ERICSSON

Commands: r

COMMAND DESCRIPTION

15/190 82-CRA 119 1170/1-V1 Uen J

Contents

1 Command Descriptions

Commands starting with “r” 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 ra

When entered in ND router configuration mode, the syntax is:

ra {**interval** *ra-interval* | **lifetime** *ra-lifetime* | **managed-config** | **other-config** | **suppress**}

{**no** | **default**} **ra** {**interval** | **lifetime** | **managed-config** | **other-config** | **suppress**}

When entered in ND router interface configuration mode, the syntax is:

ra {**enable** | **interval** *ra-interval* | **lifetime** *ra-lifetime* | **managed-config** | **other-config** | **suppress**}

{**no** | **default**} **ra** {**enable** | **interval** | **lifetime** | **managed-config** | **other-config** | **suppress**}

1.1.1 Purpose

Configures options and settings for router advertisement (RA) messages.

1.1.2 Command Mode

- ND router configuration
- ND router interface configuration

1.1.3 Syntax Description

enable	Enables the sending of RA messages for this Neighbor Discovery (ND) router interface. This keyword is not available in ND router configuration mode.
interval *ra-interval*	Optional. RA interval between transmissions (in seconds). The range of values is 5 to 600; the default value is 200 seconds.
lifetime *ra-lifetime*	Optional. RA lifetime (in seconds). The range of values is 30 to 36,000; the default value is 1,800 seconds.
managed-config	Optional. Sets the managed-address configuration flag in RA messages to TRUE; the default value is not set (FALSE).
other-config	Optional. Sets the other-stateful configuration flag in RA messages to TRUE; the default value is not set (FALSE).
suppress	Optional. Specifies that RA messages be suppressed; the default value is not suppressed.

1.1.4 Default

RA messages are not configured for any ND router or ND router interface.

1.1.5 Usage Guidelines

Use the **ra** command to configure options and settings for RA messages. In ND router configuration mode, this command configures RA for all interfaces; in ND router interface mode, it configures RA for this ND router interface. If specified, the interface parameters override the global parameters. Enter this command multiple times to configure more than one parameter.

Use the **no** or **default** form of this command to remove RA messages from the configuration for this ND router or ND router interface.

1.1.6 Examples

The following example shows how to configure RA for this ND router with a retransmission interval of `60` seconds and a lifetime of six minutes (`360` seconds):

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#router nd
[local]Redback(config-nd)#ra interval 60
[local]Redback(config-nd)#ra lifetime 360
```

The following example shows how to suppress RA for the `int1` ND router interface:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#router nd
[local]Redback(config-nd)#interface int1
[local]Redback(config-nd-if)#ra suppress
```

1.2 ra-interval

ra-interval*ra-interval*

default ra-interval

1.2.1 Purpose

Configures the interval between transmissions for Router Advertisement (RA) messages.

1.2.2 Command Mode

ND profile configuration

1.2.3 Syntax Description

ra-interval	RA interval between transmissions (in seconds). The range of values is 0 to 600. A value of 0 disables unsolicited RA messages.

1.2.4 Default

The RA interval default value is 200 seconds.

1.2.5 Usage Guidelines

Use the **ra-interval** command to configure the interval between transmissions for RA messages.

Use the **default** form of this command to return the RA interval to its default value.

1.2.6 Examples

The following example shows how to configure RA for the ND profile **ndprofile7** with a retransmission interval of 90 seconds:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#nd profile ndprofile7
[local]Redback(config-nd-profile)#ra-interval 90
```

1.3 ra-lifetime

ra-lifetime*ra-lifetime*

defaultra-lifetime

1.3.1 Purpose

Configures the lifetime of Router Advertisement (RA) messages.

1.3.2 Command Mode

ND profile configuration

1.3.3 Syntax Description

ra-lifetime	RA lifetime (in seconds). The range of values is 30 to 36,000; the default value is 1,800.

1.3.4 Default

The RA lifetime default value is 1,800 seconds.

1.3.5 Usage Guidelines

Use the **ra-lifetime** command to configure the lifetime of RA messages.

Use the **default** form of this command to return the RA lifetime parameter to its default value.

1.3.6 Examples

The following example configures an RA lifetime of `24000` seconds in the ND profile **ndprofile7**:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#nd profile ndprofile7
[local]Redback(config-nd-profile)#ra-lifetime 24000
```

1.4 ra-managed-config

ra-managed-config

{**no** | **default**} **ra-managed-config**

1.4.1 Purpose

Enables the managed-address configuration flag in Router Advertisement (RA) messages.

1.4.2 Command Mode

ND profile configuration

1.4.3 Syntax Description

This command has no keywords or arguments.

1.4.4 Default

By default, the managed-address configuration flag is disabled.

1.4.5 Usage Guidelines

Use the **ra-managed-config** command to enable the managed-address configuration flag in RA messages. This setting indicates that the IPv6 address is available through the DHCPv6 protocol. If the managed-address

configuration flag is disabled, the IPv6 address may or may not be available through this protocol.

Use the **no** or **default** form of this command to return the managed-address configuration flag to the default value.

1.4.6 Examples

The following example shows how to enable the managed-address configuration flag in RA messages for the ND profile **ndprofile7**:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#nd profile ndprofile7
[local]Redback(config-nd-profile)#ra-managed-config
```

1.5 ra-on-link

ra-on-link

{**no** | **default**} **ra-on-link**

1.5.1 Purpose

Enables the on-link flag for IPv6 prefixes in Router Advertisement (RA) messages.

1.5.2 Command Mode

ND profile configuration

1.5.3 Syntax Description

This command has no keywords or arguments.

1.5.4 Default

By default, the on-link flag for IPv6 prefixes is enabled.

1.5.5 Usage Guidelines

Use the **ra-on-link** command to enable the on-link flag for IPv6 prefixes in RA messages. This setting indicates that the IPv6 prefix advertised in RA messages is associated with an interface on a specific link. If the on-link flag is disabled, the prefix is not associated with an interface on a specific link.

Use the **no** form of this command to disable the on-link flag. Use the **default** form of this command to return the on-link flag to the default value.

1.5.6 Examples

The following example shows how to enable the on-link flag for IPv6 prefixes in RA messages for the ND profile **ndprofile7**:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#nd profile ndprofile7
[local]Redback(config-nd-profile)#ra-on-link
```

1.6 ra-other-config

ra-other-config

{**no** | **default**} **ra-other-config**

1.6.1 Purpose

Enables the other-config flag in Router Advertisement (RA) messages.

1.6.2 Command Mode

ND profile configuration

1.6.3 Syntax Description

This command has no keywords or arguments.

1.6.4 Default

By default, the Other-Config flag is enabled.

1.6.5 Usage Guidelines

Use the **ra-other-config** command to enable the other-config flag in RA messages. This setting indicates that other configuration information is available in the RA messages. If the other-config flag is disabled, other configuration information is not available.

Use the **no** form of this command to disable the other-config flag. Use the **default** form of this command to return the other-config flag to the default value.

1.6.6 Examples

The following example shows how to enable the other-config flag in RA messages for the ND profile **ndprofile7**:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#nd profile ndprofile7
[local]Redback(config-nd-profile)#ra-other-config
```

1.7 radius accounting algorithm

radius accounting algorithm {**first** | **round-robin**}

default radius accounting algorithm

1.7.1 Purpose

Specifies a load-balancing algorithm to use among multiple RADIUS accounting servers.

1.7.2 Command Mode

context configuration

1.7.3 Syntax Description

first	Specifies that the first configured RADIUS server is always queried first.
round-robin	Specifies that RADIUS servers are queried in round-robin fashion.

1.7.4 Default

The SmartEdge router uses the first configured RADIUS server first.

1.7.5 Usage Guidelines

Use the **radius accounting algorithm** command to specify a load-balancing algorithm to use among multiple RADIUS accounting servers.

Use the **default** form of this command to reset the load-balancing algorithm to use the first configured RADIUS server first.

1.7.6 Examples

The following example shows how to set the load-balancing algorithm to `round-robin`:

```
[local]Redback(config-ctx)#radius accounting algorithm round-robin
```

1.8 radius accounting deadtime

`radius accounting deadtime` ***`interval`***

`default radius accounting deadtime`

1.8.1 Purpose

Sets the interval during which the SmartEdge router treats a nonresponsive RADIUS accounting server as “dead.”

1.8.2 Command Mode

context configuration

1.8.3 Syntax Description

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

1.8.4 Default

The waiting interval is five minutes.

1.8.5 Usage Guidelines

Use the **`radius accounting deadtime`** command to set the interval during which the SmartEdge router treats a nonresponsive RADIUS accounting server as “dead.” During the interval, the SmartEdge router tries to reach another RADIUS accounting server; after the interval expires, the SmartEdge router tries again to reach the accounting server. If there is no response, the RADIUS accounting server remains marked as “dead” and the timer is set again to the configured interval.

If you disable this feature (with the 0 value), the SmartEdge router never waits but attempts to reach the server immediately.

Note: You must configure at least one RADIUS accounting server using the **radius accounting server** command (in context configuration mode) prior to entering this command.

Use the **default** form of this command to specify the default interval.

1.8.6 Examples

The following example shows how to set the deadtime interval to 10 minutes:

```
[local]Redback(config-ctx)#radius accounting deadtime 10
```

1.9 radius accounting max-outstanding

radius accounting max-outstanding *requests*

{**no** | **default**} **radius accounting max-outstanding**

1.9.1 Purpose

Modifies the number of simultaneous outstanding accounting requests that can be sent by the SmartEdge router to RADIUS accounting servers.

1.9.2 Command Mode

context configuration

1.9.3 Syntax Description

requests	Number of simultaneous outstanding requests per RADIUS server in the current context. The range of values is 1 to 256.

1.9.4 Default

The number of simultaneous outstanding accounting requests sent by the SmartEdge router is 256.

1.9.5 Usage Guidelines

Use the **radius accounting max-outstanding** to modify the number of simultaneous outstanding accounting requests that can be sent by the SmartEdge router to RADIUS accounting servers.

Use this command if the RADIUS servers cannot handle the default of 256 simultaneous outstanding accounting requests that the SmartEdge router can send to RADIUS accounting servers configured within the context.

Use the **no** or **default** form of this command to reset the maximum number of allowable outstanding requests to 256.

1.9.6 Examples

The following example shows how to limit the number of simultaneous outstanding requests to `128`:

```
[local]Redback(config-ctx)#radius accounting max-outstanding 128
```

1.10 radius accounting max-retries

radius accounting max-retries ***retries***

default radius accounting max-retries

1.10.1 Purpose

Modifies the number of retransmission attempts the SmartEdge router makes to a RADIUS server in the event that no response is received from the server within the timeout period.

1.10.2 Command Mode

context configuration

1.10.3 Syntax Description

retries	Number of times the SmartEdge router retransmits a RADIUS accounting packet. The range of values is 1 to 2,147,483,647; the default value is 3.

1.10.4 Default

The SmartEdge router sends three retransmissions.

1.10.5 Usage Guidelines

Use the **radius accounting max-retries** command to modify the number of retransmission attempts the SmartEdge router makes to a RADIUS accounting server in the event that no response is received from the server within the timeout period.

If an acknowledgment is not received, each successive, configured server is tried (wrapping from the last server to the first, if necessary) until the maximum number of retransmissions is reached.

Use the **default** form of this command to reset the number of retries to 3.

1.10.6 Examples

The following example set the retransmit value to 5:

```
[local]Redback(config-ctx)#radius accounting max-retries 5
```

The following example shows how to reset the retransmit value to the default of 3:

```
[local]Redback(config-ctx)#default radius accounting max-retries
```

1.11 radius accounting send-acct-on-off

radius accounting send-acct-on-off

{**no** | **default**} **radius accounting send-acct-on-off**

1.11.1 Purpose

Enables the sending of “accounting on” and “accounting off” messages to all RADIUS accounting servers that are configured in the current context.

1.11.2 Command Mode

context configuration

1.11.3 Syntax Description

This command has no keywords or arguments.

1.11.4 Default

Accounting on and accounting off messages are sent.

1.11.5 Usage Guidelines

Use the **radius accounting send-acct-on-off** command to enable the sending of accounting on and accounting off messages to all RADIUS accounting servers that are configured in the current context.

The SmartEdge router sends messages to RADIUS accounting servers in various circumstances:

- An Accounting-On message is sent when you enable RADIUS accounting in a context; this message is sent to all RADIUS accounting servers configured within the context. This type of message is also sent when you add a new RADIUS accounting server; however, the message is only sent to the newly added RADIUS accounting server.
- An Accounting-off message is sent when you disable RADIUS accounting within a context; this message is sent to all RADIUS accounting servers configured in a context. If you remove a single RADIUS accounting server, the message is only sent to the newly removed RADIUS accounting server.

  Note: The SmartEdge router attempts to send a single accounting on message when more than one type of RADIUS accounting is enabled. For example, if you enable both subscriber accounting and L2TP accounting, the SmartEdge router sends a single accounting on message to each RADIUS accounting server, even if you enable L2TP accounting at a later time.

  Similarly, the accounting off message is not sent until you have disabled all types of RADIUS accounting.

Use the **no** form of this command to prevent the SmartEdge router from sending these messages.

Use the **default** form of this command to return the system to its default behavior.

1.11.6 Examples

The following example shows how to disable the sending of accounting on and off messages to all other RADIUS accounting servers in the `local` context:

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

[local]Redback(config-ctx)#no radius send-acct-on-off
```

1.12 radius accounting server

radius accounting server {*ip-addr* | *hostname*} {**key** *key* | **encrypted-key** *key*} [{**oldports** | **port** *udp-port*}] [**max** *requests*]

no radius accounting server

1.12.1 Purpose

Configures the IP address or hostname of a RADIUS accounting server.

1.12.2 Command Mode

context configuration

1.12.3 Syntax Description

ip-addr	IP address of the RADIUS accounting server.
hostname	Hostname of the RADIUS accounting server. Domain Name System (DNS) must be enabled to use the *hostname* argument.
key *key*	Authentication key used when communicating with the accounting server.
encrypted-key *key*	Alphanumeric string representing the encrypted authentication key used when communicating with the RADIUS accounting server.
oldports	Optional. Designates the old RADIUS User Datagram Protocol (UDP) port 1646.
port *udp-port*	Optional. RADIUS accounting UDP port. The range of values is 1 to 65,536; the default value is 1813.

1.12.4 Default

RADIUS accounting server hostnames and IP addresses are not preconfigured. The UDP accounting port is 1813.

1.12.5 Usage Guidelines

Use the **radius accounting server** command to configure the IP address or hostname of a RADIUS accounting server. Use this command multiple times to configure up to five RADIUS accounting servers per context. To use the *hostname* argument, you must enable DNS; for more information.

Note: To enable accounting to be performed by RADIUS, you must also enter the **aaa accounting subscriber** command (in context configuration mode); for more information, see *Configuring Bridging*.

Use the **no** form of this command to delete a previously configured RADIUS accounting server.

1.12.6 Examples

The following example shows how to configures a RADIUS accounting server IP address of `10.3.3.3` with the key, `secret`, using port `4445` for accounting:

```
[local]Redback(config-ctx)#radius accounting server 10.3
.3.3 key secret port 4445
```

1.13 radius accounting server-timeout

radius accounting server-timeout ***interval***

default radius accounting server-timeout

1.13.1 Purpose

Sets the time interval the SmartEdge router waits before marking a nonresponsive RADIUS accounting server as “dead.”

1.13.2 Command Mode

context configuration

1.13.3 Syntax Description

interval	Time period that the SmartEdge router checks back for successful responses, after an individual RADIUS request times out, before treating the accounting server as “dead.” The range of values is 0 to 2147483647 seconds; the default value is 60 seconds.

1.13.4 Default

The maximum time interval is 60 seconds.

1.13.5 Usage Guidelines

Use the **radius accounting server-timeout** command to set the time interval the SmartEdge router waits before marking a non-responsive RADIUS accounting server as “dead.”

The SmartEdge router marks a RADIUS accounting server as “dead” when no response is received to any RADIUS requests during the time period specified by the ***interval*** argument. Setting the value to 0 disables this feature; in this case, no RADIUS accounting server is marked as “dead.”

Use the **default** form of this command to specify the default interval.

1.13.6 Examples

The following example shows how to set the waiting interval to 80 seconds:

```
[local]Redback(config-ctx)#radius accounting server-timeout 80
```

1.14 radius accounting timeout

radius accounting timeout ***timeout***

default radius accounting timeout

1.14.1 Purpose

Sets the maximum time the SmartEdge router waits for a response from a RADIUS accounting server before assuming that a packet is lost, or that the RADIUS accounting server is unreachable.

1.14.2 Command Mode

context configuration

1.14.3 Syntax Description

timeout	Timeout period in seconds. The range of values is 1 to 2147483647; the default value is 10 seconds.

1.14.4 Default

The maximum time is 10 seconds.

1.14.5 Usage Guidelines

Use the **radius accounting timeout** command to set the maximum time the SmartEdge router waits for a response from a RADIUS accounting server before assuming that a packet is lost, or that the RADIUS accounting server is unreachable.

Use the **default** form of this command to specify the default interval.

1.14.6 Examples

The following example shows how to set the timeout interval to `30` seconds:

```
[local]Redback(config-ctx)#radius accounting timeout 30
```

1.15 radius algorithm

radius algorithm {**first** | **round-robin**}

default radius algorithm

1.15.1 Purpose

Specifies the algorithm to use among multiple RADIUS servers.

1.15.2 Command Mode

context configuration

1.15.3 Syntax Description

first	Specifies that the first configured RADIUS server is always queried first.
round-robin	Specifies that the RADIUS servers are queried in round-robin fashion, enabling load balancing.

1.15.4 Default

The SmartEdge router queries the first configured server first.

1.15.5 Usage Guidelines

Use the **radius algorithm** command to specify the algorithm to use among multiple RADIUS servers.

Use the **default** form of this command to reset the SmartEdge router to query the first configured RADIUS server first.

1.15.6 Examples

The following example shows how to set the algorithm to `round-robin`:

```
[local]Redback(config-ctx)#radius algorithm round-robin
```

1.16 radius attribute acct-delay-time

radius attribute acct-delay-time

{**no** | **default**} **radius attribute acct-delay-time**

1.16.1 Purpose

Sends the Acct-Delay-Time attribute in the RADIUS Accounting-Request packets for the current context regardless of whether the SmartEdge router had a delay in sending the accounting record to the RADIUS server.

1.16.2 Command Mode

context configuration

1.16.3 Syntax Description

This command has no keywords or arguments.

1.16.4 Default

The Acct-Delay-Time attribute is only sent in RADIUS Accounting-Request packets for the current context, if there is a delay in sending the accounting record.

1.16.5 Usage Guidelines

Use the **radius attribute acct-delay-time** command to send the Acct-Delay-Time attribute in RADIUS Accounting-Request packets for the

current context regardless of whether the SmartEdge router had a delay in sending the accounting record to the RADIUS server. If there is no delay, the SmartEdge router sets the Acct-Delay-Time attribute to 0. By default, the Acct-Delay-Time attribute is sent in RADIUS Accounting-Request packets for the current context only if there is a delay in sending the accounting record to the RADIUS server.

Standard RADIUS attribute 41, Acct-Delay-Time, is described in *RADIUS Attributes*.

Use the **no** or **default** form of this command to reset the SmartEdge router behavior to the default condition.

1.16.6 Examples

The following example shows how to configure the SmartEdge router to send the Acct-Delay-Time attribute in RADIUS Accounting-Request packets:

```
[local]Redback(config-ctx)#radius attribute acct-delay-time
```

1.17 radius attribute acct-session-id

radius attribute acct-session-id access-request

{**no** | **default**} **radius attribute acct-session-id access-request**

1.17.1 Purpose

Sends the Acct-Session-Id attribute in RADIUS Access-Request packets for the current context.

1.17.2 Command Mode

context configuration

1.17.3 Syntax Description

access-request	Specifies that the attribute is to be sent in Access-Request packets.

1.17.4 Default

The Acct-Session-Id attribute is only sent in Accounting-Request packets.

1.17.5 Usage Guidelines

Use the **radius attribute acct-session-id** command to send the Acct-Session-Id attribute in RADIUS Access-Request packets for the current context.

This command affects only subscriber sessions, not administrator sessions.

Standard RADIUS attribute 41, Acct-Session-Id, is described in *RADIUS Attributes*.

Use the **no** or **default** form of this command to disable the sending of the Acct-Session-Id attribute in Access-Request packets.

1.17.6 Examples

The following example shows how to configure the SmartEdge router to send the Acct-Session-Id attribute in RADIUS `access-request` packets:

```
[local]Redback(config-ctx)#radius attribute acct-ses
sion-id access-request
```

1.18 radius attribute acct-terminate-cause remap

radius attribute acct-terminate-cause remap

{**no** | **default**} **radius attribute acct-terminate-cause remap**

1.18.1 Purpose

Enables the remapping of Ericsson AB account termination error codes and accesses terminate error cause configuration mode.

1.18.2 Command Mode

global configuration

1.18.3 Syntax Description

This command has no keywords or arguments.

1.18.4 Default

Remapping of account termination error codes is disabled.

1.18.5 Usage Guidelines

Use the **radius attribute acct-terminate-cause remap** command to enable the remapping of Ericsson AB account termination error codes and access terminate error cause configuration mode. By default, the SmartEdge router maps a Ericsson AB termination error code to a RADIUS Attribute 49 (Acct-Terminate-Cause) terminate cause error code, which it sends in RADIUS Accounting-Stop packets. RADIUS attribute 49 terminate cause error codes and their definitions are included in RFC 2866, ***RADIUS Accounting***. *RADIUS Attributes* lists the default mapping of Ericsson AB account termination error codes to RADIUS attribute 49 error codes.

Use the **no** or **default** form of this command to remove the remapping of all Ericsson AB account termination error codes.

1.18.6 Examples

The following example shows how to enable the remapping of Ericsson AB account termination error codes:

```
[local]Redback(config)#radius attribute acct-terminate-cause remap

[local]Redback(config-term-ec)#
```

1.19 radius attribute acct-tunnel-connection l2tp-call-serial-num

radius attribute acct-tunnel-connection l2tp-call-serial-num

{**no** | **default**} **radius attribute acct-tunnel-connection l2tp-call-serial-num**

1.19.1 Purpose

Sends a Layer 2 Tunneling Protocol (L2TP) call serial number type value in the Acct-Tunnel-Connection attribute in RADIUS packets for the current context, when the SmartEdge router is functioning as an L2TP access concentrator (LAC).

1.19.2 Command Mode

context configuration

1.19.3 Syntax Description

This command has no keywords or arguments.

1.19.4 Default

When functioning as a LAC, the SmartEdge router sends an L2TP session ID type value in the Acct-Tunnel-Connection attribute.

1.19.5 Usage Guidelines

Use the **radius attribute acct-tunnel-connection l2tp-call-serial-num** command to send an L2TP call serial number ID type value in the Acct-Tunnel-Connection attribute in the RADIUS packets for the current context, when the SmartEdge router is functioning as a LAC. This enables the RADIUS server to correlate the ID type values received from the SmartEdge router and those received from L2TP network server (LNS) devices when it attempts to authenticate Point-to-Point Protocol over Ethernet (PPPoE) sessions. (LNS) devices send L2TP call serial numbers in the Acct-Tunnel-Connection attribute by default.)

This command affects only subscriber sessions, not administrator sessions.

Standard RADIUS attribute 68, Acct-Tunnel-Connection, is described in *RADIUS Attributes*.

Use the **no** or **default** form of this command to remove a tunnel with the RADIUS server from either a LAC or LNS.

1.19.6 Examples

The following example shows how to configure the SmartEdge router, when functioning as a LAC, to send the L2TP call serial number in the Acct-Tunnel-Connection attribute to the RADIUS server:

```
[local]Redback(config-ctx)#radius attribute acct-tunnel-c
onnection l2tp-call-serial-num
```

1.20 radius attribute calling-station-id

To specify the format for the automatically generated ID string, use the following syntax:

radius attribute calling-station-id {[{**media atm** | **media eth**}] **format** {***agent*** | **description** | **hostname** ***agent*** | **slot-port** ***agent***}

no radius attribute calling-station-id {[{**media atm** | **media eth**}] **format**

default radius attribute calling-station-id

To specify that a separator character be prepended to the Calling-Station-Id attribute string in RADIUS packets, use the following syntax:

radius attribute calling-station-id prepend-separator

[**no** | **default**] **radius attribute calling-station-id prepend-separator**

To pad the virtual path identifier (VPI) or virtual channel identifier (VCI) value with zeros to make a 4-character string, use the following syntax:

radius attribute calling-station-id pvc-pad

[**no** | **default**] **radius attribute calling-station-id pvc-pad**

To use a character that separates the elements of the attribute string, use the following syntax:

radius attribute calling-station-id separator *separator*

[**no** | **default**] **radius attribute calling-station-id**

1.20.1 Purpose

Using the specified format, sends the Calling-Station-Id attribute in RADIUS Access-Request and Accounting-Request packets for the current context.

1.20.2 Command Mode

context configuration

1.20.3 Syntax Description

agent	**agent-circuit-id** [**non-ascii**] [**agent-remote-id** [**non-ascii**]] | **agent-remote-id** [**non-ascii**] The **non-ascii**, **agent-circuit-id**, and **agent-remote-id** keywords are described separately in this table.
media atm	Uses the Asynchronous Transfer Mode (ATM) media format for the automatically generated Calling-ID string.

media eth	Uses the Ethernet media format for the automatically generated Calling-ID string.
format	Indicates a particular format to be applied.
agent-circuit-id	Specifies that the format or the type of the information for the Calling-Station-Id attribute is the circuit agent ID. Optional only when specifying the **slot-port** keyword.
agent-remote-id	Optional. Specifies that the format or the type of the information for the Calling-Station-Id attribute is Agent-Remote-Id. Optional only when specifying the **agent-circuit-id** keyword.
description	Specifies a circuit description format using the information configured with the **description** command in the configuration mode for the circuit with the hostname prepended to it.
hostname	Prepends the SmartEdge router hostname to the contents of the Calling-Station-Id attribute in RADIUS packets. The hostname is either the one that has been configured using the **system hostname** command (in context configuration mode), or the default hostname, “Redback”.
non-ascii	Available in context configuration mode. Specifies one of the following translations when you use RADIUS with option 82: • The agent circuit ID is translated into binary format: **agent-circuit-id** [**non-ascii**] • The agent remote ID is translated into binary format: **agent-remote-id** [**non-ascii**]. • The agent circuit ID and agent remote IDs are both translated into binary format **agent-circuit-id** [**non-ascii**] [**agent-remote-id** [**non-ascii**]]. The default translation is the agent circuit ID and agent remote ID into hexadecimal format.
slot-port	Specifies a slot number/port number format that has the hostname prepended to it.
prepend-separator	Optional. Specifies that a separator character be prepended to the Calling-Station-Id attribute string in RADIUS packets. The separator character to append depends on which character is used for the **separator** keyword.

pvc-pad	Pads the virtual path identifier (VPI)/virtual channel identifier (VCI) value with zeros to make a 4-character string.
separator ***separator***	Character that separates the elements of the attribute string. The default separator character is the number symbol (**#**). You can change this default.

1.20.4 Default

The Calling-Station-Id attribute is not sent.

1.20.5 Usage Guidelines

Use the **radius attribute calling-station-id** command to send the Calling-Station-Id attribute, using the specified format, in RADIUS Access-Request and Accounting-Request packets for the current context.

If you specify the **media** keyword, you can customize the format for ATM or Ethernet subscribers or for both. The default format is valid for all circuit types.

If you specify the **agent-circuit-id** keyword, you can also specify the **agent-remote-id** keyword.

If you specify the **agent-circuit-id non-ascii** keywords, you can also specify the **agent-remote-id non-ascii** keywords.

For Dynamic Host Configuration Protocol (DHCP) clients, the information for the Calling-Station-Id attribute is extracted from the suboption1 information in option 82 of the DHCP request packet; for Point-to-Point Protocol over Ethernet (PPPoE) clients, the information is extracted in the PPPoE Active Discovery Request (PADR) packet.

If the **agent-circuit-id** keyword is specified, but the circuit agent ID information is not present in the DHCP request packet or in the PADR packet sent by the client, the SmartEdge router inserts the “Agent-Circuit-Id Not Present” string.

If the **agent-remote-id** keyword is specified, but the remote agent ID information is not present in the DHCP request packet or in the PADR packet sent by the client, the SmartEdge router inserts the “Agent-Remote-Id Not Present” string.

For Asynchronous Transfer Mode (ATM) permanent virtual circuits (PVCs), the format for the **slot-port** keyword is `#Hostname#slot/port#VPI#VCI`; the **description** format is `#Hostname#VC description#VPI#VCI`.

Note: If the **description** keyword is used, but the description of the ATM PVC itself has not been configured using the **description** command (in ATM PVC configuration mode), the SmartEdge router defaults to the **slot-port** format.

For virtual LANs (VLANs), the formats for the **slot-port** keyword and **description** keyword, respectively, are:

- `#Hostname#slot/port#Vlan-ID`
- `#Hostname#Vlan description#Vlan-ID`

Use the **no** form of this command to disable the sending of the Calling-Station-Id attribute.

Use the **default** form of this command to specify the default separator. To change the default separator character, specify the separator keyword and character to use as the separator.

1.20.6 Examples

The following example sends the Calling-Station-Id attribute using the `slot-port` format and inserts `agent-circuit-id` and `agent-remote-id` information into Access-Request and Accounting-Request packets:

```
[local]Redback(config-ctx)#radius attribute calling-station-id format
slot-port agent-circuit-id agent-remote-id separator #
```

The format in which the Calling-Station-Id attribute is sent for VLAN connections is as follows:

```
hostname#slot#port#(VLAN ID)#(Agent-Circuit-Id)#(Agent-Remote-Id)
```

The following example configures the context so that the Calling-Station-Id attribute is sent in Access-Request and Accounting-Request packets using a slash (/) as the separator character:

```
[local]Redback(config-ctx)#radius attribute calling-
station-id separator /
```

1.21 radius attribute filter-id

radius attribute filter-id direction {**in** | **out** | **both** | **none**}

{**no** | **default**} **radius attribute filter-id**

1.21.1 Purpose

Specifies the behavior of the SmartEdge router when it receives a RADIUS Filter-Id attribute that does not specify a direction and there is an access control list (ACL) applied to the circuit.

1.21.2 Command Mode

context configuration

1.21.3 Syntax Description

direction	Specifies the direction of the packets to which the ACL is applied.
in	Applies the ACL to inbound packets only.
out	Applies the ACL to outbound packets only.
both	Applies the ACL to inbound and outbound packets.
none	Ignores the Filter-Id attribute and does not apply the ACL to packets in either direction.

1.21.4 Default

If the Filter-Id attribute does not include a direction, the SmartEdge router applies the ACL to outbound packets only.

1.21.5 Usage Guidelines

Use the **radius attribute filter-id** command to specify the behavior of the SmartEdge router when it receives a RADIUS Filter-Id attribute that does not specify a direction and there is an ACL applied to the circuit. The choice of behavior depends on the nature of the ACL and the type of data that is exchanged.

The following sequence determines how the SmartEdge router applies the ACL:

- If the Filter-Id attribute includes a direction, it is honored.
- If the Filter-Id attribute does not include a direction, and you have configured this command, the SmartEdge router determines the direction from the configuration for this command.
- If the Filter-Id attribute does not include a direction, and this command is not configured, the SmartEdge router applies the ACL to outbound packets only (the default condition).

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

1.21.6 Examples

The following example shows how to specify that the ACL be applied to inbound packets only:

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

[local]Redback(config-ctx)#radius attribute filter-id in
```

1.22 radius attribute nas-identifier

radius attribute nas-identifier *arbitrary-string*

{**no** | **default**} **radius attribute nas-identifier***arbitrary-string*

1.22.1 Purpose

Includes the network access server (NAS)-Identifier attribute in RADIUS Access-Request and Accounting-Request packets sent by the SmartEdge router.

1.22.2 Command Mode

context configuration

1.22.3 Syntax Description

arbitrary-string	Indicates the value for the NAS system. Alphanumeric string of up to 255 characters.

1.22.4 Default

The NAS-Identifier attribute is not sent.

1.22.5 Usage Guidelines

Use the **radius attribute nas-identifier** command to include the NAS-Identifier attribute in RADIUS Access-Request and Accounting-Request packets sent by the SmartEdge router.

Standard RADIUS attribute 32, NAS-Identifier, is described in *RADIUS Attributes*

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

1.22.6 Examples

The following example shows how to configure the NAS-Identifier in RADIUS Access-Request and Accounting-Request packets sent by the SmartEdge router:

```
[local]Redback(config-ctx)#radius attribute nas-identifi
er somearbritrarystring
```

1.23 radius attribute nas-ip-address

radius attribute nas-ip-address interface ***if-name***

{**no** | **default**} **radius attribute nas-ip-address**

1.23.1 Purpose

Includes the network access server (NAS)-IPv4 Address attribute in RADIUS Access-Request and Accounting-Request packets sent by the SmartEdge router.

1.23.2 Command Mode

context configuration

1.23.3 Syntax Description

interface ***if-name***	Interface name. Uses the primary IPv4 address associated with the interface as the source IPv4 address sent in RADIUS packets. If the interface is not configured or is unreachable, the IPv4 address of the outgoing interface is used instead as the source IPv4 address for packets.

1.23.4 Default

The NAS-IPv4 Address attribute is not sent.

1.23.5 Usage Guidelines

Use the **radius attribute nas-ip-address** command to include the IPv4 NAS-IP-Address attribute in RADIUS Access-Request and Accounting-Request packets sent by the SmartEdge router.

Standard RADIUS attribute 4, NAS-IP-Address, is described in *RADIUS Attributes*.

Use the **no** or **default** form of this command to reset the SmartEdge router behavior so that the NAS-IP-Address attribute for IPv4 is not included.

1.23.6 Examples

The following example shows how to send the primary IPv4 address for interface `ether21` as the source IPv4 address in RADIUS Access-Request and Accounting-Request packets sent by the SmartEdge router:

```
[local]Redback(config-ctx)#radius attribute nas-ip-add
ress interface ether21
```

1.24 radius attribute nas-port

radius attribute nas-port format {**agent-remote-id** | **physical** | **slot-port** | **session-info**} [**no-pseudo**]

{**no** | **default**} **radius attribute nas-port format**

1.24.1 Purpose

Modifies the format of the network access server (NAS)-Port attribute, which is sent in RADIUS Access-Request and Accounting-Request packets for the current context.

1.24.2 Command Mode

context configuration

1.24.3 Syntax Description

format	Indicates a particular attribute string format is to be applied.
agent-remote-id	Specifies that the content of the NAS-Port attribute is a 32-bit remote agent ID.

physical	Provides slot, port, virtual path identifier (VPI), and virtual channel identifier (VCI) in the NAS-Port attribute sent to the RADIUS server. For Asynchronous Transfer Mode (ATM) circuits and PPP over Ethernet (PPPoE) over ATM sessions, the attribute format is ***slot-port-vpi-vci***, such that: • ***slot***—***SSSS*** (4 bits) • ***port***—***PPPP*** (4 bits) • ***vpi***—***CCCCCCCC*** (8 bits) • ***vci***—***CCCCCCCCCCCCCCCC*** (16 bits) For Ethernet and virtual LAN (VLAN) circuits, the attribute format depends on whether the session is connected through an untagged Ethernet port, a VLAN, or a stacked VLAN circuit: For untagged Ethernet, the format is ***slot/port:unused***, such that: • ***slot***—***SSSS*** (4 bits) • ***port***—***PPPP*** (4 bits) • ***unused***—***XXXXXXXXXXXXXXXXXXXXXXXX*** (24 bits) For VLAN circuits, the format is ***slot***/***port***:***vlan-id***, such that: • ***slot***—***SSSS*** (4 bits) • ***port***—***PPPP*** (4 bits) • ***zero***—***000000000000*** (12 bits) • ***vlan-id***—***CCCCCCCCCCCC*** (12 bits) For Stacked VLAN circuits, the format is slot/port:***SvlanID-CvlanID***, such that: • ***slot***—***SSSS*** (4 bits) • ***port***—***PPPP*** (4 bits) • ***SvlanID***—***SSSSSSSSSSSS*** (12 bits) • ***CvlandID***—***CCCCCCCCCCCC*** (12 bits)

slot-port	Provides slot, port, and channel information in the NAS-Port attribute sent to the RADIUS server. The attribute format is ***slot-port-channel***, such that: • ***slot***—***SSSSSSSS*** (8 bits) • ***port***—***PPPPPPPP*** (8 bits) • ***channel***—***CCCCCCCCCCCCCCCC*** (16 bits) If no channel exists, the ***channel*** argument contains zeros. This is the default format for standard RADIUS attribute 5, NAS-Port.
session-info	Provides slot, port, and session information in the NAS-Port attribute sent to the RADIUS server. For ATM circuits, the attribute format is ***slot-port-vpi-vci***, such that: • ***slot***—***SSSS*** (4 bits) • ***port***—***PPPP*** (4 bits) • ***vpi***—***CCCCCCCC*** (8 bits) • ***vci***—***CCCCCCCCCCCCCCCC*** (16 bits) For PPPoE over ATM, Ethernet, and VLAN circuits, the format is ***slot-port-unused-pppoe_session***, such that: • ***slot***—***SSSS*** (4 bits) • ***port***—***PPPP*** (4 bits) • ***unused***—***XXXXXXXX*** (8 bits) • ***session***—***CCCCCCCCCCCCCCCC*** (16 bits)
no-pseudo	Enables formatting for sessions that are not Layer 2 Tunneling Protocol (L2TP) network server (LNS) or L2TP access concentrator (LAC) sessions.

1.24.4 Default

Standard RADIUS attribute 5, NAS-Port, is sent in the slot-port format. L2TP circuits (LNS or LAC), use “pseudo” formatting.

1.24.5 Usage Guidelines

Use the **radius attribute nas-port** command to modify the format of the NAS-Port attribute, which is sent in RADIUS Access-Request and Accounting-Request packets for the current context.

Use the **radius attribute nas-port** command with the **no-pseudo** keyword to remove “pseudo” formatting on L2TP circuits (LNS or LAC).

The standard RADIUS attribute 5, NAS-Port, is described in *RADIUS Attributes*.

Use the **no** or **default** form of this command to revert to the default behavior.

1.24.6 Examples

The following example shows how to send the attribute NAS-Port using the `slot-port` format in RADIUS Access-Request and Accounting-Request packets for the `local` context:

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

[local]Redback(config-ctx)#radius attribute nas-port format slot-port
```

1.25 radius attribute nas-port-id

radius attribute nas-port-id {**format** {**agent-circuit-id** [**agent-remote-id**] | **all** | **hostname** {**agent-circuit-id** [**agent-remote-id**]} | **physical** | **agent-remote-id**} | **modified-agent-circuit-id [prefix-lg-description]** | **prepend-separator** | **separator** ***separator***}

no radius attribute nas-port-id format

default radius attribute nas-port-id {**format** | **separator** ***separator***}

1.25.1 Purpose

Modifies the format of the network access server (NAS)-Port-Id attribute, which is sent in RADIUS Access-Request and Accounting-Request packets for the current context.

1.25.2 Command Mode

context configuration

1.25.3 Syntax Description

format	Indicates a particular format to be applied.
agent-circuit-id	Specifies that the format or the type of the information for the NAS-Port-Id attribute is the circuit agent ID.
agent-remote-id	Optional. Specifies that the format or the type of the information for the Calling-Station-Id attribute is Agent-Remote-Id. Optional only when specifying the **agent-circuit-id** keyword.
hostname	Prepends the SmartEdge router hostname to the contents of the NAS-Port-Id attribute in RADIUS packets. The hostname is either the one that has been configured using the **system hostname** command (in context configuration mode), or the default hostname, “Redback”.
all	Specifies a format that includes the physical circuit and session information. This is the default format.
physical	Specifies a format that includes the physical circuit only.
modified-agent-circuit-id	Specifies that the format or the type of the information for the NAS-Port-Id attribute is a modified form of the circuit agent ID.
prefix-lg-description	Optional. Specifies that a text string description of the access link group is to be used as a prefix to the NAS-Port-Id attribute.
prepend-separator	Optional. Specifies that a separator character be prepended to the NAS-Port-Id attribute string in RADIUS packets. The separator character to append depends on which character is used for the **separator** keyword.
separator *separator*	Character to use to separate the elements of the attribute string. The default separator character is the number symbol (**#**). You can change this default.

1.25.4 Default

Standard RADIUS attribute 87, NAS-Port-Id, is sent using the **all** format.

1.25.5 Usage Guidelines

Caution!

Use the **radius attribute nas-port-id** command to modify the format of the NAS-Port-Id attribute, which is sent in RADIUS Access-Request and Accounting-Request packets for the current context.

Risk of interoperability loss. The NetOp Policy Manager (PM) requires the default format setting for this command to assimilate the RADIUS attribute information. To avoid loss of interoperability with NetOp PM, use this command with its default setting only.

If you specify the **agent-circuit-id** keyword, you can also specify the **agent-remote-id** keyword.

For Dynamic Host Configuration Protocol (DHCP) clients, the information for the NAS-Port-Id attribute is extracted from the suboption1 information in option 82 of the DHCP request packet; for Point-to-Point Protocol over Ethernet (PPPoE) clients, the information is extracted in the PPPoE Active Discovery Request (PADR) packet.

If the **agent-circuit-id** keyword is specified, but the circuit agent ID information is not present in the DHCP request packet or in the PADR packet sent by the client, the SmartEdge router inserts the “Agent-Circuit-Id Not Present” string.

If the **agent-remote-id** keyword is specified, but the remote agent ID information is not present in the DHCP request packet or in the PADR packet sent by the client, the SmartEdge router inserts the “Agent-Remote-Id Not Present” string.

If you specify the **all** keyword, the physical circuit information includes the slot, port, circuit identifier, and session identifier; the format in which the NAS-Port-Id attribute is sent is: ***slot/port*** [**vpi-vci** ***vpi vci*** | **vlan-id** [***tunl-vlan-id*:**]***pvc-vlan-id***] [**pppoe** ***sess-id*** | **clips** ***sess-id***]

The circuit identifier can be the virtual path identifier (VPI) with the virtual channel identifier (VCI), or it can be the virtual LAN (VLAN) identifier, depending on the type of circuit.

If you specify the **physical** keyword, the format in which the NAS-Port-Id attribute is sent is: ***slot/port*** [**vpi-vci** ***vpi vci*** | **vlan-id** [***tunl-vlan-id*:**]***pvc-vlan-id***].

If you specify the **modified-agent-circuit-id** keyword, the system inserts the specific subscriber line information in the NAS-Port-ID attribute. Line information includes: ***slot/port*** [**vpi-vci** ***vpi vci*** | **vlan-id**

[*tunl-vlan-id*:]*pvc-vlan-id*] which is prepended to the subscriber identification fields.

To indicate that a text string description of the access link group is to be used as a prefix to the NAS-Port-Id attribute using the **description** command, specify the **format**, **modified-agent-circuit-id**, and **prefix-lg-description** keywords with the **radius attribute nas-port-id** command. For more information about the **description** command, see the *Command List*.

Standard RADIUS attribute 87, NAS-Port-Id, and vendor-specific attributes (VSAs) 96 provided by Ericsson AB, Agent-Remote-Id, and 97, Agent-Circuit-Id, are described in *RADIUS Attributes*.

Use the **no** or **default** form of this command to reset the format for the NAS-Port-Id attribute to the **all** format.

Use the **default** form of this command to specify the default separator. To change the default separator character, specify the **separator** keyword and character to use as the separator.

1.25.6 Examples

The following example shows how to send the NAS-Port-Id attribute using the `physical` format in RADIUS Access-Request and Accounting-Request packets for the `local` context:

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

[local]Redback(config-ctx)#radius attribute nas-port-id format physical
```

1.26 radius attribute nas-port-type

radius attribute nas-port-type *port-type*

{**no** | **default**} **radius attribute nas-port-type** *port-type*

1.26.1 Purpose

Modifies the value for the network access server (NAS)-Port-Type attribute sent in RADIUS Access-Request and Accounting-Request packets.

1.26.2 Command Mode

- ATM profile configuration
- dot1q profile configuration

- link-group configuration
- port configuration

1.26.3 Syntax Description

port-type

Value that represents the type of connection the subscriber has to the network access server (NAS) through which it is authenticated. The range of values is 0 to 255. Values 0 to 19 are defined in Table 1.

The default value is either 0 or 5, indicating an asynchronous connection through a console port or a virtual connection through a transport protocol, respectively.

1.26.4 Default

The Nas-Port-Type attribute is sent in RADIUS Access-Request and Accounting-Request packets. The value is either 0 or 5, depending on how the subscriber is connected to its authenticating NAS.

1.26.5 Usage Guidelines

Use the **`radius attribute nas-port-type`** command to modify the value for the NAS-Port-Type attribute sent in RADIUS Access-Request and Accounting-Request packets. You can set this attribute for an ATM profile, an 802.1Q profile, a link group, or a port.

Note: For link groups, the NAS-Port-Type attribute can be set only at the link-group level, not for the constituent ports. Any ports that already have the NAS-Port-Type attribute configured cannot be added to a link group until this configuration is removed from the port.

Table 1 lists the definitions of the values for the ***`port-type`*** argument.

Table 1 Values for the port-type Argument

Value	Definition
0	async
1	sync
2	ISDN (sync)
3	ISDN (async V120)
4	ISDN (async V110)
5	Virtual
6	PIAFS (wireless ISDN used in Japan)

Table 1 Values for the port-type Argument

Value	Definition
7	HDLC (clear-channel)
8	X.25
9	X.75
10	G3_Fax (G.3 Fax)
11	SDSL (symmetric DSL)
12	ADSL_CAP (asymmetric DSL Carrierless Amplitude Phase Modulation)
13	ADSL_DMT (asymmetric DSL, Discrete Multi-Tone)
14	IDSL (ISDN digital subscriber line)
15	Ethernet
16	xDSL (digital subscriber line of unknown type)
17	Cable
18	Wireless (wireless—other)
19	Wireless_802_11 (wireless—IEEE 802.11)

Standard RADIUS attribute 61, NAS-Port-Type, is described in *RADIUS Attributes*.

Use the **no** or **default** form of this command to reset the SmartEdge router behavior to the default condition.

1.26.6 Examples

The following example shows how to modify the NAS-Port-Type attribute in RADIUS Access-Request and Accounting-Request packets to type 4 (ISDN) for an ATM profile:

```
[local]Redback(config)#atm profile ATMPROFILE
[local]Redback(config-atm-profile)#radius attribute nas-port-type 4
```

1.27 radius attribute nas-port-type from-encapsulation

radius attribute nas-port-type from-encapsulation

no radius attribute nas-port-type from-encapsulation

1.27.1 Purpose

Enables authentication, authorization, and accounting (AAA) to set the value for the network access server (NAS)-Port-Type attribute sent in RADIUS Access-Request and Accounting-Request packets according to the L2TP bearer type.

1.27.2 Command Mode

context configuration

1.27.3 Syntax Description

from-encapsulation	The NAS-Port-Type attribute value is taken from one of the following: • L2TP_BEARER_TYPE_DIGITAL (nas-port-type is Port_Type_ISDN_Sync;) • L2TP_BEARER_TYPE_ANALOG (nas-port-type is Port_Type_Async;) • L2TP_BEARER_TYPE_ETHERNET (nas-port-type is Port_Type_Virtual;)

1.27.4 Default

None

1.27.5 Usage Guidelines

For an LNS, to take the NAS port type from the subscriber's original ingress port type on the LAC, use the **radius attribute nas-port-type from-encapsulation** command.

Standard RADIUS attribute 61, NAS-Port-Type, is described in *RADIUS Attributes*.

Use the **no** form of this command to disable this behavior on the SmartEdge router.

1.27.6 Examples

The following example enable AAA in the local context to determine the NAS-Port-Type attribute from the L2TP bearer type:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#radius attribute nas-port-type from-encapsulation
```

1.28 radius attribute nas-ipv6-address

radius attribute nas-ipv6-address interface ***if-name***

{**no** | **default**} **radius attribute nas-ipv6-address**

1.28.1 Purpose

Includes the network access server (NAS)-IPv6-Address attribute in RADIUS Access-Request and Accounting-Request packets sent by the SmartEdge router.

1.28.2 Command Mode

context configuration

1.28.3 Syntax Description

interface ***if-name***	Interface name. Uses the primary IPv6 address associated with the interface as the source IPv6 address sent in RADIUS packets. If the interface is not configured or is unreachable, the IPv6 address of the outgoing interface is used instead as the source IPv6 address for packets.

1.28.4 Default

The NAS-IPv6-Address attribute is not sent.

1.28.5 Usage Guidelines

Use the **radius attribute nas-ipv6-address** command to include the NAS-IPv6-Address attribute in RADIUS Access-Request and Accounting-Request packets sent by the SmartEdge router.

Standard RADIUS attribute 4, NAS-IPv6-Address, is described in *RADIUS Attributes*.

Use the **no** or **default** form of this command to reset the SmartEdge router behavior so that the NAS-IPv6-Address attribute is not included.

If using RADIUS to authenticate a subscriber, configure the NAS-IPV6-Address to match the IPv6 address of the NAS interface. When a subscriber is successfully authenticated, the radius server returns an IPv6 prefix /64 to the BRAS. If the circuit is configured for dual-stack, an IPv4 address is also returned.

1.28.6 Examples

The following example shows how to send the primary IPv6 address for interface ABC as the source IPv6 address in RADIUS Access-Request and Accounting-Request packets sent by the SmartEdge router, where the interface ABC is already defined in the configuration as follows:

```
interface ABC
  ipv6 address 2001:1001:3001:4001:5001:6001:7001:7001/127
```

```
[local]Redback(config-ctx)#radius attribute nas-ipv6
-address interface ABC
```

Following is sample output sent in the RADIUS packet.

```
NAS-IPv6-Address = 2001:1001:3001:4001:5001:6001:7001:7001
```

1.29 radius attribute username

radius attribute username encaps clips {strip-mac-delimite r | {[prefix *prefix-string*] suffix *suffix-string*}}

{no | default} radius attribute username encaps clips {strip-mac-delimiter | {[prefix *prefix-string*] suffix *suffix-string*}}

1.29.1 Purpose

Specifies the format of the username attribute sent in RADIUS messages.

1.29.2 Command Mode

context configuration

1.29.3 Syntax Description

encaps clips	Specifies that this command applies to the username sent in RADIUS messages for CLIPS subscribers.
prefix *prefix-string*	Sets the username prefix to the value specified.

suffix ***suffix-string***	Sets the username suffix to the value specified.
strip-mac-delimiter	Strips the MAC delimiter characters from the username.

1.29.4 Default

The SmartEdge router does not specify RADIUS username attribute options.

1.29.5 Usage Guidelines

Use the **radius attribute username** command to specify the format of the username attribute sent in RADIUS messages.

Use the **no** or **default** form of this command to return the specified username attribute to the default.

1.29.6 Examples

The following example shows how to specify the MAC deliminator characters are stripped from the username in RADIUS messages for CLIPS subscribers:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#radius attribute username encaps clips strip-mac-delimiter
```

1.30 radius attribute vendor-specific

radius attribute vendor-specific Redback {**mac-address separator** ***char*** | **salt-encrypted-attr** {**authen-server** | **coa-server**} | **ipv4-address-release-control text** ***arbitrary-string***}

{**no** | **default**} **radius attribute vendor-specific Redback** {**mac-address separator** ***char*** | **salt-encrypted-attr** {**authen-server** | **coa-server**} | **ipv4-address-release-control text** ***arbitrary-string***}

1.30.1 Purpose

Specifies the character the SmartEdge router uses to separate the fields in the specified RADIUS attribute, whether attributes can be encrypted, or whether to enable IPv4 address save mode for the context.

1.30.2 Command Mode

context configuration

1.30.3 Syntax Description

Redback	Specifies Redback as the vendor.
mac-address	Specifies vendor-specific attribute (VSA) 145 provided by Ericsson AB, Mac-Addr, as the attribute.
separator *char*	Character to be used as a separator. The default is hyphen (-).
salt-encrypted-attr	Allows encrypted VSA attributes.
authen-server	Allows encrypted VSAs in Access-Response packets.
coa-server	Allows encrypted VSAs in CoA-Request packets.
ipv4-address-release -control	Specifies VSA 213 provided by Ericsson AB, IPv4-Address-Release-Control, as the attribute. Enables IPv4 address save mode.
text *arbitrary-string*	Arbitrary string for the IPv4-Address-Release-Control VSA. Changes to the configured string are applied immediately. The string can be up to 63 characters.

1.30.4 Default

The SmartEdge router uses the hyphen (-) as a separator, the VSAs can be encrypted, and IPv4 address save mode is disabled.

1.30.5 Usage Guidelines

Use the **radius attribute vendor-specific** command to specify the character the SmartEdge router uses to separate the fields in the specified RADIUS attribute, whether attributes can be encrypted, or whether to enable IPv4 address save mode.

IPv4 address save mode allows service providers to conserve IPv4 addresses by issuing them to subscribers from their public, shared IPv4 address pool as required. When not in use, subscribers can release IPv4 addresses back into the provider's pool. For more information about IPv4 address save mode, see *Configuring to Conserve IPv4 Addresses for Dual-Stack Subscribers* in *Configuring IPv6 Subscriber Services*.

Use the **no** or **default** form of this command to specify the default character as the separator.

1.30.6 Examples

The following example specifies the colon (:) as the separator character:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#radius attribute vendor-specific Redback mac-address separator :
```

The following example enables IPv4 address save mode for the context:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#radius attribute vendor-specific Redback ipv4-address-release-control text saveaddr
```

1.31 radius await-acct-on-ack

radius await-acct-on-ack

{**no** | **default**} **radius await-acct-on-ack**

1.31.1 Purpose

In global configuration mode, enables the SmartEdge router to wait to receive an acknowledgement of the receipt of the Accounting-On message that it sent to the RADIUS accounting servers after the router reboots. The router waits for the acknowledgement message (ACK message) before it begins to send authentication and accounting requests to the RADIUS servers.

1.31.2 Command Mode

Global configuration

1.31.3 Syntax Description

This command has no keywords or arguments.

1.31.4 Default

By default, after the SEOS router reboots, an Accounting-On message is sent to the RADIUS servers. However, the router does not wait for an acknowledgement of the receipt of this message from the servers before sending the authentication and accounting requests to the RADIUS servers.

1.31.5 Usage Guidelines

Use the **radius await-acct-on-ack** command in global configuration mode to enable the SmartEdge router to wait to receive an acknowledgement of the receipt of the Accounting-On message that it sent to the RADIUS accounting servers after the router reboots. This command ensures that the router waits for an ACK message before it begins to send authentication and accounting requests to the servers.

If the SmartEdge router has not yet received an ACK message from the RADIUS server within 10 seconds of sending the Accounting-On message, the router checks again. If after 30 retries, and still no ACK message is received, the router sends authentication and accounting requests to the RADIUS servers. The interval time between retries is 10 seconds. The maximum time the router waits for a response from a RADIUS server before assuming that a packet is lost, or that the RADIUS server is unreachable is 5 minutes.

If more than one RADIUS accounting server is configured, and the router receives an ACK message from any one of the servers, the router considers this message as an acknowledgement of the receipt of the Accounting-On message. If “two-stage accounting” is configured, an ACK message from the global accounting server (configured in context local) is the one that is counted as the acknowledgement of the receipt of the Accounting-On message.

Note: The **radius await-acct-on-ack** command does not take effect, if local authentication with RADIUS accounting is configured.

Use the **no** or **default** form of this command to return the SmartEdge router to its default setting of not waiting to receive an acknowledgement of receipt of the Accounting-On message from the RADIUS servers before the router sends authentication and accounting requests to the servers.

1.31.6 Examples

The following example shows how to configure the **radius await-acct-on-ack** command:

```
[local]Redback(config)#radius await-acct-on-ack

[local]Redback(config)#
```

1.32 radius coa server

radius coa server {*ip-addr* | *hostname*} {**key** *key* | **encrypted-key** *key*} [**port** *udp-port*]

no radius coa server {*ip-addr* | *hostname*}

1.32.1 Purpose

Configures the IP address or hostname of a RADIUS Change of Authorization (CoA) server.

1.32.2 Command Mode

context configuration

1.32.3 Syntax Description

ip-addr	IP address of the RADIUS CoA server.
hostname	Hostname of the RADIUS CoA server. The Domain Name System (DNS) must be enabled in order to use the ***hostname*** argument.
key ***key***	Alphanumeric string indicating the secret key that must be shared with the RADIUS CoA server. If multiple subscriber sessions share the same key, all sessions are affected by a CoA change.
encrypted-key ***key***	Alphanumeric string representing the encrypted secret key that must be shared with the RADIUS CoA server. If multiple subscriber sessions share the same key, all sessions are affected by a CoA change.
port ***udp-port***	Optional. RADIUS CoA server User Datagram Protocol (UDP) port. The range of values is 1 to 65,536. If no port is specified, port 3799 is used is for CoA messages. The ***udp-port*** value indicates the CoA port.

1.32.4 Default

RADIUS CoA server hostnames and IP addresses are not preconfigured. Port 3799 is the User Datagram Protocol (UDP) CoA port.

1.32.5 Usage Guidelines

Use the **radius coa server** command to configure the IP address or hostname of a RADIUS CoA server. You can enter this command multiple times to configure up to 36 RADIUS CoA servers per context. RADIUS CoA servers configured in a non-local context can change session settings only for subscribers in the same context. CoA servers configured in the local context can change settings for all subscribers.

To use the ***hostname*** argument, DNS must be enabled; for more information.

Note: To enable authentication to be performed by RADIUS, you must also enter the **aaa authentication subscriber** command (in context configuration mode); for more information, see *Configuring Bridging*.

The RADIUS CoA server can use one or more of the identifiers listed in Table 2 to identify a subscriber session.

Table 2 RADIUS CoA Session Identifiers

Identifier	Notes
Username	For global authentication, use the RBN CONTEXT-NAME VSA to search the appropriate context. If this attribute is not specified, only the local context is searched.
Acct-Session-ID	This identifier is unique across all contexts.
IP-Address	This identifier is unique only within a context. Using this identifier returns all sessions in all contexts with the specified IP address. For global authentication, use the RBN CONTEXT-NAME VSA to search the appropriate context. If this attribute is not specified, only the local context is searched.
Agent-Circuit-ID	This identifier is unique across all contexts.
Agent-Remote-ID	This identifier is unique across all contexts.

For CoA disconnect messages, specify at least one keyword in Table 2. For all other CoA messages, specify at least one keyword in Table 2, as well as one or more attributes to change. If multiple keywords are specified, all specified keywords must match subscriber session attributes.

RADIUS Attributes lists the RADIUS attributes supported by the SmartEdge router. If a CoA message contains any unsupported attributes, the request fails. RADIUS CoA and disconnect features are described in RFC 3576, *Dynamic Authorization Extensions to Remote Authentication Dial-In User Service (RADIUS)*. If the specified keyword matches multiple subscriber sessions, and the requested change is successful for only a subset of the sessions, all successful changes are preserved. The SmartEdge router sends a negative acknowledgement (NAK). When an attempt to modify any other attribute fails, the subscriber session is terminated.

Use the **no** form of this command to delete a previously configured RADIUS CoA server.

1.32.6 Examples

The following example shows how to configure a RADIUS CoA server IP address of `10.3.3.3` with the key, `secret`, using port `4444` for CoA messages:

```
[local]Redback(config-ctx)#radius coa server 10.3.3.3 key secret port 4444
```

The following example configures 36 CoA servers in the local context:

```
[local]Redback(config-ctx)#radius coa server 1.1.1.1 encrypted-key 6840234DDC4A32D4 port 3799
[local]Redback(config-ctx)#radius coa server 1.1.1.2 encrypted-key 6840234DDC4A32D4 port 3800
[local]Redback(config-ctx)# radius coa server 1.1.1.3 encrypted-key 6840234DDC4A32D4 port 3801
[local]Redback(config-ctx)# radius coa server 1.1.1.4 encrypted-key 6840234DDC4A32D4 port 3802
[local]Redback(config-ctx)# radius coa server 1.1.1.5 encrypted-key 6840234DDC4A32D4 port 3803
[local]Redback(config-ctx)# radius coa server 1.1.1.6 encrypted-key 6840234DDC4A32D4 port 3804
[local]Redback(config-ctx)# radius coa server 1.1.1.7 encrypted-key 6840234DDC4A32D4 port 3805
[local]Redback(config-ctx)# radius coa server 1.1.1.8 encrypted-key 6840234DDC4A32D4 port 3806
[local]Redback(config-ctx)# radius coa server 1.1.1.9 encrypted-key 6840234DDC4A32D4 port 3807
[local]Redback(config-ctx)# radius coa server 1.1.1.10 encrypted-key 6840234DDC4A32D4 port 3808
[local]Redback(config-ctx)# radius coa server 1.1.1.11 encrypted-key 6840234DDC4A32D4 port 3809
[local]Redback(config-ctx)# radius coa server 1.1.1.12 encrypted-key 6840234DDC4A32D4 port 3810
[local]Redback(config-ctx)# radius coa server 1.1.1.13 encrypted-key 6840234DDC4A32D4 port 3811
[local]Redback(config-ctx)# radius coa server 1.1.1.14 encrypted-key 6840234DDC4A32D4 port 3812
[local]Redback(config-ctx)# radius coa server 1.1.1.15 encrypted-key 6840234DDC4A32D4 port 3813
[local]Redback(config-ctx)# radius coa server 1.1.1.16 encrypted-key 6840234DDC4A32D4 port 3814
[local]Redback(config-ctx)# radius coa server 1.1.1.17 encrypted-key 6840234DDC4A32D4 port 3815
[local]Redback(config-ctx)# radius coa server 1.1.1.18 encrypted-key 6840234DDC4A32D4 port 3816
[local]Redback(config-ctx)# radius coa server 1.1.1.19 encrypted-key 6840234DDC4A32D4 port 3817
[local]Redback(config-ctx)# radius coa server 1.1.1.20 encrypted-key 6840234DDC4A32D4 port 3818
[local]Redback(config-ctx)# radius coa server 1.1.1.21 encrypted-key 6840234DDC4A32D4 port 3819
[local]Redback(config-ctx)# radius coa server 1.1.1.22 encrypted-key 6840234DDC4A32D4 port 3820
[local]Redback(config-ctx)# radius coa server 1.1.1.23 encrypted-key 6840234DDC4A32D4 port 3821
[local]Redback(config-ctx)# radius coa server 1.1.1.24 encrypted-key 6840234DDC4A32D4 port 3822
[local]Redback(config-ctx)# radius coa server 1.1.1.25 encrypted-key 6840234DDC4A32D4 port 3823
[local]Redback(config-ctx)# radius coa server 1.1.1.26 encrypted-key 6840234DDC4A32D4 port 3824
[local]Redback(config-ctx)# radius coa server 1.1.1.27 encrypted-key 6840234DDC4A32D4 port 3825
[local]Redback(config-ctx)# radius coa server 1.1.1.28 encrypted-key 6840234DDC4A32D4 port 3826
[local]Redback(config-ctx)# radius coa server 1.1.1.29 encrypted-key 6840234DDC4A32D4 port 3827
[local]Redback(config-ctx)# radius coa server 1.1.1.30 encrypted-key 6840234DDC4A32D4 port 3828
[local]Redback(config-ctx)# radius coa server 1.1.1.31 encrypted-key 6840234DDC4A32D4 port 3829
[local]Redback(config-ctx)# radius coa server 1.1.1.32 encrypted-key 6840234DDC4A32D4 port 3830
[local]Redback(config-ctx)# radius coa server 1.1.1.33 encrypted-key 6840234DDC4A32D4 port 3831
[local]Redback(config-ctx)# radius coa server 1.1.1.34 encrypted-key 6840234DDC4A32D4 port 3832
[local]Redback(config-ctx)# radius coa server 1.1.1.35 encrypted-key 6840234DDC4A32D4 port 3833
[local]Redback(config-ctx)# radius coa server 1.1.1.36 encrypted-key 6840234DDC4A32D4 port 3834
```

1.33 radius deadtime

radius deadtime ***interval***

{**no** | **default**} **radius deadtime**

1.33.1 Purpose

Sets the interval during which the SmartEdge router treats a nonresponsive RADIUS server as “dead.”

1.33.2 Command Mode

context configuration

1.33.3 Syntax Description

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

1.33.4 Default

The waiting interval is five minutes.

1.33.5 Usage Guidelines

Use the **radius deadtime** command to set the interval during which the SmartEdge router treats a nonresponsive RADIUS server as “dead.” During the interval, the SmartEdge router tries to reach another RADIUS server; after the interval expires, the SmartEdge router tries again to reach the server. If there is no response, the RADIUS server remains marked as “dead” and the timer is set again to the configured interval.

If you disable this feature (with the 0 value), the SmartEdge router never waits but attempts to reach the server immediately.

Note: You must configure at least one RADIUS server using the **radius server** command (in context configuration mode) prior to entering this command.

Use the **default** form of this command to specify the default interval.

1.33.6 Examples

The following example shows how to set the deadtime interval to `10` minutes:

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

1.34 radius max-outstanding

radius max-outstanding ***requests***

{**no** | **default**} **radius max-outstanding**

1.34.1 Purpose

Modifies the number of simultaneous outstanding requests that can be sent by the SmartEdge router to RADIUS servers.

1.34.2 Command Mode

context configuration

1.34.3 Syntax Description

requests	Number of simultaneous outstanding requests per RADIUS server in the current context. The range of values is 1 to 256.

1.34.4 Default

The maximum number of allowable outstanding requests is 256.

1.34.5 Usage Guidelines

Use the **radius max-outstanding** command to modify the number of simultaneous outstanding requests the SmartEdge router can send to RADIUS servers.

Use the **no** or **default** form of this command to reset the maximum number of outstanding requests to 256.

1.34.6 Examples

The following example shows how to limit the number of simultaneous outstanding requests to `128`:

```
[local]Redback(config-ctx)#radius max-outstanding 128
```

1.35 radius max-retries

radius max-retries ***retries***

default radius max-retries

1.35.1 Purpose

Modifies the number of retransmission attempts the SmartEdge router makes to a RADIUS server in the event that no response is received from the server within the timeout period.

1.35.2 Command Mode

context configuration

1.35.3 Syntax Description

retries	Number of retransmission attempts the SmartEdge router will make. The range of values is 1 to 2,147,483,647; the default value is 3.

1.35.4 Default

The SmartEdge router makes three retransmission attempts.

1.35.5 Usage Guidelines

Use the **radius max-retries** command to modify the number of retransmission attempts the SmartEdge router makes to a RADIUS server in the event that no response is received from the server within the timeout period.

You set the timeout period with the **radius timeout** command (in context configuration mode).

If an acknowledgment is not received, each successive server is tried (wrapping from the last server to the first, if necessary) until the maximum number of retransmissions is reached.

Use the **default** form of this command to specify the default number of retries.

1.35.6 Examples

The following example shows how to set the retransmit value to 5:

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

The following example resets the retransmit value to the default (3):

```
[local]Redback(config-ctx)#default radius max-retries
```

1.36 radius policy

In global configuration mode, the syntax is:

radius policy name *pol-name*

no radius policy name *pol-name*

In context configuration mode, the syntax is:

radius policy *pol-name*

no radius policy *pol-name*

default radius policy

1.36.1 Purpose

In global configuration mode, creates or modifies a RADIUS policy and accesses RADIUS policy configuration mode; in context configuration mode, assigns a RADIUS policy to the context.

1.36.2 Command Mode

context configuration

global configuration

1.36.3 Syntax Description

pol-name	Name of the RADIUS policy being assigned.
name *pol-name*	Name of the RADIUS policy being created or modified.

1.36.4 Default

No RADIUS policy is created or assigned to a context.

1.36.5 Usage Guidelines

Use the **radius policy** command in global configuration mode to create or modify a RADIUS policy and access RADIUS policy configuration mode; use it in context configuration mode to assign a RADIUS policy to the context.

The SmartEdge router supports RADIUS policy filter functionality in both inbound and outbound directions. The RADIUS policy specifies which RADIUS attributes and vendor-specific attributes (VSAs) are to be removed from RADIUS Access-Request, Access-Accept, CoA, and various Accounting-Request messages (Accounting-On, Accounting-Off, Accounting-Start, Accounting-Stop, and Accounting-Update). Up to 30 policy filters can be defined in global context. Use the **attribute** command (in RADIUS policy configuration mode) to specify the attributes to be removed from the messages. Up to 500 attributes can be filtered in a single RADIUS policy.

Use the **no** form of this command in global configuration mode to delete the policy; use it in context configuration mode to remove the policy from the context configuration.

Use the **default** form of this command in context configuration mode to return the context to the default condition of having no configured RADIUS policy.

1.36.6 Examples

The following example creates the `custom` RADIUS policy:

```
[local]Redback(config)#radius policy name custom
[local]Redback(config-rad-policy)#
```

The following example assigns the `custom` RADIUS policy to the `gold-isp` context:

```
[local]Redback(config)#context gold-isp
[local]Redback(config-ctx)#radius policy custom
```

1.37 radius route-download algorithm

radius route-download algorithm { **first** | **round-robin** }

1.37.1 Purpose

Specifies the load-balancing algorithm to use when multiple servers are configured. The default algorithm is **first**; the request is sent to the first available radius server. When **round-robin** is configured, the subsequent request is sent to the next available server.

1.37.2 Command Mode

context configuration

1.37.3 Syntax Description

first	The request is sent to the first available server.
round-robin	Requests are sent to RADIUS route download servers.

1.37.4 Default

The default setting for this command is `first`.

1.37.5 Usage Guidelines

Specify round-robin route download algorithm if you have configured multiple servers using the radius route-download server command.

1.37.6 Examples

The example shows how to set the route download algorithm to **round-robin**.

```
[local]Redback(config-ctx)#radius route-download algorithm round-robin
```

1.38 radius route-download deadtime

radius route-download deadtime ***timeout***

1.38.1 Purpose

The interval, in minutes, to consider a RADIUS route-download server unavailable before declaring it available again. If not specified, the default value is used.

1.38.2 Command Mode

context configuration

1.38.3 Syntax Description

timeout	Interval specified in minutes. The range is 0 to 65535.

1.38.4 Default

The default radius route download deadtime value is 5 minutes.

1.38.5 Examples

The following example shows how to configure the RADIUS route-download deadtime.

```
[local]Redback(config-ctx)#radius route-download deadtime 10
```

1.39 radius route-download max-retries

radius route-download max-retries ***max-retries***

1.39.1 Purpose

The maximum number of times a route download request is retried.

1.39.2 Command Mode

context configuration

1.39.3 Syntax Description

max-retries	Indicate the maximum number of route download request attempts. The range is 1 to 2147483647.

1.39.4 Default

The default number of retries is 3.

1.39.5 Examples

The example shows how to set the maximum number of times route download retries are attempted, in this case, 255 times:

```
[local]Redback(config-ctx)#radius route-download max-retries 255
```

1.40 radius route-download server

radius route-download server ***server-address*** **key** | **encrypted-key** ***key*** [**port** ***port-num***]

no radius route-download server ***server-address*** **key** | **encrypted-key** ***key*** [**port** ***port-num***]

1.40.1 Purpose

Designates a RADIUS route download server, and configures the shared key (encrypted or in plain text) and optional port. If the port is not specified, port 1812 is used.

This command is a prerequisite for all the route download commands.

1.40.2 Command Mode

context configuration

1.40.3 Syntax Description

server *server-address*	The IP address or hostname of the RADIUS server to be used as a route download server.
key | **encrypted-key** *key*	The shared key or encrypted key to be used to authenticate with the RADIUS server. The SmartEdge router sends this key in plain text or encrypted form to the RADIUS server.
port *port-num*	The UDP port to be used for RADIUS communications. If the port is not specified, the SmartEdge router uses the well-known port for RADIUS (port 1812).

1.40.4 Default

The default value for the port is 1812.

1.40.5 Usage Guidelines

This command is a prerequisite for all the route download commands.

When using the **no** form of this command, keep the following guidelines in mind:

- The **no radius route-download server** command with no additional keywords deletes the existing single server address.
- To delete multiple servers, you must enter each server address individually using the **no radius route-download server** *server-address* command.
- To delete multiple ports on the same server or multiple ports on multiple servers, you must enter each server address individually and each port individually using the **no radius route-download server** *server-address* **key** | **encrypted-key** *key* [**port** *port-num*] command.

Defining the ports is optional, however, you must always indicate a plain-text or encrypted key.

1.40.6 Examples

This example shows how to configure multiple ports on the same route-download server.

```
[local]Redback(config-ctx)#radius route-download server 10.18.18.33
encrypted-key redback port 1850
[local]Redback(config-ctx)#radius route-download server 10.18.18.33
encrypted-key redback port 1860
```

This example shows how to configure different route-download servers and identifies the ports to use on each server.

```
[local]Redback(config-ctx)#radius route-download server 10.18.18.33
encrypted-key redback port 1850
[local]Redback(config-ctx)#radius route-download server 10.18.18.60
encrypted-key redback port 1850
```

This example shows how to configure two different servers using an encrypted key (first configuration command) and a plain-text key (second configuration command).

```
[local]Redback(config-ctx)#radius route-download server 10.18.18.60
encrypted-key redback port 1829
[local]Redback(config-ctx)#radius route-download server 10.18.18.33
key redback port 1850
```

1.41 radius route-download server-timeout

radius route-download server-timeout ***timeout***

1.41.1 Purpose

The interval, in seconds, to wait for a response from the RADIUS server before declaring it unavailable. If not specified, the default value is used.

1.41.2 Command Mode

context configuration

1.41.3 Syntax Description

timeout	Timeout value in seconds. The range is 30 to 65535 seconds.

1.41.4 Default

The default value is 60 seconds.

1.41.5 Examples

The following example shows a server timeout value of `40` seconds:

```
[local]Redback(config-ctx)#radius route-download server-timeout 40
```

1.42 radius route-download timeout

`radius route-download timeout` ***`timeout`***

1.42.1 Purpose

The interval, in seconds, to wait before retrying a route download request. If no timeout value is specified, the default interval is used.

1.42.2 Command Mode

context configuration

1.42.3 Syntax Description

`timeout`	Timeout value in seconds. The range is 1 to 2147483647.

1.42.4 Default

The default route download timeout value is 10 seconds.

1.42.5 Examples

The following example shows how to configure `20` seconds as the route download timeout value:

```
[local]Redback(config-ctx)#radius route-download timeout 20
```

1.43 radius server

radius server {*ip-addr* | *hostname*} {**key** *key* | **encrypted-key** *key*} [**CoA-server**] [{**oldports** | **port** *udp-port*}]

no radius server {*ip-addr* | *hostname*}

1.43.1 Purpose

Configures the IP address or hostname of a RADIUS server.

1.43.2 Command Mode

context configuration

1.43.3 Syntax Description

ip-addr	IP address of the RADIUS server.
hostname	Hostname of the RADIUS server. The Domain Name System (DNS) must be enabled in order to use the *hostname* argument.
key *key*	Alphanumeric string indicating the authentication key that must be shared with the RADIUS server.
encrypted-key *key*	Alphanumeric string representing the encrypted authentication key that must be shared with the RADIUS server.
CoA-server	Optional. Uses the RADIUS server as a Change of Authorization (CoA) server.
oldports	Optional. Uses the RADIUS User Datagram Protocol (UDP) ports 1645 for authentication.
port *udp-port*	Optional. RADIUS authentication UDP port. The range of values is 1 to 65,536. If no port is specified, port 1812 is used for authentication. The *udp-port* value indicates the authentication port.

1.43.4 Default

RADIUS server hostnames and IP addresses are not preconfigured. 1812 is the UDP authentication port.

1.43.5 Usage Guidelines

Use the **radius server** command to configure the IP address or hostname of a RADIUS server. You can use this command multiple times to configure up to five RADIUS servers per context.

To use the ***hostname*** argument, DNS must be enabled.

Note: To enable authentication to be performed by RADIUS, you must also enter the **aaa authentication subscriber** command (in context configuration mode); for more information, see *Configuring Bridging*.

If you specify the optional **CoA-server** keyword, the same port that is used for authentication is also used for CoA messages.

The RADIUS CoA server can use one or more of the keywords listed in Table 2 to identify a subscriber session. For information on CoA interactions, see the **radius coa server** command.

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

1.43.6 Examples

The following example shows how to configure a RADIUS server with an IP address of `10.3.3.3` with the key, `secret`, using ports `4444` for authentication:

```
[local]Redback(config-ctx)#radius server 10.3.3.3 key secret port 4444
```

1.44 radius server-timeout

radius server-timeout ***interval***

default radius server-timeout

1.44.1 Purpose

Sets the time interval the SmartEdge router waits before marking a non-responsive RADIUS server as “dead.”

1.44.2 Command Mode

context configuration

1.44.3 Syntax Description

interval	Number of seconds after which the SmartEdge router checks for successful responses after an individual RADIUS request times out, before treating the server as “dead.” The range of values, in seconds, is 0 to 2,147,483,647; the default value is 60.

1.44.4 Default

The maximum time interval is 60 seconds.

1.44.5 Usage Guidelines

Use the **radius server-timeout** command to set the time interval the SmartEdge router waits before marking a non-responsive RADIUS accounting server as “dead.”

The SmartEdge router marks a RADIUS server as “dead” when no response is received to any RADIUS requests during the time period specified by the ***interval*** argument. Setting the value to 0 disables this feature; in this case, no RADIUS server is marked as “dead.”

Use the **default** form of this command to specify the default interval.

1.44.6 Examples

The following example sets the waiting interval to `80` seconds:

```
[local]Redback(config-ctx)#radius server-timeout 80
```

1.45 radius service profile

radius service profile ***prof-name***

no radius service profile ***prof-name***

1.45.1 Purpose

Creates or selects a RADIUS-guided service profile and accesses service profile configuration mode.

1.45.2 Command Mode

context configuration

1.45.3 Syntax Description

prof-name	Name of a service profile.

1.45.4 Default

No RADIUS-guided service profiles exist.

1.45.5 Usage Guidelines

Use the **radius service profile** command to create or select a RADIUS-guided service profile and access service profile configuration mode.

A RADIUS service profile specifies various service conditions and is used to activate services and establish service conditions for that subscriber session. It is these service conditions against which the service data in a CoA Request and Access Response message is matched. You can specify as many as 16 conditions in a service profile.

Use the **no** form of this command to delete the RADIUS-guided service profile from the configuration.

1.45.6 Examples

The following example shows how to create the `redirect` service profile in the `local` context and accesses service profile configuration mode:

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

[local]Redback(config-ctx)#radius service profile redirect

[local]Redback(config-svc-profile)#
```

1.46 radius source-port

radius source-port ***port-num num-ports***

no radius source-port

1.46.1 Purpose

In context configuration mode, enables the SmartEdge router to ignore the source port sent by the RADIUS server in Access-Response messages. In global configuration mode, increases the number of outstanding requests for each RADIUS server by sending requests using a different source port value.

1.46.2 Command Mode

- context configuration
- global configuration

1.46.3 Syntax Description

port-num	Port number. The range of values is 1,024 to 65,535.
num-ports	Number of ports. The range of values is 1 to 10.

1.46.4 Default

This feature is disabled.

1.46.5 Usage Guidelines

In context configuration mode, use the **radius source-port** command to enable the SmartEdge router to ignore the source port sent by the RADIUS server in Access-Response messages. In this configuration mode, this command refers to the source port that the RADIUS server uses when sending a RADIUS Access-Response message to the SmartEdge router.

In global configuration mode, use the **radius source-port** command to increase the number of outstanding requests for each RADIUS server by sending requests using a different source port value. In this configuration mode, this command refers to the source port that the SmartEdge router uses when sending a RADIUS Access-Request message to a RADIUS server.

Use the **no** form of this command to return to the default number of outstanding requests.

1.46.6 Examples

The following example shows how to configure a port number of `2000` and sets the number of ports to `5`:

```
[local]Redback(config)#radius source-port 2000 5
```

1.47 radius strip-domain

radius strip-domain

no radius strip-domain

1.47.1 Purpose

Strips the domain portion of a structured username before relaying an authentication request to a RADIUS server.

1.47.2 Command Mode

context configuration

1.47.3 Syntax Description

This command has no keywords or arguments.

1.47.4 Default

The entire username, including the domain name, is sent to the RADIUS server.

1.47.5 Usage Guidelines

Use the **radius strip-domain** command to strip the domain portion of a structured username before relaying an authentication request to a RADIUS server. The username can be either a subscriber name or administrator name.

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

1.47.6 Examples

The following example shows how to prevent the domain portion of the structured username from being sent to the RADIUS server for authentication:

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

1.48 radius timeout

radius timeout ***timeout***

default radius timeout

1.48.1 Purpose

Sets the maximum time the SmartEdge router waits for a response from a RADIUS server before assuming that a packet is lost, or that the RADIUS server is unreachable.

1.48.2 Command Mode

context configuration

1.48.3 Syntax Description

timeout	Timeout period in seconds. The range of values is 1 to 2,147,483,647; the default value is 10 seconds.

1.48.4 Default

The maximum time is 10 seconds.

1.48.5 Usage Guidelines

Use the **radius timeout** command to set the maximum time the SmartEdge router waits for a response from a RADIUS server before assuming that a packet is lost, or that the RADIUS server is unreachable.

Use the **default** form of this command to specify the default interval.

1.48.6 Examples

The following example shows how to set the timeout interval to `30` seconds:

```
[local]Redback(config-ctx)#radius timeout 30
```

1.49 range

range {***ip-addr/prefix-length*** | ***ipv6-addr/prefix-length***} [**not-advertise**]

no range {***ip-addr/prefix-length*** | ***ipv6-addr/prefix-length***} [**not-advertise**]

1.49.1 Purpose

Summarizes interarea routes advertised by an area border router (ABR).

1.49.2 Command Mode

- OSPF area configuration

- OSPF3 area configuration

1.49.3 Syntax Description

ip-addr/prefix-length	Specifies the IP address, in the form ***A.B.C.D***, and the prefix length, separated by the slash (**/**) character. The range of values for the ***prefix-length*** argument is 0 to 32.
ipv6-addr/prefix-length	Specifies the IP Version 6 (IPv6) address, in the form ***A:B:C:D:E:F:G:H***, and the prefix length, separated by the slash (**/**) character. The range of values for the ***prefix-length*** argument is 0 to 128.
not-advertise	Optional. Prevents the specified route from being advertised in interarea route summarizations

1.49.4 Default

Route address ranges for interarea route summarization are not specified.

1.49.5 Usage Guidelines

Use the **range** command to summarize interarea routes advertised by an ABR.

Use the optional **not-advertise** keyword to prevent the specified route from being advertised in route summarizations.

Use the **no** form of this command to disable route summarization for a particular summary range. All individual routes contained in the summary range will be advertised to other areas.

1.49.6 Examples

The following example shows how to advertise routes that fall into the range `10.1.0.0 255.255.0.0` in interarea route summaries (one each of the other areas):

```
[local]Redback(config-ospf-area)#range 10.1.0.0 255.255.0.0
```

1.50 range (DHCP)

range ***start-ip-addr end-ip-addr*** [**threshold** [**falling** ***min-threshold***] [**rising** ***max-threshold***] [**trap**] [**log**]]

no range ***start-ip-addr end-ip-addr***

1.50.1 Purpose

Assigns a range of IP addresses to this Dynamic Host Configuration Protocol (DHCP) subnet.

1.50.2 Command Mode

DHCP subnet configuration

1.50.3 Syntax Description

start-ip-addr	Starting IP address of the range.
end-ip-addr	Ending IP address of the range.
threshold	Optional. Enables threshold monitoring and reporting at the range level.
falling ***min-threshold***	Optional. Threshold for the minimum falling number of available leases at which point a trap or a log message is sent if configured.
rising ***max-threshold***	Optional. Threshold for the maximum rising number of available leases.
trap	Optional. Sends a Simple Network Management Protocol (SNMP) trap on reaching the threshold value.
log	Optional. Sends a log message on reaching the threshold value.

1.50.4 Default

No range of IP addresses is assigned to any subnet.

1.50.5 Usage Guidelines

Use the **range** command to assign a range of IP addresses to this DHCP subnet.

The values of the ***start-ip-addr*** and ***end-ip-addr*** arguments must be within the subnet of IP addresses that you have assigned to this subnet using the **subnet** command (in DHCP server configuration mode).

Use the optional **threshold** keyword to enable the monitoring and reporting of available leases at the range level and specify rising and falling values that can trigger an SNMP trap and log message.

You can enter either or both of the **falling** ***min-threshold*** and **rising** ***max-threshold*** constructs in any order. You can enter either or both of the **trap** and **log** keywords in any order for either construct.

Use the **no** form of this command to delete the range from the subnet configuration.

1.50.6 Examples

The following example shows how to assign a range of IP addresses to the `sub2` subnet; it also enables the monitoring and reporting of available leases for this subnet and triggers an SNMP trap when the number of available leases is decreasing and reaches `100`:

```
[local]Redback(config)#context dhcp
[local]Redback(config-ctx)#dhcp server policy
[local]Redback(config-dhcp-server)#subnet 13.1.1.1/24 name sub2
[local]Redback(config-dhcp-subnet)#range 13.1.1.50 13.1.1.100 threshold falling 100 trap
```

1.51 rate

rate [**informational**] ***kbps*** {**burst** ***bytes*** | **time-burst** ***msec***} [{**excess-burst** ***bytes*** | **time-excess-burst** ***msec*** [**counters**] **[hierarchical-counters]**

no rate

1.51.1 Purpose

Sets the rate, burst tolerance, and excess burst tolerance for traffic on the circuit, port, or subscriber record to which the quality of service (QoS) policy is attached, or for a policy group, policy access control list (ACL), or class-definition class of traffic for that policy.

1.51.2 Command Mode

- metering policy configuration

- policy group class configuration
- policing policy configuration

1.51.3 Syntax Description

informational	Optional. Specifies the rate to be used by the system only to calculate a percentage rate for a policy group class when you specify the class rate as a percentage. The effect is that the overall circuit is not rate limited.
kbps	Rate in kbps. The range of values is 5 to 10,000,000.
burst ***bytes***	Burst tolerance in bytes. The range of values is 1 to 4,250,000,000.
time-burst ***msec***	Burst tolerance expressed in milliseconds rather than bytes. The effective burst allowance is calculated as rate multiplied by time.
excess-burst ***bytes***	Optional. Excess burst tolerance in bytes. The range of values is 1 to 4,250,000,000.
time-excess-burst ***msec***	Optional. Excess-burst tolerance expressed in milliseconds rather than bytes. The effective burst allowance is calculated as rate multiplied by time.
counters	Optional. Enables statistics collection for packets that conform to or exceed the rate.
hierarchical-counters	Optional. Enables statistics collection for packets that are dropped on child circuits subject to this policy due to hierarchical inheritance.

1.51.4 Default

No rate is enforced by default.

1.51.5 Usage Guidelines

Use the **rate** command to set the rate, burst tolerance, and excess burst for traffic on the circuit , port, or subscriber record to which the QoS policy is attached, or for a policy group class of traffic for that policy. If entered in metering or policing policy configuration mode, this command accesses policy rate configuration mode; if entered in policy group class configuration mode, this command accesses policy class rate configuration mode.

Use the **informational** keyword to specify that the policy rate will not be used to enforce an overall circuit rate limit, but will be used only to calculate

the class rate if you specify the rate for a policy group class as a percentage of the policy rate, using the **rate percentage** command (in policy group class configuration mode). This keyword is not available in policy group class configuration mode.

Use the **excess-burst** ***bytes*** construct to optionally configure the excess burst tolerance. The burst tolerance and excess burst tolerance are thresholds that can be used to determine the traffic rate at which packets can be dropped or marked. Use the **time-burst** ***msec*** and **time-excess-burst** ***msec*** constructs to specify the burst and excess burst as time intervals.

When configuring a policing or metering burst allowance, it is recommended that you configure no less than 1 ms of traffic at the desired rate. For example, for a rate of 100 Mbps (`(100,000,000 bits/sec * 0.001s) / 8 bits/byte = 12,500 bytes`), configure the rate and burst as follows:

```
rate 100000 burst 12500
```

or

```
rate 100000 burst time-burst 1
```

It is recommended that you configure the burst allowance to be greater than or equal to the MTU size of the link, or for best TCP throughput, a multiple of the MTU size equal to the desired TCP optimal window size.

For acceptable performance of TCP traffic through a policing or metering rate-limiter, configure a burst allowance that is approximately one second of traffic at the desired rate. For example, for a rate of 100 Mbps (`(100,000,000 bits/sec * 1.0s) / 8 bits/byte = 12,500,000 bytes`), configure the rate and burst as follows:

```
rate 100000 burst 12500000
```

or

```
rate 100000 burst time-burst 1000
```

For optimal performance of TCP traffic through a policing or metering rate-limiter, configure an excess-burst allowance so that exceeding traffic (traffic exceeding the QoS rate and burst tolerance) is dropped on a random basis and violating traffic (traffic violating the rate and burst tolerance) is always dropped, as in the following example:

```
rate 100000 burst time-burst 1000 time-burst 2500
```

For information about dropping or marking packets when the traffic rate exceeds the burst tolerance, but does not exceed the excess burst tolerance, see the **exceed** commands. For information about dropping or marking packets when the traffic rate exceeds the excess burst tolerance, see the **violate** commands.

QoS rate customization allows unique rates to be enforced for different circuits bound to the same QoS policy. You can specify these properties in a variety of ways including:

- RADIUS VSAs 196, 156, and 157 (see the *RADIUS Attributes* document)
- The **rate circuit** command
- The **policy-refresh** command
- ANCP/TR-101 (see the *Configuring ANCP* document and the DSL Forum TR-101 documentation)
- Multicast QoS rate adjustments (see the *Configuring IP Multicast* document)

However, before specifying a customized metering or policing rate for a circuit at the circuit or class level, a corresponding base rate must already have been specified in the policing or metering policy to which the circuit is bound. Similarly, if you remove a circuit or class-level rate from a policy by using the **no rate** command, all circuits using that policy stop enforcing the corresponding rate limit, including those circuits for which the rate has been modified through customization.

Use the **counters** keyword to log statistics related to packets that conform to or exceed the rate. For a circuit with a noninherited metering or policing policy (neither the **inherit** nor **hierarchical** keyword is specified), the **counters** keyword enables statistics collection based on enforcement of this rate at the individual circuit level. For a parent circuit that propagates a metering or policing policy to its children through the **inherit** or **hierarchical** keyword, the **counters** option enables statistics collection based on collective metering or policing policy enforcement at the parent circuit level. In other words, the statistics collected for the parent circuit (where the policy is configured) reflect the totals for enforcement of this rate on this circuit and all of its children that are subject to this policy through inheritance.

Use the **hierarchical-counters** keyword to log statistics related to packets dropped on circuits subject to this rate due to the policy configured on a parent circuit with the **hierarchical** keyword specified. The **hierarchical-counters** keyword enables counters on each child circuit subject to this rate through hierarchical inheritance. These counters record the number of drops on the child circuit due to enforcement of the parent circuit policy rate.

Use the **no** form of this command to specify the default traffic rate and burst tolerance.

1.51.6 Examples

The following example marks all traffic conforming to the configured policy rate with expedited forwarding (`ef`) and marks traffic that exceeds the policy rate with default forwarding (`df`):

```
[local]Redback(config)#qos policy GE-in policing

[local]Redback(config-policy-policing)#rate 6000000 burst 9000000 counters

[local]Redback(config-policy-rate)#conform mark dscp ef

[local]Redback(config-policy-rate)#exceed mark dscp df
```

By including the **counters** keyword in the **rate** command, you can use the **show circuit counters** command (in any mode) with the **detail** keyword to display the number of packets that conform to the rate and the number of packets that exceed the rate.

1.52 rate (MDRR policy configuration)

rate ***kbps*** **burst** ***bytes***

no rate

1.52.1 Purpose

Specifies a maximum transmit scheduling rate to be enforced for an entity (circuit, port, or subscriber session) subject to the queuing policy.

1.52.2 Command Mode

MDRR policy configuration

1.52.3 Syntax Description

kbps	Rate in kilobits per second. The range is 56 to 10,000,000.
burst ***bytes***	Burst tolerance in bytes. This construct is available for MDRR policies only. The range of values is 1 to 8,000,000.

1.52.4 Default

No maximum scheduling rate is enforced for an entity (circuit, port, or subscriber session) subject to the queuing policy.

1.52.5 Usage Guidelines

Use the **rate** command to specify a maximum transmit scheduling rate to be enforced for an entity (circuit, port, or subscriber session) subject to the queuing policy.

Use the **no** form of this command to remove the specified scheduling rate.

1.52.6 Examples

The following example shows how to specify a maximum scheduling rate of **600** Mbps for the MDRR policy **MDRRpolicy1**:

```
[local]Redback(config)#qos policy MDRRpolicy1 mdrr
[local]Redback(config-policy-mdrr)#rate 60000 burst 50000
```

1.53 rate (PWFQ)

For priority weighted fair queuing (PWFQ) policies, the command syntax is:

rate {**maximum** | **minimum**} ***kbps***

no rate {**maximum** | **minimum**}

1.53.1 Purpose

Specifies a maximum transmit scheduling rate to be enforced or a minimum transmit scheduling rate to be allocated for an entity (circuit, port, or subscriber session) subject to the queuing policy.

1.53.2 Command Mode

PWFQ policy configuration

1.53.3 Syntax Description

kbps	Rate in kilobits per second. The range of values is 64 to 1,000,000.
maximum	Specifies a maximum scheduling rate value to be enforced for an entity (circuit, port, or subscriber session) subject to the queuing policy.
minimum	Specifies a desired minimum scheduling rate value to be allocated for an entity (circuit, port, or subscriber session) subject to the queuing policy. Under congestion, circuits are given scheduling preference until they have received the specified minimum rate.

1.53.4 Default

No maximum scheduling rate is enforced and no minimum rate allocated for an entity (circuit, port, or subscriber session) subject to the queuing policy.

1.53.5 Usage Guidelines

Use the **rate** command with the **maximum** keyword to specify a maximum transmit scheduling rate to be enforced for an entity (circuit, port, or subscriber session) subject to the queuing policy. Use the **rate** command with the **minimum** keyword to specify a minimum transmit scheduling rate to be allocated for an entity (circuit, port, or subscriber session) subject to the queuing policy.

For PWFQ policies:

- You must specify the maximum rate for the policy using the **rate** command; otherwise, you cannot attach the policy to any traffic-managed port, or any of the 802.1Q tunnels, or permanent virtual circuits (PVCs) configured on it.
- You cannot specify a minimum rate if you intend to specify a relative weight for this policy, using the **weight** command (in PWFQ policy configuration mode) and attach the policy to any traffic-managed port, or any of the 802.1Q tunnels, or PVCs configured on it. The **rate mimimum** and **qos weight** commands are mutually exclusive under a single PWFQ policy.
- The maximum and minimum rates, if both are specified, are compared to ensure that the minimum value is always less than the maximum value.
- If the number of links in an MLPPP bundle is equal to or less than 8, set the maximum transmit scheduling rate for the bundle according to the following formulas, where ***num-links*** is the number of links in an MLPPP bundle:
  - rate (PWFQ) = ***num-links*** * 2048 kbits/sec (for unframed E1 links)
  - rate (PWFQ) = ***num-links*** * 1984 kbits/sec (for framed E1 links)
  - rate (PWFQ) = ***num-links*** * 1544 kbits/sec (for T1 links)

  If you exceed the recommended maximum, the system throughput might be significantly degraded.
- If the number of links in an MLPPP bundle is greater than 8, set the maximum transmit scheduling rate for the bundle according to the following formulas, where ***num-links*** is the number of links in an MLPPP bundle:
  - rate (PWFQ) = ***num-links*** * 1925 kbits/sec (for unframed E1 links)
  - rate (PWFQ) = ***num-links*** * 1864 kbits/sec (for framed E1 links)
  - rate (PWFQ) = ***num-links*** * 1451 kbits/sec (for T1 links)

  If you exceed the recommended maximum, the system throughput might be significantly degraded.

Use the **no** form of this command to remove the specified minimum or maximum rate.

1.53.6 Examples

The following example shows how to specify a maximum rate of **600** Mbps and a minimum rate of **100** Mbps for the PWFQ policy **GE-in**:

```
[local]Redback(config)#qos policy GE-in pwfq
[local]Redback(config-policy-pwfq)#rate maximum 600000
[local]Redback(config-policy-pwfq)#rate minimum 100000
```

1.54 rate-adjust dhcp pwfq

rate-adjust dhcp pwfq ***kbps*** **priority-group** ***group-num***

no rate-adjust dhcp pwfq ***kbps*** **priority-group** ***group-num***

1.54.1 Purpose

Adjusts the enforcement of a priority weighted fair queuing (PWFQ) policy on a circuit based on whether the subscriber is granted a Dynamic Host Configuration Protocol (DHCP) lease.

1.54.2 Command Mode

subscriber configuration

1.54.3 Syntax Description

kbps	Rate in kilobits per second. The range of values is 1 to 1000000.
group-num	TM queue priority group number. The range of values is 0 to 7.

1.54.4 Default

No DHCP-based rate adjustments are applied to the subscriber.

1.54.5 Usage Guidelines

Use the **rate-adjust dhcp pwfq** command to adjust how a PWFQ policy is enforced on a circuit based on whether the subscriber is granted a DHCP lease. When a lease request is granted to a device on a circuit that has this attribute applied, the enforced bandwidth for the specified traffic-management (TM) queue priority group rate is decremented by the specified amount in (kilobits per second) kbps. If there is no TM queue priority group rate configured for the

policy, the rate is less than the minimal enforceable value (64 kbps), or the rate adjustment is not applied to the subscriber.

Once applied, the rate adjustment persists until the DHCP lease is released or expires. At this time, the rate enforced is restored to its full configured value.

This command might be useful for an IPTV in which Remote Multicast Replication (RMR) is being used. When a set-top box (STB) configured as a static subscriber on an 802.1q VLAN comes online and requests an IP address, the PWFQ policy enforced on the VLAN can be adjusted to account for the multicast bandwidth required for IPTV traffic.

Note: To use this command, you must have a quality of service (QoS) PWFQ policy bound to the subscriber session circuit. The dot1q PVC must be static bound with the "bind subscriber" command to the subscriber session. The policy must include an absolute rate value configured for the specified TM queue priority group. You cannot use percentage to specify the rate. For information about the **`qos policy pwfq`** and **`queue priority-group`** commands, see *Command List*.

Use the **`no`** form to remove currently configured DHCP rate adjustment commands and return the subscriber record to the default state (no rate adjustments will be made in response to DHCP lease events).

1.54.6 Examples

The following example shows how to adjust a PWFQ policy for subscriber `stb1`:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#subscriber name stb1
[local]Redback(config-sub)#password pass
[local]Redback(config-sub)#dhcp max-addr 1
[local]Redback(config-sub)#rate-adjust dhcp pwfq 3000 priority-group 3

[local]Redback(config)#port ethernet 1/4
[local]Redback(config-port)#encapsulation dot1q
[local]Redback(config-port)#dot1q pvc 111 encapsulation multi
[local]Redback(config-dot1q-pvc)#bind subscriber stb1 local
[local]Redback(config-dot1q-pvc)#qos policy queuing test
[local]Redback(config-dot1q-pvc)#circuit protocol pppoe
[local]Redback(config-dot1q-pvc)#bind authentication pap chap context local maxi
```

1.55 rate-calculation

`rate-calculation exclude layer-2-overhead`

`no rate-calculation exclude layer-2-overhead`

1.55.1 Purpose

Specifies that rate calculation is to exclude the size of Layer 2 overhead for the Layer 3 circuit on which a policy is applied.

1.55.2 Command Mode

- metering policy configuration
- policing policy configuration

1.55.3 Syntax Description

exclude	Sets rate-limit calculation exclusions.
layer-2-overhead	Specifies that Layer 2 overhead be excluded when calculating rate-limits.

1.55.4 Default

Rate calculations consider the size of the entire Layer 2 frame.

1.55.5 Usage Guidelines

Use the **rate-calculation** command to specify that rate calculation excludes the size of Layer 2 overhead for Layer 3 circuits on which a rate-limiting policy is applied. In this case, the size of the rate-limited packet equals the size of the Layer 3 packet.

The **rate-calculation** command is global across the entire policy, implying that it applies to the overall circuit level and all classes under the policy.

Use the **no** form of this command to return to the default behavior.

1.56 rate circuit

rate circuit {**in** | **out**} ***kbps*** **burst** ***bytes*** [**excess-burst** ***bytes***] [**queuing-burst** ***bytes***]

{**no** | **default**} **rate circuit** {**in** | **out**}

1.56.1 Purpose

Specifies a different rate for a circuit that has a quality of service (QoS) metering, policing, modified deficit round-robin (MDRR), or priority weighted fair queuing (PWFQ) policy attached to it.

1.56.2 Command Mode

- ATM PVC configuration
- dot1q PVC configuration
- CLIPS PVC configuration
- Frame Relay PVC configuration
- link group configuration
- port configuration

1.56.3 Syntax Description

in	Overrides the policy rate specified in the policy attached to this circuit for incoming packets.
out	Overrides the policy rate specified in the policy attached to this circuit for outgoing packets.
kbps	Rate in kilobits per second. The range of values is 5 to 10,000,000.
burst ***bytes***	Burst tolerance in bytes. The range of values is 1 to 4,250,000,000.
excess-burst ***bytes***	Optional. Excess burst tolerance in bytes. The range of values is 1 to 4,250,000,000.
queuing-burst ***bytes***	Optional. Queuing burst tolerance in bytes. The range of values is 1 to 8,000,000. This construct is only available for MDRR policies. By default, the queuing burst value is not set. If the queuing burst is not specified, only the rate of the MDRR binding is modified by the **rate circuit out** command.

1.56.4 Default

The circuit rate is based on the policy rate as specified by the attached QoS policy.

1.56.5 Usage Guidelines

Use the **rate circuit** command to specify a different rate for a circuit that has a QoS metering, policing, MDRR, or PWFQ policy attached to it. The rate that you specify for the circuit overrides the rates specified by the attached metering, policing, MDRR, and PWFQ policies.

This command allows you to attach the same policy to a number of circuits, but specify a different rate for each circuit.

This command is not supported for dynamic 802.1Q permanent virtual circuits (PVCs).

Note: Configuring the rate circuit command on an ATM port is not supported. To limit ATM traffic, configure this command on ATM PVCs.

Note: The application of a different rate in either direction occurs only while you have attached the appropriate QoS policy to the circuit.

Consider the following rules when configuring rate limiting for circuits:

- When a dynamic circuit does not share its handle with a static circuit:
  - The dynamic binding is rate limited by the Dynamic QoS Parameter (DQP).
  - The static binding is rate limited as configured by the **`rate circuit out`** command.
- When a dynamic and static circuit share the same handle and the binding is attached to the dynamic (or subscriber) circuit, rate limiting configuration that is done on the static circuit (using the **`rate circuit out`** command) is not applied to the dynamic circuit binding.

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

1.56.6 Examples

The following example shows how to changes the rate for port `1` to 2,000 kbps:

```
[local]Redback(config)#port ethernet 4/1
[local]Redback(config-port)#qos policy metering example2
[local]Redback(config-port)#rate circuit out 2000
```

The following example changes the rate for an 802.1q PVC on port `1`. Here, the **`rate circuit out`** command modifies the rate of the metering binding `my-meter1` and MDRR binding `my-mdrr1` to 1000 kbps. The burst on the metering binding is set to 2000 kbps, and the queuing burst on the MDRR binding is set to 1200 kbps:

```
[local]Redback(config)#port ethernet 5/1
[local]Redback(config-port)#encapsulation dot1q
[local]Redback(config-port)#dot1q pvc 1
[local]Redback(config-port)#qos policy metering my-meter1
[local]Redback(config-port)#qos policy metering my-mdrr1
[local]Redback(config-port)#rate circuit out 1000 burst 2000 queuing-burst 1200
```

1.57 rate-factor

rate-factor *percent*

{no | default} rate-factor

1.57.1 Purpose

Defines the percentage of bandwidth for a specific access-line type that is unavailable to traffic on the circuit, port, or subscriber record to which the quality of service (QoS) policy is attached.

1.57.2 Command Mode

- overhead profile configuration
- overhead type configuration

1.57.3 Syntax Description

percent	Percentage of overhead for this access-line type. The range of values is 1 to 100; the default value is 0.

1.57.4 Default

Overhead on the access line is 0%, which allows full bandwidth usage.

1.57.5 Usage Guidelines

Use the **rate-factor** command to define the percentage of bandwidth for a specific access-line type that is unavailable to traffic on the circuit, port, or subscriber record to which the QoS policy is attached.

Use the `no` or `default` form of this command to remove the percentage from the access-line configuration and assume for rate enforcement calculation purposes that all the bandwidth specified by the `qos rate maximum` command in the PWFQ policy is available to traffic on the circuit.

Note: The maximum rate set by the **qos rate** command (in port configuration mode) is the rate at which the port, 802.1Q tunnel, or 802.1Q permanent virtual circuit (PVC) operates; any priority weighted fair queuing (PWFQ) queue or circuit with a PWFQ policy is limited by the rate specified by that command for the circuit. Also, the sum of all traffic on the port carried by the queues belonging to the circuits or subscribers is limited to the rate specified by that command.

1.57.6 Examples

The following example shows how to configure an overhead profile for `example1`, and set the default rate factor to `15`, a reserve value to `8`, and the encapsulation type to `pppoa-llc`. After you set the overhead profile with default values, you configure `adsl1` and `vdsl1` with custom encapsulation and reserve values with a rate factor of 20%:

```
[local]Redback(config)#qos profile example1 overhead
[local]Redback(config-profile-overhead)#rate-factor 15
[local]Redback(config-profile-overhead)#encaps-access-line pppoa-llc
[local]Redback(config-profile-overhead)#reserved 8
[local]Redback(config-profile-overhead)#type adsl1
[local]Redback(config-type-overhead)#rate-factor 20
```

1.58 rate-limit ccod

rate-limit ccod ***rate-limit*** **burst** ***burst-limit***

{**no** | **default**} **rate-limit ccod**

1.58.1 Purpose

Enables rate limiting and specifies the rate and burst limits for Point-to-Point Protocol (PPP) over Ethernet over Asynchronous Transfer Mode (PPPoEoA) Active Discovery Initiation (PADI) packets and PPP over Asynchronous Transfer Mode (PPPoA) Configure-Request packets that arrive at the SmartEdge router over circuit creation on demand (CCOD) Asynchronous Transfer Mode (ATM) permanent virtual circuits (PVCs).

1.58.2 Command Mode

card configuration

1.58.3 Syntax Description

rate-limit	Maximum rate in packets per second (pps) at which the packets can be received. The range of values is 0 to 1,000; the default value is 120 on the XCRP4 Controller card.
burst ***burst-limit***	Maximum number of packets that can be received during a short burst, in pps. The range of values is 0 to 1,000; the default value is 175 on the XCRP4 Controller card.

1.58.4 Default

Rate limiting for PPPoEoA PADI packets and PPPoA Configure-Request packets is enabled using the default burst and rate values.

1.58.5 Usage Guidelines

Use the **rate-limit ccod** command to enable rate limiting and specify the rate and burst limits for PPPoEoA PADI packets and PPPoA Configure-Request packets that arrive at the SmartEdge router on on-demand ATM PVCs. By specifying the rate and burst limit values, you can establish finer control over the rate of these kinds of subscriber sessions. For example, applying card-level rate-limiting can improve the bringup rate for PPPoA and PPPoEoA subscribers when many subscribers attempt to connect simultaneously.

Use the **show rate-limit card** command (in any mode) to display the current configuration of rate limiting; see the *Command List*.

Note: The rate limit and burst limit values cannot be configured independently.

Use the **default** form of this command to enable rate limiting with the default rate and burst limits.

Use the **no** form of this command to disable rate limiting.

1.58.6 Examples

The following example shows how to configure the rate limit of PPPoEoA PADI packets and PPPoA Configure-Request packets to `500` and the burst limit to `999`:

```
[local]Redback(config-card)#rate-limit ccod 500 burst 999
```

1.59 rate-limit circuit dhcp

rate-limit circuit dhcp ***num*** **interval** ***interval-value*** **drop-interval** ***drop-interval-value*** {**per-mac** | **per-mac-and-relay**}

no rate-limit circuit dhcp ***num*** **interval** ***interval-value*** **drop-interval** ***drop-interval-value*** {**per-mac** | **per-mac-and-relay**}

1.59.1 Purpose

Limits the number of Dynamic Host Configuration Protocol (DHCP) packets that the system accepts in an interval for each MAC address or each unique combination of MAC address and DHCP relay server address on a circuit.

1.59.2 Command Mode

Global configuration

1.59.3 Syntax Description

num	Number of DHCP packets allowed on each circuit during the specified interval. The range of values is 1 to 255.
interval ***interval-value***	Specifies the interval, in seconds, during which the system counts the packets. The range of values is 1 to 127.
drop-interval ***drop-interval-value***	Specifies the interval, in seconds, during which packets are dropped, if the allowed number of messages was exceeded in the previous interval. The range of values is 1 to 127.
per-mac	Limits the rate of the packets based on each MAC address on a circuit.
per-mac-and-relay	Limits the packets based on a unique combination of the MAC address and DHCP relay server address on a circuit. Use the **per-mac-and-relay** keyword when you have an RG that uses duplicate MAC addresses.

1.59.4 Default

The SmartEdge router does not limit the number of DHCP packets that it accepts on a circuit.

1.59.5 Usage Guidelines

Use the **rate-limit circuit dhcp** command to limit the number of DHCP packets that the system accepts in an interval for each MAC address or each unique combination of MAC address and DHCP relay server address on a circuit.

The operating system does not distinguish between DHCP Discover, Request, Release, NAK, or Inform messages when it limits the number DHCP packets on a circuit.

The **rate-limit circuit dhcp** command is supported only on ATM and Ethernet traffic cards that are PPA2-based and above.

The **rate-limit circuit dhcp** command supports access link groups but not Ethernet or 802.1Q link groups. That is, the **rate-limit circuit dhcp** command affects only incoming DHCP packet traffic on the access side and includes both regular PVCs and link-group aggregated PVCs.

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

1.59.6 Examples

The following example shows how to accept 155 DHCP messages for the unique combination of the MAC address and DHCP relay server address in every 3-second interval. If more than 155 DHCP packets are received during the interval, all DHCP packets are dropped for 4 seconds starting at the time when the limit was exceeded. For example, if the 156th packet is received at 2 seconds into the 3-second interval, then the count for the drop interval starts at 2 seconds and stops at 6 seconds. Following that, 155 DHCP packets are allowed for the unique combination of the MAC address and DHCP relay server address for the next 3-second interval:

```
[local]Redback(config)#rate-limit circuit dhcp 155 interval
3 drop-interval 4 per-mac-and-relay
```

1.60 rate-limit circuit dhcpv6

rate-limit circuit dhcpv6 ***num*** **interval** ***interval-value*** **drop-interval** ***drop-interval***

no rate-limit circuit dhcpv6 ***num*** **interval** ***interval-value*** **drop-interval** ***drop-interval***

1.60.1 Purpose

Limits the number of Dynamic Host Configuration Protocol Version 6 (DHCPv6) packets that the system accepts in an interval on a circuit.

1.60.2 Command Mode

Global configuration

1.60.3 Syntax Description

num	Number of DHCPv6 packets allowed on each circuit during the specified interval. The range of values is 1 to 255.
interval ***interval-value***	Specifies the interval, in seconds, during which the system counts the packets. The range of values is 1 to 127.
drop-interval ***drop-interval-value***	Specifies the interval, in seconds, during which packets are dropped, if the allowed number of messages was exceeded in the previous interval. The range of values is 1 to 127.

1.60.4 Default

The SmartEdge router does not limit the number of DHCPv6 packets that it accepts on a circuit.

1.60.5 Usage Guidelines

Use the **rate-limit circuit dhcpv6** command to limit the number of DHCPv6 packets that the system accepts in an interval on each circuit. When enabled, rate limiting is performed on DHCPv6 packets arriving on every circuit , or virtual circuit in the case of PPPoE, even though it is configured on the parent circuit at the card level.

The operating system does not distinguish between types of messages when it limits the number DHCPv6 packets on a circuit.

This command applies only to line cards that support IPv6 single-stack or dual-stack packets.

The **rate-limit circuit dhcpv6** command supports access link groups but not Ethernet or 802.1Q link groups. That is, this command affects only incoming DHCPv6 packet traffic on the access side and includes both regular PVCs and link-group aggregated PVCs.

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

1.60.6 Examples

The following example shows how to accept 155 DHCPv6 messages in every 3-second interval on each circuit. If more than 155 DHCPv6 packets are received during the interval, all DHCPv6 packets are dropped for 4 seconds, starting at the time when the limit was exceeded. For example, if the 156th packet is received at 2 seconds into the 3-second interval, then the count for the drop interval starts at 2 seconds and stops at 6 seconds. Following that, 155 DHCPv6 packets are for the circuit for the next 3-second interval:

```
[local]Redback(config)#rate-limit circuit dhcpv6 155
interval 3 drop-interval 4
```

1.61 rate-limit circuit nd

rate-limit circuit nd ***num*** **interval** ***interval-value***
drop-interval ***drop-interval***

no rate-limit circuit nd ***num*** **interval** ***interval-value***
drop-interval ***drop-interval***

1.61.1 Purpose

Limits the number of Neighbor Discovery (ND) packets that the system accepts in an interval on a circuit.

1.61.2 Command Mode

Global configuration

1.61.3 Syntax Description

num	Number of ND packets allowed on each circuit during the specified interval. The range of values is 1 to 255.
interval ***interval-value***	Specifies the interval, in seconds, during which the system counts the packets. The range of values is 1 to 127.
drop-interval ***drop-interval-value***	Specifies the interval, in seconds, during which packets are dropped, if the allowed number of messages was exceeded in the previous interval. The range of values is 1 to 127.

1.61.4 Default

The SmartEdge router does not limit the number of ND packets that it accepts on a circuit.

1.61.5 Usage Guidelines

Use the **rate-limit circuit nd** command to limit the number of ND packets that the system accepts in an interval on each circuit. When enabled, rate limiting is performed on ND packets coming on every circuit, or virtual circuit in the case of PPPoE, even though it is configured on the parent circuit at the card level.

The operating system does not distinguish between types of messages when it limits the number ND packets on a circuit.

This command applies only to line cards that support IPv6 single-stack or dual-stack packets.

The **rate-limit circuit nd** command supports access link groups but not Ethernet or 802.1Q link groups. That is, this command affects only incoming ND packet traffic on the access side and includes both regular PVCs and link-group aggregated PVCs.

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

1.61.6 Examples

The following example shows how to accept 155 ND messages in every 3-second interval on each circuit. If more than 155 ND packets are received during the interval, all ND packets are dropped for 4 seconds starting at the time when the limit was exceeded. For example, if the 156th packet is received at 2 seconds into the 3-second interval, then the count for the drop interval starts at 2 seconds and stops at 6 seconds. Following that, 155 ND packets are for the circuit for the next 3-second interval:

```
[local]Redback(config)#rate-limit circuit nd 155 inte
rval 3 drop-interval 4
```

1.62 rate-limit dhcp

rate-limit dhcp *rate-limit* **burst** *burst-limit*

{**no** | **default**} **rate-limit dhcp**

1.62.1 Purpose

Enables rate limiting and specifies the rate and burst limits for Dynamic Host Configuration Protocol (DHCP) packets that arrive at the SmartEdge router.

1.62.2 Command Mode

card configuration

1.62.3 Syntax Description

rate-limit	Maximum rate in packets per second (pps) at which the packets can be received. The range of values is 0 to 4294967295 pps; the default value is 4294967295 pps.
burst ***burst-limit***	Maximum number of packets that can be received during a short burst. The range of values is 0 to 4294967295 pps; the default value is 4294967295 pps.

1.62.4 Default

Rate limiting for packets is enabled using the default rate and burst values.

1.62.5 Usage Guidelines

Use the **rate-limit dhcp** command to enable rate limiting and specify the rate and burst limits for DHCP packets that arrive at the SmartEdge router. By specifying the rate and burst limit values, you can establish finer control over the rate of these kinds of subscriber sessions.

Use the **show rate-limit card** command (in any mode) to display the current configuration of rate limiting. This command is described in the Command List.

Note: You cannot configure the rate limit and burst limit values independently.

Table 3 shows the traffic cards supported for the **rate-limit dhcp** command.

Table 3 Traffic Cards Supported for the rate-limit dhcp Command

Type	Traffic Cards Supported
ATM	• ATM OC-3c/STM-1c (2-port) • ATM OC-3c/STM-1c (8-port) • ATM OC-12c/STM-4c (2-port)
Ethernet	• Gigabit Ethernet (4-port) • Advanced Gigabit Ethernet (4-port) • Gigabit Ethernet 3 (4-port) • Gigabit Ethernet 1020 (10-port • Gigabit Ethernet 1020 (20-port) • Gigabit Ethernet (5-port) • Gigabit Ethernet (20-port) • Gigabit Ethernet DDR (10-port) • 10 Gigabit Ethernet (1-port) • 10 Gigabit Ethernet (4-port) • 10 Gigabit Ethernet/OC-192c DDR (1-port)

Use the **no** form of this command to disable rate limiting.

Use the **default** form of this command to set the rate and burst limits to default values.

1.62.6 Examples

The following example shows how to configure the rate limit for DHCP packets to `500` and the burst limit to `999`:

```
[local]Redback(config-card)#rate-limit dhcp 500 burst 999
```

1.63 rate-limit dhcpv6

rate-limit dhcpv6 ***rate-limit*** **burst** ***burst-limit***

{no | default} rate-limit dhcpv6

1.63.1 Command Mode

card configuration

1.63.2 Purpose and Usage Guidelines

Use the **rate-limit dhcpv6** command to specify the rate and burst limits for DHCPv6 packets that arrive at the SmartEdge router. This command limits the impact of DHCPv6-packet denial of service (DoS) attacks.

This command applies only to line cards that support IPv6 single-stack or dual-stack packets.

1.63.3 Syntax Description

rate-limit	Maximum rate in packets per second (pps) at which the DHCPv6 packets can be received. The range of values is 0 to 4,294,967,295; the default value is 4,294,967,295.
burst-limit	Maximum number of DHCPv6 packets that can be received during a short burst, in pps. The range of values is 0 to 4,294,967,295; the default value is 4,294,967,295.

1.63.4 Default

The **default** form of this command sets the ***rate-limit*** and ***burst-limit*** to 4,294,967,295 pps.

The **no** form of this command removes the rate limit.

1.63.5 Examples

The following example shows how to configure the rate limit of DHCPv6 packets to `500` and the burst limit to `999`:

```
[local]Redback(config-card)#rate-limit dhcpv6 500 burst 999
```

1.64 rate-limit nd

rate-limit nd *rate-limit* burst *burst-limit*

{no | default} rate-limit nd *rate-limit* burst *burst-limit*

1.64.1 Command Mode

card configuration

1.64.2 Purpose and Usage Guidelines

Use the **rate-limit nd** command to specify the rate and burst limits for ND packets that arrive at the SmartEdge router. This command limits the impact of ND-packet denial of service (DoS) attacks.

1.64.3 Syntax Description

rate-limit	Maximum rate in packets per second (pps) at which ND packets can be received. The range of values is 0 to 4,294,967,295; the default value is 4,294,967,295.
burst-limit	Maximum number of ND packets that can be received during a short burst, in pps. The range of values is 0 to 4,294,967,295; the default value is 4,294,967,295.

1.64.4 Default

The **default** form of this command sets the ***rate-limit*** and ***burst-limit*** to 4,294,967,295 pps.

The **no** form of this command removes the rate limit.

1.64.5 Examples

The following example shows how to configure the rate limit of ND packets to `500` and the burst limit to `999`:

```
[local]Redback(config-card)#rate-limit nd 500 burst 999
```

1.65 rate-limit padi

rate-limit padi ***rate-limit*** **burst** ***burst-limit***

{**no** | **default**} **rate-limit padi**

1.65.1 Purpose

Enables rate limiting and specifies the rate and burst limits for Point-to-Point Protocol over Ethernet (PPPoE) Active Discovery Initiation (PADI) packets that arrive at the SmartEdge router.

1.65.2 Command Mode

card configuration

1.65.3 Syntax Description

rate-limit	Maximum rate in packets per second (pps) at which the packets can be received. The range of values is 0 to 1000; the default value is 75.
burst ***burst-limit***	Maximum number of packets that can be received during a short burst, in pps. The range of values is 0 to 1000; the default value is 100.

1.65.4 Default

Rate limiting for PADI packets is enabled using the default burst and rate values.

1.65.5 Usage Guidelines

Use the **rate-limit padi** command to enable rate limiting and specify the rate and burst limits for PADI packets that arrive at the SmartEdge router. By specifying the rate and burst limit values, you can establish finer control over the rate of these kinds of subscriber sessions.

If both PADI and PPP LCP Configure-Request rate limiting are enabled on the same card, the first protocol to be processed is rate limited. For example, if a PADI packet arrives on a PPPoEoA circuit, the PADI packet is rate limited first; if the PADI packet is allowed to pass through, no further rate limiting is applied. If the circuit has PPPoA encapsulation, only the first LCP Configure-Request packet is rate limited.

Use the **show rate-limit card** command (in any mode) to display the current configuration of rate limiting. The **show rate-limit card** command is described in the Command List.

Note: The rate limit and burst limit values cannot be configured independently.

Use the **no** form of this command to disable rate limiting.

Use the **default** form of this command to enable rate limiting with the default rate and burst limits.

1.65.6 Examples

The following example shows how to configure the rate limit of PADI packets to `500` and the burst limit to `999`:

```
[local]Redback(config-card)#rate-limit padi 500 burst 999
```

1.66 rate-limit ppp-lcp-confreq

rate-limit ppp-lcp-confreq ***rate-limit*** **burst** ***burst-limit***

{**no** | **default**} **rate-limit ppp-lcp-confreq**

1.66.1 Purpose

Enables rate limiting and specifies the rate and burst limits for Point-to-Point (PPP) Link Control Protocol (LCP) Configure-Request packets that arrive at the SmartEdge router on static Asynchronous Transfer Mode (ATM) permanent virtual circuits (PVCs).

1.66.2 Command Mode

card configuration

1.66.3 Syntax Description

rate-limit	Maximum rate in packets per second (pps) at which the PPP LCP Configure-Request packets can be received. The range of values is 0 to 1000; the default value is 350 on the XCRP4 Controller card.
burst ***burst-limit***	Maximum number of PPP LCP Configure-Request packets that can be received during a short burst, in pps. The range of values is 0 to 1000; the default value is 500 on the XCRP4 Controller card.

1.66.4 Default

Rate limiting for PPP LCP Configure-Request packets is enabled using the default burst and rate values.

1.66.5 Usage Guidelines

Use the **rate-limit ppp-lcp-confreq** command to enable rate limiting and specify the rate and burst limits for PPP LCP Configure-Request packets that arrive at the SmartEdge router on static ATM PVCs. By specifying the rate and burst limit values, you can establish finer control over the rate of these kinds of subscriber sessions. For example, applying card-level rate-limiting can improve the bringup rate for PPP over ATM (PPPoA) and PPP over Ethernet over ATM (PPPoEoA) subscribers when many subscribers attempt to connect simultaneously.

If both PADI and PPP LCP Configure-Request rate limiting are enabled on the same card, the first protocol to be processed is rate limited. For example, if a PADI packet arrives on a PPPoEoA circuit, the PADI packet is rate limited first; if the PADI packet is allowed to pass through, no further rate limiting is applied. If the circuit has PPPoA encapsulation, only the first LCP Configure-Request packet is rate limited.

Use the **show rate-limit card** command (in any mode) to display the current configuration of rate limiting. The **show rate-limit card** command is described in the *Command List*.

Note: The rate limit and burst limit values cannot be configured independently.

Use the **no** form of this command to disable rate limiting.

Use the **default** form of this command to enable rate limiting with the default rate and burst limits.

1.66.6 Examples

The following example shows how to configure the rate limit of PPP LCP Configure-Request packets to `500` and the burst limit to `999`:

```
[local]Redback(config-card)#rate-limit ppp-lcp-confreq 500 burst 999
```

1.67 rate percentage

rate percentage *percent-rate* [**counters**]

no rate percentage

1.67.1 Purpose

Assigns a percentage of the overall policy rate to this class of traffic on the circuit, or port, or subscriber record to which the quality of service (QoS) policy is attached and accesses policy class rate configuration mode.

1.67.2 Command Mode

- policy group class configuration
- PWFQ policy configuration

1.67.3 Syntax Description

percent-rate	Relative class rate, as a percentage of the policy rate, for this class.
counters	Optional. Logs statistics related to packets that conform to or exceed the rate.

1.67.4 Default

No rate percentage is specified for this class.

1.67.5 Usage Guidelines

Use the **rate percentage** command to assign a percentage (a relative class rate) of the overall policy rate to this class of traffic on the circuit, or port, or subscriber record to which the QoS policy is attached, and access policy class rate configuration mode. The percentage applies to the policy rate, burst, and excess burst values.

Use the **no** form of this command to remove the rate percentage from this class configuration.

Note: The maximum rate set by the **qos rate** command (in port configuration mode) is the rate at which the port, 802.1Q tunnel, or 802.1Q permanent virtual circuit (PVC) operates; any priority weighted fair queuing (PWFQ) queue or circuit with a PWFQ policy is limited by the rate specified by that command for the circuit. Also, the sum of all traffic on the port carried by the queues belonging to the circuits or subscribers is limited to the rate specified by that command.

For priority weighted fair queuing (PWFQ) queues with a PWFQ policy, the sum of all TM queue priority group rates for a node can oversubscribe the configured global policy rate.

1.67.6 Examples

The following example assign `25%` of the policy rate to the `realtime` class:

```
[local]Redback(config)#qos policy rate-incoming policing
[local]Redback(config-policy-policing)#rate informational 6000000 burst 10000 counters
[local]Redback(config-policy-policing)#access-group Class local
[local]Redback(config-policy-policy-acl)#class realtime
[local]Redback(config-policy-policy-acl-class)#rate percentage 25
```

By including the **counters** keyword in the **rate percentage** command, you can use the **show circuit counters** command (in any mode) with the **detail** keyword to display the number of packets that conform to the rate percentage and the number of packets that exceed that rate percentage.

1.68 rbak-term-ec

rbak-term-ec *term-error-code* **ietf-attr-49** *error-code*

no rbak-term-ec *term-error-code*

1.68.1 Purpose

Remaps a Redback account (session) termination error code to a different RADIUS attribute 49 (Acct-Terminate-Cause) error code.

1.68.2 Command Mode

terminate error cause configuration

1.68.3 Syntax Description

term-error-code	Redback account termination error code to be remapped.
ietf-attr-49 *error-code*	Attribute 49 error code to which the Redback termination error code is remapped.

1.68.4 Default

No Redback account termination error codes are remapped.

1.68.5 Usage Guidelines

Use the **rbak-term-ec** command to remap a Redback account (session) termination error code to a different RADIUS attribute 49 (Acct-Terminate-Cause) error code. *RADIUS Attributes* lists the default

mapping of Redback account termination error codes to RADIUS attribute 49 (Acct-Terminate-Cause) error codes. RADIUS attribute 49 error codes and their definitions are included in RFC 2866, ***RADIUS Accounting***.

Use the **no** form of this command to specify the default RADIUS attribute 49 error code for the specified Redback account termination error code.

1.68.6 Examples

The following example remaps Redback account termination code `24` (Authentication failed) from its default RADIUS attribute 49 error code 17 (User error), to the RADIUS attribute 49 error code `2` (network access server [NAS] error).

```
[local]Redback(config)#radius attribute acct-terminate-cause remap
[local]Redback(config-term-ec)#rbak-term-ec 24 ieft-attr-49 2
```

1.69 reachable-time

reachable-time ***duration***

{**no** | **default**} **reachable-time**

1.69.1 Purpose

Specifies the value for the Reachable Time field in Router Advertisement (RA) messages.

1.69.2 Command Mode

- ND router configuration
- ND router interface configuration

1.69.3 Syntax Description

duration	Value for the Reachable Time field (in milliseconds). The range of values is 0 to 3,600,000; the default value is 0 (unspecified).

1.69.4 Default

The duration is unspecified in any RA messages.

1.69.5 Usage Guidelines

Use the **reachable-time** command to specify the value for the Reachable Time field in RA messages. This value is the time this Neighbor Discovery (ND) router or ND router interface assumes that a neighbor is reachable. In ND router configuration mode, this command specifies the global value for all interfaces; in ND router interface mode, it specifies the value for this ND router interface. If specified, the parameters for an interface override the global parameters.

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

1.69.6 Examples

The following example specifies a reachable time of `1800` milliseconds for all interfaces for the ND router:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#router nd
[local]Redback(config-nd-if)#reachable-time 1800
```

The following example specifies a reachable time of `3600` milliseconds for the `int1` ND router interface:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#router nd
[local]Redback(config-nd)#interface int1
[local]Redback(config-nd-if)#reachable-time 3600
```

1.70 reauthorize

reauthorize {**bulk** *index-num* | **session** *sess-id* | **username** *subscriber*}

1.70.1 Purpose

Modifies subscriber attributes in real time during an active session, using RADIUS authentication.

1.70.2 Command Mode

Exec (10)

1.70.3 Syntax Description

bulk ***index-num***	Index number of the reauthorization record in the user database on the RADIUS server. The SmartEdge router attaches the name of the context in which reauthorization occurs to the ***index-num*** value. The range of values is 1 to 65,535.
session ***sess-id***	RADIUS accounting ID attribute (Acct-Session-Id) value that identifies an active subscriber session.
username ***subscriber***	Structured subscriber name in the form ***sub-name@ctx-name***. You can specify a string of up to 253 characters, including the separator character. The separator character and format that you specify are the defaults, as shown, or are defined by the **aaa username-format** command (in global configuration mode).

1.70.4 Default

None

1.70.5 Usage Guidelines

Use the **reauthorize** command to modify subscriber attributes in real time during an active session. Reauthorization does not require Point-to-Point Protocol (PPP) renegotiation and does not interrupt or drop the session.

Note: Before entering this command, you must first enable subscriber reauthorization with the **aaa reauthorization bulk** command (in context configuration mode).

Note: The SmartEdge router appends the context name to the subscriber name when sending reauthorization messages; for example, `joe@local`.

Note: This command and the **policy-refresh** command (in exec mode) each perform the same function. However, the **policy-refresh** command uses the command-line interface (CLI) to perform the update instead of the RADIUS authentication process.

Table 4 lists the standard RADIUS attributes that are reauthorized when you enter this command.

Table 4 Standard RADIUS Attributes Supported by Reauthorization

#	Attribute Name	Description
11	Filter-Id	Filters inbound or outbound traffic through an access control list (ACL).
25	Class	Forwards the information sent by the RADIUS server to the SmartEdge router, without interpretation, in subsequent accounting messages to the RADIUS accounting server for that subscriber session.
27	Session-Timeout	Sets the in-service time allowed before termination of the session.
28	Idle-Timeout	Sets the idle time allowed before termination of the session.

Table 5 lists the vendor-specific attributes (VSAs) provided by Ericsson AB that are reauthorized when you enter this command.

Table 5 Vendor VSA Attributes Supported by Reauthorization

#	VSA Name	Description
33	Mcast-Send	Defines whether the subscriber can send multicast packets.
34	Mcast-Receive	Defines whether the subscriber can receive multicast packets.
35	Mcast-MaxGroups	Specifies the maximum number of multicast groups of which the subscriber can be a member.
87	QoS-Policy-Policing	Attaches a QoS policing policy to the subscriber session.
88	QoS-Policy-Metering	Attaches a QoS metering policy to the subscriber session.
89	QoS-Policy-Queuing	Attaches a QoS queuing (scheduling) policy to the subscriber session.
90	Igmp-Service-Profile-Id	Applies an IGMP service profile to the subscriber session.
92	Forward-Policy	Attaches an in or out forward policy to the subscriber session.
101	Shaping-Profile-Name	Indicates the name of the ATM shaping profile.
102	Bridge-Profile-Name	Indicates the name of the bridge profile.
107	HTTP-Redirect-Profile-Name	Indicates the name of the HTTP redirect profile.
113	Session-Traffic-Limit	Specifies that inbound or outbound traffic be limited.

For details about these attributes, see *RADIUS Attributes*.

Note: You must configure at least one RADIUS server in the current context before any messages can be sent to it. To configure the server, use the **radius server** command (in context configuration mode); for more information, see *Configuring RADIUS*.

Note: To enable RADIUS authentication, you must use the **aaa authentication subscriber** command (in context configuration mode).

1.70.6 Examples

The following example displays a subscriber record on a RADIUS server. The subscriber has requested a new service that is translated to a particular session `timeout` value:

```
#reauth of absolute timeout

reauth-501@ABC User-Password==”redback”

    Service-Type=Outbound-User,

    Reauth_String=”2;pppoe1@local;27;1000;”
```

Before entering the **reauthorize** command, the subscriber record appears as follows:

```
[local]Redback>show subscribers active


pppoe1@local

    Circuit 13/1 vpi-vci 0 33

    Internal Circuit 13/1:1023:63/1/2/22

    Current port-limit unlimited

    ip address 10.1.1.4
```

The following example reauthorizes the subscriber session, `pppoe1@local`, after the new value for the RADIUS attribute 27 has been sent to the RADIUS server:

```
[local]Redback#reauthorize username pppoe1@local
[local]Redback>show subscribers active

pppoe1@local
    Circuit 13/1 vpi-vci 0 33
    Internal Circuit 13/1:1023:63/1/2/22
    Current port-limit unlimited
    ip address 10.1.1.4
    timeout absolute 1000
```

1.71 receive

receive {**permit** | **deny**}

no receive {**permit** | **deny**}

1.71.1 Purpose

Configures the setting in the IGMP snooping profile that controls the ability of the associated circuits to receive multicast data.

1.71.2 Command Mode

- IGMP snooping profile configuration

1.71.3 Syntax Description

permit	Permits circuits to receive multicast data.
deny	Does not permit circuits to receive multicast data.

1.71.4 Default

The receipt of multicast data is permitted on all circuits.

1.71.5 Usage Guidelines

Use the **receive** command to configure the setting in the IGMP snooping profile that controls the ability of the associated circuits to receive multicast data. The configuration applies to all circuits that are associate with the specified IGMP snooping profile.

Use the **no** form of this command to return the IGMP snooping profile to the default setting in which the receipt of multicast data is permitted on all circuits.

1.71.6 Examples

The following example shows how to disable the receipt of multicast data by all circuits attached to an IGMP snooping profile called `sanjose1`:

```
[local]Redback#configure

[local]Redback(config)#igmp snooping profile sanjose1

[local]Redback(config-igmp-snooping-profile)#receive deny
```

The following example shows how to permit the receipt of multicast data by all circuits attached to an IGMP snooping profile called `sanjose1`:

```
[local]Redback#configure

[local]Redback(config)#igmp snooping profile sanjose1

[local]Redback(config-igmp-snooping-profile)#receive permit
```

1.72 receiver

receiver *ip-addr* {**primary** | **secondary**} **mechanism** {**ftp** | **sftp** | **scp**} **login** *login-name* {**password** *password* | **encrypted** *password* | **nopassword**}

no receiver *ip-addr* {**primary** | **secondary**}

1.72.1 Purpose

Specifies the remote file servers where bulkstats files for this policy are stored.

1.72.2 Command Mode

bulkstats configuration

1.72.3 Syntax Description

ip-addr	IP address of the bulkstats file server.
primary	Specifies that the file server is the primary receiver.
secondary	Specifies that the file server is the secondary receiver.
mechanism ftp	Specifies the file transfer method as File Transfer Protocol (FTP).
mechanism sftp	Specifies the file transfer method as Secure Shell FTP (SFTP).
mechanism scp	Specifies the file transfer method as Secure Copy Protocol (SCP).
login ***login-name***	Login name to be used for file transfer.
password ***password***	Password to be used with the logon name.
encrypted ***password***	Encrypted password to be entered with the logon name. (The password is encrypted while saving the configuration.)
nopassword	Specifies that a password is not required with the logon name.

1.72.4 Default

No server is specified to receive bulkstats.

1.72.5 Usage Guidelines

Use the **receiver** command to specify the remote file servers where bulk statistics (bulkstats) files for this policy are stored.

If a transfer to the primary file server that receives bulkstats fails, a transfer to the secondary receiver is immediately attempted. If the transfer to the secondary receiver fails, the SmartEdge router waits five minutes before making another attempt. Retries continue every five minutes until a transfer succeeds.

Note: Whenever a transfer to any bulkstats file server fails, a Simple Network Management Protocol (SNMP) trap is generated.

Use the **no** form of this command to delete a previously configured bulkstats remote file server. If you use the **no** form of this command while bulkstats

collection is running, no data is transmitted to the deleted file server until you define a new bulkstats file server.

1.72.6 Examples

The following example identifies the server at IP address, `198.168.145.99`, as the primary bulkstats file server; the logon account is `snmp` and its password is `snmp`:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#bulkstats policy bulk
[local]Redback(config-bulkstats)#receiver 198.168.145.99 primary mechanism ftp login snmp
password snmp
```

To see how this information displays, see the example for the **`show bulkstats`** command in *Configuring Bulkstats*.

1.73 record-route

`record-route`

`no record-route`

1.73.1 Purpose

Configures a Resource Reservation Protocol (RSVP) label-switched path (LSP) to actively record the routes through which the LSP forwards packets.

1.73.2 Command Mode

RSVP LSP configuration

1.73.3 Syntax Description

This command has no keywords or arguments.

1.73.4 Default

Route information is recorded.

1.73.5 Usage Guidelines

Use the **record-route** command to configure an RSVP LSP to actively record the routes through which the LSP forwards packets.

Use the **show rsvp lsp** command to display the detailed output containing information about the recorded route, which you can use for troubleshooting purposes, and to prevent routing loops.

Use the **no** form of this command to disable route recording for the RSVP LSP.

1.73.6 Examples

The following example configures the LSP, `test07`, to actively record the routes through which it forwards packets:

```
[local]Redback(config-ctx)#router rsvp
[local]Redback(config-rsvp)#lsp test07
[local]Redback(config-rsvp-lsp)#record-route
```

1.74 redirect destination circuit

redirect destination circuit ***dest-name***

no redirect destination

1.74.1 Purpose

Redirects packets to an output destination.

1.74.2 Command Mode

- forward policy configuration
- policy group class configuration

1.74.3 Syntax Description

dest-name	Output destination for redirected traffic.

1.74.4 Default

Packets are not redirected.

1.74.5 Usage Guidelines

Use the **redirect destination circuit** command to redirect packets to an output destination.

The destination name is the one that you specified for the circuit using the **forward output** command (in ATM PVC, Frame Relay PVC, GRE tunnel, or port configuration mode).

Use the **no** form of this command to disable the redirecting of packets.

Note: The **redirect destination circuit** command is only supported on Layer 3 circuits.

1.74.6 Examples

The following example redirects traffic to the output destination circuit OD15:

```
[local]Redback#config
[local]Redback(config)#forward policy RedirectPolicy
[local]Redback(config-policy-frwd)#redirect destination circuit OD15
```

1.75 redirect destination local

redirect destination local

no redirect destination

1.75.1 Purpose

In forward policy configuration mode, redirects packets not associated with a class to the HTTP server on the controller card.

In policy ACL configuration mode, redirects only packets associated with a class to the HTTP server on the controller card.

1.75.2 Command Mode

- forward policy configuration

- policy ACL class configuration

1.75.3 Syntax Description

This command has no keywords or arguments.

1.75.4 Default

Packets are not redirected.

1.75.5 Usage Guidelines

In forward policy configuration mode, use the **redirect destination local** command to redirect packets not associated with a class to the HTTP server on the controller card. In policy ACL configuration mode, use the **redirect destination local** command to redirect only packets associated with a class to the HTTP server on the controller card.

Use the **no** form of this command to disable the redirecting of packets.

1.75.6 Examples

The following example configures the forward policy, `Business-Redirect`, which redirects packets associated with the class, `Redirect`, to the HTTP server on the controller card:

```
[local]Redback(config)#forward policy Business-Redirect
[local]Redback(config-policy-frwd)#redirect destination local
[local]Redback(config-policy-frwd)#access-group bus-redirect local
[local]Redback(config-policy-group)#class Redirect
[local]Redback(config-policy-group)#redirect destination local
```

1.76 redirect destination next-hop

redirect destination {**ip** | **ipv6**} **nexthop** {*ipaddr* | **default**}

no redirect destination {*ip* | *ipv6*}

1.76.1 Purpose

Redirects packets to the specified IP address or to the packets' default destination IP address according to the routing table.

1.76.2 Command Mode

- forward policy configuration
- policy group class configuration

1.76.3 Syntax Description

ip-addr...	For IPv4, one to eight next-hop IP addresses in order of priority, including up to one default entry. Each entry in the list is an IP address in the form ***A.B.C.D***. The unspecified IPv4 address (0.0.0.0) or broadcast address (255.255.255.255) cannot be used in a redirect to IPv4 next hop. For IPv6, one or two IPv6 addresses and a single default argument. The command will be rejected if: • IPv4 next-hop list contains 0.0.0.0 or 255.255.255 • IPv6 next-hop list contains the " ::" address, the fe80::/10 prefix or duplicate addresses For IPv4 address duplicates are allowed, but using duplicates does not make sense.
default	Specifies that the packet's destination IP address should be used to forward the packet according to the routing table. When the **default** keyword is active, the packet is routed and not redirected.

1.76.4 Default

Packets are not redirected.

1.76.5 Usage Guidelines

Use the **redirect destination [ip|ipv6] next-hop** command to redirect packets to the specified IP address or to the packets' default destination IP address according to the routing table.

If an address is unreachable, then the next lower priority address is tried. From time to time, the system will try to return to the highest priority entry

available. The **default** keyword can be used in the next-hop list instead of an IP address to indicate that the destination IP address from the packet should be used when all higher priority next hops are unreachable. The **default** keyword can be first in the list, which means redirecting packets only when the normal route is unreachable.

Dual (IPv4 and IPv6) redirect action for the next-hop mode is supported in the default forward class. No other combination is allowed. The IPv4 and IPv6 redirect to next-hop commands can be set in any order. The definition of dual (IPv4 and IP6) action in a regular forward class is not supported. Instead, the user must use a distinct rule set in the IPv4 and IPv6 Policy ACLs that matches corresponding IPv4 and IPv6 forward classes.

Note: To modify the list of next-hop entries using the same IP type (IPv4 to IPv4 or IPv6 to IPv6), re-enter the entire **redirect destination [ip|ipv6] next-hop** command. Any attempted modification of the redirect to next-hop to a different IP type is rejected.

To modify a list of next hop entries to a different IP type (IPv4 to IPv6 or IPv6 to IPv4), first explicitly cancel the redirection using the **no redirect destination [ip|ipv6] next-hop** command. Then, enter the modified **redirect destination [ip|ipv6] next-hop** command.

Use the **no** form of this command to disable the redirecting of packets.

Note: The **redirect destination [ip|ipv6] next-hop** command is only supported on Layer 3 circuits.

1.76.6 Examples

The following example redirects traffic to the next-hop IP address, `10.1.1.1`. If that address is unreachable, the SmartEdge router redirects traffic to the next-hop IP address, `10.1.2.1`. If both addresses are unreachable, traffic is routed normally:

```
[local]Redback#config
[local]Redback(config)#forward policy RedirectPolicy
[local]Redback(config-policy-frwd)#redirect destination next-hop 10.1.1.1 10.1.2.1 default
```

The following example routes traffic normally. If the route is unavailable, traffic is redirected to the next-hop IP address, `10.1.1.1`:

```
[local]Redback#config
[local]Redback(config)#forward policy RedirectPolicy
[local]Redback(config-policy-frwd)#redirect destination next-hop default 10.1.1.1
```

The following example redirects traffic to the next-hop IP address, `192.1.1.1`. If that address is unreachable, the SmartEdge router attempts to redirect traffic

to the next-hop IP address, `10.1.1.1`.When it reaches the last IPv6 next-hop address that is unreachable, the PBR drops the packet and sends a destination unreachable ICMP message to the source:

```
[local]Redback#config
[local]Redback(config)#forward policy RedirectPolicy
[local]Redback(config-policy-frwd)#redirect destination next-hop 192.1.1.1 10.1.1.1
```

In the following example a forward policy for IPv6 traffic is configured with one IPv6 next-hop. If this IPv6 addresses is unreachable, traffic is dropped and a destination unreachable ICMPv6 message is sent to the source

```
[local]Redback#config
[local]Redback(config)#forward policy Demov6
[local]Redback(config-policy-frwd)#redirect destinatio
n ipv6 next-hop 3:5::5
```

The following example shows a definition of dual (IPv4 and IPv6) redirect action for the next-hop mode in the default forward class.

```
[local]Redback#config
[local]Redback(config)#forward policy Demo
[local]Redback(config-policy-frwd)#redirect destination ip next-hop 3.4.5.6
[local]Redback(config-policy-frwd)#redirect destination ipv6 next-hop 2001:0f68:::
[local]Redback(config-policy-frwd)#ip access-group Acl-demo-ipv4 local
 [local]Redback(config-policy-frwd)#ipv6 access-group Acl-demo-ipv6 local
[local]Redback(config-policy-group)#class http-ipv4
[local]Redback(config-policy-group)#redirect destination next-hop 10.1.1.1
[local]Redback(config-policy-group)#class http-ipv6
[local]Redback(config-policy-group)#redirect destination ipv6 next-hop 2001:0f68::
```

1.77 redistribute (BGP, IPv4)

redistribute { [**connected** | ipsec | **isis** *instance* [**level-1** | **level-2**] | **nat** | **ospf** *instance* [**internal** | [**external**] [**nssa-external**] | **rip** *instance* | **static** [**dvsr**] | **subscriber** [**address** | **static**]} [**route-map** *map-name*]

no redistribute {**connected** | ipsec |**isis** *instance* [**level-1** | **level-2**] | **nat** | **ospf** *instance* [**internal** | [**external**] [**nssa-external**] | **rip** *instance* | **static** [**dvsr**] | **subscriber** [**address** | **static**]} [**route-map** *map-name*]

1.77.1 Purpose

Redistributes IPv4 routes learned through other routing protocols into the Border Gateway Protocol (BGP) routing domain.

1.77.2 Command Mode

BGP address family configuration

1.77.3 Syntax Description

connected	Redistributes routes from directly attached networks into the BGP routing domain.
ipsec	Redistributes IPsec routes into the BGP routing domain.
isis *instance*	Intermediate System-to-Intermediate System (IS-IS) instance name. Redistributes routes from the specified IS-IS routing instance into the BGP routing domain.
level-1	Optional. Specifies IS-IS level 1 routing.
level-2	Optional. Specifies IS-IS level 2 routing.
nat	Redistributes network address translation (NAT) routes into the BGP routing domain.
ospf *instance*	Open Shortest Path First (OSPF) instance ID. Redistributes routes from the specified OSPF routing instance into the BGP routing domain. The range of values is 1 to 65,535.
internal	Optional. Redistributes OSPF internal routes into the BGP routing domain.
external	Optional. Redistributes OSPF external routes into the BGP routing domain.
nssa-external	Optional. Redistributes OSPF not-so-stubby-area (NSSA) routes into the BGP routing domain.
rip *instance*	Routing Information Protocol (RIP) instance name. Redistributes routes from the specified RIP routing instance into the BGP routing domain.
static	When entered without the **subscriber** keyword, redistributes static routes into the BGP routing domain. When entered with the optional **subscriber** keyword, redistributes static subscriber routes into the BGP routing domain.
dvsr	Redistributes the dynamically verified static routing (DVSR) subtype of static routes into the BGP routing domain.
subscriber	Redistributes subscriber addresses, ND, DHCP-PD, and subscriber address routes into the BGP routing domain.

address	Used with the **subscriber** keyword. Redistributes subscriber address routes into the BGP routing domain.
route-map *map-name*	Optional. Route map name. Applies a previously configured route map. If this option is not specified, all routes from the specified protocol are redistributed with their default attributes into the BGP routing domain.

1.77.4 Default

Routes learned by other protocols are not redistributed into the BGP routing domain.

1.77.5 Usage Guidelines

Use the **redistribute** command to redistribute IPv4 routes learned through other routing protocols into the BGP routing domain. Redistributed routes are advertised to all BGP neighbors for the address family.

Note: The default route, 0.0.0.0, is not redistributed. Use the **network** command in BGP address family configuration mode to advertise the default route.

You must enter multiple **redistribute** commands to redistribute routes from several different kinds of routing protocols into the BGP routing domain.

Use the **no** form of this command to disable the specified type of route redistribution.

1.77.6 Examples

The following example redistributes external OSPF routes from OSPF instance `100` into the BGP routing domain as unicast routes. The static route `192.200.201.0/24` is redistributed into the BGP routing domain as unicast routes with the community attribute of `100:100`:

```
[local]Redback(config-ctx)#route-map static-to-bgp
[local]Redback(config-route-map)#ip address prefix-list static-to-bgp-prefix
[local]Redback(config-route-map)#set community 100:100
[local]Redback(config-route-map)#exit
[local]Redback(config-ctx)#ip prefix-list static-to-bgp-prefix
[local]Redback(config-prefix-list)#permit 192.200.201.0/24

.

.

.

[local]Redback(config-ctx)#router bgp 100
[local]Redback(config-bgp)#address-family ipv4 unicast
[local]Redback(config-bgp-af)#redistribute ospf 100 external
[local]Redback(config-bgp-af)#redistribute static route-map static-to-bgp
```

The following example redistributes IPsec routes through BGP for non VPN context:

```
[local]Redback(config-isis-if-af)#router bgp 150
[local]Redback(config-bgp)#address-family ipv4 unicast
[local]Redback(config-bgp-af)#redistribute ipsec
[local]Redback(config-bgp-af)#neighbor 15.0.0.1 internal
[local]Redback(config-bgp-neighbor)#update-source to_router1
[local]Redback(config-bgp-neighbor)#address-family ipv4 unicast
```

The following example redistributes IPsec routes through BGP for VPN context:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#router bgp 1
[local]Redback(config)#context vpn1 vpn-rd 1:1
[local]Redback(config-ctx)#router bgp  vpn
[local]Redback(config-bgp)#address-family ipv4 unicast
[local]Redback(config-ctx)#redistribute ipsec
```

1.78 redistribute (BGP, IPv6)

redistribute {**connected** | **isis** *instance* [**level-1** | **level-2**] | **nat** | **ospf3** *instance* [**internal** | [**external**] [**nssa-external**] | **ripng** *instance* | **static** [**dvsr**] | **subscriber** {**address** | **dhcp-pd** | **nd** | **static**}} [**route-map** *map-name*]

no redistribute {**connected** | **isis** *instance* [**level-1** | **level-2**] | **nat** | **ospf3** *instance* [**internal** | [**external**] [**nssa-external**] | **ripng** *instance* | **static** [**dvsr**] | **subscriber** {**address** | **dhcp-pd** | **nd** | **static**}} [**route-map** *map-name*]

1.78.1 Purpose

Redistributes IPv6 routes learned through other routing protocols into the Border Gateway Protocol (BGP) routing domain.

1.78.2 Command Mode

BGP address family configuration

1.78.3 Syntax Description

connected	Redistributes routes from directly attached networks into the BGP routing domain.
isis *instance*	Intermediate System-to-Intermediate System (IS-IS) instance name. Redistributes routes from the specified IS-IS routing instance into the BGP routing domain.
level-1	Optional. Specifies IS-IS level 1 routing.

level-2	Optional. Specifies IS-IS level 2 routing.
nat	Redistributes network address translation (NAT) routes into the BGP routing domain.
ospf3 ***instance***	Open Shortest Path First version 3 (OSPFv3) instance ID. Redistributes routes from the specified OSPFv3 routing instance into the BGP routing domain. The range of values is 1 to 65,535.
internal	Optional. Redistributes OSPFv3 internal routes into the BGP routing domain.
external	Optional. Redistributes OSPFv3 external routes into the BGP routing domain.
nssa-external	Optional. Redistributes OSPFv3 not-so-stubby-area (NSSA) routes into the BGP routing domain.
ripng ***instance***	Routing Information Protocol (RIP) instance name. Redistributes routes from the specified RIP routing instance into the BGP routing domain.
static	When entered without the **subscriber** keyword, redistributes static routes into the BGP routing domain. When entered with the optional **subscriber** keyword, redistributes static subscriber routes into the BGP routing domain.
dvsr	Redistributes the dynamically verified static routing (DVSR) subtype of static routes into the BGP routing domain.
subscriber	Redistributes subscriber addresses, ND, DHCP-PD, and subscriber address routes into the BGP routing domain.
address	Used with the **subscriber** keyword. Redistributes subscriber address routes into the BGP routing domain.
dhcp-pd	Used with the **subscriber** keyword. Redistributes subscriber routes corresponding to DHCPv6 delegated IPv6 prefixes.
nd	Used with the **subscriber** keyword. Redistributes subscriber routes corresponding to ND-assigned IPv6 prefixes.
route-map ***map-name***	Optional. Route map name. Applies a previously configured route map. If this option is not specified, all routes from the specified protocol are redistributed with their default attributes into the BGP routing domain.

1.78.4 Default

Routes learned by other protocols are not redistributed into the BGP routing domain.

1.78.5 Usage Guidelines

Use the **redistribute** command to redistribute IPv6 routes learned through other routing protocols into the BGP routing domain. Redistributed routes are advertised to all BGP neighbors for the address family.'

You must enter multiple **redistribute** commands to redistribute routes from several different kinds of routing protocols into the BGP routing domain.

Use the **no** form of this command to disable the specified type of route redistribution.

1.78.6 Examples

The following example redistributes external OSPFv3 routes from OSPFv3 instance `100` into the BGP routing domain as unicast routes. The static route `2001:101:101:303::/64` is redistributed into the BGP routing domain as unicast routes with the community attribute of `100:100`:

```
[local]Redback(config-ctx)#route-map static-to-bgp
[local]Redback(config-route-map)#ip address prefix-list static-to-bgp-prefix
[local]Redback(config-route-map)#set community 100:100
[local]Redback(config-route-map)#exit
[local]Redback(config-ctx)#ip prefix-list static-to-bgp-prefix
[local]Redback(config-prefix-list)#permit 2001:101:101:303::/64
.
.
.
[local]Redback(config-ctx)#router bgp 100
[local]Redback(config-bgp)#address-family ipv6 unicast
[local]Redback(config-bgp-af)#redistribute ospfv3 100 external
[local]Redback(config-bgp-af)#redistribute static route-map static-to-bgp
```

1.79 redistribute (IS-IS, IPv4)

redistribute {**bgp** *asn* | **connected** | ipsec | **isis** *instance-name* | **nat** | **ospf** *instance-id* [**match** {**external-type-1** | **external-type-2** | **inter-area** | **intra-area** | **nssa-external-type-1**| **nssa-external-type-2**}] | **rip** *instance-name* | **static** [**dvsr**] | **subscriber** [**address** | **static**]} [**level-1** | **level-2**] [**metric** *metric*] [**metric-type** {**internal** | **external**}] [**route-map** *map-name*]

no redistribute {**bgp** *asn* | **connected** | ipsec | **isis** *instance-name*| **nat** | **ospf** *instance-id* [**match** {**external-type-1** | **external-type-2**

| **inter-area** | **intra-area** | **nssa-external-type-1**| **nssa-external-type-2**}] | **rip** ***instance-name*** | **static** [**dvsr**] | **subscriber** [**address** | **static**]} [**level-1** | **level-2**] [**metric** ***metric***] [**metric-type** {**internal** | **external**}] [**route-map** ***map-name***]

1.79.1 Purpose

Redistributes IPv4 routes learned through external routing protocols into the Intermediate System-to-Intermediate System (IS-IS) routing instance.

1.79.2 Command Mode

IS-IS address family configuration

1.79.3 Syntax Description

bgp ***asn***	Border Gateway Protocol (BGP) autonomous system number (ASN). Redistributes routes from BGP into the IS-IS routing instance. The range of values for the ***asn*** argument is 1 to 65,535.
ipsec	Redistributes IPsec routes into the IS-IS routing instance.
connected	Redistributes IPsec routes into the IS-IS routing instance.
isis ***instance-name***	IS-IS instance name. Redistributes routes from the specified IS-IS routing instance into the current IS-IS routing instance.
nat	Redistributes network address translation (NAT) routes into the IS-IS routing instance.
ospf ***instance-id***	Open Shortest Path First (OSPF) instance ID. Redistributes routes from the specified OSPF routing instance into the IS-IS routing instance. The range of values is 1 to 65,535.

match	Redistributes only those OSPF routes matching the specified route type into the IS-IS routing instance. Use one of the following keywords to specify the type of OSPF route to redistribute: • **external-type-1**—Redistributes OSPF external type-1 routes only. • **external-type-2**—Redistributes OSPF external type-2 routes only. • **inter-area**—Redistributes OSPF inter-area routes only. • **intra-area**—Redistributes OSPF intra-area routes only. • **nssa-external-type-1**—Redistributes Not-So-Stubby Area (NSSA) external type-1 routes only. • **nssa-external-type-2**—Redistributes NSSA external type-2 routes only.
external-type-1	Redistributes only type-1 external OSPF routes into the IS-IS routing instance.
external-type-2	Redistributes only type-2 external OSPF routes into the IS-IS routing instance.
inter-area	Redistributes only inter-area OSPF routes into the IS-IS routing instance.
intra-area	Redistributes only intra-area OSPF routes into the IS-IS routing instance.
nssa-external-type-1	Redistributes only NSSA type-1 OSPF routes into the IS-IS routing instance.
nssa-external-type-2	Redistributes only NSSA type-2 OSPF routes into the IS-IS routing instance.
rip *instance-name*	Routing Information Protocol (RIP) instance name. Redistributes routes from the specified RIP routing instance into the IS-IS routing instance.
static	When entered without the **subscriber** keyword, redistributes static routes into the IS-IS routing instance. When entered with the optional **subscriber** keyword, redistributes static subscriber routes into the IS-IS routing instance.
dvsr	Optional. Redistributes dynamically verified static routing (DVSR) subtype of static routes into the IS-IS routing instance.
subscriber	Redistributes subscriber addresses, ND, DHCP-PD, and subscriber address routes into the IS-IS routing instance.
address	Used with the **subscriber** keyword. Redistributes subscriber address routes into the IS-IS routing instance

level-1	Optional. Redistributes only level 1 routes into the IS-IS routing instance.
level-2	Optional. Redistributes only level 2 routes into the IS-IS routing instance independently.
metric ***metric***	Optional. Metric assigned to the redistributed routes. The range of values is 0 to 16,777,215; the default metric is 0. When no metric is configured by the **redistribute command**, the metric specified in the **route-map** command option is used as the internal prefix in the IS-IS domain. In addition, if no metric is specified in the **route-map** command, the original route metric is used as the internal prefix in the IS-IS domain.
metric-type	Optional. Assigns a metric type to the redistributed routes; the default metric type is internal.
internal	Assigns an internal metric type to redistributed routes. When the system receives an LSP with an internal metric type, the total cost is the cost the route from itself to the redistributing system plus the advertised cost to reach the destination.
external	Assigns an external metric type to redistributed routes. When the system receives a link-state protocol data unit (LSP) with an external metric type, it considers only the advertised cost to reach the destination
route-map ***map-name***	Optional. Route map name. Applies a previously configured route map that filters the routes that are redistributed into the IS-IS routing instance. If this option is not specified, all routes from the specified protocol are redistributed into the IS-IS routing instance.

1.79.4 Default

Routes learned by other protocols are not redistributed into the IS-IS routing instance.

1.79.5 Usage Guidelines

Use the **redistribute** command to redistribute IPv4 routes learned through external protocols into the IS-IS routing instance.

You must enter multiple **redistribute** commands to redistribute routes from several different kinds of routing protocols into the IS-IS routing instance.

Use the **no** form of this command to disable redistribution into the IS-IS routing instance.

1.79.6 Examples

The following example shows how to redistribute static IP routes into an IS-IS level-1 area with an advertised metric of `10`. The internal metric type is used by default:

```
[local]Redback(config-ctx)#router isis ip-backbone
[local]Redback(config-isis)#address-family ipv4 unicast
[local]Redback(config-isis-af)#redistribute static level-1 metric 10
```

The following example shows how to redistribute IPsec routes through IS-IS. Route map and distance are explicitly required to redistribute IPsec. The metric value (`+5` in this example) is added to IPsec tunnel cost. IPsec tunnel cost is configured in the configuration, `ip route traffic-selector-guided`.

```
[local]Redback(config-ctx)#ip prefix-list p1
[local]Redback(config-prefix-list)#seq 10 permit any
[local]Redback(config-prefix-list)#route-map r1 permit 1
[local]Redback(config-route-map)#match ip address prefix-list p1
[local]Redback(config-route-map)#set metric +5
[local]Redback(config-route-map)#router isis 2
[local]Redback(config-isis)#net 47.0001.1111.2222.3333.00
[local]Redback(config-isis)#address-family ipv4 unicast
[local]Redback(config-isis-af)#redistribute ipsec route-map r1
[local]Redback(config-isis-af)#interface traffic_interface
[local]Redback(config-isis-if)#address-family ipv4 unicast
```

1.80 redistribute (IS-IS, IPv6)

redistribute {**bgp** *asn* | **connected** | **isis** *instance-name* | **nat** | **ospf3** *instance-id* [**match** {**external-type-1** | **external-type-2** | **inter-area** | **intra-area** | **nssa-external-type-1**| **nssa-external-type-2**}] | **ripng** *instance-name* | **static** [**dvsr**] | **subscriber** {**address** | **dhcp-pd** | **nd** | **static**}} [**level-1** | **level-2**] [**metric** *metric*]
[**metric-type** {**internal** | **external**}] [**route-map** *map-name*]

no redistribute {**bgp** *asn* | **connected** | **isis** *instance-name* | **nat** | **ospf3** *instance-id* [**match** {**external-type-1** | **external-type-2** | **inter-area** | **intra-area** | **nssa-external-type-1**| **nssa-external-type-2**}] | **ripng** *instance-name* | **static** [**dvsr**] | **subscriber** {**address** | **dhcp-pd** | **nd** | **static**}} [**level-1** | **level-2**] [**metric** *metric*] [**metric-type** {**internal** | **external**}] [**route-map** *map-name*]

1.80.1 Purpose

Redistributes IPv6 routes learned through external routing protocols into the Intermediate System-to-Intermediate System (IS-IS) routing instance.

1.80.2 Command Mode

IS-IS address family configuration

1.80.3 Syntax Description

bgp ***asn***	Border Gateway Protocol (BGP) autonomous system number (ASN). Redistributes routes from BGP into the IS-IS routing instance. The range of values for the ***asn*** argument is 1 to 65,535.
connected	Redistributes routes from directly attached networks into the IS-IS routing instance.
isis ***instance-name***	IS-IS instance name. Redistributes routes from the specified IS-IS routing instance into the current IS-IS routing instance.
nat	Redistributes network address translation (NAT) routes into the IS-IS routing instance.
ospf3 ***instance-id***	OSPF Version 3 (OSPFv3) instance ID. Redistributes routes from the specified OSPFv3 routing instance into the IS-IS routing instance. The range of values is 1 to 65535.
match	Redistributes only those OSPFv3 routes matching the specified route type into the IS-IS routing instance. Use one of the following keywords to specify the type of OSPFv3 route to redistribute: • **external-type-1**—Redistributes OSPFv3 external type-1 routes only. • **external-type-2**—Redistributes OSPFv3 external type-2 routes only. • **inter-area**—Redistributes OSPFv3 interarea routes only. • **intra-area**—Redistributes OSPF intra-area routes only. • **nssa-external-type-1**—Redistributes Not-So-Stubby Area (NSSA) external type-1 routes only. • **nssa-external-type-2**—Redistributes NSSA external type-2 routes only.
external-type-1	Redistributes only type-1 external OSPFv3 routes into the IS-IS routing instance.
external-type-2	Redistributes only type-2 external OSPFv3 routes into the IS-IS routing instance.
inter-area	Redistributes only inter-area OSPFv3 routes into the IS-IS routing instance.
intra-area	Redistributes only intra-area OSPFv3 routes into the IS-IS routing instance.

nssa-external-type-1	Redistributes only NSSA type-1 OSPFv3 routes into the IS-IS routing instance.
nssa-external-type-2	Redistributes only NSSA type-2 OSPFv3 routes into the IS-IS routing instance.
rip *instance-name*	Routing Information Protocol (RIP) instance name. Redistributes routes from the specified RIP routing instance into the IS-IS routing instance.
static	When entered without the **subscriber** keyword, redistributes static routes into the IS-IS routing instance. When entered with the optional **subscriber** keyword, redistributes static subscriber routes into the IS-IS routing instance.
dvsr	Optional. Redistributes dynamically verified static routing (DVSR) subtype of static routes into the IS-IS routing instance.
subscriber	Redistributes subscriber addresses, ND, DHCP-PD, and subscriber address routes into the IS-IS routing instance.
dhcp-pd	Used with the **subscriber** keyword. Redistributes subscriber routes corresponding to DHCPv6 delegated IPv6 prefixes.
nd	Used with the **subscriber** keyword. Redistributes subscriber routes corresponding to ND-assigned IPv6 prefixes.
address	Used with the **subscriber** keyword. Redistributes subscriber address routes into the IS-IS routing instance
level-1	Optional. Redistributes only level 1 routes into the IS-IS routing instance.
level-2	Optional. Redistributes only level 2 routes into the IS-IS routing instance independently.
metric *metric*	Optional. Metric assigned to the redistributed routes. The range of values is 0 to 16,777,215; the default metric is 0. When no metric is configured by the **redistribute command**, the metric specified in the **route-map** command option is used as the internal prefix in the IS-IS domain. In addition, if no metric is specified in the **route-map** command, the original route metric is used as the internal prefix in the IS-IS domain.
metric-type	Optional. Assigns a metric type to the redistributed routes; the default metric type is internal.
internal	Assigns an internal metric type to redistributed routes. When the system receives an LSP with an internal metric type, the total cost is the cost the route from itself to the redistributing system plus the advertised cost to reach the destination.

external	Assigns an external metric type to redistributed routes. When the system receives a link-state protocol data unit (LSP) with an external metric type, it considers only the advertised cost to reach the destination
route-map *map-name*	Optional. Route map name. Applies a previously configured route map that filters the routes that are redistributed into the IS-IS routing instance. If this option is not specified, all routes from the specified protocol are redistributed into the IS-IS routing instance.

1.80.4 Default

Routes learned by other protocols are not redistributed into the IS-IS routing instance.

1.80.5 Usage Guidelines

Use the **redistribute** command to redistribute IPv6 routes learned through external protocols into the IS-IS routing instance.

You must enter multiple **redistribute** commands to redistribute routes from several different kinds of routing protocols into the IS-IS routing instance.

Use the **no** form of this command to disable redistribution into the IS-IS routing instance.

1.80.6 Examples

The following example shows how to redistribute static IPv6 routes into an IS-IS level-1 area with an advertised metric of `10`. The internal metric type is used by default:

```
[local]Redback(config-ctx)#router isis ip-backbone
[local]Redback(config-isis)#address-family ipv6 unicast
[local]Redback(config-isis-af)#redistribute static level-1 metric 10
```

1.81 redistribute (OSPF)

redistribute {**bgp** *asn* | **connected** | ipsec | **isis** *instance* [**level-1** | **level-2**] | **nat** | **ospf** *instance* [**external** [**type-1** | **type-2**]] [**inter-area**] [**intra-area**] [**nssa** [**type-1** | **type-2**]] | **rip** *instance* | **static** [**dvsr**] | **subscriber** [**address** | **static**]} [**metric** *metric*] [**metric-type** *type*] [**route-map** *map-name*] [**tag** *tag*]

no redistribute {**bgp** *asn* | **connected** | ipsec | **isis** *instance* [**level-1** | **level-2**] | **nat** | **ospf** *instance* [**external** [**type-1** | **type-2**]] [**inter-area**] [**intra-area**] [**nssa** [**type-1** | **type-2**]] | **rip** *instance*

| **static** [**dvsr**] | **subscriber** [**address** | **static**]} [**metric** ***metric***] [**metric-type** ***type***] [**route-map** ***map-name***] [**tag** ***tag***]

1.81.1 Purpose

Redistribute routes learned from other protocols into the Open Shortest Path First (OSPF) routing instance.

1.81.2 Command Mode

OSPF router configuration

1.81.3 Syntax Description

bgp ***asn***	Border Gateway Protocol (BGP) autonomous system number (ASN). Redistributes routes from the specified BGP autonomous system (AS) into the OSPF routing instance. The range of values for the ***asn*** argument is 1 to 65,535.
connected	Redistributes routes from directly attached networks into the OSPF routing instance.
ipsec	Redistributes IPsec routes into the OSPF routing instance.
isis ***instance***	Intermediate System-to-Intermediate System (IS-IS) instance name. Redistribute routes from the specified IS-IS routing instance into the OSPF routing instance.
level-1	Optional. Redistributes IS-IS level 1 routes only.
level-2	Optional. Redistributes IS-IS level 2 routes only.
nat	Redistributes network address translation (NAT) routes into the OSPF routing instance.
ospf ***instance***	OSPF instance ID. Redistributes routes from another OSPF routing instance into the current OSPF routing instance. The range of values for the ***instance*** argument is 1 to 65,535.
external	Optional. Redistributes only external OSPF routes.
type-1	Optional. Redistributes only Type 1 external OSPF routes.
type-2	Optional. Redistributes only Type 2 external OSPF routes.
inter-area	Optional. Redistributes only interarea OSPF routes.
intra-area	Optional. Redistributes only intraarea OSPF routes.
nssa	Optional. Redistributes only OSPF NSSA routes.
type-1	Optional. Redistributes only OSPF NSSA Type 1 routes.
type-2	Optional. Redistributes only OSPF NSSA Type 2 routes.

rip *instance*	Routing Information Protocol (RIP) instance name. Redistributes routes from the specified RIP routing instance into the current OSPF routing instance.
static	When entered without the **subscriber** keyword, redistributes static routes into the OSPF routing instance. When entered with the optional **subscriber** keyword, redistributes static subscriber routes into the OSPF routing instance.
dvsr	Optional. Redistributes the dynamically verified static routing (DVSR) subtype of static routes into the OSPF routing instance.
subscriber	Redistributes subscriber addresses and subscriber address routes into the OSPF routing instance.
address	Used with the **subscriber** keyword. Redistributes subscriber address routes into the OSPF routing instance.
metric *metric*	Optional. Cost of the redistributed routes. The range of values is 0 to 16,777,215; the default value is 20.
metric-type *type*	Optional. Metric type assigned to the redistributed routes. The *type* argument specifies one of the following metric types: • **1**—Specifies a Type 1 metric type. • **2**—Specifies a Type 2 metric type.
route-map *map-name*	Optional. Route map name. Modifies the attributes of redistributed routes using the specified route map.
tag *tag*	Optional. Route tag used to redistribute routes. An unsigned 32-bit integer, the range of values is 1 to 4,294,967,295; the default value is 0.

1.81.4 Default

Routes learned by other protocols are not redistributed into the OSPF routing instance.

1.81.5 Usage Guidelines

Use the **redistribute** command to redistribute routes learned from other protocols into the OSPF routing instance.

Note: IPv6 routes cannot be redistributed into an OSPF routing instance.

You must enter multiple **redistribute** commands to redistribute routes from several different kinds of routing protocols into the OSPF routing instance.

Use the **no** form of this command to disable redistribution of the specified routing protocol or method.

1.81.6 Examples

The following example shows how to redistribute RIP into an OSPF routing instance:

```
[local]Redback(config-ospf)#redistribute rip
```

The following example shows how to redistribute IPsec routes through an OSPF routing instance:

```
[local]Redback(config-ctx)#router ospf 1
[local]Redback(config-ospf)#fast-convergence
[local]Redback(config-ospf)#area 0.0.0.0
[local]Redback(config-ospf-area)#interface traffic_interface
[local]Redback(config-ospf-if)#redistribute ipsec
[local]Redback(config-ospf)#redistribute static
```

1.82 redistribute (OSPFv3)

redistribute {**bgp** *asn* | **connected** | **isis** *instance* [**level-1** | **level-2**] | **nat** | **ospf3** *instance* [**external** [**type-1** | **type-2**]] [**inter-area**] [**intra-area**] [**nssa** [**type-1** | **type-2**]] | **ripng** *instance* | **static** [**dvsr**] | **subscriber** {**address** | **dhcp-pd** | **nd** | **static**}} [**metric** *metric*] [**metric-type** *type*] [**route-map** *map-name*] [**tag** *tag*]

no redistribute {**bgp** *asn* | **connected** | **isis** *instance* [**level-1** | **level-2**] | **nat** | **ospf3** *instance* [**external** [**type-1** | **type-2**]] [**inter-area**] [**intra-area**] [**nssa** [**type-1** | **type-2**]] | **ripng** *instance* | **static** [**dvsr**] | **subscriber** {**address** | **dhcp-pd** | **nd** | **static**}} [**metric** *metric*] [**metric-type** *type*] [**route-map** *map-name*] [**tag** *tag*]

1.82.1 Purpose

Redistribute routes learned from other protocols into the Open Shortest Path First version 3 (OSPFv3) routing instance.

1.82.2 Command Mode

OSPF3 router configuration

1.82.3 Syntax Description

bgp *asn*	Border Gateway Protocol (BGP) autonomous system number (ASN). Redistributes routes from the specified BGP autonomous system (AS) into the OSPFv3 routing instance. The range of values for the ***asn*** argument is 1 to 65,535.
connected	Redistributes routes from directly attached networks into the OSPFv3 routing instance.
isis *instance*	Intermediate System-to-Intermediate System (IS-IS) instance name. Redistribute routes from the specified IS-IS routing instance into the OSPFv3 routing instance.
level-1	Optional. Redistributes IS-IS level 1 routes only.
level-2	Optional. Redistributes IS-IS level 2 routes only.
nat	Redistributes network address translation (NAT) routes into the OSPFv3 routing instance.
ospf3 *instance*	OSPF instance ID. Redistributes routes from another OSPFv3 routing instance into the current OSPFv3 routing instance. The range of values for the ***instance*** argument is 1 to 65,535.
external	Optional. Redistributes only external OSPFv3 routes.
type-1	Optional. Redistributes only Type 1 external OSPFv3 routes.
type-2	Optional. Redistributes only Type 2 external OSPFv3 routes.
inter-area	Optional. Redistributes only interarea OSPFv3 routes.
intra-area	Optional. Redistributes only intraarea OSPFv3 routes.
nssa	Optional. Redistributes only OSPFv3 NSSA routes.
type-1	Optional. Redistributes only OSPFv3 NSSA Type 1 routes.
type-2	Optional. Redistributes only OSPFv3 NSSA Type 2 routes.
ripng *instance*	Routing Information Protocol (RIP) instance name. Redistributes routes from the specified RIP routing instance into the current OSPFv3 routing instance.
static	When entered without the **subscriber** keyword, redistributes static routes into the OSPFv3 routing instance. When entered with the optional **subscriber** keyword, redistributes static subscriber routes into the OSPFv3 routing instance.
dvsr	Optional. Redistributes the dynamically verified static routing (DVSR) subtype of static routes into the OSPFv3 routing instance.
subscriber	Redistributes subscriber addresses, ND, DHCP-PD, and subscriber address routes into the OSPFv3 routing instance.

address	Used with the **subscriber** keyword. Redistributes subscriber address routes into the OSPFv3 routing instance.
dhcp-pd	Used with the **subscriber** keyword. Redistributes subscriber routes corresponding to DHCPv6 delegated IPv6 prefixes. The **dhcp-pd** keyword is available for the OSPFv3 routing instance only.
nd	Used with the **subscriber** keyword. Redistributes subscriber routes corresponding to ND-assigned IPv6 prefixes. The **nd** keyword is available for the OSPFv3 routing instance only.
metric *metric*	Optional. Cost of the redistributed routes. The range of values is 0 to 16,777,215; the default value is 20.
metric-type *type*	Optional. Metric type assigned to the redistributed routes. The *type* argument specifies one of the following metric types: • **1**—Specifies a Type 1 metric type. • **2**—Specifies a Type 2 metric type.
route-map *map-name*	Optional. Route map name. Modifies the attributes of redistributed routes using the specified route map.
tag *tag*	Optional. Route tag used to redistribute routes. An unsigned 32-bit integer, the range of values is 1 to 4,294,967,295; the default value is 0.

1.82.4 Default

Routes learned by other protocols are not redistributed into the OSPFv3 routing instance.

1.82.5 Usage Guidelines

Use the **redistribute** command to redistribute routes learned from other protocols into the OSPFv3 routing instance.

Note: You can redistribute IPv6 unicast routes from RIPng into an OSPFv3 routing instance. IPv4 routes cannot be redistributed into an OSPFv3 routing instance.

You must enter multiple **redistribute** commands to redistribute routes from several different kinds of routing protocols into the OSPFv3 routing instance.

Use the **no** form of this command to disable redistribution of the specified routing protocol or method.

1.82.6 Examples

The following example shows how to redistribute RIP into the OSPFv3 routing instance:

```
[local]Redback(config-ospf3)#redistribute rip
```

1.83 redistribute (RIP)

redistribute {**bgp** *asn* | **connected** | ipsec |**isis** *instance* [**level-1** | **level-2** | **level-1-2**] |**nat** | **ospf** *instance* [**match** {**external-type-1** | **external-type-2** | **inter-area** | **intra-area** | **nssa-external-type-1**| **nssa-external-type-2**}] | **rip** *instance* | **static** [**dvsr**] | **subscriber** [**address** | **static**]} [**metric** *metric*] [**route-map** *map-name*]

no redistribute {**bgp** *asn* | **connected** | ipsec | **isis** *instance* | **nat** | **ospf** *instance* [**match** {**external-type-1** | **external-type-2** | **inter-area** | **intra-area** | **nssa-external-type-1**| **nssa-external-type-2**}] | **rip** *instance* | **static** [**dvsr**] | **subscriber** [**address** | **static**]} [**metric** *metric*] [**route-map** *map-name*]

1.83.1 Purpose

Redistributes routes learned from other routing protocols into the Routing Information Protocol (RIP) routing instance.

1.83.2 Command Mode

RIP router configuration

1.83.3 Syntax Description

bgp *asn*	Border Gateway Protocol (BGP) autonomous system number (ASN). Redistributes routes from the specified BGP autonomous system (AS) into the RIP routing instance. The range of values for the *asn* argument is 1 to 65,535.
connected	Redistributes directly attached networks into the RIP routing instance.
ipsec	Redistributes IPsec routes into the RIP routing instance.

isis *instance*	Intermediate System-to-Intermediate System (IS-IS) instance name. Redistributes routes from the specified IS-IS instance into the RIP routing instance.
level-1	Optional. Redistributes IS-IS level 1 routes only.
level-2	Optional. Redistributes IS-IS level 2 routes only.
level-1-2	Optional. Redistributes IS-IS level 1 and level 2 routes.
nat	Redistributes network address translation (NAT) routes into the RIP routing instance.
ospf *instance*	Open Shortest Path First (OSPF) instance ID. Redistributes routes from the specified OSPF routing instance into the RIP routing instance. The range of values is 1 to 65,535.
match	Redistributes only those OSPF routes matching the specified route type into the RIP routing instance. Use one of the following keywords to specify the type of OSPF route to redistribute: • **external-type-1**—Redistributes OSPF external type-1 routes only. • **external-type-2**—Redistributes OSPF external type-2 routes only. • **inter-area**—Redistributes OSPF inter-area routes only. • **intra-area**—Redistributes OSPF intra-area routes only. • **nssa-external-type-1**—Redistributes Not-So-Stubby Area (NSSA) external type-1 routes only. • **nssa-external-type-2**—Redistributes NSSA external type-2 routes only.
external-type-1	Redistributes only type-1 external OSPF routes into the RIP routing instance.
external-type-2	Redistributes only type-2 external OSPF routes into the RIP routing instance.
inter-area	Redistributes only inter-area OSPF routes into the RIP routing instance.
intra-area	Redistributes only intra-area OSPF routes into the RIP routing instance.
nssa-external-type-1	Redistributes only NSSA type-1 OSPF routes into the RIP routing instance.
nssa-external-type-2	Redistributes only NSSA type-2 OSPF routes into the RIP routing instance.
rip *instance*	RIP instance name. Redistributes routes from another RIP routing instance into the current RIP routing instance.

static	When entered without the **subscriber** keyword, redistributes static routes into the RIP routing instance. When entered with the optional **subscriber** keyword, redistributes static subscriber routes into the RIP routing instance.
dvsr	Optional. Redistributes the dynamically verified static routing (DVSR) subtype of static routes into the RIP routing instance.
subscriber	Redistributes subscriber addresses and subscriber address routes into the RIP routing instance.
address	Used with the **subscriber** keyword. Redistributes subscriber address routes into the RIP routing instance.
metric ***metric***	Optional. Metric used for the redistributed route. The range of values is 0 to 16. If no metric is specified, the metric configured with the **default-metric** command is used in RIP router configuration mode. If the **default-metric** command has not been configured, the default metric for redistributed routes is 0.
route-map ***map-name***	Optional. Route map name. Applies the conditions of the specified route map to routes that are redistributed into the RIP routing instance.

1.83.4 Default

Routes learned by other protocols are not redistributed into the RIP routing instance.

1.83.5 Usage Guidelines

Use the **redistribute** command to redistribute routes learned from other routing protocols into the RIP routing instance.

You must enter multiple **redistribute** commands to redistribute routes from several different kinds of routing protocols into the RIP routing instance.

Use the **no** form of this command to disable the specified type of route redistribution.

1.83.6 Examples

The following example shows how to redistribute static routes into RIP routing instance, `rip001`:

```
[local]Redback(config-ctx)#router rip rip001[local]Redback(config-rip)#redistribute static
```

The following example shows how to redistribute IPsec routes through a RIP routing instance. Route map and distance are explicitly required to redistribute IPsec. The metric value (+2 in this example) is added to IPsec tunnel cost, which is configured in configuration, ip route traffic-selector-guided.

```
[peer-responder]Redback(config-ctx)#router rip r1
[peer-responder]Redback(config-rip)#default-metric 1
[peer-responder]Redback(config-rip)#interface traffic_interface
[peer-responder]Redback(config-rip-if)#redistribute ipsec route-map r1
[peer-responder]Redback(config-rip)#ip prefix-list p1
[peer-responder]Redback(config-prefix-list)#seq 10 permit any
[peer-responder]Redback(config-prefix-list)#route-map r1 permit 1
[peer-responder]Redback(config-route-map)#match ip address prefix-list p1
[peer-responder]Redback(config-route-map)#set metric +2
```

The following example shows how to prevent routes from directly attached networks from being redistributed into RIP routing instance, rip001:

```
[local]Redback(config-ctx)#router rip rip001[local]Redback(config-rip)#no redistribute connected
```

1.84 redistribute (RIPng)

redistribute {**bgp** *asn* | **connected** | **isis** *instance* [**level-1** | **level-2** | **level-1-2**] |**nat** | **ospf3** *instance* [**match** {**external-type-1** | **external-type-2** | **inter-area** | **intra-area** | **nssa-external-type-1**| **nssa-external-type-2**}] | **ripng** *instance* | **static** [**dvsr**] | **subscriber** {**address** | **dhcp-pd** | **nd** | **static**}} [**metric** *metric*] [**route-map** *map-name*]

no redistribute {**bgp** *asn* | **connected** | **isis** *instance* | **nat** | **ospf3** *instance* [**match** {**external-type-1** | **external-type-2** | **inter-area** | **intra-area** | **nssa-external-type-1**| **nssa-external-type-2**}] | **ripng** *instance* | **static** [**dvsr**] | **subscriber** {**address** | **dhcp-pd** | **nd** | **static**}} [**metric** *metric*] [**route-map** *map-name*]

1.84.1 Purpose

Redistributes routes learned from other routing protocols into the Routing Information Protocol next generation (RIPng) routing instance.

1.84.2 Command Mode

RIPng router configuration

1.84.3 Syntax Description

bgp ***asn***	Border Gateway Protocol (BGP) autonomous system number (ASN). Redistributes routes from the specified BGP autonomous system (AS) into the RIP routing instance. The range of values for the ***asn*** argument is 1 to 65,535.
connected	Redistributes directly attached networks into the RIPng routing instance.
isis ***instance***	Intermediate System-to-Intermediate System (IS-IS) instance name. Redistributes routes from the specified IS-IS instance into the RIPng routing instance.
level-1	Optional. Redistributes IS-IS level 1 routes only.
level-2	Optional. Redistributes IS-IS level 2 routes only.
level-1-2	Optional. Redistributes IS-IS level 1 and level 2 routes.
nat	Redistributes network address translation (NAT) routes into the RIPng routing instance.
ospf3 ***instance***	Open Shortest Path First version 3 (OSPFv3) instance ID. Redistributes routes from the specified OSPF routing instance into the RIPng routing instance. The range of values is 1 to 65,535.
match	Redistributes only those OSPFv3 routes matching the specified route type into the RIPng routing instance. Use one of the following keywords to specify the type of OSPFv3 route to redistribute: • **external-type-1**—Redistributes OSPFv3 external type-1 routes only. • **external-type-2**—Redistributes OSPFv3 external type-2 routes only. • **inter-area**—Redistributes OSPFv3 inter-area routes only. • **intra-area**—Redistributes OSPFv3 intra-area routes only. • **nssa-external-type-1**—Redistributes Not-So-Stubby Area (NSSA) external type-1 routes only. • **nssa-external-type-2**—Redistributes NSSA external type-2 routes only.

external-type-1	Redistributes only type-1 external OSPFv3 routes into the RIPng routing instance.
external-type-2	Redistributes only type-2 external OSPFv3 routes into the RIPng routing instance.
inter-area	Redistributes only interarea OSPFv3 routes into the RIPng routing instance.
intra-area	Redistributes only intra-area OSPFv3 routes into the RIPng routing instance.
nssa-external-type-1	Redistributes only NSSA type-1 OSPFv3 routes into the RIPng routing instance.
nssa-external-type-2	Redistributes only NSSA type-2 OSPFv3 routes into the RIPng routing instance.
ripng *instance*	RIPng instance name. Redistributes routes from another RIPng routing instance into the current RIPng routing instance.
static	When entered without the **subscriber** keyword, redistributes static routes into the RIPng routing instance. When entered with the optional **subscriber** keyword, redistributes static subscriber routes into the RIPng routing instance.
dvsr	Optional. Redistributes the dynamically verified static routing (DVSR) subtype of static routes into the RIPng routing instance.
subscriber	Redistributes subscriber addresses, ND, DHCP-PD, and subscriber address routes into the RIPng routing instance.
address	Used with the **subscriber** keyword. Redistributes subscriber address routes into the RIPng routing instance.
dhcp-pd	Used with the **subscriber** keyword. Redistributes subscriber routes corresponding to DHCPv6 delegated IPv6 prefixes.
nd	Used with the **subscriber** keyword. Redistributes subscriber routes corresponding to ND-assigned IPv6 prefixes.

metric ***metric***	Optional. Metric used for the redistributed route. The range of values is 0 to 16. If no metric is specified, the metric configured with the **default-metric** command is used in RIPng router configuration mode. If the **default-metric** command has not been configured, the default metric for redistributed routes is 0.
route-map ***map-name***	Optional. Route map name. Applies the conditions of the specified route map to routes that are redistributed into the RIPng routing instance.

1.84.4 Default

Routes learned by other protocols are not redistributed into the RIPng routing instance.

1.84.5 Usage Guidelines

Use the **redistribute** command to redistribute routes learned from other routing protocols into the RIPng routing instance.

You must enter multiple **redistribute** commands to redistribute routes from several different kinds of routing protocols into the RIPng routing instance.

Use the **no** form of this command to disable the specified type of route redistribution.

1.84.6 Examples

The following example shows how to redistribute static routes into RIPng routing instance, `rip001`:

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

The following example prevents routes from directly attached networks from being redistributed into RIPng routing instance, `rip002`:

```
[local]Redback(config-ctx)#router ripng rip002
[local]Redback(config-rip)#no redistribute connected
```

1.85 redistribute bgp (LDP)

[no] redistribute bgp [route-map *map-name*]

1.85.1 Purpose

Redistributes iBGP routes to LDP.

1.85.2 Command Mode

router ldp configuration

1.85.3 Syntax Description

route-map *map-name*	Optional. Use a route-map and specify the name.

1.85.4 Default

iBGP routes are not redistributed outside autonomous systems (ASs).

1.85.5 Usage Guidelines

Use the **redistribute bgp route-map** command to redistribute iBGP routes to LDP (for example, to create an LSP between ASs that are spanned by an L3VPN. Use the optional **route-map *map-name*** construct if there is more than one route to be redistributed.

1.85.6 Examples

The following example redistribute the iBGP route map AS-211 into LDP.

```
[local]Redback(config-ctx)#router ldp
[local]Redback(config-ldp)#redistribute bgp
route-map AS-211
```

1.86 redistribute subscriber aggregate

redistribute subscriber aggregate [**route-map *map-name***]

no redistribute subscriber aggregate

1.86.1 Purpose

Redistributes downloaded routes into BGP as IPv4 unicast addresses. Unless you explicitly configure this command, BGP will not get these routes from the routing information base (RIB).

1.86.2 Command Mode

BGP address family configuration

1.86.3 Syntax Description

route-map ***map-name***	Optional. Route map name. Applies a previously configured route map. If this option is not specified, all routes from the specified protocol are redistributed with their default attributes into the BGP routing domain.

1.86.4 Usage Guidelines

Use this command to redistribute subscriber aggregate routes to BGP. Subscriber aggregate routes are preprovisioned prefixes that BGP announces to its peers before the subscribers come up. The subscriber aggregate prefixes are installed in the routing table with a **null0** next hop so that no packet flows are forwarded until the subscriber is up. For scalability, do not redistribute other forms of subscriber routes to BGP, including "redistribute subscriber," "redistribute subscriber static," and "redistribute subscriber address." The subscriber aggregate prefix should be shorter than the actual subscriber prefix (which typically are /32 prefixes). Otherwise, if only **redistribute subscriber aggregate** is configured, BGP withdraws the prefix when the subscriber comes up.

1.86.5 Examples

The following example shows how to configure the **redistribute subscriber aggregate** command in router bgp mode and address-family submode. These commands make the RIB redistribute any routes of type **SUB P** (subscriber pool) to BGP, in the desired context and under the specified address family:

```
[local]Redback(config-ctx)#router bgp 1
[local]Redback(config-bgp)#address-family ipv4 unicast
[local]Redback(config-bgp-af)#redistribute subscriber aggregate
```

1.87 redundancy-mode

redundancy-mode {master-slave | independent}

no redundancy-mode {master-slave | independent}

1.87.1 Purpose

Enables either the master-slave or independent L2VPN XC redundancy mode for all redundant XC pairs that have the current L2VPN profile attached.

1.87.2 Command Mode

L2VPN profile peer configuration

1.87.3 Syntax Description

master-slave	Enables master-slave redundancy mode on the XCs that have the current L2VPN profile attached.
independent	Enables independent redundancy mode on the XCs that have the current L2VPN profile attached.

1.87.4 Default

Redback mode (redundancy is not Muley-signaled).

1.87.5 Usage Guidelines

Use the **redundancy-mode** command to enable either the master-slave or independent L2VPN XC redundancy mode for all redundant XC pairs that have the current L2VPN profile attached.

In master-slave mode, the hub node serves as the master endpoint that selects which L2VPN XC to use for forwarding. The status of the active L2VPN XC is communicated to the slave node through the signalling mechanism. The slave node inherits its state from the master endpoint. For example, if the master endpoint of the L2VPN XC is active, the slave node is active; if the master endpoint is on standby, the slave endpoint is on standby. With master-slave L2VPN redundancy mode, only XC redundancy is achieved

In independent redundancy mode, the L2VPN XC endpoint nodes independently select which L2VPN XC is used for forwarding. Each node advertises its forwarding state over each L2VPN XC in a set. Each endpoint compares the local and remote status of its L2VPN XC, and activates the L2VPN XC that is active at both endpoints.

Use the **no** form of this command to remove L2VPN redundancy configuration from a specified L2VPN XC profile.

1.87.6 Examples

The following example shows how to enable master-slave redundancy mode for all redundant XC pairs that have the L2VPN profile called `ms-prof` attached:

```
[local]Redback(config)#l2vpn profile ms-prof
[local]Redback(config-l2vpn-xc-profile)#peer 100.100.100.1
[local]Redback(config-l2vpn-xc-profile-peer)#redundancy-mode master-slave
```

The following example shows how to enable independent redundancy mode for all redundant XC pairs that have the L2VPN profile called `in-prof` attached:

```
[local]grumpy(config)#l2vpn profile in-prof
[local]grumpy(config-l2vpn-xc-profile)#peer 100.100.100.1
[local]grumpy(config-l2vpn-xc-profile-peer)#redundancy-mode independent
```

1.88 refresh-interval

`refresh-interval` ***`interval`***

1.88.1 Purpose

Configures the frequency of generating refresh messages.

1.88.2 Command Mode

RSVP interface configuration

1.88.3 Syntax Description

`interval`	Frequency, in seconds, at which refresh messages are generated. The range of values is 1 to 65535.

1.88.4 Default

Refresh messages are generated every 30 seconds.

1.88.5 Usage Guidelines

Use the **`refresh-interval`** command to configure the frequency of generating refresh messages.

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 shows how to set the refresh-interval timing parameter to `45` seconds:

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

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

[local]Redback(config-rsvp-if)#refresh-interval 45
```

1.89 registration max-lifetime

`registration max-lifetime` ***`seconds`***

`no registration max-lifetime`

1.89.1 Purpose

Specifies the maximum lifetime registration for any mobile node (MN) that uses this foreign agent (FA) instance.

1.89.2 Command Mode

Mobile IP interface configuration

1.89.3 Syntax Description

`seconds`	Maximum lifetime registration. The range of values is 1 to 65535 seconds. The default value is 1800 seconds (30 minutes).

1.89.4 Default

The maximum lifetime registration is 1800 seconds (30 minutes).

1.89.5 Usage Guidelines

Use the **registration max-lifetime** command to specify the maximum lifetime registration for any MN that uses this FA instance.

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

1.89.6 Examples

The following example shows how to specify a maximum registration lifetime of 60 minutes (3600 seconds) with the FA instance in this context:

```
[local]Redback(config)#context fa
[local]Redback(config-ctx)#router mobile-ip
[local]Redback(config-mip)#interface mn-access
[local]Redback(config-mip-if)#registration max-lifetime 3600
```

1.90 registration max-lifetime (HA)

registration max-lifetime ***seconds***

no registration max-lifetime

1.90.1 Purpose

Specifies the registration maximum lifetime for any mobile node (MN) that uses this home agent (HA) instance.

1.90.2 Command Mode

HA configuration

1.90.3 Syntax Description

seconds	Registration maximum lifetime. The range of values is 1 to 65535 seconds.

1.90.4 Default

The registration maximum lifetime default is 1800 seconds (30 minutes).

1.90.5 Usage Guidelines

Use the **registration max-lifetime** command to specify the registration maximum lifetime for any MN that uses this HA instance.

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

1.90.6 Examples

The following example shows how to specify a registration maximum lifetime of 60 minutes (3600 seconds) for the HA instance in this context:

```
[local]Redback(config)#context ha
[local]Redback(config-ctx)#router mobile-ip
[local]Redback(config-mip)#home-agent
[local]Redback(config-mip-ha)#registration max-lifetime 3600
```

1.91 release download

release download *url*

1.91.1 Purpose

Installs an alternate software release image or a modular patch on the system.

1.91.2 Command Mode

Exec (10)

1.91.3 Syntax Description

url	URL of a pre-existing configuration file. See the Usage Guidelines section for the format of this argument.

1.91.4 Default

None

1.91.5 Usage Guidelines

Use the **release download** command to install either an alternate software release image or a modular patch on the system.

Note: The **release download** command works in the local context.

The **release download** *url* downloads a new software release image to the alternate system partition.

The new software release image remains inactive until the **release upgrade** command (in exec mode) activates it and installs it in the active system partition. If the system already has an image in the alternate partition, you are prompted for confirmation to allow the system to overwrite it.

Keep the following guidelines (seeTable 6) for the *url* argument in this command.

Table 6 url Argument

Syntax for *url* Argument	Description
[/*device*][/*directory*]/*filename*.*ext*][(1)]	Use when referring to a file on the local file system. The value for the *device* argument can be **flash**, or if a mass-storage device is installed, **md**. If you do not specify the *device* argument, the default value is the device in the current working directory. If you do not specify the *directory* argument, the default value is the current directory. Directories can be nested. The value for the *filename* argument can be up to 256 characters in length.
protocol://*username*[:*passwd*]@{*ip-addr* | *hostname*}[//*directory*]/*filename.ext*	Use when downloading from a remote server. The *protocol* argument is **ftp** or **scp**; that is, File Transfer Protocol (FTP) or Secured Copy Protocol (SCP), respectively. The *username*[:*passwd*] construct specifies the user and an optional password. The *ip-addr* argument is the IP address of the server, and the *hostname* argument is the hostname of the server. If a username is not specified, the SmartEdge router sends the username for the administrator account for the current logon session. The *username*[:*passwd*] construct specifies the user and an optional password. The *ip-addr* argument is the IP address of the server, and the *hostname* argument is the hostname of the server. If a username is not specified, the SmartEdge router sends the username for the administrator account for the current logon session. Use double slashes (//*directory*) if the pathname to the directory on the remote server is an absolute pathname; use a single slash (/*directory*) if it is a relative pathname (under the hierarchy of username account home directory).

*(1) The value for the **filename** argument can be up to 256 characters in length. You can only use the **hostname** argument if Domain Name System (DNS) is enabled with the ip domain-lookup, **ip domain-name**, and **ip name-servers** commands (in context configuration mode); see the Command List.*

1.91.6 Examples

The following example shows how to install a SmartEdge OS image:

```
[local]Redback#release download ftp://guest@10.13.49.10//images/REL_6_4_1/SEOS-6
```

1.92 release download modular

release download modular *url*

1.92.1 Purpose

Installs a modular patch on the system.

1.92.2 Command Mode

Exec (10)

1.92.3 Syntax Description

modular	Optional.
url	URL of a pre-existing configuration file. See the Usage Guidelines section for the format of this argument.

1.92.4 Default

None

1.92.5 Usage Guidelines

Caution!

Risk of system crash. Before using this command, to reduce the risk, contact technical support and verify the patch version, the current SmartEdge router version, and the hardware component versions.

Use the **release download modular** command to install a modular patch on the system. The command places a downloaded patch file on the active partition rather than the alternate boot system partition. Subscriber sessions remain active while the traffic card Packet Processing ASIC (PPA) software is upgraded with the new patch release. If the active partition has insufficient space, an informational log message is generated.

Note: The **release download modular** command works in local context.

Keep the following guidelines (Table 7) for the ***url*** argument in this command.

Table 7 url Argument

Syntax for ***url*** Argument	Description
[**/*device***][**/*directory***]**/*filename.ext***][(1)]	Use when referring to a file on the local file system. The value for the ***device*** argument can be **flash**, or if a mass-storage device is installed, **md**. If you do not specify the ***device*** argument, the default value is the device in the current working directory. If you do not specify the ***directory***argument, the default value is the current directory. Directories can be nested. The value for the ***filename*** argument can be up to 256 characters in length.
protocol**://*username***[**:*passwd***]**@**{***ip-addr*** | ***hostname***}[**//*directory***]**/*filename.ext***	Use when downloading from a remote server. The ***protocol*** argument is **ftp** or **scp**; that is, File Transfer Protocol (FTP) or Secured Copy Protocol (SCP), respectively. The ***username***[**:*passwd***] construct specifies the user and an optional password. The ***ip-addr*** argument is the IP address of the server, and the ***hostname*** argument is the hostname of the server. If a username is not specified, the SmartEdge router sends the username for the administrator account for the current logon session. The ***username***[**:*passwd***] construct specifies the user and an optional password. The ***ip-addr*** argument is the IP address of the server, and the ***hostname*** argument is the hostname of the server. If a username is not specified, the SmartEdge router sends the username for the administrator account for the current logon session. Use double slashes (**//*directory***) if the pathname to the directory on the remote server is an absolute pathname; use a single slash (**/*directory***) if it is a relative pathname (under the hierarchy of username account home directory).

*(1) The value for the **filename** argument can be up to 256 characters in length. You can only use the **hostname** argument if Domain Name System (DNS) is enabled with the **ip domain-lookup**, **ip domain-name**, and **ip name-servers** commands (in context configuration mode); see the Command List.*

1.92.6 Examples

The following example shows how to install a SmartEdge OC modular patch:

```
[local]Redback#release download modular ftp://guest@10.13.49.10//images/REL_6_4_1/
```

1.93 release erase

release erase

1.93.1 Purpose

Manually erases the alternate image on the system.

1.93.2 Command Mode

Exec (10)

1.93.3 Syntax Description

This command has no keywords or arguments.

1.93.4 Default

None

1.93.5 Usage Guidelines

Use the **release erase** command to manually erase the alternate image on the system. You cannot use this command if the system is configured to use the alternate image upon reload.

1.93.6 Examples

The following example shows how to erase the alternate system image:

```
[local]Redback#release erase

The following "alternate" release will be erased:

Version SE800-2.4.4.0.158-Release
Built on Wed Mar  5 10:00:02 PST 2003
Copyright (C) 1998-2003, Redback Networks Inc. All rights reserved.

Are you sure you wish to erase this release? (y/n) y

Erasing the "alternate" release...
```

1.94 release sync

release sync

1.94.1 Purpose

Forces a data synchronization of the system image on the standby controller card with the system image on the primary partition of the active controller card. The standby controller card reboots twice during the process.

1.94.2 Command Mode

Exec (10)

1.94.3 Syntax Description

This command has no keywords or arguments.

1.94.4 Default

None

1.94.5 Usage Guidelines

Use the **`release sync`** command to force a data synchronization of the system image on the standby controller card with the system image on the primary partition of the active controller card.

Note: This command should be used only if you suspect the system image on the standby controller card has been corrupted, such as might occur when a software bug prevents successful automatic synchronization. During normal operations, the system image on the standby controller is automatically synchronized with the active controller and this command is not necessary.

Caution!

During synchronization, the system operates without controller card redundancy. Any failure of the active controller card in this condition causes the system to reboot and loose all active sessions and dynamic routing information.

Caution!

Do not remove the active or standby controller card or reboot the system during this operation.

The **`release sync`** process is not affected by traffic card installation and removal; the active controller, and hence the system, continues to forward traffic and detect and notify the administrator of any faults that occur while the standby controller is being data synchronized (the FAIL LED is blinking).

Use the **`show redundancy`** command (in any mode) to determine whether the key processes on the active and standby controller cards are data synchronized. Both processes and files must be synchronized for full redundancy.

Note: The SmartEdge 100 router does not support this command.

1.94.6 Examples

The following example shows how to force a data synchronization of the system image on the primary partition of the standby controller card with those on the active controller card:

```
[local]Redback#release sync

Apr 10 00:47:09: %DLM-6-INFO: Asked sync client to sync to running image

Apr 10 00:47:10: %LOG-6-PRI_STANDBY: Apr 10 00:47:10:
%DLM-6-INFO: Reloading xcrp for user requested release sync

Apr 10 00:47:10: %LOG-6-PRI_STANDBY: Apr 10 00:47:10: %DLM-6-INFO: Reloading xcr

Apr 10 00:47:10: %LOG-6-PRI_STANDBY: Apr 10 00:47:10: %ALAPI-6-INFO: XCRP in slo

Apr 10 00:47:17 Redback /netbsd: VX Redundancy state:  ###  VX_M2M_LINKUP FALSE

Apr 10 00:47:17 Redback /netbsd: Notifying all processes of M2MDOWN status
```

1.95 release upgrade

release upgrade [**in-service**] [{**at** *at-time*} | {**in** *in-time*}]

no release upgrade

1.95.1 Purpose

Replaces the currently running SmartEdge router (on the primary memory partition) with the SmartEdge router stored on the alternate memory partition.

1.95.2 Command Mode

Exec (10)

1.95.3 Syntax Description

at *at-time*	Optional. Specified time at which to perform the release upgrade. The value for the **at** *at-time* construct is in a *yyyy:mm:dd:hh:mm*[*:ss*] format, where *yyyy* = year, *mm* = month, *dd* = day, *hh* = hour, *mm* = minute, and [*:ss*] is optional seconds.
in *in-time*	Optional. Number of minutes to wait before performing the release upgrade. The value for the **in** *in-time* construct is in a *dd:hh:mm* format, where *dd* = day, *hh* = hour, and *mm* = minute.

1.95.4 Default

None.

1.95.5 Usage Guidelines

Use the **release upgrade** command to replace the currently running SmartEdge OS image (on the primary memory partition) with the image stored on the alternate memory partition. Enter the command from the console port on the active controller card to view the progress of the upgrade operation.

If you enter the **release upgrade** command, the system goes out of service to restart and any subscriber sessions disconnect. After the upgrade, several minutes pass before the standby controller card, if present, automatically synchronizes with the active controller card.

Keep the following guidelines in mind when you use this command:

- After an upgrade is complete, press **Enter** to display the `login` prompt.
- If the router reloads before the specified time, the **at** ***at-time*** or **in** ***in-time*** construct is voided.
- If the upgrade fails (for example, if no valid release image is in the alternate partition), the upgrade fails and a message is logged. Use the **show log** command (in exec mode) to display the stored messages.
- The **release upgrade** command does not cause a switchover to the standby controller card if one is not present.
- Use **no release upgrade** to cancel any pending upgrade previously scheduled using the **at** or **in** keyword options.

1.95.6 Examples

The following example shows how to configure the system to reload using the alternate installed image:

```
[local]Redback#release upgrade
```

```
The system will reboot and the following release will become active:

Version SEOS-5.0.5-Release
Built on Mon Jan  09 01:30:02 PST 2006
Copyright (C) 1998-2006, Redback Networks Inc. All rights reserved.

Are you sure you wish to continue? (y/n) y

Setting boot partition to "alternate"...

The "reload" command will reboot all cards on this system
Do you want to save the current configuration? (y/n) y
.
.
.
Configuration complete
% Startup configuration processing took: 33 seconds
```

1.96 release upgrade modular

release upgrade modular

1.96.1 Purpose

Upgrades the active system image with a modular patch release that has been downloaded to the active partition by the **release download** command.

1.96.2 Command Mode

Exec (10)

1.96.3 Syntax Description

This command has no keywords or arguments.

1.96.4 Default

None

1.96.5 Usage Guidelines

Caution!

Risk of system crash. Before using this command, to reduce the risk, contact technical support and verify the patch version, current operating system version, and hardware component versions.

Use the **release upgrade modular** command to upgrade the active system image with a modular patch release that has been downloaded to the active partition by the **release download** command.

The system remains in service and does not restart, and subscriber sessions remain connected. The primary and alternate boot partitions of the system are not switched.

On the SmartEdge 100 router, subscriber sessions are expected to remain active while the SmartEdge 100 software upgrades to the software modular patch release.

The patch file version is displayed during the upgrade, and the system loads and prompts for confirmation before it proceeds. When installation completes, an informational log message is sent, indicating the success or failure of the operation.

Keep the following guidelines in mind when you use this command:

- After an upgrade is complete, press **Enter** to display the `login` prompt.
- If the upgrade fails (for example, if no valid release image is in the alternate partition), the upgrade fails and a message is logged. Use the **show log** command (in exec mode) to display the stored messages.

1.96.6 Examples

The following example illustrates a successful application of the **release sync in-service** command:

```
[local]Redback#release upgrade modular

[local]Redback#
```

1.97 reload

reload

1.97.1 Purpose

Reloads the system software on the active controller card, and then on the standby controller card.

1.97.2 Command Mode

Exec (15)

1.97.3 Syntax Description

This command has no keywords or arguments.

1.97.4 Default

None

1.97.5 Usage Guidelines

Use the **reload** command to reload the system software on the active controller card, and then on the standby controller card. When you enter this command, the system performs minimal housekeeping, then reloads as if powered off and then powered on again. The system prompts you to confirm the reload. Type **y** to proceed with the reload, or **n** to cancel the reload.

During the reload sequence for a SmartEdge router, the traffic cards are held in low-power mode until the SmartEdge router determines which slot has the active controller card. After the active controller card (and the standby controller card, if it is installed) are initialized, the SmartEdge router then determines if a power capacity check is needed. If the chassis has a single controller card or the active and standby controller cards are identical, the traffic cards are initialized starting with the lowest-numbered slot.

However, if the controller cards do not match, the SmartEdge router performs a power capacity check. Starting with the lowest-numbered traffic card slot, each installed traffic card is initialized and the available power is recalculated.

The SmartEdge router leaves the traffic card in low-power mode if not enough available power exists to initialize it.

Note: If the active and standby controller versions are different, the SmartEdge router allocates power for both controller cards, initializes them, and issues a controller mismatch alarm.

The SmartEdge router always reserves enough power during system configuration so that if the system has only a single controller card installed, a standby controller card of the same type can be installed at a later time.

During the reload sequence for a SmartEdge 100 chassis, the media interface cards (MICs) are held in low-power mode until after the controller carrier card is initialized. Then they are initialized starting with the lowest-numbered MIC slot.

Note: The SmartEdge 100 router does not have a standby controller card.

1.97.6 Examples

The following example reloads the system software on the active controller card, and then on the standby controller card:

```
[local]Redback#reload
```

1.98 reload card

reload card {**all** | ***slot***}

1.98.1 Purpose

Reloads the I/O carrier card in the SmartEdge 100 chassis, a traffic or services card in the specified slot, or all traffic cards in any SmartEdge chassis except the SmartEdge 100 chassis.

1.98.2 Command Mode

Exec (15)

1.98.3 Syntax Description

all	Reloads all traffic cards in the chassis.
slot	Chassis slot number of the traffic or services card or I/O carrier card. The range of values is: • SmartEdge 100 router—2. • SmartEdge 400 router—1 to 4. • SmartEdge 800 or SmartEdge 1200 router—1 to 6 and 9 to 14.

1.98.4 Default

None

1.98.5 Usage Guidelines

Use the **reload card** command to reload the I/O carrier card in the SmartEdge 100 chassis, traffic, service, or storage card in the specified slot, or all traffic cards in any SmartEdge chassis except the SmartEdge 100 chassis.

To reload the active controller card or the controller carrier card, use the **reload** command. If the system has a standby controller card, and a change in software release or configuration on the active controller card is detected after it has been reloaded, the system reloads the standby controller card so that it mirrors (is synchronized with) the active controller card.

Note: If the traffic card that is currently installed in any slot is not the same type as that specified by the slot configuration by using either the **card** command or **port** command (in global configuration mode), the SmartEdge router does not initialize the card; instead it is held in low-power mode with its components in reset mode. If the card types are identical, the system initializes the card.

1.98.6 Examples

The following example shows how to reload the traffic card in slot `1`:

```
[local]Redback#reload card 1
```

1.99 reload disk

reload disk ***slot_num disk_num***

1.99.1 Command Mode

Exec

1.99.2 Syntax Description

slot_num	Chassis slot number of the SSE card.
disk_num	Disk number on the SSE card. Values: 1 or 2.

1.99.3 Default

None.

1.99.4 Usage Guidelines

Gracefully shuts down the specified SSE disk and reloads the SSE disk. This command is equivalent to removing and reinserting the disk.

Note: CPG supports a single hard disk for each SSE card.

If you issue this command on the active SSE card during data synchronization on any partition, the following warning message appears: `Executing the command during data synchronization on any of the partitions will cause data corruption.`

1.99.5 Examples

```
[local]Redback#reload disk 2 2
```

1.100 reload fpga

reload fpga {***slot*** | **mic*mic-slot***}

1.100.1 Purpose

Reloads the code in the field-programmable gate array (FPGA) on a particular traffic card.

1.100.2 Command Mode

Exec (15)

1.100.3 Syntax Description

slot	Chassis slot number of the traffic card to reload. The range of values is 1 to 14.
mic ***mic-slot***	Reloads the code in the FPGA on the media interface card (MIC) in the specified slot. The range of values is 1 to 2. This option applies to the ATM OC MIC only.

1.100.4 Default

None

1.100.5 Usage Guidelines

Use the **reload fpga** command to reload the code in the FPGA on a particular traffic card. This command also upgrades the code should it be required for a new software release.

To use this command, the card must have been configured in the specified slot on the SmartEdge router. On the SmartEdge 100 router, both the MIC and the carrier card must be configured.

Note: On the SmartEdge 100 router, only ATM OC MICs support this command. For all other SmartEdge 100 MICs, FPGAs are loaded from a SmartEdge OS file when the system is booted.

Caution!

Risk of data loss. Depending on the traffic card type, it takes three to ten minutes for the **reload fpga** command to successfully upgrade the FPGA. Interrupting the upgrade can leave the traffic card inoperable. To reduce the risk, do not interrupt the process in the middle of an FPGA upgrade.

1.100.6 Examples

The following example shows how to reload the FPGA on the traffic card in slot 4:

```
[local]Redback#reload fpga 4
```

The following example shows how to reload the FPGA on the ATM OC MIC in slot 2:

```
[local]Redback#reload fpga mic 2
```

1.101 reload mic

reload mic {**1** | **2**}

1.101.1 Purpose

Reloads the specified media interface card (MIC) and all associated components and reformats the compact-flash (CF) card installed in the external slot of the SmartEdge 100 chassis.

1.101.2 Command Mode

Exec (15)

1.101.3 Syntax Description

1	Reloads the MIC with ports 2/3 to 2/14.
2	Reloads the MIC with ports 2/15 to 2/26.

1.101.4 Default

None

1.101.5 Usage Guidelines

Use the **reload mic** command to reload the specified MIC and all associated components and reformat the CF card installed in the external slot of the SmartEdge 100 chassis. Traffic on the specified MIC is momentarily interrupted, while traffic on the unspecified MIC and native Gigabit Ethernet (GE) ports remains unaffected. This command does not affect the Packet Processing ASIC (PPA), forwarding path field-programmable gate array (FPGA), or unspecified MIC.

Use the **reload** command (in exec mode) to reload the entire SmartEdge 100 router. For information about the **reload** command, see the *Command List*.

1.101.6 Examples

The following example shows how to reload the MIC with ports 2/15 to 2/26:

```
[local]Redback#reload mic 2
```

1.102 reload standby

reload standby

1.102.1 Purpose

Reloads the system software on the standby controller card only.

1.102.2 Command Mode

Exec (15)

1.102.3 Syntax Description

This command has no keywords or arguments.

1.102.4 Default

None

1.102.5 Usage Guidelines

Use the **reload standby** command to reload the system software on the standby controller card only.

Note: The SmartEdge 100 router does not have a standby controller card.

1.102.6 Examples

The following example shows how to reload the system software on the standby controller card:

```
[local]Redback#reload standby
```

1.103 reload switch-over

reload switch-over

1.103.1 Purpose

Reloads the system software on the active controller card and, if the standby controller card is ready, cause the standby to become the active controller card.

1.103.2 Command Mode

Exec (15)

1.103.3 Syntax Description

This command has no keywords or arguments.

1.103.4 Default

None

1.103.5 Usage Guidelines

Use the **reload switch-over** command to reload the system software on the active controller card, and if the standby controller card is ready, cause the standby to become the active controller card.

If the standby is not ready, this command performs the same function as the **reload** command. Both controller cards ***reload,*** and the current active controller card remains active.

Note: The SmartEdge 100 router does not have a standby controller card.

1.103.6 Examples

The following example shows how to reload the system software on the active controller card, and if the standby controller card is ready, cause the standby to become the active controller card:

```
[local]Redback#reload switch-over
```

1.104 remote-as

remote-as {***asn*** | ***nn*:*nn***}

no remote-as {***asn*** | ***nn*:*nn***}

1.104.1 Purpose

Configures the autonomous system number (ASN) of the external Border Gateway Protocol (eBGP) neighbor.

1.104.2 Command Mode

BGP neighbor configuration

1.104.3 Syntax Description

asn	ASN in integer format. The range of values is 1 to 65535. The subrange of 64512 to 65535 is reserved for private ASNs.
nn*:*nn	ASN in 4-byte integer format, where the first ***nn*** indicates the two higher-order bytes and the second ***nn*** denotes the two lower-order bytes.

1.104.4 Default

None

1.104.5 Usage Guidelines

Use the **remote-as** command to configure the ASN of the eBGP neighbor.

Use the **no** form of this command to remove the ASN.

1.104.6 Examples

The following example shows how to assign ASN `4001` to the eBGP neighbor at IP address `102.201.2.45`:

```
[local]Redback(config-ctx)#router bgp 100
[local]Redback(config-bgp)#neighbor 102.201.2.45 external
[local]Redback(config-bgp-neighbor)#remote-as 4001
```

1.105 remote-encap

remote-encap {**1qtunnel** | **bridge1483** | **dot1q** | **ethernet**}

no remote-encap {**1qtunnel** | **bridge1483** | **dot1q** | **ethernet**}

1.105.1 Purpose

Specifies the encapsulation type used at the remote end of any XCs that have the specified L2VPN profile attached.

1.105.2 Command Mode

L2VPN profile peer configuration

1.105.3 Syntax Description

1qtunnel	Specifies 802.1Q tunnel encapsulation.
bridge1483	Specifies ATM RFC 1483 bridged encapsulation.
dot1q	Specifies 802.1Q Ethernet encapsulation.
ethernet	Specifies Ethernet encapsulation.

1.105.4 Default

No encapsulation is configured for the remote end of an XC.

1.105.5 Usage Guidelines

Use the **remote-encap** command to specify the encapsulation type used at the remote end of any XCs that have the specified L2VPN profile attached.

Use the **no** form of this command to remove the encapsulation configuration for the remote end of an XC from an L2VPN profile.

1.105.6 Examples

The following example shows how to specify that 802.1Q tunnel encapsulation is used at the remote end of any XCs that have the L2VPN profile called `802tun` attached:

```
[local]Redback(config)#l2vpn profile 802tun
[local]Redback(config-l2vpn-xc-profile)#peer 100.100.100.1
[local]Redback(config-l2vpn-xc-profile-peer)#remote-encap 1qtunnel
```

1.106 remotefile

remotefile format *format-string* [*OS-variable*] [*OS-variable*] ...

no remotefile format

1.106.1 Purpose

Specifies the format of the filename and the location of the bulkstats collection files that are stored on remote file servers.

1.106.2 Command Mode

bulkstats configuration

1.106.3 Syntax Description

format	Specifies the format of the filename for the bulkstats collection files.
format-string	Describes the format strings used to format the remote filename for the bulkstats collection files. Format strings can contain anything or nothing as a label for a SmartEdge OS variable. They follow the C programming language printf() function syntax and must be enclosed in quotation marks.
OS-variable	Optional. SmartEdge OS system variable. describes the supported variables.

1.106.4 Default

No filename format is defined for bulkstats collection files for any policy.

1.106.5 Usage Guidelines

Use the **remotefile** command to specify the format of the filename and the location of the bulkstats collection files that are stored on remote file servers.

Table 8 describes the format strings used to format the remote filename.

Table 8 Supported Variables for the remotefile Command

Variable	Description	Type
context	Context name	String
date	Today's date in *YYYYMMDD* format	String

Table 8 Supported Variables for the remotefile Command

Variable	Description	Type
epochtime	Time of day in epoch format (seconds since January 1, 1970)	Integer
hostname	Hostname as specified in the configuration file	String
policy	Bulkstats policy name	String
sysuptime	System uptime in seconds	Integer
timeofday	Time of day in ***HHMMSS*** format (using a 24-hour clock)	String

You cannot change the remote filename or location 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 **receiver** command.

Use the **no** form of this command to delete information about the format of the remote filename and location used to store bulkstats data for this policy.

1.106.6 Examples

The following example shows how to specify the format of the filename where the bulkstats data for the `bulk` policy is to be stored:

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

[local]Redback(config-ctx)#bulkstats policy bulk


[local]Redback(config-bulkstats)#remotefile format  "Bulks
tats/%s_%s" hostname timeofday
```

The file is specified as Bulkstats/***hostname_HHMMSS*** where the ***hostname*** argument is the name configured for the SmartEdge router and the ***HHMMSS*** argument is the hour, minute, and second (24-hour clock) of the transfer.

To see how this information displays, see the example for the **show bulkstats** command in *Configuring Bulkstats*.

1.107 remove-private-as

remove-private-as

no remove-private-as

1.107.1 Purpose

Removes private autonomous system numbers (ASNs) from routes that are advertised to the Border Gateway Protocol (BGP) neighbor address family or peer group address family.

1.107.2 Command Mode

- BGP neighbor address family configuration
- BGP peer group address family configuration

1.107.3 Syntax Description

This command has no keywords or arguments.

1.107.4 Default

The ASNs are not removed.

1.107.5 Usage Guidelines

Use the **remove-private-as** command to remove private ASNs from routes that are advertised to the BGP neighbor address family or peer group address family.

Use the **no** form of this command to send private ASNs.

1.107.6 Examples

The following example advertises BGP unicast routes to the neighbor at IP address `102.21.2.45`. Any ASNs contained in these routes are removed:

```
[local]Redback(config-ctx)#router bgp 100
[local]Redback(config-bgp)#neighbor 102.201.2.45 external
[local]Redback(config-bgp-neighbor)#address-family ipv4 unicast
[local]Redback(config-bgp-peer-af)#remote-as 200
[local]Redback(config-bgp-peer-af)#remove-private-as
```

1.108 rename

rename ***current-url new-url*** [**-noconfirm**]

1.108.1 Purpose

Renames a file or directory on the local file system.

1.108.2 Command Mode

Exec (10)

1.108.3 Syntax Description

current-url	Current URL of the file (or directory) that is to be renamed.
new-url	URL of the file (or directory) after renaming.
-noconfirm	Optional. Replaces an existing file (or directory) without asking for confirmation.

1.108.4 Default

None

1.108.5 Usage Guidelines

Use the **rename** command to rename a file or directory on the local file system. The ***current-url*** and ***new-url*** arguments use the following form:

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

The value for the ***device*** argument can be **flash**, or if a mass-storage device is installed, **md**. If you do not specify the ***device*** argument, the default value is the device in the current working directory. If you do not specify the ***directory*** argument, the default value is the current directory. Directories can be nested. The value for the ***filename*** argument can be up to 256 characters in length.

This command works only for renaming files and directories on a single local file system device; that is, the URLs must be identical, except for the ***filename.ext*** argument. The command fails if the values of the ***current-url***and ***new-url*** arguments are identical; this is the URLs are identical.

A file with the new name must not already exist; that is, the SmartEdge router does not overwrite an existing file on the local file system without first seeking confirmation. Use the **-noconfirm** optional keyword to avoid the confirmation prompt.

1.108.6 Examples

The following example shows how to rename the file, `redback.bin`, to `old.bin`:

```
[local]Redback#rename /flash/redback.bin /flash/old.bin
```

1.109 replay-tolerance

replay-tolerance*seconds*

no replay-tolerance

1.109.1 Purpose

Configures the tolerance for timestamp-based replay protection used between the home agent (HA) instance and the registering mobile nodes (MN).

1.109.2 Command Mode

HA configuration

1.109.3 Syntax Description

seconds	Tolerance for timestamp-based replay protection used between the HA instance and registering MNs. The range of values is 4 to 255 seconds.

1.109.4 Default

The default for tolerance for timestamp-based replay protection is 7 seconds.

1.109.5 Usage Guidelines

Use the **replay-tolerance** command to configure the tolerance for timestamp-based replay protection used between the HA instance and the registering MN. The **replay-tolerance** command specifies the number of seconds that the HA instance timestamp and MN timestamp can be different. When the HA instance discovers that this difference is greater than the number of seconds specified, it rejects the MN registration.

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

1.109.6 Examples

The following example shows how to configure a timestamp-based replay tolerance of `10` seconds for this HA instance:

```
[local]Redback(config)#context ha
[local]Redback(config-ctx)#router mobile-ip
[local]Redback(config-mip)#home-agent
[local]Redback(config-mip-ha)#replay-tolerance 10
```

1.110 report

report {**tx-speed** *tx-kbps* | **rx-speed** *rx-kbps*}

{**no** | **default**} **report** {**tx-speed** | **rx-speed**}

1.110.1 Purpose

Specifies the transmit and receive speeds to be included in the IEFT standard, Layer 2 Tunneling Protocol (L2TP) Rx Connect Speed attribute-value pair (AVP) 24 and Tx Connect Speed AVP 38 for any Asynchronous Transfer Mode (ATM) permanent virtual circuit (PVC) that references this ATM profile.

1.110.2 Command Mode

- ATM profile configuration
- dot1q profile configuration

1.110.3 Syntax Description

tx-speed *tx-kbps*	Transmit speed, in kbps, to be included in L2TP AVP 38; the range of values is 1 to 4,294,967,296.
rx-speed *rx-kbps*	Receive speed, in kbps, to be included in L2TP AVP 24; the range of values is 1 to 4,294,967,296.

1.110.4 Default

The RxConnect Speed is the port speed. The TxConnect Speed depends on the ATM traffic class specified for the profile; see Table 9.

1.110.5 Usage Guidelines

Use the **report** command to specify the receive and transmit speeds to be included in the IEFT standard L2TP Rx Connect Speed AVP 24 and Tx Connect Speed AVP 38 for any ATM PVC that references this ATM profile.

Use the **no** or **default** form of this command to report default values in L2TP AVPs 24 and 38. Table 9 lists the default values for the TxConnect speed; for all traffic classes except UBR, the reported default value is the value of the specified argument in the **shaping** command (in ATM profile configuration mode).

Table 9 Default Values for TxConnect Speed

ATM Traffic Class	Default TxConnect Speed Reported
CBR	Value of the ***rate*** argument
UBR	Port speed
UBR pcr	Value of the ***pcr*** argument
UBRe	Value of the ***pcr*** argument
VBR-nrt	Value of the ***scr*** argument
VBR-rt	Value of the ***scr*** argument

1.110.6 Examples

The following example shows how to specify the receive and transmit speeds as `2400` kbps in an ATM profile, `low_rate`:

```
[local]Redback(config)#atm profile low_rate
[local]Redback(config-atm-profile)#shaping vbr-nrt pcr 2500 cdvt 20 scr 2400 bt 10
[local]Redback(config-atm-profile)#report tx-speed 2500
[local]Redback(config-atm-profile)#report rx-speed 2500
```

1.111 resequence as-path-list

resequence as-path-list ***apl-name***

1.111.1 Purpose

Assigns new sequence numbers to existing entries in the specified autonomous system (AS) path list so that entries are in increments of 10.

1.111.2 Command Mode

context configuration

1.111.3 Syntax Description

apl-name	Name of the AS path list to be resequenced.

1.111.4 Default

Sequence numbers are assigned by the system in increments of 10.

1.111.5 Usage Guidelines

Use the **resequence as-path-list** command to assign new sequence numbers to existing entries in the specified AS path list so that entries are in increments of 10.

This command is useful when you have manually assigned sequence numbers and have no room to insert new entries in between existing entries. You can manually assign sequence numbers using the **seq** ***seq-num*** construct in the **as-path-list** command in context configuration mode.

Note: Two resequence commands, **resequence ip access-list** and **resequence policy access-list**, are not included in this document. For more information on these commands, see the *Command List*.

1.111.6 Examples

The following example shows how to resequence entries in the AS path list, `filter1`, by increments of 10:

```
[local]Redback(config-ctx)#resequence as-path-list filter1
```

1.112 resequence community-list

resequence community-list ***cl-name***

1.112.1 Purpose

Assigns new sequence numbers to existing entries in the specified community list so that entries are in increments of 10.

1.112.2 Command Mode

context configuration

1.112.3 Syntax Description

cl-name	Name of the community list to be resequenced.

1.112.4 Default

Sequence numbers are assigned by the system in increments of 10.

1.112.5 Usage Guidelines

Use the **resequence community-list** command to assign new sequence numbers to existing entries in the specified community list so that entries are in increments of 10.

This command is useful when you have manually assigned sequence numbers and have no room to insert new entries in between existing entries. You can manually assign sequence numbers using the **seq** ***seq-num*** construct in the **community-list**command in context configuration mode.

Note: Two resequence commands, **resequence ip access-list** and **resequence policy access-list**, are not included in this document. For more information on these commands, see the *Command List*.

1.112.6 Examples

The following example shows how to resequence entries in the community list, `cl012`, by increments of 10:

```
[local]Redback(config-ctx)#resequence community-list cl012
```

1.113 resequence ext-community-list

resequence ext-community-list ***ecl-name***

1.113.1 Purpose

Assigns new sequence numbers to existing entries in the specified extended community list so that entries are in increments of 10.

1.113.2 Command Mode

context configuration

1.113.3 Syntax Description

ecl-name	Name of the extended community list to be resequenced.

1.113.4 Default

Sequence numbers are assigned by the system in increments of 10.

1.113.5 Usage Guidelines

Use the **resequence ext-community-list** command to assign new sequence numbers to existing entries in the specified extended community list so that entries are in increments of 10.

This command is useful when you have manually assigned sequence numbers and have no room to insert new entries in between existing entries. You can manually assign sequence numbers using the **seq** ***seq-num*** construct in the **ext-community-list** command in context configuration mode.

1.113.6 Examples

The following example shows how to resequence entries in the extended community list, `ecl05`, by increments of 10:

```
[local]Redback(config-ctx)#resequence ext-community-list ecl05
```

1.114 resequence ip access-list

resequence ip access-list ***acl-name***

1.114.1 Purpose

Reassigns sequence numbers to the entries in the specified IP access control list (ACL) to be in increments of 10.

1.114.2 Command Mode

context configuration

1.114.3 Syntax Description

acl-name	Name of the ACL to be resequenced.

1.114.4 Default

No resequencing is performed.

1.114.5 Usage Guidelines

Use the **`resequence ip access-list`** command to reassign sequence numbers to the entries in the specified IP ACL to be in increments of 10. This command is useful if manually assigned sequence numbers have left no room between entries for additional entries.

1.114.6 Examples

The following example shows how to resequence the statements in the ACL, `fremont1`:

```
[local]Redback(config-ctx)#resequence ip access-list fremont1
```

1.115 resequence ip prefix-list

`resequence ip prefix-list` ***`pl-name`***

1.115.1 Purpose

Assigns new sequence numbers to existing entries in the specified IP prefix list so that entries are in increments of 10.

1.115.2 Command Mode

context configuration

1.115.3 Syntax Description

pl-name	Name of the IP prefix list to be resequenced.

1.115.4 Default

Sequence numbers are assigned by the system in increments of 10.

1.115.5 Usage Guidelines

Use the **resequence ip prefix-list** command to assign new sequence numbers to existing entries in the specified IP prefix list so that entries are in increments of 10.

This command is useful when you have manually assigned sequence numbers and have no room to insert new entries in between existing entries. You can manually assign sequence numbers using the **seq** *seq-num* construct in the **ip prefix-list** command in context configuration mode.

1.115.6 Examples

The following example shows how to resequence entries in the prefix list, `pl226`, by increments of 10:

```
[local]Redback(config-ctx)#resequence ip prefix-list pl226
```

1.116 resequence ipv6 prefix-list

resequence ipv6 prefix-list *ipv6-pl-name*

1.116.1 Purpose

Assigns new sequence numbers to existing entries in the specified IP Version 6 (IPv6) prefix list so that entries are in increments of 10.

1.116.2 Command Mode

context configuration

1.116.3 Syntax Description

ipv6-pl-name	Name of the IPv6 prefix list to be resequenced.

1.116.4 Default

Sequence numbers are assigned by the system in increments of 10.

1.116.5 Usage Guidelines

Use the **resequence ipv6 prefix-list** command to assign new sequence numbers to existing entries in the specified IPv6 prefix list so that entries are in increments of 10.

This command is useful when you have manually assigned sequence numbers and have no room to insert new entries in between existing entries. You can manually assign sequence numbers using the **seq** *seq-num* construct in the **ipv6 prefix-list** command in context configuration mode.

1.116.6 Examples

The following example shows how to resequence entries in the prefix list, ipv6p65, by increments of 10:

```
[local]Redback(config-ctx)#resequence ipv6 prefix-list ipv6pl65
```

1.117 resequence policy access-list

resequence policy access-list *acl-name*

1.117.1 Purpose

Reassigns sequence numbers to the entries in the specified policy access control list (ACL) to be in increments of 10.

1.117.2 Command Mode

context configuration

1.117.3 Syntax Description

acl-name	Name of the ACL to be resequenced.

1.117.4 Default

No resequencing is performed.

1.117.5 Usage Guidelines

Use the **resequence policy access-list** command to reassign sequence numbers to the entries in the specified policy ACL to be in increments of 10. This command is useful if manually assigned sequence numbers have left no room between entries for additional entries.

1.117.6 Examples

The following example shows how to resequence the statements in the policy ACL, `oakland2`:

```
[local]Redback(config-ctx)#resequence policy access-list oakland2
```

1.118 resequence route-map

resequence route-map ***map-name***

1.118.1 Purpose

Assigns new sequence numbers to existing entries in the specified route map so that entries are in increments of 10.

1.118.2 Command Mode

context configuration

1.118.3 Syntax Description

map-name	Name of the route map to be resequenced.

1.118.4 Default

Sequence numbers are assigned by the system in increments of 10.

1.118.5 Usage Guidelines

Use the **resequence route-map** command to assign new sequence numbers to existing entries in the specified route map so that entries are in increments of 10.

This command is useful when you have manually assigned sequence numbers and have no room to insert new entries in between existing entries. You can manually assign sequence numbers using the **seq** ***seq-num*** construct in the **route-map** command in context configuration mode.

1.118.6 Examples

The following example shows how to resequence entries in the route map, `rm045`, by increments of 10:

```
[local]Redback(config-ctx)#resequence route-map rm045
```

1.119 reserved

reserved ***bytes***

no reserved

1.119.1 Purpose

Specifies the number of additional nonstandard Layer 1 overhead bytes reserved, per packet, for a specific access-line type.

1.119.2 Command Mode

- overhead profile configuration
- overhead type configuration

1.119.3 Syntax Description

bytes	Number of reserved bytes, per packet, for the specified access-line type. The range of values is 1 to 255; the default value is 0.

1.119.4 Default

No additional nonstandard Layer 1 overhead bytes are reserved.

1.119.5 Usage Guidelines

Use the **reserved** command to specify the number of additional nonstandard Layer 1 overhead bytes reserved, per packet, for a specific access-line type.

Use the **no** form of this command to remove the specified bytes, per packet from the access-line configuration.

1.119.6 Examples

The following example shows how to configure an overhead profile for `example1`, and set the encapsulation type to `pppoa-llc`. After you set the default values, you set the data type to `adsl`, the rate factor to `20`, and the reserved value to `16`:

```
[local]Redback(config)#qos profile example1 overhead
[local]Redback(config-profile-overhead)#encaps-access-line pppoa-llc
[local]Redback(config-profile-overhead)#reserved 8
[local]Redback(config-profile-overhead)#type adsl1
[local]Redback(config-type-overhead)#rate-factor 20
[local]Redback(config-type-overhead)#reserved 16
```

1.120 res-prefix

res-prefix ***res-prefix***

no res-prefix

1.120.1 Purpose

Configures part of the resource prefix.

1.120.2 Command Mode

SNMP alarm model configuration

1.120.3 Syntax Description

res-prefix	An OID used to construct the value of alarmActiveResou rceId. The OID becomes a prefix part to the constructed alarmActiveResourceId. The remaining IDs are obtained from the OID matched by using the **vb-subtree** command.

1.120.4 Default

None

1.120.5 Usage Guidelines

Use the **res-prefix** command in conjunction with the **vb-subtree** command to create the value of resource prefix for the alarm you are configuring. The value of the resource prefix is determined by appending any indices created by the **vb-subtree** command. If the value of **res-prefix** is not set, then the prefix outlined by **vb-subtree** is used as the resource prefix.

Use the **no** form of this command to remove this portion of the resource prefix.

1.120.6 Examples

The following example shows how to configure the resource prefix with a value of **AlarmID**:

```
[local]jazz#config
[local]jazz(config)#snmp alarm model 1 state clear
[local]jazz(config-snmp-alarmmodel)#no vb-subtree
[local]jazz(config-snmp-alarmmodel)#res-prefix AlarmID
[local]jazz(config-snmp-alarmmodel)#exit
```

1.121 restricted

restricted

{**no** | **default**} **restricted**

1.121.1 Purpose

Specifies that circuits (including Virtual Private LAN Services (VPLS) circuits) to which this profile is assigned are restricted to accepting only source packets from statically allowed medium access control (MAC) addresses.

1.121.2 Command Mode

bridge profile configuration

1.121.3 Syntax Description

This command has no keywords or arguments.

1.121.4 Default

Circuits are not restricted.

1.121.5 Usage Guidelines

Use the **restricted** command to specify that circuits (including VPLS circuits) to which this profile is assigned are restricted to accepting only packets from statically allowed MAC addresses. Learning is not possible on restricted circuits.

This command causes all MAC addresses previously learned for a circuit to which this profile is assigned to be erased. It also prevents learning of MAC addresses on the circuit, because packets from unknown MAC addresses are dropped before they are learned.

Use the **no** or **default** form of this command to remove the restriction from the profile.

1.121.6 Examples

The following example shows how to specify that the MAC addresses be restricted for any circuit to which this profile is assigned:

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

1.122 retain-ibgp-routes

retain-ibgp-routes

{**no** | **default**} **retain-ibgp-routes**

1.122.1 Purpose

Forces the Border Gateway Protocol (BGP) neighbor to retain routes from an internal BGP (iBGP) peer when the peer has restarted, provided the peer supports a graceful restart.

1.122.2 Command Mode

BGP neighbor configuration

1.122.3 Syntax Description

This command has no keywords or arguments.

1.122.4 Default

The command is disabled.

1.122.5 Usage Guidelines

Use the **retain-ibgp-routes** command to force the BGP neighbor to retain routes from an iBGP peer when the peer has restarted, provided the peer supports a graceful restart.

By default, routes are not retained for an iBGP peer after the peer restarts unless all iBGP peers support a graceful restart. However, in some network topologies, it may be desirable and feasible to retain the routes for an iBGP peer, even if not all iBGP peers support a graceful restart.

Use the **no** or **default** form of this command to disable this feature.

1.122.6 Examples

The following example shows how to force the BGP neighbor, `10.1.1.1`, to retain routes from an iBGP peer once the peer has restarted, provided the peer supports a graceful restart:

```
[local]Redback(config-bgp)#neighbor 10.1.1.1 internal
[local]Redback(config-bgp-neighbor)#retain-ibgp-routes
```

1.123 retransmit-interval

retransmit-interval ***interval***

{**no** | **default**} **retransmit-interval**

1.123.1 Purpose

Modifies the interval at which link-state advertisements (LSAs) retransmissions are sent out through the specified interface, sham link, or virtual link.

1.123.2 Command Mode

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

1.123.3 Syntax Description

interval	Interval, in seconds, at which LSA transmissions are sent. The range of values is 1 to 65535; the default value is 5.

1.123.4 Default

LSA retransmissions are sent every five seconds.

1.123.5 Usage Guidelines

Use the **retransmit-interval** command to modify the interval at which LSA retransmissions are sent out through the specified interface, sham link, or virtual link.

When a SmartEdge router sends LSAs to neighbors, it expects to receive an acknowledgment packet within a set amount of time. If the SmartEdge router does not receive an acknowledgment, it retransmits the LSA.

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

1.123.6 Examples

The following example shows how to configure an OSPF interface to retransmit LSAs every 7 seconds:

```
[local]Redback(config-ospf-if)#retransmit-interval 7
```

1.124 retry

retry ***count***

(**no** | **default**} **retry**

1.124.1 Purpose

Specifies the number of times an unacknowledged control message is retransmitted to a Layer 2 Tunneling Protocol (L2TP) peer before the tunnel is brought down.

1.124.2 Command Mode

L2TP peer configuration

1.124.3 Syntax Description

count	Number of times an unacknowledged control message is retransmitted to a peer. The range of values is 1 to 100; the default value is 10.

1.124.4 Default

An unacknowledged control message is retransmitted ten times.

1.124.5 Usage Guidelines

Use the **retry** command to specify the number of times an unacknowledged control message is retransmitted to an L2TP peer before the tunnel is brought down. You may want to increase the value from the default of 10 if the L2TP media is not reliable.

Use the **no** or **default** form of this command to set the number of retransmissions to the default.

1.124.6 Examples

The following example shows how to configure the peer so that unacknowledged control messages are retransmitted five times before the tunnel is brought down:

```
[local]Redback(config-ctx)#l2tp-peer name peer1
[local]Redback(config-l2tp)#retry 5
```

1.125 revert (APS)

revert [***wtr-interval***]

{**no** | **default**} **revert**

1.125.1 Purpose

Sets the switching algorithm to revertive switching and the wait-to-restore (WTR) interval for an Automatic Protection Switching (APS) or Multiplex Section Protection (MSP) group.

1.125.2 Command Mode

APS configuration

1.125.3 Syntax Description

wtr-interval	Optional. Time to wait before reverting to the working port after it is up. The range of values is 1 to 60 minutes; the default value is 5.

1.125.4 Default

The switching algorithm is nonrevertive.

1.125.5 Usage Guidelines

Use the **revert** command to set the switching algorithm to revertive switching and the WTR value for an APS/MSP group.

If you specify this command without the optional ***wrt-interval*** argument, the system uses the default value.

Use the **no** form of this command to set the switching algorithm to nonrevertive switching, that is, an infinite WRT.

Use the **default** form of this command to set the WTR to 5 minutes.

1.125.6 Examples

The following example shows how to set the switching algorithm to revertive with a WRT of 3 minutes:

```
[local]Redback(config)#aps group lab48 pos
[local]Redback(config-aps)#revert 3
```

1.126 revert (SSE)

revert

no | **default revert**

1.126.1 Purpose

Configures the redundant group to always use the primary SSE or disk as active when available.

1.126.2 Command Mode

SSE group configuration

1.126.3 Syntax Description

This command has no keywords or arguments.

1.126.4 Default

Redundancy is nonrevertive by default.

1.126.5 Usage Guidelines

This command can only be configured for network-redundant SSE groups. Configures the redundant group to always use the primary SSE or disk as active when available.

On primary SSE failover, the secondary takes on the active redundancy state and continues to support data transaction on the SSE group. Configure the **revert** command to use the primary as the active device when it becomes available again.

If configured, the primary reverts to the active device when the following conditions are met:

- The primary unit is functioning
- All configured partitions are functioning
- Data is synchronized on all mirrored partitions
- The secondary active was not previously manually switched over. A manual switchover prevents revert from occurring immediately after a manual switchover is completed; in this case you must manually perform the switchover to switch the primary back to active.

Use the **no** or **default** form of this command to return to the default of state of nonrevertive.

1.126.6 Examples

The following example shows how to configure the redundant group to always use the primary SSE or disk as active when available.

```
[local]Redback(config)#sse group sse_group_1 network-redundant
[local]Redback(config-SE-group)#revert
```

1.127 revocation

revocation [**mobile-notify** *condition*] [**timeout** *seconds*] [**retransmit** *num*]

no revocation [**mobile-notify** *condition*] [**timeout** *seconds*] [**retransmit** *num*]

1.127.1 Purpose

Configures registration revocation for this foreign agent (FA) instance.

1.127.2 Command Mode

FA configuration

1.127.3 Syntax Description

mobile-notify *condition*	Optional. Specifies the conditions for which the SmartEdge router notifies mobile nodes (MNs) that their Mobile IP service has been revoked, according to one of the following keywords: • **always**—Always notify the MNs. • **never**—Never notify the MNs. • **home-dictate**—Notify the MNs based on the home-agent (HA) preference specified by the setting I-bit in received registration revocation requests and replies. This is the default.
timeout *seconds*	Number of seconds between registration revocation messages. The range of values is 1 to 100; the default value is 7.
retransmit *num*	Number of times the SmartEdge router transmits registration revocation messages. The range of values is 1 to 100; the default value is 3.

1.127.4 Default

Registration revocation is not configured for any FA instance.

1.127.5 Usage Guidelines

Use the **revocation** command to configure registration revocation for this FA instance. For more information, see RFC 3543, ***Registration Revocation in Mobile IPv4***.

Use the **no** form of this command to remove the registration from the configuration for this FA instance.

1.127.6 Examples

The following example shows how to configure this FA instance to `always` notify the MNs when service is revoked:

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

[local]Redback(config-ctx)#router mobile-ip

[local]Redback(config-mip)#foreign-agent

[local]Redback(config-mip-fa)#revocation mobile-notify always
```

1.128 revocation (HA)

revocation [**mobile-notify** {**always** | **never** | **foreign-dictate**}] [**timeout** *seconds*] [**retransmit** *num*]

no revocation [**mobile-notify** *condition*] [**timeout** *seconds*] [**retransmit** *num*]

1.128.1 Purpose

Configures registration revocation as described in RFC 3543*Registration Revocation in Mobile IPv4*, for this home agent (HA) instance. Registration revocation is negotiated between the HA instance and its foreign agent (FA) peers.

1.128.2 Command Mode

HA configuration

1.128.3 Syntax Description

mobile-notify *condition*	Optional. Specifies the conditions for which the HA instance negotiates I-bit support with its FA peers when the mobile node (MN) registers, according to one of the following keywords: • **always**—Always notify the MN when Mobile IP services have been revoked, except when the MN is no longer receiving service from the FA peer. This is the default. • **never**—Never notify the MN that Mobile IP services have been revoked. • **foreign-dictate**—Does not negotiate I-bit support with the FA peer when the MN registers. The FA peer determines whether to notify the MN.
timeout *seconds*	Number of seconds between registration revocation retransmissions. A registration revocation request is retransmitted to the FA peer when an acknowledgement is not received. The range of values is 1 to 100; the default value is 7.
retransmit *num*	Number of times the SmartEdge router retries transmission registration revocation messages. The range of values is 1 to 100; the default value is 3.

1.128.4 Default

Registration revocation is not configured for any HA instance.

1.128.5 Usage Guidelines

Use the **revocation** command to configure registration revocation, as described in RFC 3543, Registration Revocation in Mobile IPv4, for this HA instance. Registration revocation is negotiated between the HA instance and its FA peers.

Note: To use registration revocation, you must configure authentication with the revocation command. If authentication is not enabled for the FA peer, registration revocation is not negotiated for registrations received from that peer. For more information about authentication, see the **authentication** command (in HA configuration or FA peer configuration mode).

Use the **no** form of this command to disable support for registration revocation for the HA instance.

1.128.6 Examples

The following example shows how to enable registration revocation support for the HA instance. Registration revocation I-bit support is negotiated with the FA peer and the MN is never notified that Mobile IP services have been revoked:

```
[local]Redback(config)#context ha
[local]Redback(config-ctx)#router mobile-ip
[local]Redback(config-mip)#home-agent
[local]Redback(config-mip-ha)#revocation mobile-notify never
```

1.129 rmdir

rmdir ***url***

1.129.1 Purpose

Removes a directory from the local file system.

1.129.2 Command Mode

Exec (10)

1.129.3 Syntax Description

url	URL of the directory to be removed.

1.129.4 Default

None

1.129.5 Usage Guidelines

Use the **rmdir** command to remove a directory on the local file system.

When referring to a directory on the local file system, the URL takes the following form:

[***/device***][***/directory***]...[***/directory***]

The value for the ***device*** argument can be **flash**, or if a mass-storage device is installed, **md**. If you do not specify the ***device*** argument, the default value is the device in the current working directory. If you do not specify the ***directory*** argument, the default value is the current directory. Directories can be nested. The value for the ***filename*** argument can be up to 256 characters in length.

Before you remove a directory, you must remove all files from the directory using the **delete** command.

1.129.6 Examples

The following example shows how to remove the top-level directory, `backups`, from the flash file system:

```
[local]Redback#rmdir /flash/backups
```

1.130 rmon alarm

rmon alarm ***index object-id interval*** {**absolute** | **delta**} **rising-threshold** ***value*** [***event-index***] **falling-threshold** ***value*** [***event-index***] [**owner** ***owner-name***]

{no | default} rmon alarm ***index***

1.130.1 Purpose

Defines a Remote Monitoring (RMON) alarm and associates it with the RMON event that reports the alarm when its criteria are met.

1.130.2 Command Mode

Global configuration

1.130.3 Syntax Description

index	Index that uniquely identifies an alarm event with an entry in the alarm table in the RMON Management Information Base (RMON-MIB).
object-id	Object ID (OID) of the MIB object to be monitored.
interval	Sampling time in seconds. The range of values is 1 to 2,147,483,647.
absolute	Compares the actual object value against the threshold value.
delta	Compares the difference between successive samples of the object value against the threshold value.
rising-threshold ***value***	Value at which an alarm event is triggered.
event-index	Optional. Index of the entry in the event table in the RMON-MIB that is associated with the alarm event.
falling-threshold ***value***	Value at which an alarm event is triggered.
owner ***owner-name***	Optional. Name of the alarm owner.

1.130.4 Default

No RMON alarms are defined.

1.130.5 Usage Guidelines

Use the **rmon alarm** command to define an RMON alarm and to associate it with the RMON event that reports the alarm when its criteria are met.

Keep the following guidelines in mind when you use the **rmon alarm** command:

- Enable the Simple Network Management Protocol (SNMP) server using the **snmp server** command (in global configuration mode) before you use this command.
- Before you define RMON alarms, define the RMON events that will describe and report the RMON alarms when they occur. Use the **rmon event** command (in global configuration mode) and save the index identifiers of the event entries, because the **rmon alarm** command uses them.
- You can use this command multiple times to define multiple RMON alarms.
- If you configure an RMON alarm on an invalid OID, SNMP warning log messages are generated. Use the **no** form of the **rmon alarm** command to remove the invalid alarm configuration; otherwise, the SNMP daemon removes the invalid RMON alarm entry after 50 minutes.

Use the **no** or **default** form of this command to delete an entry from the RMON alarm table.

1.130.6 Examples

The following example shows how to configure entries in the RMON events table with index identifiers `11` and `12`. Then it shows how to define an RMON alarm that triggers when the difference between successive 60-second samples of the `ipForwDatagrams` alarm rises faster than `3,000,000` or drops faster than `1,000,000`:

```
[local]Redback(config)#rmon event 11 log notify owner gold.isp.net description
"packets per second rising too quickly in context gold.isp.net"

[local]Redback(config)#rmon event 12 log notify owner gold.isp.net description
"packets per second falling too quickly in context gold.isp.net"

[local]Redback(config)#rmon alarm 1 ipForwDatagrams.0 60 delta rising-threshold 3000000 11
falling-threshold 1000000 12 owner gold.isp.net
```

1.131 rmon event

rmon event *index* [**log**] [**notify**] [**owner** *owner-name*] [**description** *text*]

{**no** | **default**} **rmon event** *index*

1.131.1 Purpose

Defines a Remote Monitoring (RMON) event.

1.131.2 Command Mode

Global configuration

1.131.3 Syntax Description

index	Index that uniquely identifies an event with an entry in the event table in the RMON Management Information Base (RMON-MIB).
log	Optional. Specifies that the event generates an entry in the RMON-MIB log table.
notify	Optional. Specifies that the event generates an SNMP notification.
owner ***owner-name***	Optional. Owner of the event.
description ***text***	Optional. Description of the event.

1.131.4 Default

No RMON events are defined

1.131.5 Usage Guidelines

Use the **rmon event** command to define an RMON event and optionally to provide a description of the event.

You must enable the SNMP server using the **snmp server** command (in global configuration mode) before you use this command.

If notification is enabled using the **notify** keyword, the SNMP notification is sent to the destination obtained from the SNMP-NOTIFICATION-MIB and the SNMP-TARGET-MIB, as configured by one or more **snmp target** or **snmp notify-target** commands as either an SNMP trap or inform protocol data unit (PDU).

Use the **no** or **default** form of this command to delete an entry from the RMON event table.

1.131.6 Examples

The following example shows an RMON event that is saved in the SNMP log table and sends an SNMP notification:

```
[local]Redback(config)#rmon event 1 log notify owner gold.isp.net description "packets per second too
high in context gold.isp.net"
```

1.132 robust

robust ***packet-number***

1.132.1 Purpose

Configures the number of IGMP packets that can be lost before group membership for a specified Ethernet bridge expires.

1.132.2 Command Mode

IGMP snooping bridge configuration

1.132.3 Syntax Description

packet-number	Expected packet loss for this bridge. The range of values is 2 to 7; the default value is 2.

1.132.4 Default

The default expected packet loss for a bridge is 2 packets.

1.132.5 Usage Guidelines

Use the **robust** command to configure the number of IGMP packets that can be lost before group membership for a specified Ethernet bridge expires.

The packets that can be lost are IGMP reports. If a host fails to respond to a membership query for two successive intervals, that host is dropped from the outgoing circuit list.

1.132.6 Examples

The following example shows how to configure the expected packet loss for a bridge called `br-sj-1` to be `4` packets:

```
local]Router(config)#context sj1

[local]Router(config-ctx)#bridge br-sj-1

[local]Router(config-bridge)#igmp snooping

[local]Router(config-igmp-snooping)#robust 4
```

1.133 route-map (BGP)

route-map *map-name* {**in** | **out**}

no route-map *map-name* {**in** | **out**}

1.133.1 Purpose

Applies a route map that modifies Border Gateway Protocol (BGP) attributes or filters BGP routes received from or sent to the BGP neighbor or peer group.

1.133.2 Command Mode

- BGP neighbor address family configuration
- BGP peer group address family configuration

1.133.3 Syntax Description

map-name	Name of the route map.
in	Applies the route map to incoming BGP routes sent from the BGP neighbor.
out	Applies the route map to outgoing BGP routes sent to the BGP neighbor.

1.133.4 Default

A route map is not applied to a BGP neighbor.

1.133.5 Usage Guidelines

Use the **route-map** command to apply a route map that modifies BGP attributes or to filter BGP routes sent to or received from the BGP neighbor or peer group. Use the **in** keyword to modify attributes or filter incoming routes received from the neighbor or peer group. Use the **out** keyword to modify attributes or filter outgoing routes sent to the neighbor.

Use the **route-map** command in context configuration mode to determine the attribute modifications and filtering conditions of the applied route map.

Currently, route map changes automatically take effect, and issuing the **clear bgp neighbor** *ip-addr* **soft** [**in** | **out**] command in exec mode to update a route map can cause updates to be unnecessarily sent; therefore, it is not recommended.

To aggregate multiple policy changes, such as the route map, the operating system performs the automatic update 15 seconds after any routing policy has changed.

Note: If the remote peer does not support the BGP route refresh capability, an inbound policy change for the peer will result in an automatic hard reset of the session.

Use the **no** form of this command to remove a route map.

1.133.6 Examples

The following example shows how to deny unicast BGP routes `10.0.0.0/8` (and more-specific routes) sent from the unicast BGP neighbor at IP address `102.210.210.1`. All other routes to this neighbor have the community attribute set to `100:14499`. Only multicast BGP routes `204.16.16.0/24` are sent to the multicast BGP neighbor at IP address `68.68.68.68`:

```
[local]Redback(config-ctx)#route-map rmap-20 deny 10
[local]Redback(config-route-map)#match ip address prefix-list prefix-deny-10
[local]Redback(config-route-map)#exit
[local]Redback(config-ctx)#route-map rmap-20 permit 20
[local]Redback(config-route-map)#set community 100:14499
[local]Redback(config-route-map)#exit
[local]Redback(config-ctx)#route-map rmap-30 permit 10
[local]Redback(config-route-map)#match ip address prefix-list prefix-permit-300
[local]Redback(config-route-map)#exit
[local]Redback(config-ctx)#ip prefix-list prefix-deny-10
[local]Redback(config-prefix-list)#permit 10.0.0.0/8 le 32
[local]Redback(config-prefix-list)#exit
[local]Redback(config-ctx)#ip prefix-list prefix-permit-300
[local]Redback(config-prefix-list)#permit 204.16.16.0/24
[local]Redback(config-prefix-list)#exit
.
.
.
[local]Redback(config-ctx)#router bgp 100
[local]Redback(config-bgp)#neighbor 102.210.210.1 external
[local]Redback(config-bgp-neighbor)#remote-as 200
[local]Redback(config-bgp-neighbor)#address-family ipv4 unicast
[local]Redback(config-bgp-peer-af)#route-map rmap-200 in
[local]Redback(config-bgp-peer-af)#exit
[local]Redback(config-bgp-neighbor)#exit
[local]Redback(config-bgp)#neighbor 68.68.68.68 external
[local]Redback(config-bgp-neighbor)#remote-as 300
[local]Redback(config-bgp-neighbor)#send community
[local]Redback(config-bgp-neighbor)#address-family ipv4 multicast
[local]Redback(config-bgp-peer-af)#route-map rmap-300 out
```

1.134 route-map (routing policies)

route-map *map-name* [*seq-num*] [**deny** *seq-num* | **permit** *seq-num*] | [**description** *text*]

no route-map *map-name* [*seq-num*] [deny seq-num | **permit** *seq-num*] | [**description**]

1.134.1 Purpose

Creates a route map for policy routing and enters route map configuration mode.

1.134.2 Command Mode

context configuration

1.134.3 Syntax Description

map-name	Descriptive name for the route map.
seq-num	Optional. Sequence number for the route map entry, relative to other route map entries in the same route map. Route map entries are tested in order of ascending sequence number; that is, the route map entry with the lowest sequence number is examined first when Border Gateway Protocol (BGP) routes are tested. The range of values is 1 to 4294967295; the default value is 10 greater than the largest sequence number of any route map entry in the route map.
deny ***seq-num***	Optional. Sequence number for the route map entry. The range of values is 1 to 4294967295. Routes using the specified sequence number are denied.
permit ***seq-num***	Optional. Sequence number for the route map entry. The range of values is 1 to 4294967295. Routes using the specified sequence number are permitted.
description ***text***	Optional. Description of the route map. No ***text*** argument is specified when the **description** keyword is used with the **no** form of this command.

1.134.4 Default

The action is permit. If not specified, the sequence number is 10 greater than the largest sequence number for a route map entry with the same ***map-name*** argument.

1.134.5 Usage Guidelines

Use the **route-map** command to create a route map for policy routing and enter route map configuration mode. Use this command in conjunction with the **match** commands in route map configuration mode to specify the conditions under which a route is accepted or rejected by the routing application that is using the route map. If the route entry indicates permit, the **set** commands can be used to modify the accepted routes attributes. Additionally, the **continue** clause can be used to modify execution flow in the route map.

Route map entries are tested in ascending order. For a route to match a particular route map entry, all match conditions must be satisfied. A route map entry with no match conditions can be used to unconditionally change a route's attributes by applying set actions.

Note: A reference to a route map that does not exist, or does not contain any configured entries, implicitly matches and permits all routes.

Use the **no** form of this command to delete a specific route map entry or to delete the entire route map. Because there can be only one description for a route map, when you use the **no** form of this command to delete the route map description, it is not necessary to include the ***text*** argument.

1.134.6 Examples

The following example shows how to redistribute static routes with destination addresses that match the IP access list `acc03` into the BGP routing process. The **set** command is used to modify the metric of selected routes:

```
[local]Redback(config-ctx)#ip prefix-list acc03
[local]Redback(config-prefix-list)#permit 81.1.0.0/16 le 32
[local]Redback(config-prefix-list)#permit 77.0.0.0/8 le 32
[local]Redback(config-prefix-list)#exit
[local]Redback(config-ctx)#route-map rmap1 permit 10
[local]Redback(config-route-map)#match ip address prefix-list acc03
[local]Redback(config-route-map)#set metric 10
[local]Redback(config-route-map)#exit
[local]Redback(config-ctx)#router bgp 65012
[local]Redback(config-bgp)#address-family ipv4 unicast
[local]Redback(config-addrfamily)#redistribute static route-map rmap1
```

1.135 route-origin

route-origin ***ext-com***

no route-origin

1.135.1 Purpose

Identifies the specific site from where a route has originated.

1.135.2 Command Mode

BGP address family configuration

1.135.3 Syntax Description

ext-com	Site of origin extended community value used to uniquely identify a site within internally connected multiple Virtual Private Network (VPN) sites. The site of origin extended community value can be expressed in either of the following formats: • ***asn*:*nnnn***, where ***asn*** is the autonomous system number, ***nnnn*** is either a 32-bit integer or a 16-bit integer, depending on the size of the ASN. You can specify the ASN as either a two-byte (two-octet) or four-byte (four-octet) integer. A value of 65535 or lower is interpreted as a two-byte integer, unless you add an **L** suffix (for example, **125L**), in which case it is interpreted as a four-byte integer. A value larger than 65535 is always interpreted as a four-byte integer, and the **L** suffix is optional. If the ASN is two-bytes, then ***nnnn*** is a 32-bit integer. If the ASN is four-bytes, then ***nnnn*** is a 16-bit integer. • ***ip-addr*:*nn***, where ***ip-addr*** is the IP address in the form ***A.B.C.D*** and ***nn*** is a 16-bit integer.

1.135.4 Default

No site of origin is specified.

1.135.5 Usage Guidelines

Use the **route-origin** command identify the specific site from where a route has originated.

When routes are received by a provider edge (PE) router, the route's route-origin attribute is checked against the route origin associated with the VPN for the receive site. Received routes are rejected if the route origin values are the same. This prevents the readvertisement of routes back to their originating sites.

Note: The **route-origin** command is useful only when Border Gateway Protocol (BGP) is used for PE-to-customer edge (CE) routing.

Use the **no** form of this command to remove the route-origin attribute from a route.

1.135.6 Examples

The following example shows how to configure routes originating from context `foo` to carry route origin `100:300` as part of the extended community attribute when they are advertised to other PE routers:

```
[local]Redback(config)#context foo vpn-rd 10.11.12.13:100
[local]Redback(config-ctx)#router bgp vpn
[local]Redback(config-bgp)#address-family ipv4 unicast
[local]Redback(config-bgp-af)#route-origin 100:300
[local]Redback(config-bgp-af)#export route-target 10.11.12.13:100
[local]Redback(config-bgp-af)#import route-target 100:100 10.11.12.13:100
```

1.136 router ancp

router ancp

no router ancp

1.136.1 Purpose

Creates the Access Node Control Protocol (ANCP) router and accesses ANCP configuration mode.

1.136.2 Command Mode

context configuration

1.136.3 Syntax Description

This command has no keywords or arguments.

1.136.4 Default

The ANCP router does not exist.

1.136.5 Usage Guidelines

Use the **router ancp** command to create the ANCP router and access ANCP configuration mode. The ANCP router is always created in the `local` context.

Use the **no** form of this command to delete the ANCP router and close all ANCP sessions; however, digital subscriber line (DSL) information learned from the sessions is not removed.

1.136.6 Examples

The following example creates the ANCP router in the `local` context and accesses ANCP configuration mode:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#router ancp
[local]Redback(config-ancp)#
```

1.137 router bfd

router bfd

no router bfd

1.137.1 Purpose

Creates a Bidirectional Forwarding Detection (BFD) instance and enters BFD router configuration mode.

1.137.2 Command Mode

context configuration

1.137.3 Syntax Description

This command has no keywords or arguments.

1.137.4 Default

No BFD instances are configured.

1.137.5 Usage Guidelines

Use the **router bfd** command to create a BFD instance and enter BFD router configuration mode.

Use the **no** form of this command to disable the BFD instance.

1.137.6 Examples

The following example shows how to create a BFD instance on the context, `local`, and enters BFD router configuration mode:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#router bfd
[local]Redback(config-bfd)#
```

1.138 router bgp

router bgp {*asn* | *nn*:*nn*}

no router bgp {*asn* | *nn*:*nn*}

1.138.1 Purpose

Configures a Border Gateway Protocol (BGP) routing instance using an autonomous system number (ASN) and enters BGP router configuration mode.

1.138.2 Command Mode

context configuration

1.138.3 Syntax Description

asn	ASN in integer format. The range of values is 1 to 65535. The subrange of 64512 to 65535 is reserved for private ASNs.
nn:nn	ASN in 4-byte integer format, where the first ***nn*** indicates the two higher-order bytes and the second ***nn*** denotes the two lower-order bytes.

1.138.4 Default

BPG routing is not enabled.

1.138.5 Usage Guidelines

Use the **router bgp** command to configure a BGP routing instance using an ASN, and to enter BGP configuration mode.

Use the **no** form of this command to disable the BGP routing instance.

1.138.6 Examples

The following example shows how to enable BGP routing for ASN `321` and enter BGP router configuration mode:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#router bgp 321
[local]Redback(config-bgp)#
```

1.139 router bgp vpn

router bgp vpn

1.139.1 Purpose

Configures a Border Gateway Protocol (BGP) routing instance in a Virtual Private Network (VPN) context and enters BGP configuration mode.

1.139.2 Command Mode

context configuration

1.139.3 Syntax Description

This command has no keywords or arguments.

1.139.4 Default

None

1.139.5 Usage Guidelines

Use the **`router bgp vpn`** command to configure a BGP routing instance in a VPN context, and enter BGP configuration mode. A BGP instance is always required within a VPN context for the following reasons:

- Customer routes must be distributed into BGP so they can be advertised across the internal BGP (iBGP) sessions that connect provider edge (PE) routers. Customer routes can be distributed into BGP either statically or from other active routing protocols.
- Route targets must also be configured within BGP address family configuration mode.

BGP does not function properly in a VPN context until it is first configured in the local context. Even though an autonomous system number (ASN) is not used when configuring a BGP instance in a VPN context, this instance uses the ASN from the BGP instance in the local context for peering with customer edge (CE) routers.

When configuring BGP peering sessions within a VPN context, only external neighbor sessions can be configured, because peering in a VPN context must only be configured with CE routers. Furthermore, the only permitted address family is IP Version 4 (IPv4) unicast, and peer groups cannot be configured.

1.139.6 Examples

The following example shows how to configure a BGP routing instance within a VPN context, and redistribute static routes from a customer into BGP:

```
[local]Redback(config)#context vpncontext vpn-rd 701:3
[local]Redback(config-ctx)#router bgp vpn
[local]Redback(config-bgp)#address-family ipv4 unicast
[local]Redback(config-bgp-af)#redistribute static
```

The following example shows how to configure a BGP peering session with a CE router:

```
[local]Redback(config)#context vpncontext vpn-rd 701:3
[local]Redback(config-ctx)#router bgp vpn
[local]Redback(config-bgp)#neighbor 205.1.2.2 external
[local]Redback(config-bgp-neighbor)#remote-as 100
[local]Redback(config-bgp-neighbor)#address-family ipv4 unicast
```

1.140 router-dead-interval

router-dead-interval *interval*

{**no** | **default**} **router-dead-interval**

1.140.1 Purpose

Modifies the amount of time the Open Shortest Path First (OSPF) or OSPF Version 3 (OSPFv3) process waits to receive a Hello packet from a neighbor before determining that the neighbor is not operational.

1.140.2 Command Mode

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

1.140.3 Syntax Description

interval	Amount of time, in seconds, that the OSPF or OSPFv3 process waits to receive a Hello packet. The range of values is 1 to 65,535. The value must be the same for all routers on a common network.

1.140.4 Default

The interval is 40 seconds for broadcast and point-to-point (P2P) networks, and 120 seconds for point-to-multipoint (P2MP) and nonbroadcast multiaccess (NBMA) networks.

1.140.5 Usage Guidelines

Use the **router-dead-interval** command to modify the amount of time the OSPF or OSPFv3 process waits to receive a Hello packet from a neighbor before determining that the neighbor is not operational. The OSPF router dead interval can be configured on a specific interface, sham link, or virtual link

If a Hello packet is not received within the configured amount of time, the OSPF or OSPFv3 process modifies its topology database to indicate that the neighbor is not operational.

The OSPF router dead interval value must be the same for all routers on a common network. The value must be greater than that of the Hello interval to avoid destroying adjacencies when the neighbor router is operational.

The following restrictions apply to the **router-dead-interval** command:

- After the **fast-hello** command is configured, you cannot use the **router-dead-interval** command until the **fast-hello** command has been disabled.
- After the **router-dead-interval** command has been configured, you cannot use the **fast-hello** command until the **router-dead-interval** command has been disabled.

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

1.140.6 Examples

The following example shows how to configure an OSPF interface to wait 60 seconds without receiving a Hello packet from its neighbor before determining that the neighbor is not operational:

```
[local]Redback(config-ospf-if)#router-dead-interval 60
```

1.141 route-reflector-client

route-reflector-client

no route-reflector-client

1.141.1 Purpose

Configures the internal Border Gateway Protocol (iBGP) neighbor (or peer group) as a route reflector client for the BGP address family.

1.141.2 Command Mode

- BGP neighbor address family configuration
- BGP peer group address family configuration

1.141.3 Syntax Description

This command has no keywords or arguments.

1.141.4 Default

The neighbor is not configured as a route reflector client.

1.141.5 Usage Guidelines

Use the **route-reflector-client** command to configure the iBGP neighbor (or peer group) for the specified address family as a route reflector client. No other configuration is required for an iBGP neighbor to act as a route reflector client.

Together, a route reflector and its clients form a cluster. If there is more than one route reflector in a cluster, all route reflectors in that cluster should be configured with the same ID through the **cluster-id** command. If there is no cluster ID, the router ID is used.

Note: This command cannot be enabled on a BGP neighbor that is part of a peer group because this feature cannot be customized for individual members inside of a peer group.

Use the **no** form of this command to remove the route reflector client specification from the iBGP neighbor.

1.141.6 Examples

The following example shows how to configure the iBGP neighbor at IP address, `102.210.210.1`, as a route reflector client for the unicast address family:

```
[local]Redback(config-ctx)#router bgp 100
[local]Redback(config-bgp)#neighbor 102.210.210.1 internal
[local]Redback(config-bgp-neighbor)#remote-as 100
[local]Redback(config-bgp-neighbor)#address-family ipv4 unicast
[local]Redback(config-bgp-peer-af)#route-reflector-client
```

1.142 router-id (BGP)

router-id ***ip-addr***

no router-id ***ip-addr***

1.142.1 Purpose

Configures a fixed Border Gateway Protocol (BGP) router ID for the SmartEdge router.

1.142.2 Command Mode

BGP router configuration

1.142.3 Syntax Description

ip-addr	IP address of the SmartEdge router.

1.142.4 Default

The router ID is the IP address of a loopback interface, if one is configured. If a loopback interface is not configured, the interface with the highest IP address is used as the router ID.

1.142.5 Usage Guidelines

Use the **router-id** command to configure a fixed BGP router ID for the SmartEdge router.

Caution!

Risk of dropped connection. When you change a router ID, any active peering sessions using the current router ID are dropped. To reduce the risk, avoid changing the router ID when peering sessions are actively running.

Note: If a context does not contain any IPv4 address configuration and BGP is being used, you must configure the **router-id** command in the context or routing protocol instance level. If you configure a context with only IPv6 addresses and no IPv4 addresses and run BGP in that context, BGP does not establish a relationship with any neighbors if the **router-id** command is not configured.

Use the **no** form of this command to remove the fixed router ID.

1.142.6 Examples

The following example shows how to configure a fixed BGP router ID of `10.10.1.1`:

```
[local]Redback(config-ctx)#router bgp 64001
[local]Redback(config-bgp)#router-id 10.1.1.1
```

1.143 router-id (contexts)

router-id ***ip-addr***

no router-id

1.143.1 Purpose

Configures a global router ID for the SmartEdge router.

1.143.2 Command Mode

context configuration

1.143.3 Syntax Description

ip-addr	IP address of the interface to be used as the router ID.

1.143.4 Default

A global router ID is not preconfigured.

1.143.5 Usage Guidelines

Use the **router-id** command to configure a global router ID for the SmartEdge router.

The global router ID in context configuration mode provides a consistent router ID for use by all routing protocols; however, if the router ID is configured as part of an individual routing protocol, such as the Open Shortest Path First (OSPF) protocol or the Border Gateway Protocol (BGP), it will take precedence over the global router ID in context configuration mode.

Note: The global router ID must be configured for the Resource Reservation Protocol (RSVP) to operate correctly.

Note: If a context does not contain any IPv4 address configuration and BGP is being used, you must configure the **router-id** command in the context or routing protocol instance level. If you configure a context with only IPv6 addresses and no IPv4 addresses and run BGP in that context, BGP does not establish a relationship with any neighbors if the **router-id** command is not configured.

Use the **no** form of this command to remove a global router ID.

1.143.6 Examples

The following example shows how to configure the IP address, `193.25.105.83`, as the global router ID in context configuration mode:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#router-id 193.25.105.83
```

1.144 router-id (LDP)

router-id ***ip-addr***

no router-id ***ip-addr***

1.144.1 Purpose

Configures the interface to be used as the Label Distribution Protocol (LDP) router ID.

1.144.2 Command Mode

LDP router configuration

1.144.3 Syntax Description

ip-addr	IP address in the form **A.B.C.D**.

1.144.4 Default

By default, the SmartEdge router determines the LDP router ID in the following sequence:

1 If a fixed LDP router ID configured through the **router-id** command in LDP configuration mode, it is used.

2 If a fixed LDP router ID is not configured, the configured loopback interface with the highest IP address is used as the LDP router ID.

3 If a loopback interface is not configured, the operational IS-IS or OSPF interface with the highest IP address is used as the LDP router ID.

1.144.5 Usage Guidelines

Use the **router-id** command to configure the interface to be used as the LDP router ID.

Caution!

Risk of traffic interruption. Because the router ID is used as the transport IP address for establishing a Transmission Control Protocol (TCP) connection, changing the router ID causes an active LDP session to be torn down, and then re-established. To reduce the risk, do not change the router ID when an LDP session is active.

Note: We recommend that you configure a loopback interface that is advertised by the Open Shortest Path First (OSPF) or Intermediate System-to-Intermediate System (IS-IS) routing instance to ensure that the LDP router ID is always reachable.

Use the **no** form of this command to return the system to its default behavior.

1.144.6 Examples

The following example shows how to configure the interface, `ldp-routerID`, as the LDP router ID:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#router isis isis-backbone
[local]Redback(config-isis)#net 49.2222.0010.0100.1001.00
[local]Redback(config-isis)#exit
[local]Redback(config-ctx)#interface ldp-routerID
[local]Redback(config-ctx)#ip address 10.1.1.1 255.255.255.0
[local]Redback(config-if)#isis router isis-backbone
[local]Redback(config-if)#exit
[local]Redback(config-ctx)#router ldp
[local]Redback(config-ldp)#router-id 10.1.1.1
```

1.145 router-id (OSPF)

router-id ***ip-addr***

no router-id

1.145.1 Purpose

Configures a fixed Open Shortest Path First (OSPF) or OSPF Version 3 (OSPFv3) router ID for the SmartEdge router.

1.145.2 Command Mode

- OSPF router configuration

- OSPF3 router configuration

1.145.3 Syntax Description

ip-addr	IP address of the interface to be used as the router ID.

1.145.4 Default

A router ID is not preconfigured.

1.145.5 Usage Guidelines

Use the **router-id** command to configure a fixed OSPF or OSPFv3 router ID for the SmartEdge router.

OSPF or OSPFv3 uses the router ID to identify the originating router for packets and link-state advertisements (LSAs). If the OSPF or OSPFv3 router ID is not configured, OSPF or OSPFv3 chooses the lowest loopback interface address. If there are no loopback interfaces, OSPF or OSPFv3 chooses the lowest interface address. The default OSPF or OSPFv3 router ID is selected when OSPF or OSPFv3 is started initially or restarted using the **process restart** command (in exec mode). For information on the **process restart** command, see the *Command List*.

Use the **no** form of this command to remove a router ID.

1.145.6 Examples

The following example shows how to configure the IP address, `193.25.105.83`, as the router ID:

```
[local]Redback(config-ospf)#router-id 193.25.105.83
```

1.146 router isis

router isis ***instance-name***

no router isis ***instance-name***

1.146.1 Purpose

Creates an Intermediate System-to-Intermediate System (IS-IS) instance and enters IS-IS router configuration mode.

1.146.2 Command Mode

context configuration

1.146.3 Syntax Description

instance-name	IS-IS instance name.

1.146.4 Default

No instance of IS-IS is configured.

1.146.5 Usage Guidelines

Use the **router isis** command to create an IS-IS instance and to enter IS-IS router configuration mode. To enable the IS-IS routing process, you must assign a network entity title (NET) to the instance. Use the **net** command in IS-IS router configuration mode.

A context can have multiple IS-IS instances. No more than one instance of IS-IS can operate on a single interface. To enable IS-IS on an interface, use the **interface**command in IS-IS router configuration mode.

Use the **no** form of this command to delete the IS-IS instance.

Caution!

Risk of IS-IS configuration settings loss. The **no router isis** command removes the IS-IS instance and all related configuration settings, which is different from deleting the last NET. Deleting the last NET disables the IS-IS instance while preserving all configuration information. To reduce the risk, delete the last NET.

1.146.6 Examples

The following example shows how to configure the `ip-backbone` IS-IS instance and assigns it a NET of `47.001.002.002.002.00`:

```
[local]Redback(config-ctx)#router isis ip-backbone
[local]Redback(config-isis)#net 47.0001.0002.0002.0002.00
```

1.147 router ldp

```
router ldp

no router ldp
```

1.147.1 Purpose

Enables a Label Distribution Protocol (LDP) routing instance for a context and enters LDP router configuration mode.

1.147.2 Command Mode

context configuration

1.147.3 Syntax Description

This command has no keywords or arguments.

1.147.4 Default

LDP routing is disabled.

1.147.5 Usage Guidelines

Use the router ldp command to enable an LDP routing instance for context, and to enter LDP router configuration mode. Our implementation of LDP follows the LDP specification as described in RFC 3036, ***LDP Specification***.

For the context in which you configure LDP, you must also:

- Configure an Multiprotocol Label Switching (MPLS) routing instance.
- Enable MPLS on the interface on which you plan to enable LDP.

You may also need to enable an Interior Gateway Protocol (IGP), such as Open Shortest Path First (OSPF) or Intermediate System-to-Intermediate System (IS-IS), on the interface.

To ensure that the LDP router ID is always reachable, we recommend that you also configure a loopback interface that is advertised by the IGP, such as OSPF or IS-IS, routing instance.

Note: For the commands used to configure an IGP routing instance and interface, such as IS-IS or OSPF, see the *Command List*.

Note: LDP configuration is supported in the local context only.

Use the **no** form of this command to disable LDP routing for the context.

1.147.6 Examples

The following example shows how to enable an LDP routing instance for the `local` context and enter `LDP` router configuration mode:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#router ldp
[local]Redback(config-ldp)#
```

1.148 router mobile-ip

router mobile-ip

no router mobile-ip

1.148.1 Purpose

Enables Mobile IP services in this context and accesses Mobile IP configuration mode.

1.148.2 Command Mode

context configuration

1.148.3 Syntax Description

This command has no keywords or arguments.

1.148.4 Default

Mobile IP services are not enabled in any context.

1.148.5 Usage Guidelines

Use the **router mobile-ip** command to enable Mobile IP services in this context and access Mobile IP configuration mode.

Use the **no** form of this command to disable Mobile IP services in this context.

1.148.6 Examples

The following example shows how to enable Mobile IP services in the `fa` context:

```
[local]Redback(config)#context fa
[local]Redback(config-ctx)#router mobile-ip
[local]Redback(config-mip)#
```

1.149 router mpls

`router mpls`

`no router mpls`

1.149.1 Purpose

Enables Multiprotocol Label Switching (MPLS) routing within a context and enters MPLS router configuration mode.

1.149.2 Command Mode

context configuration

1.149.3 Syntax Description

This command has no keywords or arguments.

1.149.4 Default

MPLS routing is disabled.

1.149.5 Usage Guidelines

Use the **`router mpls`** command to enable MPLS routing within a context and enter MPLS router configuration mode.

Note: MPLS is supported in the local context only.

Use the **`no`** form of this command to disable MPLS routing.

1.149.6 Examples

The following example shows how to enable MPLS routing and enter MPLS router configuration mode:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#router mpls
[local]Redback(config-mpls)#
```

1.150 router mpls-static

router mpls-static

no router mpls-static

1.150.1 Purpose

Enables static Multiprotocol Label Switching (MPLS) routing within a context and enters MPLS static router configuration mode.

1.150.2 Command Mode

context configuration

1.150.3 Syntax Description

This command has no keywords or arguments.

1.150.4 Default

Static MPLS routing is disabled.

1.150.5 Usage Guidelines

Use the **router mpls-static** command to enable static MPLS routing within a context and enter MPLS static router configuration mode.

Note: MPLS is supported in the local context only.

Use the **no** form of this command to disable static MPLS routing.

1.150.6 Examples

The following example shows how to enable static MPLS routing and enter MPLS static router configuration mode:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#router mpls-static
[local]Redback(config-mpls-static)#
```

1.151 router msdp

router msdp

no router msdp

1.151.1 Purpose

Enables Multicast Source Discovery Protocol (MSDP) within a context and enters MSDP router configuration mode.

1.151.2 Command Mode

context configuration

1.151.3 Syntax Description

This command has no keywords or arguments.

1.151.4 Default

MSDP is disabled.

1.151.5 Usage Guidelines

Use the **router msdp** command to enable MSDP within a context and enter MSDP router configuration mode.

Use the **no** form of this command to disable MSDP within a context.

1.151.6 Examples

The following example shows how to enable MSDP and enters MSDP router configuration mode:

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

1.152 router nd

router nd

no router nd

1.152.1 Purpose

Creates or selects a Neighbor Discovery (ND) router and accesses ND router configuration mode.

1.152.2 Command Mode

context configuration

1.152.3 Syntax Description

This command has no keywords or arguments.

1.152.4 Default

No ND router is created.

1.152.5 Usage Guidelines

Use the **router nd** command to create or select an ND router and access ND router configuration mode. You can create a single ND router in each context.

Use the **no** form of this command to remove the ND router from the configuration; the **no** form also removes the ND-specific configuration from any interfaces in this context.

1.152.6 Examples

The following example shows how to create an ND router in the `local` context:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#router nd
```

1.153 router ospf

router ospf ***instance***

no router ospf ***instance***

1.153.1 Purpose

Configures an Open Shortest Path First (OSPF) routing instance and enters OSPF router configuration mode.

1.153.2 Command Mode

context configuration

1.153.3 Syntax Description

instance	Instance ID. The range of values is 1 to 65,535.

1.153.4 Default

OSPF routing is disabled.

1.153.5 Usage Guidelines

Use the **router ospf** command to configure an OSPF routing instance and to enter OSPF router configuration mode.

Use the **no** form of this command to disable OSPF routing.

1.153.6 Examples

The following example shows how to configure the OSPF instance, `105`, and enter OSPF router configuration mode:

```
[local]Redback(config-ctx)#router ospf 105
[local]Redback(config-ospf)#
```

1.154 router ospf3

router ospf3 ***instance-id***

no router ospf3 ***instance-id***

1.154.1 Purpose

Creates an Open Shortest Path First Version 3 (OSPFv3) routing instance and enters OSPF3 router configuration mode.

1.154.2 Command Mode

context configuration

1.154.3 Syntax Description

instance-id	Instance ID. The range of values is 1 to 65,535.

1.154.4 Default

OSPFv3 routing is disabled.

1.154.5 Usage Guidelines

Use the **router ospf3** command to create an OSPFv3 routing instance and to enter OSPF3 router configuration mode.

Use the **no** form of this command to disable OSPFv3 routing.

1.154.6 Examples

The following example shows how to configure the OSPFv3 instance, `105`, and enter OSPF3 router configuration mode.

```
[local]Redback(config-ctx)#router ospf3 105
[local]Redback(config-ospf3)#
```

1.155 router-priority

router-priority ***priority***

default router-priority

1.155.1 Purpose

Modifies the Open Shortest Path First (OSPF) or OSPF Version 3 (OSPFv3) preference for the SmartEdge router to act as the designated router on a network.

1.155.2 Command Mode

- OSPF interface configuration
- OSPF3 interface configuration

1.155.3 Syntax Description

priority	Priority setting. The range of values is 0 to 255; the default value is 1.

1.155.4 Default

The priority value is 1.

1.155.5 Usage Guidelines

Use the **router-priority** command to modify the OSPF or OSPFv3 preference for the SmartEdge router to act as the designated router on a network.

Enter any value greater than or equal to 1 to indicate that the SmartEdge router can act as the designated router. The router with the highest priority is used as the designated router for the network if there is not a designated router already on the network. If two routers have the same priority value, the router with the higher router ID is the designated router for the network; see the **router-id** command.

A value of 0 causes the router to never act as the designated router.

Use the **default** form of this command to return the priority to the default value of 1.

1.155.6 Examples

The following example shows how to set the router priority to 2:

```
[local]Redback(config-ospf-if)#router-priority 2
```

1.156 router rip

router rip ***instance***

no router rip ***instance***

1.156.1 Purpose

Creates an instance of the Routing Information Protocol (RIP) routing process and enters RIP router configuration mode.

1.156.2 Command Mode

context configuration

1.156.3 Syntax Description

instance	RIP instance name.

1.156.4 Default

The RIP routing process is disabled.

1.156.5 Usage Guidelines

Use the **router rip** command to creates an instance of the RIP routing process and to enter RIP router configuration mode. Each RIP instance has its own routing table. You can configure multiple RIP instances

To configure a RIP instance on an interface, use the **rip router**, **rip listen**, or **rip supply** command in interface configuration mode.

Use the **no** form of this command to disable an instance of the RIP routing process.

1.156.6 Examples

The following example shows how to enable the RIP instance, `rip001`, and enter RIP router configuration mode:

```
[local]Redback(config-ctx)#router rip rip001
[local]Redback(config-rip)#
```

1.157 router ripng

router ripng *instance-id*

no router ripng *instance-id*

1.157.1 Purpose

Creates an instance of the Routing Information Protocol next generation (RIPng) routing process and enters RIPng router configuration mode.

1.157.2 Command Mode

context configuration

1.157.3 Syntax Description

instance-id	RIPng instance ID.

1.157.4 Default

The RIPng routing process is disabled.

1.157.5 Usage Guidelines

Use the **router ripng** command to create an instance of the RIPng routing process and to enter RIPng router configuration mode. Each RIPng instance has its own routing table. You can configure multiple RIPng instances.

Use the **no** form of this command to disable an instance of the RIPng routing process.

1.157.6 Examples

The following example shows how to enable the RIPng instance, `ripng001`, and enter RIPng router configuration mode:

```
[local]Redback(config-ctx)#router ripng ripng001
[local]Redback(config-ripng)#
```

1.158 router rsvp

router rsvp

no router rsvp

1.158.1 Purpose

Enables Resource Reservation Protocol (RSVP) routing within a context and enters RSVP router configuration mode.

1.158.2 Command Mode

context configuration

1.158.3 Syntax Description

This command has no keywords or arguments.

1.158.4 Default

RSVP is disabled.

1.158.5 Usage Guidelines

Use the **router rsvp** command to enable RSVP routing within a context and enter RSVP router configuration mode.

Use the **no** form of this command to disable RSVP routing within a context.

1.158.6 Examples

The following example shows how to enable RSVP routing and enter RSVP router configuration mode:

```
[local]Redback(config)#context isp35
[local]Redback(config-ctx)#router rsvp
[local]Redback(config-rsvp)#
```

1.159 route-target filter

```
route-target filter

no route-target filter
```

1.159.1 Purpose

Enables automatic Border Gateway Protocol (BGP) route target community filtering.

1.159.2 Command Mode

BGP address family configuration

1.159.3 Syntax Description

This command has no keywords or arguments.

1.159.4 Default

Denies all incoming IP Version 4 (IPv4) Virtual Private Network (VPN) routes that are not imported into any VPN context, if the local router is not configured as a route reflector.

1.159.5 Usage Guidelines

Use the **route-target filter** command to enable automatic BGP route target community filtering. This command configures the local router, if it is not configured as a route reflector, to ignore all VPN routes received that are not imported into any VPN context.

Note: For BGP route target filtering to work properly, you must first use the **address-family ipv4 vpn** command to specify the use of VPN-IPv4 prefixes for the BGP instance.

You can control the number of IPv4 VPN routes that the local autonomous system border router (ASBR) advertise to the remote ASBR by configuring a community for exportable routes on the inbound interface of the provider edge (PE) router, and configuring a community based filter on the outbound interface of the local ASBR to advertise only routes that match the community.

Use the **no** form of this command to allow the local router to accept all BGP IPv4 VPN routes. Accepting all IPv4 VPN routes is the desired behavior for a router configured as an ASBR for inter-autonomous system (AS) VPNs.

1.159.6 Examples

The following example shows how to configure a local router to accept all received IPv4 VPN routes:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#router bgp 100
[local]Redback(config-bgp)#address-family ipv4 vpn
[local]Redback(config-bgp-af)#no route-target filter
```

1.160 rpf-interface

rpf-interface ***interface1 interface2***

no rpf-interface

1.160.1 Purpose

Identifies both the active and backup Reverse Path Forwarding (RPF) interfaces that establish sessions over which the source device sends multicast join requests.

1.160.2 Command Mode

pim dual join configuration mode

1.160.3 Syntax Description

interface1	Active RPF interface that establishes a session over which the source device sends multicast join requests.
interface2	Backup RPF interface that establishes a session over which the source device sends multicast join requests.

1.160.4 Default

None

1.160.5 Usage Guidelines

Use the **rpf-interface** command to identify both the active and backup RPF interfaces that establish sessions over which the source device sends multicast join requests.

Use the **no** form of this command to set the active and backup RPF interfaces to the previously set values.

1.160.6 Examples

The following example shows how to set the RPF interface `int1` as the active link and RPF interface `int2` as the backup link for multicast join request sessions:

```
[local]Redback(config)#context local
[local]Redback(config-ctx)#pim dual-join group 225.100.1.1 source 192.110.30.6
[local]Redback(config-pim-dual-join)#rpf-interface int1 int2
```

1.161 rro-prefix-type

rro-prefix-type {**router-id** | **interface**}

no rro-prefix-type {**router-id** | **interface**}

1.161.1 Purpose

Configures the Resource Reservation Protocol (RSVP) record route object (RRO) IP prefix type.

1.161.2 Command Mode

RSVP router configuration

1.161.3 Syntax Description

router-id	Uses the router ID as the IP prefix when sending an RRO.
interface	Uses the outbound interface IP address when sending an RRO.

1.161.4 Default

The router ID is used as the IP prefix type when sending an RRO.

1.161.5 Usage Guidelines

Use the **rro-prefix-type** command to configure the RSVP RRO IP prefix type. You can change the IP prefix inside an RRO to be either the router ID or the interface IP address. This can be used for Multiprotocol Label Switching (MPLS) fast reroute for node protection and interarea node protection. During MPLS fast reroute, the point of local repair (PLR) router needs to match the bypass label-switched path (LSP) egress address with the IP prefix inside the RRO of the next-next-hop node.

1.161.6 Examples

The following example shows how to configure the RSVP RRO to use the outbound interface IP address when sending an RRO:

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

[local]Redback(config-rsvp)#rro-prefix-type interface
```