



Implementing RSVP for MPLS-TE

Resource Reservation Protocol (RSVP) is a signaling protocol that enables systems to request resource reservations from the network. RSVP processes protocol messages from other systems, processes resource requests from local clients, and generates protocol messages. As a result, resources are reserved for data flows on behalf of local and remote clients. RSVP creates, maintains, and deletes these resource reservations.


MPLS Traffic Engineering (MPLS-TE) learns the topology and resources available in a network and then maps traffic flows to particular paths based on resource requirements and network resources such as bandwidth. MPLS TE builds a unidirectional tunnel from a source to a destination in the form of a label switched path (LSP), which is then used to forward traffic. MPLS-TE uses RSVP to signal LSPs.

- [Setting up MPLS LSP Using RSVP, on page 1](#)
- [Overview of RSVP for MPLS-TE Features, on page 2](#)
- [Bandwidth Reservation Percentage, on page 2](#)
- [Caveats for Out-of-Sequence , on page 2](#)
- [Keychain Configuration For RSVP Authentication, on page 3](#)
- [Configuring RSVP for MPLS-TE, on page 3](#)
- [RSVP for MPLS-TE Features - Details, on page 14](#)
- [Additional References , on page 17](#)

Setting up MPLS LSP Using RSVP

The following figure shows how RSVP sets up an LSP from router R1 through router R4 that can be used for TE in an MPLS environment.

Figure 1: MPLS LSP Using RSVP



The LSP setup is initiated when the LSP head node sends path messages to the tail node. The Path messages reserve resources along the path to each node, and creates path states associated with the session on each node. When the tail node receives a path message, it sends a reservation (RESV) message with a label back to the

previous node. The reservation state in each router is considered as a soft state, which means that periodic PATH and RESV messages must be sent at each hop to maintain the state.

When the reservation message arrives at the previous node, it causes the reserved resources to be locked and forwarding entries are programmed with the MPLS label sent from the tail-end node. A new MPLS label is allocated and sent to the next node upstream. When the reservation message reaches the head node, the label is programmed and the MPLS data starts to flow along the path.

Overview of RSVP for MPLS-TE Features

This section provides an overview of the various features of RSVP for MPLS-TE.

RSVP is automatically enabled on interfaces on which MPLS-TE is configured. For MPLS-TE LSPs with bandwidth, the RSVP bandwidth has to be configured on the interfaces. There is no need to configure RSVP, if all MPLS-TE LSPs have zero bandwidth.

RSVP Graceful restart ensures high availability and allows RSVP TE enabled routers to recover RSVP state information from neighbors after a failure in the network.

RSVP requires that the path and reservation state that are set up during LSP signaling must be refreshed by periodically sending refresh messages. Refresh messages are used to synchronize the state between RSVP neighbors and to recover from lost RSVP messages. RSVP refresh reduction feature includes support for reliable messages which are transmitted rapidly when the messages are lost. Summary refresh messages contain information to refresh multiple states and reduces the number of messages required to refresh states.

RSVP messages can be authenticated to ensure that only trusted neighbors can set up reservations.

For detailed information about RSVP for MPLS-TE features, see the *RSVP for MPLS-TE Features- Details* topic.

Bandwidth Reservation Percentage

The Bandwidth Reservation Percentage allows the RSVP interface bandwidth to be specified as percentages of the link's physical bandwidth.

For more information on configuring RSVP bandwidth, refer the *Implementing MPLS Traffic Engineering* chapter in the *MPLS Configuration Guide for Cisco 8000 Series Routers*. For more information on commands for configuring RSVP bandwidth, refer the *RSVP Infrastructure Commands* chapter in the *MPLS Command Reference for Cisco 8000 Series Routers*.

Caveats for Out-of-Sequence

These caveats are listed for out-of-sequence:

- When RSVP messages traverse multiple interface types with different maximum transmission unit (MTU) values, some messages can become out-of-sequence if they are fragmented.
- Packets with some IP options may be reordered.
- Change in QoS configurations may lead to a transient reorder of packets.
- QoS policies can cause a reorder of packets in a steady state.

Because all out-of-sequence messages are dropped, the sender must retransmit them. Because RSVP state timeouts are generally long, out-of-sequence messages during a transient state do not lead to a state timeout.

Keychain Configuration For RSVP Authentication

Before implementing RSVP authentication, you must configure a keychain first. The name of the keychain must be the same as the one used in the keychain configuration. For more information about configuring keychains, see *System Security Configuration Guide for Cisco 8000 Series Routers*.



Note RSVP authentication supports only keyed-hash message authentication code (HMAC) type algorithms.

Configuring RSVP for MPLS-TE

RSVP requires coordination among several routers, establishing exchange of RSVP messages to set up LSPs. Depending on the requirements, RSVP requires some basic configuration described in the following topics:

Configuring Traffic Engineering Tunnel Bandwidth

To configure traffic engineering tunnel bandwidth, you must first set up TE tunnels and configure the reserved bandwidth per interface (there is no need to configure bandwidth for the data channel or the control channel).

Cisco IOS XR software supports two MPLS DS-TE modes: Prestandard and IETF.



Note For prestandard DS-TE you do not need to configure bandwidth for the data channel or the control channel. There is no other specific RSVP configuration required for this application. When no RSVP bandwidth is specified for a particular interface, you can specify zero bandwidth in the LSP setup if it is configured under RSVP interface configuration mode or MPLS-TE configuration mode.

Confirming DiffServ-TE Bandwidth

In RSVP global and subpools, reservable bandwidths are configured per interface to accommodate TE tunnels on the node. Available bandwidth from all configured bandwidth pools is advertised using IGP. RSVP signals the TE tunnel with appropriate bandwidth pool requirements.

Configuration Example

In this example, the **bandwidth** command sets the total reservable bandwidth, the maximum RSVP bandwidth available for a flow and the sub-pool bandwidth for the HundredGigE 0/0/0/3 interface.

```
Router# configure
Router(config)# rsvp interface HundredGigE0/0/0/3
Router(config-rsvp-if)# bandwidth 1000 mbps 100 mbps sub-pool 150 mbps
Router(config-rsvp-if)# commit
```

Global, Interface, and Neighbor Authentication Modes

You can configure global defaults for all authentication parameters including key, window size, and lifetime. These defaults are inherited when you configure authentication for each neighbor or interface. However, you can also configure these parameters individually on a neighbor or interface basis, in which case the global values (configured or default) are no longer inherited.



Note RSVP uses the following rules when choosing which authentication parameter to use when that parameter is configured at multiple levels (interface, neighbor, or global). RSVP goes from the most specific to least specific; that is, neighbor, interface, and global.

Global keys simplify the configuration and eliminate the chances of a key mismatch when receiving messages from multiple neighbors and multiple interfaces. However, global keys do not provide the best security.

Interface keys are used to secure specific interfaces between two RSVP neighbors. Because many of the RSVP messages are IP routed, there are many scenarios in which using interface keys are not recommended. If all keys on the interfaces are not the same, there is a risk of a key mismatch for the following reasons:

- When the RSVP graceful restart is enabled, RSVP hello messages are sent with a source IP address of the local router ID and a destination IP address of the neighbor router ID. Because multiple routes can exist between the two neighbors, the RSVP hello message can traverse to different interfaces.
- When the RSVP fast reroute (FRR) is active, the RSVP Path and Resv messages can traverse multiple interfaces.
- When Generalized Multiprotocol Label Switching (GMPLS) optical tunnels are configured, RSVP messages are exchanged with router IDs as the source and destination IP addresses. Since multiple control channels can exist between the two neighbors, the RSVP messages can traverse different interfaces.

Neighbor-based keys are particularly useful in a network in which some neighbors support RSVP authentication procedures and others do not. When the neighbor-based keys are configured for a particular neighbor, you are advised to configure all the neighbor's addresses and router IDs for RSVP authentication.

Configuring RSVP Message Authentication Globally

The RSVP authentication feature permits neighbors in an RSVP network to use a secure hash algorithm to authenticate all RSVP signaling messages digitally. The authentication is accomplished on a per-RSVP-hop basis using an RSVP integrity object in the RSVP message. The integrity object includes a key ID, a sequence number for messages, and keyed message digest.

You can globally configure the values of authentication parameters including the key-chain, time interval that RSVP maintains security associations with other trusted RSVP neighbors (life time) and maximum number of RSVP authenticated messages that can be received out of sequence (window size). These defaults are inherited for each neighbor or interface.

Configuration Example

In this example, authentication parameters are configured globally on a router. The authentication parameters including authentication key-chain, lifetime, and window size are configured. A valid key-chain should be configured before performing this task.

```
Router# configure
Router(config)# key chain mpls-keys
Router(config-mpls-keys)# commit
Router(config-mpls-keys)# exit
Router(config)# rsvp authentication
Router(config-rsvp-auth)# key-source key-chain mpls-keys
Router(config-rsvp-auth)# life-time 2000
Router(config-rsvp-auth)# window-size 33
```

Verification

Verify the configuration of authentication parameters using the following command.

```
Router# show rsvp authentication detail

RSVP Authentication Information:
  Source Address:          3.0.0.1
  Destination Address:    3.0.0.2
  Neighbour Address:      3.0.0.2
  Interface:              HundredGigabitEthernet 0/0/0/3
  Direction:              Send
  LifeTime:                2000 (sec)
  LifeTime left:          1305 (sec)
  KeyType:                 Static Global KeyChain
  Key Source:              mpls-keys
  Key Status:              No error
  KeyID:                   1
  Digest:                  HMAC MD5 (16)
  window-size:            33
Challenge:                 Not supported
  TX Sequence:             5023969459702858020 (0x45b8b99b00000124)
  Messages successfully authenticated: 245
  Messages failed authentication: 0
```

Configuring RSVP Authentication for an Interface

You can individually configure the values of RSVP authentication parameters including key-chain, life time, and window size on an interface. Interface specific authentication parameters are used to secure specific interfaces between two RSVP neighbors.

Configuration Example

This example configures authentication key-chain, life time for the security association, and window size on an interface. A valid key-chain should be already configured to use it as part of this task.

```
Router# configure
Router(config)# rsvp interface HundredGigE0/0/0/3
Router(config-rsvp-if)# authentication
Router(config-rsvp-if-auth)# key-source key-chain mpls-keys
Router(config-rsvp-if-auth)# life-time 2000
Router(config-rsvp-if-auth)# window-size 33
Router(config-rsvp-if-auth)# commit
```

Cisco IOS XR Release 7.11.1 introduces support to disable RSVP authentication.

Verification

Verify the configuration of authentication parameters using the following command.

```
Router# show rsvp authentication detail
```

```

RSVP Authentication Information:
  Source Address:      3.0.0.1
  Destination Address: 3.0.0.2
  Neighbour Address:   3.0.0.2
  Interface:           HundredGigabitEthernet 0/0/0/3
  Direction:           Send
  LifeTime:            2000 (sec)
  LifeTime left:       1305 (sec)
  KeyType:             Static Global KeyChain
  Key Source:          mpls-keys
  Key Status:          No error
  KeyID:               1
  Digest:              HMAC MD5 (16)
  window-size:         33
  Challenge:           Not supported
  TX Sequence:         5023969459702858020 (0x45b8b99b00000124)
  Messages successfully authenticated: 245
  Messages failed authentication:      0

```

Configuring RSVP Authentication on a Neighbor

You can individually configure the values of RSVP authentication parameters including key-chain, life time, and window size on a neighbor.

Configuration Example

This example configures the authentication key-chain, life time for the security association, and window size on a RSVP neighbor. A valid key-chain should be already configured to use it as part of this task.

```

Router# configure
Router(config)# rsvp neighbor 10.0.0.1 authentication
Router(config-rsvp-nbor-auth)# key-source key-chain mpls-keys
Router(config-rsvp-nbor-auth)# life-time 2000
Router(config-rsvp-nbor-auth)# window-size 33
Router(config-rsvp-nbor-auth)# commit

```

Verification

Verify the configuration of authentication parameters using the following command.

```

Router# show rsvp authentication detail

RSVP Authentication Information:
  Neighbour Address:      10.0.0.1
  Interface:              HundredGigabitEthernet 0/0/0/3
  Direction:              Send
  LifeTime:                2000 (sec)
  LifeTime left:          1205 (sec)
  KeyType:                 Static Global KeyChain
  Key Source:              mpls-keys
  Key Status:              No error
  KeyID:                   1
  Digest:                  HMAC MD5 (16)
  window-size:             33
  Challenge:               Not supported

```

RSVP Authentication by Using All the Modes: Example

The configuration example shows how to perform the following functions:

- Authenticates all RSVP messages.
- Authenticates the RSVP messages to or from 10.0.0.1 by setting the keychain for the **key-source key-chain** command to `nbr_keys`, SA lifetime is set to 3600, and the default window-size is set to 1.
- Authenticates the RSVP messages not to or from 10.0.0.1 by setting the keychain for the **key-source key-chain** command to `default_keys`, SA lifetime is set to 3600, and the window-size is set 64 when using GigabitEthernet0/6/0/0; otherwise, the default value of 1 is used.

```

rsvp
interface GigabitEthernet0/6/0/0
 authentication
   window-size 64
 !
 !
neighbor 10.0.0.1
 authentication
   key-source key-chain nbr_keys
 !
 !
 authentication
   key-source key-chain default_keys
   life-time 3600
 !
 !

```



Note If a keychain does not exist or contain valid keys, this is considered a configuration error because signaling fails. However, this can be intended to prevent signaling. For example, when using the above configuration, if the `nbr_keys` does not contain valid keys, all signaling with 10.0.0.1 fails.

Configuring Graceful Restart

RSVP graceful restart provides a mechanism to ensure high availability (HA), which allows detection and recovery from failure conditions for systems running Cisco IOS XR software, and ensures non-stop forwarding services. RSVP graceful restart is based on RSVP hello messages and allows RSVP TE enabled routers to recover RSVP state information from neighbors after a failure in the network. RSVP uses a Restart Cap object (RSVP RESTART) in hello messages in which restart and recovery times are specified to advertise the restart capability of a node. The neighboring node helps a restarting node by sending a Recover Label object to recover the forwarding state of the restarting node.

You can configure standard graceful restart which is based on node-id address based hello messages and also interface-based graceful restart which is interface-address based hello messages.

Configuration Example

In this example, RSVP-TE is already enabled on the router nodes on a network and graceful restart needs to be enabled on the router nodes for failure recovery. Graceful restart is configured globally to enabled node-id address based hello messages and also on a router interface to support interface-address based hello messages.

```

Router# configure
Router(config)# rsvp
Router(config-rsvp)# signalling graceful-restart
Router(config-rsvp)# interface HundredGigabitEthernet 0/0/0/3

```

```
Router(config-rsvp-if)# signalling hello graceful-restart interface-based
Router(config-rsvp-if)# commit
```

Verification

Use the following commands to verify that graceful restart is enabled.

```
Router# show rsvp graceful-restart

Graceful restart: enabled Number of global neighbors: 1
Local MPLS router id: 192.168.55.55
Restart time: 60 seconds Recovery time: 120 seconds
Recovery timer: Not running
Hello interval: 5000 milliseconds Maximum Hello miss-count: 4

Router# show rsvp graceful-restart neighbors detail

Neighbor: 192.168.77.77 Source: 192.168.55.55 (MPLS)
Hello instance for application MPLS
Hello State: UP (for 00:20:52)
Number of times communications with neighbor lost: 0
Reason: N/A
Recovery State: DONE
Number of Interface neighbors: 1
address: 192.168.55.0
Restart time: 120 seconds Recovery time: 120 seconds
Restart timer: Not running
Recovery timer: Not running
Hello interval: 5000 milliseconds Maximum allowed missed Hello messages: 4
```

Change the Restart-Time: Example

The example shows how to change the restart time that is advertised in hello messages sent to neighbor nodes.

```
rsvp signalling graceful-restart restart-time 200
```


Configuring Refresh Reduction

Table 1: Feature History Table

Feature Name	Release Information	Feature Description
Set Global RSVP Message Retransmission Interval	Release 24.1.1	<p>During Fast Reroute (FRR), an RSVP router sends multiple messages to neighbors. If a neighbor fails to acknowledge the messages due to an overload of RSVP message processing or a high frequency of failures, RSVP retransmits the messages, which can result in network congestion. You can now set a longer RSVP message retransmission interval to provide sufficient processing time for neighbors, reduce signaling overhead, and prevent network congestion.</p> <p>You can set this interval for all directly connected neighbors at once or remote neighbors connected through backup tunnels. Previously, you could only enable this option per interface.</p> <p>The feature introduces these changes:</p> <p>CLI: signalling refresh reduction reliable retransmit-time (RSVP configuration)</p> <p>YANG Data Model: Cisco-IOS-XR-ip-rsvp-cfg.yang (see GitHub, YANG Data Models Navigator)</p>

RSVP Refresh Reduction improves the reliability of Resource Reservation Protocol (RSVP) signaling to enhance network performance and message delivery and it is enabled by default. Refresh reduction is used with a neighbor only if the neighbor supports it. You can also disable refresh reduction on an interface if you want.

This feature ensures reliable delivery of RSVP messages when network traffic is disrupted. To ensure that its message is delivered to its neighbor, RSVP requests the neighbor to send an acknowledgment message by a given time duration. If it doesn't receive the acknowledgment, it resends the message and doubles its current wait time. After 5 attempts, RSVP stops retransmitting the message to the neighbor.

Configuration Example

The example shows how to configure the various parameters available for the refresh reduction feature.

The following parameters are configured to change their default values:

- refresh interval
- number of refresh messages a node can miss
- retransmit time
- acknowledgment hold time
- acknowledgment message size
- refresh message summary size

```

Router# configure
Router(config)# rsvp
Router(config-rsvp)# interface HundredGigabitEthernet 0/0/0/3
Router(config-rsvp-if)# signalling refresh interval 40
Router(config-rsvp-if)# signalling refresh missed 6
Router(config-rsvp-if)# signalling refresh reduction reliable retransmit-time 2000
Router(config-rsvp-if)# signalling refresh reduction reliable ack-hold-time 1000
Router(config-rsvp-if)# signalling refresh reduction reliable ack-max-size 1000
Router(config-rsvp-if)# signalling refresh reduction summary max-size 1500
Router(config-rsvp-if)# commit

```

Set Global RSVP Message Retransmission Interval

During FRR, RSVP reroutes tunnels with failed links. With the default RSVP message retransmission interval in use, when a large number of LSPs are rerouted simultaneously, an RSVP router sends a large number of messages to its neighbor – whether connected directly or through a backup tunnel – and this may overwhelm its processing load. While the neighbor is still processing messages, the RSVP router retransmits the earlier messages, further overwhelming it and causing network congestion.

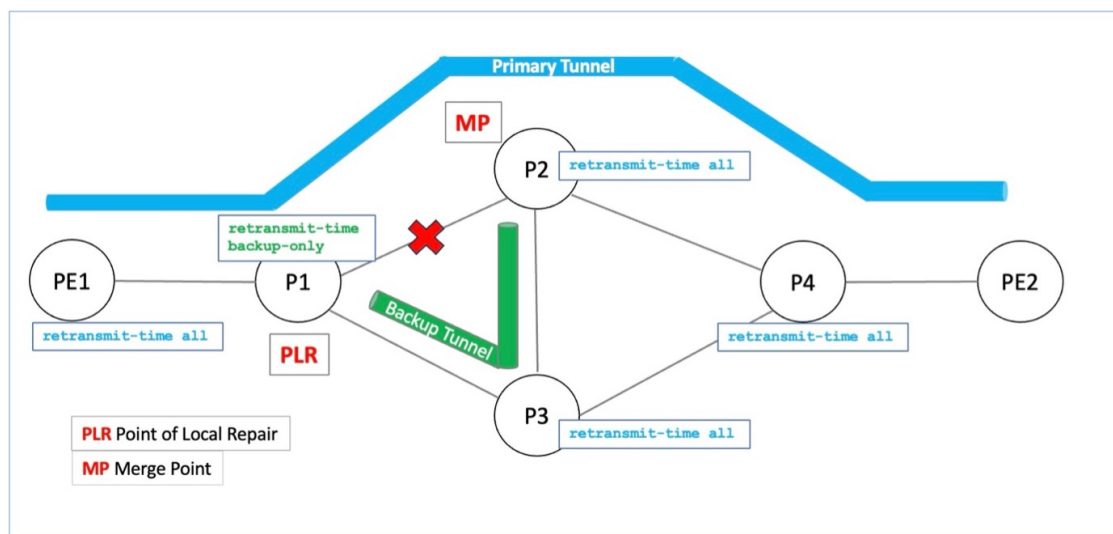
This feature allows you to configure a global timer with a longer retransmission interval so that retransmissions are relatively relaxed and provide more time for the neighbor to acknowledge messages.

To set the retransmission interval for messages sent over the backup tunnel during FRR, use the **backup-only** configuration which applies the RSVP-global configuration on backup tunnels. Another advantage of this configuration is that the interval is automatically applied to auto-backup tunnels, which are created dynamically.

Alternatively, to set the retransmission interval for all RSVP messages sent with reliable messaging (Path, Resv, PathTear, ResvTear, etc), use the **all** configuration which applies the RSVP-global configuration on all interfaces

Configuration Example for Set Global RSVP Message Retransmission Interval

Consider the following network topology.



Pointers:

- MPLS-TE, RSVP and Refresh-Reduction are enabled on all routers.

- The primary tunnel's path is PE1-P1-P2-P4-PE2.
PE1 is the headend router, P1-to-P4 are Provider routers, and PE2 is the tail-end router.
- P1 is a potential Point of Local Repair (PLR) and headend of the backup tunnel with the path P1-P3-P2.
PLR: A router that creates a backup tunnel to handle link failures in the primary tunnel.
- P2 is the Merge Point that merges the backup tunnel with the primary tunnel.

When P1 detects a link failure between P1-P2, it initiates FRR and reroutes the traffic through the backup tunnel.

In the following sections, the **backup-only** configuration is enabled on P1, and the **all** configuration is enabled on PE1, P2, P3 and P4.

PE1 Configuration

A retransmission interval of 4000 milliseconds is set for all RSVP-enabled PE1 interfaces:

```
PE1(config)# rsvp
PE1(config-rsvp)# signalling refresh reduction reliable retransmit-time all 4000
PE1(config-rsvp)# interface HundredGigE 0/0/0/0
PE1(config-rsvp-if)# bandwidth 200000
PE1(config-rsvp-if)# signalling refresh interval 180
PE1(config-rsvp-if)# signalling refresh reduction bundle-max-size 1400
PE1(config-rsvp-if)# commit
```

Similarly, enable configurations on P2, P3 and P4.

Verification

```
PE1# show rsvp interface detail

INTERFACE: HundredGigE0/0/0/0
..
      Retransmit: 4000ms. (Global)
```

RSVP message retransmission count:

```
PE1# show rsvp counters messages
HundredGigE0/0/0/0  Recv      Xmit
                   ..      ..
      Retransmit           8
..
```

RSVP message retransmission to its neighbors:

```
PE1# show rsvp neighbors detail

Global Neighbor: 10.1.1.12
Interface Neighbor: 10.2.2.10
..
      Retransmitted messages: 34
```

Running Configuration

```
rsvp
 interface HundredGigE0/0/0/0
   bandwidth 200000
   signalling refresh interval 180
```

```

    signalling refresh reduction bundle-max-size 1400
    !
    signalling refresh reduction reliable retransmit-time all 4000

```

P1 Configuration

A retransmission interval of 5000 milliseconds is set for P1's backup tunnel interface.



Note This configuration is also applicable for auto-backup tunnels.

```

P1(config)# rsvp
P1(config-rsvp)# signalling refresh reduction reliable retransmit-time backup-only 5000
P1(config-rsvp)# interface hundredGigE 0/0/0/0
P1(config-rsvp-if)# bandwidth 200000
P1(config-rsvp-if)# signalling refresh interval 180
P1(config-rsvp-if)# signalling refresh reduction bundle-max-size 1400
P1(config-rsvp-if)# commit

```

Verification

```

P1# show rsvp interface detail

INTERFACE: HundredGigE0/0/0/0
..
                Retransmit: 5000ms. (Global)

```

Running Configuration

```

rsvp
 interface HundredGigE0/0/0/0
   bandwidth 200000
   signalling refresh interval 180
   signalling refresh reduction bundle-max-size 1400
   !
   signalling refresh reduction reliable retransmit-time backup-only 5000

```

Change the Hello Interval: Example

The example shows how to change the interval at which RSVP graceful restart hello messages are sent per neighbor, and change the number of hellos missed before the neighbor is declared down.

```

rsvp signalling hello graceful-restart refresh interval 4000
rsvp signalling hello graceful-restart refresh misses 4

```

Disable Refresh Reduction: Example

If the peer node does not support refresh reduction, or for any other reason you want to disable refresh reduction on an interface, the example shows how to disable refresh reduction on that interface.

```

Router(config)# rsvp
Router(config-rsvp)# interface hundredGigE 0/0/0/0
Router(config-rsvp-if)# signalling refresh reduction disable

```

RSVP Prefix Filtering

Two procedures are provided to show how RSVP Prefix Filtering is associated:

Configuring ACL Based Prefix Filtering

You can configure extended access lists (ACLs) to forward, drop, or perform normal processing on RSVP router-alert (RA) packets. For each incoming RSVP RA packet, RSVP inspects the IP header and attempts to match the source or destination IP addresses with a prefix configured in an extended ACL. If there is no explicit permit or explicit deny, the ACL infrastructure returns an implicit deny by default. By default, RSVP processes the packet if the ACL match yields an implicit (default) deny.

Configuration Example

This example configures ACL based prefix filtering on RSVP RA packets. When RSVP receives a RA packet from source address 10.0.0.1 it is forwarded and packets destined to the IP address 172.16.0.1 are dropped.

```
Router# configure
Router(config)# ipv4 access-list rsvpac1
Router(config-ipv4-acl)# 10 permit ip host 10.0.0.1 any
Router(config-ipv4-acl)# 20 deny ip any host 172.16.0.1

Router# configure
Router(config)# rsvp
Router(config-rsvp)# signalling prefix-filtering access-list rsvp-acl
Router(config-rsvp)# commit
```

Verification

Verify the configuration of ACL based prefix filtering

```
Router# show rsvp counters prefix-filtering access-list rsvp-acl
```

ACL:rsvp-acl	Forward	Local	Drop	Total
Path	0	0	0	0
PathTear	0	0	0	0
ResvConfirm	0	0	0	0
Total	0	0	0	0

Configuring RSVP Packet Dropping

You can configure extended access lists (ACLs) to forward, drop, or perform normal processing on RSVP router-alert (RA) packets. By default, RSVP processes the RA packets even if the ACL match yields an implicit deny. You can configure RSVP to drop RA packets when the ACL matches results in an implicit deny.

Configuration Example

This example configures ACL based prefix filtering on RSVP RA packets. When RSVP receives a RA packet from source address 10.0.0.1 it is forwarded and packets destined to the IP address 172.16.0.1 are dropped. RA packets are dropped if the ACL matches results in an implicit deny.

```
Router# configure
Router(config)# ipv4 access-list rsvpac1
Router(config-ipv4-acl)# 10 permit ip host 10.0.0.1 any
Router(config-ipv4-acl)# 20 deny ip any host 172.16.0.1
Router(config-ipv4-acl)# exit
```

```
Router(config)# rsvp
Router(config-rsvp)# signalling prefix-filtering default-deny-action drop
Router(config-rsvp)# commit
```

Verification

Verify the configuration of RSVP packet drop using the following command.

```
Router# show rsvp counters prefix-filtering access-list rsvpac1
```

ACL: rsvpac1	Forward	Local	Drop	Total
Path	4	1	0	5
PathTear	0	0	0	0
ResvConfirm	0	0	0	0
Total	4	1	0	5

Enabling RSVP Traps

By implementing the RSVP MIB, you can use SNMP to access objects belonging to RSVP. You can also specify two traps (NewFlow and LostFlow) which are triggered when a new flow is created or deleted. RSVP MIBs are automatically enabled when you turn on RSVP, but you need to enable RSVP traps.

Configuration Example

This example shows how to enable RSVP MIB traps when a flow is deleted or created and also how to enable both the traps.

```
Router# configure
Router(config)# snmp-server traps rsvp lost-flow
Router(config)# snmp-server traps rsvp new-flow
Router(config)# snmp-server traps rsvp all
Router(config)# commit
```

Eliminating Security Associations for RSVP Authentication

To eliminate RSVP authentication SA's, use the **clear rsvp authentication** command. To eliminate RSVP counters for each SA, use the **clear rsvp counters authentication** command.

RSVP for MPLS-TE Features - Details

RSVP Graceful Restart Operation

RSVP graceful restart is based on RSVP hello messages. Hello messages are exchanged between the router and its neighbor nodes. Each neighbor node can autonomously issue a hello message containing a hello request object. A receiver that supports the hello extension replies with a hello message containing a hello acknowledgment (ACK) object. If the sending node supports state recovery, a Restart Cap object that indicates a node's restart capability is also carried in the hello messages. In the Restart Cap object, the restart time and the recovery time is specified. The restart time is the time after a loss in Hello messages within which RSVP hello session can be re-established. The recovery time is the time that the sender waits for the recipient to re-synchronize states after the re-establishment of hello messages.

For graceful restart, the hello messages are sent with an IP Time to Live (TTL) of 64. This is because the destination of the hello messages can be multiple hops away. If graceful restart is enabled, hello messages

(containing the restart cap object) are sent to an RSVP neighbor when RSVP states are shared with that neighbor. If restart cap objects are sent to an RSVP neighbor and the neighbor replies with hello messages containing the restart cap object, the neighbor is considered to be graceful restart capable. If the neighbor does not reply with hello messages or replies with hello messages that do not contain the restart cap object, RSVP backs off sending hellos to that neighbor. If a hello Request message is received from an unknown neighbor, no hello ACK is sent back.

RSVP Authentication

Network administrators need the ability to establish a security domain to control the set of systems that initiates RSVP requests. The RSVP authentication feature permits neighbors in an RSVP network to use a secure hash to sign all RSVP signaling messages digitally, thus allowing the receiver of an RSVP message to verify the sender of the message without relying solely on the sender's IP address.

The signature is accomplished on a per-RSVP-hop basis with an RSVP integrity object in the RSVP message as defined in RFC 2747. The integrity object includes a key ID, a sequence number for messages, and keyed message digest. This method provides protection against forgery or message modification. However, the receiver must know the security key used by the sender to validate the digital signature in the received RSVP message. Network administrators manually configure a common key for each RSVP neighbor on the shared network. The sending and receiving systems maintain a security association for each authentication key that they share. For detailed information about different security association parameters, see the **Security Association Parameters** table.

You can configure global defaults for all authentication parameters including key, window size, and lifetime. These defaults are inherited when you configure authentication for each neighbor or interface. However, you can also configure these parameters individually on a neighbor or interface basis, in which case the global values (configured or default) are no longer inherited.

Interface and neighbor interface modes unless explicitly configured, inherit the parameters from global configuration mode as follows:

- Window-size is set to 1.
- Lifetime is set to 1800.
- key-source key-chain command is set to none or disabled.

The following situations explain how to choose between global, interface, or neighbor configuration modes:

- Global configuration mode is optimal when a router belongs to a single security domain (for example, part of a set of provider core routers). A single common key set is expected to be used to authenticate all RSVP messages.
- Interface, or neighbor configuration mode, is optimal when a router belongs to more than one security domain. For example, a provider router is adjacent to the provider edge (PE), or a PE is adjacent to an edge device. Different keys can be used but not shared.

A security association (SA) is a collection of information that is required to maintain secure communications with a peer. The following table lists the main parameters that defines a security association

Table 2: Security Association Parameters

Security Association Parameter	Description
src	IP address of the sender.

Security Association Parameter	Description
dst	IP address of the final destination.
interface	Interface of the security association.
direction	Send or receive type of the security association.
Lifetime	Expiration timer value that is used to collect unused security association data.
Sequence Number	Last sequence number that was either sent or accepted (dependent of the direction type).
key-source	Source of keys for the configurable parameter.
keyID	Key number (returned from the key-source) that was last used.
Window Size	Specifies the maximum number of authenticated messages that can be received out of order.
Window	Specifies the last <i>window size</i> value sequence number that is received or accepted.

MPLS-TE LSP OOR

The MPLS-TE LSP OOR function adds capability for the RSVP-TE control plane to track the LSP scale of transit routers, so that it can take a specific set of (pre-configured) actions when threshold limits are crossed, and inform other routers in the network. MPLS-TE keeps track of the number of transit LSPs set up through the router. The limits do not apply to ingress and egress LSP routers since they are driven by explicit configuration. In other words, the configuration determines how many egress or ingress LSPs a router has. For midpoint routers, the number is a function of the topology, the links metrics, and links' bandwidth.

State Transition Triggers - The LSP OOR state transition is triggered by checking the total transit LSP count and the unprotected count. If either count crosses the threshold, the state transition is triggered. If both counts cross the limit, the more critical state is chosen. Each limit will have a value for the *Yellow* threshold and a value for the *Red* threshold. When these thresholds are crossed, the configured MPLS-TE LSP OOR actions take effect. Similarly, the transition to *Green* state occurs when the LSP numbers drop.

LSP OOR State Dampening - The reason for LSP OOR State Dampening is that the number of accepted LSPs would be at the threshold and once an LSP is deleted, the state goes back from Red to Yellow, and a new LSP is setup and the state goes back to Red.

The solution is to introduce dampening when there is a state transition from Red to Yellow or from Yellow to Green. Whenever the transit number of LSPs crosses down a threshold, a timer is started for 10 seconds. After the timer expires, the new state is computed and moved to it. The timer is stopped if the transit number threshold is crossed (up) again. The transition from a state to a more severe state is not dampened.

Low and High Priority LSPs - When the LSP OOR is in yellow or red state, new high priority LSPs will not preempt low priority LSPs. Preemption can still occur but only for bandwidth reasons. In other words, if the router is in Red state where one of the actions is to reject any new LSP, the new high-priority LSPs are rejected even if there is an established low-priority LSP. The low-priority LSP is not removed to make room for the high-priority one.

Configuration Limit - Setting the configured limit to a value that is smaller than the current number of LSPs will trigger state transition but will not cause existing LSPs to be deleted or preempted. Setting the configured limit to a value that is larger than the current number of LSPs takes the node out of LSP OOR state. When an LSP cannot be admitted due to LSP OOR, the LSRs send Path Error messages to the LERs.

Event Logging - This is generated when the system transitions across OOR states, such as a resource change into an *yellow* or *red* state. Reporting level for *Red* is critical (1), and for *yellow* is warning (4). The following example shows that the count has crossed the threshold of 5000.

```
RP/0/RP1/CPU0:May 15 17:05:48 PDT: te_control[1034]: %ROUTING-MPLS_TE-4-LSP_OOR :

Transit LSP resources changed to Yellow.
Total transit: configured threshold 5000; actual count 5001;
Unprotected transit: configured threshold 4294967295; actual count 0
```

When the resource comes out of OOR, it will report as *green*.

Configuration Example

```
mpls traffic-eng
 lsp-oor
  green
    action accept reopt-lsp
    action flood available-bw 20
    recovery-duration
    action admit lsp-min-bw X -- > (in kbps, a lower limit than yellow and red state)

  yellow
    transit-all threshold 75000
    action accept reopt-lsp
    action flood available-bw 0
    action admit lsp-min-bw Y

  red
    transit-all threshold 90000
    action flood available-bw 0
    action admit lsp-min-bw Z
```

The LSP OOR threshold values are set to yellow as 75000 and red as 90000. When these thresholds are crossed, corresponding actions are applied to all the TE interfaces.



Note The default values of the above thresholds are infinite.

When the LSP OOR *yellow* state is reached, the **accept reopt-lsp** action, **flood available-bw 0** action and **admit lsp-min-bw** actions are activated. This allows headend routers to reoptimize existing LSPs through, but doesn't allow new LSPs to get established. Also, MPLS-TE advertises zero bandwidth out of all interfaces, making this transit router less preferable for new LSPs. To handle a sudden burst of new LSPs that get signaled, the **action admit lsp-min-bw** function ensures only a small number of high bandwidth LSPs get provisioned through the affected router. When the red threshold state is crossed, the **flood available-bw 0** and **admit lsp-min-bw** actions prevent any additional or reoptimized transit LSPs from getting set up through the affected router.

Additional References

For additional information related to RSVP, refer to the following references:

Related Documents

Related Topic	Document Title
RSVP Infrastructure Commands	<i>RSVP Infrastructure Commands</i> module in <i>MPLS Command Reference for Cisco 8000 Series Routers</i> .
MPLS Traffic Engineering Commands	<i>MPLS Traffic Engineering commands</i> module in <i>MPLS Command Reference for Cisco 8000 Series Routers</i> .

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport