following example configures the policy to take no action on packets that exceed the configured excess burst tolerance:

```
[local]Redback(config)#qos policy protection1 policing
[local]Redback(config-policy-policing)#rate 10000 burst 100000
excess-burst 120000
[local]Redback(config-policy-rate)#violate no-action
```

1.94 virtual-address

virtual-address *ip-addr*

no virtual-address *ip-addr*

1.94.1 Purpose

Configures the virtual IP address for the Virtual Router Redundancy Protocol (VRRP) interface.

1.94.2 Command Mode

VRRP configuration

1.94.3 Syntax Description

ip-addr	Virtual IP address.

1.94.4 Default

None

1.94.5 Usage Guidelines

Use the **virtual-address** command to configure the virtual IP address for the VRRP interface. You can configure multiple virtual IP addresses for a single VRRP instance.

Note: For a VRRP owner router, the virtual address must be match one of the interface IP addresses on which the owner VRRP is configured.

Caution!

Risk of conflicting IP addresses. Static Address Resolution Protocol (ARP) configuration takes precedence over a VRRP association of a virtual medium access control (MAC) address with a virtual address. To reduce the risk, avoid configuring static ARP entries for VRRP virtual addresses.

Use the **no** form of this command to remove the virtual IP address.

1.94.6 Examples

The following example configures a router running VRRP on interface `eth1` and assigns a virtual IP address of `10.1.1.2:`

```
[local]Redback(config-ctx)#interface eth1
[local]Redback(config-if)#ip address 10.1.1.2/24
[local]Redback(config-if)#vrrp 1 owner
[local]Redback(config-vrrp)#virtual-address 10.1.1.2
```

1.95 virtual-link

virtual-link {*transit-id* | *transit-addr*} *virtual-endpoint-addr*

no virtual-link {*transit-id* | *transit-addr*} *virtual-endpoint-addr*

1.95.1 Purpose

In OSPF area configuration mode, creates an Open Shortest Path First (OSPF) virtual link through the specified transit area and enters OSPF virtual link configuration mode.

In OSPF3 area configuration mode, creates an OSPF Version 3 (OSPFv3) virtual link through the specified transit area and enters OSPF3 virtual link configuration mode.

1.95.2 Command Mode

- OSPF area configuration
- OSPF3 area configuration

1.95.3 Syntax Description

transit-id	Transit area ID for the virtual link specified as a 32-bit number.
transit-addr	Transit area IP address for the virtual link in the form *A.B.C.D*.
virtual-endpoint-addr	Router ID of the virtual link endpoint in the form *A.B.C.D*.

1.95.4 Default

There are no predefined virtual links for the area.

1.95.5 Usage Guidelines

Use the **virtual-link** command in OSPF area configuration mode to create an OSPF virtual link through the specified transit area and enters OSPF virtual link configuration mode.

Use the **virtual-link** command in OSPF3 area configuration mode to create an OSPFv3 virtual link through the specified transit area and enters OSPF3 virtual link configuration mode.

Virtual links can be configured between any two backbone routers that have an interface to a common non-backbone area. Virtual links belong to the backbone. The protocol treats two routers joined by a virtual link as if they were connected by an unnumbered point-to-point backbone network.

Virtual links can only be configured in the backbone area (area ID=0), and the transit area cannot be the backbone area.

Use the **no** form of this command to remove the virtual link.

For more information on OSPF virtual links, see RFC 2328, *OSPF Version 2*.

For more information on OSPFv3 virtual links, see RFC 2740, *OSPF for IPv6*.

1.95.6 Examples

The following example configures a virtual link through area 1, with a virtual link endpoint of 30.30.30.30, and enters OSPF virtual link configuration mode:

```
[local]Redback(config-ospf)#router ospf 1
[local]Redback(config-ospf)#area 0
[local]Redback(config-ospf-area)#virtual-link 1 30.30.30.30
[local]Redback(config-ospf-virt-link)#
```

1.96 vpls

vpls

no vpls

1.96.1 Purpose

Enables Virtual Private LAN Services (VPLS) on a bridge and enters VPLS configuration mode.

1.96.2 Command Mode

Bridge configuration

1.96.3 Syntax Description

This command has no keywords or arguments.

1.96.4 Default

VPLS is not enabled on the bridge.

1.96.5 Usage Guidelines

Use the **vpls** command to enable VPLS on a bridge and enter VPLS configuration mode.

Use the **no** form of this command to disable VPLS on the bridge.

1.96.6 Examples

The following example enables VPLS on the `to-pe4` bridge and enter VPLS configuration mode:

```
[local]Redback#config
[local]Redback(config)#context local
[local]Redback(config-ctx)#bridge to-pe4
[local]Redback(config-bridge)#vpls
```

1.97 vpls profile

vpls profile ***prof-name***

```
no vpls profile prof-name
```

1.97.1 Purpose

Creates a new Virtual Private LAN Services (VPLS) profile, or selects an existing one for modification, and enters VPLS profile configuration mode.

1.97.2 Command Mode

Global configuration

1.97.3 Syntax Description

prof-name	Name of the VPLS profile (maximum of 40 characters).

1.97.4 Default

None

1.97.5 Usage Guidelines

Use the **vpls profile** command to create a new VPLS profile, or select an existing one for modification, and enter VPLS profile configuration mode. VPLS profiles are used to configure one or more neighbors to which a VPLS instance can establish peering connections. All neighbors configured within a VPLS profile are referenced by the VPLS profile name, which is unique in the system.

The VPLS profile is referenced from the VPLS instance configuration. Multiple VPLS instances can apply (share) the same VPLS profile. If a profile is updated, then all instances of its usage use the changed attributes. Conflicts arising, due to the updated VPLS profile in the VPLS instances, do not result in rejecting the VPLS profile or the updates; the individual VPLS instances handle these conditions.

1.97.6 Examples

The following example creates the `prof-123` VPLS profile and enters VPLS profile configuration mode:

```
[local]Redback#config
[local]Redback(config)#vpls profile prof-123
[local]Redback(config-vpls-profile)#
```

1.98 vpn

vpn [**domain-id** ***ip-addr***] {**domain-tag** ***tag-name*** | **local-as** ***asn***}

no vpn

1.98.1 Purpose

Enables an Open Shortest Path First (OSPF) instance within a Virtual Private Network (VPN) context to treat redistributed Border Gateway Protocol (BGP) routes as VPN routes.

1.98.2 Command Mode

OSPF router configuration

1.98.3 Syntax Description

domain-id ***ip-addr***	Optional. Domain ID value. Used to determine whether redistributed BGP routes should be treated as VPN routes and be handled differently than an OSPF instance configured within a VPN context; the default value is 0.
domain-tag ***tag-name***	Domain tag. Used for type 5 link-state advertisements (LSAs) corresponding to redistributed BGP routes within the VPN domain. Either the ***tag-name*** or ***asn*** argument must be specified.
local-as ***asn***	Autonomous system number (ASN), 2-byte. Used to formulate the tag for type 5 LSAs corresponding to redistributed BGP routes with the same VPN. Either the ***tag-name*** or ***asn*** argument must be specified, but the ***tag-name*** argument overrides the use of the ***asn*** argument to formulate the tag.

1.98.4 Default

OSPF VPN treatment of routes is disabled.

1.98.5 Usage Guidelines

Use the **vpn** command to enable an OSPF instance within a VPN context to treat redistributed BGP routes as VPN routes.

When a customer edge (CE) site is connected to multiple areas, the CE router’s connection to a provider edge (PE) router should be in area 0 to allow correct handling of summary LSAs.

Note: The **vpn** command is useful only when OSPF is used for PE-to-CE routing.

Use the **no** form of this command to disable the OSPF VPN treatment of routes.

1.98.6 Examples

The following example configures an OSPF instance within a VPN context to treat redistributed BGP routes with domain IDs equal to `1.1.1.1` as VPN routes:

```
[local]Redback(config-ospf)#vpn domain-id 1.1.1.1 domain-tag 0xfeedacee
```

1.99 vpn-context

vpn-context ***ctx-name***

no vpn-context ***ctx-name***

1.99.1 Purpose

Specifies the context in which the IP-in-IP tunnel or Generic Routing Encapsulation (GRE) tunnel to this home agent (HA) peer is terminated.

1.99.2 Command Mode

HA peer configuration

1.99.3 Syntax Description

ctx-name	Context in which the IP-in-IP tunnel or GRE tunnel to this HA peer is terminated and in which the IP routes are added for the mobile nodes (MNs) that are registered with this HA peer.

1.99.4 Default

None

1.99.5 Usage Guidelines

Use the **vpn-context** command to specify the context in which the IP-in-IP tunnel or GRE tunnel to this HA peer is terminated. The HA peers can share a context if they use public IP addresses or if their private IP addresses do not overlap. HA peers with overlapping private IP addresses must each have their own context.

Use the **no** form of this command to specify the default condition.

1.99.6 Examples

The following example specifies the `ha-vpn1` context for the MNs associated with this HA peer:

```
[local]Redback(config)#context fa
[local]Redback(config-ctx)#router mobile-ip
[local]Redback(config-mip)#foreign-agent
[local]Redback(config-mip-fa)#ha-peer 172.16.2.1
[local]Redback(config-mip-hapeer)#vpn-context ha-vpn1
```

1.100 vrrp

vrrp ***router-id*** {**owner** | **backup**}

no vrrp ***router-id***

1.100.1 Purpose

Configures a virtual router as an owner or backup router, assigns a Virtual Router Redundancy Protocol (VRRP) ID and enters VRRP configuration mode.

1.100.2 Command Mode

Interface configuration

1.100.3 Syntax Description

router-id	Virtual router ID. The range of values is 1 to 255.
owner	Configures the virtual router as an owner.
backup	Configures the virtual router as a backup in the event an owner virtual router goes down.

1.100.4 Default

None

1.100.5 Usage Guidelines

Use the **vrrp** command to configure a virtual router as an owner or backup router, assign a VRRP ID, and to enter VRRP configuration mode.

For more information on VRRP, see RFC 2338, *Virtual Router Redundancy Protocol*.

Note: Each virtual router corresponding to an interface that is bound to 802.1Q circuits and that uses the same Ethernet port must have a unique virtual router ID. If multiple interfaces are bound to 802.1Q circuits associated with the same Ethernet port, and there are virtual routers with duplicate router identifiers, only one of the virtual routers will be operational.

Use the **no** form of this command to remove the virtual router.

1.100.6 Examples

The following example configures an owner virtual router with a VRRP ID of `23`:

```
[local]Redback(config-if)#vrrp 23 owner
[local]Redback(config-vrrp)#
```

1.101 weight

weight ***weight***

no weight

1.101.1 Purpose

Assigns a relative weight that is used to calculate a traffic ratio for all circuits to which you attach this policy.

1.101.2 Command Mode

PWFQ policy configuration

1.101.3 Syntax Description

weight	Relative weight that is assigned to any circuit to which you attach this policy. The range of values is 1 to 4096.

1.101.4 Default

All circuits to which this policy is attached have the same weight.

1.101.5 Usage Guidelines

Use the **weight** command to assign a relative weight that is used to calculate a traffic ratio for all circuits to which you attach a policy.

You can assign a relative weight or a minimum absolute rate using the **rate** command (in PWFQ policy configuration mode), but you cannot do both. The relative weight and minimum absolute rate are mutually exclusive.

You can, however, assign a relative weight (using this command) and a maximum absolute rate using the **rate** command (in PWFQ policy configuration mode).

Use the **no** form of this command to specify the default condition.

1.101.6 Examples

The following example shows how to assign guaranteed bandwidth to two different policies bound to two 802.1Q PVCs that share the same port. First, you configure two 802.1Q PVCs, set encapsulation to `dot1q`, and configure policy queuing. Next, you bind one 802.1Q PVC to QoS policy A and the other 802.1Q PVC to QoS policy B. You want to guarantee 1200/(1200+1500)% of the available bandwidth to QoS policy A and 1500/(1200 + 1500)% of the available bandwidth to QoS policy B:

```
[local]Redback(config)#qos policy A pwfq
[local]Redback(config-policy-pwfq)#weight 1200
[local]Redback(config)#qos policy B pwfq
[local]Redback(config-policy-pwfq)#weight 1500
```

1.102 xc

For a cross-connection group, the command syntax is:

xc {**lg** *lg-name-in* | **lg id** *id-num-in* | *slot-in/port-in*} *circuit-id-in* [**through** *end-circuit-id-in*] [*circuit-type*] **to**

{**lg** *lg-name-out* | **lg id** *id-num-out* | *slot-out/port-out*} *circuit-id-out* [**through** *end-circuit-id-out*] [*circuit-type*] [**interworking**] [**vc-id** *vc-id*] [**profile** *profile-name*] [**backup**]

no xc {**lg** *lg-name-in* | **lg id** *id-num-in* | *slot-in/port-in*} *circuit-id-in* [**through** *end-circuit-id-in*] [*circuit-type*] **to**

```
{lg lg-name-out | lg id id-num-out | slot-out/port-
out} circuit-id-out [ through end-circuit-id-out]
[circuit-type][interworking] [vc-id vc-id] [profile
profile-name] [backup]
```

1.102.1 Purpose

Creates a cross-connection between an two parent circuits, two child circuits, or a parent circuit and a child circuit.

1.102.2 Command Mode

XC group configuration

1.102.3 Syntax Description

`lg`	Specifies that the circuit to be cross-connected is a constituent circuit in an access link group.
`lg` *`lg-name-in`*	Specifies the name of an access link group to be cross-connected inbound.
`lg id` *`id-num-in`*	Specifies the ID of an access link group be cross-connected inbound. The link-group ID is automatically assigned by the SmartEdge router, and is displayed when you enter the **`show link-group detail`** command.
`slot-in`	Chassis slot number of the traffic card with the ATM or 802.1Q PVC for which a cross-connection is to be specified.
`port-in`	Port number of the circuit for which a cross-connection is to be specified.
`circuit-id-in`	Identifier for the circuit to be cross-connected, according to one of the constructs listed in Table 7.
`through` *`end-circuit-id-in`*	Optional. Last circuit identifier in a range of circuits to be cross-connected.

circuit-type	Optional. Circuit type for which a cross-connection is to be specified. Required only if the specified circuit is an IPv6oE or PPPoE child circuit. Not specified when the specified circuit is encapsulated as **bridge1483**, **1qtunnel**, **dot1q**, or **raw**, or **route1483**. If the circuit is a child circuit type, according to one of the following keywords: • **ipv6oe**—Specifies an IP Version 6 over Ethernet (IPv6oE)-encapsulated circuit. • **pppoe**—Specifies a Point-to-Point Protocol (PPP) over Ethernet (PPPoE)-encapsulated circuit. If omitted, cross-connects all child circuits on the specified circuit.
to	Indicates the start of the outbound circuit specification.
lg *lg-name-out*	Specifies the name of an access link group to be cross-connected outbound.
lg id *id-num-out*	Specifies the ID of an access link group be cross-connected outbound.
slot-out	Chassis slot number of the traffic card for which a cross-connection is to be specified.
port-out	Port number of the circuit for which a cross-connection is to be specified.
circuit-id-out	Identifier of the circuit to be cross-connected, according to one of the values listed in Table 7.
through *end-circuit-id-out*	Optional. Last circuit identifier for a range of circuits to be cross-connected.
interworking	Optional. Specifies an interworking cross-connection.
vc-id *vc-id*	Optional. Specify the VC identifier associated with the cross-connection.
profile *profile-name*	Optional. Associate an L2VPN profile to an XC. When you attach an L2VPN profile to an XC, the LSP that is specified in the profile is automatically mapped to the XC you are configuring.
backup	Enters primary L2VPN XC configuration mode, where you can configure the VC ID and peer address for a backup XC.

1.102.4 Default

No cross-connections are defined.

1.102.5 Usage Guidelines

Use the **`xc`** command to create a cross-connection between two parent circuits, two child circuits, or a parent circuit and a child circuit. The parent circuits can be Asynchronous Transfer Mode (ATM) permanent virtual circuits (PVCs), 802.1Q PVCs, or 802.1Q tunnels. Table 8 lists all possible supported combinations of parent and child circuits. You must include the **`bind bypass`** command (in ATM PVC or dot1q PVC configuration mode) before you make the cross-connection.

Note: To create a cross connection for a L2VPN, see *Configuring L2VPN*.

Note: All cross-connections are bidirectional. In the syntax statements, the suffixes *`in`* and *`out`* serve only to distinguish the two circuits to be cross-connected; either circuit can be designated the *`in`* circuit when you configure the cross-connection.

Note: The SmartEdge 100 router limits the value of the *`slot`* argument to 2.

Note:

The value for the *`port`* argument on the SmartEdge 100 router is either of the following:

- For a native port, it is 1 or 2.
- For a MIC port, it depends on the MIC and MIC slot in which it is installed.

Table 7 lists the values for the *`circuit-id`* argument.

Table 7 Values for the circuit-id Argument

Construct	Description
`vlan-id` *`vlan-id`*	Virtual LAN (VLAN) tag value for an 802.1Q tunnel or PVC. The *`vlan-id`* argument is one of the following constructs: • *`pvc-vlan-id`*—VLAN tag value of a PVC that is not within an 802.1Q tunnel. • *`tunl-vlan-id`*—VLAN tag value of a tunnel. • *`tunl-vlan-id:pvc-vlan-id`*—VLAN tag value for the tunnel followed by the VLAN tag value for the PVC within the tunnel. The range of values for either VLAN tag value is 1 to 4095.
`vpi-vci` *`vpi vci`*	Virtual path identifier (VPI) and virtual circuit identifier (VCI) for an ATM PVC. The range of values is 0 to 255 and 1 to 65535, respectively.

Table 7 Values for the circuit-id Argument

Construct	Description
pppoe [*pppoe-id*]	PPPoE encapsulation and optional session identifier
dlci *pppoe-id*	DLCI number or ID

If you specify the **through** *end-vci-in* and **through** *end-vci-out* constructs, the number of ATM PVCs in the input range must match the number specified by the output range.

If you specify the **through** *end-vlan-in* and **through** *end-vlan-out* constructs, the number of 802.1Q PVCs in the input range must match the number specified by the output range.

Note: If you are cross-connecting a range of PVCs, you can use multiple **xc** commands to remove specific PVCs from the range of cross-connected PVCs. See the second example in Section 1.102.6 on page 154.

Use the **lg** *lg-name-in* and **lg** *lg-name-out* constructs to specify a constituent circuit in an access link group; Table 8 lists the types of cross-connections that you can configure with 802.1Q PVCs and tunnels within an access link group.

Table 8 Access Link Group Cross-Connections

Circuit Inside Access Link Group	Circuit Outside Access Link Group
802.1Q PVC	802.1Q PVC
802.1Q PVC inside an 802.1Q tunnel	802.1Q PVC inside an 802.1Q tunnel
802.1Q PVC	802.1Q PVC inside an 802.1Q tunnel
802.1Q PVC inside an 802.1Q tunnel	802.1Q PVC
802.1Q PVC	ATM RFC 1483 bridged PVC
802.1Q PVC inside an 802.1Q tunnel	ATM RFC 1483 bridged PVC

The two circuits need not be constituent circuits in the same access link group.

The *lg-name-in* and *lg-name-out* arguments specify the name of the access link groups that you created with the **link-group** command (in global configuration mode).

Table 9 lists the supported combinations of parent and child circuit types and the traffic that is cross-connected. You can cross-connect tunnels even if they contain PVCs within them.

Note: Any PPPoE or IPv6oE child circuit on a parent can be cross-connected to a child on another parent with the same type of encapsulation. The inbound IPoE parent circuits are usually terminated and routed to any IP-type port or circuit, such as a Gigabit Ethernet port, but can be cross-connected instead, or the traffic can be dropped. IPv6oE child circuits must be cross-connected.

Note: Any circuit with raw encapsulation must be cross-connected.

Note: You cannot bind ATM or 802.1Q circuits with raw encapsulation to either an interface or subscriber.

Note: The default traffic type for any PVC or tunnel is IP over Ethernet (IPoE).

Table 9 Supported Cross-Connections and Their Encapsulations

Parent Circuit Type for Bidirectional Cross-Connection	First Parent Circuit Encapsulation	Second Parent Circuit Encapsulation	Cross-Connected Circuit Traffic

ATM PVC-to-ATM PVC	`bridge1483`
	`pppoe`
	`raw`
	`route1483`
		`bridge1483`
		`pppoe`
		`raw`
		`route1483`
			Parent-to-parent
			Parent-to-parent
			Parent-to-parent
			Parent-to-parent

	`multi`	`multi`	IPoE-to-IPoE parent-to-parent IPv6oE-to-IPv6oE child-to-child PPPoE-to-PPPoE child-to-child
	`multi`	`pppoe`	PPPoE-to-IPoE child-to-parent
ATM PVC-to-802.1Q PVC	`bridge1483`	`dot1q` `raw`	Parent-to-parent Parent-to-parent
	`multi`	`dot1q`	IPoE-to-IPoE parent-to-parent IPv6oE-to-IPoE child-to-parent PPPoE-to-IPoE child-to-parent

Table 9 Supported Cross-Connections and Their Encapsulations

Parent Circuit Type for Bidirectional Cross-Connection	First Parent Circuit Encapsulation	Second Parent Circuit Encapsulation	Cross-Connected Circuit Traffic
	`multi`	`multi`	IPoE-to-IPoE parent-to-parent IPv6oE-to-IPv6oE child-to-child PPPoE-to-PPPoE child-to-child
	`multi`	`pppoe`	PPPoE-to-PPPoE child-to-parent
	`pppoe`	`dot1q`	PPPoE-to-IPoE parent-to-parent
	`pppoe`	`pppoe`	PPPoE-to-PPPoE parent-to-parent
	`pppoe`	`multi`	PPPoE-to-PPPoE parent-to-child
ATM PVC-to-802.1Q PVC in 802.1Q tunnel	`bridge1483`	`dot1q` `raw`	Parent-to-IPoE parent Parent-to-parent
	`multi`	`dot1q`	IPoE-to-IPoE parent-to-parent IPv6oE-to-IPoE child-to-parent PPPoE-to-IPoE child-to-parent
	`multi`	`multi`	IPoE-to-IPoE parent-to-parent IPv6oE-to-IPv6oE child-to-child PPPoE-to-PPPoE child-to-child
	`multi`	`pppoe`	PPPoE-to-PPPoE child-to-parent
	`pppoe`	`dot1q`	PPPoE-to-IPoE parent-to-parent
	`pppoe`	`pppoe`	PPPoE-to-PPPoE parent-to-parent

Table 9 Supported Cross-Connections and Their Encapsulations

Parent Circuit Type for Bidirectional Cross-Connection	First Parent Circuit Encapsulation	Second Parent Circuit Encapsulation	Cross-Connected Circuit Traffic
	`pppoe`	`multi`	PPPoE-to-PPPoE parent-to-child
802.1Q PVC-to-802.1Q PVC	`dot1q`	`dot1q`	IPoE-to-IPoE parent-to-parent
	`multi`	`dot1q`	IPoE-to-IPoE parent-to-parent IPv6oE-to-IPoE child-to-parent PPPoE-to-IPoE child-to-parent
	`multi`	`multi`	IPoE-to-IPoE parent-to-parent IPv6oE-to-IPv6oE child-to-child PPPoE-to-PPPoE child-to-child
802.1Q PVC-to-802.1Q PVC	`raw`	`raw`	Parent-to-parent
802.1Q PVC-to-802.1Q PVC in 802.1Q tunnel	`dot1q`	`dot1q`	IPoE-to-IPoE parent-to-parent
	`multi`	`dot1q`	IPoE-to-IPoE parent-to-parent IPv6oE-to-IPoE child-to-parent
	`multi`	`multi`	IPoE-to-IPoE parent-to-parent IPv6oE-to-IPv6oE child-to-child
	`raw`	`raw`	Parent-to-parent
802.1Q PVC in 802.1Q tunnel-to-802.1Q PVC in 802.1Q tunnel	`1qtunnel`	`1qtunnel`	IPoE-to-IPoE parent-to-parent
	`multi`	`1qtunnel`	IPoE-to-IPoE parent-to-parent IPv6oE-to-IPoE child-to-parent

Table 9 Supported Cross-Connections and Their Encapsulations

Parent Circuit Type for Bidirectional Cross-Connection	First Parent Circuit Encapsulation	Second Parent Circuit Encapsulation	Cross-Connected Circuit Traffic
	`multi`	`multi`	IPoE-to-IPoE parent-to-parent IPv6oE-to-IPv6oE child-to-child
	`raw`	`raw`	IPoE-to-IPoE parent-to-parent

The traffic that flows through the cross-connection between the inbound and outbound circuits depends on the encapsulation specified for the inbound and outbound circuits:

- For parent circuits, the encapsulation type specified for the circuit filters the type of packets passed through the cross-connection, with only the inbound type of encapsulated packets being passed to the outbound circuit and only the outbound type of encapsulated packets being accepted by the outbound circuit.

- For child circuits, the **`circuit protocol`** command (in protocol configuration mode), acts as the filter, with only the specified type of encapsulated packets being passed from the inbound child circuit to the outbound child circuit and only the outbound type of encapsulated packets being accepted by the outbound child circuit.

Use the **`interworking`** keyword only if you are cross-connecting an ATM PVC with RFC 1483, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*-routed encapsulation to a multiprotocol 802.1Q PVC. In this instance, only the IPv4 packets are forwarded to the 802.1Q PVC.

Use the **`show bypass`** command (in any mode) to display information about one or more cross-connected circuits.

Use the **`no`** form of this command to delete the cross-connection between two circuits on one or more ATM PVCs or 802.1Q PVCs.

1.102.6 Examples

The following example shows how to specify cross-connections between the inbound PPPoE circuits on a range of ATM PVCs with VCIs `1` to `10` on port `1` of the ATM card in slot `3` and the outbound PPPoE circuits on a range of ATM PVCs with VCIs `101` to `110` on port 1 of the ATM card in slot `9`; both ranges of PVCs use VPI `32`. In this example, only PPPoE-encapsulated packets are passed from the inbound circuits and only PPPoE-encapsulated packets are accepted by the outbound circuits:

```
[local]Redback(config-xc-group)#xc 3/1 vpi-vci 1 101 through 110 pppoe
to 9/1 vpi-vci 2 101 through 110 pppoe
```

The following example shows how to remove the VCI `101` from the range of PVCs specified in the preceding example:

```
[local]Redback(config-xc-group)#xc 3/1 vpi-vci 1 102 through 110 pppoe
to 9/1 vpi-vci 2 101 through 110 pppoe
```

The following example shows how to create two ATM PVCs on ATM OC ports with an existing profile, `ubr`, and encapsulated with `raw` mode, and cross-connect them. In this example, any type of packet is passed from the inbound circuit and accepted by the outbound circuit:

```
[local]Redback(config)#port atm 3/1
[local]Redback(config-atm-oc)#atm pvc 0 32 profile ubr encapsulation raw
[local]Redback(config-atm-pvc)#exit
[local]Redback(config-atm-oc)#exit
[local]Redback(config)#port atm 4/2
[local]Redback(config-atm-oc)#atm pvc 1 55 profile ubr encapsulation raw
[local]Redback(config-atm-pvc)#bind bypass
[local]Redback(config-atm-pvc)#exit
[local]Redback(config-atm-oc)#exit
[local]Redback(config-xc-group)#xc 3/1 vpi-vci 0 32 to 4/2 vpi-vci 1 55
```

The following example shows how to create a multiprotocol ATM PVC on an ATM OC port, an 802.1Q PVC, and cross-connects the two circuits. In this example, all dot1q-encapsulated packets are passed from the inbound circuit but only IPoE-encapsulated packets are accepted by the outbound circuit:

```
!Create the ATM PVC and its IPoE circuit
[local]Redback(config)#port atm 3/1
[local]Redback(config-atm-oc)#atm pvc 2 115 profile test encapsulation
multi
[local]Redback(config-atm-pvc)#bind bypass
[local]Redback(config-atm-pvc)#exit
[local]Redback(config-atm-oc)#exit
!Create the 802.1Q PVC
[local]Redback(config)#port ethernet 2/1
[local]Redback(config-port)#encapsulation dot1q
[local]Redback(config-port)#dot1q pvc 1
[local]Redback(config-dot1q-pvc)#bind bypass
[local]Redback(config-dot1q-pvc)#exit
[local]Redback(config-port)#exit

!Cross-connect the ATM PVC to the 802.1Q PVC
[local]Redback(config-xc-group)#xc 3/1 vpi-vci 2 115 to 2/1 vlan-id 1
```

The following example shows how to create an ATM PVC with RFC 1483 routed encapsulation on an ATM OC port, a multiprotocol 802.1Q PVC, and an interworking cross-connection between the two circuits. Only inbound IPoE (IPv4) packets are forwarded from the inbound circuit to the outbound circuit:

```
!Create the ATM PVC and its IPoE circuit
[local]Redback(config)#port atm 3/1
[local]Redback(config-atm-oc)#atm pvc 3 110 profile test encapsulation
route1483
[local]Redback(config-atm-pvc)#bind bypass
[local]Redback(config-atm-pvc)#exit
[local]Redback(config-atm-oc)#exit

!Create the 802.1Q PVC
[local]Redback(config)#port ethernet 2/1
[local]Redback(config-port)#encapsulation dot1q
[local]Redback(config-port)#dot1q pvc 3 encapsulation multi
[local]Redback(config-dot1q-pvc)#bind bypass
[local]Redback(config-dot1q-pvc)#exit
[local]Redback(config-port)#exit

!Cross-connect the ATM PVC to the 802.1Q PVC
[local]Redback(config-xc-group)#xc 3/1 vpi-vci 3 110 to 2/1 vlan-id 5
12-13-interworking
```

The following example shows how to create two cross-connected 802.1Q PVC circuits with raw encapsulation:

```
! Create the first 802.1Q PVC circuits with raw encapsulation
[local]Redback(config)#port ethernet 1/1
[local]Redback(config-port)#encapsulation dot1q
[local]Redback(config-port)#dot1q pvc 100 encapsulation raw
[local]Redback(config-dot1q-pvc)#bind bypass
[local]Redback(config-dot1q-pvc)#exit
[local]Redback(config-port)exit

!Create the second 802.1Q PVC circuits with raw encapsulation
[local]Redback(config)#port ethernet 2/2
[local]Redback(config-port)#encapsulation dot1q
[local]Redback(config)#dot1q pvc 200 encapsulation raw
[local]Redback(config-dot1q-pvc)#bind bypass
[local]Redback(config-dot1q-pvc)#exit
[local]Redback(config-port)exit

!Cross-connect the two circuits
[local]Redback(config-xc-group)#xc 1/1 vlan-id 100 to 2/2 vlan-id 200
```

1.103 xc ds0s

xc ds0s *slot/port*:[*ds3-channel*:]*ds1-channel*:*ds0-channel-group* [**through** *ds0-channel-group*] **ces** {[**udp** *sport:dport* **profile** *name*] | [**vpn-label** *static-label* **profile** *profile-name*]}

no xc ds0s *slot/port*:[*ds3-channel*:]*ds1-channel*:*ds0-channel-group* [**through** *ds0-channel-group*] **ces** {[**udp** *sport:dport* **profile** *name*] | [**vpn-label** *static-label* **profile** *profile-name*]}

1.103.1 Purpose

Configures the pseudowire cross connect for a CESoPSN circuit.

1.103.2 Command Mode

L2VPN XC Group Config Mode.

1.103.3 Syntax Description

slot:port	Slot and port of the cross connect. Combination of source and destination UDP port.
ds3-channel	Optional DS3 channel of the cross connect.
ds1-channel	The DS1 channel of the cross connect.
ds0-channel-group	The DS0 channel group ID of the cross connect. This is the lowest numbered timeslot in the group.
through *ds0-channel-group*	The end of the range of the DS0 channel group for the cross connect.
udp *sport:dport*	The combination of source and destination UDP port of the cross connect. Source port range is 1024 to 65535. Destination port range is 1 to 65535.
vpn-label *static-label*	The static VPN label for the cross connect. Range is from 4096 to 65535
profile *profile-name*	The profile to reference. This is optional if peer is specified.

1.103.4 Default

No cross-connect is created.

1.103.5 Usage Guidelines

If SONET port is VT1.5 mapped, then ***slot/port*:[*ds1-channel* | *e1-channel*]:*ds0-channel-group***.

On executing keyword **through**, the **sport** and **dport** will be incremented relatively based on the DS0 channel groups in the **through** range. The range provided here by the **through** keyword represents each independent DS0 channel group. Configure channel groups appropriately, as the system does not check internally on existence of a channel group. If the channel group does not exist, the PW is DOWN.

For example **2/1:1:1:1 through 3 udp 4:5** would indicate that there are three independent DS0 channel groups, that is, DS0 channel group 1 , DS0 channel group 2 and DS0 group 3; and each DS0 channel group will have its independent cross connects with UDP ports 4:5, 5:6, and 6:7 respectively. No overlap or punching of the range is allowed.

1.103.6 Examples

The following example shows how to configure a cross connect:

```
[local]Redback(config)#context local
Redback(config-ctx)#l2vpn
Redback(config-l2vpn)#xc-group name
Redback(config-l2vpn-xc-group)#xc ds0s 3/4:1:1:1 ces udp 1024:5005 profile name1
Redback(config-l2vpn-xc-group)#xc ds0s 3/4:1:1:3 through 8 ces udp 9000:1006 profile name1
Redback(config-l2vpn-xc-group)#xc ds0s 3/4:3:1:1 ces vpn-label 4097 profile name1
```

1.104 xc [ds1 | e1]

xc [ds1 | e1] ***slot/port*****:[*****ds3-channel*****:]*****e1/ds1-channel*** **[through** ***e1/ds1-channel*****] ces {[udp** ***sport:dport*** **profile** ***name*****] | [vpn-label** ***static-label*** **profile** ***profile-name*****]}**

no xc [ds1 | e1] ***slot/port*****:[*****ds3-channel*****:]*****e1/ds1-channel*** **[through** ***e1/ds1-channel*****] ces {[udp** ***sport:dport*** **profile** ***name*****] | [vpn-label** ***static-label*** **profile** ***profile-name*****]}**

1.104.1 Purpose

Configures the pseudowire cross connect for a SAToP circuit.

1.104.2 Command Mode

L2VPN XC Group Config Mode.

1.104.3 Syntax Description

ds1	Cross connect a DS1 attachment circuit.
e1	Connect a E1 attachment circuit.
slot:port	Slot and port of the cross connect.
ds3-channel	Optional DS3 channel of the cross connect.
e1/ds1-channel	The DS1 channel of the cross connect.
udp ***sport:dport***	The combination of source and destination UDP port of the cross connect. Source port range is 1024 to 65535. Destination port range is 1 to 65535.
vpn-label ***static-label***	The static VPN label for the cross connect. Range is from 4096 to 65535
profile ***profile-name***	The profile to reference. This is optional if peer is specified.

1.104.4 Default

No cross-connect is created.

1.104.5 Usage Guidelines

If SONET port is VT1.5 mapped, then ***slot/port*:[*ds1-channel* | *e1-channel*]**.

On executing keyword **through**, the **sport** and **dport** will be incremented relatively based on the DS1 channels in the **through** range. The range provided here by **through** keyword represents each independent DS1 channel.

For example **2/1:1:1 through 3 udp 4:5** would indicate that there are three independent DS1 channels, that is, DS1 channel 1 , DS1 channel 2 and DS1 channel 3; and each DS1 channel will have its independent cross connects with UDP ports 4:5, 5:6, and 6:7 respectively. No overlap or punching of the range is allowed.

1.104.6 Examples

The following example shows how to configure various types of SAToP cross connects:

```
[local]Redback(config)#context local
Redback(config-ctx)#l2vpn
Redback(config-l2vpn)#xc-group name
Redback(config-l2vpn-xc-group)#xc ds1 3/4:1:1 ces udp 1030:1040 profile name1
Redback(config-l2vpn-xc-group)#xc ds1 3/4:1:1 through 8 ces udp 1050:1060 profile name
Redback(config-l2vpn-xc-group)#xc ds1 3/4:3:1 ces vpn-label 4556 profile name1
Redback(config-l2vpn-xc-group)#xc ds1 3/4:3:2 through 8 ces vpn-label 4556 profile name1
```

1.105 xc-group (global)

xc-group{**default** | ***group-name***}

no xc-group{**default** | ***group-name***}

1.105.1 Purpose

Creates an empty group of cross-connected circuits or selects an existing one and accesses XC group configuration mode.

1.105.2 Command Mode

Global configuration

1.105.3 Syntax Description

default	Selects the default XC group.
group-name	Name of the XC group to be created or selected.

1.105.4 Default

None

1.105.5 Usage Guidelines

Use the **xc-group** command to create an new empty group of cross-connected circuits or select an existing one and access XC group configuration mode.

Use the **no** form of this command to remove the specified XC group from the configuration; removing the group also removes the cross-connections that are members of it.

1.105.6 Examples

The following example shows how to create the `one` XC group and access the XC group configuration mode:

```
[local]Redback(config)#xc-group one
[local]Redback(config-xc-group)#
```

1.106 xc-group (L2VPN)

xc-group{*group-name* | **default**}

no xc-group{*group-name* | **default**}

1.106.1 Purpose

Creates a Layer 2 Virtual Private Network (L2VPN) cross-connection group and enters L2VPN XC group configuration mode.

1.106.2 Command Mode

L2VPN configuration

1.106.3 Syntax Description

group-name	L2VPN cross-connection group name.
default	Creates the default L2VPN cross-connection group.

1.106.4 Default

None

1.106.5 Usage Guidelines

Use the **xc-group** command to create an L2VPN cross-connection group and enter L2VPN XC group configuration mode.

Use the **no** form of this command to delete an L2VPN cross-connection group and all configured cross-connections within the group.

1.106.6 Examples

The following example creates the L2VPN cross-connection group, group2, and enters L2VPN XC group configuration mode:

```
[local]Redback(config-ctx)#l2vpn
[local]Redback(config-l2vpn)#xc-group group2
[local]Redback(config-l2vpn-xc-group)#
```

1.107 xc (L2VPN vc-id)

For an ATM pseudowire cross-connection hosting a single PVC:

xc *slot/port*[**:***chan-num*][**:***sub-chan-num*] [*circuit-id*] **vc-id** *vc-id* **peer** *peer-addr* [**remote encap** *type*] [**exp-bits** *bits-num*] [**cell encap** {**nto1-vcc** | **nto1-vpc**}]

no xc *slot/port*[**:***chan-num*][**:***sub-chan-num*] [*circuit-id*] **vc-id** *vc-id* **peer** *peer-addr* [**remote encap** *type*] [**exp-bits** *bits-num*] [**cell encap** {**nto1-vcc** | **nto1-vpc**}]

For an ATM pseudowire cross-connection hosting multiple PVCs:

xc *slot/port* [*circuit-id*] [**n-to-1**] **vc-id** *vc-id* **peer** *peer-addr* [**cell encap** {**nto1-vcc** | **nto1-vpc**}]

no xc *slot/port* [*circuit-id*] [**n-to-1**] **vc-id** *vc-id* **peer** *peer-addr* [**cell encap** {**nto1-vcc** | **nto1-vpc**}

For an Ethernet pseudowire cross-connection:

xc *slot/port* [*circuit-id*] **vc-id** *vc-id* **peer** *peer-addr* [**remote encap** *type*]

no xc *slot/port* [*circuit-id*] **vc-id** *vc-id* **peer** *peer-addr* [**remote encap** *type*]

For an Ethernet transport-enabled pseudowire cross-connection:

```
xc slot/port circuit-id transport vc-id vc-id peer peer-addr [remote encap type]
```

```
no xc slot/port circuit-id transport vc-id vc-id peer peer-addr [remote encap type]
```

For a link group pseudowire cross-connection:

```
xc lg link-group circuit-id vc-id vc-id peer peer-addr [remote encap type]
```

```
no xc lg link-group circuit-id vc-id vc-id peer peer-addr [remote encap type]
```

For a link group transport pseudowire cross-connection:

```
xc lg link-group circuit-id transport vc-id vc-id peer peer-addr [remote encap type]
```

```
no xc lg link-group circuit-id transport vc-id vc-id peer peer-addr [remote encap type]
```

In all the preceding syntax descriptions, an L2VPN profile can be substituted for the peer arguments; that is, you can specify {**profile** *profile-name* [**backup**]} instead of the {**peer** *peer-addr* [*optional-peer-arguments*]}. In general, to associate an L2VPN profile to an XC, use the following syntax:

```
xc {lg link-group | slot/port} [transport] vc-id vc-id profile profile-name [backup]
```

```
no xc {lg link-group | slot/port} [transport] vc-id vc-id profile profile-name [backup]
```

1.107.1 Purpose

Creates a Label Distribution Protocol (LDP) Layer 2 Virtual Private Network (L2VPN) cross-connection.

1.107.2 Command Mode

L2VPN XC group configuration

1.107.3 Syntax Description

cell encap	Optional. Enables ATM cell mode encapsulation on the specified pseudowire cross-connection. Select **nto1-vcc** cell mode encapsulation or **nto1-vpc** cell mode encapsulation.
chan-num	Optional. Channel number on the port for which a cross-connection is to be specified. The range of values is 0 to 32,767. For Asynchronous Transfer Mode (ATM) OC cards, a default channel number of 1 must be specified.
circuit-id	Optional. Layer 2 (L2) circuit ID. Depending on the type of circuit being cross-connected, the L2 circuit ID takes one of the following constructs: • **vpi-vci** *vpi vci*—ATM permanent virtual circuit (PVC). Specifies the virtual path identifier (VPI) and virtual channel identifier (VCI). For *vpi*, acceptable values are 0 to 255. For *vci*, acceptable values are 1 to 65,535. • **vpi-vci** *vpi start-vci* **through** *end-vci*—Range of ATM PVCs. Specifies the VPI and the range of VCIs. For *vpi*, acceptable values are 0 to 255. For *start-vci* and *end-vci*, acceptable values are 1 to 65,535. • **vlan-id** *pvc-vlan-id*—Virtual LAN (VLAN) 802.1Q PVC that is not within an 802.1Q tunnel. Specifies the PVC VLAN tag. For *pvc-vlan-id*, acceptable values are 1 to 4,095. • **vlan-id** {*start-pvc-vlan-id* **through** *end-pvc-vlan-id* | **any**}—Range of VLAN 802.1Q PVCs that are not within an 802.1Q tunnel. Specifies the range of PVC VLAN tags. For *start-pvc-vlan-id* and *end-pvc-vlan-id* , acceptable values are 1 to 4,095. The keyword **any** is equivalent to **1 through 4095**. • **vlan-id** *tunl-vlan-id*:*pvc-vlan-id*—VLAN 802.1Q PVC that is within an 802.1Q tunnel. Specifies the VLAN tag for the tunnel followed by the PVC VLAN tag. For *tunl-vlan-id* and *pvc-vlan-id*, acceptable values are 1 to 4,095. • **vlan-id** *tunl-vlan-id*:{*start-pvc-vlan-id* **through** *end-pvc-vlan-id* | **any**}—Range of VLAN 802.1Q PVCs that are within an 802.1Q tunnel. Specifies the VLAN tag for the tunnel followed by the range of PVC VLAN tags. For *tunl-vlan-id*, *start-pvc-vlan-id*, and *end-pvc-vlan-id;* acceptable values are 1 to 4,095. The keyword **any** is equivalent to **1 through 4095**. • **dlci** *dlci*—Data-link connection identifier (DLCI) for the Frame Relay PVC. The range of values for the *dlci* argument is 16 to 991. For Ethernet ports with no 802.1Q PVCs, no circuit descriptor is specified.
exp-bits *bits-num*	Optional. EXP bits to be used for transport. The range of the *bits-num* argument values is 0 to 7.

lg ***link-group***	Access link group for which a cross-connection is to be specified.
n-to-1	Optional. Selects *n*-to-1 circuits over one pseudowire, where *n* is equal to the number of circuits to be carried on the pseudowire.
nto1-vcc	Selects *n*-to-1 virtual channel connection (VCC) cell mode encapsulation, where *n* is equal to 1. Required for Cell (AAL2) PW configuration.
nto1-vpc	Selects *n*-to-1 virtual path connection (VPC) cell mode encapsulation, where *n* is equal to 1. Required for Cell (AAL2) PW configuration.
peer ***peer-addr***	IP address of the remote peer provider edge (PE) router.
port	Card port number of the port for which a cross-connection is to be specified.
remote-encap ***type***	Encapsulation type of the remote end of the cross-connection when the ***type*** is different from the near end. The ***type*** argument, which specifies the encapsulation of the remote end, can take any one of the following values: • **1qtunnel**—Specifies the 802.1Q tunnel encapsulation type. • **bridged1483**—Specifies the RFC 1483 bridged encapsulation type. • **dot1q**—Specifies the 802.1Q Ethernet encapsulation type. • **ethernet**—Specifies Ethernet encapsulation.
slot	Chassis slot number with the port for which a cross-connection is to be specified.
sub-chan-num	Optional. Subchannel number on the port for which a cross-connection is to be specified. The range of values is 0 to 255.
transport	Specifies the 802.1Q PVC is transport-enabled.
vc-id ***vc-id***	Virtual circuit (VC) identifier associated with the LDP L2VPN cross-connection. The range of the ***vc-id*** argument value is 0 to 4,294,967,295.
profile ***profile-name***	Alternate option to specifying an L2VPN peer. Associates an L2VPN profile to a cross-connection. When you attach an L2VPN profile to an XC, the LSP that is specified in the profile is automatically mapped to the XC you are configuring.
backup	Enters primary L2VPN XC configuration mode, where you can configure the VC ID and peer address for a backup cross-connection.

1.107.4 Default

None

1.107.5 Usage Guidelines

Use the **xc vc-id** command to create an LDP L2VPN cross-connection.

When creating a cross-connection to a remote circuit that uses an encapsulation type that is different than the encapsulation type of the local circuit, use the **`remote-encap`** keyword to specify the encapsulation type used at the remote end of the cross-connection.

Note:

The SmartEdge router supports the following encapsulation interconnectivity:

- ATM RFC 1483 bridged to Ethernet
- ATM RFC 1483 bridged to dot1q
- ATM RFC 1483 routed to Ethernet
- ATM RFC 1483 routed to dot1q

ATM-to-ATM pseudowire cross-connections support the following encapsulation types:

- AAL5
- *n*-to-1 cell mode

  Note: ATM cell mode encapsulation is not supported on the SmartEdge 100 router media interface cards (MICs).

For ATM OC cards, you must specify a default channel number of 1 in the **`xc vc-id`** command; for example, if the card is an ATM-OC3c/STM-1c, then you must specify a default channel number of 1.

Keep the following information in mind when configuring ATM cell mode pseudowire cross-connections:

- ATM pseudowire cross-connections support *n*-to-1 cell encapsulation on VCCs and VPCs, where *n* is equal to 1. In VPC mode, you must configure an explicit range of VCIs for the ATM cells to be transported over the pseudowire.
- The SmartEdge router supports up to 16,000 PVCs and 16,000 VCIs for each ATM card.

  Note: Cell mode encapsulation is supported only on ATM cards running in the default **`vc-fair`** mode. Cell mode-encapsulated PVCs cannot be provisioned if the ATM card is set to **`atm-priority`** mode or **`ip-priority`** mode. For more information about setting the ATM card mode with the **`atm mode`** command, see *Configuring Cards*.

Use the **`no`** form of this command to delete all LDP L2VPN cross-connections.

1.107.6 Examples

The following example shows how to create an LDP L2VPN cross-connection between an ATM PVC and the remote peer PE router, `101.1.1.1`:

```
[local]Redback(config-ctx)#l2vpn
[local]Redback(config-l2vpn)#xc-group group2
[local]Redback(config-l2vpn-xc-group)#xc 12/1 vpi-vci 200 1256
vc-id 2 peer 101.1.1.1
[local]Redback(config-l2vpn-xc-group)#
```

The following example shows how to create an LDP L2VPN cross-connection between an 802.1Q PVC and the remote peer PE router, `101.1.1.1`:

```
[local]Redback(config-ctx)#l2vpn
[local]Redback(config-l2vpn)#xc-group group2
[local]Redback(config-l2vpn-xc-group)#xc 12/1 vlan-id 200 vc-id
2 peer 101.1.1.1
[local]Redback(config-l2vpn-xc-group)#
```

The following example shows how to create an LDP L2VPN cross-connection between an Frame Relay PVC and the remote peer PE router, `101.1.1.2`:

```
[local]Redback(config-ctx)#l2vpn
[local]Redback(config-l2vpn)#xc-group group2
[local]Redback(config-l2vpn-xc-group)#xc 12/1 dlci 101 vc-id 2 peer 101.1.1.2
[local]Redback(config-l2vpn-xc-group)#
```

The following example shows how to create an LDP L2VPN cross-connection between an Ethernet port and the remote peer PE router, `101.1.1.3`:

```
[local]Redback(config-ctx)#l2vpn
[local]Redback(config-l2vpn)#xc-group group2
[local]Redback(config-l2vpn-xc-group)#xc 12/1 vc-id 2 peer 101.1.1.3
[local]Redback(config-l2vpn-xc-group)#
```

The following example shows how to create an LDP L2VPN cross-connection between an Ethernet port and a remote circuit that uses 802.1Q PVC encapsulation:

```
[local]Redback(config-ctx)#l2vpn
[local]Redback(config-l2vpn)#xc-group group2
[local]Redback(config-l2vpn-xc-group)#xc 12/1 vc-id 2 peer 101.1.1.3 remote-encap dot1q
[local]Redback(config-l2vpn-xc-group)#
```

The following example shows how to create an ATM cell mode LDP L2VPN pseudowire cross-connection that hosts a single PVC. In this example, an ATM cross connection with a VC ID of 2000 is created on port 1, VPI 0, and VCI 32 on the card in slot 4. The IP address of the remote peer PE is 111.111.111.111:

```
[local]Redback(config-ctx)#l2vpn
[local]Redback(config-l2vpn)#xc-group group2
[local]Redback(config-l2vpn-xc-group)#xc 4/1 vpi-vci 0 32 vc-id 2000 peer 111.111.111.111 cell-encap nto1-vcc
```

The following example shows how to create an ATM cell mode LDP L2VPN pseudowire cross-connection that hosts multiple PVCs:

```
[local]Redback(config-ctx)#l2vpn
[local]Redback(config-l2vpn)#xc-group group2
[local]Redback(config-l2vpn-xc-group)#xc 4/1 vpi-vci 0 32 through 8000 n-to-1 vc-id 1 peer 1.1.1.1 cell-encap
```

The following example illustrates how to configure an Ethernet transport-enabled pseudowire cross-connection. The **link-group**, **port**, and **context** commands are entered in the global configuration mode:

```
link-group group2 access economical
 encapsulation dot1q
 dot1q pvc transport 100 through 200
  l2vpn local
port ethernet 2/10
 no shutdown
 encapsulation dot1q
 link-group group2
context local
!
 l2vpn
  xc-group default
!  LDP circuit bindings
  xc lg group2 vlan-id 100 through 200 transport vc-id 100 peer 180.180.180.180
```

The following example shows how to attach the L2VPN profile `l2vpn-prof1` to the cross-connection on VC `102` and port `1` of the card in slot `3`:

```
[local]Redback(config-xc-group)#xc 3/1 vc-id 102 profile l2vpn-prof1
```

1.108 xc (L2VPN vpn-label)

For an Ethernet pseudowire cross-connection:

xc *slot/port*[**:***chan-num*[**:***sub-chan-num*]] [*circuit-id*] **vpn-label** *label* **peer** *peer-addr* [**remote encap** *type*] [**exp-bits** *bits-num*]

no xc *slot/port*[**:***chan-num*[**:***sub-chan-num*]] [*circuit-id*] **vpn-label** *label* **peer** *peer-addr* [**remote encap** *type*] [**exp-bits** *bits-num*]

For an Ethernet transport-enabled pseudowire cross-connection:

xc *slot/port*[**:***chan-num*[**:***sub-chan-num*]] *circuit-id* **transport vpn-label** *label* **peer** *peer-addr* [**remote encap** *type*] [**exp-bits** *bits-num*]

no xc *slot/port*[**:***chan-num*][**:***sub-chan-num*]] *circuit-id* **transport vpn-label** *label* **peer** *peer-addr* [**remote encap** *type*] [**exp-bits** *bits-num*]

For a link group pseudowire cross-connection:

xc lg *link-group circuit-id* **vpn-label** *label* **peer** *peer-addr* [**remote encap** *type*] [**exp-bits** *bits-num*]

no xc lg *link-group circuit-id* **vpn-label** *label* **peer** *peer-addr* [**remote encap** *type*] [**exp-bits** *bits-num*]

For a link group transport pseudowire cross-connection:

xc xc lg ***link-group circuit-id*** **transport vpn-label** ***label*** **peer** ***peer-addr*** [**remote encap** ***type***] [**exp-bits** ***bits-num***]

no xc lg ***link-group circuit-id*** **transport vpn-label** ***label*** **peer** ***peer-addr*** [**remote encap** ***type***] [**exp-bits** ***bits-num***]

1.108.1 Purpose

Creates a static Layer 2 Virtual Private Network (L2VPN) cross-connection.

1.108.2 Command Mode

L2VPN XC group configuration

1.108.3 Syntax Description

slot	Chassis slot number with the port for which a cross-connection is to be specified.
port	Card port number of the port for which a cross-connection is to be specified.
chan-num	Optional. Channel number on the port for which a cross-connection is to be specified. The range of values is 0 to 32,767. For Asynchronous Transfer Mode (ATM) OC cards, a default channel number of 1 must be specified.
sub-chan-num	Optional. Subchannel number on the port for which a cross-connection is to be specified. The range of values is 0 to 255.

circuit-id	Optional. Layer 2 (L2) circuit ID. Depending on the type of circuit being cross-connected, the L2 circuit ID takes one of the following constructs: • **vpi-vci** *vpi vci*—ATM permanent virtual circuit (PVC). Specifies the virtual path identifier (VPI) and virtual channel identifier (VCI). For *vpi*, acceptable values are 0 to 255. For *vci*, acceptable values are 1 to 65,535. • **vpi-vci** *vpi start-vci* **through** *end-vci*—Range of ATM PVCs. Specifies the VPI and the range of VCIs. For *vpi*, acceptable values are 0 to 255. For *start-vci* and *end-vci*, acceptable values are 1 to 65,535. • **vlan-id** *pvc-vlan-id*—Virtual LAN (VLAN) 802.1Q PVC that is not within an 802.1Q tunnel. Specifies the PVC VLAN tag. For *pvc-vlan-id*, acceptable values are 1 to 4,095. • **vlan-id** {*start-pvc-vlan-id* **through** *end-pvc-vlan-id* | **any**}—Range of VLAN 802.1Q PVCs that are not within an 802.1Q tunnel. Specifies the range of PVC VLAN tags. For *start-pvc-vlan-id* and *end-pvc-vlan-id*, acceptable values are 1 to 4,095. The keyword **any** is equivalent to **1 through 4095**. • **vlan-id** *tunl-vlan-id***:***pvc-vlan-id*—VLAN 802.1Q PVC that is within an 802.1Q tunnel. Specifies the VLAN tag for the tunnel followed by the PVC VLAN tag. For *tunl-vlan-id* and *pvc-vlan-id*, acceptable values are 1 to 4,095. • **vlan-id** *tunl-vlan-id*:{*start-pvc-vlan-id* **through** *end-pvc-vlan-id* | **any**}—Range of VLAN 802.1Q PVCs that are within an 802.1Q tunnel. Specifies the VLAN tag for the tunnel followed by the range of PVC VLAN tags. For *tunl-vlan-id*, *start-pvc-vlan-id*, and *end-pvc-vlan-id;* acceptable values are 1 to 4,095. The keyword **any** is equivalent to **1 through 4095**. • **dlci** *dlci*—Data-link connection identifier (DLCI) for the Frame Relay PVC. The range of values for the *dlci* argument is 16 to 991. For Ethernet ports with no 802.1Q PVCs, no circuit descriptor is specified.
transport	Specifies the 802.1Q PVC is transport-enabled.
label	Inner label associated with the static L2VPN cross-connection. The range of the *label* argument values is 4096 to 65535.
peer *peer-addr*	IP address of the remote peer provider edge (PE) router.

remote-encap *type*	Specifies that a different encapsulation type is used at the remote end of the cross-connection. The *type* argument specifies one of the following encapsulation types: • **1qtunnel**—Specifies the 802.1Q tunnel encapsulation type. • **bridged1483**—Specifies the RFC 1483 bridged encapsulation type. • **dot1q**—Specifies the 802.1Q Ethernet encapsulation type. (1)
exp-bits *bits-num*	Optional. EXP bits to be used for transport. The range of the *bits-num* argument values is 0 to 7.

*(1) The **remote-encap type** construct is required only for connections that use ATM and Ethernet interworking. If both ends of the connect use the same encapsulation type, you do not need to specify the **remote-encaptype** construct.*

1.108.4 Default

None

1.108.5 Usage Guidelines

Use the **xc vpn-label** command to create a static L2VPN cross-connection.

For ATM OC cards, you must specify default channel number of 1 in the **xc vpn-label** command; for example, if the card is an ATM-OC3c/STM-1c, then you must specify a default channel number of 1.

Use the **no** form of this command to delete all static L2VPN cross-connections.

1.108.6 Examples

The following example creates a static L2VPN cross-connection between an ATM PVC and the remote peer PE router, `192.168.1.1:`

```
[local]Redback(config-ctx)#l2vpn
[local]Redback(config-l2vpn)#xc-group group2
[local]Redback(config-l2vpn-xc-group)#xc 12/1 vpi-vci 10 12 vpn-label
 5000 peer 101.1.1.1
[local]Redback(config-l2vpn-xc-group)#
```

The following example creates a static L2VPN cross-connection between an 802.1Q PVC and the remote peer PE router, `192.168.1.1:`

```
[local]Redback(config-ctx)#l2vpn
[local]Redback(config-l2vpn)#xc-group group2
[local]Redback(config-l2vpn-xc-group)#xc 12/1 vlan-id 200 vpn-label 5000
 peer 101.1.1.1
[local]Redback(config-l2vpn-xc-group)#
```

The following example creates a static L2VPN cross-connection between an Frame Relay PVC and the remote peer PE router, `101.1.1.2`:

```
[local]Redback(config-ctx)#l2vpn
[local]Redback(config-l2vpn)#xc-group group2
[local]Redback(config-l2vpn-xc-group)#xc 12/1 dlci 101 vpn-label 5000
peer 101.1.1.2
[local]Redback(config-l2vpn-xc-group)#
```

The following example creates a static L2VPN cross-connection between an Ethernet port and the remote peer PE router, `101.1.1.3`:

```
[local]Redback(config-ctx)#l2vpn
[local]Redback(config-l2vpn)#xc-group group2
[local]Redback(config-l2vpn-xc-group)#xc 12/1 vpn-label 5000 peer
101.1.1.3
[local]Redback(config-l2vpn-xc-group)#
```

1.109 yellow-alarm

yellow-alarm {**detection** | **generation**}

{**no** | **default**} **yellow-alarm** {**detection** | **generation**}

1.109.1 Purpose

Enables the detection or generation of yellow alarms on the DS-1 channel.

1.109.2 Command Mode

DS-1 configuration

1.109.3 Syntax Description

detection	Enables yellow-alarm detection.
generation	Enables yellow-alarm generation.

1.109.4 Default

Detection and generation of yellow alarms are enabled.

1.109.5 Usage Guidelines

Use the **yellow-alarm** command to enable the detection or generation of yellow alarms on a DS-1 channel.

Use the **no** form of this command to disable the specified yellow alarm function.

Use the **default** form of this command to enable the specified yellow alarm function with its default values.

1.109.6 Examples

The following example shows how to disable yellow alarm detection on a DS-1 channel:

```
[local]Redback(config)#port ds1 4/1:1:1
[local]Redback(config-ds1)#no yellow-alarm detection
```