Configuring MPLS

SYSTEM ADMINISTRATOR GUIDE

46/1543-CRA 119 1170/1-V1 Uen F

Copyright

Disclaimer

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

Trademark List

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

Contents

1 Overview

This document provides an overview of Multiprotocol Label Switching (MPLS) traffic engineering (TE) and describes the tasks and commands used to configure, monitor, troubleshoot, and administer MPLS features through the SmartEdge router.

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.

The following sections provide an overview of MPLS-TE, RSVP, and MPLS features supported by the SmartEdge router.

1.1 MPLS Architecture

The SmartEdge router supports Multiprotocol Label Switching (MPLS), which is a method for efficiently forwarding packets through a network. MPLS operates across an interface in a local context where MPLS is enabled.

In a conventional network, routers forward packets through the network, from one router to the next, with each router making an independent forwarding decision by analyzing the packet header. This conventional approach to forwarding packets has become insufficient to support current networking demands.

With MPLS, the complete analysis of the packet header is performed only once, when it enters an MPLS-enabled network. At each incoming (ingress) point of the network, packets are assigned a label by an edge label-switched router (LSR). Packets are forwarded along a label-switched path (LSP) where each LSR makes forwarding decisions based on the label information. At each hop, the LSR swaps the existing label for a new label that tells the next hop how to forward the packet. At the outgoing (egress) point, an edge LSR removes the label, and forwards the packet to its destination. MPLS uses Resource Reservation Protocol (RSVP) or the Label Distribution Protocol (LDP) to communicate labels and their meanings among LSRs.

An LSP is a specific traffic path through an MPLS-enabled network, and can be signaled or static. RSVP LSPs are dynamic. You specify the ingress LSR and the egress LSR, but the next hops through the network are determined using RSVP or LDP, which assign labels in LSRs based on information from existing routing protocols. However, you can also use the `source-path` command (in RSVP LSP configuration mode) to assign an explicit route (a list of specific hops through a network) to an RSVP LSP. RSVP LSPs can usually change according to changes in network conditions, but an RSVP LSP with an assigned source path fails if changing network conditions make it topologically impossible. With static LSPs, you manually specify the ingress LSR, all next-hop LSRs, and the egress LSR. It cannot change with changes in network conditions. Figure 1 shows a static LSP through a simple MPLS-enabled network. A packet enters the network at the ingress LSR A, is forwarded to the next-hop LSRs C and D, and exits the network through the egress LSR E.

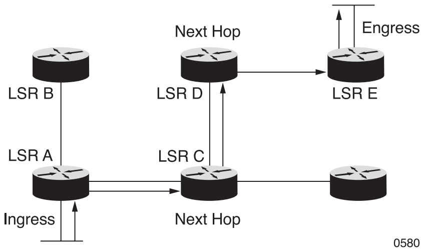

Figure 1 Static LSP in a Simple MPLS-Enabled Network

The SmartEdge router supports the ability to configure a standby RSVP LSP as a backup to a primary LSP, or to another standby (backup) LSP in what is known as a backup-to-backup LSP configuration. Backup paths and backup-to-backup paths can be generated by using Constrained Shortest Path First (CSPF) path calculation. For more information about CSPF, see CSPF.

Backup and backup-to-backup RSVP LSP tunnels are configured on a one-to-one, end-to-end path backup basis. Failover occurs in a hierarchy as follows:

- If the primary LSP fails, traffic fails over to the backup LSP.
- If the backup LSP also fails, traffic fails over to the backup-to-backup LSP.

In both cases, failover is revertive; in other words, if the higher-order LSP becomes active again, traffic reverts to the higher-order LSP.

To enable MPLS forwarding, you must enable an interface for MPLS by creating an MPLS instance, and adding an interface to it. To enable RSVP signaling, you must enable one or more interfaces for RSVP by creating an RSVP instance and adding an interface to it. For static LSPs, RSVP does not need to be enabled; however, for RSVP LSPs, both RSVP and MPLS must be enabled. If MPLS is not properly enabled on the correct interface, RSVP does not come up and MPLS forwarding does not work.

Note: MPLS is supported in the local context only.

For information about troubleshooting MPLS, see *Troubleshooting MPLS*.

1.2 MPLS QoS

The quality of service (QoS) feature for the SmartEdge router uses the Differentiated Services Code Point (DSCP) value to classify and mark ingress IP packets. At each transit node, the DSCP value is used to select the per-hop behavior (PHB) that determines the scheduling treatment and, in some cases, drop probability for each packet.

QoS DSCP can also be used over MPLS networks by copying the three most significant DSCP bits into the EXP field of MPLS labels at label imposition time.

The default MPLS QoS behavior adheres to the following rules:

- If there are two labels (tunnel and VPN labels) then the DSCP bits are copied into the EXP field of both labels. If penultimate hop popping is enabled, the tunnel label is taken off at the penultimate hop. The egress router will then use the VPN label EXP bits for egress queueing decisions. If there is no VPN label, then the egress router uses the DSCP value.
- If access control list (ACL)-based QoS or policing is used to change the DSCP at the ingress router, then bits 0–2 of this new value must be copied into the EXP field.
- The DSCP value is never changed after the ingress router, even if the EXP value in the tunnel or VPN label is changed.

The SmartEdge router provides commands that allow you to change the default MPLS QoS behavior to accommodate situations, such as VPN configurations, where you may want to change the way the DSCP bits are handled.

For information about configuring MPLS QoS, see *Configuring Circuits for QoS*.

1.3 MPLS TTL

The time-to-live (TTL) field in the IP packet header indicates how many hops a packet can travel before being dropped. The TTL value is decremented by one at each hop, until it reaches zero, and the packet is dropped; however, there needs to be a mechanism to ensure that the TTL field is decremented whenever a packet is labeled and forwarded through an MPLS LSP.

The default behavior for the SmartEdge router ensures that the TTL value is properly decremented by performing the following operations:

- At the ingress LSR, the IP TTL field is propagated to the MPLS TTL field located in the label header.
- The MPLS TTL field is decremented at each hop in the LSP.

- At the egress LSR, the MPLS TTL field replaces the IP TTL field, and the label is popped.

The SmartEdge router provides commands that allow you to change the default MPLS TTL behavior to accommodate situations, such as VPN configurations, where you may want to change the way the TTL field is handled.

1.4 MPLS over MLPPP

For static links (that is, for nonsubscriber links), the SmartEdge router MPLS implementation recognizes multilink Point-to-Point Protocol (MLPPP) constituents. The label manager (LM) installs and maintains a single adjacency for the bundle, and traffic is balanced by a round-robin algorithm across the constituents. This makes the aggregate bandwidth of the MLPPP bundle available to the traffic stream and helps conserve equal-cost multipath (ECMP) paths when the number of constituents increases.

1.5 Next-Hop Fast Reroute

Next-hop fast reroute (NFRR) is a feature that allows you to quickly reroute IP and MPLS traffic in the event of a link failure or a node failure. This is done by creating a bypass RSVP LSP for link protection or node protection.

A bypass LSP is no different from any other RSVP LSP, except that it does not carry traffic under normal conditions. When a link or node failure is detected, traffic is quickly rerouted onto a bypass RSVP to circumvent the failure. Traffic enters the headend router of a bypass RSVP, which is called the point of local repair (PLR), and exits the tail end router of the bypass RSVP LSP, which is called the merge point (MP). Any type of traffic intended to use the next hop can be switched onto the bypass LSP.

The SmartEdge router can apply a bypass LSP for RSVP-TE signaled LSPs in which both the protected LSP and bypass LSP can be over a LAG. The following sections provide information on the different types of NFFR.

1.5.1 NFRR for Link Protection

A bypass RSVP LSP for link protection reroutes traffic when a link failure is detected between an LSR and the next-hop LSR. Figure 2 shows an example where a bypass RSVP LSP has been created to protect against a link failure. The bypass RSVP LSP is created on LSR A, which is also the PLR, and when the IP address 20.20.20.2 is unreachable across LSP 1, the bypass RSVP LSP provides a path to reroute traffic to LSR B, which is also the MP. Traffic then continues across LSP 1 to LSR C and LSR D.

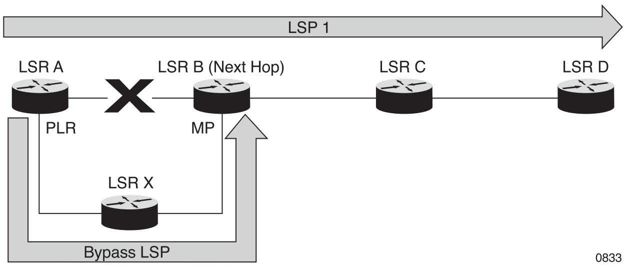

Figure 2 Next-Hop Fast Reroute for Link Protection

Note: When creating a bypass RSVP LSP for link protection you, must specify only the LSR to protect against.

1.5.2 NFRR for Node Protection

A bypass RSVP LSP for node protection reroutes traffic when a next-hop LSR failure is detected. Figure 3 shows an example where a bypass RSVP LSP has been created to protect against a node failure. The bypass RSVP LSP is created on LSR A, which is also the PLR, and when LSR B failure is detected, the bypass RSVP LSP provides a path to reroute traffic to LSR C, which is LSR A's next-next hop and the MP. Traffic then continues across LSP 1 to LSR D.

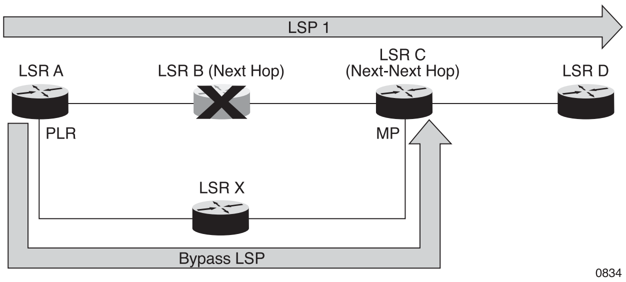

Figure 3 Next-Hop Fast Reroute for Node Protection

Note: When creating a bypass RSVP LSP for node protection, you must specify the LSR to protect against and the next-next-hop LSR.

1.5.3 NFRR for Link Aggregation

PPA2 cards support NFRR for RSVP-TE signaled LSPs in which both the protected LSP and bypass LSP can be over an Ethernet link group (network-facing).

Traffic protection is supported at the point of local repair (PLR) and is applied to the following traffic, where an RSVP-TE tunnel is the next hop:

- Static and BGP IPv4 routes
- IGP shortcuts
- IP/LDP Local Protection control and data traffic
- PW (L2VPN and VPLS) and L3VPN

The link group may contain up to eight constituent links, which may be over multiple line cards. Up to 50 LSPs with bypass are supported. Both link and node protection NFRR are supported, but they are not supported simultaneously for a given next hop.

This feature is supported on PPA2 cards only.

1.6 RSVP Fast-Reroute with Multiple Bypass LSPs

FRR supports the protection of an RSVP LSP with a bypass LSP. A bypass LSP is preestablished to protect an LSP that traverses either a specific link (link bypass LSP) or node (node bypass LSP). This feature allows for very fast path protection of an LSP if a failure occurs in its original path. In this release, the SmartEdge router supports a new function that chooses the best preferred bypass LSP when multiple candidate bypass LSPs protect the same address.

The selection of a preferred bypass LSP is hierarchical and based on the following criteria in the order shown:

1. A node bypass LSP is preferred over a link bypass LSP.

2. If multiple bypass LSP candidates still exist after criterion 1 is applied, a bypass LSP whose next hop is over a link-group is the most preferred bypass LSP.

3. If multiple LSP candidates still exist after criterion 2 is applied, a bypass LSP whose next hop is through a different line card than the protected LSP is the most preferred bypass LSP.

4. If multiple LSP candidates still exist after criterion 3 is applied, the first bypass LSP to become active is the most preferred bypass LSP.

During an XC reboot or when LSPs are initially configured, the first bypass LSP to become active is chosen as the preferred LSP. When subsequent bypass LSPs come up or the current preferred bypass LSP is brought down, the selection of the preferred bypass LSP is recalculated. This same calculation is also performed when a protected LSP becomes active.

If a new bypass LSP is chosen as the preferred bypass LSP, it replaces the old bypass LSP. This process happens in the following sequence:

1. A new node bypass LSP becomes active.

2. All active LSPs are checked by the SmartEdge OS.

   If an unprotected LSP that allows protection exists, it gets assigned to the new bypass LSP.

   If an LSP is already assigned to a bypass LSP, the selection of the preferred bypass LSP is recalculated. If the new bypass LSP is the preferred bypass LSP, the old bypass LSP is removed and the new preferred bypass LSP assumes control. If the new bypass LSP is not preferred, the LSP remains assigned to the original bypass LSP.

Only the preferred bypass LSP is communicated outside the RSVP process; therefore, other processes (such as the LM, RIB, or PPA) learn of only the preferred bypass LSP, even if other secondary (nonpreferred) bypass LSPs exist that protect the same address. The secondary bypass LSPs remain in a pool that is maintained by RSVP. These secondary bypass LSPs remain inactive until the preferred bypass LSP goes down, at which point a new preferred bypass LSP is automatically chosen from the pool. Only the preferred bypass LSPs that are being used to protect other LSPs maintain an active state during the restart process.

1.7 RSVP-TE

Resource Reservation Protocol (RSVP) traffic engineering (TE) is an extension to RVSP for establishing LSP paths in MPLS networks. RSVP-TE works with routing protocols to reserve resources across the network based on network constraint parameters, such as available bandwidth and the number of explicit hops. It allows you to allocate resources along the path. For information about RSVP-TE, see RFC 3209, *RSVP-TE: Extensions to RSVP for LSP Tunnels*.

1.8 RSVP Graceful Restart

RSVP graceful restart enables the SmartEdge router and its neighbors to continue forwarding packets without disrupting network traffic when a neighbor is down. When RSVP graceful restart is enabled, the SmartEdge router initiates graceful restart when a neighbor is down. During graceful restart, all RSVP LSPs between the SmartEdge router and the restarting neighbor are maintained. The SmartEdge router uses RSVP Hello messages to determine if a neighbor is down. Use the **`hello interval`** and **`hello keep-multiplier`** commands in RSVP interface configuration mode to enable and configure RSVP Hello messages.

Note: RSVP Hello messages are disabled by default. They must be enabled on the interface and the neighbor to which it connects for RSVP graceful restart to function properly.

Be aware that when RSVP graceful restart is enabled with the **`graceful-restart`** command, helper mode is also enabled by default. When RSVP graceful restart helper mode is enabled, the SmartEdge router assists the restarting neighbor by maintaining the LSPs between the SmartEdge router and the RSVP neighbors.

Use the **`no graceful-restart helper`** command to disable RSVP graceful restart helper mode on a system. When graceful restart helper mode is disabled, the SmartEdge router does not help the restarting neighbor.

Graceful restart continues until all prior adjacencies are established or the recovery timer (which is configured with the `graceful-restart` command) expires.

1.9 CSPF

The SmartEdge router supports the Constrained Shortest Path First (CSPF) algorithm. CSPF works with link-state routing protocols such as OSPF and IS-IS to automate the provisioning of label-switched paths (LSPs) in core networks and differentiate levels of quality of service (QoS) in Layer 2 and Layer 3 service applications. It calculates LSPs in an MPLS TE domain on a single path from the headend (ingress) router to the tailend (egress) router, based on a set of criteria called constraints. These constraints include bandwidth requirements for an LSP, the cost of the link to the destination, and other criteria. The head- end router uses CSPF to calculate the best path; other routers in the domain do not calculate CSPF. When you use CSPF, the ingress and egress routers must be part of the same MPLS TE domain.

When the constraints are not met, CSPF removes these links from its network topology and then runs the Shortest Path First (SPF) algorithm on the remaining LSP links. For example, if the bandwidth constraint is not met, all the TE LSP links that do not meet this constraint are removed from the network topology. CSPF then calculates the SPF on the remaining LSP links and generates a list of nodes called a path list. This path list of explicit routes provides the shortest path (that meets the constraints) through the network to the destination.

The CSPF calculation creates an ordered set of IP addresses mapping to the next-hop addresses in the TE LSP, and the Explicit Route Object (ERO), which defines the path that RSVP must take from the ingress router to the egress router. CSPF sends these results to RVSP; RSVP then uses this path to signal the LSP to reserve resources.

Figure 4 illustrates the LSP between the ingress and egress routers in an Interior Gateway Protocol (IGP) domain.

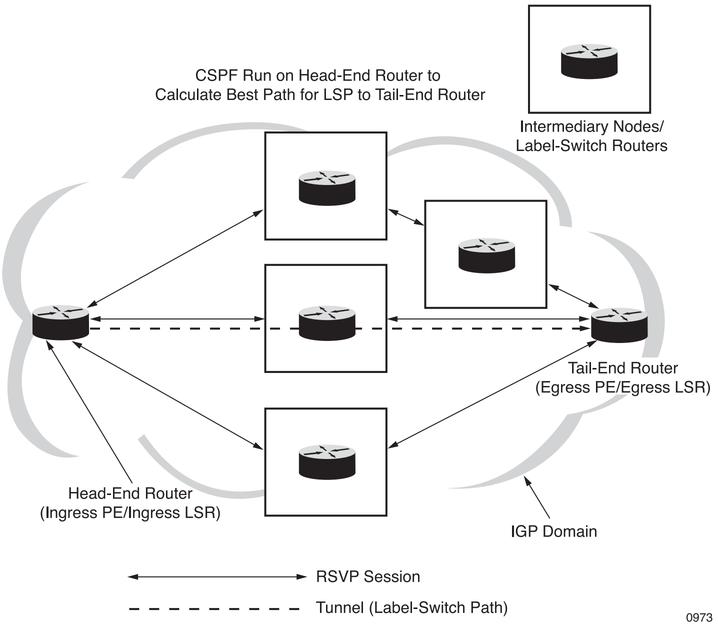

Figure 4   LSP in an IGP Domain

To calculate the best path to the destination, CSPF uses a traffic engineering database (TED), the cost of the link to the destination, and the following constraints on an LSP:

- Administrative group
- Exclude links
- Hop limit
- Minimum bandwidth
- Priority

The TED is a link-state database of TE-enabled links for the IGP routing domain. For information about these constraints, see Configuring CSPF Calculation.

Any LSP, including fast-reroute (FRR) bypass paths can be calculated by using CSPF.

Note: To use CSPF, either Intermediate System-to-Intermediate System (IS-IS) or the Open Shortest Path First (OSPF) protocol must be used with traffic engineering extensions enabled.

Before configuring CSPF on your router, you must use the `router-id` command in context configuration mode to configure a router ID.

The sections that follow provide more information about configuring CSPF.

1.9.1 CSPF Administrative Attributes

With CSPF, you configure administrative groups that are associated with an LSP. You typically use link colors as values when configuring administrative groups. Each value is associated with a specific class that you define. You can define up to 32 link attributes: 32 values (0 to 31). The path names and their corresponding values must be the same on all routers within a single MPLS TE domain.

Use the following commands to configure administrative groups:

- *admin-group*
- *attribute (RSVP)*

1.9.2 CSPF Tiebreakers

To determine tiebreakers between paths, CSPF chooses paths in the following order:

1. The path with greatest minimum available bandwidth.
2. The path with lowest hop count.
3. A randomly selected path (the top path on the path list).

1.9.3 Supported Standards for CSPF

The following standards are supported for CSPF:

- RFC3630—***Traffic Engineering Extensions to OSPF Version 2***
- RFC 2627—***QoS Routing Mechanisms and OSPF Extensions***
- RFC 2702—***Requirements for Traffic Engineering Over MPLS***
- draft-ietf-isis-traffic-02.txt—***IS-IS extensions for Traffic Engineering***

2 Configuration and Operations Tasks

Note: In this section, the command syntax in the task tables displays only the root command.

For information about how to troubleshoot MPLS, see the *Troubleshooting MPLS*.

To configure MPLS and MPLS-related features, perform the tasks described in the following sections.

2.1 Configuring MPLS

To configure MPLS, perform the tasks described in the following sections:

2.1.1 Create an MPLS Routing Instance

To create an MPLS routing instance, perform the tasks described in Table 1.

Table 1 Create an MPLS Routing Instance

Task	Root Command	Notes
Create an MPLS routing instance, and access MPLS router configuration mode.	*router mpls*	Enter this command in context configuration mode. Note that MPLS is supported in the local context only.
Enable MPLS on an interface.	*interface (MPLS and RSVP)*	Enter this command in MPLS router configuration mode.
Configure MPLS TTL.	For the complete list of tasks used to configure MPLS TTL, see Configure the MPLS TTL.	—

2.1.2 Configure the MPLS TTL

To configure the MPLS TTL, perform the tasks described in Table 2. Enter all commands in MPLS router configuration mode.

Table 2 Configure the MPLS TTL

Task	Root Command	Notes
Enable transit routers to decrement the MPLS TTL by 1 at each hop.	*decrement ttl*	The default behavior of the SmartEdge router is to decrement the MPLS TTL by 1 at each hop, so the **`decrement ttl`** command is used to return the router to its default behavior after it has been changed by the **`no`** form of this command.
Enable the propagation of the IP TTL to the MPLS tunnel label TTL at the ingress router.	*propagate ttl ip-to-mpls*	The default behavior of the SmartEdge router is to propagate of the MPLS tunnel label TTL to the IP TTL at the egress router, so the **`propagate ttl ip-to-mpls`** command is used to return the router to its default behavior after it has been changed by the **`no`** form of this command.
Enable the propagation of the MPLS tunnel label TTL to the IP TTL at the egress router.	*propagate ttl mpls-to-ip*	The default behavior of the SmartEdge router is to propagate of the MPLS tunnel label TTL to the IP TTL at the egress router, so the **`propagate ttl mpls-to-ip`** command is used to return the router to its default behavior after it has been changed by the **`no`** form of this command.

2.2 Configuring MPLS Static

To configure MPLS static, perform the tasks described in the following sections.

2.2.1 Create an MPLS Static Routing Instance

To create an MPLS static routing instance, perform the task described in Table 3.

Table 3 Configure an MPLS Static Routing Instance

Task	Root Command	Notes
Create an MPLS static routing instance, and enter MPLS static router configuration mode.	*router mpls-static*	Enter this command in context configuration mode. Note that MPLS is supported in the local context only.

2.2.2 Configure an MPLS Static Interface

To configure an MPLS static interface, perform the tasks described in Table 4. Enter all commands in MPLS static interface configuration mode, unless otherwise noted.

Table 4 Configure a Static Interface

Task	Root Command	Notes
Enable MPLS static on an interface, and access MPLS interface configuration mode.	*interface (MPLS and RSVP)*	Enter this command in MPLS static router configuration mode.
Configure a static MPLS label-action mapping for an intermediate LSR.	*label-action*	Use the following command syntax: `label-action in-label-num [php egress-addr | pop | swap out-label-num next-hop-addr]` Use the `swap` keyword to replace the incoming label with the outgoing label.
Configure a static MPLS label-action mapping for an egress LSR.	*label-action*	Use the following command syntax: `label-action in-label-num pop` Use the `pop` keyword to remove the top label in the label stack.

2.2.3 Configure an MPLS Static LSP

To configure an MPLS static LSP, perform the tasks described in Table 5. Enter all commands in MPLS static LSP configuration mode, unless otherwise noted.

Table 5 Configure an MPLS Static LSP

Task	Root Command	Notes
Create a static LSP and enter MPLS static LSP configuration mode.	*lsp*	Enter this command in MPLS static router configuration mode.
Associate a description with a static LSP.	*description (MPLS static and RSVP LSPs)*	—
Configure a next-hop entry for a static LSP.	*next-hop*	—

Table 5 Configure an MPLS Static LSP

Task	Root Command	Notes
Specify the IP address of the egress LSR in a static LSP.	*egress*	An egress LSR is the last router in the chain of routers that constitute an LSP; see Figure 1.
Configure the outgoing label number for a static LSP.	*out-label*	—

2.3 Configuring RSVP

To configure RSVP, perform the tasks described in the following sections:

2.3.1 Create an RSVP Routing Instance

To create an RSVP routing instance, perform the tasks described in Table 6. Enter all commands in RSVP router configuration mode, unless otherwise noted.

Table 6 Create an RSVP Routing Instance

Task	Root Command	Notes
Create an RSVP routing instance within a context and enter RSVP router configuration mode.	*router rsvp*	Enter this command in context configuration mode.
(Optional) Enable BFD on the IP interface that is configured for RSVP.	*bfd*	BFD detection is disabled on RSVP links by default. You must enable BFD at both the router RSVP level and for any individual neighbor that requires FRR. For more information about configuring BFD on RSVP LSPs, see *Configuring BFD*.

Table 6 Create an RSVP Routing Instance

Task	Root Command	Notes
Enable an egress router to advertise an explicit null label (value 0), in place of an implicit null label (value 3), to the penultimate hop router.	*explicit-null (RSVP)*	By default, RSVP advertises an implicit null label for directly connected prefixes. An implicit null label causes the upstream router to perform penultimate hop popping (PHP), and the implicit null label is not transmitted on the egress router. In some cases, such as QoS enforcement, PHP may not be desirable. In those cases, using the **explicit-null** command causes the egress router to advertise an explicit null label in place of an implicit null label for directly connected prefixes, which forces the upstream router to transmit packets with an explicit null label on the last hop.
Set the amount of time that RSVP waits after a failed interface comes back up before traffic is switched back to the primary LSP from a bypass LSP.	*frr-auto-revert-delay*	When the delay-interval value is changed, and it is lower than the delay interval set for any existing bypass RSVP LSPs that are scheduled to switch back to their primary LSPs, then their delay timer is reset to the new, lower value. From Release 2.6.5.2 to Release 5.0.3.1 of the SmartEdge router, when the NFRR auto-revert delay is enabled, traffic is automatically switched to the primary LSP after the specified delay interval has elapsed. Starting with Release 5.0.3.2, after the delay interval has elapsed, a new instance of the primary LSP must be established before traffic is switched to it; otherwise, traffic continues to use the bypass LSP. If the interface goes down before the specified delay interval has elapsed, the traffic does not switch back to the primary LSP, but continues to use the bypass LSP. The delay timer is in effect only when a previously failed interface comes back up and stays up. Use the **no** form of this command to disable the NFRR auto-revert delay. If the NFRR auto-revert delay is disabled, then all existing bypass LSPs do not switch back to their primary LSPs, even if their delay timer has started.

Table 6 Create an RSVP Routing Instance

Task	Root Command	Notes
Enable RSVP LSPs to serve as Interior Gateway Protocol (IGP) shortcuts to nodes in a network.	*igp-shortcut*	When RSVP LSPs are enabled to serve as IGP shortcuts, Open Shortest Path First (OSPF) includes the RSVP LSPs in their Shortest Path First (SPF) calculation when determining the shortest-path tree to all nodes in a network. In order for the shortcuts to work, you must also configure the **`mpls igp-shortcut`** command in OSPF configuration mode. The **`igp-shortcut`** command (in RSVP router configuration mode) enables all RSVP LSPs for the specified RSVP routing instance to serve as IPG shortcuts. To enable only a specific RSVP LSP to serve as an IGP shortcut, enter this command in RSVP LSP configuration mode.
Enable the use of RSVP LSPs for tunneling LDP data and control traffic.	*tunnel-shortcut*	To configure LDP over RSVP, use the **`tunnel-shortcut`** command in RSVP router configuration mode in conjunction with the **`tunnel-shortcut`** command in LDP router configuration mode and the **`mpls tunnel-shortcut`** command in OSPF router configuration mode. For more information, see *Enabling LDP over RSVP*.
Enable the generation of RSVP-INFO messages when any RSVP LSP changes state.	*log-lsp-up-down*	The generation of RSVP-INFO messages cannot be disabled using the **`no terminal monitor`** command. Use the **`no log-lsp-up-down`** command to disable the generation of RSVP-INFO messages.
Configure the RSVP record route object (RRO) IP prefix type.	*rro-prefix-type*	Enter this command in RSVP router configuration mode. You can change the IP prefix inside an RRO to be either the router ID or the interface IP address. This is used for node protection and interarea node protection. During NFRR, the PLR LSR needs to match the bypass RSVP LSP egress IP address with the IP prefix inside the RRO of the next-next-hop node.

2.3.2 Configure an RSVP LSP

To configure an RSVP LSP, perform the tasks described in Table 7. Enter all commands in RSVP LSP configuration mode, unless otherwise noted.

Note: Depending on the command syntax you use for the `lsp` command in RSVP router configuration mode, you can create a standard or backup RSVP LSP.

Table 7 Configure an RSVP LSP

#	Task	Root Command	Notes
1.	Create a standard RSVP LSP and enter RSVP LSP configuration mode.	*lsp*	Enter this command in RSVP router configuration mode. Use the following command syntax: **lsp** ***lsp-name***
2.	Create a backup RSVP LSP and enter RSVP LSP configuration mode.	*lsp*	Enter this command in RSVP router configuration mode. Use the following command syntax: **lsp** ***lsp-name*** **backup-for** ***lsp-name***
3.	Create a backup to the backup RSVP LSP you created in step 2 and enter RSVP LSP configuration mode.	*lsp*	Enter this command in RSVP router configuration mode. Use the following command syntax: **lsp** ***lsp-name*** **backup-for** ***lsp-name*** In this case, replace the ***lsp-name*** argument with the name of the backup LSP you configured in step 2.
4.	Associate a description with an RSVP LSP.	*description (MPLS static and RSVP LSPs)*	—

Table 7 Configure an RSVP LSP

#	Task	Root Command	Notes
5.	Configure the LSP to be exclusive.	*exclusive*	When a mapped RSVP LSP is configured to be exclusive, it supports PW traffic only. Only primary LSPs can be configured to be exclusive. Mapped bypass and backup LSPs inherit exclusivity from the primary LSP. For more information about mapped RSVP LSPs, see one of the following sections, as appropriate: • *L2VPN XC-to-MPLS LSP Mapping* in *Configuring L2VPN*. • *Pseudowire Load Balancing* in *Configuring VPLS*.
6.	Enable an RSVP LSP to serve as an IGP shortcut to nodes in a network.	*igp-shortcut*	When a RSVP LSP is enabled to serve as an IGP shortcut, Open Shortest Path First (OSPF) can include the RSVP LSP in their Shortest Path First (SPF) calculation when determining the shortest-path tree to all nodes in a network. In order for IGP shortcuts to work, you must also configure the **mpls igp-shortcut** command in OSPF router configuration mode. This command (in RSVP LSP configuration mode) enables the specified RSVP LSP to serve as an IPG shortcut. To enable all RSVP LSPs for an RSVP routing instance to serve as IGP shortcuts, enter this command in RSVP router configuration mode.

Table 7 Configure an RSVP LSP

#	Task	Root Command	Notes
7.	Enable the use of the RSVP LSP for tunneling.	*tunnel-shortcut*	To configure LDP over RSVP, use the `tunnel-shortcut` command in RSVP LSP configuration mode in conjunction with the `tunnel-shortcut` command in LDP router configuration mode and the `mpls tunnel-shortcut` command in OSPF router configuration mode. For more information, see *Enabling LDP over RSVP*.
8.	Specify the IP address of the ingress LSR in an RSVP LSP.	*ingress*	An ingress LSR is the first router in the chain of routers that constitute an LSP; see Figure 1. An ingress IP address does not have to be specified for an RSVP LSP. If it is not specified, the IP address of the interface used to reach the egress IP address is used. If the interface changes, the ingress IP address will also change; however, if an ingress IP address is specified, then the specified address is always used.
9.	Specify the IP address of the egress LSR in an RSVP LSP.	*egress*	An egress LSR is the last router in the chain of routers that constitute an LSP; see Figure 1.
10.	Permit an LSP to be protected by a bypass RSVP LSP.	*local-protection*	When configured, the LSP advertises to the ingress and transit nodes that a bypass RSVP LSP can be used to provide MPLS fast reroute protection. This configuration affects both ingress LSR and the transit LSRs of the LSP operation.

Table 7 Configure an RSVP LSP

#	Task	Root Command	Notes
11.	Assign a configured explicit route to an LSP.	*source-path*	Before you can assign a source path to an LSP, you must configure an explicit route to use as the source path. Use the **`explicit-route`** command in MPLS router configuration mode to indicate a list of specific hops through a network that you want for your LSP, and then use the **`source-path`** command to assign that explicit route to your LSP.
12.	Configure an RSVP LSP to actively record the routes through which it forwards packets.	*record-route*	You can use the recorded route information for troubleshooting and to prevent routing loops.
13.	Enable or disable an RSVP LSP.	*shutdown (RSVP LSP)*	Use the **`no`** form of this command to enable an existing RSVP LSP.

2.3.3 Configure a Bypass RSVP LSP

To configure a bypass RSVP LSP, perform the tasks described in Table 8. Enter all commands in RSVP LSP configuration mode, unless otherwise noted.

Note: Depending on the command syntax you use for the **`lsp`** command in RSVP router configuration mode, you can create a bypass RSVP for one of the following protection schemes:

- Link protection
- Node protection

Table 8 Configure a Bypass RSVP LSP

Task	Root Comma nd	Notes
Create a bypass RSVP LSP for link protection and enter RSVP LSP configuration mode.	*lsp*	Enter this command in RSVP router configuration mode. Use the following command syntax: **`lsp`** *`lsp-name`* **`bypass`** *`ip-addr`*

Table 8 Configure a Bypass RSVP LSP

Task	Root Comma nd	Notes
Create a bypass RSVP LSP for node protection and enter RSVP LSP configuration mode.	*lsp*	Enter this command in RSVP router configuration mode. Use the following command syntax: `lsp lsp-name bypass ip-addr node-protect-lsp-egress ip-addr`
Associate a description with a bypass RSVP LSP.	*description (MPLS static and RSVP LSPs)*	—
Configure a bypass RSVP LSP to match the next-next-hop interface IP address.	*fast-reroute*	If the next-next-hop node does not use the router ID in the RSVP RRO, the PLR LSR can optionally configure the bypass LSP to match a known next-next-hop interface IP address. This is also useful in the case of interarea node protection.
Enable an RSVP LSP to serve as an IGP shortcut to nodes in a network.	*igp-shortcut*	When a RSVP LSP is enabled to serve as an IGP shortcut, link-state protocols, such as IS-IS and OSPF, include the RSVP LSP in their Shortest Path First (SPF) calculation when determining the shortest-path tree to all nodes in a network. This command (in RSVP LSP configuration mode) enables the specified RSVP LSP to serve as an IPG shortcut. To enable all RSVP LSPs for an RSVP routing instance to serve as IGP shortcuts, enter this command in RSVP router configuration mode.

Table 8 Configure a Bypass RSVP LSP

Task	Root Command	Notes
Specify the IP address of the ingress LSR in a bypass RSVP LSP.	*ingress*	An ingress LSR is the first router in the chain of routers that constitute an LSP; see Figure 1. An ingress IP address does not have to be specified for an RSVP LSP. If it is not specified, the IP address of the interface used to reach the egress IP address is used. If the interface changes, the ingress IP address will also change; however, if an ingress IP address is specified, then the specified address is always used.
Specify the IP address of the egress LSR in a bypass RSVP LSP.	*egress*	An egress LSR is the last router in the chain of routers that constitute an LSP; see Figure 1.
Assign a configured explicit route to an LSP.	*source-path*	Before you can assign a source path to an LSP, you must configure an explicit route to use as the source path. Use the **`explicit-route`** command in MPLS router configuration mode to indicate a list of specific hops through a network that you want for your LSP, and then use the **`source-path`** command to assign that explicit route to your LSP.
Configure a bypass RSVP LSP to actively record the routes through which it forwards packets.	*record-route*	You can use the recorded route information for troubleshooting, and to prevent routing loops.
Enable or disable a bypass RSVP LSP.	*shutdown (RSVP LSP)*	Use the **`no`** form of this command to enable an existing RSVP LSP.
Apply a constraint.	*constraint*	Enter this command in RSVP LSP configuration mode.
Configure the administrative groups to exclude or include in the LSP.	*admin-group*	—

2.3.4 Configure an Explicit Route

When an LSP is configured to use an explicit route, it uses the path determined by that explicit route. If the path defined by the explicit route is not topologically possible, because either the network is partitioned or insufficient resources are available, the LSP fails. No alternate paths can be used. If the LSP succeeds, it continues to use the explicit route.

To configure an explicit route, perform the tasks described in Table 9.

Table 9 Configure an Explicit Route

Task	Root Command	Notes
Create an explicit route and access RSVP explicit route configuration mode.	*explicit-route*	Enter this command in RSVP router configuration mode.
Configure a next-hop entry for an RSVP explicit route.	*next-hop*	Enter this command in RSVP explicit route configuration mode.

2.3.5 Configure an RSVP Interface

To configure an RSVP interface, perform the tasks described in Table 10.

Table 10 Configure an RSVP Interface

Task	Root Command	Notes
Enable RSVP on an interface, and access RSVP interface configuration mode.	*interface (MPLS and RSVP)*	Enter this command in RSVP router configuration mode.
Enable authentication for an RSVP interface.	*monitor ospf interface*	Enter this command in RSVP interface configuration mode. Key chains allow you to control authentication for SmartEdge router routing protocols. Neighboring routers using RSVP to exchange reservation and path messages must utilize an accepted key ID and key string. If multiple key IDs have been configured, the one with the most recent send time exceeding the current time is used. All key IDs that have not expired and that have a receive time exceeding the current time are accepted. Routes within the same area are not required to use the same authentication key ID. However, if two routers directly exchange updates, they must have the same authentication key ID.
Configure the RSVP reservation state lifetime.	For the complete list of tasks used to configure the RSVP reservation state lifetime, see Configure the RSVP Reservation State Lifetime.	
Configure RSVP graceful restart.	For the complete list of tasks used to configure RSVP graceful restart, see Configure RSVP Graceful Restart.	

2.3.6 Configure the RSVP Reservation State Lifetime

When RSVP is enabled, refresh messages are frequently generated and sent 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

To configure an RSVP reservation state lifetime, perform the tasks described in Table 11. Enter all commands in RSVP interface configuration mode, unless otherwise noted.

Table 11 Configure the RSVP Reservation State Lifetime

Task	Root Command	Notes
Configure the frequency of generating refresh messages.	*refresh-interval*	Before you can specify the lifetime of a reservation state using the refresh-interval command, you must ensure that the keep-multiplier timing parameter has also been specified.
Configure the RSVP keep-multiplier timing parameter.	*keep-multiplier*	Before you can specify the lifetime of a reservation state using the **`keep-multiplier`** command, you must ensure that the refresh-interval timing parameter has also been specified.

2.3.7 Configure RSVP Graceful Restart

Keep the following in mind before enabling and configuring RSVP graceful restart on your system:

- When RSVP graceful restart is first enabled on an RSVP routing instance, the RSVP graceful restart helper mode is automatically enabled by default, and the restart and recovery timers use their default values.

- When RSVP graceful restart is enabled on an RSVP instance, all RSVP LSPs that are configured under that instance take on the same timer values. However, you can use the **`graceful-restart restart-time`** and **`graceful-restart recovery-time`** commands in RSVP interface configuration mode to administratively configure an individual RSVP LSP to use a restart and recovery time that is different from all of the other LSPs in the RSVP instance.

- RSVP graceful restart and recovery timers can be configured on RSVP routing instance as well as on a particular RSVP interface. Be aware that the configuration of the restart and recovery timers in RSVP interface configuration mode takes precedence over any restart and recovery timer configuration that is performed in RSVP router configuration mode.

- Helper mode can be enabled and disabled in RSVP router configuration mode only. Be aware that helper behavior is the same for all neighbors in the RSVP routing instance. For example, if helper mode is disabled on an RSVP instance, helper mode is also disabled on the neighbors in that RSVP instance.

- The maximum helper mode restart and recovery timers are used to determine how long the RSVP instance waits for a neighbor to restart and recover after a restart. If a neighbor advertises restart and recovery timer values that are different from the locally configured maximum helper restart and recovery timer values, the smaller timer values take precedence. If you do not configure the maximum helper restart and recovery timers on the SmartEdge router, the router uses the restart and recovery timer values advertised by the neighbor.

- RSVP Hello messages are disabled by default. They must be enabled on the interface and the neighbor to which it connects for RSVP graceful restart to function properly.

Note: Before you can configure graceful restart for an RSVP routing instance, you need to create and configure an RSVP routing instance, as described in Create an RSVP Routing Instance.

To enable and configure RSVP graceful restart on an RSVP instance, perform the tasks described in Table 12. To configure RSVP graceful restart and recovery timer values on an individual RSVP interface, perform the tasks described in Table 13.

Table 12 Enable and Configure RSVP Graceful Restart on an RSVP Routing Instance

Task	Root Command	Notes
Enter RSVP router configuration mode for the RSVP routing instance in which you want to enable RSVP graceful restart.	*router rsvp*	Enter this command in context configuration mode.

Table 12 Enable and Configure RSVP Graceful Restart on an RSVP Routing Instance

Task	Root Command	Notes
Enable RSVP graceful restart.	*graceful-restart (RSVP)*	Enter this command in RSVP router configuration mode. RSVP graceful restart uses RSVP Hello messages to determine if a neighbor is down, and if it should initiate graceful restart procedures. Use the **`hello interval`** and **`hello keep-multiplier`** commands in RSVP interface configuration mode to enable and configure RSVP Hello messages. The default RSVP graceful restart time is 30 seconds and the default RSVP graceful restart recovery time is 60 seconds.
(Optional) Modify the graceful restart time.	*graceful-restart restart-time* ***`seconds`***	Use the **`graceful-restart restart-time`** command in RSVP router configuration mode to modify the graceful restart time globally.
(Optional) Modify the graceful restart recovery time.	*graceful-restart recovery-time* ***`seconds`***	Use the **`graceful-restart recovery-time`** ***`seconds`*** command in RSVP router configuration mode to modify the graceful restart recovery time globally.
(Optional) Disable RSVP graceful restart helper mode.	*no graceful-restart helper*	RSVP graceful restart helper mode is enabled by default when RSVP graceful restart is enabled with the **`graceful-restart`** command. When graceful restart is disabled (with the **`no graceful-restart`** command), RSVP graceful restart helper mode is automatically disabled, as well. Use the **`graceful-restart helper`** command to reenable helper mode on a system where it is disabled. Be aware that helper mode cannot be enabled on a system where RSVP graceful restart is disabled.

Table 12 Enable and Configure RSVP Graceful Restart on an RSVP Routing Instance

Task	Root Command	Notes
(Optional) Specifies the maximum time interval (in seconds) that the SmartEdge router waits for a neighbor to come up after graceful restart before considering the neighbor to be down.	*graceful-restart maximum_helper_recovery-time* ***interval***	Replace the ***interval*** argument with the number of seconds the SmartEdge router waits for a neighbor to come up after graceful restart before considering the neighbor to be down. When the helper recovery timer expires, the status information for the neighbor expires and is no longer retained. Range is 20 to 3600. If you do not configure the maximum helper recovery timer on the SmartEdge router, the router uses the restart timer value advertised by the neighbor. To configure the helper recovery timer, you must first enable RSVP helper mode by issuing the **graceful-restart helper** command.
(Optional) Specifies the maximum time interval (in seconds) that the SmartEdge router retains information about the state of a neighbor after graceful restart.	*graceful-restart maximum_helper_restart-time* ***interval***	Replace the ***interval*** argument with the number of seconds the SmartEdge router retains information about the state of a neighbor after graceful restart. When the helper restart timer expires, the LSPs between the SmartEdge router and the restarting neighbor are no longer retained. The range of values for the ***interval*** argument is 10 to 1800. If you do not configure the maximum helper restart timer on the SmartEdge router, the router uses the restart timer value advertised by the neighbor. To configure the helper restart timer, you must first enable RSVP helper mode by issuing the **graceful-restart helper** command.

Table 13 describes how to configure RSVP graceful restart on an individual RSVP interface.

Table 13 Configure RSVP Graceful Restart on an RSVP Interface

Task	Root Command	Notes
Enter RSVP router configuration mode for the RSVP routing instance that contains the RSVP LSP you want to configure.	*router rsvp*	Enter this command in context configuration mode.
Access RSVP interface configuration mode for the interface you want to configure.	*interface (MPLS and RSVP)*	Enter this command in RSVP router configuration mode.
Configure the interval at which RSVP Hello messages are sent out from the specified interface.	*hello interval (RSVP)*	RSVP Hello messages are disabled by default. They must be enabled on the interface and the neighbor to which it connects for RSVP graceful restart to function properly.
(Optional). Configure the number of lost RSVP Hello messages that can be missed by a neighbor before it declares that the peer adjacency is down.	*hello keep-multiplier*	—
(Optional) Modify the graceful restart time for the interface.	*graceful-restart restart-time* ***`seconds`***	Use the **`graceful-restart restart-time`** command in RSVP interface configuration mode to modify the graceful restart time on a specific interface.
(Optional) Modify the graceful restart recovery time for the interface.	*graceful-restart recovery-time* ***`seconds`***	Use the **`graceful-restart recovery-time`** ***`seconds`*** command in RSVP interface configuration mode to modify the graceful restart recovery time on a specific interface.

2.4 Configuring CSPF Calculation

To configure CSPF calculation, perform the tasks described in Table 14. Enter all commands in RSVP router configuration mode, unless otherwise noted.

Table 14 Configure CSPF Calculation

Task	Root Command	Notes
Configure a global router ID (IP address) for the SmartEdge router.	*router-id (contexts)*	Enter this command in context configuration mode. You must configure a global router ID for CSPF to work properly.
Configure the subscription bandwidth on a interface.	*over-subscription-factor*	—
Configure the amount of time for the headend (ingress) router waits between attempts to establish the primary path.	*timer retry*	—
Create an explicit route and access RSVP explicit route configuration mode.	*explicit-route*	For CSPF, the explicit route name is used by the **`dynamic-path`** command.
Configure a next-hop entry for an RSVP explicit route.	*next-hop*	Enter this command in RSVP explicit route configuration mode.
Configure RSVP-TE link attributes for traffic engineering.	*attribute (RSVP)*	Enter this command in RSVP link attribute configuration mode. By default, no administrative attributes are defined.
Specify the bandwidth to be reserved.	*bandwidth*	Enter this command in RSVP interface configuration mode. Use the **`show rsvp interface`** command to display the reserved bandwidth at all eight priority levels.
Configure administrative attributes on an interface.	*admin-group*	Enter this command in RSVP interface configuration mode. An administrative group is defined in the **`attribute`** command.
Configure the TE metric for an interface.	*te-metric*	—
Create a constraint that is applied on a TE LSP.	*constraint*	Enter this command in RSVP constraint configuration mode. The **`admin-group`**, **`exclude`**, **`hop-limit`**, **`minimum-bandwidth`**, and **`priority`** commands define the constraint that is applied to the TE LSP.
Specify which administrative groups are valid on the LSP.	*admin-group*	—

Table 14 Configure CSPF Calculation

Task	Root Command	Notes
Configure which nodes and links to exclude from the CSPF path computation.	*exclude*	—
Configure the maximum number of routers that the tunnel can traverse, including ingress and egress routers.	*hop-limit*	—
Configure the required minimum bandwidth requirement for an LSP to be applied.	*minimum-bandwidth*	—
Configure the priority of a tunnel.	*priority (RSVP)*	You typically configure the setup priority and hold priority to the same value.
Specify a dynamic path that CSPF applies to the LSP.	*dynamic-path*	Enter this command in RSVP LSP configuration mode. The dynamic path name is the explicit route name that you define using the **`explicit-route`** command.
Apply the constraints to the TE LSP defined at the RSVP router configuration level.	*constraint*	—
Enable CSPF.	*cspf*	—

2.5 MPLS Operations

To manage MPLS functions, perform the appropriate tasks described in Table 15. Enter the **`show`** commands in any mode; enter the **`clear`** and **`debug`** commands in exec mode.

For information about troubleshooting MPLS, see *Troubleshooting MPLS*.

Table 15 MPLS Operations Tasks

Task	Root Command
Clear Resource Reservation Protocol (RSVP) counter information.	*clear rsvp counters*
Enable the generation of debug messages for Constrained Shortest Path First (CSPF).	*debug cspf*

Table 15   MPLS Operations Tasks

Task	Root Command
Enable the generation of debug messages for label manager (LM) activities.	*debug lm*
Initiate a Multiprotocol Label Switching (MPLS) ping across a Label Distribution Protocol (LDP) label-switched path (LSP).	*ping mpls ldp*
Initiates an MPLS ping to a medium access control (MAC) address in a Virtual Private LAN Services (VPLS) network.	*ping mpls mac-address*
Test the status of a pseudo-wire.	*ping mpls pw*
Initiate an MPLS ping across an RSVP LSP.	*ping mpls rsvp*
Display the current MPLS configuration information for the current context.	*show configuration mpls*
Display CSPF database information.	*show cspf database*
Display the current RSVP configuration information for the current context.	*show configuration rsvp*
Display LM information.	*show mpls*
Display MPLS interface information.	*show mpls interface*
Enable debug messages for static MPLS. You can filter the output by configuration, general, ISM, label-action, label management, LSP, and next-hop event messages.	*debug mpls-static*
Display the current MPLS static configuration information for the current context.	*show configuration mpls-static*
Display static MPLS label mapping information.	*show mpls-static label-action*
Display static MPLS LSP information. Displays the LSP state, name, ID, next-hop, endpoint, and out-label. If you add the detail keyword, it also provides the LSP circuit handle, the outgoing circuit handle, and the outgoing interface grid.	*show mpls-static lsp*
Display LM process information.	*show process lm*
Display RSVP counter information.	*show rsvp counters*
Display RSVP debug information.	*show rsvp debug*
Display the explicit route information.	*show rsvp explicit-route*
Display RSVP interface information.	*show rsvp interface*
Display RSVP LSP information.	*show rsvp lsp*
Trace the LSP route that packets take when traveling to the specified destination.	*traceroute mpls*

3 Configuration Examples

The following sections provide MPLS configuration examples.

For information about troubleshooting MPLS, see *Troubleshooting MPLS*.

3.1 MPLS Static LSP

The following example illustrates three routers configured to create an MPLS static LSP tunnel between **LSR_A** and **LSR_C**, using **LSR_B** as a next hop. Figure 5 shows the network topology for the configuration.

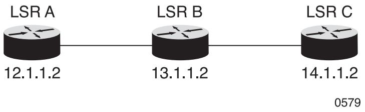

Figure 5 MPLS Static LSP Tunnel Network Topology

The configuration for **LSR_A** is as follows:

```
[local]LSR_A#config
[local]LSR_A(config)#context local
[local]LSR_A(config-ctx)#router mpls-static
[local]LSR_A(config-mpls-static)#lsp new
[local]LSR_A(config-mpls-static-lsp)#next-hop 13.1.1.2
[local]LSR_A(config-mpls-static-lsp)#out-label 30
[local]LSR_A(config-mpls-static-lsp)#egress 14.1.1.2
[local]LSR_A(config-mpls-static-lsp)#end
```

The configuration for **LSR_B** is as follows:

```
[local]LSR_B#config
[local]LSR_B(config)#context local
[local]LSR_B(config-ctx)#router mpls-static
[local]LSR_B(config-mpls-static)#interface foo
[local]LSR_B(config-mpls-static-if)#label-action 30 swap 37 14.1.1.2
[local]LSR_B(config-mpls-static-if)#end
```

The configuration for **LSR_C** is as follows:

```
[local]LSR_C#config
[local]LSR_C(config)#context local
[local]LSR_C(config-ctx)#router mpls-static
[local]LSR_C(config-mpls-static)#interface foo
[local]LSR_C(config-mpls-static-if)#label-action 37 pop
[local]LSR_C(config-mpls-static-if)#end
```

3.2 RSVP LSP

The following example illustrates three routers configured to create an RSVP LSP tunnel between **LSR A** and **LSR C**, using **LSR B** as a next hop. Figure 6 shows the network topology for the configuration.

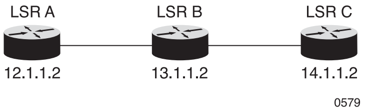

Figure 6 RSVP LSP Tunnel Topology

The configuration for **LSR_A** is as follows:

```
[local]LSR_A#config
[local]LSR_A(config)#context local
[local]LSR_A(config-ctx)#router rsvp
[local]LSR_A(config-rsvp)#interface foo
[local]LSR_A(config-rsvp-if)#exit
[local]LSR_A(config-rsvp)#explicit-route two
[local]LSR_A(config-rsvp-explicit-route)#next-hop 13.1.1.2
[local]LSR_A(config-rsvp-explicit-route)#next-hop 14.1.1.2
[local]LSR_A(config-rsvp-explicit-route)#exit
[local]LSR_A(config-rsvp)#lsp newtest
[local]LSR_A(config-rsvp-lsp)#ingress 12.1.1.2
[local]LSR_A(config-rsvp-lsp)#egress 14.1.1.2
[local]LSR_A(config-rsvp-lsp)#source-path two
[local]LSR_A(config-rsvp-lsp)#end
```

The configuration for **LSR B** is as follows:

```
[local]LSR_B#config
[local]LSR_B(config)#context local
[local]LSR_B(config-ctx)#router rsvp
[local]LSR_B(config-rsvp)#interface foo
[local]LSR_B(config-rsvp-if)#end
```

The configuration for **LSR_C** is as follows:

```
[local]LSR_C#config

[local]LSR_C(config)#context local

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

[local]LSR_C(config-rsvp)#interface foo

[local]LSR_C(config-rsvp-if)#end
```

3.3 RSVP Graceful Restart

The following example shows how to enable RSVP graceful restart helper mode, set the helper recovery timer to 200 seconds, and set the helper restart timer to 100 seconds:

```
[local]jazz(config)#context local
[local]jazz(config-ctx)#router rsvp
[local]Redback(config-rsvp)#graceful-restart
[local]Redback(config-rsvp)#graceful-restart maximum_helper_recovery-time 200
[local]Redback(config-rsvp)#graceful-restart maximum_h
elper_restart-time 100
```

3.4 CSPF Calculation

The following example illustrates four routers configured to use CSPF to generate paths to establish LSPs based upon link coloring. Figure 7 shows the network topology for the configuration:

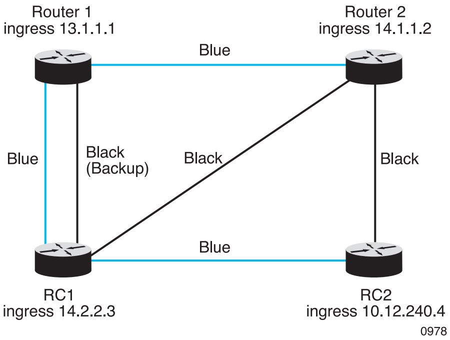

Figure 7 CSPF Example

The example that follows configures router RBAK1.

Configure a router ID:

```
[local]Redback#configure

[local]Redback(config)#context local

[local]Redback(config-ctx)#router-id 193.25.105.83
```

Configure a constraint:

```
[local]Redback#configure

[local]Redback(config)#context local

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

[local]Redback(config-rsvp)#log-lsp-up-down
```

Define link attributes:

```
[local]Redback(config-rsvp)#link-attributes

[local]Redback(config-rsvp-link-attr)#attribute BLUE 1

[local]Redback(config-rsvp-link-attr)#attribute BLACK 2

[local]Redback(config-rsvp-link-attr)#exit
```

Create BLUE_PATH constraint:

```
[local]Redback(config-rsvp)#constraint BLUE_PATH

[local]Redback(config-rsvp-constr)#admin group include-all BLUE

[local]Redback(config-rsvp-constr)#exit
```

Create BLACK_PATH constraint:

```
[local]Redback(config-rsvp)#constraint BLACK_PATH

[local]Redback(config-rsvp-constr)#admin group include-all BLACK

[local]Redback(config-rsvp-constr)#exit

[local]Redback(config-rsvp)#interface if.1
```

Assign admin group to interface:

```
[local]Redback(config-rsvp-if)#admin-group BLUE

[local]Redback(config-rsvp-if)#exit

[local]Redback(config-rsvp)#interface if.5
```

Assign admin group to interface:

```
[local]Redback(config-rsvp-if)#admin-group BLUE

[local]Redback(config-rsvp-if)#exit

[local]Redback(config-rsvp)#interface if.7

[local]Redback(config-rsvp)#interface loopback0
```

Create an RSVP LSP between RBAK1 and RBAK2:

```
[local]Redback(config-rsvp)#lsp RBAK1_RBAK2_BLUE

[local]Redback(config-rsvp-lsp)#ingress 13.1.1.1

[local]Redback(config-rsvp-lsp)#egress 14.1.1.2
```

Apply constraint to LSP:

```
[local]Redback(config-rsvp-lsp)#constraint BLUE_PATH

[local]Redback(config-rsvp-lsp)#exit
```

Create an RSVP LSP between RBAK1 and RC1:

```
[local]Redback(config-rsvp)#lsp RBAK1_RC1_BLUE

[local]Redback(config-rsvp-lsp)#ingress 13.1.1.1

[local]Redback(config-rsvp-lsp)#egress 14.2.2.3
```

Apply constraint to RSVP LSP:

```
[local]Redback(config-rsvp-lsp)#constraint BLUE_PATH
```

Create a backup RSVP LSP between RBAK1 and RC1. This LSP acts as backup for RBAK1_RC1_BLUE:

```
[local]Redback(config-rsvp)#lsp RBAK1_RC1_BLACK

[local]Redback(config-rsvp-lsp)#ingress 13.1.1.1

[local]Redback(config-rsvp-lsp)#egress 14.2.2.3
```

Apply constraint to backup RSVP LSP:

```
[local]Redback(config-rsvp-lsp)#constraint BLACK_PATH
```

Assign the LSP RBAK1_RC1_BLACK as a backup for LSP RBAK1_RC1_BLUE.

```
[local]Redback(config-rsvp)#lsp RBAK1_RC1_BLACK
backup-for RBAK1_RC1_BLUE
```

Apply constraint to RSVP LSP:

```
[local]Redback(config-rsvp-lsp)#constraint BLACK_PATH
```

The example that follows configures router RBAK2.

Configure a router ID:

```
[local]Redback#configure

[local]Redback(config)#context local

[local]Redback(config-ctx)#router-id 193.25.102.84
```

Configure constraint for RBAK 2:

```
[local]Redback#configure

[local]Redback(config)#context local

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

[local]Redback(config-rsvp)#log-lsp-up-down

[local]Redback(config-rsvp)#graceful-restart
```

Define link attributes:

```
[local]Redback(config-rsvp)#link-attributes
[local]Redback(config-rsvp-link-attr)#attribute BLUE 1
[local]Redback(config-rsvp-link-attr)#attribute BLACK 2
[local]Redback(config-rsvp-link-attr)#exit
```

Create BLUE_PATH constraint:

```
[local]Redback(config-rsvp)#constraint BLUE_PATH
[local]Redback(config-rsvp-constr)#admin group include-all BLUE
[local]Redback(config-rsvp-constr)#exit
```

Create BLACK_PATH constraint:

```
[local]Redback(config-rsvp)#constraint BLACK_PATH
[local]Redback(config-rsvp-constr)#admin group include-all BLACK
[local]Redback(config-rsvp-constr)#exit
[local]Redback(config-rsvp)#interface if.15
```

Assign admin group to interface:

```
[local]Redback(config-rsvp-if)#admin-group BLUE BlACK
[local]Redback(config-rsvp-if)#exit
[local]Redback(config-rsvp)#interface if.17
[local]Redback(config-rsvp)#interface if.18
```

Assign admin group to interface:

```
[local]Redback(config-rsvp-if)#admin-group BlACK
[local]Redback(config-rsvp-if)#exit
[local]Redback(config-rsvp)#interface if.16
```

Assign admin group to interface:

```
[local]Redback(config-rsvp-if)#admin-group BlACK
[local]Redback(config-rsvp-if)#exit
[local]Redback(config-rsvp)#interface if.25
[local]Redback(config-rsvp)#interface loopback0
```

Create an RSVP LSP between RBAK2 and RC1:

```
[local]Redback(config-rsvp)#lsp RBAK2_RC1_BLACK
[local]Redback(config-rsvp-lsp)#ingress 14.1.1.2
[local]Redback(config-rsvp-lsp)#egress 14.2.2.3
```

Apply constraint to RSVP LSP:

```
[local]Redback(config-rsvp-lsp)#constraint BLACK_PATH
[local]Redback(config-rsvp-lsp)exit
```

Create an RSVP LSP between RBAK2 and RBAK1:

```
[local]Redback(config-rsvp)#lsp Connect_RBAK2_To_RBAK1_BLUE
[local]Redback(config-rsvp-lsp)#ingress 14.1.1.2
[local]Redback(config-rsvp-lsp)#egress 13.1.1.1
```

Apply constraint to RSVP LSP:

```
[local]Redback(config-rsvp-lsp)#constraint BLACK_PATH
[local]Redback(config-rsvp-lsp)exit
```

Create RSVP LSP between RBAK2 and RC2:

```
[local]Redback(config-rsvp)#lsp Connect_RBAK2_To_RC2_BLACK
[local]Redback(config-rsvp-lsp)#ingress 14.1.1.2
[local]Redback(config-rsvp-lsp)#egress 10.12.240.4
```

Apply constraint to RSVP LSP:

```
[local]Redback(config-rsvp-lsp)#constraint BLACK_PATH
```

The example that follows configures router RC1.

Configure a router ID:

```
[local]Redback#configure
[local]Redback(config)#context local
[local]Redback(config-ctx)#router-id 193.25.101.81
```

Create a constraint:

```
[local]Redback#configure
[local]Redback(config)#context local
[local]Redback(config-ctx)#router rsvp
[local]Redback(config-rsvp)#log-lsp-up-down
[local]Redback(config-rsvp)#graceful-restart
```

Define link attributes:

```
[local]Redback(config-rsvp)#link-attributes
[local]Redback(config-rsvp-link-attr)#attribute BLUE 1
[local]Redback(config-rsvp-link-attr)#attribute BLACK 2
[local]Redback(config-rsvp-link-attr)#exit
```

Create BLUE_PATH constraint:

```
[local]Redback(config-rsvp)#constraint BLUE_PATH
[local]Redback(config-rsvp-constr)#admin group include-all BLUE
[local]Redback(config-rsvp-constr)#exit
```

Create BLACK_PATH constraint:

```
[local]Redback(config-rsvp)#constraint BLACK_PATH

[local]Redback(config-rsvp-constr)#admin group include-all BLACK

[local]Redback(config-rsvp-constr)#exit

[local]Redback(config-rsvp)#interface if.2

[local]Redback(config-rsvp)#interface if.3

[local]Redback(config-rsvp)#interface if.4
```

Assign admin group to interface:

```
[local]Redback(config-rsvp-if)#admin-group BLUE

[local]Redback(config-rsvp-if)#exit

[local]Redback(config-rsvp)#interface if.17

[local]Redback(config-rsvp)#interface if.19

[local]Redback(config-rsvp)#interface if.3

[local]Redback(config-rsvp-if)#admin-group BLACK
```

Assign admin group to interface:

```
[local]Redback(config-rsvp-if)#exit

[local]Redback(config-rsvp)#interface if.7

[local]Redback(config-rsvp)#interface if.8

[local]Redback(config-rsvp)#interface loopback0
```

Create an RSVP LSP between RC1 and RC2:

```
[local]Redback(config-rsvp)#lsp Connect_RC1_To_RC2_BLUE

[local]Redback(config-rsvp-lsp)#ingress 14.2.2.3

[local]Redback(config-rsvp-lsp)#egress 10.12.240.4
```

Apply constraint to RSVP LSP:

```
[local]Redback(config-rsvp-lsp)#constraint BLUE_PATH

[local]Redback(config-rsvp-lsp)#exit
```

Create RSVP LSP between RC1 and RBAK2:

```
[local]Redback(config-rsvp)#lsp Connect_RC1_To_RBAK2_BLACK

[local]Redback(config-rsvp-lsp)#ingress 14.2.2.3

[local]Redback(config-rsvp-lsp)#egress 14.1.1.2
```

Apply constraint to RSVP LSP:

```
[local]Redback(config-rsvp-lsp)#constraint BLACK_PATH

[local]Redback(config-rsvp-lsp)#exit
```

Create RSVP LSP between RC1 to RBAK1:

```
[local]Redback(config-rsvp)#lsp Connect_RC1_To_RBAK1_BLACK

[local]Redback(config-rsvp-lsp)#ingress 14.2.2.3

[local]Redback(config-rsvp-lsp)#egress 13.1.1.1
```

Apply constraint to RSVP LSP:

```
[local]Redback(config-rsvp-lsp)#constraint BLACK_PATH
```

The example that follows configures router RC2.

Configure a router ID:

```
[local]Redback#configure

[local]Redback(config)#context local

[local]Redback(config-ctx)#router-id 193.25.100.80
```

Create a constraint for RC2:

```
[local]Redback#configure

[local]Redback(config)#context local

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

[local]Redback(config-rsvp)#log-lsp-up-down

[local]Redback(config-rsvp)#graceful-restart
```

Define link attributes:

```
[local]Redback(config-rsvp)#link-attributes

[local]Redback(config-rsvp-link-attr)#attribute BLUE 1

[local]Redback(config-rsvp-link-attr)#attribute BLACK 2

[local]Redback(config-rsvp-link-attr)#exit
```

Create BLUE_PATH constraint:

```
[local]Redback(config-rsvp)#constraint BLUE_PATH

[local]Redback(config-rsvp-constr)#admin group include-all BLUE

[local]Redback(config-rsvp-constr)#exit
```

Create BLACK_PATH constraint:

```
[local]Redback(config-rsvp)#constraint BLACK_PATH

[local]Redback(config-rsvp-constr)#admin group include-all BLACK

[local]Redback(config-rsvp-constr)#exit

[local]Redback(config-rsvp)#interface if.2

[local]Redback(config-rsvp)#interface if.3

[local]Redback(config-rsvp)#interface if.4
```

Assign admin group to interface:

```
[local]Redback(config-rsvp-if)#admin-group BLUE

[local]Redback(config-rsvp-if)#exit

[local]Redback(config-rsvp)#interface interface.7
```

Assign admin group to interface:

```
[local]Redback(config-rsvp-if)#admin-group BLACK

[local]Redback(config-rsvp-if)#exit

[local]Redback(config-rsvp)#interface loopback0
```

Create RSVP LSP between RC1 and RBAK2:

```
[local]Redback(config-rsvp)#lsp Connect_RC2_To_RBAK2_BLACK

[local]Redback(config-rsvp-lsp)#ingress 10.12.240.4

[local]Redback(config-rsvp-lsp)#egress 14.1.1.2
```

Apply constraint to RSVP LSP:

```
[local]Redback(config-rsvp-lsp)# constraint BLACK_PATH

[local]Redback(config-rsvp-lsp)#exit
```

Create RSVP LSP between RC2 and RC1:

```
[local]Redback(config-rsvp)#lsp Connect_RC2_To_RC1_BLUE

[local]Redback(config-rsvp-lsp)#ingress 10.12.240.4

[local]Redback(config-rsvp-lsp)#egress 14.2.2.3
```

Apply constraint to RSVP LSP:

```
[local]Redback(config-rsvp-lsp)#constraint BLUE_PATH

[local]Redback(config-rsvp-lsp)#exit
```

Glossary

MLPPP
Multilink PPP. An extension to PPP that allows a router to use more than one physical link for communication.

MP
Multilink PPP or Merge Point.

Merge Point: The point at which traffic exits the tail end router of a bypass RSVP LSP.

MLPPP: an extension to PPP that allows a router to use more than one physical link for communication.

PLR
Point of local repair.

Traffic enters the headend router of a bypass RSVP, which is called the point of local repair (PLR), and exits the tail end router of the bypass RSVP LSP, which is called the merge point (MP).