ERICSSON

Commands: ip through li

COMMAND DESCRIPTION

10/190 82-CRA 119 1170/1-V1 Uen K

Contents

1 Command Descriptions

Commands starting with “ip” through commands starting with “li” 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 ip access-group (circuits)

`ip access-group` ***`"acl-name1 acl-name2 acl-name3..."`*** **`context`** ***`context-name`*** {**`in`** | **`out`**} [**`count`**]

`no ip access-group` ***`"acl-name1 acl-name2 acl-name3..."`*** **`context`** ***`context-name`*** {**`in`** | **`out`**}

1.1.1 Purpose

Applies from one to ten IP access control lists (ACL) to packets associated with the current circuit or port.

1.1.2 Command Mode

- dot1q PVC configuration
- port configuration
- dot1q child circuit configuration

1.1.3 Syntax Description

acl-name	Name of the IP ACL to apply to the circuit or port, which can be up to 39 alphanumeric characters long. You can specify up to 10 ACL names in the command. Enclose multiple ACL names within quotation marks and separate each ACL name with one or more spaces as shown in the syntax. The total number of characters for the ACL names must not exceed 255 (average of 24 characters per name). The colon character (:) is not allowed in ACL names.
context ***context-name***	Specifies the name of context of the circuit.
in	Specifies that the ACL is to be applied to incoming packets.
out	Specifies that the ACL is to be applied to outgoing packets.
count	Optional. Enables ACL packet counting.

1.1.4 Default

No ACL is applied.

1.1.5 Usage Guidelines

Use the **ip access-group** command to apply an IP ACL to packets associated with the current circuit or port, filtering the flow of traffic. If you configure multiple ACLs to an IP access group, the SmartEdge router combines the ACLs in order of appearance within the IP access group to produce a specific filtering behavior. The SmartEdge router appends an implicit ***deny ip any any*** rule after all configured rules complete.

The SmartEdge router ignores conditional ACLs referenced in an access group.

Note: Applying an ACL has no effect if the named ACL has not yet been defined. All packets are permitted as if no restrictions were in place.

If an access group for an interface has multiple ACLs, some of the ACLs can be unconfigured; however any unconfigured ACLs have no (zero) rules. Only the configured ACLs in the access group apply to traffic.

When you use the **count** keyword, the system keeps track of the number of matches that occur. By default, counting of packets is disabled.

Caution!

Risk of performance loss. Enabling the count and log functions can affect system performance. To reduce the risk, exercise caution when enabling these features on a production system.

To disable packet counting, enter the **ip access-group** command again, omitting the **count** keyword.

Note:

The following restrictions and limitations affect the application of the **ip access-group** command to layer 2 circuits:

- No support for packet logging when IP ACL filters are applied to layer 2 circuits.
- No support for dynamic ACLs.
- No support for 802.1Q PVCs with raw encapsulation.
- No support for inheritance of IP ACL filters; that is, IP ACL applied to the outer VLAN filters all traffic on the VLAN except the traffic that is going to an inner VLAN.

The following list shows the Layer 2 configurations to which you can apply IP ACL filters:

- Ethernet port
- Cross-connected individual VLAN-based circuit or range VLAN-based circuits
- Cross-connected VLAN-based circuit or range of VLAN-based circuits in an 801.Q tunnel
- Cross-connected PPPoE child circuit
- Cross-connected VLAN-based aggregated circuit in an access link group
- L2VPN port
- VLAN-based circuit or range of circuits attached to an L2VPN
- VLAN-based circuit or range of circuits in an 801.Q tunnel attached to an L2VPN
- VLAN-based aggregated circuit in an access link group attached to an L2VPN

- Port bound to an VPLS bridge
- VLAN-based aggregated circuit in an access link-group attached to a VPLS enabled bridge
- VLAN-based circuit or range of circuits in an 801.Q tunnel bound to a VPLS-enabled bridge
- VLAN-based circuit or range of circuits to a VPLS-enabled bridge

The **ip access-group (circuits)** command does not support the following Layer 2 circuits:

- Raw encapsulation VLAN-based circuits
- Transport-enabled circuits
- ATM circuits

Use the **no** form of this command to remove an applied IP ACL from association with the interface. Enter empty quotations marks (“ ”) to remove all associated ACL names. If you want to delete one or more (but not all) ACLs, enter their names in quotation marks.

1.1.6 Examples

IP ACL filters can be applied to cross-connected Ethernet VLANs using the **ip access-group** command. The following example shows a cross-connected VLAN (in the **CTX_XC** context) to which the **ACL_XC** ACL has been applied:

```
[local]Redback(config)#port ethernet 2/15
[local]Redback(config-port)#encapsulation dot1q
[local]Redback(config-port)#dot1q pvc 10
[local]Redback(config-dot1q-pvc)#ip access-group ACL_XC context CTX_XC in count
[local]Redback(config-dot1q-pvc)#bind bypass
!
[local]Redback(config)#port ethernet 2/16
[local]Redback(config-port)#encapsulation dot1q
[local]Redback(config-port)#dot1q pvc 100
[local]Redback(config-dot1q-pvc)#bind bypass
!
[local]Redback(config)#xc-group default
[local]Redback(config-xc-group)xc 2/15 vlan-id 10 to 2/16 vlan-id 100
```

To see the completed configuration of the access group, enter the **show access-group detail** command in **CTX_XC** context mode. The **show access-group detail** command displays the **Circuit [L2]** flag in brackets to indicate the corresponding interface is a layer 2 port or circuit:

```
[local]Redback(config-ctx)#show access-group detail
IP Filter ACL  : ACL_XC
ACL context    : CTX_XC
Circuit [L2]   : 2/15 vlan-id 10
Direction      : In          ACL status  : No access-list
Count          : Rules       Log         : No
Number of rules: 2
```

This example shows an IP ACL filter applied to a VLAN configured in an access link group bound to a VPLS-enabled bridge:

```
[local]Redback(config)#vpls profile toSE2
[local]Redback(config-vpls-profile)#neighbor 2.2.2.2
!
[local]Redback(config)#context local
[local]Redback(config-ctx)#interface bridge1 bridge
[local]Redback(config-if)#bridge name vplsSE1
!
[local]Redback(config)#context local
[local]Redback(config-ctx)#bridge vplsSE1
[local]Redback(config-bridge)#vpls
[local]Redback(config-vpls)#profile toSE2 pw-id 10
!
[local]Redback(config)#link-group LAG access
[local]Redback(config-link-group)#encapsulation dot1q
[local]Redback(config-link-group)#dot1q pvc 10
[local]Redback(config-dot1q-pvc)#ip access-group ACL_1 context local in count
[local]Redback(config-dot1q-pvc)#bind interface bridge1 local
!
[local]Redback(config)#port ethernet 4/5
[local]Redback(config-port)#link-group LAG
```

1.2 ip access-group (interfaces and subs)

[no]ip access-group ***"acl-name1 acl-name2 acl-name3..."*** **{in | out} [count] [log]**

1.2.1 Command Mode

- Interface configuration
- Subscriber configuration

1.2.2 Syntax Description

acl-name	Name of the IP ACL to apply to the interface, which can be up to 39 alphanumeric characters long. You can specify up to 10 ACL names in the command. Enclose multiple ACL names within quotation marks and separate each ACL name with one or more spaces as shown in the syntax. The total number of characters for the ACL names must not exceed 255 for interface mode and 253 for subscriber mode (average of 24 characters per name). The colon character (:) is not allowed in ACL names.
in	Specifies that the ACL is to be applied to incoming packets.
out	Specifies that the ACL is to be applied to outgoing packets.
count	Optional. Enables ACL packet counting. Not available in subscriber configuration mode.
log	Optional. Enables ACL packet logging. Not available in subscriber configuration mode.

1.2.3 Default

No ACL is applied.

1.2.4 Usage Guidelines

Use the **ip access-group** command to apply one to ten IP access control lists (ACLs) to packets associated with an interface or subscriber, restricting the flow of traffic through the SmartEdge router. If you configure multiple ACLs to an IP access group, the SmartEdge router combines the ACLs in order of appearance within the IP access group to produce a specific filtering behavior. If you configure a dynamic filter ACL for a subscriber, the SmartEdge router applies the rules of the combined ACL and then the dynamic filter ACL. The SmartEdge router appends an implicit ***deny ip any any*** rule after all configured rules complete.

The SmartEdge router ignores conditional ACLs referenced in an access group.

Note: Applying an ACL to an interface has no effect if the named ACL has not yet been defined. All packets are permitted as if no restrictions were in place.

If an access group for an interface has multiple ACLs, some of the ACLs can be unconfigured; however any unconfigured ACLs have no (zero) rules. Only the configured ACLs in the access group apply to traffic.

When you use the **count** keyword, the system keeps track of the number of matches that occur. When you use the **log** keyword, the system keeps track of the number of packets that were denied. By default, counting and logging of packets is disabled.

Caution!

Risk of performance loss. Enabling the count and log functions can affect system performance. To reduce the risk, exercise caution when enabling these features on a production system.

To disable packet counting or logging, enter the **ip access-group** command again, omitting the **count** or **log** keyword.

Use the **no** form of this command to remove an applied IP ACL from association with the interface. Enter empty quotations marks (“ ”) to remove all associated ACL names. If you want to delete one or more (but not all) ACLs, enter their names in quotation marks.

1.2.5 Examples

The following example shows how to apply the IP ACLs, **WebCacheACL and SmartFilter**, to the interface, **topgun**, and enable both packet counting and logging:

```
[local]Redback(config)#context fighter
[local]Redback(config-ctx)#interface topgun
[local]Redback(config-if)#ip access-group "WebCacheACL SmartFilter" in log count
```

The following example shows how to apply the ACLs, **WebCacheACL and SmartFilter**, to the subscriber, **joe:**

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#subscriber name joe
[local]Redback(config-sub)#ip access-group "WebCacheACL SmartFilter" out
```

1.3 ip access-group (policy)

ip access-group [*acl-name context-name*]

no ip access-group [*acl-name context-name*]

1.3.1 Purpose

Applies an IPv4 access control list (ACL) to packets associated with the current forward, metering, or policing policy; that is, the specified ACL defines class matching criteria for the current policy.

1.3.2 Command Mode

- Forward policy configuration
- Policing policy configuration
- Metering policy configuration

1.3.3 Syntax Description

acl-name	Specifies the name of a IPv4 policy ACL created using the `policy access-list` command (in context configuration mode). This parameter and the accompanying *`context-name`* parameter are optional only if the current policy is configured with the `radius-guided` option/keyword.
context-name	Specifies the name of the context in which the policy ACL was defined.

1.3.4 Default

None

1.3.5 Usage Guidelines

Use the **`ip access-group`** (policy) command to apply a policy ACL to the current class-based policy (forward, QoS metering, or policing policy) and enter policy-group configuration mode.

If the class-based policy was defined as RADIUS-guided, the policy ACL that it references can be dynamic, static, or both:

- A dynamic policy ACL is one that the SmartEdge router applies to the class-based policy for a particular subscriber session using the rules specified in an instance of vendor-specific attribute (VSA) 164 that the RADIUS server supplies in an access-response or COA message for the subscriber. For inbound forward policies only, a dynamic ACL may be alternatively specified using VSA 196 attribute `fwd-in-access-group`. In these cases, the **`ip access-group`** command may be used without specifying the name of a statically configured policy ACL.

- A static policy ACL is a locally configured policy access-list whose name must be explicitly specified. If you include the ***acl-name*** argument, you must also include the ***context-name*** argument when you apply a static policy ACL to a forward policy or QoS policy.

  You can apply a dynamic policy ACL in addition to a static policy ACL. If VSA 164 was used to apply the dynamic ACL, the static policy ACL takes precedence over the dynamic policy ACL; a locally configured access-list's rules are evaluated before those from a dynamic ACL specified via VSA 164. If VSA 196 `fwd-in-access-group`. is used to apply a dynamic ACL to a subscriber, the dynamic policy ACL supersedes and replaces the locally configured access-list, if applicable.

  If the class-based policy is not defined as RADIUS-guided, the policy ACL that it references must be static, and the **`ip access-group`** command must specify the locally configured access-list's name and context.

Note: If a forward policy, Network Address Translation (NAT) policy, or quality of service (QoS) policy references a policy ACL that does not exist, the reference is ignored.

Warning!

The system does not warn you if you enter the **`ip access-group`** command with the name of a static policy ACL that does not exist in the specified context; that is, the **`policy access-list`** command has not configured a static policy ACL with the matching name in the specified context. If the ACL does not exist when the class-based policy is referenced by a static circuit (namely, a port, channel, PVC, or VLAN), the ACL reference is ignored and applicable traffic is treated as if it matched none of the classes specified in the policy. If the ACL does not exist when the class-based policy is referenced by a subscriber session circuit (namely, PPPoA, PPPoE, or CLIPs), the subscriber's applicable policy reference will be rejected, which may in turn cause the subscriber session to be terminated (unless the **`session-action failure always-up`** command is configured for the subscriber).

1.3.6 Examples

The following example shows how to create an IPv4 policy ACL, `class_ipv4` in the `local` context and attaches it to the QoS policing policy, `POL1`. The ACL defines two classes, `B` and `C`, that are referenced by `POL1`:

Note: The command line `qos policy POL1 policing` creates a policy that is not RADIUS guided. When the current policy is not RADIUS guided, the ***acl-name context-name*** arguments are not optional.

```
[local]Redback#config
Enter configuration commands, one per line, 'end' to exit
[local]Redback#context local
[local]Redback(config-ctx)#policy access-list class-ipv4
!
!
[local]Redback(config-access-list)#seq 10 permit ip any 15.1.0.0 0.0.255.255 class B
[local]Redback(config-access-list)#seq 20 permit ip any 15.2.0.0 0.0.255.255 class C
[local]Redback(config-access-list)#exit
[local]Redback(config)#qos policy POL1 policing                          //See NOTE
[local]Redback(config-policy-policing)#ip access-group class_ipv4 local
[local]Redback(config-policy-group)#class B
[local]Redback(config-policy-group-class)#rate 50 burst 20000 counters
[local]Redback(config-policy-class-rate)#conform mark dscp af11
[local]Redback(config-policy-class-rate)#exit
[local]Redback(config-policy-group-class)#exit
[local]Redback(config-policy-group)#class C
[local]Redback(config-policy-group-class)#rate 200 burst 90000 counters
[local]Redback(config-policy-class-rate)#conform mark dscp ef
[local]Redback(config-policy-class-rate)#commit
```

1.4 ip access-list

ip access-list *acl-name* [**ssh-and-telnet-acl**]

no ip access-list *acl-name* [**ssh-and-telnet-acl**]

1.4.1 Purpose

Configures an IP access control list (ACL) and enters access control list configuration mode.

1.4.2 Command Mode

Context configuration

1.4.3 Syntax Description

acl-name	Name of the ACL. Must be unique within the context.
ssh-and-telnet-acl	Optional. Specifies that the ACL applies to Telnet and Secure Shell (SSH) traffic. The Telnet or SSH ACL applies only to the remote host address (or IP source address). The IP destination address is ignored.

1.4.4 Default

None

1.4.5 Usage Guidelines

Use the **ip access-list** command to configure an IP ACL and enter access control list configuration mode, where you can define statements using the **permit** and **deny** commands; the SmartEdge OS automatically sets the order of the statements. If you want to configure the order of the statements, use the **seq permit** and **seq deny** commands.

Note: In every ACL, there is an automatic **deny any any** statement at the end of the list. This automatic statement blocks all traffic not explicitly allowed, although it does not appear in the output of the *show configuration acl* command. It could block valid access to a context; for example, in the local context, it could block administrator access to the Ethernet management port. To allow administrator access, add a statement to explicitly allow access from authorized sources; for example, you could add a **permit ip any any** or **permit ip** ***src src-wildcard dest dest-wildcard*** statement to the ACL.

Use the **resequence ip access-list** command to reorder a filtering ACL

When the IP ACL is created and its conditions have been set, you can apply the list to any of these entities:

- An interface to restrict the flow of traffic through the SmartEdge router with the **ip access-group** command (in interface configuration mode).
- Local inbound traffic coming into the SmartEdge operating system kernel with the **admin-access-group** command (in context configuration mode).
- A subscriber record, named profile, or default profile with the **ip access-group** command (in subscriber configuration mode).
- Inbound SSH and Telnet traffic with the **service** command (in context configuration mode).
- An interface enabled with reverse path forwarding (RPF) to allow packets that fail the RPF check but match the ACL to pass through with the **ip verify unicast source** command (in interface configuration mode).

A reference to an IP ACL that does not exist or does not contain any configured entries implicitly matches and permits all packets.

Use the **no** form of this command to remove an ACL from the configuration.

1.4.6 Examples

The following example shows how to create an IP ACL, **WebCacheACL:**

```
[local]Redback(config-ctx)#ip access-list WebCacheACL
[local]Redback(config-access-list)#
```

1.5 ip address (interface)

ip address *ip-addr* {**application** | {*netmask* | */prefix-length*} **[secondary] [tag** *tag*]}

{**no** | **default**} **ip address** *ip-addr* {**application** | {*netmask* | */prefix-length*} **[secondary] [tag** *tag*]}

1.5.1 Purpose

Assigns a primary IP address, and optionally, one or more secondary IP addresses, to an interface.

1.5.2 Command Mode

Interface configuration

1.5.3 Syntax Description

ip-addr	Primary or secondary IP address of the interface.
application	Specifies a second IP address for to the current interface that can be used for applications such as Inter-Chassis Redundancy (ICR) operation. This option allows ICR active and backup SmartEdge systems to use different IP addresses as the giaddr from identically addressed multibind interfaces. The different giaddr IP addresses are a requirement that makes it possible for the identically addressed multibind interfaces of multiple SmartEdge systems in the ICR operation configuration to communicate with a DHCP server. The application address must be one of the addresses in the subnet configured for the multibind interface in order to be used for giaddr. Any ARP requests received for application addresses are dropped. For additional information about the **giaddr** command, see the *Command List*. Up to 16 application addresses per interface may be configured.
netmask	Network mask for the associated IP network.

prefix-length	Prefix length for the associated IP address. The range of values is 0 to 32
secondary	Optional. Configures the address as a secondary IP address on the interface.
tag ***tag***	Optional. Route tag for the IP address. An unsigned 32-bit integer, the range of values is 1 to 4,294,967,295; the default value is 0.

1.5.4 Default

No IP address is assigned to an interface.

1.5.5 Usage Guidelines

Use the **ip address** command to assign a primary IP address, and optionally, one or more secondary IP addresses, to an interface. This assignment enables IP services on an interface.

Use the ***ip-addr*** argument and either the ***netmask*** or ***/prefix-length*** construct to assign the interface a primary IP address and netmask or prefix length. For nonloopback interfaces, use the **bind interface** command (in port configuration mode) to bind a circuit to the interface on which IP services are enabled.

Note: The Address Resolution Protocol (ARP) is enabled by default on broadcast-capable interfaces.

Use the optional **secondary** keyword to designate an IP address as a secondary IP address for the interface. You can configure up to 15 secondary addresses for each primary interface. Interface costs configured for routing protocols apply to secondary IP addresses in the same manner that they apply to primary IP addresses. Secondary IP addresses are treated as locally attached networks.

If Routing Information Protocol (RIP) split horizon is enabled on an interface that is configured with multiple IP addresses, a single update sourced by the primary IP address is sent advertising only the major networks. If split horizon is disabled, multiple updates sourced from each address on the interface are sent and all subnets are advertised.

Use the optional **tag** ***tag*** construct to assign a route tag to the IP address. If you do not include this construct, the value 0 is assigned as the route tag.

Assigning a route tag allows you to propagate the connected route for the interface to other protocols such as Border Gateway Protocol (BGP) and Open Shortest Path First (OSPF), using a route map with a match condition that specifies the route tag value. For more information about route tags and the routing policy commands to manage them, see *Configuring Routing Policies*.

When configuring an OSPF interface, use the **ip address** command first to establish the interface, and then enable OSPF on it by using the **interface** command in OSPF area configuration mode; see *Configuring OSPF*. The primary IP address of the interface must belong to the area in which OSPF is enabled. In addition, only neighbors on the primary address subnet can be OSPF peers.

Caution!

Risk of IP service loss. Removing the primary IP address disables all IP services for that address on the specified interface. Disabling IP services deletes a corresponding OSPF interface from the running configuration. To reduce the risk, do not remove a primary IP address for an OSPF interface, unless you have configured a secondary IP address for the OSPF interface, or intend to delete it.

Use the **bind interface** command (in link configuration mode) to statically bind a port, channel, permanent virtual circuits (PVCs), 802.1Q tunnel, link group, Generic Routing Encapsulation (GRE) tunnel circuit, or overlay tunnel circuit to a previously created interface in the specified context. No data can flow through a port, channel, PVC, 802.1Q tunnel, child circuit, link group, or tunnel circuit until it is bound to an interface. Both the interface and the specified context must exist before you enter the **bind interface** command. If either is missing, an error message displays. For more information on **bind interface** command, see the *Command List*.

Note: When adding interfaces to a context follow these limitations:

- Do not configure more than one interface IP address in a context on the same subnet.
- The host portion of an interface IP address cannot be 0 or the subnet for a broadcast IP address.

Use the **no** or **default** form of this command to remove an IP address from an interface. You must remove all secondary IP addresses before you can remove the primary IP address.

1.5.6 Examples

The following example shows how to assign an IP address and netmask to the **enet1** interface:

```
[local]Redback(config-ctx)#interface enet1
[local]Redback(config-if)#ip address 10.4.5.2/24
```

The following example shows how to configure two noncontiguous Classless InterDomain Routing (CIDR) blocks for the **downstream** interface:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#interface downstream
[local]Redback(config-if)#ip address 10.0.0.1/24
[local]Redback(config-if)#ip address 11.0.0.1/24 secondary
```

The following example shows how to bind port **3/1** to the **downstream** interface using either IP address:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#interface downstream
[local]Redback(config-if)#ip address 10.0.0.2/28
[local]Redback(config-if)#ip address 11.0.0.2/28 secondary
[local]Redback(config-if)#exit
[local]Redback(config-ctx)#exit
[local]Redback(config)#port ether 3/1
[local]Redback(config-port)#bind interface downstream local
```

1.6 ip address (subscriber)

ip address {*ip-addr* [*netmask* | */prefix-length*] | **pool** [**name** *name*]}

no ip address {*ip-addr* [*netmask* | */prefix-length*] | **pool**}

1.6.1 Purpose

Assigns an IP address to the subscriber record or profile.

1.6.2 Command Mode

Subscriber configuration

1.6.3 Syntax Description

ip-addr	IP address for the subscriber record or profile.
netmask	Optional. Network mask for the IP address. You must enter a mask of at least 24 bits; that is, a mask in the range of 255.255.255.0 to 255.255.255.255.
prefix-length	Optional. Prefix length. The range of values is 0 to 32.
pool	Indicates that the subscriber will be assigned an IP address from a locally managed IP pool. Required if configuring a default subscriber profile.
name ***name***	Optional. Name of an IP pool or an interface with a named or unnamed IP pool.

1.6.4 Default

None

1.6.5 Usage Guidelines

Use the **ip address** command to assign an IP address to the subscriber record or profile. To specify a range of contiguous IP addresses, use the optional ***netmask*** argument. For Point-to-Point Protocol (PPP)-encapsulated circuits, only the first available IP address in a subscriber record is used for address negotiation. For subscriber circuits using RFC 1483 bridged encapsulation, entries are added to the host table for any such IP addresses.

You can specify either an IP address or an IP pool, but not both. You must use the **pool** keyword to configure a default subscriber profile. The **name** ***name*** construct is either the name of a named IP pool (created with the ***pool-name*** argument) or the name of an interface (created with the ***if-name*** argument).

When binding a subscriber circuit that has been configured with the **bind authentication** command (in subscriber configuration mode), and the local or Remote Authentication Dial-In User Service (RADIUS) subscriber record specifies an IP pool or interface name, the SmartEdge router first checks for an available IP address in the IP pool specified in the record. If the pool does not exist, it then looks for an interface with that name. If there are no unnamed IP pools associated with the interface, the binding for the subscriber circuit fails. For more information on the **bind authentication** command (in subscriber configuration mode), see the *Command List*.

If this subscriber will be a user of clientless IP service selection (CLIPS), or if this named or default subscriber profile is intended for such subscribers, follow these guidelines:

- For static CLIPS circuits, a subscriber record or its assigned profile must have one and only one IP address. If you enter this command more than once for a subscriber record or profile, only the last IP address is applied to the static CLIPS circuit.
- For dynamic CLIPS circuits, do not use this command to assign an IP address; instead, use the **dhcp max-addrs** command (in subscriber configuration mode) and specify **1** as the value for the ***max-num*** argument. For more information about the **dhcp max-addr** command, see the *Command List*.

  Note: To create a pool of IP addresses for an interface, use the **ip pool** command (in interface configuration mode); to assign an IP address to an interface, use the **ip address** command (in interface configuration mode).

Any IP address assigned to a subscriber must fall within the address and netmask range configured for an interface in the context to which the subscriber is to be bound; otherwise, the binding fails. The same is true of IP addresses that are returned by RADIUS servers and that are to be assigned to subscribers.

Note: If you are authenticating a subscriber using the RADIUS, the subscriber record is ignored.

To assign an IP pool address to the subscriber using RADIUS, configure the RADIUS server to return either 255.255.255.254 or 0.0.0.0 as the value for attribute 8, Framed-IP-Address. These values allow the subscriber to be assigned any available IP address from any pool configured within the context.

If you specify a named IP pool, configure the RADIUS server to return the name of the pool in the vendor-specific attribute (VSA) 36 provided by Ericsson AB, IP-Address-Pool-Name.

Use the **no** form of this command to remove an IP address from a subscriber record.

1.6.6 Examples

The following example shows how to define the IP address, **10.1.1.7**, for a subscriber, **host1:**

```
[local]Redback(config-ctx)#subscriber name host1

[local]Redback(config-sub)#ip address 10.1.1.7
```

The next example shows how to define two IP addresses, **10.1.1.14** and **10.1.1.15**, for a subscriber, **host2:**

```
[local]Redback(config-ctx)#subscriber name host2
[local]Redback(config-sub)#ip address 10.1.1.14
[local]Redback(config-sub)#ip address 10.1.1.15
```

The following example shows how to define eight IP addresses, **10.1.1.32** to **10.1.1.39**, for a subscriber, **host8:**

```
[local]Redback(config-ctx)#subscriber name host8
[local]Redback(config-sub)#ip address 10.1.1.32 255.255.255.248
```

1.7 ip-address (RFlow)

ip-address *ip-v4-address* **context** *context-name*

no ip-address *ip-v4-address* **context** *context-name*

1.7.1 Purpose

Specifies the IP address of the external collector to which you want to export flow records.

1.7.2 Command Mode

Flow collector configuration

1.7.3 Syntax Description

ip-v4-address	Specifies the IP address of the external collector to which you want to export flow records.
context *context-name*	Identifies the context that hosts the interface to the external collector.

1.7.4 Default

None

1.7.5 Usage Guidelines

Use the **ip-address** command in flow collector configuration mode to specify the IP address of the external collector to which you want to export flow records.

Use the **no** form of this command to deny the exporting of flow records to an external collector.

1.7.6 Examples

The following example shows how to configure an external collector called **c1** to receive exported flow records from the SmartEdge router:

```
[local]Redback#configure

[local]Redback(config)#context foo

[local]Redback(config-ctx)#flow collector c1

[local]Redback(config-flow-collector)#ip-address 172.21.31.121 context ctx1
```

1.8 ip arp

ip arp *ip-addr mac-addr* [**alias**]

no ip arp *ip-addr mac-addr* [**alias**]

1.8.1 Purpose

Associates an IP address with a medium access control (MAC) address and creates a corresponding entry in the Address Resolution Protocol (ARP) table.

1.8.2 Command Mode

Context configuration

1.8.3 Syntax Description

ip-addr	Host IP address in the form ***A.B.C.D***.
mac-addr	MAC address of the host in the form ***hh:hh:hh:hh:hh:hh***.
alias	Optional. Configures the system to respond to ARP requests for the IP address.

1.8.4 Default

No entry is created in the ARP table.

1.8.5 Usage Guidelines

Use the **ip arp** command to associate an IP address with a MAC address and create a corresponding entry in the ARP table.

Note: If you enter both this command and the **ip subscriber arp** command (in subscriber configuration mode) and specify the same IP address and MAC address, the most recently updated command takes precedence. Only the circuit and interface are updated in the ARP table.

Use the **no** form of this command to remove an entry from the configuration and from the ARP table.

1.8.6 Examples

The following example shows how to associates IP address, **31.22.213.124**, with the MAC address, **00:30:23:32:12:82**, and creates a corresponding entry in the ARP table:

```
[local]Redback(config)#context local

[local]Redback(config-ctx)#ip arp 31.22.213.124 00:30:23:32:12:82
```

1.9 ip arp arpa

ip arp arpa

{**no** | **default**} **ip arp arpa**

1.9.1 Purpose

Enables the standard Address Resolution Protocol (ARP) on this interface.

1.9.2 Command Mode

Interface configuration

1.9.3 Syntax Description

This command has no keywords or arguments.

1.9.4 Default

Standard ARP is enabled.

1.9.5 Usage Guidelines

Use the **ip arp arpa** command to enable standard ARP on this interface.

Use the **no** form of this command to disable standard ARP on this interface.

Use the **default** form of this command to enable standard ARP on this interface.

1.9.6 Examples

The following example shows how to disable standard ARP on the **toToronto** interface at IP address, **10.20.1.1:**

```
[local]Redback(config-ctx)#interface toToronto
[local]Redback(config-if)#ip address 10.20.1.1 255.255.255.0
[local]Redback(config-if)#no ip arp arpa
```

1.10 ip arp delete-expired

ip arp delete-expired

{**no** | **default**} **ip arp delete-expired**

1.10.1 Purpose

Enables the automatic deletion of expired dynamic Address Resolution Protocol (ARP) entries associated with this interface from the ARP table.

1.10.2 Command Mode

Interface configuration

1.10.3 Syntax Description

This command has no keywords or arguments.

1.10.4 Default

Automatic deletion is disabled.

1.10.5 Usage Guidelines

Use the **ip arp delete-expired** command to enable the automatic deletion of expired dynamic ARP entries associated with this interface from the ARP table. Entries are deleted after they have been in the ARP table for the amount of time specified by the **ip arp timeout** command (in interface configuration mode). If the **ip arp timeout** command is not configured, the default value of 3,600 seconds (60 minutes) is used.

If you do not enable automatic deletion of expired dynamic ARP entries, expired entries are treated differently depending on the value of the ***seconds*** argument in the **ip arp timeout** command. If the value of the ***seconds*** argument is greater than 70, an ARP entry is refreshed unless no ARP reply is received in response to the refresh request packet. In that case, the entry is removed from the cache. If the value of the ***seconds*** argument is less than 70, expired entries are removed from the cache.

Use the **no** or **default** form of this command to disable the automatic deletion of expired entries.

1.10.6 Examples

The following example shows how to configure the system to automatically delete expired dynamic ARP entries on the `toBoston` interface at IP address, `10.30.2.1`:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#interface toBoston
[local]Redback(config-if)#ip address 10.30.2.1 255.255.255.0
[local]Redback(config-if)#ip arp delete-expired
```

1.11 ip arp maximum incomplete-entries

ip arp maximum incomplete-entries ***num-entries***

{**no** | **default**} **ip arp maximum incomplete-entries**

1.11.1 Purpose

Sets a maximum allowable number of incomplete entries for subscriber circuits that can exist in the Address Resolution Protocol (ARP) table for the context.

1.11.2 Command Mode

Context configuration

1.11.3 Syntax Description

num-entries	Maximum number of incomplete entries in the ARP table. The range of values is 1 to 4,294,967,295; the default value is 4,294,967,295.

1.11.4 Default

The maximum number of incomplete entries for subscriber circuits in the ARP table is 4,294,967,295.

1.11.5 Usage Guidelines

Use the **ip arp maximum incomplete-entries** command to set a maximum allowable number of incomplete entries for subscriber circuits that can exist in the ARP table for the context.

When requesting the medium access control (MAC) address that corresponds to a particular IP address, the SmartEdge router creates an incomplete entry in the ARP table and sends an ARP request packet. On reply, the entry is updated and complete.

Use the **no** or **default** form of this command to return to the default setting of a maximum of 4,294,967,295 incomplete entries for subscriber circuits in the ARP table.

1.11.6 Examples

The following example shows how to limit the number of incomplete entries in the ARP table to **250** for the **local** context:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#ip arp maximum 250
```

1.12 ip arp proxy-arp

ip arp proxy-arp [**always**]

{**no** | **default**} **ip arp proxy-arp**

1.12.1 Purpose

Enables the proxy Address Resolution Protocol (ARP) on this interface.

1.12.2 Command Mode

Interface configuration

1.12.3 Syntax Description

always	Optional. Indicates that proxy ARP must be functional for multiple hosts on the same circuit.

1.12.4 Default

Proxy ARP is disabled.

1.12.5 Usage Guidelines

Use the **ip arp proxy-arp** command to enable proxy ARP on this interface. When enabled, the SmartEdge router acts as an ARP proxy for hosts that are not on the same interface as the ARP request sender.

Note: You must enable standard ARP on this interface before you can enable proxy ARP; by default, standard ARP is enabled.

Proxy ARP and secured ARP are mutually exclusive services for an interface; enabling either service for an interface automatically disables the other service for that interface.

Use the **always** keyword to enable proxy ARP for multiple hosts that reside on the same circuit; if not specified, this capability is limited to hosts on individual circuits.

Use the **no** or **default** form of this command to disable proxy ARP on this interface.

Note: To disable only the support for multiple hosts on the same circuit, you must first disable proxy ARP, and then enable it without the **always** keyword.

1.12.6 Examples

The following example shows how to enable proxy ARP on the **fromBoston** interface at IP address, **10.2.3.4**, for all hosts on the circuit:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#interface fromBoston
[local]Redback(config-if)#ip address 10.2.3.4 255.255.255.0
[local]Redback(config-if)#ip arp proxy-arp always
```

1.13 ip arp secured-arp

ip arp secured-arp [**always**]

{**no** | **default**} **ip arp secured-arp**

1.13.1 Purpose

Enables the secured Address Resolution Protocol (ARP) on a specified interface.

1.13.2 Command Mode

Interface configuration

1.13.3 Syntax Description

always	Optional. Indicates that proxy ARP must be functional for multiple hosts on the same circuit.

1.13.4 Default

Secured ARP is disabled.

1.13.5 Usage Guidelines

Use the **ip arp secured-arp** command to enable secured ARP on a specified interface.

Note: You must enable standard ARP on this interface before you can enable secured ARP; by default, standard ARP is enabled.

Secured ARP and proxy ARP are mutually exclusive services for an interface; enabling either service for an interface automatically disables the other service for the same interface.

Use the **always** keyword to enable secured ARP for multiple hosts that reside on the same circuit; if not specified, this capability is limited to hosts on individual circuits.

When secured ARP is enabled, ARP requests received on an interface are not answered unless the request comes from the circuit known to contain the requesting host. ARP requests are sent by the interface only on the circuit known to contain the target host, and are not flooded to all circuits bound to an interface.

Use the **no** or **default** form of this command to disable secured ARP on this interface.

Note: To disable only the support for multiple hosts on the same circuit, you must first disable secured ARP, and then enable it without the **always** keyword.

1.13.6 Examples

The following example shows how to enable secured ARP on the interface, **sec-arp**, at IP address, **10.1.1.1**, for all hosts on the circuit:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#interface sec-arp
[local]Redback(config-if)#ip address 10.1.1.1 255.255.255.0
[local]Redback(config-if)#ip arp secured-arp always
```

1.14 ip arp timeout

ip arp timeout *seconds*

{**no** | **default**} **ip arp timeout**

1.14.1 Purpose

Configures how long Address Resolution Protocol (ARP) entries remain in the ARP table before automatic deletion (if configured).

1.14.2 Command Mode

Interface configuration

1.14.3 Syntax Description

seconds	Number of seconds after which an ARP entry is deleted from the ARP table. The range of values is 0 to 4,294,967; the default value is 3,600.

1.14.4 Default

ARP entries remain in the table for 3,600 seconds (1 hour).

1.14.5 Usage Guidelines

Use the **ip arp timeout** command to specify how long ARP entries remain in the ARP table.

If you do not use the **ip arp delete-expired** command (in interface configuration mode) to enable the automatic deletion of expired dynamic ARP entries, expired entries are treated differently depending on the value of the ***seconds*** argument in the **ip arp timeout** command. If the value of the ***seconds*** argument is greater than 70, an ARP entry is refreshed unless no ARP reply is received in response to the refresh request packet. In that case, the entry is removed from the cache. If the value of the ***seconds*** argument is less than 70, expired entries are removed from the cache.

Use the **no** or **default** form of this command to restore the timeout setting to its default value of 3,600 seconds.

1.14.6 Examples

The following example shows how to set the ARP timeout value for the **toToronto** interface at IP address, **10.30.2.1**, to two hours (**7200** seconds):

```
[local]Redback(config-ctx)#interface toToronto

[local]Redback(config-if)#ip address 10.30.2.1 255.255.255.0

[local]Redback(config-if)#ip arp timeout 7200
```

1.15 ip clear-df

ip clear-df

{**no** | **default**} **ip clear-df**

1.15.1 Purpose

Specifies that the IP header Don't Fragment (DF) flag should be ignored in any packet that is to be transmitted on this outbound interface when that packet is too large to be forwarded to a device with a smaller maximum transmission unit (MTU) than is required by the packet.

1.15.2 Command Mode

Interface configuration

1.15.3 Syntax Description

This command has no keywords or arguments.

1.15.4 Default

The IP header DF flag is honored.

1.15.5 Usage Guidelines

Use the **ip clear-df** command to specify that the IP header DF flag should be ignored in any packet that is to be transmitted on this outbound interface when that packet is too large to be forwarded to a device with a smaller MTU than is required by the packet. In this case, the DF flag is cleared in the resulting fragmented packets. The DF flag is not affected in packets that are not too large for the MTU of the device to which they are transmitted.

If you enter the **clear-df** command (in tunnel configuration mode) for a tunnel circuit, instead of this command, the DF flag is cleared in all packets that are transmitted on that Generic Routing Encapsulation (GRE) tunnel circuit. If you run both commands, the **clear-df** command takes precedence for that GRE tunnel circuit, and clears the DF flag in all packets transmitted on that tunnel circuit. For more information about the **clear-df** command (in tunnel configuration mode), see the *Command List*.

Use the **no** or **default** form of this command to honor the DF flag in all packets.

1.15.6 Examples

The following example shows how to specify that the DF flag should be ignored in large packets:

```
[local]Redback(config)#context isp1
[local]Redback(config-ctx)#interface large-packets
[local]Redback(config-if)#ip clear-df
```

1.16 ip domain-lookup

ip domain-lookup

no ip domain-lookup

1.16.1 Purpose

Enables the SmartEdge router to use Domain Name System (DNS) resolution to look up hostname-to-IP address mappings in the host table for the context.

1.16.2 Command Mode

Context configuration

1.16.3 Syntax Description

This command has no keywords or arguments.

1.16.4 Default

DNS lookup is disabled.

1.16.5 Usage Guidelines

Use the **ip domain-lookup** command to enable the SmartEdge router to use DNS resolution to look up hostname-to-IP address mappings in the host table for the context.

This command allows a user to ping or Telnet to a host using a hostname, instead of having to know the host's specific IP address. When a command references a hostname, the SmartEdge router consults the local host table to obtain the hostname-to-IP address mapping. If the information is not in the

local host table, the SmartEdge router generates a DNS query to resolve the hostname.

For DNS resolution to function, one or more DNS servers must be specified using the **ip name-servers** command. Hostnames that are statically entered into the local host table using the **ip host** command are also used for DNS resolution.

Use the **no** form of this command to disable DNS resolution lookup.

1.16.6 Examples

The following example shows how to enable DNS resolution:

```
[local]Redback(config-ctx)#ip domain-lookup
```

1.17 ip dmz

ip dmz source *ip-addr nat-addr* **context** *ctx-name*

no ip dmz source *ip-addr nat-addr* **context** *ctx-name*

1.17.1 Purpose

Configures the source and Network Address Translation (NAT) IP addresses for a demilitarized zone (DMZ) host server.

1.17.2 Command Mode

NAT policy configuration

1.17.3 Syntax Description

source *ip-addr*	Original source IP address for the DMZ host server on the private network.
nat-addr	NAT address. The IP address of the DMZ host server on the public network to which the source IP address is mapped.
context *ctx-name*	Name of the context in which the NAT address of the DMZ host server is defined for the interface that is used to forward packets after the source IP address is translated.

1.17.4 Default

No DMZ host server is configured.

1.17.5 Usage Guidelines

Use the **ip dmz** command to configure a DMZ host server.

Use the **no** form of this command to remove the DMZ host server from the configuration.

1.17.6 Examples

The following example shows how to configure a DMZ host server with an internal network address, **10.1.1.1**, and an external network address, **201.1.1.1**,which are defined in the **local** context:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#nat policy policy1
[local]Redback(config-policy-nat)#ip dmz source 10.1.1.1 201.1.1.1 context local
```

1.18 ip domain-name

ip domain-name *name*

no ip domain-name *name*

1.18.1 Purpose

Creates a Domain Name System (DNS) name (or alias) for the context.

1.18.2 Command Mode

Context configuration

1.18.3 Syntax Description

name	Name (or alias) of the domain for the context.

1.18.4 Default

No domain names are created for the context.

1.18.5 Usage Guidelines

Use the **ip domain-name** command to create a domain name (or alias) for the context.

You can create up to six domain names for each context.

Use the **no** form of this command to remove the domain name (or alias) from the configuration.

1.18.6 Examples

The following example shows how to create a domain name for the **local** context, **Ericsson.com:**

```
[local]Redback(config-ctx)#ip domain-name redback.com
```

1.19 ip host (context)

ip host ***hostname ip-addr***

no ip host ***hostname ip-addr***

1.19.1 Purpose

Creates a static hostname-to-Internet Protocol version 4 (IPv4) address Domain Name System (DNS) mapping in the host table for the context.

1.19.2 Command Mode

Context configuration

1.19.3 Syntax Description

hostname	Name of the host.
ip-addr	IPv4 address of the host.

1.19.4 Default

No static mappings are preconfigured.

1.19.5 Usage Guidelines

Use the **ip host** command to create a static hostname-to-IPv4 address DNS mapping in the host table for the context.

You can create up to 64 static entries in the host table. The SmartEdge router always consults the host table prior to generating a DNS lookup query.

Use the **no** form of this command to remove the specified static entry. Specifying a new IPv4 address for an existing hostname removes the previously specified IPv4 address.

1.19.6 Examples

The following example shows how to statically map the hostname, **hamachi**, to the IPv4 address, **192.168.42.105:**

```
[local]Redback(config-ctx)#ip host hamachi 192.168.42.105
```

1.20 ip host (port)

ip host *ip-addr*[*/prefix-length* | *mac-addr*]

no ip host *ip-addr*[*/prefix-length* | *mac-addr*]

1.20.1 Purpose

Associates an IPoE-encapsulated (the default) or 802.1Q-encapsulated Ethernet port with the IP address and medium access control (MAC) address of the remote host on the circuit.

1.20.2 Command Mode

Port configuration

1.20.3 Syntax Description

ip-addr	IP address of the host on this circuit in the form *A.B.C.D*.
prefix-length	Optional. Destination subnet. The range of values is 0 to 32.
mac-addr	Optional. MAC address of the remote host on this circuit in the form *hh:hh:hh:hh:hh:hh*.

1.20.4 Default

No IP host address is associated with the port

1.20.5 Usage Guidelines

Use the ip host command to associate an IPoE-encapsulated (the default) or 802.1Q-encapsulated Ethernet port with the IP address and MAC address of the remote host on the circuit. Configuring the port with the **ip host** command enables the use of multibind interfaces (connecting multiple hosts to one subnet) for the default circuit of the port.

Note: You can only use this command on the default circuit of the specified port.

You can associate multiple IP entries with a port by repeating the **ip host** command with different arguments. You may add up to four IP hosts for each port.

You can use this command on any Ethernet port that can be bound to multibind interfaces, including all ports with IPoE encapsulation and 802.1Q encapsulation, and including the Ethernet management port. However, you cannot use this command on ports whose interfaces are bound through a link-group.

To configure IP addresses for 802.1Q PVCs on the port, use the **ip host** command in dot1q pvc configuration mode.

Note: This command is also documented in *Configuring Circuits* for permanent virtual circuits (PVCs) and in *Configuring Single Circuit Tunnels* for single-circuit tunnels.

Use the **no** form of this command to remove IP associations from the configuration of this port.

1.20.6 Examples

The following example shows how to associate an Ethernet port with the IP address of the host on the PVC:

```
[local]Redback(config)#port ethernet 12/1

[local]Redback(config-port)#bind interface hello-int2 hello

[local]Redback(config-port)#ip host 10.10.10.14/24
```

1.21 ip host (PVC)

ip host *ip-addr*[*/prefix-length* | *mac-addr*]

no ip host *ip-addr*[*/prefix-length* | *mac-addr*]

1.21.1 Purpose

Associates an 802.1Q, Asynchronous Transfer Mode (ATM), or Frame Relay permanent virtual circuit (PVC) with the IP address and medium access control (MAC) address of the remote host on the circuit.

1.21.2 Command Mode

- ATM PVC configuration
- dot1q PVC configuration
- Frame Relay PVC configuration
- link PVC configuration

1.21.3 Syntax Description

ip-addr	IP address of the host on this circuit in the form *A.B.C.D*.
prefix-length	Optional. Destination subnet. The range of values is 0 to 32.
mac-addr	Optional. MAC address of the remote host on this circuit in the form *hh:hh:hh:hh:hh:hh*.(1)

*(1) The **mac-addr** argument applies to only IP over Ethernet (IPoE) ATM PVCs; that is, to ATM PVCs with multiprotocol encapsulation.*

1.21.4 Default

No IP host address or MAC address is associated with the PVC.

1.21.5 Usage Guidelines

Use the **ip host** command to associate an 802.1Q, an ATM, or a Frame Relay PVC with the IP address of the host on the circuit.

Use this command only for an 802.1Q, an ATM, or a Frame Relay PVC that you intend to bind to an interface.

You can associate multiple IP entries with an ATM PVC by repeating the **ip host** command using different arguments. Up to 64 IP hosts may be added for each ATM PVC.

Note: This command is available only for individual PVCs; you cannot enter it if you have created or selected a range of PVCs. You must first select the individual PVC before you can enter this command.

Note: This command is not available for an 802.1Q or ATM PVC that you intend to cross-connect.

Note: The ***mac-addr*** argument is not available for a Frame Relay PVC or for an ATM PVC for which you have specified **route1483** encapsulation.

Use the **no** form of this command to delete the association.

Note: This command is also documented in *Configuring ATM, Ethernet, and POS Ports* and in *Configuring Single Circuit Tunnels*

1.21.6 Examples

The following example shows how to associate an ATM PVC on an ATM OC port with the IP address of the host on the PVC:

```
[local]Redback(config)#port atm 2/1
[local]Redback(config-atm-oc)#atm pvc 3 32 profile 1.vbrrt encapsulation route1483
[local]Redback(config-atm-pvc)#bind interface foo local
[local]Redback(config-atm-pvc)#ip host 10.10.10.14/24
```

The following example shows how to create a multiprotocol ATM PVC on an ATM OC port and, because it is not to be cross-connected, associate an IP address and MAC address with it, and bind it to an interface:

```
[local]Redback(config)#port atm 2/1
[local]Redback(config-atm-oc)#atm pvc 4 210 profile cbr1 encapsulation multi
[local]Redback(config-atm-pvc)#bind interface ip-out local
[local]Redback(config-atm-pvc)#ip host 1.1.1.4 00:30:88:01:01:01
```

1.22 ip icmp

ip icmp suppress packet-too-big

{**no** | **default**} **ip icmp**

1.22.1 Purpose

Specifies that the Internet Control Message Protocol (ICMP) Destination Unreachable packet-too-big message should be suppressed when any packet that is to be transmitted on this interface has its Don't Fragment (DF) flag set, and is too large to be forwarded without fragmentation.

1.22.2 Command Mode

Interface configuration

1.22.3 Syntax Description

suppress packet-too-big	Suppresses the generation of the ICMP Destination Unreachable packet-too-big message.

1.22.4 Default

ICMP Destination Unreachable packet-too-big messages are generated.

1.22.5 Usage Guidelines

Use the **ip icmp** command to specify that the ICMP Destination Unreachable packet-too-big message should be suppressed when any packet that is to be transmitted on this interface has its DF flag set, and is too large to be forwarded without fragmentation.

Use the **no** or **default** form of this command to generate ICMP Destination Unreachable packet-too-big messages.

1.22.6 Examples

The following example shows how to suppress the Destination Unreachable packet-too-big messages:

```
[local]Redback(config)#context isp1
[local]Redback(config-ctx)#interface large-packets
[local]Redback(config-if)#ip icmp suppress packet-too-big
```

1.23 ip igmp service-profile

ip igmp service-profile *prof-name*

no ip igmp service-profile *prof-name*

1.23.1 Purpose

Enables an existing Internet Group Management Protocol (IGMP) service profile on a single subscriber record, a named subscriber profile, or a default subscriber profile.

1.23.2 Command Mode

Subscriber configuration

1.23.3 Syntax Description

prof-name	Name of the IGMP service profile enabled on the subscriber profile.

1.23.4 Default

None

1.23.5 Usage Guidelines

Use the **ip igmp service-profile** command to enable a existing IGMP service profile on a single subscriber record, a named subscriber profile, or a default subscriber profile. The service profile used is determined in the following order:

- Subscriber profile
- Default subscriber profile
- Service profile configured on the subscriber's parent interface

If a service profile is not defined in the subscriber record, it inherits the service profile from the default subscriber profile. If the default subscriber profile is not configured with a service profile, the service profile configured on the interface is used.

Use the **no** form of this command to disable the service profile on the subscriber.

1.23.6 Examples

The following example shows how enable the IGMP service profile, **sp04**, on the **default** subscriber profile:

```
[local]Redback(config-ctx)#subscriber default
[local]Redback(config-sub)#ip igmp service-profile sp04
```

1.24 ip interface

ip interface name ***if-name***

no ip interface name ***if-name***

1.24.1 Purpose

Configure hosts to use a specific Dynamic Host Configuration Protocol (DHCP) interface to acquire address information for a subscriber's circuit.

1.24.2 Command Mode

Subscriber configuration

1.24.3 Syntax Description

name ***if-name***	DHCP interface name.

1.24.4 Default

The subscriber is bound to the first available DHCP interface.

1.24.5 Usage Guidelines

Use the **ip interface** command to configure hosts to use a specific DHCP interface to acquire address information for a subscriber's circuit.

You must enable the specified interface for DHCP proxy or DHCP relay using the **dhcp proxy** or **dhcp relay** command (in interface configuration mode), respectively.

You must use the **dhcp max-addr** command (in subscriber configuration mode) to enable hosts to acquire address information for the subscriber's circuit.

Use the **no** form of this command to restore the default condition where the subscriber is bound to the first available DHCP interface.

1.24.6 Examples

The following example shows how to create an interface and specifies that hosts use the DHCP **if-dhcp** interface to acquire address information for the circuit used by the **sub-dhcp** subscriber:

```
[local]Redback(config-ctx)#interface name if-dhcp
[local]Redback(config-if)#ip address 10.1.1.1 255.255.255.0
[local]Redback(config-if)#dhcp relay
[local]Redback(config-if)#exit
[local]Redback(config-ctx)#subscriber name sub-dhcp
[local]Redback(config-sub)#dhcp max-addr 3
[local]Redback(config-sub)#ip interface name if-dhcp
```

1.25 ip martian

ip martian *ip-addr/prefix-length* [**eq** *eq-value*] [**ge** *ge-value*] [**le** *le-value*]

no ip martian *ip-addr/prefix-length* [**eq** *eq-value*] [**ge** *ge-value*] [**le** *le-value*]

1.25.1 Purpose

Adds custom IP martian addresses to the list of default martian IP addresses in the routing table.

1.25.2 Command Mode

Context configuration

1.25.3 Syntax Description

ip-addr/prefix-length	IP address (in the form ***A.B.C.D***) and prefix length, separated by the slash (**/**) character. The range of values for the ***prefix-length*** argument is 0 to 32.
eq ***eq-value***	Optional. Equal to value. The ***eq-value*** argument specifies the length of the mask to be matched; the **eq** keyword indicates that the mask length must exactly match the specified value. The range of values for the ***eq-value*** argument is 1 to 32.
ge ***ge-value***	Optional. Greater than or equal to value. The ***ge-value*** argument specifies the length of the mask to be matched; the **ge** keyword indicates that all masks of a length greater than or equal to the specified value will match. The range of values for the ***ge-value*** argument is 1 to 32.
le ***le-value***	Optional. Less than or equal to value. The ***le-value*** argument specifies the length of the mask to be matched; the **le** keyword indicates that all masks of a length less than or equal to the specified value will match. The range of values for the ***le-value*** argument is 1 to 32.

1.25.4 Default

For IPv4, the martian addresses of 0.0.0.0/8 and 127.0.0.0/8 are installed in the routing table.

1.25.5 Usage Guidelines

Use the **ip martian** command to add custom IP martian addresses to the list of default martian IP addresses in the routing table.

IP martian addresses are host or network addresses about which all routing information is ignored. IP martian addresses are typically advertised by misconfigured routers using dynamic protocols.

Use the **no** form of this command to remove a configured IP martian address from the routing table.

1.25.6 Examples

The following example shows how to configure a martian address of **10.1.0.0/20** for the **local** context. Routes matching this prefix are ignored:

```
[local]Redback(config-ctx)#ip martian 10.1.0.0/20
```

1.26 ip maximum-routes

ip maximum-routes[**multicast**] [**vpn**] ***route-limit*** [**log-only** | **threshold** ***value***] [**mid-threshold** ***value***]

1.26.1 Purpose

Configures an upper limit for the number of routes installed in an IP routing table.

1.26.2 Command Mode

Context configuration

1.26.3 Syntax Description

multicast	Optional. Sets the maximum route limit for unicast routes in a multicast topology.
vpn	Optional. Sets the maximum route limit for all non-local context unicast routing tables. When the **vpn** keyword is used in the local context, it specifies a default maximum route setting that automatically applies to all non-local contexts; however, if the **ip maximum-route** command is used in a specific non-local context, then it overrides the default maximum route setting.
route-limit	Maximum number of routes allowed in the IP routing table. If this limit is reached and a protocol instance attempts to add additional routes, the routes are rejected, a warning is triggered, and the protocol instance is shut down. Range of values is 1 to 4294967295.
log-only	Optional. Configures the route limit as an advisory limit. An advisory limit triggers only a warning, and additional routes are not rejected.
threshold*value*	Optional. Threshold value for the mandatory limit that triggers a warning. Range of values is 1 to 100.
mid-threshold ***value***	Optional. Threshold value for the mid-level limit that triggers a warning. Range of values is 1 to 100.

1.26.4 Default

No maximum limit is set.

1.26.5 Usage Guidelines

Use the **ip maximum-routes** command to configure an upper limit for the number of routes installed in an IP routing table.

A route limit sets an upper limit for the number of prefixes installed in a routing table; for example, you can use a route limit to limit the number of routes received from the customer edge (CE) router in a Virtual Private Network (VPN) context.

There are two modes for route limits: advisory (log only) and mandatory. An advisory limit only triggers warnings; a mandatory limit triggers warnings, rejects any additional routes after the maximum is reached, and shuts down the offending protocol instance (the instance that exceeds the limit). When the maximum is exceeded, you can clear the condition either by reconfiguring the route limit or by using the **clear ip route** command with the **maximum-routes** keyword. Clearing the condition also re-enables any routing protocol instances that were shut down when the maximum route limit was exceeded.

Use the **vpn** keyword in the local context to specify a default maximum route setting that automatically applies to all non-local contexts. To override the default maximum route setting, use the **ip maximum-route** command in the non-local context that you want to configure.

1.26.6 Examples

The following example shows how to configure a (mandatory) upper limit of **500** routes for the IP routing table:

```
[local]Redback#config context ip
[local]Redback(config-ctx)#ip maximum-routes 500
```

The following example shows how to :

- Print a warning log when the number of routes in the context reaches 800.
- Clear all routes and notifies all routing protocols when the number of routes in the context reaches 1000. No new routes will be accepted until a new maximum-routes configuration is set.

```
[local]Redback(config-ctx)#ip maximum-routes 1000 threshold 80
```

The following example configures the system to log a warning when 1000 routes are reached:

Note: No action is required and the number of routes is allowed to grow.

```
[local]Redback(config-ctx)#ip maximum-routes 1000 log-only
```

The following example shows how to :

- Print a warning log when the number of routes in the context reaches 800.
- Clear all routes and notifies all routing protocols when the number of routes in the context reaches 1000 or 50. No new routes will be accepted until a new maximum-routes configuration is set.

```
[local]Redback(config-ctx)#ip maximum-routes 1000 threshold 80 mid-threshold 50
```

Note: If configuring IPv4 and IPv6 maximum-routes to be used together or if using IPv6 maximum-routes only, refer to Section 1.66.6 on page 107 for configuration examples.

1.27 ip mstatic

ip mstatic ***source-ip-addr/prefix-length*** {***rpf-ip-addr*** | ***rpf-if-name***} [***distance***]

no ip mstatic ***source-ip-addr/prefix-length*** {***rpf-ip-addr*** | ***rpf-if-name***} [***distance***]

1.27.1 Purpose

Configures a static route for multicast reverse path forwarding (RPF) lookup.

1.27.2 Command Mode

Context configuration

1.27.3 Syntax Description

source-ip-addr/prefix-length	IP address of the multicast source (in the form ***A.B.C.D***) and prefix length, separated by the slash (**/**) character. The range of values for the ***prefix-length*** argument is 0 to 32.
rpf-ip-addr	IP address of the RPF neighbor or route.

rpf-if-name	Interface name used for the RPF lookup.
distance	Optional. Administrative distance assigned to the static route used for RPF lookup. The range of values for the ***distance*** argument is 1 to 255.

1.27.4 Default

None

1.27.5 Usage Guidelines

Use the **ip mstatic** command to configure a static route for multicast RPF lookup.

Use the **no** form of this command to delete a static route for multicast RPF lookup.

1.27.6 Examples

The following example shows how to configure a static route for multicast RPF lookup with the source IP address 192.168.100.0 and a prefix length of 24. The IP address of the RPF neighbor is 192.168.101.1. The route uses the RPF neighbor IP address to perform the RPF lookup and is assigned an administrative distance of 110:

```
[local]Redback(config)#context isp1

[local]Redback(config-ctx)#ip mstatic 192.168.100.0/24 192.168.101.1 110
```

1.28 ip mtu

ip mtu ***bytes***

{no | default} ip mtu

1.28.1 Purpose

Sets the maximum transmission unit (MTU) size for traffic sent on the circuit to which the interface is bound.

1.28.2 Command Mode

Interface configuration

1.28.3 Syntax Description

bytes	MTU size in bytes. The range of values is 256 to 16,384.

1.28.4 Default

MTU for the media type of the port or circuit to which the interface is bound.

1.28.5 Usage Guidelines

Use the **ip mtu** command to set the MTU size for traffic sent on the circuit to which the interface is bound. If an IP packet exceeds the MTU configured for an interface, the system fragments that packet.

Note: This command does not apply to loopback interfaces.

An interface does not have an MTU size until either one is explicitly configured using the **ip mtu** command, or a circuit is bound to the interface. If no MTU size is configured, the MTU size is the same as that of the bound circuit. If an IP MTU is explicitly configured, the resulting IP MTU is calculated. It is the lesser of the configured IP MTU and the circuit MTU.

Note: In an OSPF routing scenario, the MTU value must be explicitly configured on a multibind interface to match the MTU on the remote end. Unlike the behavior for a single-bind interface, there is no default value for this parameter. If the MTU value is not configured on a multibind interface, the OSPF neighbor does not come up.

Use the **no** or **default** form of this command to remove the IP MTU and use the MTU of the bound circuit.

1.28.6 Examples

The following example sets the maximum IP packet size for the **atm1** interface to **300** bytes:

```
[local]Redback(config-ctx)#interface to_sj1

[local]Redback(config-if)#ip mtu 300
```

1.29 ip multicast boundary

ip multicast boundary *acl-name*

no ip multicast boundary *acl-name*

1.29.1 Purpose

Configures an administratively scoped boundary for multicast routing.

1.29.2 Command Mode

Interface configuration

1.29.3 Syntax Description

acl-name	Name of the access control list (ACL) that controls the range of group addresses affected by the boundary.

1.29.4 Default

None

1.29.5 Usage Guidelines

Use the **ip multicast boundary** command to configure an administratively scoped boundary for multicast routing. This boundary prevents forwarding of multicast data packet destined for group addresses denied by the ACL.

Use the **no** form of this command to remove the multicast boundary from the interface.

1.29.6 Examples

The following example shows how to configure an administratively scoped boundary for multicast using ACL **20:**

```
[local]Redback(config-ctx)#interface enet01
[local]Redback(config-if)#ip multicast boundary 20
```

1.30 ip multicast receive

ip multicast receive {**permit** | **deny**}

no ip multicast receive

1.30.1 Purpose

Configures the multicast receive permissions for a subscriber record, a named subscriber profile, or a default subscriber profile.

1.30.2 Command Mode

Subscriber configuration

1.30.3 Syntax Description

permit	Allows the subscriber to receive multicast traffic.
deny	Denies the subscriber the ability to receive multicast traffic.

1.30.4 Default

The multicast receive permission is set to permit.

1.30.5 Usage Guidelines

Use the **ip multicast receive** command to configure the multicast receive permissions for a subscriber record, a named subscriber profile, or a default subscriber profile. Permission attributes are applied in the following order:

- Subscriber profile
- Default subscriber profile
- System defaults

If a permission is not defined in the subscriber, it inherits the value of the permission from the default subscriber profile. If the permission is not defined in the default subscriber profile, the system default values are used.

For multicast routing to function on subscribers, you must use the **pim sparse-mode passive** command in interface configuration mode to enable Protocol Independent Multicast Sparse-Mode (PIM-SM) on the interface.

For multicast routing to function on subscribers, you must use the **pim sparse-mode passive** command in interface configuration mode to enable Protocol Independent Multicast Sparse Mode (PIM-SM) on the interface.

Use the **no** form of this command to delete receive permissions for the profile to which the command is applied.

1.30.6 Examples

The following example shows how to setets receive permissions to **permit** for the default subscriber profile:

```
[local]Redback(config-ctx)#subscriber default
[local]Redback(config-sub)#ip multicast receive permit
```

The following example shows how to set receive permissions to **deny** for subscriber **freddy:**

```
[local]Redback(config-ctx)#subscriber name freddy
[local]Redback(config-sub)#ip multicast receive deny
```

1.31 ip multicast send

ip multicast send {**permit** [**unsolicit**] | **deny**}

no ip multicast send

1.31.1 Purpose

Configures the multicast send permissions for a subscriber record, a named subscriber profile, or a default subscriber profile.

1.31.2 Command Mode

Subscriber configuration

1.31.3 Syntax Description

permit	Allows the subscriber to send multicast traffic.
unsolicit	Optional. Used in conjunction with the **permit** keyword to indicate that the subscriber is allowed to send unsolicited multicast traffic.
deny	Denies the subscriber the ability to send multicast traffic.

1.31.4 Default

The multicast send permission is set to deny.

1.31.5 Usage Guidelines

Use the **ip multicast send** command to configure the multicast send permissions for a subscriber record, a named subscriber profile, or a default subscriber profile.

If the **permit** keyword is used without the **unsolicit** keyword, the subscriber must join a group prior to sending unsolicited multicast data. If used together (**permit unsolicit**), a subscriber is allowed to send unsolicited multicast traffic. Permissions are examined in the following order:

- Subscriber profile
- Default subscriber profile
- System defaults.

If a permission is not defined in the subscriber profile, it inherits the value of the permission from the default subscriber profile. If the permission is undefined in the default subscriber profile, the system default values are used.

For multicast routing to function on subscribers, you must use the **pim sparse-mode** command in interface configuration mode to enable Protocol Independent Multicast Sparse-Mode (PIM-SM) on the interface.

For multicast routing to function on subscribers, you must use the **pim sparse-mode passive** command in interface configuration mode to enable Protocol Independent Multicast Sparse Mode (PIM-SM) on the interface.

Use the **no** form of this command to delete all send permissions for the profile. Deleting the permissions in a subscriber profile causes the system to use the permissions from the default subscriber profile. If no such permissions exist in the default subscriber profile, the system default is used.

1.31.6 Examples

The following example shows how to configure the default subscriber profile with the permission to send multicast traffic; however, subscriber **mike** is denied sending multicast traffic:

```
[local]Redback(config-ctx)#subscriber default
[local]Redback(config-sub)#ip multicast send permit
[local]Redback(config-sub)#exit
[local]Redback(config-ctx)#subscriber name mike
[local]Redback(config-sub)#ip multicast send deny
```

The following example (using the **no** form) shows how to delete send permissions in the default subscriber profile; however, the system default for multicast send is permit, so the subscriber **jane** can send and receive multicast traffic:

```
[local]Redback(config-ctx)#subscriber default
[local]Redback(config-sub)#no ip multicast send
[local]Redback(config-sub)#exit
[local]Redback(config-ctx)#subscriber name jane
[local]Redback(config-sub)#ip address 10.10.1.4
[local]Redback(config-sub)#exit
```

1.32 ip name-servers

ip name-servers *primary-ip-addr* [*secondary-ip-addr*]

no ip name-servers

1.32.1 Purpose

Specifies the Internet Protocol version 4 (IPv4) address of a primary (and, optionally, a secondary) Domain Name System (DNS) server.

1.32.2 Command Mode

Context configuration

1.32.3 Syntax Description

primary-ip-addr	IPv4 address of the primary DNS server.
secondary-ip-addr	Optional. IPv4 address of the secondary DNS server.

1.32.4 Default

No DNS server IPv4 addresses are preconfigured.

1.32.5 Usage Guidelines

Use the **ip name-servers** command to specify the IPv4 address of a primary (and, optionally, a secondary) DNS server.

For DNS resolution to function, you must configure domain-name lookup using the **ip domain-lookup** command (in context configuration mode), and there must be an IP route to the DNS servers.

Use the **no** form of this command to remove the specified DNS server association. If you delete the primary DNS server, any configured secondary DNS server becomes the primary server.

1.32.6 Examples

The following command shows how to configure an association with a primary DNS server at IPv4 address, **128.215.33.47**, and a secondary server at IPv4 address, **196.145.92.33:**

```
[local]Redback(config-ctx)#ip name-servers 128.215.33.47 196.145.92.33
```

The following command shows how to remove the primary DNS server, making the server that was previously the secondary into the primary:

```
[local]Redback(config-ctx)#no ip name-servers 128.215.33.47
```

1.33 ip nat

ip nat *pol-name* [**acl-counters**]

no ip nat *pol-name* [**acl-counters**]

1.33.1 Purpose

Attaches a Network Address Translation (NAT) policy to packets received or transmitted on any circuit bound to the specified interface and optionally enables ACL counters for the policy..

1.33.2 Command Mode

Interface configuration

1.33.3 Syntax Description

pol-name	NAT policy name.
acl-counters	Enables ACL counters for circuits bound to the interface in which the policy is attached.

1.33.4 Default

None

1.33.5 Usage Guidelines

Use the **ip nat** command to attach a NAT policy to packets received or transmitted on any circuit bound to the specified interface. You can use the optional **acl-counters** keyword, to enable ACL counters.

Use the **no** form of this command to remove the NAT policy from the interface.

1.33.6 Examples

The following example shows how to translate an IP source address for the **p1** NAT policy and apply the policy to packets traveling across the **pos1** interface:

```
[local]Redback(config-ctx)#nat policy p1
[local]Redback(config-policy-nat)#ip static in source 10.1.2.3 32.32.32.32
[local]Redback(config-policy-nat)#exit
[local]Redback(config-ctx)#interface pos1
[local]Redback(config-if)#ip nat p1
```

1.34 ip nat pool

ip nat pool *pool-name* [**napt** [**multibind** | **paired-mode**]] [**logging**]

no ip nat pool *pool-name* [**napt** [**multibind** | **paired-mode**]] [**logging**]

1.34.1 Command Mode

Context configuration

1.34.2 Syntax Description

pool-name	NAT pool name.
napt	Optional. Enables support for translation of Transmission Control Protocol/User Datagram Protocol (TCP/UDP) ports.
multibind	Optional. Enables the NAT pool to be applied to multibind interfaces. The **multibind** and **paired-mode** keywords are mutually exclusive.
paired-mode	Specifies that each subscriber associated with the pool is always assigned the same external IP address. If you use the **paired-mode** keyword, the pool is still applied to multibind interfaces.
logging	Optional with the NAPT keyword. Enables NAT pool logging.

1.34.3 Default

None

1.34.4 Usage Guidelines

Use the **ip nat pool** command to configure a Network Address Translation (NAT) pool. You can also enable NAPT multibind mode and/or paired mode for the pool.

You can also use the **logging** keyword to enable NAT pool logging.

The command also enters NAT pool configuration mode where you can configure pool attributes such as an access group to classify dynamic NAT, the action to take with default traffic, or ICMP notification.

Use the **no** form of this command to remove a NAT pool. It is not possible to change the type of a pool. You must delete and recreate it to change the type.

1.34.5 Examples

The following example shows how to configure a NAT pool, `NAT-POOL-BASIC`, with 14 IP addresses (`171.71.71.4` to `171.71.71.7` and `171.71.71.101` to `171.71.71.110`):

```
[local]Redback(config-ctx)#ip nat pool NAT-POOL-BASIC
[local]Redback(config-nat-pool)#address 171.71.71.4 255.255.255.252
[local]Redback(config-nat-pool)#address 171.71.71.101 to 171.71.71.110
```

The following example shows how to configure an enhanced NAT pool, `nat-pool-1` with NAPT paired-mode, and logging enabled. It also configures the logging profile to be associated with the pool, and the address range for the pool (which excludes well-known port numbers:

```
[local]rock1200(config-ctx)#ip nat pool nat-pool-1 napt paired-mode logging
[local]rock1200(config-nat-pool)#logging-profile nat-log-profile context nat-context
[local]rock1200(config-nat-pool)#paired-mode subscriber over-subscription 32 port-limit 4096
[local]rock1200(config-nat-pool)#address 100.1.1.1 to 100.1.1.2
[local]rock1200(config-nat-pool-record)#exclude well-known
```

1.35 ip pool (context configuration)

ip pool {**falling-threshold** *num* {**trap** [**log**] | **log**} | **options use-class-c-bcast-addrs**}

no ip pool {**falling-threshold** | **options use-class-c-bcast-addrs**}

1.35.1 Purpose

Specifies context-specific falling-threshold parameters or includes Class C network and broadcast IP addresses in IP pools in the context.

1.35.2 Command Mode

Context configuration

1.35.3 Syntax Description

falling-threshold *num*	Threshold value for creating a falling-threshold crossing event. The range of values is 0 to 4,294,967,295.
trap	Reports the falling-threshold event with a Simple Network Management Protocol (SNMP) event.
log	Logs the falling-threshold event. Optional only if you specify the **trap** keyword.
options use-class-c-bcast-addrs	Allows Class C network (.0) and broadcast (.255) IP addresses in all configured IP pools in this context.

1.35.4 Default

No threshold parameters are defined for any context; Class C network and broadcast IP addresses are excluded.

1.35.5 Usage Guidelines

Use the **ip pool** command (in context configuration mode) to specify falling-threshold parameters or to include Class C network and broadcast IP addresses in IP pools for the context.

The falling-threshold parameters provide an alert when the number of available IP addresses for all IP pools in the context is reduced to the value specified. This value is unaffected if any threshold for an individual IP pool is altered.

Use the **falling-threshold** *num* construct to specify the total number of available IP addresses in all pools in the context, for which a falling-threshold crossing event is generated. A crossing event occurs only when the total number of available IP addresses in all pools in the context equals the value specified. If the number of available IP addresses becomes greater than the value specified, and then drops again to the value, a second falling-threshold crossing event is generated.

If you specify the **falling-threshold** *num* construct and the threshold parameters already exist, the current falling threshold parameters are set to the new values, or are added to the definition of the context if they did not previously exist. If you specify a value that is larger than the sum of all IP addresses in all IP pools in the context, no threshold event can occur at the context level. To remove the threshold, specify **0** for the *num* argument.

You can specify that the falling-threshold crossing event be reported with an SNMP trap, a log message, or both the trap and the log message.

By default, network (.0) and broadcast (.255) IP addresses are excluded in any IP pool of Class C IP addresses, even when that pool is supernetted; you must specify the **options use-class-c-bcast-addrs** construct to include the intervening Class C network and broadcast addresses in the range. For example:

- If you do not specify this option, and you configure the pool with an IP address of 192.200.100.0/23, IP addresses 192.200.100.0, 192.200.100.255, 192.200.101.0, and 192.200.101.255 are excluded in the pool.
- If you do not specify this option, 192.200.100.255 and 192.200.101.0 are included.

For more information about guidelines for IP addresses in IP pools and the description of the **ip pool** command (in interface configuration mode), see the *Command List*.

Use the **no** form of this command to remove context-specific threshold parameters to exclude intervening Class C network and broadcast IP addresses in any IP pool in the context.

1.35.6 Examples

The following example specifies that an SNMP trap and a log message be generated for the **isp1.net** context when the available IP addresses in all IP pools in the context equals **1,000:**

```
[local]Redback(config)#context isp1.net

[local]Redback(config-ctx)#ip pool falling-threshold 1000 trap log
```

1.36 ip pool (interface configuration)

ip pool *ip-addr* {*netmask* | */ prefix-length* | **to** *ip-addr*} [**name** *pool-name*] [**falling-threshold** *num* {**trap** [**log**] | **log**}]

{**no** | **default** } **ip pool** *ip-addr* {*netmask* | */ prefix-length* | **to** *ip-addr*} [**name** *pool-name*] [**falling-threshold** *num* {**trap** [**log**] | **log**}]

1.36.1 Purpose

Creates or modifies a pool of IP addresses for an interface to allow a subscriber on a Point-to-Point Protocol (PPP) or PPP over Ethernet (PPPoE)-encapsulated circuit to be assigned any available IP address from the pool.

1.36.2 Command Mode

Interface configuration

1.36.3 Syntax Description

ip-addr	Starting IP address of the IP pool in the form ***A.B.C.D***.
netmask	Network mask for the associated IP network in the form ***A.B.C.D***. The range of values is 255.255.0.0 to 255.255.255.255.
prefix-length	Prefix length. The range of values is 16 to 32.
to ***ip-addr***	Ending address of the IP pool.
name ***pool-name***	Optional. Name for the IP pool; a string with up to 31 characters.
falling-threshold ***num***	Optional. Threshold value for creating a falling-threshold crossing event. The range of values is 0 to 65,535; if omitted, the default value is 0.
trap	Reports the falling-threshold event with a Simple Network Management Protocol (SNMP) event.
log	Logs the falling-threshold event; this keyword is optional if you specify the **trap** keyword.

1.36.4 Default

No IP pool is created for any interface.

1.36.5 Usage Guidelines

Use the **ip pool** command (in interface configuration mode) to create or modify a pool of IP addresses for an interface to allow a subscriber on a PPP- or PPPoE-encapsulated circuit to be assigned an IP address from the pool. The interface must have been created using the **interface** command (in context configuration mode) with the **multibind** keyword.

You can use IP pools to provide addresses for the Dynamic Host Configuration Protocol (DHCP) server; specifically, if no range of values is specified for a DHCP subnet, the DHCP server takes the IP addresses from the IP pool defined by the **interface** command (in context configuration mode). This IP pool can be used by the DHCP server and PPP subscribers on the same interface.

Note: This command does not apply to loopback interfaces.

To create the pool, specify an IP address within the range for the pool and either the netmask or the prefix length. You can enter this command multiple times if you are configuring a last-resort interface.

The number of available IP addresses in a pool is decremented whenever an IP address is assigned from the pool and incremented when it is returned to the pool.

If you use the RADIUS to authenticate subscribers, follow these guidelines:

- You must ensure that the RADIUS server is configured to return attribute 8, Framed-IP-Address, with a value of 255.255.255.254 or 0.0.0.0. These values allow the subscriber to be assigned any available IP address from any pool configured within the context.
- If you create a named pool, you must ensure that the RADIUS server is configured to return VSA 36 provided by Ericsson, IP-Address-Pool-Name, with the name of the IP pool.

The name that you specify for the IP pool (the ***pool-name*** argument) can be the name an interface created with the **interface** command (in context configuration mode), but it must be unique among all named IP pools within the context.

The falling-threshold parameters provide an alert when the number of available IP addresses in the pool is reduced to the value specified.

Use the **to** ***ip-addr*** construct to select a range of IP addresses for the IP pool.

Use the **falling-threshold** ***num*** construct to specify the number of available IP addresses in the pool for which a falling-threshold crossing event is generated. A crossing event occurs only when the number of available IP addresses in the pool equals the value specified. If the number of available IP addresses becomes greater than the value specified and then drops again to the value, a second falling-threshold crossing event is generated.

If you specify the **falling-threshold** ***num*** construct and the IP pool already exists, the current falling-threshold parameters are set to the new values, or are added to the definition of the IP pool if they did not previously exist. If you enter the **ip pool** command without the falling-threshold parameters and the IP pool already exists, the threshold is removed.

You can specify that the falling-threshold crossing event be reported with an SNMP trap, a log message, or both the trap and the log message.

For information about configuring context-specific falling-threshold parameters or including Class C network and broadcast IP addresses in IP pools in the context, see *Configuring Contexts and Interfaces* and for information about the **ip pool** command (in context configuration mode), see the *Command List*.

Use the **no** or **default** form of this command to delete the IP address pool for the specified starting IP address or all IP pools created in the interface.

1.36.6 Examples

The following example shows how to create a named IP pool for the interface **isp1.net** context and specifies that both an SNMP **trap** and a **log** message be generated when the number of available IP addresses in the pool equals **22:**

```
[local]Redback(config)#context isp1.net
[isp1.net]Redback(config-ctx)#interface isp1.net multibind
[isp1.net]Redback(config-if)#ip address 10.1.1.1 255.255.255.0
[isp1.net]Redback(config-if)#ip pool 10.1.1.1 255.255.255.0 name ip-pool1  falling-threshold 22 trap log
```

The following example shows how to create a named IP pool for the **isp1.net** context and specifies a range of IP addresses for the IP pool using the **to** ***ip-addr*** construct:

```
[local]Redback(config)#context isp1.net
[isp1.net]Redback(config-ctx)#interface isp1.net multibind
[isp1.net]Redback(config-if)#ip address 10.1.1.1/24
[isp1.net]Redback(config-if)#ip pool 10.1.1.2 to 10.1.1.100
```

1.37 ip prefix-list

ip prefix-list ***pl-name***

no ip prefix-list ***pl-name***

1.37.1 Purpose

Creates an IP prefix list used to filter routes and enters IP prefix list configuration mode.

1.37.2 Command Mode

Context configuration

1.37.3 Syntax Description

pl-name	IP prefix list name.

1.37.4 Default

There are no preconfigured IP prefix lists.

1.37.5 Usage Guidelines

Use the **ip prefix-list** command to create an IP prefix list used to filter routes and to enter IP prefix list configuration mode where you can define conditions using the **permit** and **deny** commands.

Note: A reference to an IP prefix list that does not exist, or does not contain any configured entries, implicitly matches and permits all IP prefixes.

Use the **no** form of this command to remove an IP prefix list.

1.37.6 Examples

The following example shows how to create the IP prefix list, **list102**, and enter IP prefix list configuration mode:

```
[local]Redback(config-ctx)#ip prefix-list list102
[local]Redback(config-prefix-list)#
```

1.38 ip profile

ip profile ***profile-name***

no ip profile ***profile-name***

1.38.1 Purpose

Attaches an RFlow profile to an external collector.

1.38.2 Command Mode

Flow collector configuration

1.38.3 Syntax Description

profile-name	Name of the RFlow profile that you want to attach to the external collector.

1.38.4 Default

None

1.38.5 Usage Guidelines

Use the **ip profile** command in flow collector configuration mode to attach an RFlow profile to an external collector. You can attach a maximum of 10 RFlow profiles to each external collector.

Use the **no** form of this command to remove an RFlow profile attachment from an external collector.

1.38.6 Examples

The following example shows how to attach an RFlow profile called **p1** to an external collector called **c1:**

```
[local]Redback#configure
[local]Redback(config)#context foo
[local]Redback(config-ctx)#flow collector c1
[local]Redback(config-flow-collector)#ip profile p1
```

1.39 ip route

ip route *ip-addr/prefix-length* {*next-hop-ip-addr* | *next-hop-if-name* | **null0** | **context** *ctx-name*} [**connected**] [**bfd**] [**dvsr** *dvsr-profile-name* [**verify-address** *verify-addr*]] [**cost** *cost*] [**distance** *distance*] [**permanent**] [**tag** *tag*] [**description** *text*]

no ip route *ip-addr/prefix-length* {*next-hop-ip-addr* | *next-hop-if-name* | **null0** | **context** *ctx-name*} [**connected**] [**bfd**] [**dvsr** *dvsr-profile-name* [**verify-address** *verify-addr*]] [**cost** *cost*] [**distance** *distance*] [**permanent**] [**tag** *tag*] [**description** *text*]

1.39.1 Purpose

Configures one or more static routes when the system is not configured to dynamically select a route to the destination.

1.39.2 Command Mode

Context configuration

1.39.3 Syntax Description

ip-addr/prefix-length	IP address (in the form ***A.B.C.D***) and prefix length, separated by the slash (**/**) character. The range of values for the ***prefix-length*** argument is 0 to 32.
next-hop-ip-addr	IP address of the next hop that can be used to reach the network.
next-hop-if-name	Interface name of the next hop that can be used to reach the network.
null0	Creates a null interface to prevent routing loops.
context ***ctx-name***	Another context, which can be used as a next hop to reach a network.
connected	Optional. Specifies that the IP next hop should be on the connected circuit subnet.
bfd	Optional. Enables Bidirectional Forwarding Detection (BFD) for the static route.
dvsr ***dvsr-profile-name***	Optional. Dynamically verified static routing (DVSR) profile name. Defines a DVSR using the specified profile name. The **dvsr** ***dvsr-profile-name*** construct cannot be used with the ***next-hop-ip-addr*** or ***next-hop-if-name*** arguments, or the **null0** or **permanent** keywords.
verify-address ***verify-addr***	Optional. Host IP address the DVSR route should verify. If the **verify-address** ***verify-addr*** construct is not configured, the ***next-hop-ip-addr*** or ***next-hop-if-name*** argument will be used for the verification.
cost ***cost***	Optional. Cost of the route. The range of values is 0 to 15.
distance ***distance***	Optional. Administrative distance assigned to the route. The range of values is 1 to 255.
permanent	Optional. Indicates that the route cannot be removed, even if the interface is shut down.
tag ***tag***	Optional. Route tag used as a match value for controlling redistribution through route maps. An unsigned 32-bit integer, the range of values is 1 to 4,294,967,295; the default value is 0.
description ***text***	Optional. Description for the static route.

1.39.4 Default

None

1.39.5 Usage Guidelines

Use the **ip route** command to configure one or more static routes when the system is not configured to dynamically select a route to the destination.

A static route can be overridden by a dynamically learned route with a lower administrative distance.

Use the **null0** keyword to prevent routing loops. A null interface is always up and can never forward or receive traffic. The null interface provides an alternative method of filtering traffic. You can avoid the overhead involved with access control lists by directing undesired network traffic to the null interface.

Note: The OSPF and Intermediate System-to-Intermediate System (IS-IS) routing processes always create a route to a null interface when summarizing a group of routes.

Use the **context** ***ctx-name*** construct to forward traffic to another routing context (next-hop context). The **context** ***ctx-name*** construct can be used to configure VPN customer Internet access or inter-VPN routing leaks. The next-hop context must be a different routing context than the one to which the static route belongs. If the next-hop context does not exist, and the **service multiple-contexts** command is enabled on the router, the context is created. Intercontext static routing between two non-local contexts is not allowed unless the **service inter-context routing** command is enabled on the router. The prefix using the next-hop context is considered to be valid only if the next-hop context has the routes that are being covered by this prefix. This prefix is installed in the Router Information Base (RIB) only if the next-hop context can reach those networks.

Intercontext routing enables the creation of routing sessions between peers that belong to different contexts that are not connected by a physical port, eliminating the need for a physical link between the contexts.

Peering between the routes can be used to advertise routes to the peer in the other context. The BGP peer in the local context acts as a route reflector for the peer in the non-local context and reflects Internal BGP (iBGP) routes to it. Route-maps can also be used to filter routes being sent and received between BGP peers with an intercontext BGP connection. The route-map functionality also works in intercontext routing scenarios.

Use the **bfd** keyword to enable BFD for a static route. BFD is a simple Hello protocol can detect communication failures in less than one second. When BFD detects a communication failure to the next hop specified for a static route (that has BFD enabled), the static route is withdrawn. By default, BFD is disabled for all static routes.

Use the **dvsr** ***dvsr-profile-name*** construct to configure a static route with DVSR capability. A DVSR route needs to reference an existing DVSR profile by name. Protocol redistribution can be specified through the **redistribute static dvsr** command to only import DVSR-capable routes. The verify-host address of the DVSR route is by default the next-hop IP address of the route. If

the DVSR verify-host is not the same as the next-hop IP address, make sure that there is a route to reach the verify-host address. In addition, the next hop of that route must be the same as the next hop of the DVSR route.

Use the **no** form of this command to remove static routes.

1.39.6 Examples

The following example routes packets for network **20.0.0.0/8** to the device at IP address **121.109.3.4** if dynamic information with administrative distance less than **110** is not available:

```
[local]Redback(config-ctx)#ip route 20.0.0.0/8 121.109.3.4 distance 110
```

The following example shows how to configure a null interface for network **172.0.0.0/8**:

```
[local]Redback(config-ctx)#ip route 172.0.0.0/8 null0
```

The following example shows how to route packets for network **129.108.0.0/16** to the device at IP address **129.108.6.6**:

```
[local]Redback(config-ctx)#ip route 129.108.0.0/16 129.108.6.6
```

The following example shows how to configure a static route from the **local** context using context, **vpn-abc**, as the next hop context:

```
[local]Redback(config-ctx)#ip route 12.1.1.0/24 context vpn-abc
```

The following examples show how to enable intercontext routing for the creation of routing sessions between peers that belong to different contexts that are not connected by a physical port, eliminating the requirement of an actual physical link between the contexts. The BGP peer in the local context acts as a route reflector (RR) for the peer in the non-local context and reflects iBGP routes to it

The following shows the configuration for local context A.

```
[local]Redback(config)#context local

[local]Redback(config-ctx)#no ip domain-lookup

[local]Redback(config-ctx)#router-id 8.8.8.8

 [local]Redback(config-ctx)#interface int_to_ctxB loopbackrouter-id 8.8.8.8

 [local]Redback(config-if)#ip address 1.2.3.4/32
!
  [local]Redback(config-ctx)#router bgp 1

 [local]Redback(config-bfd)# neighbor 4.3.2.1 internal

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

!
[local]Redback(config-ctx)#ip route 4.3.2.1/32 context ctxB
```

The following shows the configuration for non-local context B.

```
[local]Redback(config)#context ctxB

 [local]Redback(config-ctxb)#interface int_to_local1 loopback

 [local]Redback(config-if)#ip address 4.3.2.1/32

!
[local]Redback(config-ctxb)#router bgp 1
[local]Redback(config-bfd)# address-family ipv4 unicast

 !
 [local]Redback(config-bfd)# neighbor 1.2.3.4 internal

 [local]Redback(config-bfd)# address-family ipv4 unicast
!
 [local]Redback(config-ctxb)#route-reflector-client

[local]Redback(config-ctxb)#route-map rp1 in

!
[local]Redback(config-if)#ip route 1.2.3.4/32 context local
```

1.40 ip soft-gre

ip soft-gre [**source** *src-addr*]

no ip soft-gre [**source** *src-addr*]

1.40.1 Purpose

Enables soft-Generic Routing Encapsulation (GRE) tunneling on the specified context.

1.40.2 Command Mode

Context configuration

1.40.3 Syntax Description

source *src-addr*	Optional. Source address for the soft GRE tunnel. The IP address is in the form ***A.B.C.D.***

1.40.4 Default

Soft GRE tunneling is disabled.

1.40.5 Usage Guidelines

Use the **ip soft-gre** command to enable soft GRE tunneling on the specified context.

Encapsulating packets with GRE from an ingress provider edge (PE) router to an egress PE router is called soft GRE tunneling. Soft GRE tunnels are not Interior Gateway Protocol (IGP) visible links, and routing adjacencies are not supported across these tunnels. As a result, soft GRE tunnels have little in common with traditional (hard) GRE tunnels. The tunnel exists only in the sense of GRE encapsulation and decapsulation.

Only the ingress PE router and the egress PE router need to support the soft GRE functionality, and the PE routers can span over multiple autonomous systems.

Using soft GRE tunnels to transport Multiprotocol Label Switching (MPLS)-encapsulated packets is called Border Gateway Protocol/MPLS Virtual Private Network (BGP/MPLS VPN) over GRE, and is used to offer BGP/MPLS VPN service when a portion of a network does not have label switching enabled. BGP/MPLS VPN over GRE does not require preconfiguration of the remote GRE endpoint. These endpoints are the BGP next-hop addresses of the VPN routes, and are learned dynamically through BGP.

Using soft GRE tunnels to transport Layer 2 Virtual Private Network (L2VPN)-encapsulated packets is called L2VPN over GRE, and can be used instead of a Multiprotocol Label Switching (MPLS) tunnel in the backbone. L2VPN over GRE does not require preconfiguration of the remote GRE

endpoint. The GRE tunnel endpoint is the remote PE's address to which the L2VPN packets are being transported.

Use the **no** form of this command to disable soft GRE on the specified context.

1.40.6 Examples

The following example shows how to enable soft GRE in the **local** context:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#ip soft-gre
```

1.41 ip source-address

For NetOp EMS configurations, the syntax is as follows:

ip source-address [**netop**] {**all** | [***packet-type***] [***packet-type***] ... }

no ip source-address [**netop**] {**all** | [***packet-type***] [***packet-type***] ... }

For all other configurations, the syntax is as follows:

ip source-address [**all** | {[***packet-type***] [***packet-type***] ... }]

no ip source-address [**all** | {[***packet-type***] [***packet-type***] ... }]

1.41.1 Purpose

Specifies the primary IP address of this interface as the source address for one or more types of locally generated packets or packets sent to a Dynamic Host Configuration Protocol (DHCP) server. Additionally, allows the existing node discovery feature to refer to a management-configured interface for the source address rather than the IP address of the interface determined by routing.

1.41.2 Command Mode

Interface configuration

1.41.3 Syntax Description

all	Optional. Specifies the primary IP address of this interface as the source address for all types of packets listed in Table 1.
packet-type	Optional. Type of packets in which the primary IP address of this interface is used as the source address, according to one of the keywords listed in Table 1. You can list multiple packet types, each separated by a space.

1.41.4 Default

The IP address for the interface on which the traffic is transmitted is used as the source address in locally generated packets or packets sent to a DHCP relay server.

1.41.5 Usage Guidelines

Use the **ip source-address** command to specify the primary IP address of this interface as the source address for one or more types of locally generated packets or packets sent to a DHCP relay server. The primary IP address for the interface is assigned using the **ip address** command (in interface configuration mode).

Note: Enter this command with the IP source addresses of loopback interfaces and not with IP addresses of interfaces associated with physical ports or circuits. You should not specify the IP source address of a physical port or circuit because if the port or circuit goes down, the reply packets would be disrupted.

You can specify multiple keywords in any order with this command; you can also enter the command multiple times to specify additional protocols. Table 1 lists the keywords for the types of packets in which the IP address is sent.

Table 1 Keywords for Supported Protocols and Servers

Keyword	Packet Description
dhcp-server	Specifies packets to a DHCP relay server.
ftp	Specifies File Transfer Protocol (FTP) packets.
icmp-dest-unreachable	Specifies Internet Control Message Protocol (ICMP) type 3, Destination Unreachable, packets.
icmp-time-exceeded	Specifies that all replies to ICMP type 11 packets are sourced with the defined IP address.

Table 1 Keywords for Supported Protocols and Servers

Keyword	Packet Description
netop	Specifies advertisement packets which the SmartEdge router sends as part of the automatic node discovery process with the NetOp EMS server. Allows the NetOp EMS server to reach the SmartEdge router through the IP source address set by this command and bound to traffic cards as opposed to the default management IP address of the controller card.
radius	Specifies packets to a RADIUS server.
snmp	Specifies Simple Network Management Protocol (SNMP) packets.
ssh	Specifies Secure Shell (SSH) and Secure FTP (SFTP) packets.
syslog	Specifies syslog packets.
tacacs+	Specifies Terminal Access Controller Access Control System Plus (TACACS+) packets.
telnet	Specifies Telnet packets.
tftp	Specifies Trivial FTP (TFTP) packets.

Use the **all** keyword to specify all supported protocols and servers.

By default, the local IP address for the interface on which the traffic is transmitted is included in transmitted packets. As a result, the local IP address used for packets can change from connection to connection, based on the interface that the routing algorithm has chosen to reach the destination.

For IP packets sent by IP routing protocols, including Open Shortest Path First (OSPF), Routing Information Protocol (RIP), Resource Reservation Protocol (RSVP), and the multicast protocols, but not including Intermediate System-to-Intermediate System (IS-IS), the local IP address selection is often constrained by the protocol specification so that the protocol operates correctly. When this constraint exists in the routing protocol, the IP source address included in the outgoing packet is determined by the routing protocol and not the **ip source-address** command.

Note: For the RADIUS application, use the **radius attribute nas-ip-address** command (in context configuration mode) to configure the SmartEdge router to send the IP source address in access request and accounting request packets to the RADIUS server. For more information, see *Configuring RADIUS*.

Use the **no** form of this command to use the local IP address for the interface on which the traffic is transmitted.

1.41.6 Examples

The following example specifies the IP address of the **notify** interface in the **local** context for all outgoing Telnet packets:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#interface notify
[local]Redback(config-if)#ip address 172.16.1.1/24
[local]Redback(config-if)#ip source-address telnet
```

The following example shows how to add the SNMP to the list of protocols using the IP address for the **notify** interface:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#interface notify
[local]Redback(config-if)#ip source-address snmp
```

As a result, both the Telnet and SNMP protocols use the IP address of the **notify** interface.

The following example specifies that ICMP packets will also use the IP address of the **notify** interface:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#interface notify
[local]Redback(config-if)#ip source-address icmp-dest-unreachable
```

1.42 ip source-address flow-ip

ip source-address flow-ip {***packet-type*** [***packet-type***]...}

no ip source-address flow-ip {***packet-type*** [***packet-type***]...}

1.42.1 Purpose

Configures an IP address to be the source of IP packets that are exported to an external collector.

1.42.2 Command Mode

Interface configuration

1.42.3 Syntax Description

packet-type	Type of packets in which the primary IP address of this interface is used as the source address, according to one of the keywords listed in Table 2. You can list multiple packet types, each separated by a space.

1.42.4 Default

None

1.42.5 Usage Guidelines

Use the **ip source-address flow-ip** command to configure an IP address to be the source of IP packets that are exported to an external collector.

You can specify multiple protocols in any order; you can also enter the command multiple times to specify additional protocols. Table 2 lists the keywords for the types of packets in which the IP address is sent.

Table 2 Keywords for Supported Protocols and Servers

Keyword	Packet Description
dhcp-server	Specifies packets to a DHCP relay server.
ftp	Specifies File Transfer Protocol (FTP) packets.
icmp-dest-unreachable	Specifies Internet Control Message Protocol (ICMP) type 3, Destination Unreachable, packets.
icmp-time-exceed	Specifies that all replies to ICMP type 11 packets are sourced with the defined IP address.
netop	Specifies advertisement packets that the SmartEdge router sends as part of the automatic node discovery process with the NetOp EMS server. The NetOp EMS server can reach the SmartEdge router through the IP source address set by this command and bound to traffic cards, as opposed to the default management IP address of the controller card.
radius	Specifies packets to a RADIUS server.
snmp	Specifies Simple Network Management Protocol (SNMP) packets.
ssh	Specifies Secure Shell (SSH) and Secure FTP (SFTP) packets.

Table 2 Keywords for Supported Protocols and Servers

Keyword	Packet Description
syslog	Specifies syslog packets.
tacacs+	Specifies Terminal Access Controller Access Control System Plus (TACACS+) packets.
telnet	Specifies Telnet packets.
tftp	Specifies Trivial FTP (TFTP) packets.

Use the **no** form of this command to remove an interface as a source for sending packets to the external collector.

Note: For more information about using the **ip source-address flow-ip** command, see the *Command List*.

1.42.6 Examples

The following example shows how to configure the **rflow2** interface to send packets to the external collector:

```
[local]Redback#configure
[local]Redback(config)#context local
[local]Redback(config-ctx)#interface rflow2
[local]Redback(config-if)# ip source-address flow-ip ftp telnet
```

1.43 ip source-validation

ip source-validation

no ip source-validation

1.43.1 Purpose

Enables IP source-address validation (SAV), which denies all IP packets from address sources that are not reachable through a subscriber's associated circuit.

1.43.2 Command Mode

Subscriber configuration

1.43.3 Syntax Description

This command has no keywords or arguments.

1.43.4 Default

IP SAV is disabled.

1.43.5 Usage Guidelines

Note: Generally, reverse path forwarding (RPF), IP SAV, and ingress filtering all refer to the same functionality

Use the **ip source-validation** command to enable IP SAV. IP SAV denies all IP packets from address sources that are not reachable through the subscriber's associated circuit. You can use this command to prevent source address spoofing.

Use the **no** form of this command to disable IP SAV.

1.43.6 Examples

The following example shows how to enable IP SAV for the subscriber, **bart:**

```
[local]Redback(config-ctx)#subscriber name bart
[local]Redback(config-sub)#ip source-validation
```

1.44 ip static in

For source IP address translation:

ip static in source *ip-addr nat-addr* [**context** *ctx-name*]

no ip static in source *ip-addr nat-addr* [**context** *ctx-name*]

For source IP address and TCP or UDP translation:

ip static in {**tcp** | **udp**} **source** *ip-addr port nat-addr nat-port* [**context** *ctx-name*]

no ip static in {**tcp** | **udp**} **source** *ip-addr port nat-addr nat-port* [**context** *ctx-name*]

1.44.1 Purpose

Translates the source IP address in the private network, and optionally, Transmission Control Protocol/User Datagram Protocol (TCP/UDP) ports, of incoming packets on the interface to which the Network Address Translation (NAT) policy is attached. In the reverse direction, translates the destination IP address, and optionally, TCP/UDP ports, of outgoing packets on the interface.

1.44.2 Command Mode

NAT policy configuration

1.44.3 Syntax Description

source	Indicates the source information.
tcp	Indicates a TCP port.
udp	Indicates a UDP port.
ip-addr	Original source IP address.
port	Original TCP or UDP source port number. The range of values is 1 to 65,535. Required when using the **tcp** or **udp** keyword.
nat-addr	NAT address. The IP address to which the source IP address is mapped in the address translation table.
nat-port	TCP or UDP port number to which the source port number is mapped in the address translation table. The range of values is 1 to 65,535. Required when using the **tcp** or **udp** keyword.
context ***ctx-name***	Optional. Context name. Required for intercontext forwarding of packets. Interfaces in the specified context are used to forward packets after addresses are translated.

1.44.4 Default

If no action is configured for the NAT policy, by default, packets are dropped.

1.44.5 Usage Guidelines

Use the **ip static in** command to translate the source IP address in the private network, and optionally, TCP/UDP ports, of incoming packets on the interface to which the NAT policy is attached. In the reverse direction, this command translates the destination IP address, and optionally, TCP/UDP ports, of outgoing packets on the interface.

Incoming packets with a source IP address that matches the ***ip-addr*** argument use the IP address specified with the ***nat-addr*** argument as their source IP

address instead. In the opposite direction, outgoing packets with a destination IP address that matches the ***nat-addr*** argument use the ***ip-addr*** argument as the destination IP address.

If the ***nat-addr*** argument overlaps an IP address in a Network Access Port Translation (NAPT) pool, the static translation takes precedence.

Use the **no** form of this command to disable the translation of the source IP address and TCP/UDP ports.

1.44.6 Examples

The following example shows how to translate the source IP address of packets received on the interface, **customer1**, to **2.2.2.2** when the original source address of the packets is **1.1.1.1**. At the same time, the destination address of packets sent out the interface are translated to **1.1.1.1** when the original destination address of the packets is **2.2.2.2:**

```
[local]Redback(config-ctx)#nat policy p2

[local]Redback(config-policy-nat)#ip static in source 1.1.1.1 2.2.2.2

[local]Redback(config-policy-nat)#exit

[local]Redback(config-ctx)#interface customer1

[local]Redback(config-if)#ip address 1.1.1.254/24

[local]Redback(config-if)#ip nat p2
```

1.45 ip static out

ip static out source ***ip-addr nat-addr***

no ip static outsource ***ip-addr nat-addr***

1.45.1 Purpose

Translates the source IP address in the private network of outgoing packets on the interface to which the Network Address Translation (NAT) policy is applied, and in the reverse direction, translates the destination IP address of incoming packets on the interface.

1.45.2 Command Mode

NAT policy configuration

1.45.3 Syntax Description

source	Indicates the source information.
ip-addr	Original source IP address.
nat-addr	NAT address. The IP address to which the source IP address is mapped in the address translation table.

1.45.4 Default

If no action is configured for the NAT policy, packets are dropped.

1.45.5 Usage Guidelines

Use the **ip static out** command to translate the source IP address in the private network of outgoing packets on the interface to which the NAT policy is applied, and in the reverse direction, to translate the destination IP address of incoming packets on the interface.

Outgoing packets with a source IP address that match the ***ip-addr*** argument use the IP address specified with the ***nat-addr*** argument as their source IP address instead. In the opposite direction, incoming packets with a destination IP address that matches the ***nat-addr*** argument use the ***ip-addr*** argument as the destination IP address.

Use the **no** form of this command to disable the translation of the IP address.

1.45.6 Examples

The following example shows how to translate the IP source address of packets sent out the interface, **pos1**, to **10.30.40.50** when the original source address of the packets is **64.64.64.64**. At the same time, the destination address of packets coming into the interface are translated to **64.64.64.64** when the destination address of the packets is **10.30.40.50:**

```
[local]Redback(config-ctx)#nat policy p1
[local]Redback(config-policy-nat)#ip static out source 64.64.64.64 10.30.40.50
[local]Redback(config-policy-nat)#exit
[local]Redback(config-ctx)#interface pos1
[local]Redback(config-if)#ip nat p1
```

1.46 ip subscriber arp

ip subscriber arp ***ip-addr mac-addr***

no ip subscriber arp ***ip-addr***

1.46.1 Purpose

Creates an entry in the Address Resolution Protocol (ARP) cache for a subscriber whose host cannot (or is not configured to) respond to ARP requests.

1.46.2 Command Mode

Subscriber configuration

1.46.3 Syntax Description

ip-addr	IP address of the subscriber's host.
mac-addr	Medium access control (MAC) address of the subscriber's host.

1.46.4 Default

None

1.46.5 Usage Guidelines

Use the **`ip subscriber arp`** command to create an entry in the ARP cache for a subscriber whose host cannot (or is not configured to) respond to ARP requests.

Note: This command is available only if you are configuring a named subscriber record and is only relevant for circuits with RFC 1483 bridged-encapsulation.

Note: If you enter both the **`ip subscriber arp`** and the **`ip arp`** commands (in subscriber and context configuration modes, respectively), and specify the same IP address and MAC address, the most recently updated command takes precedence. Only the circuit and interface are updated in the ARP table.

Use the **`no`** form of this command to remove the specified entry.

1.46.6 Examples

The following example shows how to configure an ARP cache entry for a host with IP address, **10.1.1.1**, and hardware address, **d3:9f:23:46:77:13**, for the **NoGrokARPs** subscriber. The entry is installed into the ARP cache of the appropriate interface when the circuit is brought up:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#subscriber name NoGrokARPs
[local]Redback(config-sub)#ip address 10.1.1.1
[local]Redback(config-sub)#ip subscriber arp 10.1.1.1 d3:9f:23:46:77:13
```

1.47 ip subscriber route

ip subscriber route ***ip-addr*** {***netmask*** | ***/prefix-length***} [***next-hop-ip-addr***]

no ip subscriber route ***ip-addr*** {***netmask*** | ***/prefix-length***} [***next-hop-ip-addr***]

1.47.1 Purpose

Assigns one or more static IP routes to a subscriber's configuration.

1.47.2 Command Mode

Subscriber configuration

1.47.3 Syntax Description

ip-addr	IP address of the target network or subnet.
netmask	Network mask where the 1 bits indicates the network, or subnet, and the 0 bits indicate the host portion of the network address provided.
prefix-length	Prefix length. The range of values is 0 to 32. Optional when specified in conjunction with the ***next-hop-ip-addr*** argument.
next-hop-ip-addr	Optional. Required with RFC 1483 bridged-encapsulated circuits, and optional with other encapsulation types. IP address of a next-hop router that can reach the target network or subnet.

1.47.4 Default

None

1.47.5 Usage Guidelines

Use the **ip subscriber route** command to assign one or more static IP routes to a subscriber's configuration.

Note: This command is available only if you are configuring a named subscriber record.

To configure a default static IP route, use the ***netmask*** argument. If you use non-zero bits for the host portion of the network address, the route is not added to the routing table.

With RFC 1483 bridged encapsulation, a valid next-hop address and interface are required. If you are not using RFC 1483 bridged encapsulation, you can omit the next-hop address, but the route is not added to the routing table, unless the subscriber's circuit has one of the encapsulation types that does not require a next-hop address to be configured: Asynchronous Transfer Mode (ATM) Route1483, Layer 2 Tunneling Protocol (L2TP), Point-to-Point Protocol (PPP) over ATM (PPPoA), or PPP over Ethernet (PPPoE).

Use the **no** form of this command to delete a static route from the subscriber's configuration.

The routes for multiple protocols, including subscriber routes, have default routing distance values. When routing multiple routes with the same destination, the route with the lowest distance value is preferred.

Unlike the distance values for Border Gateway Protocol (BGP), Open Shortest Path First (OSPF), and Routing Information Protocol (RIP) routes, the distance values for directly connected, static IP, and subscriber routes cannot be modified. They always take the default distance values, as shown in Table 3.

Table 3 Protocol Default Distance Values

Protocol	Default Distance Value
Directly connected	0
Static IP	1
Subscriber IP host	15
Subscriber IP route	16

1.47.6 Examples

The following example assigns the IP route, **216.199.130.160 255.255.255.224**, to the subscriber, **SamQ:**

```
[local]Redback(config-ctx)#subscriber name SamQ
[local]Redback(config-sub)#ip address 10.1.2.3
[local]Redback(config-sub)#ip subscriber route 216.199.130.160 255.255.255.224
```

1.48 ip tcp mss

ip tcp mss replace [*dir*] *mss-size*

no ip tcp mss replace [*dir*] [*mss-size*]

1.48.1 Purpose

Changes the value of the maximum segment size (MSS) field in the TCP header to prevent fragmentation.

1.48.2 Command Mode

Interface configuration

1.48.3 Syntax Description

replace	Replace the value of the MSS field in the TCP header with the specified value.
dir	Optional. Identifies the direction of the traffic for which you are specifying a maximum segment size: • **in**—To specify an MSS for ingress traffic. • **out**—To specify an MSS for egress traffic. If you do not specify a direction, the MSS applies to both directions.
mss-size	Maximum segment size of a datagram in bytes. This value must be between 216 and 16,384 bytes and replaces the value of the MSS field in the TCP header.

1.48.4 Default

Packets for ingress and egress traffic pass unaltered through the SmartEdge router.

1.48.5 Usage Guidelines

Use the **ip tcp mss** command to replace the value of the MSS field in the TCP header to prevent fragmentation. Specify the maximum size of ingress and egress traffic in bytes.

The system does not replace MSS value in the datagram if the MSS value is bigger than the one found in the datagram. MSS replacement applies only to TCP SYN packets.

To set a different MSS for ingress traffic and egress traffic, enter the command twice—once for ingress traffic and once for egress traffic. To set the same MSS for both ingress and egress traffic, do not specify the direction. If you set an MSS for only one direction, no MSS is set for the other direction and the packets for that direction pass unaltered through the SmartEdge router.

Use the **no** form of this command to delete the current MSS configuration. Packets for ingress and egress traffic pass unaltered through the SmartEdge router.

1.48.6 Examples

The following example shows how to configure the **seattle-p2p** interface with an MSS of **1420** bytes for both ingress and egress traffic:

```
[local]Redback(config-ctx)#interface seattle-p2p
[local]Redback(config-if)#ip tcp mss replace 1420
```

1.49 ip to qos

ip {*dscp-value* | **all**} **to qos** *pd-value*

default qos {*dscp-value* | **all**}

1.49.1 Purpose

Translates Differentiated Services Code Point (DSCP) values into packet descriptor (PD) quality of service (QoS) values on ingress.

1.49.2 Command Mode

Class map configuration

1.49.3 Syntax Description

dscp-value	An integer from 0 to 63 representing the contents of the most significant 6 bits of the IP header type of service (ToS) field. You can enter the value in decimal or hexadecimal format, for example **16** or **0x10**. You can also enter a standard DSCP marking label as defined in the **`violate mark dscp`** command.
`all`	Maps all valid values for the source value to the specified target value. Any existing configuration for the classification map is overridden.
pd-value	An integer from 0 to 63 (6 bits), with the packet priority encoded in 3 higher-order bits and the packet drop precedence in the 3 lower-order bits. You can enter the value in decimal or hexadecimal format, for example **16** or **0x10**. You can also enter a standard DSCP marking label as defined in the **`violate mark dscp`** command. The scale used by this command for packet priority, from 0 (lowest priority) to 7 (highest priority), is the relative inverse of the scale used by the **`mark priority`** command. For details on this command, see the *Command List*.

1.49.4 Default

None

1.49.5 Usage Guidelines

Use the **`ip to qos`** command to define ingress mappings from IP header values to PD QoS values.

If you specify the **`all`** keyword, all valid IP header values are mapped to the specified QoS values. Any existing configuration for the classification map is overridden. You can use the **`all`** keyword to specify a single default value for all the mapping entries, then override that value for a subset of entries by entering subsequent mapping commands without this keyword.

Use the **`default`** form of this command to revert values for one or all map entries to their default values, where each DSCP value is mapped to the equal and equivalent PD QoS value.

1.49.6 Examples

The following example shows how to define the classification map **dscp-to-pd** for PD bits on ingress, then map all IP header values to the **af13** PD QoS value.

It overrides this default mapping for IP header DSCP values **af21** and **1**, which are mapped to PD QoS values **25** and **df** respectively:

```
[local]Redback(config)#qos class-map dscp-to-pd ip in
[local]Redback(config-class-map)#ip all to qos af13
[local]Redback(config-class-map)#ip af21 to qos 25
[local]Redback(config-class-map)#ip 1 to qos df
```

1.50 ip unnumbered

ip unnumbered *if-name*

{**no** | **default**} **ip unnumbered**

1.50.1 Purpose

Enables IP processing on an interface without assigning it an explicit IP address.

1.50.2 Command Mode

Interface configuration

1.50.3 Syntax Description

if-name	Name of the interface from which an IP address is to be borrowed.

1.50.4 Default

Interfaces do not borrow IP addresses.

1.50.5 Usage Guidelines

Use the **ip unnumbered** command to enable IP processing on an interface without assigning it an explicit IP address. This feature allows the interface to borrow the IP address of another interface.

Use the **no** or **default** form of this command to remove the ability to borrow IP addresses from another interface.

1.50.6 Examples

The following example shows how to configure the **seattle-p2p** interface to borrow an IP address from the **eth2** interface:

```
[local]Redback(config-ctx)#interface seattle-p2p
[local]Redback(config-if)#ip unnumbered eth2
```

1.51 ip verify unicast source

ip verify unicast source reachable-via {**any** | **rx**} [**allow-default**] [**access-group** ***acl-name*** [**acl-count**]]

no ip verify unicast source [**reachable-via**]

1.51.1 Purpose

Performs a reverse path forwarding (RPF) check to verify the source IP address on all incoming unicast packets on the specified interface.

1.51.2 Command Mode

Interface configuration

1.51.3 Syntax Description

reachable-via any	Specifies that the source IP address can be reached through any interface.
reachable-via rx	Specifies that the source IP address can be reached through the receiving interface.
allow-default	Optional. Allows the RPF check to look up the default route for verification. By default, an RPF check matching on the default route fails.
access-group ***acl-name***	Optional. Access control list (ACL) used to select the source IP addresses that are subject to an RPF check.
acl-count	Optional. Enables the counting of ACL entry matches.

1.51.4 Default

If **allow-default** is not configured, an RPF check matching on the default route fails.

1.51.5 Usage Guidelines

Note: Generally, RPF, IP source-address validation (SAV), and ingress filtering all refer to the same functionality.

Use the **ip verify unicast source** command to perform an RPF check to verify the source IP address on all incoming unicast packets on the specified interface.

If the packet passes the RPF check, the packet is forwarded as normal; however, if the router does not find a reverse path for the packet, the packet is dropped.

The unicast RPF check is a network security feature designed to address RFC 2827, `Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing`. That is, the Unicast RPF check feature addresses problems that are caused by the introduction of malicious IP packets with forged (spoofed) source IP addresses into a network by discarding IP packets that have no verifiable source IP address. Denial-of-Service (DoS) attacks use spoofed source IP addresses to give attackers the ability to circumvent efforts to locate or stop the attacks. Such attacks are eliminated by forwarding only packets that have source addresses that are valid and consistent with the IP routing table.

Note: Verifying the unicast source should be applied to an inbound interface at the upstream end of a connection.

Use the **no** form of this command to disable an RPF check to verify the source IP address on all incoming unicast packets at the specified interface.

1.51.6 Examples

The following example shows how to perform a unicast RPF check from interface **foo** on all unicast sources reachable by any interface:

```
[local]Redback(config-ctx)#interface foo
[local]Redback(config-if)#ip verify unicast source reachable-via any
```

1.52 ipip mtu

ipip mtu ***bytes***

no ipip mtu

1.52.1 Purpose

Sets the maximum transmission unit (MTU) for packets sent on IP-in-IP tunnels.

1.52.2 Command Mode

Dynamic Tunnel Profile configuration

1.52.3 Syntax Description

bytes	MTU size in bytes. The range of values is 256 through 1480 bytes.

1.52.4 Default

1480 bytes

1.52.5 Usage Guidelines

Use the **ipip mtu** command to set the MTU for packets for IP-in-IP tunnels. If an IP packet exceeds the MTU, the system fragments that packet.

A tunnel uses the MTU size for the interface to which the tunnel is bound to compute the tunnel MTU size, unless you explicitly configure the MTU using this command. After you configure an MTU for the tunnel, the system determines the effective MTU by comparing the configured MTU with the interface MTU and selecting the lesser of the two values.

Use the **no** form of this command to delete the configured MTU and use the interface MTU.

1.52.6 Examples

The following example shows how to set the maximum IP packet size for IP-in-IP tunnels for `prof1 to 1200 bytes:`

```
[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)#ipip mtu 1200
[local]Redback(config-mip-dyn-tun1-profile)#end
```

1.53 ipv6 access-group (interface)

ipv6 access-group ***acl-name*** **{in | out} [count]**

no ipv6 access-group ***acl-name*** **{in | out} [count]**

1.53.1 Command Mode

Interface configuration

1.53.2 Syntax Description

acl-name	One to ten names of the ACLs to be applied to the interface; separated by spaces. Each name string can be up to 39 characters. The aggregated string can be up to 255 characters.
in	Applies the ACL to incoming traffic.
out	Applies the ACL to outgoing traffic.
count	Counts the number of packets that pass through the filter. Dropped packets are not counted. This option may impact system performance and should be enabled only when diagnostic information is required.

1.53.3 Default

There are no ACLs applied to the interface.

1.53.4 Usage Guidelines

Use the **ipv6 access-group** command to apply an IPv6 filtering ACL to an interface bound to a port to filter control and administrative traffic. You can use it to control access to the Ethernet management port, or to ports on the traffic cards.

To control incoming administrative IPv6 traffic to the SmartEdge system, use the **ipv6 admin-access-group** command in each context.

Use the **no** form of the command to remove the ACL from the interface.

1.53.5 Examples

The following example shows an IPv6 ACL applied to the interface that is bound to the Ethernet management port, with counting enabled:

```
[local]Redback(config-ctx)#context local
[local]Redback(config-ctx)#interface mgt
[local]Redback(config-if)#ipv6 address 1:2:3:4:5/64
[local]Redback(config-if)#ipv6 access-group listmgt in count
```

You must have previously bound the management port to the interface, as follows:

```
...
[local]Redback(config)#port ethernet 7/1
! XCRP management ports on slot 7 and 8 are configured through 7/1
[local]Redback(config-port)#no shutdown
[local]Redback(config-port)#bind interface mgt local
```

1.54 ipv6 access-group (policy)

ipv6 access-group [*acl-name context-name*]

no ipv6 access-group [*acl-name context-name*]

1.54.1 Purpose

Applies an IPv6 access control list (ACL) to packets associated with the current forward, metering, or policing policy. That is, the specified ACL defines class matching criteria for the current policy.

1.54.2 Command Mode

- Forward policy configuration
- Policing policy configuration
- Metering policy configuration

1.54.3 Syntax Description

acl-name	Specifies the name of a IPv6 policy ACL created using the **ipv6 policy access-list** command (in context configuration mode). This parameter and the accompanying ***context-name*** parameter are optional only if the current policy is configured with the **radius-guided** option/keyword.
context-name	Specifies the name of the context in which the policy ACL was defined.

1.54.4 Default

None

1.54.5 Usage Guidelines

Use the **ipv6 access-group** (policy) command to apply a policy ACL to the current class-based policy (forward, QoS metering, or policing policy) and enter policy-group configuration mode.

The following restrictions apply:

- Do not exceed configuration of more than 12 ACLs for each RADIUS access-list.
- IPv6 ACL per rule logging of packets is not supported.

If the class-based policy was defined as RADIUS-guided, the policy ACL that it references can be dynamic, static, or both:

- A dynamic policy ACL is one that the SmartEdge router applies to the class-based policy for a particular subscriber session using the rules specified in an instance of vendor-specific attribute (VSA) 164 that the RADIUS server supplies in an access-response or COA message for the subscriber. For inbound forward policies only, a dynamic ACL may be alternatively specified using VSA 196 attribute **ipv6-fwd-in-access-group**. In these cases, the `ip6 access-group` command may be used without specifying the name of a statically configured policy ACL.
- A static policy ACL is a locally configured policy access-list whose name must be explicitly specified. If you include the *acl-name* argument, you must also include the *context-name* argument when you apply a static policy ACL to a forward policy or QoS policy.

  You can apply a dynamic policy ACL in addition to a static policy ACL. If VSA 164 was used to apply the dynamic ACL, the static policy ACL takes precedence over the dynamic policy ACL; a locally configured access-list's rules are evaluated before those from a dynamic ACL specified via VSA 164. If VSA 196 **ipv6-fwd-in-access-group**. is used to apply a dynamic ACL to a subscriber, the dynamic policy ACL supersedes and replaces the locally configured access-list, if applicable.

  If the class-based policy is not defined as RADIUS-guided, the policy ACL that it references must be static, and the **ip access-group** or **ipv6 access-group** command must specify the locally configured access-list's name and context.

Note: If a forward policy or quality of service (QoS) policy references a policy ACL that does not exist, the reference is ignored.

Warning!

The system does not warn you if you enter the `ipv6 access-group` command with the name of a static policy ACL that does not exist in the specified context; that is, the `ipv6 policy access-list` command has not configured a static policy ACL with the matching name in the specified context. If the ACL does not exist when the class-based policy is referenced by a static circuit (namely, a port, channel, PVC, or VLAN), the ACL reference is ignored and applicable traffic is treated as if it matched none of the classes specified in the policy. If the ACL does not exist when the class-based policy is referenced by a subscriber session circuit (namely, PPPoA, PPPoE, or CLIPs), the subscriber's applicable policy reference will be rejected, which may in turn cause the subscriber session to be terminated (as determined by the configuration of the `session-action failure` attribute for the subscriber).

1.54.6 Examples

The following example shows how to create an IPv6 policy ACL, `class_ipv6` in the `local` context and attaches it to the QoS policing policy, `POL1`. The ACL defines two classes, `B` and `C`, that are referenced by `POL1`:

Note: The command line `qos policy POL1 policing` creates a policy that is not RADIUS guided. When the current policy is not RADIUS guided, the *`acl-name context-name`* arguments are not optional.

```
[local]Redback#config
Enter configuration commands, one per line, 'end' to exit
[local]Redback#context local
[local]Redback(config-ctx)#ipv6 policy access-list class_ipv6
[local]Redback(config-ipv6-access-list)#seq 10 permit ipv6 any 2000:1:2:3::/64 class B
[local]Redback(config-ipv6-access-list)#seq 20 permit ipv6 any 2000:1:2:3::/64 class C
[local]Redback(config-ipv6-access-list)#exit
[local]Redback(config-ctx)#exit
[local]Redback(config)#qos policy POL1 policing
[local]Redback(config-policy-policing)#ipv6 access-group class_ipv6 local
[local]Redback(config-policy-group)#class B
[local]Redback(config-policy-group-class)#rate 50 burst 20000 counters
[local]Redback(config-policy-class-rate)#conform mark dscp af11
[local]Redback(config-policy-class-rate)#exit
[local]Redback(config-policy-group-class)#exit
[local]Redback(config-policy-group)#class C
[local]Redback(config-policy-group-class)#rate 200 burst 90000 counters
[local]Redback(config-policy-class-rate)#conform mark dscp ef
[local]Redback(config-policy-class-rate)#commit
```

1.55 ipv6 access-group (subscriber)

`[ no ]ipv6 access-group` *`acl-name`* **`{in | out}`**

1.55.1 Command Mode

subscriber configuration

1.55.2 Syntax Description

acl-name	One to ten names of the ACLs to be applied to the interface; separated by spaces. Each name string can be up to 39 characters. The aggregated string can be up to 255 characters. The ACLs must be previously created using the **ipv6 access-list** command in context configuration mode.
in	Applies the ACL to incoming traffic.
out	Applies the ACL to outgoing traffic.

1.55.3 Usage Guidelines

Use the **ipv6 access-group** command in subscriber configuration mode to apply an IPv6 ACL to a subscriber record, named profile, or default profile to filter subscriber traffic.

1.55.4 Example

The following example shows how to create an IPv6 filtering ACL, `acl2`, (in the `local` context) and applies it to the default subscriber profile:

```
[local]Redback(config-ctx)#ipv6 access-list acl2
[local]Redback(config-ipv6-access-list)#seq 10 permit ipv6 1:2:3:4::/48
[local]Redback(config-ipv6-access-list)#seq 20 permit ipv6 1:2:3:4::5/64
[local]Redback(config-ipv6-access-list)#exit
[local]Redback(config-ctx)#subscriber default
[local]Redback(config-sub)#ipv6 access-group acl2 in
```

1.56 ipv6 access-list

[no] ipv6 access-list ***list-name***

1.56.1 Command Mode

Context configuration

1.56.2 Syntax Description

list-name	Name of the ACL; an alphanumeric string up to 39 characters long.

1.56.3 Default

There are no ACLs on the context.

1.56.4 Usage Guidelines

Use the **ipv6 access-list** command to create, select, or remove an access control list (ACL) used to filter IPv6 traffic to the SmartEdge system, to subscribers (subscriber records, named profile or default profile), to the Ethernet management port, or to interfaces bound to ports on the chassis cards (also enters access control-list configuration mode).

To filter IPv6 traffic to the SmartEdge system, or to interfaces bound to ports, such as the Ethernet Management port, use the **ipv6 access-list** command to create an ACL, and then add `permit` and `deny` condition statements to the ACL. You can use the **seq permit** or **seq deny** command to explicitly order the sequence of the statements, or the **permit** or **deny** commands to have the SmartEdge OS order the statements automatically

For IPv6 ACLs, the recommended limit is 100 rules for each IPv6 ACL.

Note: In every ACL, there is an automatic **deny any any** statement at the end of the list. This automatic statement blocks all traffic not explicitly allowed, although it does not appear in the output of the *show configuration acl* command. It could block valid access to a context; for example, in the local context, it could block administrator access to the Ethernet management port. To allow administrator access, add a statement to explicitly allow access from authorized sources; for example, you could add a **permit ipv6 any any** or **permit ipv6** ***src src-wildcard dest dest-wildcard*** statement to the ACL.

Use the **resequence ip access-list** command to reorder a filtering ACL

Use the **no** form of the command to remove the ACL.

To protect control and administrative traffic, apply the ACL to each context configured on the SmartEdge router. To protect the Ethernet Management port or an interface bound to another port, apply the ACL to the interface bound to the port.

To apply an IPv6 ACL to a subscriber record, named profile, or default profile, enter the *ipv6 access-group (subscriber)* command in subscriber configuration mode.

1.56.5 Examples

The following example shows an IPv6 ACL created in the local context:

```
[local]Redback(config-ctx)#context local
[local]Redback(config-ctx)#ipv6 access-list listmgt
[local]Redback(config-ipv6-access-list)#seq 10 permit ipv6 1:2:3:4::5/48
[local]Redback(config-ipv6-access-list)#seq 20 permit ipv6 1:2:3:4::5/64
```

1.57 ipv6 address

ipv6 address ***ip-addr/prefix-length*** [**secondary**]

{**no** | **default**} **ipv6 address** ***ip-addr/prefix-length*** [**secondary**]

1.57.1 Purpose

Assigns a primary Internet Protocol Version 6 (IPv6) address, and optionally, one or more secondary IPv6 addresses, to an interface.

1.57.2 Command Mode

Interface configuration

1.57.3 Syntax Description

ip-addr	Primary or secondary IPv6 address of the interface.
prefix-length	Prefix length for the associated IPv6 address. The range of values is 0 to 128.
secondary	Optional. Configures the address as a secondary IPv6 address on the interface.

1.57.4 Default

No IPv6 address is assigned to an interface.

1.57.5 Usage Guidelines

Use the **ipv6 address** command to assign a primary IPv6 address, and optionally, one or more secondary IPv6 addresses, to an interface. This assignment enables IPv6 services on an interface.

Use the ***ip-addr*** argument and the ***/prefix-length*** construct to assign the interface a primary IPv6 address or prefix length. For nonloopback interfaces, use the **bind interface** command (in port configuration mode) to bind a circuit to the interface on which IP services are enabled. For more information on the **bind interface** command, see the *Command List*.

Note: The Neighbor Discovery (ND) protocol is enabled by default on broadcast-capable interfaces.

Use the optional **secondary** keyword to designate a IPv6 address as a secondary IPv6 address for the interface. You can configure up to 15 secondary addresses for each primary interface. Interface costs configured for routing protocols apply to secondary IP addresses in the same manner that they apply to primary IP addresses. Secondary IP addresses are treated as locally attached networks.

If Routing Information Protocol (RIP) split horizon is enabled on an interface that is configured with multiple IP addresses, a single update sourced by the primary IPv6 address is sent that advertises only the major networks. If split horizon is disabled, multiple updates sourced from each address on the interface are sent and all subnets are advertised.

When configuring an Open Shortest Path First (OSPF) interface, use the **ipv6 address** command first to establish the interface, and then enable OSPF version 3 (OSPFv3) on it by using the **interface** command in OSPFv3 area configuration mode; see *Configuring RADIUS*. The primary IPv6 address of the interface must belong to the area in which OSPFv3 is enabled. In addition, only neighbors on the primary address subnet can be OSPFv3 peers.

Caution!

Risk of IP service loss. Removing the primary IPv6 address disables all IP services for that address on the specified interface. Disabling IPv6 services deletes a corresponding OSPFv3 interface from the running configuration. To reduce the risk, do not remove a primary IPv6 address for an OSPFv3 interface, unless you have configured a secondary IPv6 address for the OSPFv3 interface, or intend to delete it.

Use the **bind interface** command (in IPv6 tunnel configuration mode) to statically bind a port, channel, permanent virtual circuits (PVCs), 802.1Q tunnel, link group, Generic Routing Encapsulation (GRE) tunnel circuit, or overlay tunnel circuit to a previously created interface in the specified context. No data can flow through a port, channel, PVC, 802.1Q tunnel, child circuit, link group, or tunnel circuit until it is bound to an interface. For more information on **bind interface** command, see the *Command List*.

Use the **no** or **default** form of this command to remove a IPv6 address from an interface. You must remove all secondary IPv6 addresses before you can remove the primary IPv6 address.

Note: When adding multiple interfaces to a context do not configure more than one interface IP address in a context on the same subnet.

1.57.6 Examples

The following example shows how to assign an IPv6 address to the **enet1** interface:

```
[local]Redback(config-ctx)#interface enet1
[local]Redback(config-if)#ipv6 address 7001::1/64
```

The following example shows how to configure two noncontiguous blocks for the **downstream** interface:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#interface downstream
[local]Redback(config-if)#ipv6 address 7002::1/112
[local]Redback(config-if)#ipv6 address 7003::1/112 secondary
```

The following example shows how to bind the Ethernet port **3/1** to the **downstream** interface using either IPv6 address:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#interface downstream
[local]Redback(config-if)#ipv6 address 7002::1/112
[local]Redback(config-if)#ipv6 address 7003::1/112 secondary
[local]Redback(config-if)#exit
[local]Redback(config-ctx)#exit
[local]Redback(config)#port ether 3/1
[local]Redback(config-port)#bind interface downstream local
```

1.58 ipv6 admin-access-group

ipv6 admin-access-group *acl-name* {**in** | **out**} [**count**]

no ipv6 admin-access-group *acl-name* {**in** | **out**} [**count**]

1.58.1 Purpose

Applies or removes an Access Control List (ACL) in a context to filter administrative traffic to the SmartEdge system.

1.58.2 Command Mode

Context configuration

1.58.3 Syntax Description

acl-name	One to ten names of the ACLs to be applied to the interface; separated by spaces. Each name string can be up to 39 characters. The aggregated string can be up to 255 characters.
in	Applied to incoming traffic.
out	Applied to outgoing traffic.
count	Enables the SmartEdge OS to count the number of packets that passed through the filter. Dropped packets are not counted. This option may impact system performance. It is recommended to enable it only when diagnostic information is required.

1.58.4 Default

There are no ACLs in the context.

1.58.5 Usage Guidelines

Use the **ipv6 admin-access-group** command to apply an IPv6 filtering ACL to a context. To filter incoming administrative IPv6 traffic to the SmartEdge system, apply the command in each context.

Use the **no** form of the command to remove the ACL from a context.

1.58.6 Examples

The following example shows how to apply the ACL **listmgt** to the local context, without counting enabled:

```
[local]Redback(config-ctx)#context local
[local]Redback(config-ctx)#ipv6 admin-access
-group listmgt in
```

1.59 ipv6 delegated-prefix (DHCPv6 PD Prefix)

ipv6 delegated-prefix *ipv6-prefix*

no ipv6 delegated-prefix *ipv6-prefix*

1.59.1 Purpose

In a subscriber record, specifies the delegated IPv6 prefix to use for Dynamic Host Configuration Protocol version 6 (DHCPv6) prefix delegation (PD).

1.59.2 Command Mode

Subscriber record configuration

1.59.3 Syntax Description

ipv6-prefix	Specifies an IPv6 prefix (in the form A:B:C:D:E:F:G:H) and prefix length, separated by the slash (/) character. This is the IPv6 prefix that will be delegated to the subscriber through DHCPv6 PD. Range of values is 1 to 128.

1.59.4 Default

No delegated IPv6 prefixes are specified in the subscriber record.

1.59.5 Usage Guidelines

Use the **ipv6 delegated-prefix** command to specify a delegated IPv6 prefix to use for DHCPv6 PD.

Do not configure a delegated IPv6 prefix that overlaps with any other interface prefix within the same context.

Note: The DHCPv6 delegated prefix attribute is configurable only in a subscriber record.

Use the **no** version of this command to remove a DHCPv6 PD IPv6 prefix from the subscriber record.

1.59.6 Examples

The following example shows how to configure the DHCPv6 PD prefix attribute in a subscriber record called **test**. In this example, the IPv6 prefix is **2001:db8:1::/48**:

```
[local]BRAS(context)#subscriber name test

[local]BRAS(config-sub)#ipv6 delegated-prefix 2001:db8:1::/48
```

1.60 ipv6 delegated-prefix maximum

ipv6 delegated-prefix maximum ***maximum-prefixes***

{**no**} **ipv6 delegated-prefix maximum** ***maximum-prefixes***

1.60.1 Purpose

Sets the maximum number of delegated IPv6 prefixes that the router expects the DHCPv6 PD server to assign to hosts associated with the circuit.

1.60.2 Command Mode

- Default subscriber profile configuration
- Subscriber record configuration
- Subscriber profile configuration

1.60.3 Syntax Description

maximum-prefixes	Maximum number of unique IP prefixes the router expects the external DHCPv6 server to assign to hosts associated with a given subscriber circuit. The range of values is 1 to 5.

1.60.4 Default

The ***maximum-prefixes*** value is set to **1**.

1.60.5 Usage Guidelines

Use the **ipv6 delegated-prefix maximum** command to set the maximum number of delegated IPv6 prefixes that the router expects the DHCPv6 PD server to assign to hosts associated with the circuit.

Use the **no** form of this command to return the ***maximum-prefixes*** value to **1** (the default value).

1.60.6 Examples

The following example shows to configure the subscriber, **dhcpv6-test**, to expect a total of 5 IPv6 prefixes that can be assigned at any time:

```
[local]Redback(config-ctx)#subscriber name dhcpv6-test
[local]Redback(config-sub)#ipv6 delegated-prefix maximum 5
```

1.61 ipv6 framed-pool

ipv6 framed-pool [*name*]

{**no**} **ipv6 framed-pool** [*name*]

1.61.1 Purpose

Dictates that the specified subscriber obtains its IPv6 prefixes from a shared IPv6 prefix pool configured under the same context as the subscriber.

1.61.2 Command Mode

- Default subscriber profile configuration
- Subscriber record configuration
- Subscriber profile configuration

1.61.3 Syntax Description

name	Optional. Name of the a shared IPv6 prefix pool from which the subscriber obtains IPv6 prefixes.

1.61.4 Default

No shared IPv6 prefix pool is configured.

1.61.5 Usage Guidelines

Use the **ipv6 framed-pool** command to dictate that the specified subscriber obtains its IPv6 prefixes from a shared IPv6 prefix pool configured under the same context as the subscriber.

Use the **no** form of this command to remove the configuration from a subscriber dictating that the subscriber obtains its IPv6 prefixes from a shared IPv6 prefix pool.

1.61.6 Examples

The following example shows how to configure the subscriber `foo` to obtain IPv6 prefixes from a shared IPv6 prefix pool that is configured under the interface the subscriber (`blue`) is bound to:

```
[local]Redback(config)#context blue

[local]Redback(config-ctx)#subscriber name foo

[local]Redback(config-sub)#ipv6 pool
```

1.62 ipv6 framed-prefix

ipv6 framed-prefix ***ipv6-prefix***

no ipv6 framed-prefix ***ipv6-prefix***

1.62.1 Purpose

In a subscriber record, specifies the IPv6 prefix that will be assigned to subscribers using ND or static assignment.

1.62.2 Command Mode

Subscriber record configuration

1.62.3 Syntax Description

ipv6-prefix	Specifies a framed IPv6 prefix (in the form A:B:C:D:E:F:G:H) and prefix length, separated by the slash (/) character. Range of values for the prefix-length argument is 0 to 128. The framed IPv6 prefix is a unique prefix that cannot be a part of the interface IPv6 address or assigned to any other subscriber.

1.62.4 Default

No framed IPv6 prefixes are specified in the subscriber record

1.62.5 Usage Guidelines

Use the **ipv6 framed-prefix** command to specify the IPv6 prefix that will be assigned to subscribers using ND or static assignment. The framed IPv6 prefix is configurable only in a subscriber record.

Note: When configuring the **ipv6 framed-prefix** command, be sure the framed IPv6 prefix you specify does not overlap with any other interface prefix (the framed IPv6 prefix cannot be a part of the interface IPv6 address or assigned to any other subscriber).

Use the **no** version of this command to remove a framed IPv6 prefix from the subscriber record.

1.62.6 Examples

The following example shows how to configure the IPv6 framed prefix attribute in a subscriber record called **test**. In this example, the IPv6 prefix that will be assigned to subscribers using ND or static assignment is **2001:db8:b:4f::/64**:

```
[local]BRAS(context)#subscriber name test

[local]BRAS(config-sub)#ipv6 framed-prefix 2001:db8:b:4f::/64
```

1.63 ipv6 framed-route

ipv6 framed-route *ipv6-prefix* [*next-hop*] *metric*

no ipv6 framed-route *ipv6-prefix*[*next-hop*] *metric*

1.63.1 Purpose

In a subscriber record, specifies a static IPv6 route that will be installed for the subscriber.

1.63.2 Command Mode

Subscriber record configuration mode

1.63.3 Syntax Description

ipv6-prefix	Specifies an IPv6 prefix (in the form A:B:C:D:E:F:G:H) and prefix length, separated by the slash (/) character. Range of values for the prefix-length argument is 0 to 128.
next-hop	IPv6 address of a next-hop router that can reach the target network or subnet.
metric	Specifies the weight assigned for the route. Range is from 1 through 65, 536.

1.63.4 Default

No IPv6 routes are specified in the subscriber record.

1.63.5 Usage Guidelines

Use the **ipv6 framed-route** command to specify a static IPv6 route that will be installed for the subscriber.

Use the **no** version of this command to remove a static IPv6 route configuration from a subscriber record.

1.63.6 Examples

The following example shows how to configure the static (framed) IPv6 route attribute in a subscriber record called **test**. In this example, the IPv6 address for the static route is 2010:db8:b:5f::1/48, the next hop is on the IPv6 address 2002:db8:b:5f::1, and the metric is set to 1000.

```
[local]Redback(config-ctx)#subscriber name test
[local]Redback(config-sub)#ipv6 framed-route 2010:db8:b:5f::1/48 2002:db8:b:5f::1 1000
```

1.64 ipv6 host

ipv6 host ***hostname ipv6-addr***

no ipv6 host ***hostname ipv6-addr***

1.64.1 Purpose

Statically maps a hostname to an IPv6 address in the host table for Domain Name System (DNS) lookup.

1.64.2 Command Mode

Context configuration

1.64.3 Syntax Description

hostname	Name of the host.
ipv6-addr	IPv6 address of the host.

1.64.4 Default

No static mappings are preconfigured.

1.64.5 Usage Guidelines

Use the **ipv6 host** command to statically map a hostname to an IPv6 address in the host table for Domain Name System (DNS) lookup.

You can create up to 64 static entries in the host table. The SmartEdge router always consults the host table prior to generating a DNS lookup query.

Use the **no** form of this command to remove the specified static entry in the host table. Specifying a new IPv6 address for an existing hostname removes the previously specified IPv6 address.

1.64.6 Examples

The following example shows how to statically map the hostname, **hamachi**, to the IPv6 address, **2007::1:**

```
[local]Redback(config-ctx)#ipv6 host hamachi 2007::1
```

1.65 ipv6 link-local

ipv6 link-local ***ipv6-addr***

{no | default} ipv6 link-local ***ipv6-addr***

1.65.1 Purpose

Configures an IPv6 link local address.

1.65.2 Command Mode

Interface configuration

1.65.3 Syntax Description

ipv6-addr	IPv6 link-local address (in the form ***A:B:C:D:E:F:G:H.***

1.65.4 Default

The link-local address is not configured.

1.65.5 Usage Guidelines

Use the **ipv6 link-local** command to configure an IPv6 link-local address.

Use the **no** or **default** form of this command to remove an IPv6 link-local address configuration.

1.65.6 Examples

The following example shows how to configure an IPv6 link local address:

[local]Redback#**configure**

[local]Redback(config)#**context local**

[local]Redback(config-ctx)#**interface eth1**

[local]Redback(config-if)#**ipv6 link-local 1:1:1:1::100**

1.66 ipv6 maximum-routes

ipv6 maximum-routes[**multicast**] [**vpn**] ***route-limit*** [**log-only** | **threshold** ***value***] [**mid-threshold** ***value***]

1.66.1 Purpose

Configures an upper limit for the number of routes installed in an IPv6 routing table.

1.66.2 Command Mode

Context configuration

1.66.3 Syntax Description

multicast	Optional. Sets the maximum route limit for unicast routes in a multicast topology.
vpn	Optional. Sets the maximum route limit for all non-local context unicast routing tables. When the **vpn** keyword is used in the local context, it specifies a default maximum route setting that automatically applies to all non-local contexts; however, if the **ipv6 maximum-route** command is used in a specific non-local context, then it overrides the default maximum route setting.
route-limit	Maximum number of routes allowed in the IPv6 routing table. If this limit is reached and a protocol instance attempts to add additional routes, the routes are rejected, a warning is triggered, and the protocol instance is shut down. Range of values is 1 to 4294967295.
log-only	Optional. Configures the route limit as an advisory limit. An advisory limit triggers only a warning, and additional routes are not rejected.
threshold*value*	Optional. Threshold value for the mandatory limit that triggers a warning. Range of values is 1 to 100.
mid-threshold ***value***	Optional. Threshold value for the mid-level limit that triggers a warning. Range of values is 1 to 100.

1.66.4 Default

No maximum limit is set.

1.66.5 Usage Guidelines

Use the **ipv6 maximum-routes** command to configure an upper limit for the number of routes installed in an IPv6 routing table.

A route limit sets an upper limit for the number of prefixes installed in a routing table; for example, you can use a route limit to limit the number of routes received from the customer edge (CE) router in a Virtual Private Network (VPN) context.

There are two modes for route limits: advisory (log only) and mandatory. An advisory limit only triggers warnings; a mandatory limit triggers warnings, rejects any additional routes after the maximum is reached, and shuts down the offending protocol instance (the instance that exceeds the limit). When the maximum is exceeded, you can clear the condition either by reconfiguring

the route limit or by using the **clear ipv6 route** command with the **maximum-routes** keyword. Clearing the condition also re-enables any routing protocol instances that were shut down when the maximum route limit was exceeded.

Use the **vpn** keyword in the local context to specify a default maximum route setting that automatically applies to all non-local contexts. To override the default maximum route setting, use the **ipv6 maximum-route** command in the non-local context that you want to configure.

1.66.6 Examples

The following example shows how to configure maximum routes on each VPN context that does not already have maximum routes configured. A warning is logged when a maximum of 400 routes reach the mandatory limit of 80 and a mid-threshold of 50:

```
[local]Redback#config context ipv6
[local]Redback(config-ctx)#ipv6 maximum-routes vpn 400 threshold 80 mid-threshold 50
```

The following example provides an aggregated configuration for IPv4 and IPv6 maximum-routes in non-local context resulting in the following aggregated values:

- Aggregated maximum-routes: 2000
- Aggregated threshold: 1700
- Aggregated mid-threshold: 1100

```
[local]Redback(config-ctx)#ip maximum-routes 1000 threshold 80 mid-threshold 50
[local]Redback(config-ctx)#ipv6 maximum-routes 1000 threshold 90 mid-threshold 6
```

The following example provides an aggregated configuration for IPv4 and IPv6 maximum-routes in non-local context without IPv6 mid-threshold, resulting in the following aggregated values:

- Aggregated maximum-routes: 2000
- Aggregated threshold: 1700
- Aggregated mid-threshold: 0

Note: Since mid-threshold is not configured for one of the address families, the mid-threshold equals 0 (zero) for the aggregate.

```
[local]Redback(config-ctx)#ip maximum-routes 1000 threshold 80 mid-threshold 50
[local]Redback(config-ctx)#ipv6 maximum-routes 1000 threshold 90
```

The following example provides an aggregated configuration for IPv4 and IPv6 maximum-routes without IPv6 threshold and mid-threshold values set which results in the following aggregated values:

- Aggregated maximum-routes: 2000
- Aggregated threshold: 1800
- Aggregated mid-threshold: 0

```
[local]Redback(config-ctx)#ip maximum-routes 1000 threshold 80 mid-threshold 50
[local]Redback(config-ctx)#ipv6 maximum-routes 1000
```

1.67 ipv6 mtu

ipv6 mtu ***mtu-length***

{**no** | **default**} **ipv6 mtu**

1.67.1 Purpose

Configures the maximum transmission unit (MTU) datagram size for traffic sent on the circuit to which the interface is bound.

1.67.2 Command Mode

Interface configuration

1.67.3 Syntax Description

mtu-length	Datagram size in bytes allowed for the interface. Range is from 1280 to 16,384.

1.67.4 Default

None

1.67.5 Usage Guidelines

Use the **ipv6 mtu** command to configure the MTU datagram size for traffic sent on the circuit to which the interface is bound. If an IPv6 packet exceeds the MTU configured for an interface, the system fragments that packet.

Note: This command does not apply to loopback interfaces.

Use the **no** or **default** form of this command to set the IPv6 MTU datagram size for this interface back to the default value of 1280.

1.67.6 Examples

The following example shows how to set the maximum MTU datagram size for the interface `to-blue` to `1300` bytes:

```
[local]Redback(config-ctx)#interface to-blue
[local]Redback(config-if)#ipv6 mtu 1300
```

1.68 ipv6 name-servers

ipv6 name-servers ***primary-ipv6-addr*** [***secondary-ipv6-addr***]

no ipv6 name-servers

1.68.1 Purpose

Specifies the IP Version 6 (IPv6) address of a primary (and, optionally, a secondary) Domain Name System (DNS) server.

1.68.2 Command Mode

Context configuration

1.68.3 Syntax Description

primary-ipv6-addr	IPv6 address of the primary DNS server.
secondary-ipv6-addr	Optional. IPv6 address of the secondary DNS server.

1.68.4 Default

No DNS server IPv6 addresses are preconfigured.

1.68.5 Usage Guidelines

Use the **ipv6 name-servers** command to specify the IPv6 address of a primary (and, optionally, a secondary) DNS server.

For DNS resolution to function, you must configure the domain name lookup using the **ip domain-lookup** command (in context configuration mode), and there must be an IPv6 route to the DNS servers.

Use the **no** form of this command to remove the specified DNS server association. If you delete the primary DNS server, any configured secondary DNS server becomes the primary server.

1.68.6 Examples

The following command shows how to configure an association with a primary DNS server at IPv6 address, **2007::1**, and a secondary server at IPv6 address, **2007::2:**

```
[local]Redback(config-ctx)#ipv6 name-servers 2007::1 2007::
```

The following command removes the primary DNS server, making the server that was previously the secondary into the primary:

```
[local]Redback(config-ctx)#no ipv6 name-servers 2007::1
```

1.69 ipv6 nd-profile

ipv6 nd-profile *name*

no ipv6 nd-profile

1.69.1 Purpose

References an ND profile in a subscriber record or profile.

1.69.2 Command Mode

- Default subscriber profile configuration

- Subscriber record configuration
- Subscriber profile configuration

1.69.3 Syntax Description

name	Name of the ND profile you want to reference.

1.69.4 Default

Subscribers inherit ND attributes from the global default ND profile.

1.69.5 Usage Guidelines

Use the **ipv6 nd-profile** command to reference an ND profile in a subscriber record or profile. Subscribers associated with the record or profile inherit the ND attributes from the specified ND profile.

If you do not reference an ND profile in a subscriber profile or record, the SmartEdge router automatically assigns a default ND profile (called the GLOBAL_DEFAULT_PROFILE) to the subscriber circuit.

Use the **show nd profile** command to see a list of all available ND profiles available in the current context; use the **show nd profile GLOBAL_DEFAULT_PROFILE** command to see the default configuration used by the GLOBAL_DEFAULT_PROFILE.

Use the **no** version of this command to remove a referenced ND profile from a subscriber record or profile.

1.69.6 Examples

The following example shows how to apply an ND profile called **abc** to the subscriber record for a subscriber called **test**:

```
[local]BRAS(context)#subscriber name test
[local]BRAS(config-sub)#ipv6 nd-profile abc
```

1.70 ipv6 path-mtu-discovery discovery-interval

ipv6 path-mtu-discovery discovery-interval ***seconds***

{**no** | **default**} **ipv6 path-mtu-discovery discovery-interval**

1.70.1 Purpose

Globally configures IPv6 path maximum transmission unit (PMTU) negotiation on a router and the timeout value used for aging PMTUs.

1.70.2 Command Mode

Global configuration

1.70.3 Syntax Description

seconds	Number of seconds that must pass before a PMTU value is aged out and the PMTU discovery process starts over. The range of values is 1 to 2000000000 seconds. The default is 600 seconds.

1.70.4 Default

600 seconds (10 minutes).

1.70.5 Usage Guidelines

Use the **`ipv6 path-mtu-discovery discovery-interval`** command to globally configure IPv6 PMTU negotiation on a router and the timeout value used for aging PMTUs.

Note: This command is available only on systems that have an XCRP4 card installed.

When PMTU discovery is configured, the router compares the IPv6 minimum link MTU to the PMTU and uses the largest MTU value to determine the maximum size of the packets allowed on a path. After the maximum MTU is discovered, source router fragments any IPv6 packet that exceeds the MTU of the receiving node into multiple smaller packets. When the PMTU discover timer expires, the current PMTU values are purged and the IPv6 PMTU negotiation process starts over, taking into account any network changes that increased size of the maximum MTU allowed. This process conserves bandwidth by ensuring packets are transmitted at the optimum MTU.

When PMTU negotiation is disabled, a router uses the IPv6 minimum link MTU value to determine the maximum size of packets allowed on a path. This often wastes bandwidth because paths typically have a PMTU that is greater than the IPv6 minimum link, which means the packets being transferred are smaller than necessary. (The PMTU is equal to the lowest link MTU value of all links in a path.)

Note: By default, the IPv6 MTU is set to 1280 for the interface. You can use the **ipv6 mtu** command in interface configuration mode to modify the IPv6 MTU for an interface.

Use the **no** or **default** form of this command to globally disable IPv6 PMTU discovery on a router.

Warning!

Risk of data loss for IPv6 data packets created by a traffic card. IPv6 data packets created on a traffic card are fragmented based on the MTU set on the egress port. In the current release of the SmartEdge OS any ICMPv6 “Packet too big” message sent from any IPv6 router on the data path to the traffic card that created an IPv6 data packet if the fragment size exceeds the PMTU of the data path is ignored. As a result, the data packet traffic is dropped. (ICMPv6 "Packet too big" messages sent from any traffic card or IPv6 router on the data path to the control or ASE card that created an IPv6 control packet if the fragment size exceeds the PMTU are acted upon.)

To avoid this risk, ensure that the MTU set on the traffic card with the **mtu** command in port configuration mode is set to be less than the PMTU of the data path. If a network-wide PMTU policy is used, set matching port-level MTU and network-wide PMTU values.

1.70.6 Examples

The following example shows how to globally enable IPv6 PMTU negotiation on a router and configure the timeout value for aging path MTUs to be 2500 seconds:

```
[local]Redback(config)#ipv6 path-mtu-discovery discovery-interval 2500
```

1.71 ipv6 policy access-list

[no] ipv6 policy access-list ***acl-name***

1.71.1 Command Mode

context configuration

1.71.2 Syntax Description

acl-name	Policy ACL name.

1.71.3 Default

None

1.71.4 Usage Guidelines

Use the **ipv6 policy access-list** command to create or access an IPv6 policy access control list (ACL) and enter access control list configuration mode.

Use the **no** form of this command to remove the policy ACL.

1.71.5 Examples

The following example shows how to create an IPv6 policy ACL, `class_ipv6`, in the `local` context and attaches it to the QoS policing policy, `POL1`. The ACL defines two classes: `B` and `C`, that are referenced by `POL1`:

```
[local]Redback#config
Enter configuration commands, one per line, 'end' to exit
[local]Redback#context local
[local]Redback(config-ctx)#ipv6 policy access-list class_ipv6
[local]Redback(config-ipv6-access-list)#seq 10 permit ipv6 any 2000:1:2:3::/64 class B
[local]Redback(config-ipv6-access-list)#seq 20 permit ipv6 any 2000:1:2:3::/64 class C
[local]Redback(config-ipv6-access-list)#exit
[local]Redback(config-ctx)#exit
[local]Redback(config)#qos policy POL1 policing
[local]Redback(config-policy-policing)#ipv6 access-group class_ipv6 local
[local]Redback(config-policy-group)#class B
[local]Redback(config-policy-group-class)#rate 50 burst 20000 counters
[local]Redback(config-policy-class-rate)#conform mark dscp af11
[local]Redback(config-policy-class-rate)#exit
[local]Redback(config-policy-group-class)#exit
[local]Redback(config-policy-group)#class C
[local]Redback(config-policy-group-class)#rate 200 burst 90000 counters
[local]Redback(config-policy-class-rate)#conform mark dscp ef
[local]Redback(config-policy-class-rate)#commit
```

1.72 ipv6 pool

In interface configuration mode:

ipv6 pool {[**dhcpv6**] *starting-prefix/prefix_length last-prefix/prefix_length*} [**name** *pool-name*] [**threshold** {**absolute** | **percentage**} **falling** *first-threshold* {**trap** [**log**] | **log** [**trap**]} [*second-threshold* {**trap** [**log**] | **log** [**trap**]}]

[**no** | **default**] **ipv6 pool** {[**dhcpv6**] *starting-prefix/prefix_length last-prefix/prefix_length* } [**name** *pool-name*] [**threshold** {**absolute** | **percentage**} **falling** *first-threshold* {**trap** [**log**] | **log** [**trap**]} [*second-threshold* {**trap** [**log**] | **log** [**trap**]}]

In context configuration mode:

ipv6 pool {[**dhcpv6**]**threshold** {**absolute** | **percentage**} **falling** ***first-threshold*** {**trap** [**log**] | **log** [**trap**]} [***second-threshold*** {**trap** [**log**] | **log** [**trap**]}]

[**no** | **default**] **ipv6 pool** {[**dhcpv6**] [**threshold** {**absolute** | **percentage**} **falling** ***first-threshold*** {**trap** [**log**] | **log** [**trap**]} [***second-threshold*** {**trap** [**log**] | **log** [**trap**]}]

1.72.1 Purpose

In interface configuration mode, creates a pool of IPv6 prefixes for a multibind interface to allow a subscriber to be assigned any available IP prefixes from the pool and, optionally, creates a threshold value for which a crossing event occurs.

In context configuration mode, creates a threshold value for which a crossing event occurs.

1.72.2 Command Mode

- Context configuration
- Interface configuration

1.72.3 Syntax Description

dhcpv6	Optional. Configure DHCPv6 PD pool attributes.
prefix/*prefix-length*	IPv6 prefix and prefix length for the associated IPv6 address, separated by the slash (/) character. The range of values for the ***prefix-length*** argument is 32 to 64.
prefix/prefix-length	Ending address of the IPv6 pool.
name ***pool-name***	Optional. Name for the IP pool; a string with up to 31 characters.
threshold	Optional. Creates a threshold value for which a crossing event occurs.
absolute	Specifies that the threshold is an absolute value (in this case, the number of entries in an IP pool).
percentage	Specifies that the threshold is a percentage, which is calculated internally.
falling	Indicates this is a falling-threshold crossing event.

first-threshold	Sets the threshold value. Range of values for an absolute threshold is is 0 to 65,535. Range of values for a percentage threshold is 1 to 100. If this value is omitted, the default value is 0.
log	Logs the threshold event; this keyword is optional if you specify the trap keyword.
trap	Reports the threshold event with a Simple Network Management Protocol (SNMP) event.
second-threshold	Creates a second falling-threshold value. Range of values is 0 to 65,535; if omitted, the default value is 0.

1.72.4 Default

No IPv6 prefix pool exists

1.72.5 Usage Guidelines

Use the **ipv6 pool** command in interface configuration mode to create a pool of IPv6 prefixes for a multibind interface to allow a subscriber to be assigned any available IP prefixes from the pool and, optionally, a threshold value for which a crossing event occurs. In this case, the threshold applies to a particular IP pool or to all pools configured under the specified interface.

Use the **ipv6 pool** command in context configuration mode to create a threshold value for which a crossing event occurs. In this case, the threshold applied to the entire context.

DHCPv6 PD pools lease IPv6 prefixes that can be assigned to DHCPv6 subscribers; DHCPv6 PD pools are managed by DHCPv6. Shared IPv6 pools lease prefixes that can be assigned to the WAN link of the CPE router; shared prefix pools are managed by AAA.

You can specify up to two threshold events for a pool. If you specify a second threshold event, you must configure the second threshold value to be less than the first threshold value. You can specify a maximum of two actions (traps, logs, or both) for each threshold event.

Note: In the current release, the shared IPv6 prefix pool is used by ND only.

Use the **no** form of this command to remove a pool of IPv6 prefixes for a multibind interface.

1.72.6 Examples

The following example shows how to create a DHCPv6 PD pool of IPv6 prefixes under the multibind interface **i1**:

```
[local]Redback(config-ctx)#interface i1 multibind
[local]Redback(config-if)#ipv6 pool dhcpv6 3001:db8:1:1::/64 3001:db8:1:100::/64
```

1.73 ipv6 prefix-list

ipv6 prefix-list ***pl-name***

no ipv6 prefix-list ***pl-name***

1.73.1 Purpose

Creates an IP Version 6 (IPv6) prefix list used to filter routes and enters IPv6 prefix list configuration mode.

1.73.2 Command Mode

Context configuration

1.73.3 Syntax Description

pl-name	IPv6 prefix list name.

1.73.4 Default

There are no preconfigured IPv6 prefix lists.

1.73.5 Usage Guidelines

Use the **ipv6 prefix-list** command to create an IPv6 prefix list used to filter routes and to enter IPv6 prefix list configuration mode where you can define conditions using the **permit** and **deny** commands.

Note: A reference to an IPv6 prefix list that does not exist, or does not contain any configured entries, implicitly matches and permits all IPv6 prefixes.

Use the **no** form of this command to remove an IPv6 prefix list.

1.73.6 Examples

The following example shows how to create the IPv6 prefix list, **list102**, and enter IPv6 prefix list configuration mode:

```
[local]Redback(config-ctx)#ipv6 prefix-list list102
[local]Redback(config-ipv6-prefix-list)#
```

1.74 ipv6 route

ipv6 route *ipv6-addr/prefix-length* {*next-hop-ipv6-addr* | **null0**} [**cost** *cost*] [**distance** *distance*] [**permanent**] [**tag** *tag*] [**bfd**]

no ipv6 route *ipv6-addr/prefix-length* {*next-hop-ipv6-addr* | **null0**} [**cost** *cost*] [**distance** *distance*] [**permanent**] [**tag** *tag*] [**bfd**]

1.74.1 Purpose

Configures one or more static routes when the system is not configured to dynamically select a route to the destination.

1.74.2 Command Mode

Context configuration

1.74.3 Syntax Description

ipv6-addr/prefix-length	IPv6 address (in the form *A:B:C:D:E:F:G:H*) and prefix length, separated by a slash (**/**). The range of values for the *prefix-length* argument is 0 to 128.
next-hop-ipv6-addr	IPv6 address of the next hop that can be used to reach the network. You cannot configure a link-local address as the IPv6 next-hop address.
null0	Creates a null interface to prevent routing loops.
cost *cost*	Optional. Cost of the route. The range of values is 0 to 15.
distance *distance*	Optional. Administrative distance assigned to the route. The range of values is 1 to 255.
permanent	Optional. Indicates that the route cannot be removed, even if the interface is shut down.
tag *tag*	Optional. Route tag (an unsigned 32-bit integer) is used as a match value for controlling redistribution through route maps., the range of values is 1 to 4,294,967,295; the default is 0.
bfd	Optional. Enables BFD for the static route.

1.74.4 Default

None

1.74.5 Usage Guidelines

Use the **ipv6 route** command to configure one or more static routes when the system is not configured to dynamically select a route to the destination.

A static route can be overridden by a dynamically learned route with a lower administrative distance.

Note: Do not configure a link-local address as the next-hop IPv6 address. The OS rejects attempts to configure a link-local address as the next hop.

Use the **null0** keyword to prevent routing loops. A null interface is always up and can never forward or receive traffic. The null interface provides an alternative method of filtering traffic. You can avoid the overhead involved with using access control lists by directing undesired network traffic to the null interface.

Note: The Open Shortest Path First Version 3 (OSPFv3) and Intermediate System-to-Intermediate System (IS-IS) routing processes always create a route to a null interface when summarizing a group of routes.

Use the **no** form of this command to remove static routes.

1.74.6 Examples

The following example routes packets for network, 2000:8A2E:56 48:CDF7:65B3:2F29:B3D5:3995/64, to the device at IPV6 address AB34:665F:B90B:3290:EA11:2678:FFFF:3210:

```
[local]Redback(config-ctx)#ipv6 route 2000:8A2E:5648:CDF7:65B3:2F29:B3D5:3995/64
AB34:665F:B90B:3290:EA11:2678:FFFF:3210
```

The following example configures a null interface for network 665F:B90B:3290:EA11:CDF7:65B3:2F29:B3D5/128:.

```
[local]Redback(config-ctx)#ipv6 route 665F:B90B:3290:EA11:CDF7:65B3:2F29:B3D5/128 null0
```

The following example routes packets for network 2000:8A2E:5 648:CDF7:65B3:2F29:B3D5:3995/64 to the device at IP address AB34:665F:B90B:3290:EA11:2678:FFFF:3210 if dynamic information with an administrative distance less than 110 is not available:

Note: Due to space constraints, the PDF version of this example cannot show the entire command line.

```
[local]Redback(config-ctx)#ipv6 route 2000:8A2E:5648:CDF7:65B3:2F29:B3D5:3995/64
AB34:665F:B90B:3290:EA11:2678:FFFF:3210 distance 110
```

The following example shows a static route configuration with bfd.

```
[local]Redback(config-ctx)#ipv6 route 3000::1/64 2001::5 bfd
```

1.75 ipv6 source-validation

ipv6 source-validation

no ipv6 source-validation

1.75.1 Purpose

In a subscriber record or profile, enables IPv6 source-address validation (SAV), which denies all IP packets from address sources that are not reachable through a subscriber's associated circuit.

1.75.2 Command Mode

- Default subscriber profile configuration
- Subscriber record configuration
- Subscriber profile configuration

1.75.3 Syntax Description

This command has no keywords or arguments.

1.75.4 Default

IPv6 SAV is disabled.

1.75.5 Usage Guidelines

Note: Generally, reverse path forwarding (RPF), IP SAV, and ingress filtering all refer to the same functionality.

Use the **ipv6 source-validation** command to enable IPv6 SAV. IPv6 SAV denies all IP packets from address sources that are not reachable through the subscriber's associated circuit. You can use this command to prevent source address spoofing.

Use the **no** version of this command to disable IPv6 SAV.

1.75.6 Examples

The following example shows how to enable IPv6 SAV for the subscriber, **bart**:

```
[local]Redback(config-ctx)#subscriber name bart
[local]BRAS(config-sub)#ipv6 source-validation
```

1.76 ipv6 unnumbered

ipv6 unnumbered ***interface-name***

{no | default} ipv6 unnumbered

1.76.1 Purpose

Configures IPv6 to use in IP address from another interface.

1.76.2 Command Mode

Interface configuration

1.76.3 Syntax Description

interface-name	Specifies the interface from which to use the IP address for IPv6

1.76.4 Default

The IPv6 in the current interface does not use an IP address from another interface.

1.76.5 Usage Guidelines

Use the **ipv6 unnumbered** command to configure IPv6 in one interface to use an IP address from another interface.

Use the **no** or **default** form of this command to not use an IP address from another interface.

1.76.6 Examples

The following example shows how to configure IPv6 in interface **Unnumbered** to use the IP address from interface **IPv6**:

```
[local]Redback#config
[local]Redback(config)#context local
[local]Redback(config-ctx)#interface Unnumbered
[local]Redback(config-if)#ipv6 unnumbered IPv6
```

1.77 ipv6 url

ipv6 url *ipv6-url*

no ipv6 url

1.77.1 Purpose and Usage Guidelines

Configures the IPv6 URL to which the current subscriber HTTP session is to be redirected. You can add this command and the **url** command (for IPv4 subscribers) to an HTTP redirect profile to enable HTTP redirects for dual stack subscribers, or you can add only the **ipv6 url** command to a profile to redirect only IPv6 subscribers.

1.77.2 Command Mode

HTTP redirect profile configuration

1.77.3 Syntax Description

ipv6-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—IPv6 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.

1.77.4 Default

An HTTP redirect URL is not configured.

1.77.5 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.

1.77.6 Example

The following example shows how to configure HTTP redirect profile `h` in an IPv6 access-list to redirect both IPv4 and IPv6 subscribers with a message the subscribers receive as they are redirected:

```
[local]Redback(config-ipv6-access-list)#http-
redirect profile h
!
[local]Redback(config-hr-profile)#url "http://2.2.2.2:8888"

[local]Redback(config-hr-profile)#ipv6 url "http://[2000:1:2::1]:80"

[local]Redback(config-hr-profile)# message "to be
redirected to this address"
```

1.78 isp-log

isp-log

no isp-log

1.78.1 Purpose

Enables ISP logging on the system.

1.78.2 Command Mode

Exec (15)

1.78.3 Syntax Description

This command has no keywords or arguments.

1.78.4 Default

ISP logging is enabled by default.

1.78.5 Usage Guidelines

Use the **isp-log** command to enable ISP logging on the system. The system requires you to confirm that you want to enable ISP logging on the system; type **y** to enable ISP logging. The ISP log persists across switchovers and reboots.

Use the **no** form of this command to disable ISP logging. If you disable the ISP log, the system removes any existing ISP log file. To save an existing ISP log file before disabling ISP logging, use the **copy** command to extract the file from the system. If you disable the ISP log, a warning displays asking you to confirm that you want to disable ISP logging. If you type **y**, the existing ISP log file and all ISP events logged on the system are removed.

1.78.6 Examples

The following example shows how to enable the ISP-log:

```
[local]Redback#en
[local]Redback#isp-log
This command will enable ISP logging on the system.
Are you sure you wish to proceed (y/n)?y
[local]Redback#
```

The following example shows how to disable the ISP-log:

```
[local]Redback#en
[local]Redback#no isp-log
WARNING: This command will remove all ISP events logged on the system and disable ISP logging.
Are you sure you wish to proceed (y/n)?y
[local]Redback#
```

1.79 isp-log add

isp-log add ***comment***

1.79.1 Purpose

Adds a comment to the ISP log.

1.79.2 Command Mode

Exec (15)

1.79.3 Syntax Description

comment	Adds a new entry to the ISP log file and specifies the comment to add in that entry. Can contain up to 128 characters.

1.79.4 Default

None

1.79.5 Usage Guidelines

Use the **isp-log add** command to specify a comment to add to the ISP log.

1.79.6 Examples

The following example shows how to add the comment "log_enabled" to the ISP log file:

```
[local]Redback#en
[local]Redback#isp-log add log_enabled
[local]Redback#
```

1.80 isp-log size

isp-log size *size*

{no | default} isp-log size

1.80.1 Purpose

Sets the maximum size of the ISP log file.

1.80.2 Command Mode

Global configuration

1.80.3 Syntax Description

size *size*	Maximum size of the ISP log file, in KB. Enter a value between 1 and 10,240.

1.80.4 Default

The default size of the ISP log is 1,024 KB (1 MB).

1.80.5 Usage Guidelines

Use the **isp-log size** command to set the size of the ISP log file.

When the ISP log file reaches the size limit, the system logs an entry in the ISP file stating that the file is full, stops writing any future log entries to the ISP log file, and the displays the following system error message: `%SYSLOG-6-ERROR: ISP logging disabled due to file size limit being reached`.

To resume logging entries in the ISP log file, you must either:

- Extract the ISP log file by using the **copy** ***src-url dest-url*** **clear** command to allow the system to generate a new ISP log file.
- Clear the ISP log file using the **clear isp-log** command.

Caution!

If the size of the existing ISP log is larger than the value you configure using the **isp-log size** command, the system displays a warning message stating that any existing ISP log files will be cleared.

Use the **no** or **default** form of this command to set the file size to the default value of 1,024 KB.

1.80.6 Examples

The following example shows how to set the ISP log file size to 5 MB:

```
[local]Redback#config
 Enter configuration commands, one per line, 'end' to exit
[local]Redback(config)#isp-log size 5120
[local]Redback(config)#
```

1.81 is type

is type {**level-1** | **level-1-2** | **level-2-only**}

no is type

1.81.1 Purpose

Configures the Intermediate System-to-Intermediate System (IS-IS) routing level used by the SmartEdge router for the specified IS-IS instance.

1.81.2 Command Mode

IS-IS router configuration

1.81.3 Syntax Description

level-1	Specifies that the SmartEdge router operates only in the level 1 area.
level-1-2	Specifies that the SmartEdge router participates in both IS-IS level 1 and level 2 routing.
level-2-only	Specifies that the SmartEdge router operates in level 2 only.

1.81.4 Default

The SmartEdge router participates in both level 1 and level 2 routing.

1.81.5 Usage Guidelines

Use the **is type** command to configure the IS-IS routing level used by the SmartEdge router for the specified IS-IS instance.

Use the **level-1** keyword to specify level 1 routing. All other destinations are routed to the closest device running either level 2 or both levels. If the wide-style metric is enabled with the **metric-style** command, routes can be advertised from level 2 areas into the level 1 area, and devices running level 1 can select the best level 2 device on a per-destination basis.

Use the **level-1-2** keyword to specify both level 1 and level 2 routing. The database and Shortest Path First (SPF) computation for each level is independent. When the wide-metric style is enabled with the **metric-style** command, the router can advertise and summarize level 1 routes into level 2 areas and the opposite.

Use the **level-2-only** keyword to specify level 2 routing.

Use the **no** form of this command to restore the SmartEdge router to the default behavior of participating in both level 1 and level 2 routing.

1.81.6 Examples

The following example shows how to configure the SmartEdge router for IS-IS **level-2-only** routing:

```
[local]Redback(config-ctx)#router isis ip-backbone
[local]Redback(config-isis)#is type level-2-only
```

1.82 join-group

join-group *group-ip* [**source** *source-ip*]

no join-group *group-ip* [**source** *source-ip*]

1.82.1 Purpose

Statically joins a bridge to a specified multicast group.

1.82.2 Command Mode

IGMP snooping bridge configuration

1.82.3 Syntax Description

group-ip	IP address of the multicast group you want the bridge to join.
source*source-ip*	Optional. Specifies a source device that sends IGMP packets to the group. Replace the ***A.B.C.D.*** argument with the IP address of the source device.

1.82.4 Default

No groups are statically joined to the bridge.

1.82.5 Usage Guidelines

Use the **join-group** command to statically join a bridge to a specified multicast group.

Use the **no** form of this command to remove a statically joined bridge from a group.

Use this command to enable ISSU features and functions by entering the license key required by this command. You cannot access ISSU features unless you enter the correct password value required by this command. XCRP4-based SmartEdge systems require different password values. You can use the **show licenses all** command to view whether ISSU features is enabled on the system. If you have successfully enabled ISSU, the output displays **yes** next to the ISSU line under Software Features heading in the CLI.

Use the **no issu password** command to disable MPLS functions and features. A password is not required if you are disabling the license for ISSU features and functions; it is ignored if entered.

For more information ISSU process, see the `Basic System Operations Guide` for the SmartEdge router.

1.82.6 Examples

The following example shows how to statically join a group with an IP address of `234.1.2.3` to a bridge called `br1`:

```
[local]Redback#configure
Enter configuration commands, one per line, 'end' to exit
[local]Redback(config)#context local
[local]Redback(config-ctx)#bridge br1
[local]Redback(config-bridge)#igmp snooping
[local]Redback(config-igmp-snooping)#join-group 234.1.2.3
```

The following example shows how to statically join a group with an IP address of `230.1.2.3` to a bridge called `br2` and specify a device with an IP address of `122.1.2.3` as a source:

```
[local]Redback#configure
Enter configuration commands, one per line, 'end' to exit
[local]Redback(config)#context local
[local]Redback(config-ctx)#bridge br2
[local]Redback(config-bridge)#igmp snooping
[local]Redback(config-igmp-snooping)#igmp join-group 230.1.2.3 source 122.1.2.3
```

1.83 keepalive (ANCP)

keepalive interval ***seconds*** **retry** ***retry-num***

{**no** | **default**} **keepalive**

1.83.1 Purpose

Configures the parameters for sending and receiving keepalive messages to and from Access Node Control Protocol (ANCP) neighbor peers.

1.83.2 Command Mode

ANCP configuration

1.83.3 Syntax Description

interval ***seconds***	Number of seconds between keepalive messages sent to ANCP neighbor peers. The range of values is 1 to 25; the default value is 10 seconds.
retry ***retry-num***	Number of missing keepalive messages permitted from an ANCP neighbor peer before the session is disconnected. The range of values is 1 to 10; the default value is 3.

1.83.4 Default

The interval value is 10 seconds; the retry value is 3.

1.83.5 Usage Guidelines

Use the **keepalive** command to configure the parameters for sending and receiving keepalive messages to and from ANCP neighbor peers.

The SmartEdge router keeps track of the number of missing keepalive messages from each ANCP neighbor peer. If the number missing messages exceeds that specified by the **retry** ***retry-num*** construct, it disconnects the session for that peer.

Caution!

Risk of performance loss. When the system has many active General Switch Management Protocol (GSMP) peer sessions and the value of the ***seconds*** argument in the **keepalive** command syntax is less than 10, the system might incur a loss of performance. To minimize the risk under these conditions, change the value of the ***seconds*** argument to 10 or greater.

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

1.83.6 Examples

In the following example, the SmartEdge router sends keepalive messages to ANCP neighbor peers every **5** seconds. It disconnects the session to an ANCP neighbor peer if it does not receive **10** keepalive messages from that peer:

```
[local]Redback(config-ancp)#keepalive interval 5 retries 10
```

1.84 keepalive (channel)

keepalive [**check-interval** {**minutes** | **seconds**} ***time***] [**retries** ***retry-num***]

no keepalive

default keepalive [**check-interval**] [**retries**]

1.84.1 Purpose

Enables the keepalive function on a DS-1 channel on a channelized DS-3 channel or port, clear-channel DS-3 channel or port, E3 port, E1 channel or port, or DS-0 channel group on a channelized E1 channel or port that is encapsulated with Cisco High-Level Data Link Control (HDLC).

1.84.2 Command Mode

- DS-0 group configuration
- DS-1 configuration
- DS-3 configuration
- E1 configuration
- E3 configuration

1.84.3 Syntax Description

check-interval	Optional. Sets the time interval between keepalive checks.
minutes	Specifies that the unit of measure for the ***time*** argument is minutes.
seconds	Specifies that the unit of measure for the ***time*** argument is seconds; this is the default.

time	Time in either minutes or seconds (depending on the preceding keyword) between keepalive checks. The range of values is 1 to 60 minutes, or 1 to 300 seconds; the default value is 10 seconds.
retries ***retry-num***	Optional. Number of times the system is to retry an unsuccessful keepalive check. The range of values is 2 to 10; the default value is 3.

1.84.4 Default

The keepalive function is enabled with an interval of 10 seconds and 3 messages.

1.84.5 Usage Guidelines

Use the **keepalive** command to enable the keepalive function on a DS-1 channel on a channelized DS-3 channel or port, clear-channel DS-3 channel or port, E3 port, E1 channel or port, or DS-0 channel group on a channelized E1 channel or port that is encapsulated with Cisco HDLC.

This command specifies the interval between keepalive messages and the number of unconfirmed messages, either keepalive or packets, before declaring that the connection is broken:

- If the remote end does not have the keepalive function enabled, the connection is declared broken after the specified number of keepalive messages have been sent and are unconfirmed.
- If the remote end does have the keepalive function enabled, the connection is declared broken after the specified number of packet or keepalive messages have been sent and are unconfirmed.
- The interval must be the same on both ends of the connection.

Use the **no** form of this command to disable the keepalive function.

Use the **default** form of this command or enter the **keepalive** command without keywords to set the interval and number of messages to their defaults.

Note: This command is also described in *Configuring ATM, Ethernet, and POS Ports* for Packet over SONET/SDH (POS) ports.

1.84.6 Examples

The following example shows how to set the keepalive interval to **20** and the number of unconfirmed messages to **5** on clear-channel DS-3 channel **1:**

```
[local]Redback(config)#port ds3 3/1:1
[local]Redback(config-ds3)#encapsulation cisco-hdlc
[local]Redback(config-ds3)#keepalive check-interval seconds 20 retries
```

1.85 keepalive (LDP)

keepalive {holdtime *seconds* | **interval** *seconds***}**

no keepalive {holdtime *seconds* | **interval** *seconds***}**

1.85.1 Purpose

Enables the configuration of intervals controlling Label Distribution Protocol (LDP) liveliness detection and LDP keepalive packet transmission for a SmartEdge® router Label Switched Router (LSR).

1.85.2 Command Mode

LDP router configuration

1.85.3 Syntax Description

holdtime *seconds*	Number of seconds after which an inactive LDP session will be terminated and the corresponding TCP session will be closed. Inactivity is defined as not receiving any LDP packets from the neighbor. The range of values is 45 to 3600 seconds. The default value is 90 seconds.
interval *seconds*	Number of seconds between successive transmissions of keepalive packets. Keepalive packets are only sent in the absence of other LDP packets transmitted over the LDP session. The range of values is 15 to 1200 seconds. The default value is 30 seconds.

1.85.4 Default

The default holdtime is 90 seconds. The default interval is 30 seconds. The system enforces specification of an interval that is at least half the holdtime (whether the holdtime is defaulted to 90 seconds or specified explicitly).

1.85.5 Usage Guidelines

Use the **keepalive** command to enable the configuration of intervals controlling LDP liveliness detection and LDP keepalive packet transmission for a SmartEdge® router LSR. The **holdtime** option sets the number of seconds after which an inactive LDP session will be terminated and the corresponding TCP session will be closed. The interval option sets the number of seconds between successive transmission of keepalive packets.

Use the **no** or **default** form of this command to return to the default holdtime of 90 seconds or to return to the default interval of 30 seconds.

1.85.6 Examples

The following example shows how to set the number of seconds after which an inactive LDP session will be terminated to **45** and set the number of seconds between successive transmissions of keepalive packets to **100**:

```
[local]Redback(config-ldp)#keepalive holdtime 45

[local]Redback(config-ldp)#keepalive interval 100
```

1.86 keepalive (POS)

keepalive [**check-interval** {**minutes** | **seconds** | **milliseconds**} *time*] [**retries** *retry-num*]

no keepalive

default keepalive [**check-interval**] [**retries**]

1.86.1 Purpose

Enables the keepalive function on a Packet over SONET/SDH (POS) port that is encapsulated with Cisco High-Level Data Link Control (HDLC).

1.86.2 Command Mode

Port configuration

1.86.3 Syntax Description

check-interval	Optional. Sets the time interval between keepalive checks.
minutes *time*	Time in minutes between keepalive checks. The range of values is 1 to 60 minutes.
seconds *time*	Time in seconds between keepalive checks. The range of values is 1 to 300 seconds.
milliseconds *time*	Time in milliseconds between keepalive checks. The range of values is 30 to 1000 milliseconds, rounded to the nearest 10 milliseconds.
retries *retry-num*	Optional. Number of times the system is to retry an unsuccessful keepalive check. The range of values is 2 to 10; the default value is 3.

1.86.4 Default

The keepalive function is enabled with an interval of 10 seconds and three retries.

1.86.5 Usage Guidelines

Use the **keepalive (POS)** command to enable the keepalive function on a POS port that is encapsulated with Cisco HDLC. This command specifies the interval between keepalive messages and the number of unconfirmed messages, either keepalive or packets, before declaring that the connection is broken:

- If the remote end does not have the keepalive function enabled, the connection is declared broken after the specified number of keepalive messages have been sent.
- If the remote end does have the keepalive function enabled, the connection is declared broken after the specified number of packet or keepalive messages have been sent and are unconfirmed.
- The interval must be the same on both ends of the connection.

  Note: The keepalive function is disabled on a port in an Automatic Protection Switching (APS) group when the traffic status of the port is Standby.

Use the **no** form of this command to disable the keepalive function.

Use the **default** form of this command or enter the command without keywords to specify the default values for the interval and number of messages.

1.86.6 Examples

The following example shows how to specify the keepalive interval as **20** and the number of unconfirmed messages as **5** on a POS port:

```
[local]Redback(config)#port pos 1/8
[local]Redback(config-port)#encapsulation cisco-hdlc
[local]Redback(config-port)#keepalive check-interval seconds 20 retries 5
```

1.87 keepalive (tunnel)

keepalive [*seconds* [*retry-num*]]

{**no** | **default**} **keepalive**

1.87.1 Purpose

Enables the sending of keepalive packets on Generic Routing Encapsulation (GRE) tunnels and specifies the interval and number of retries.

1.87.2 Command Mode

Tunnel configuration

1.87.3 Syntax Description

seconds	Optional. Number of seconds between the sending of a keepalive packet. The range of values is 1 to 32,766; the default value is 10.
retry-num	Optional. Number of times a keepalive packet is sent without response before the tunnel is brought down. The range of values is 2 to 254; the default value is 4.

1.87.4 Default

The sending of keepalive packets is disabled.

1.87.5 Usage Guidelines

Use the **keepalive** command to enable the sending of keepalive packets on GRE tunnels and specify the interval between keepalive packets and the number of retries.

Use the **no** form of this command to disable the sending of keepalive packets.

Use the **default** form of this command to specify the default values for the ***seconds*** argument and the ***retry-num*** argument.

1.87.6 Examples

The following example shows how to enable the sending of keepalive packets with the default values for the ***seconds*** and ***retry-num*** arguments:

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

1.88 keep-multiplier

keep-multiplier ***multiplier***

1.88.1 Purpose

Configures the Resource Reservation Protocol (RSVP) keep-multiplier timing parameter.

1.88.2 Command Mode

RSVP interface configuration

1.88.3 Syntax Description

multiplier	Multiplier used for calculating the lifetime of a reservation state. The range of values is 1 to 255.

1.88.4 Default

The default keep-multiplier value is 3.

1.88.5 Usage Guidelines

Use the **keep-multiplier** command to configure the RSVP keep-multiplier timing parameter.

When RSVP is enabled, refresh messages are sent periodically so that reservation states in neighboring nodes do not expire. The lifetime of a reservation state is determined by using two interrelated timing parameters: the keep-multiplier and the refresh-interval. Use the following formula to determine the lifetime of a reservation state:

Lifetime = (keep-multiplier + 0.5) * 1.5 * refresh-interval

1.88.6 Examples

The following example configures the keep-multiplier timing parameter to **15:**

```
[local]Redback(config-ctx)#router rsvp
[local]Redback(config-rsvp)#interface rsvp05
[local]Redback(config-rsvp-if)#keep-multiplier 15
```

1.89 key-chain

key-chain ***key-chain-name*** **key-id** ***key-id***

no key-chain ***key-chain-name*** [**key-id** ***key-id***]

1.89.1 Purpose

Creates a new key chain with a key, or creates a key within an existing key chain, and enters key chain configuration mode.

1.89.2 Command Mode

Context configuration

1.89.3 Syntax Description

key-chain-name	Name of the key chain.
key-id	Identification number of a key within the chain. The range of values is 1 to 65,535. Must be unique within the key chain. Optional only when deleting a key chain.

1.89.4 Default

No key chains are created.

1.89.5 Usage Guidelines

Use the **key-chain key-id** command to create a new key chain with a key, or to create a key within an existing key chain, and to enter key chain configuration mode.

Key chains allow you to control authentication keys used by various routing protocols in the system. Currently, the SmartEdge router supports the use of key chains with the Mobile IP services, Open Shortest Path First (OSPF), intermediate-system-to-intermediate-system (IS-IS), and Virtual Router Redundancy Protocol (VRRP) routing protocols.

For information about the **authentication** command used with the **key-chain key-id** command for routing protocols, see the *Command List*. For information about the **authentication** command that is used with the **key-chain key-id** command for Mobile IP services, see the *Command List* .

Use the **no** form of this command with the **key-id** ***key-id*** construct to remove a key from the key chain configuration. Use the **no** form of this command without the optional construct to remove the entire key chain.

1.89.6 Examples

The following example shows how to create a new key chain, **superkeychain**, and create three keys within it (IDs **200**, **201**, **202**), each with its own string and lifetime:

```
[local]Redback(config-ctx)#key-chain superkeychain key-id 200
[local]Redback(config-key-chain)#key-string di492jffs
[local]Redback(config-key-chain)#accept-lifetime 2001:01:01:01:01 duration 10000
[local]Redback(config-key-chain)#send-lifetime 2001:01:01:01:01 infinite
[local]Redback(config-key-chain)#key-chain superkeychain key-id 201
[local]Redback(config-key-chain)#key-string 7744kkciao
[local]Redback(config-key-chain)#accept-lifetime 2001:01:01:01:1 infinite
[local]Redback(config-key-chain)#send-lifetime 2001:01:01:01:01
[local]Redback(config-key-chain)#key-chain superkeychain key-id 202
[local]Redback(config-key-chain)#key-string secret222
[local]Redback(config-key-chain)#accept-lifetime 2001:01:01:01:01 2002:01:01:00:00
[local]Redback(config-key-chain)#send-lifetime 2001:01:01:01:01 infinite
```

In this example, you do not have to exit from key chain configuration mode before you enter the **key-chain** command because commands from the next highest mode in the hierarchy (context configuration mode, in this case) are accepted in any configuration mode.

1.90 key-chain description

key-chain ***key-chain-name*** **description** ***text***

no key-chain *key-chain-name* [**description** *text*]

1.90.1 Purpose

Configures a key chain name and description.

1.90.2 Command Mode

Context configuration

1.90.3 Syntax Description

key-chain-name	Name of the key chain.
text	Alphanumeric text description to be associated with the key chain. Optional only when deleting a key chain.

1.90.4 Default

No key chains are created.

1.90.5 Usage Guidelines

Use the **key-chain description** command to configure a key chain name and description.

Only one description can be associated with a single key chain. To update a description, issue this command with the new description; the old description is overwritten.

Use the **no** form of this command with the **description** ***text*** construct to remove a description from the key chain configuration. Use the **no** form of this command without the optional construct to delete the entire key chain.

1.90.6 Examples

The following example shows how to configure **key01** with a text description specifying **3 keys ospf only:**

```
[local]Redback(config-ctx)#key-chain key01 description 3 keys ospf only
```

1.91 key-string

key-string {*string* | **hex** *hex-string*}

no key-string {*string* | **hex** *hex-string*}

1.91.1 Purpose

Configures a string for the specified key.

1.91.2 Command Mode

Key chain configuration

1.91.3 Syntax Description

string	Alphanumeric string.
hex *hex-string*	Hexadecimal string. Must be composed of valid hexadecimal characters (A-F, a-f, 0-9) and may be preceded by an optional 0x or 0X. The 0x or 0X is not included in the stored key string.

1.91.4 Default

No key string is configured.

1.91.5 Usage Guidelines

Use the **key-string** command to configure a string for the specified key. A string is equivalent to a password and is encrypted in the output of the **show configuration** command. In the output of the **show key-chain** command, the key string is shown both encrypted and unencrypted. You can replace an existing key string by using the **key-string** command again, specifying a new string.

The SmartEdge router stores hexadecimal strings left justified in the key string with the remaining characters set to 0x0.

Use the **no** form of this command to remove the key string from the configuration.

1.91.6 Examples

The following example shows how to configure **7744kkciao** as the string for the key chain, **secretkeychain:**

```
[local]Redback(config-ctx)#key-chain secretkeychain key-id 200

[local]Redback(config-key-chain)#key-string 7744kkciao
```

1.92 l2protocol-tunnel

l2protocol-tunnel

no l2protocol-tunnel

1.92.1 Purpose

Sets the Layer 2 Protocol tunnel attribute in the spanning tree profile, and enables circuits assigned the spanning-tree profile to send bridge protocol data units (BPDUs) using the group MAC address.

1.92.2 Command Mode

Spanning-tree profile configuration

1.92.3 Syntax Description

This command has no keywords or arguments.

1.92.4 Default

The associated port is not enabled for sending BPDUs.

1.92.5 Usage Guidelines

Use the **l2protocol-tunnel** command to set the Layer 2 Protocol tunnel attribute in the spanning-tree profile, and enable circuits assigned the spanning-tree profile to send BPDUs using the group MAC address.

BPDUs sent through the Layer 2 Protocol tunnel go to the group MAC address. All other circuits send BPDUs to the standard MAC address.

Use the **group-mac-address** command (in spanning-tree configuration mode) to set the group MAC destination address.

1.92.6 Examples

The following example illustrates the creation of the spanning-tree profile `womp` in which the **`l2protocol-tunnel`** command is set to enable the associated ports to send BPDUs through the Layer 2 Protocol tunnel.

The **`spanning-tree profile`** command (port configuration mode) then assigns the spanning-tree profile to an Ethernet port.

In the last part of the configuration, the **`group-mac-address`** command (in bridge configuration mode) specifies the destination MAC address for BPDUs sent through the Layer 2 Protocol tunnel:

```
[local]Redback(config)#spanning-tree profile womp
[local]Redback(config-stp-prof)#l2protocol-tunnel
[local]Redback(config-stp-prof)#exit
[local]Redback(config)#port ethernet 1/1
[local]Redback(config-port)#spanning-tree profile womp
[local]Redback(config-ctx)#bridge isp3
[local]Redback(config-bridge)#description Bridge for all traffic to ISP3
[local]Redback(config-bridge)#aging-time 18000
[local]Redback(config-bridge)#spanning-tree
[local]Redback(config-bridge-stp)#group-mac-address 01.80.C2.00.00.02
```

1.93 l2tp

`l2tp` [**`all`**]{**`encrypted 1`** | **`password`**} ***`password`***

`no l2tp` [**`all`**]

1.93.1 Purpose

Enables Layer 2 Tunneling Protocol (L2TP) features and functions.

1.93.2 Command Mode

Software license configuration

1.93.3 Syntax Description

`all`	Optional. Enables all L2TP features and functions; this is the default.
`encrypted 1`	Specifies that the password that follows is encrypted.

password	Specifies that the password that follows is not encrypted
password	Paid license password that is required to enable L2TP features and functions. The ***password*** argument is unique for L2TP and is provided at the time the software license is paid.

1.93.4 Default

L2TP features and functions are disabled.

1.93.5 Usage Guidelines

Use the **l2tp** command to enable L2TP features and functions. You can specify the ***password*** argument in either encrypted or unencrypted form. Neither form displays by the **show configuration** ***command*** command (in any mode).

Use the **no** form of this command to disable L2TP features and functions. A password is not required if you are disabling the license for any of the L2TP features and functions; it is ignored if entered.

1.93.6 Examples

The following example shows how to license L2TP features and functions. The password is in an unencrypted form:

```
[local]Redback(config-license)#l2tp all password l2tp-password
```

1.94 l2tp admin

l2tp admin {**down peer** ***peer-name*** [**seconds** ***seconds***] | **up peer** ***peer-name***}

1.94.1 Purpose

Marks a Layer 2 Tunneling Protocol (L2TP) peer as up (“alive”) or down (“dead”).

1.94.2 Command Mode

Exec

1.94.3 Syntax Description

down	Marks the L2TP peer as “dead”; no new sessions are assigned to this peer.
peer ***peer-name***	Name of the L2TP peer to be marked.
seconds ***seconds***	Optional. Number of seconds for which the L2TP peer is marked as “dead”; the range of values is 1 to 60,000.
up	Marks the L2TP peer as “alive”; new sessions are assigned to this peer.

1.94.4 Default

No L2TP peer is marked as down or up.

1.94.5 Usage Guidelines

Use the **l2tp admin** command to mark an L2TP peer as up (“alive”) or down or (“dead”).

Use the **down** keyword to gracefully remove a peer from the configuration. Use the **seconds** ***seconds*** construct to specify the interval after which the peer is restored to the “alive” state.

1.94.6 Examples

The following example shows how to mark the L2TP peer, **ira**, as “dead”:

```
[local]Redback#l2tp admin ira down
```

1.95 l2tp admin test

l2tp admin test peer ***peer-name*** {**hello** | **ses-setup** [**count** [***num***]] | **tunl-setup**}

1.95.1 Purpose

Performs Layer 2 Tunneling Protocol (L2TP) peer testing.

1.95.2 Command Mode

Exec

1.95.3 Syntax Description

peer ***peer-name***	Name of the L2TP peer to be tested.
hello	Sends Hello message to any idle tunnels.
ses-setup	Performs session testing.
count	Optional. Specifies the number of sessions to set up.
num	Optional. Number of sessions to set up; the range of values is 1 to 10,000.
tunl-setup	Performs tunnel testing.

1.95.4 Default

None

1.95.5 Usage Guidelines

Use the **l2tp admin test** command to perform L2TP peer testing. You can test any idle tunnels to a peer using Hello messages, test the tunnels, and test sessions on tunnels.

Use the **ses-setup** keyword to test if an L2TP network server (LNS) peer allows a session to be set up without the need for a client to connect to it. The L2TP session is established, but the Point-to-Point Protocol (PPP) is not negotiated for the session with the peer; as a result, the peer times out and closes the session.

Use the **tunl-setup** keyword to test if a tunnel can be created to a peer; this test validates the remote IP address and if configured, the local IP address, specified by the **l2tp-peer** command (in context configuration mode), and the tunnel authorization key, specified by the **tunnel-auth key** command (in L2TP peer configuration mode).

1.95.6 Examples

The following example shows how to test the L2TP peer, **ira**, using a **hello** message:

```
[local]Redback#l2tp admin test peer ira hello
```

1.96 l2tp avp

l2tp avp {**rx-speed** | **tx-speed**} **source** {**dslam** | **qos** | **report**}

no l2tp avp {**rx-speed** | **tx-speed**}

1.96.1 Purpose

Enables population of the Layer 2 Tunneling Protocol (L2TP) Receive (Rx) Connect Speed or Transmit (Tx) Connect Speed attribute-value pair (AVP) from a custom source.

1.96.2 Command Mode

Context configuration

1.96.3 Syntax Description

rx-speed	Populates the L2TP Rx Connect Speed AVP.
tx-speed	Populates the L2TP Tx Connect Speed AVP.
source	Specifies the source used to populate the AVP.
dslam	Populates the AVP with the vendor tag value.
qos	Populates the AVP with the QoS policy value.
report	If the traffic card has a profile available, populates the AVP with the value from the profile; otherwise, populates the AVP with the port speed.

1.96.4 Default

The default behavior uses the **report** keyword, and populates the AVP with either the value from the profile or the port speed, depending on the type of traffic card used.

1.96.5 Usage Guidelines

Use the **l2tp avp** command to enable the population of L2TP Rx Connect Speed (38) or L2TP Tx Connect Speed (24) AVPs from a custom source.

If you choose the **report** keyword, this command populates the AVP value with the rate from the circuit profile. For instance, if you are configuring an ATM circuit, use the **report** command (in ATM profile configuration mode) to set the RX speed and Tx speed. This is then picked up by the **l2tp avp** command.

Use the **no** form of this command to reset the behavior to the default settings.

1.96.6 Examples

The following example shows how to populate the rx-speed from the circuit profile:

```
[local]Redback(config-ctx)#l2tp avp rx-speed source report
```

1.97 l2tp avp calling-number

To set the format of the subscriber-dialed calling number, use the following syntax:

```
l2tp avp calling-number format {all | {[hostname]
[pppoe-id] [slot-port] [use-CLID] [virtual-id]}
```

```
{no | default} l2tp avp calling-number format
```

To set the separator character used in the subscriber-dialed calling number, use the following syntax:

```
l2tp avp calling-number separator character
```

```
{no | default} l2tp avp calling-number separator
```

1.97.1 Purpose

Specifies the subscriber calling information to be passed to a Layer 2 Tunneling Protocol (L2TP) network server (LNS) in a Dialed Number Identification Service (DNIS) attribute-value pair (AVP).

1.97.2 Command Mode

Context configuration

1.97.3 Syntax Description

all	Includes all options with the exception of **use-CLID**; this setting is the default.
hostname	Optional. Includes the currently configured hostname of the router.
pppoe-id	Optional. Includes the session ID of the incoming Point-to-Point Protocol over Ethernet (PPPoE) session.
slot-port	Optional. Includes the slot number and port number of the incoming circuit.

use-CLID	Optional. Populates AVP #22 with the same information that is sent to RADIUS when using the **radius attribute calling-station-id** command (in context configuration mode).
virtual-id	Optional. Includes the virtual path identifier (VPI), virtual channel identifier (VCI), or virtual LAN ID (VLAN ID) of the incoming circuit.

1.97.4 Default

All options are sent to the peer with the exception of **use-CLID**. The separator character is a space.

1.97.5 Usage Guidelines

Use the **l2tp avp calling-number format** command to specify what subscriber calling information is passed to an LNS in a DNIS AVP.

Note: An L2TP access concentrator (LAC) sends an AVP only if the **dnis generate** command (in L2TP peer configuration mode) is configured and enabled under the peer.

Use the **no** or **default** form of this command to send all options to the peer and interpret the space as the separator character.

1.97.6 Examples

The following example shows how to display all information (hostname, slot, and port, PPPoE ID, and virtual ID):

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#l2tp avp calling-number format all
[local]Redback(config-ctx)#
```

The following example shows how to display only the hostname:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#l2tp calling-number format hostname
```

1.98 l2tp avp nas-port-id format all

l2tp avp nas-port-id format all [**include-mac**]

no l2tp avp nas-port-id format all

1.98.1 Purpose

Enables a SmartEdge router configured as an LAC to propagate physical port information that is compatible with an SMS router configured as an LNS. Also, enables a SmartEdge router configured as an LNS to propagate physical port information that is compatible with an SMS router configured as an LAC.

1.98.2 Command Mode

Context configuration

1.98.3 Syntax Description

include-mac	Instructs the SmartEdge router to also include the PPPoE client's MAC address in the building of AVP #49.

1.98.4 Default

When the SmartEdge router is configured as an LAC, AVP#49 is not sent to the LNS. When the SmartEdge router is configured as an LNS, the fixed string, `256/17, is sent to the RADIUS server.`

1.98.5 Usage Guidelines

Use the **l2tp avp nas-port-id format all** command to enable the SmartEdge router configured as an LAC to propagate physical port information that is compatible with an SMS router configured as an LNS. Also, use this command to enable the SmartEdge router configured as an LNS to propagate physical port information that is compatible with an SMS router configured as an LAC.

On a SmartEdge router configured as an LAC, the **l2tp avp nas-port-id format all** command specifies sending the Vendor Specific L2TP AVP #49 (NAS-Port-ID attribute) to the LNS with physical port information in clear text format compatible with an SMS LNS.

The default behavior for the SmartEdge router configured as the LAC is that AVP #49 is not sent to the LNS.

On a SmartEdge router configured as an LNS, the **l2tp avp nas-port-id format all** command specifies sending AVP #49 (when and if it is received from the LAC) to the RADIUS server.

The default behavior for the SmartEdge router configured as the LNS is to send the fixed string, `256/17`, to the RADIUS server.

An example of the physical port information that could be found in AVP #49 follows:

```
3/1 vpi-vci 7 308 pppoe 287
```

Use the **no** form of this command to reset the behavior to the default settings.

1.98.6 Examples

The following example shows how to use the **l2tp avp nas-port-id format all** command:

```
[local]Redback(config-ctx)#l2tp avp nas-port-id format all
```

1.99 l2tp clear-radius-peer

l2tp clear-radius-peer ***time-inactive***

{**no** | **default**} **l2tp clear-radius-peer**

1.99.1 Purpose

Enables any Layer 2 Tunneling Protocol (L2TP) peer configured by a Remote Authentication Dial-In User Service (RADIUS) server in this context to be automatically removed from memory after it is marked inactive.

1.99.2 Command Mode

Context configuration

1.99.3 Syntax Description

time-inactive	Time, in minutes, that a peer can be inactive before being removed from memory. The range of values is 5 to 300.

1.99.4 Default

No time limit is in effect; no inactive RADIUS-configured peers are cleared from memory.

1.99.5 Usage Guidelines

Use the **l2tp clear-radius-peer** command to enable any L2TP peer configured by a RADIUS server in this context to be automatically removed from memory after it is marked inactive. A RADIUS-configured peer is marked as inactive if:

- The session count is 0.
- The peer is not labeled “dead”; it is alive or its deadtime has expired.
- The time interval since the last session was terminated or since the peer was initially created, if no sessions have been active, is equal to or greater than the time specified by the ***time-inactive*** argument.

If a RADIUS-configured peer is inactive, it is cleared from memory.

Use the **no** or **default** form of this command to remove the time limit.

1.99.6 Examples

The following example shows how to set the inactive time limit to **10** minutes:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#l2tp clear-radius-peer 10
```

1.100 l2tp deadtime

l2tp deadtime ***minutes***

{**no** | **default**} **l2tp deadtime**

1.100.1 Purpose

Sets the minimum amount of time for which any “dead” Layer 2 Tunneling Protocol (L2TP) peer that is configured in the context and that is not a member of a peer group is ignored.

1.100.2 Command Mode

Context configuration

1.100.3 Syntax Description

minutes	Minimum number of minutes that a peer is marked as “dead”. The range of values is 1 to 100; the default value is 2.

1.100.4 Default

The deadtime is set to 5 minutes.

1.100.5 Usage Guidelines

Use the **l2tp deadtime** command to set the minimum amount of time that any “dead” L2TP peer that is configured in the context and that is not a member of a peer group is ignored. You can use this command to control the deadtime for peers created by the Remote Authentication Dial-In User Service (RADIUS).

A peer is labeled “dead” after it is determined that a new tunnel cannot be established to the peer. This feature prevents a troubled L2TP peer from being inundated with connection attempts without disconnecting the peer altogether. It also allows you to identify troubled peers.

A peer remains labeled as “dead” until a new session is established to it as follows:

- After the deadtime is expired and a connection request arrives, the peer is again considered as a destination.
- If a connection attempt is not made to the peer (the peer is not selected as the destination), the “dead” label is not removed.
- If a connection attempt is made and is successful, the “dead” label is removed from the peer; if the attempt is not successful, the deadtime is again applied to the peer.

  Note: Current sessions to the peer are not brought down if the peer should be labeled “dead”. Only attempts to add new tunnels are affected.

A “dead” peer is labeled as “dead” in the output of the **show l2tp peer** command (in any mode) for at least the length of time indicated in the ***minutes*** argument.

Use the **no** or **default** form of this command to set the deadtime to 2 minutes.

1.100.6 Examples

The following example shows how to set the number of deadtime minutes to **10** for any L2TP peer that is not a member of a peer group in the context:

```
[local]Redback(config-ctx)#l2tp deadtime 10
```

1.101 l2tp fragment

l2tp fragment {**l2tp-packet** | **user-packet**}

{**no** | **default**} **l2tp fragment**

1.101.1 Purpose

Specifies the type of fragmentation used for Layer 2 Tunneling Protocol (L2TP) packets that are sent downstream and that need fragmentation.

1.101.2 Command Mode

Context configuration

1.101.3 Syntax Description

l2tp-packet	Fragments the encapsulating packet after the L2TP header is added; this is the default.
user-packet	Fragments the user data packet before the L2TP header is added.

1.101.4 Default

Fragmentation occurs after the L2TP header is added.

1.101.5 Usage Guidelines

Use the **l2tp fragment** command to specify the type of fragmentation for L2TP packets that are sent downstream.

It is more efficient to fragment the user data packet, because it is reassembled on the user's computer; fragmenting the L2TP packet requires that the L2TP access concentrator (LAC) must reassemble the packet, which takes more processing time.

Use the **no** or **default** form of this command to specify fragmentation after the L2TP header is added.

1.101.6 Examples

The following example shows how to enable fragmentation for user data packets before the L2TP header is added:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#l2tp fragment user-packet
```

1.102 l2tp-group

l2tp-group name ***l2tp-group-name*** **description** ***text***

no l2tp-group name ***l2tp-group-name*** **description** ***text***

1.102.1 Purpose

Creates a group of Layer 2 Tunneling Protocol (L2TP) tunnels to L2TP network servers (LNSs) among which Point-to-Point Protocol (PPP) sessions are parceled out, and enters L2TP group configuration mode.

1.102.2 Command Mode

Context configuration

1.102.3 Syntax Description

name ***l2tp-group-name***	Name of the L2TP group being created. L2TP group names must be unique from other L2TP group names, peer names, and domain aliases in the context.
description ***text***	Text description of the L2TP group.

1.102.4 Default

No L2TP group is created.

1.102.5 Usage Guidelines

Use the **l2tp-group** command to create a group of L2TP tunnels to LNSs (peers) among which PPP sessions are parceled out, and enter L2TP group configuration mode. All peers in a group must be defined (with the **l2tp-peer** command in context configuration mode) within the same context as the group itself. It is part of the LAC configuration.

PPP sessions are distributed among tunnels in a group according to the algorithm specified for the group with the **algorithm** command in L2TP group configuration mode.

A group name that is created with the **l2tp-group** command can be entered as the ***l2tp-peer-name*** or ***tunl-name*** argument value for the **tunnel name** command in subscriber configuration mode.

Peer names, group names, and domain aliases for those names must be unique within the context in which they are created.

Use the **description** ***text*** construct to delete all description information for the L2TP group, peer, or domain alias.

Use the **no** form of this command to disband the L2TP group and delete all references to it by the L2TP peers that formed the group and to delete all description information.

1.102.6 Examples

The following example shows how to create an L2TP group, **group1**:

```
[local]Redback(config-ctx)#l2tp-group name group1

[local]Redback(config-l2tp-group)#
```

1.103 l2tp-peer

l2tp-peer {**default** | **name** ***l2tp-peer-name*** **media udp-ip remote** {**ip** ***ip-addr*** | **dns** ***dns-name***} | **unnamed**} [**local** ***ip-addr***]

no l2tp-peer {**default** | **name** ***l2tp-peer-name*** | **unnamed**}

1.103.1 Purpose

Creates a Layer 2 Tunneling Protocol (L2TP) peer, either an L2TP access concentrator (LAC) or an L2TP network server (LNS), a default peer, or an anonymous (unnamed) peer, or selects one for modification, in the current context, and enters L2TP peer configuration mode.

1.103.2 Command Mode

Context configuration

1.103.3 Syntax Description

default	Creates a default L2TP tunnel.
name ***l2tp-peer-name***	Name of the L2TP peer that the peer supplies as a hostname in Start-Control-Connection-Request (SCCRQ) packets sent to the SmartEdge router.
media udp-ip	Specifies that the tunnel is User Datagram Protocol (UDP) IP-encapsulated.
remote ip ***ip-addr***	IP address of the L2TP peer.

remote dns ***dns-name***	Domain Name System (DNS) name of the L2TP peer.
unnamed	Creates an anonymous L2TP peer.
local ***ip-addr***	Optional. Local IP address. When configuring a LAC, the ***ip-addr*** argument requests the IP address of the LAC. When configuring an LNS, the ***ip-addr*** argument requests the IP address of the LNS.

1.103.4 Default

No L2TP named, default, or anonymous peer is created.

1.103.5 Usage Guidelines

Use the **l2tp-peer** command to create an L2TP peer, a default peer, or an anonymous peer, or select one for modification, in the current context, and enter L2TP peer configuration mode.

Use the **default** keyword to create a set of defaults that apply to any L2TP peer in the current context. Each configured attribute for the default peer is included in all L2TP peer configurations in the context. However, if you configure a named or anonymous peer, attribute values that you specify for that peer override the values set for the default peer.

If you specify the **name** ***l2tp-peer-name*** construct, the L2TP peer name must be unique from other L2TP peer names, group names, and domain aliases within the context.

When configuring the SmartEdge router as a LAC, the ***l2tp-peer-name*** argument is the name or the domain alias for the LNS at the other end of the tunnel; it represents the peer in the hostname attribute of packets exchanged in L2TP. When configuring it as an LNS, the ***l2tp-peer-name*** argument is the name of the LAC.

The name of the L2TP peer that the peer supplies as a hostname in SCCRQ packets.

Use the **remote ip** ***ip-addr*** construct to specify the IP address for the LNS; use the **remote dns** ***dns-name*** construct to specify the DNS name for the LNS. Use the **local** ***ip-addr*** construct to specify the IP address for the LAC.

You can assign an alias for the L2TP peer name with the **domain** command in L2TP peer configuration mode. Peer names, group names, and domain aliases must be unique within the context. For example, if a peer is named “isp”, no other peer, group, or alias can also be named “isp” within the context.

Note: The peer name for the SmartEdge router is frequently the hostname for the SmartEdge router (by default, Redback). If you are configuring a new system, you may want to rename the SmartEdge router. To change the hostname of a SmartEdge router, enter the **system hostname** command in global configuration mode. For more information about this command, see the *Command List*.

Be aware that if the SmartEdge router is currently in service and you change its hostname, any authentication based on the previous definition fails.

Instead of using the SmartEdge router hostname as the peer name, you can create another hostname to use as a peer name; to create another hostname, enter the **local-name** command in L2TP peer configuration mode.

Note: This command supports multiple L2TP tunnels that are identically named. This is commonly the case when Microsoft Windows clients are the L2TP peers.

Use the **unnamed** keyword to configure how the system responds to anonymous peers. Use the anonymous peer configuration for any incoming SCCRQ packets that contain a hostname not found in the local L2TP peer configurations, or for peers configured by a Remote Authentication Dial-In User Service (RADIUS) server.

To configure the parameters for an anonymous L2TP peer, you can use all the L2TP configuration mode commands, except for **domain**. We recommend that you use the **tunnel-auth** command in L2TP configuration mode, to accept all incoming peer requests that contain a specific tunnel password. In addition, we recommend that you restrict the use of this peer to the SmartEdge router using the **function** command in L2TP configuration mode with the **lns-only** keyword. Otherwise, outgoing calls might be placed on anonymous peers.

Use the **no** form of this command to delete the default peer or an existing L2TP peer in the current context.

1.103.6 Examples

The following example shows how to create an L2TP peer, **lac1.net**, in the **local** context:

```
[local]Redback(config-config)#context local
[local]Redback(config-ctx)#l2tp-peer name lac1.net media udp-ip remote ip 10.5.5.5
[local]Redback(config-l2tp)#
```

The following example shows how to create a default L2TP tunnel for tunnels in the **local** context:

```
[local]Redback(config-config)#context local
[local]Redback(config-ctx)#l2tp-peer default
[local]Redback(config-l2tp)#
```

1.104 l2tp proxy-auth

l2tp proxy-auth

{**no** | **default**} **l2tp proxy-auth**

1.104.1 Purpose

Enables proxy authentication for Layer 2 Tunneling Protocol (L2TP) access concentrator (LAC) peers.

1.104.2 Command Mode

Context configuration

1.104.3 Syntax Description

This command has no keywords or arguments.

1.104.4 Default

Proxy authentication is disabled.

1.104.5 Usage Guidelines

Use the **l2tp proxy-auth** command to enable proxy authentication for LAC peers.

Use the **no** or **default** form of this command to disable proxy authentication for LAC peers.

1.104.6 Examples

The following example shows how to enable proxy authentication for LAC peers:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#l2tp proxy-auth
```

1.105 l2tp radius-peer

l2tp radius-peer ***use-server-auth-id***

default l2tp radius-peer ***use-server-auth-id***

1.105.1 Purpose

Enables the Layer 2 Tunneling Protocol (L2TP) daemon to use the Tunnel-Server-Auth-ID (91) Remote Authentication Dial-In User Service (RADIUS) attribute as its peer name if the Tunnel-Assignment-ID (82) RADIUS attribute is not present.

1.105.2 Command Mode

Context configuration

1.105.3 Syntax Description

use-server-auth-id	Specifies that the L2TP daemon use the Tunnel-Server-Auth-ID (91) RADIUS attribute.

1.105.4 Default

The default uses the Tunnel-Assignment-ID (82) RADIUS attribute as the peer name.

1.105.5 Usage Guidelines

Use the **l2tp radius-peer** command with the **use-server-auth-id** keyword to enable the L2TP daemon to use the Tunnel-Server-Auth-ID (91) RADIUS attribute as its peer name if the Tunnel-Assignment-ID (82) RADIUS attribute is not present.

Use the **default** form of this command to return to using the Tunnel-Assignment-ID (82) RADIUS attribute as the peer name.

Note: For modes relevant to RADIUS attributes, see *RADIUS Attributes*.

1.105.6 Examples

The following example shows how to enable the L2TP daemon to use the Tunnel-Server-Auth-ID (91) RADIUS attribute as its peer name:

```
[local]Redback(config)#l2tp radius-peer Tunnel-Assignment-ID
```

1.106 l2tp renegotiate lcp

l2tp renegotiate lcp {**always** | **never** | **on-mismatch**}

{**no** | **default**} **l2tp renegotiate lcp**

1.106.1 Purpose

Specifies the conditions under which the SmartEdge router, when acting as a Layer 2 Tunneling Protocol (L2TP) network server (LNS) renegotiates the Link Control Protocol (LCP) options with an L2TP access concentrator (LAC).

1.106.2 Command Mode

Context configuration

1.106.3 Syntax Description

always	Renegotiates regardless of any LCP or Authentication packets received.
never	Does not ever renegotiate.
on-mismatch	Renegotiates if the received proxy LCP options do not match the configured options. This is the default.

1.106.4 Default

Renegotiates if the received proxy LCP options do not match the configured options.

1.106.5 Usage Guidelines

Use the **l2tp renegotiate lcp** command to specify the conditions under which the SmartEdge router, when acting as an LNS, renegotiates with a LAC.

As part of L2TP session establishment, a LAC might send proxy-lcp and proxy-auth options (LCP and Authentication packets it received from its client) in one of its messages to the SmartEdge router. In this case, the SmartEdge router, acting as an LNS, might receive all the necessary LCP information without negotiating directly with the client. However, if a proxy LCP packet is not received, then the SmartEdge router renegotiates the LCP, depending on the conditions specified by this command.

Use the **always** keyword to support those situations for which renegotiation is required, regardless of the information received from the client.

Use the **never** keyword to support those Point-to-Point Protocol (PPP) clients that cannot successfully establish a session if renegotiation occurs. In this case, the SmartEdge router attempts to use proxy-LCP information as much as possible. That is, it accepts non-critical values, even on mismatch. But it does not tolerate authentication problems or a lack of a proxy LCP.

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

1.106.6 Examples

The following example shows how to specify that no renegotiation take place:

```
[local]Redback(config)#context local
[local]Redback(config)#l2tp renegotiate lcp never
```

1.107 l2tp renegotiate mru

l2tp renegotiate mru ***mru-size***

{**no** | **default**} **l2tp renegotiate mru**

1.107.1 Purpose

Specifies the maximum receive unit (MRU) size used during renegotiation.

1.107.2 Command Mode

Context configuration

1.107.3 Syntax Description

mru-size	Size in bytes of MRU used during renegotiations.

1.107.4 Default

MRU size is to-be-determined bytes.

1.107.5 Usage Guidelines

Use the **l2tp renegotiate mru** command to specify the MRU size in bytes that the SmartEdge router, when acting as an LNS, uses in renegotiations with a LAC.

Use the **no** or **default** form of this command to use the default MRU size.

1.107.6 Examples

The following example shows how to specify that the SmartEdge router in the local context uses 4234 bytes in L2TP renegotiations with a LAC:

```
[local]Redback(config)#context local
[local]Redback(config)#l2tp renegotiate mru 4234
```

1.108 l2tp strict-deadtime

l2tp strict-deadtime

{**no** | **default**} **l2tp deadtime**

1.108.1 Purpose

Enables the strict enforcement of the deadtime, even if all Layer 2 Tunneling Protocol (L2TP) peers are labeled dead.

1.108.2 Command Mode

Context configuration

1.108.3 Syntax Description

This command has no keywords or arguments.

1.108.4 Default

Strict enforcement of the deadtime is disabled.

1.108.5 Usage Guidelines

Use the **l2tp strict-deadtime** command to enable the strict enforcement of the deadtime, even if all L2TP peers are labeled dead. You can use this command to control connection attempts to dead peers that are created by the Remote Authentication Dial-In User Service (RADIUS).

A peer is labeled dead after it is determined that a new tunnel cannot be established to the peer. This feature controls connection requests as follows:

- If strict deadtime is disabled:

When a connection request arrives and all candidate peers for that destination are labeled dead, the SmartEdge router attempts to make a connection to one of the dead peers, even if the deadtime has not expired for any of them.

- If strict deadtime is enabled:

  No connection attempt is made until the deadtime for at least one candidate peer has expired.

Use the **no** or **default** form of this command to disable strict enforcement of the deadtime.

1.108.6 Examples

The following example shows how to enable the strict enforcement of the deadtime for all L2TP peers in the context:

```
[local]Redback(config-ctx)#l2tp strict-deadtime
```

1.109 l2vpn

l2vpn

no l2vpn

1.109.1 Purpose

Enters L2VPN configuration mode.

1.109.2 Command Mode

Context configuration

1.109.3 Syntax Description

This command has no keywords or arguments.

1.109.4 Default

None

1.109.5 Usage Guidelines

Use the **l2vpn** command to enter L2VPN configuration mode.

Note: L2VPNs are supported in the local context only. You cannot enter L2VPN configuration mode in a non-local context.

Use the **no** form of this command to delete all configured Layer 2 Virtual Private Network (L2VPN) cross-connections.

1.109.6 Examples

The following example shows how to change the command mode from **context** configuration to **L2VPN** configuration:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#l2vpn
[local]Redback(config-l2vpn)#
```

1.110 l2vpn (ctx-name)

l2vpn ***ctx-name***

no l2vpn

1.110.1 Purpose

Enables a Layer 2 (L2) circuit for Layer 2 Virtual Private Network (L2VPN) operation.

1.110.2 Command Mode

- ATM PVC configuration
- dot1q PVC configuration
- Frame Relay PVC configuration
- port configuration
- link-group configuration

1.110.3 Syntax Description

ctx-name	Name of the context in which the L2VPN is created.

1.110.4 Default

L2 circuits are not enabled for L2VPN operation.

1.110.5 Usage Guidelines

Use the **l2vpn (ctx-name)** command in any L2 circuit configuration mode to enable an L2 circuit for L2VPN operation.

Note: L2VPNs are supported in the local context only.

The use of this command in link-group configuration mode is restricted to the link-group access type.

Note: Enabling L2VPN operation is supported for on-demand 802.1Q permanent virtual circuits (PVCs), but not for on-demand Asynchronous Transfer Mode (ATM) PVCs.

Use the **no** form of this command to disable L2 circuits for L2VPN operation.

1.110.6 Examples

The following example shows how to enable an ATM PVC for L2VPN operation:

```
[local]Redback(config)#port atm 6/1
[local]Redback(config-atm)#atm pvc 1 101 profile ubr encapsulation bridge1483
[local]Redback(config-atmpvc)#l2vpn local
```

The following example shows how to enables an 802.1Q PVC for L2VPN operation:

```
[local]Redback(config)#port ethernet 3/0
[local]Redback(config-port)#encapsulation dot1q
[local]Redback(config-port)#dot1q pvc 20
[local]Redback(config-dot1q-pvc)#l2vpn local
```

The following example shows how to enable an on-demand 802.1Q PVC for L2VPN operation:

```
[local]Redback(config)#port ethernet 3/0
[local]Redback(config-port)#encapsulation dot1q
[local]Redback(config-port)#dot1q pvc on-demand 20
[local]Redback(config-dot1q-pvc)#l2vpn local
```

The following example shows how to enable a Frame Relay PVC for L2VPN operation:

```
[local]Redback(config)#port pos 3/1
[local]Redback(config-port)#frame-relay pvc 16
[local]Redback(config-frpvc)#l2vpn local
```

The following example shows how to enable an Ethernet port for L2VPN operation:

```
[local]Redback(config)#port ethernet 3/0
[local]Redback(config-port)#l2vpn local
[local]Redback(config-port)#
```

1.111 l2vpn profile

l2vpn profile ***profile-name***

no l2vpn profile ***profile-name***

1.111.1 Purpose

Create a new L2VPN profile or select an existing L2VPN profile and enter L2VPN profile configuration mode.

1.111.2 Command Mode

Global configuration

1.111.3 Syntax Description

profile-name	Name that identifies this L2VPN profile.

1.111.4 Default

None

1.111.5 Usage Guidelines

Use the **l2vpn profile** command to create a new L2VPN profile or select an existing L2VPN profile and enter L2VPN profile configuration mode.

Use the **no** form of this command to delete an L2VPN profile from your system.

1.111.6 Examples

The following example shows how to create a new L2VPN profile called ldp-profile1 and enter L2VPN profile configuration mode:

```
[local]Redback(config)#l2vpn profile ldp-profile1
[local]Redback(config-l2vpn-xc-profile)#
```

1.112 label-action

label-action ***in-label-num*** [**php** ***egress-addr*** | **pop** | **swap** ***out-label-num next-hop-addr***]

no label-action ***in-label-num*** [**php** ***egress-addr*** | **pop** | **swap** ***out-label-num next-hop-addr***]

1.112.1 Purpose

Configures a static Multiprotocol Label Switching (MPLS) label-action mapping.

1.112.2 Command Mode

MPLS static interface configuration

1.112.3 Syntax Description

in-label-num	Number of the incoming label. The range of values is 16 to 1,024.
php	Optional. Penultimate Hop Pop pops (removes) the label before forwarding the IP-only packet from the egress label-switched router (LSR). The egress LSR then forwards the packet based on its destination address.
egress-addr	Optional. IP address of the egress LSR.
pop	Optional. Pops (removes) the top label in the stack and forwards the remaining payload as either a labeled packet, or an unlabeled IP packet.
swap	Optional. Replaces the incoming label with the outgoing label, and forwards to the IP address of the next hop.
out-label-num	Optional. Number of the outgoing label. The range of values is 16 to 1,024.
next-hop-addr	Optional. IP address of the next hop.

1.112.4 Default

None

1.112.5 Usage Guidelines

Use the **label-action** command to configure a static MPLS label-action mapping for the MPLS static interface.

Label actions change the label information for labeled packets as they are forwarded through an LSR. For instance, a label can be removed from a stack of labels, a label can be swapped for another label, or the label can be completely removed from the packet.

Use the **no** form of this command to delete a static MPLS label-action mapping.

1.112.6 Examples

The following example shows how to swap the MPLS label **16** for label **24** and forward the labeled packet to the next hop **10.10.10.2:**

```
[local]Redback(config-ctx)#router mpls-static
[local]Redback(config-mpls-static)#interface isp6
[local]Redback(config-mpls-static-if)#label-action 16 swap 24 10.10.10.2
```

1.113 label-binding

[**neighbor** *ip-addr*] **label-binding prefix-list** *pl-name* {**in** | **out**}

no [**neighbor** *ip-addr*] **label-binding prefix-list** *pl-name* {**in** | **out**}

1.113.1 Purpose

Applies an IP prefix list to filter Label Distribution Protocol (LDP) label advertisements.

1.113.2 Command Mode

LDP router configuration

1.113.3 Syntax Description

neighbor *ip-addr*	Optional. Neighbor IP address. Filters label advertisements to and from the specified neighbor. If this construct is omitted, the prefix list is applied to all neighbors.
prefix-list *pl-name*	Prefix list name. Applies the filters in the specified prefix list to label advertisements. In doing so, restricts label advertisements to or from a forwarding equivalence class (FEC), or set of destinations, that are identified in the prefix list.
in	Applies the prefix list to incoming label advertisements.
out	Applies the prefix list to outgoing label advertisements.

1.113.4 Default

Labels of directly connected interfaces and labels learned from LDP neighbors are advertised.

1.113.5 Usage Guidelines

Use the **label-binding** command to apply an IP prefix list to filter LDP label advertisements.

If the LDP neighbor's transport IP address differs from its router ID, the IP address specified in the **neighbor** *ip-addr* construct must be the LDP neighbor's transport IP address.

A typical application is to apply a prefix list that restricts LDP to advertise labels for only loopback interface IP addresses. Limiting LDP label advertisements to loopback interfaces provides fast and reliable transportation of label binding information, and streamlines the efforts to build LSPs.

To filter label advertisements, you must first configure the IP prefix list through the **ip prefix-list** command in context configuration mode. For more information, see *Configuring Routing Policies*.

Use the **no** form of this command to remove LDP label advertisement filtering.

1.113.6 Examples

The following example shows how to configure the LDP instance running in the **local** context to send LDP label advertisements over loopback interface addresses only:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#ip prefix-list loopback-only
[local]Redback(config-prefix-list)#permit 0.0.0.0/0 eq 32
[local]Redback(config-prefix-list)#exit
[local]Redback(config-ctx)#router ldp
[local]Redback(config-ldp)#label-binding prefix-list loopback-only out
```

1.114 lacp

lacp ***lacp-params***

{**no** | **default**} **lacp** ***lacp-params***

1.114.1 Purpose

Configures the Link Aggregation Control Protocol (LACP) parameters for the link group.

1.114.2 Command Mode

Link group configuration

1.114.3 Syntax Description

lacp-params	LACP parameters for the link group.

1.114.4 Default

LACP parameter defaults is described in Table 4.

1.114.5 Usage Guidelines

Use the **lacp** command to configure the LACP parameters for the link group. This command applies only to access, Ethernet, and 802.1Q link groups. Table 4 lists the LACP parameters for this command.

Table 4 LACP Parameters

Parameter	Description
{**active** | **passive**}	Configures the LACP in active or passive mode. In active mode, LACP starts sending LACP control packets to the peer group port. In passive mode, the LACP starts to exchange LACP packets only after it receives an LACP packet from the partner system. By default, the LACP is not enabled in either mode.
admin-key ***num-value***	Configures the LACP administrative key. The administrative key uniquely identifies the LACP enabled link group. If this is not configured, the system generates a unique administrative key for the link group. The range of values is 32,767 to 65,535.
hold-timeout ***seconds***	Configures the LACP hold down-time. The range of values is 1 to 90 seconds; the default value is 3 seconds.
ignore-system-id	Enables the SmartEdge router to operate as the common endpoint in a multichassis link group configuration by ignoring the system ID of the connected network nodes. If enabled, you must also set the *maximum-links* command to the value `1` for its ***max-active*** argument. Multichassis link aggregation is supported only by the access, Ethernet, and 802.1Q link group types.
passive	
periodic-timeout {**long** | **short**}	Specifies the interval at which the partner system sends the port state information. The **short** timeout exchange interval is 1 second; the **long** timeout exchange interval is 30 seconds; the default behavior is short timeout.
revertible	Specifies revertible behavior for the standby port. The default behavior is revertible.

To set the LACP priority of the SmartEdge system, use the **system lacp priority** command (in global configuration mode).

To view the LACP system priority, MAC address, LACP ID, and other LACP parameters use the **show lacp** command (any mode). To view the LACP configuration, use the **show config**.

Note: If you configure an access link group for LACP, the **qos hierarchical mode strict** command is required on all PPA2 line cards (both economical and noneconomical).

Use the **no** or **default** form of this command to reset the specified parameter to its default condition.

1.114.6 Examples

The following example shows how to configure LACP parameters for the **foo** link group to **active** mode:

```
[local]Redback(config)#link-group foo access
[local]Redback(config-link-group)#lacp active
```

1.115 lacp priority

lacp priority ***priority-num:***

{**no** | **default**} **lacp priority**

1.115.1 Purpose

Specifies the Link Aggregation Control Protocol (LACP) port priority of an Ethernet port when determining the order for aggregation.

1.115.2 Command Mode

Port configuration

1.115.3 Syntax Description

priority-num	Optional. Specifies the LACP port priority of an Ethernet port when determining the order for aggregation. The range of values is 1 to 65535; the default value is 2.

1.115.4 Default

The default value of the LACP port priority on an Ethernet port is 2.

1.115.5 Usage Guidelines

Use the **lacp priority** command to configure an Ethernet port to specify the LACP priority of that port when determining the order for aggregation.

Use either the **no** or **default** form of this command to return to the port to its default behavior.

1.115.6 Examples

The following example shows how to configure Ethernet port **7** in slot **1** to set the LACP priority to **111:**

```
[local]Redback(config)#port ethernet 7/1
[local]Redback(config-port)#lacp priority 111
```

1.116 last-member-query-interval

last-member-query-interval ***interval***

{**no** | **default**} **last-member-query-interval** ***interval***

1.116.1 Purpose

Configures the interval at which the router sends Internet Group Management Protocol (IGMP) group-specific host query messages.

1.116.2 Command Mode

IGMP snooping configuration

1.116.3 Syntax Description

interval	Interval, in milliseconds, at which IGMP group-specific host query messages are sent.

1.116.4 Default

The default last member query interval is 1,000 milliseconds (1 second).

1.116.5 Usage Guidelines

Use the **last-member-query-interval** command to configure the interval at which the router sends IGMP group-specific host query messages.

Use the **no** or **default** form of this command to set the interval to the default value of 1,000 milliseconds.

1.116.6 Examples

The following example shows how to set the last member query interval to **2500** milliseconds (2.5 seconds):

```
[local]Redback(config)#context blue
[local]Redback(config-ctx)#bridge metro
[local]Redback(config-bridge)#igmp snooping
[local]Redback(config-igmp-snooping)#igmp last-member-query-interval 2500
```

1.117 ldp-igp-synchronization

ldp-igp-synchronization [**timeout** ***seconds***]

no ldp-igp-synchronization

1.117.1 Purpose

Enables Label Distribution Protocol (LDP) Interior Gateway Protocol (IGP) synchronization with Intermediate System-to-Intermediate System (IS-IS) or Open Shortest Path First (OSPF) on all interfaces.

1.117.2 Command Mode

- IS-IS router configuration
- IS-IS interface configuration
- OSPF interface configuration
- OSPF router configuration

1.117.3 Syntax Description

timeout ***seconds***	Optional. Sets the maximum time, in seconds, that the interface waits before transporting traffic without receiving LDPs notification that label exchange is completed. For IS-IS, the range of values is 5 to 60. For OSPF, the range of values is 5 to 65535. The **timeout** ***seconds*** construct is not available in IS-IS interface configuration mode or OSPF interface configuration mode.

1.117.4 Default

LDP-IGP synchronization is disabled. If a timeout is not specified, the IGP continues to advertise the maximum metric for a link indefinitely if the IGP and LDP fail to synchronize.

1.117.5 Usage Guidelines

Use the **ldp-igp-synchronization** command to enable LDP-IGP synchronization with IS-IS or OSPF on all interfaces.

LDP establishes the LSPs on the shortest path to a destination as determined by IP forwarding. For the LSP to be established, each link must have an operational adjacency and an operational LDP session, and MPLS label bindings must have been exchanged over each session. Because the LDP protocol cannot itself alert dependent services to an interruption in an LSP, the IGP can route traffic through the link before it is established or after an LDP session has closed; in either case, packet loss can occur. In this release, the SmartEdge router supports LDP-IGP synchronization, which minimizes traffic loss in this scenario. When LDP-IGP synchronization is enabled, the IGP advertises the maximum routing metric for the link until it detects that LDP has converged. After the LSP is established, the IGP advertises the configured metric for the link and the LDP and IGP are considered synchronized.

LDP-IGP synchronization is supported on a per-interface basis for IS-IS and OSPF only. Synchronization is supported on LAN interfaces, provided the LAN interfaces are point-to-point interfaces. Because LDP can be configured in just the local context, only local context IGP instances support LDP-IGP synchronization at this time.

Note: LDP-IGP synchronization is supported in the local-context only.

Although LDP-IGP is enabled for all interfaces, it can be selectively disabled on an interface by using the **no ldp-igp-syncronization** command in router IS-IS interface or router OSPF area interface configuration mode for the interface.

To view LDP-IGP synchronization states, use the **show isis interfaces** command with the **extensive** keyword.

1.117.6 Examples

The following example shows how to configure LDP-IGP synchronization with IS-IS and OSPF with a timeout interval of 35 seconds:

```
[local]Redback(config-ctx)#router isis ip-backbone
[local]Redback(config-isis)#ldp-igp-synchronization timeout 35
[local]Redback(config-isis)#exit
[local]Redback(config-ctx)#router ospf ip-ospf
[local]Redback(config-ospf)#area 0.0.0.0
[local]Redback(config-ospf-area)#ldp-igp-synchronization timeout 35
```

1.118 learning

learning

{**no** | **default**} **learning**

1.118.1 Purpose

Enables the bridge to learn medium access control (MAC) addresses.

1.118.2 Command Mode

Bridge configuration

1.118.3 Syntax Description

This command has no keywords or arguments.

1.118.4 Default

Learning is enabled.

1.118.5 Usage Guidelines

Use the **learning** command to enable the bridge to learn MAC addresses.

Use the **no** or **default** form of this command to disable learning.

1.118.6 Examples

The following example shows how to disable learning for the bridge:

```
[local]Redback(config)#context bridge
[local]Redback(config-ctx)#bridge isp1
[local]Redback(config-bridge)#no learning
```

1.119 level

level ***n***

{**no** | **default**} **level**

1.119.1 Purpose

Begins the configuration of maintenance domain (MD) level ***n***.

1.119.2 Command Mode

CFM configuration

1.119.3 Syntax Description

n	Specifies which MD level is to be configured. Enter a value from 0 to 7.

1.119.4 Default

level 0

1.119.5 Usage Guidelines

Use the **level** command to begin the configuration of MD level ***n***. Since there are eight possible values for the ***n*** argument of this command, there are eight possible MD level configuration modes.

Network customers, service providers, and operators, each view the network at their assigned MD level. Typical maintenance levels reserve the highest levels for customers (users of the end-to-end link), middle levels for service providers (managers of link segments and network edge services), and lowest levels for operators (managers of core bridges and routers).

The following illustration shows a four-level CFM system. Although the customer sees the entire CFM managed segment at the highest MD level,

the maintenance association intermediate points (MIPs) at lower levels are hidden. The triangular-shaped maintenance points are called the maintenance association endpoints (MEPs) and the oval shaped maintenance points are the MIPs:

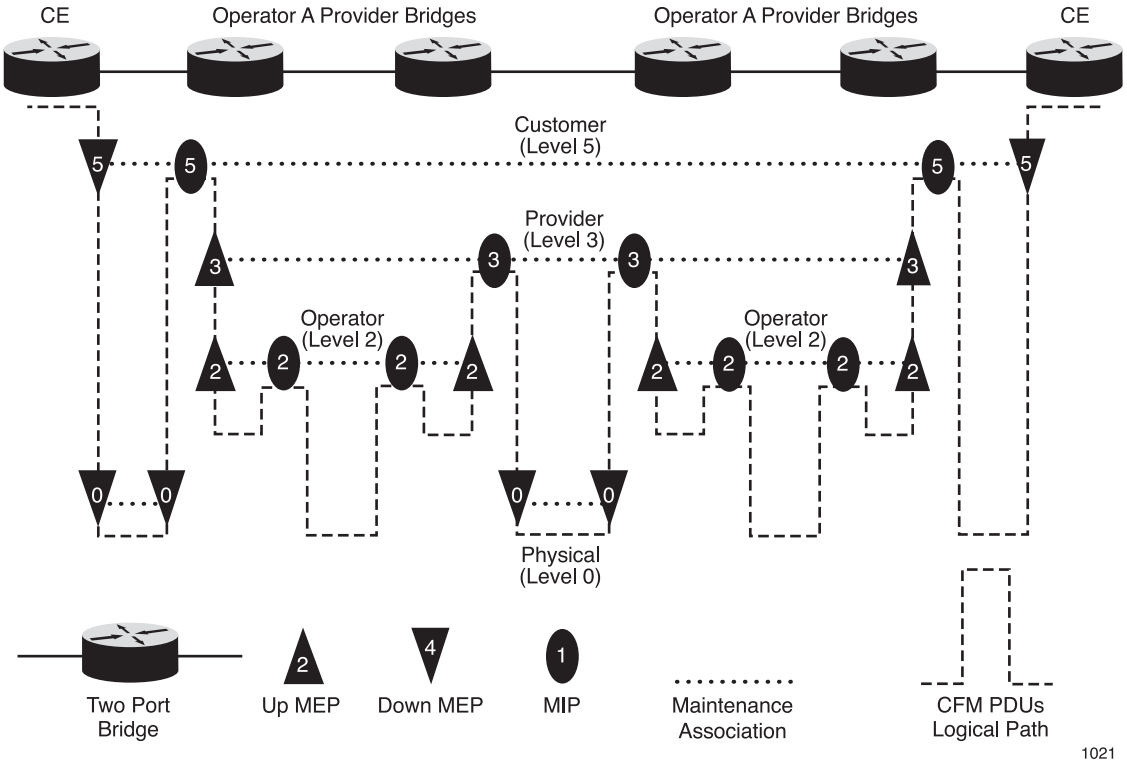

Figure 1 Four-Level CFM system

While a MEP in a lower level can be defined as a MEP or MIP in a higher level MD, a lower level MIP is always hidden from the higher MD levels.

1.119.6 Examples

The following example shows how to use this command to set the MD level to **4:**

```
[local]Redback(config)#ethernet-cfm instance-1
[local]Redback(config-ether-cfm)#level 4
[local]Redback(config-ether-cfm)#domain-name sbc.com
```

1.120 limit

limit ***kilobytes***

default limit

1.120.1 Purpose

Sets a limit on the space that is used to store bulkstats collection files on the SmartEdge router.

1.120.2 Command Mode

Bulkstats configuration

1.120.3 Syntax Description

kilobytes	Amount of space, in KB, used to store bulkstats data. The range of values is 100 to 100,000 KB. The default value is 1,024 KB.

1.120.4 Default

The limit for storing bulkstats data is 1,024 KB (or 1 MB).

1.120.5 Usage Guidelines

Use the **limit** command to set a limit on the space that is used to store bulkstats collection files on the SmartEdge router.

You cannot change the limit size while bulkstats collection is enabled; you must first disable bulkstats collection using the **collection** command in bulkstats configuration mode and then re-enable bulkstats collection after entering the **limit** command.

Caution!

Risk of data loss. If bulkstats collection is re-enabled after a new limit value has been set, data is deleted, and a new collection file is created. To reduce the risk, enter a **bulkstats force transfer** command (in exec mode) for the specified policy prior to disabling bulkstats collection so that all collected data is transferred to the bulkstats file server. For information on the **bulkstats force transfer** command, see the *Command List*.

If data collection fails or if the file size reaches the limit before collection, the oldest data is overwritten, which allows collection to continue with the most recent data saved.

Use the **default** form of this command to set the bulkstats data storage limit to 1,024 KB.

1.120.6 Examples

The following example shows how to limit the space used to store bulkstats data to **4906** KB:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#bulkstats policy bulk
[local]Redback(config-bulkstats)#limit 4906
```

1.121 link-dampening

link-dampening [up ***up-delay*** **down** ***down-delay*** **restart** ***restart-delay*****]**

{no | default} link-dampening

1.121.1 Purpose and Usage Guidelines

Enables subscribers to maintain a steady state on any Asynchronous Transfer Mode (ATM), Ethernet (including Gigabit Ethernet), or Packet over SONET/SDH (POS) port.

This command dampens the link state detection to reduce port flaps.

- *Recommendation:* In Automatic Protection Switching (APS) configurations this command be enabled only on the APS working port to ensure that path alarms do not cause the subscribers to be disconnected.
- *Recommendation:* This command should be applied only on ports configured on a subscriber-facing card.
- This command does not apply to the **shutdown** or **no shutdown** command (in ATM OC, and port configuration mode). Using these commands causes the port to go down immediately.
- Use the **show port** command with the **detail** keyword (in any mode) to display the state of link-dampening for this port.
- If you enter this command without specifying the delay times, the system uses the default values.

  Use the **no** form of this command to disable link-dampening.

  Use the **default** form of this command to configure link dampening with the default delay times.

1.121.2 Command Mode

- ATM OC configuration
- Port configuration

1.121.3 Syntax Description

up ***up-delay***	Delay in milliseconds before the SmartEdge OS declares the port is up. The range of values is 0 to 65535 milliseconds; the default is 10000 milliseconds (10 seconds). A value of 0 disables link dampening for down-to-up transitions.(1)
down ***down-delay***	Delay in milliseconds before the SmartEdge OS declares the port is down. The range of values is 0 to 65535 milliseconds; the default is 2500 milliseconds (2.5 seconds). A value of 0 disables link dampening for up-to-down transitions.(1)
restart ***restart-delay***	Specifies the delay before declaring a port is up after a restart of the system. The range of delay values is 0 to 65535 seconds; the default value is 0 seconds. *Recommendation:* Set the restart time to a value greater than 250 seconds. Typically, it takes 250 seconds or more for the SmartEdge OS to restart.

(1) If the system declares that the port is down, the port-down event is delayed for the configured delay time (`down-delay`* argument), and the subscriber sees no state change for that port. When the port comes back up, the port must be up for the configured delay time (*`up-delay`* argument) before the system declares that the port is up.*

1.121.4 Default

Disabled on all ATM, Ethernet and POS ports. The **no** form of this command disables link dampening. The **default** form of this command configures link dampening with the default delay times described in the *Syntax Description* section of this command description.

1.121.5 Examples

The following example shows how to enable subscribers to maintain a steady state on an Ethernet port with the default values of 10000 milliseconds and 2500 milliseconds for link-dampening up and down delay, respectively:

```
[local]Redback(config)#port ethernet 2/1
[local]Redback(config-port)#link-dampening
```

The following example shows how to disable the **link-dampening** command on an Ethernet port:

```
[local]Redback(config)#port ethernet 2/1
[local]Redback(config-port)#no link-dampening
```

1.122 link-group (BFD)

link-group {**single-session** | **multiple-session**}

no link-group {**single-session** | **multiple-session**}

1.122.1 Purpose

Specifies whether an interface or neighbor that is part of a link aggregation group (LAG) uses single-session or multiple-session Bidirectional Forwarding Detection (BFD).

1.122.2 Command Mode

- BFD interface configuration
- BFD neighbor configuration

1.122.3 Syntax Description

single-session	Specifies single-session BFD.
multiple-session	Specifies multiple-session BFD.

1.122.4 Default

All deployments of BFD over trunk LAGs are multiple session.

1.122.5 Usage Guidelines

Use the **link-group** command to specify whether an interface or neighbor that is part of a LAG uses single-session or multiple-session BFD.

This command applies only to dot1q and Ethernet 802.3ad link group types.

Note: Neighbor configuration takes precedence over an interface configuration. For example, if a neighbor is configured for single session BFD over a LAG, and an interface within that same LAG is configured to support multiple-session BFD, the router uses the neighbor configuration (single-session BFD).

Use the **no** form of this command to return the interface or neighbor to the default setting of multiple-session BFD.

Note: If the remote peer is not reachable over a link group, the router ignores the **link-group** command settings for the interface or neighbor.

1.122.6 Examples

The following example configures an interface to use single-session BFD:

```
[local]Redback #configure
[local]Redback(config)#context local
[local]Redback(config-ctx)#router bfd
[local]Redback(config-bfd)#interface m1
[local]Redback(config-bfd-if)#link-group single-session
```

The following example configures a neighbor to use multiple-session BFD:

```
[local]Redback #configure
[local]Redback(config)#context local
[local]Redback(config-ctx)#router bfd
[local]Redback(config-bfd)#neighbor 2.2.3.10
[local]Redback(config-bfd-nbr)#link-group multiple-session
```

1.123 link-group (Global, DS-1, E1, Port Configuration Modes)

link-group *lg-name* **[access [economical]** | **dot1q** | **ether** | **mp]**

no link-group *lg-name* **[access [economical]** | **dot1q** | **ether** | **mp]**

1.123.1 Purpose

Creates an empty link group and accesses link group configuration mode, or adds a DS-1 channel, clear-channel E1 channel, clear-channel E1 port, Fast Ethernet (FE) port, POS port, or Gigabit Ethernet (GE) port to a link group.

1.123.2 Command Mode

- DS-1 configuration
- E1 configuration
- Global configuration
- Port configuration

1.123.3 Syntax Description

lg-name	Name of the link group.
access	Specifies an access link group for FE or GE ports. Entered only when creating an access link group; omitted when adding an FE or a GE port to an existing link group.
economical	Specifies that the access link group does not maintain replicas of the circuit features of the active ports on the standby ports. In economical operation, the resources used by the standby ports are reduced, although when a standby port becomes active, a small number of packets are lost in the transition.
dot1q	Specifies a link group for FE or GE ports with 802.1Q encapsulation. Entered only when creating an 802.1Q link group; omitted when adding an FE or a GE port to an existing link group.
ether	Specifies a link group for FE or GE ports with IP-over-Ethernet (IPoE) encapsulation. Entered only when creating an Ethernet link group; omitted when adding an FE or a GE port to an existing link group.
mp	Specifies a Multilink PPP (MLPPP) link group for DS-1 channels, clear-channel E1 channels, POS ports on channelized SONET and SDH cards, or clear-channel E1 ports with Point-to-Point Protocol (PPP) encapsulation. Entered only when creating the link group; omitted when adding a channel or port to an existing MLPPP bundle.
bulkstats	Specifies that there is a bulkstats schema profile to associate with the link group.

1.123.4 Default

No link groups exist. No channels or ports are included in a newly created link group.

1.123.5 Usage Guidelines

Creates an empty link group and accesses link group configuration mode, or adds a link to an existing link group.

Do not enter the link group type keyword when adding a link to a MLPPP link group (**mp** keyword).

Use the bulkstats schema command in link group configuration mode to specify the bulkstats schema profile to associate with the link group.

The following channel and port configuration restrictions apply:

- All DS-1 channels, E1 channels, or E1 ports in a link group must be configured on the same traffic card and must have identical configurations.
- All DS-1 channels, E1 channels, POS ports, or E1 ports to be added to a Multilink PPP (MLPPP) bundle must be configured with PPP encapsulation.
- All FE or GE ports in a link group must have identical configurations with the exception of their descriptions.
- FE ports cannot be mixed with GE ports in the same link group, and you cannot mix ports on FE traffic cards unless they are configured with the same speed.
- You cannot mix ports on GE 3 (GE3), GE 1020 (GE1020), or 10 GE (10GE) traffic cards with ports on any other type of GE traffic card in an access link group.

Table 5 lists the types and numbers of ports, channels, or 802.1Q PVCs that you can add to each type of link group.

Table 5 Link Types and Bundle Sizes

Link Group Type	**Aggregated Link Type**	**Maximum Number and Type of Constituent Links**	**Comment**
802.1Q (`dot1q`)	802.1Q PVCs	• 8 FE ports at the same speed • 8 GE ports of any type at the same speed	Ports are added to the link group, not the PVCs. Untagged traffic on a port configured with 802.1Q encapsulation is also aggregated.
Access (`access`)	• FE ports • GE ports	• 8 FE ports at the same speed • 8 GE3 and GE1020 ports of either type • 8 10GE ports of the same type • 8 GE ports of any other type	You can mix GE3 and GE1020 ports, but you cannot mix either of these types with older versions of the GE traffic cards. You cannot mix 10GE ports with any other type of GE traffic card.
Ethernet (`ether`)	• FE ports • GE ports	• 8 FE ports at the same speed • 8 GE ports of any type at the same speed	

Table 5 Link Types and Bundle Sizes

Link Group Type	Aggregated Link Type	Maximum Number and Type of Constituent Links	Comment
MLPPP (**mp**)	DS-1 channels	16 channels	
MLPPP (**mp**)	Clear-channel E1 channels or ports	16 channels or ports	
MLPPP (**mp**)	Channelized SONET/SDH POS ports	• 8 ports • 30 MLPPP bundles per channelized STM-1 port, assuming an average 2 to 3 E1 links per bundle • 150 MLPPP bundles per channelized OC-12 port (assuming an average 2 to 3 T1 links per bundle)	

The following table describes the egress traffic-distribution mechanism and functional restrictions on economical access link groups.

Table 6 Economical Access Link Group Notes and Restrictions

Link Group Feature	Notes and Restrictions
Load balancing and load distribution	When an outer PVC is configured with the **dot1q pvc** command and the **replicate** keyword, the egress traffic of its inner PVCs is distributed among the ports on a per-link basis. When an outer PVC is configured with the **dot1q pvc** command and the **replicate** keyword, the egress traffic of its inner PVCs is distributed among the ports on a per-SPG-ID basis. If the **replicate** keyword is not used, the egress traffic is hashed on the link group at the circuit level; that is, the packets of any circuit egress from a single pseudocircuit on a single port. In a 802.1Q tunnel configuration, the **replicate** keyword is not supported on the configuration of the inner circuit 802.1Q PVC (C-VLAN) in an 802.1Q tunnel and is supported only on the outer circuit 802.1Q PVC tunnel (S-VLAN).

Table 6 Economical Access Link Group Notes and Restrictions

Link Group Feature	Notes and Restrictions
Maximum links	The maximum number of active links (**maximum-links** command) is eight when the link group is configured with the **economical** keyword.
Policing restriction	If an outer PVC is replicated, applying a QoS policy to it is not supported.
Replication restriction	The replicate feature is supported only on the outer PVCs in access link groups; that is, the **dot1q pvc** command with the **replicate** keyword (in link group configuration mode) is supported only if the encapsulation specified is **1qtunnel**.
PPPoE limitation	Some subscriber circuits using PPPoE might be disconnected when an active link in an economical link-group fails or is administratively shutdown. This issue occurs on PPPoE sessions that have not been fully established before the process of switching over to the standby links begins.
Circuit types	Economical access link groups support all circuit types configurable on the SmartEdge router including transport-enabled 802.1Q PVCs.

Use the **no** form of this command to delete the link group, or deleted an FE port or a GE port from a link group.

1.123.6 Examples

1.123.6.1 Create MLPPP Link Group and Bind it to an Interface

The following example shows how to create a link group as an MLPPP link group, `lg-mppp` and bind the link group to an already existing interface, `if-mppp`, interface in the `local` context

```
[local]Redback(config)#link-group lg-mppp mp
[local]Redback(config-link-group)#bind interface if-mppp local
```

1.123.6.2 Add Two DS-1 Channels to an Exisiting MLPPP Link Group

The following lines add two DS-1 channels with PPP encapsulation and associates them with the `lg-mppp` MLPPP link group (created in the preceding example):

```
[local]Redback(config)#port ds1 1/1:1
[local]Redback(config-ds1)#encapsulation ppp
[local]Redback(config-ds1)#no shutdown
[local]Redback(config-ds1)#link-group lg-mppp
[local]Redback(config-ds1)#exit
[local]Redback(config)#port ds1 1/2:1
[local]Redback(config-ds1)#encapsulation ppp
[local]Redback(config-ds1)#no shutdown
[local]Redback(config-ds1)#link-group lg-mppp
[local]Redback(config-ds1)#exit
```

1.123.6.3 Access Link Group Example

The following example shows how to create an access link group with the name `Gretzky`:

```
[local]Redback(config)#link-group Gretzky access
```

1.124 linktrace

`linktrace`

`{no | default} linktrace`

1.124.1 Purpose

The **`no`** and **`default`** forms of this command specify that the maintenance points in the current maintenance domain (MD) do not respond to link-trace messages (LTMs).

1.124.2 Command Mode

CFM configuration

1.124.3 Syntax Description

This command has no keywords or arguments.

1.124.4 Default

Maintenance points respond to LTMs, unless disabled by this command.

1.124.5 Usage Guidelines

Use the **no linktrace** or **default linktrace** command to specify that the maintenance points in the current MD do not respond to LTMs. No LTRs are sent out, but the LTMs are forwarded.

Use the **linktrace** command to enable response.

1.124.6 Examples

In the following example, the **no linktrace** command disables responses to LTMs in the **sbc** CFM instance (**sbc.com** maintenance domain):

```
[local]Redback(config)#ethernet-cfm instance-1
[local]Redback(config-ether-cfm)#level 4
[local]Redback(config-ether-cfm)#no linktrace
```

1.125 listen

listen

{**no** | **default**} **listen**

1.125.1 Purpose

Enables the specified interface to receive and process Routing Information Protocol (RIP) or RIP next generation (RIPng) packets.

1.125.2 Command Mode

- RIP interface configuration
- RIPng interface configuration

1.125.3 Syntax Description

This command has no keywords or arguments.

1.125.4 Default

After RIP or RIPng is enabled on an interface using the interface command (in RIP or RIPng router configuration mode), by default, the interface can listen to and process RIP or RIPng packets; otherwise, it cannot.

1.125.5 Usage Guidelines

Use the **listen** command to enable the specified interface to receive and process RIP or RIPng packets.

Note: This command does not apply to loopback interfaces.

Use the **no** or **default** form of this command to disable the processing of RIP or RIPng packets by an interface.

1.125.6 Examples

The following example shows how to enable the **fe01** interface to receive and process RIP packets:

```
[local]Redback(config-ctx)#router rip rip002
[local]Redback(config-rip)#interface fe01
[local]Redback(config-rip-if)#listen
```

Glossary

link group bundle
Synonym to *link group*. The bundle of constituent links that are members of the link group.

MLPPP
Multilink PPP. An extension to PPP that allows a router to use more than one physical link for communication.

MP
Multilink PPP or Merge Point.

Merge Point: The point at which traffic exits the tail end router of a bypass RSVP LSP.

MLPPP: An extension to PPP that allows a router to use more than one physical link for communication.

PWFQ
Priority weighted fair queuing.