



RSVP Fast Local Repair

The RSVP Fast Local Repair feature provides quick adaptation to routing changes occurring in global and Virtual Routing and Forwarding (VRF) domains, without the overhead of the refresh period to guarantee the quality of service (QoS) for data flows. With fast local repair (FLR), Resource Reservation Protocol (RSVP) speeds up its response to routing changes from 30 seconds to a few seconds.

- [Finding Feature Information, on page 1](#)
- [Prerequisites for RSVP FLR, on page 1](#)
- [Restrictions for RSVP FLR, on page 2](#)
- [Information About RSVP FLR, on page 2](#)
- [How to Configure RSVP FLR, on page 3](#)
- [Configuration Examples for RSVP FLR, on page 7](#)
- [Additional References, on page 11](#)
- [Feature Information for RSVP FLR, on page 12](#)
- [Glossary, on page 12](#)

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest caveats and feature information, see [Bug Search Tool](#) and the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the feature information table.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for RSVP FLR

You must configure RSVP on one or more interfaces on at least two neighboring devices that share a link within the network.

Restrictions for RSVP FLR

- RSVP FLR applies only when RSVP is used to set up resource reservations for IPv4 unicast flows; IPv4 multicast flows are not supported.
- RSVP FLR does not apply to traffic engineering (TE) tunnels and, therefore, does not affect TE sessions.
- RSVP FLR does not support message bundling.

Information About RSVP FLR

Feature Overview of RSVP FLR


RSVP FLR provides for dynamic adaptation when routing changes occur in global or VRF routing domains. When a route changes, the next PATH and RESV message refreshes establish path and reservation states along the new route. Depending on the configured refresh interval, this reroute happens in tens of seconds. However, during this time, the QoS of flows is not guaranteed because congestion may occur while data packets travel over links where reservations are not yet in place.

In order to provide faster adaptation to routing changes, without the overhead of a refresh period, RSVP registers with the Routing Information Base (RIB) and receives notifications when routes change, thereby triggering state refreshes for the affected destinations. These triggered refreshes use the new route information and, as a result, install reservations over the new path.

When routes change, RSVP has to reroute all affected paths and reservations. Without FLR, the reroute happens when refresh timers expire for the path states. With real-time applications such as VoIP and video on demand (VoD), the requirement changes and the reroute must happen, within three seconds from the triggering event such as link down or link up.

The figure below illustrates the FLR process.

Figure 1: Overview of RSVP FLR



Initial RSVP states are installed for an IPv4 unicast flow over Routers A, B, C, D, and E. Router A is the source or headend, and Router E is the destination or tailend. The data packets are destined to an address of Router E. Assume that a route change occurs, and the new path taken by the data packets is from Router A to Router B to Router F to Router D to Router E; therefore, the old and new paths differ on the segments between Routers B and D. The Router B to Router C to Router D segment is the old segment, and the Router B to Router F to Router D segment is the new segment.

A route may change because of a link or node failure, or if a better path becomes available.

RSVP at Router B detects that the route change affects the RSVP flow and initiates the FLR procedure. The node that initiates an FLR repair procedure, Router B in the figure above, is the point of local repair (PLR). The node where the new and old segments meet, Router D in the figure above, is the merge point (MP). The interfaces at the PLR and the MP that are part of the old segment are the old interfaces, and the interfaces that are part of the new segment are the new interfaces.

If a route has changed because of a failure, the PLR may not be the node that detects the failure. For example, it is possible that the link from Router C to Router D fails, and although Router C detects the failure, the route change at Router B is the trigger for the FLR procedure. Router C, in this case, is also referred to as the node that detects the failure.

The support for FLR in VRF domains means that RSVP can get a route change notification, even if there is a route change in any VRF domains, because RSVP FLR was previously supported only in the global routing domain.

Benefits of RSVP FLR

Faster Response Time to Routing Changes

FLR reduces the time that it takes for RSVP to determine that a physical link has gone down and that the data packets have been rerouted. Without FLR, RSVP may not recognize the link failure for 30 seconds when all of the sessions are impacted by having too much traffic for the available bandwidth. With FLR, this time can be significantly reduced to a few seconds.

After detecting the failure, RSVP recomputes the admission control across the new link. If the rerouted traffic fits on the new link, RSVP reserves the bandwidth and guarantees the QoS of the new traffic.

If admission control fails on the new route, RSVP does not explicate tear down the flow, but instead sends a RESVERROR message toward the receiver. If a proxy receiver is running, then RSVP sends a PATHERROR message toward the headend, in response to the RESVERROR message, indicating the admission failure. In both cases, with and without a proxy receiver, the application tears down the failed session either at the headend or at the final destination.

Until this happens, the data packets belonging to this session still flow over the rerouted segment although admission has failed and QoS is affected.

The support of FLR in VRF domains means that if there is a route change in any routing domain, RSVP can use FLR to adapt to the routing change, because RSVP FLR was previously supported only in the global routing domain.

How to Configure RSVP FLR

You can configure the RSVP FLR parameters in any order that you want.

Configuring the RSVP FLR Wait Time

SUMMARY STEPS

1. `enable`

2. **configure terminal**
3. **interface** *type slot / subslot / port*
4. **ip rsvp bandwidth** [*interface-kbps* [*single-flow-kbps*[**bc1** *kbps* | **sub-pool** *kbps*]]] **percent** *percent-bandwidth* [*single-flow-kbps*]
5. **ip rsvp signalling fast-local-repair wait-time** *interval*
6. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: <pre>Router> enable</pre>	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	configure terminal Example: <pre>Router# configure terminal</pre>	Enters global configuration mode.
Step 3	interface <i>type slot / subslot / port</i> Example: <pre>Router(config)# interface gigabitEthernet 0/0/0</pre>	Configures the interface type and enters interface configuration mode.
Step 4	ip rsvp bandwidth [<i>interface-kbps</i> [<i>single-flow-kbps</i> [bc1 <i>kbps</i> sub-pool <i>kbps</i>]]] percent <i>percent-bandwidth</i> [<i>single-flow-kbps</i>] Example: <pre>Router(config-if)# ip rsvp bandwidth 7500 7500</pre>	Enables RSVP on an interface. • The optional <i>interface-kbps</i> and <i>single-flow-kbps</i> arguments specify the amount of bandwidth that can be allocated by RSVP flows or to a single flow, respectively. Values are from 1 to 10000000. • The optional sub-pool and <i>kbps</i> keyword and argument specify subpool traffic and the amount of bandwidth that can be allocated by RSVP flows. Values are from 1 to 10000000. Note Repeat this command for each interface on which you want to enable RSVP.
Step 5	ip rsvp signalling fast-local-repair wait-time <i>interval</i> Example: <pre>Router(config-if)# ip rsvp signalling fast-local-repair wait-time 100</pre>	Configures the delay that RSVP uses before starting an FLR procedure. • Values for the <i>interval</i> argument are 1 to 2500 milliseconds (ms); the default is 0.
Step 6	end Example:	(Optional) Returns to privileged EXEC mode.

	Command or Action	Purpose
	<code>Router(config-if) # end</code>	

Configuring the RSVP FLR Repair Rate

SUMMARY STEPS

1. `enable`
2. `configure terminal`
3. `ip rsvp signalling fast-local-repair rate rate`
4. `exit`

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: <code>Router> enable</code>	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	configure terminal Example: <code>Router# configure terminal</code>	Enters global configuration mode.
Step 3	ip rsvp signalling fast-local-repair rate rate Example: <code>Router(config)# ip rsvp signalling fast-local-repair rate 100</code>	Configures the repair rate that RSVP uses for an FLR procedure. • Values for the <i>rate</i> argument are 1 to 2500 messages per second; the default is 400.
Step 4	exit Example: <code>Router(config)# exit</code>	(Optional) Returns to privileged EXEC mode.

Configuring the RSVP FLR Notifications

Perform this task to configure the number of RSVP FLR notifications.

SUMMARY STEPS

1. `enable`
2. `configure terminal`
3. `ip rsvp signalling fast-local-repair notifications number`

4. exit

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: Router> enable	Enables privileged EXEC mode. Enter your password if prompted.
Step 2	configure terminal Example: Router# configure terminal	Enters global configuration mode.
Step 3	ip rsvp signalling fast-local-repair notifications <i>number</i> Example: Router(config)# ip rsvp signalling fast-local-repair notifications 100	Configures the number of per flow notifications that RSVP processes during an FLR procedure before it suspends. Values for the <i>number</i> argument are 10 to 10000; the default is 1000.
Step 4	exit Example: Router(config)# exit	(Optional) Returns to privileged EXEC mode.

Verifying the RSVP FLR Configuration

Perform this task to verify the RSVP FLR configuration. You can use these commands in any order.

**Note**

You can use the following **show** commands in user EXEC or privileged EXEC mode.

SUMMARY STEPS

1. enable
2. show ip rsvp signalling fast-local-repair [statistics [detail]]
3. show ip rsvp interface [detail] [interface-type interface-number]
4. show ip rsvp
5. show ip rsvp sender [detail] [filter [destination ip-address | hostname] [dst-port port-number] [source ip-address | hostname] [src-port port-number]]
6. exit

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: <pre>Router> enable</pre>	(Optional) Enables privileged EXEC mode. Enter your password if prompted. Note Omit this step if you are using the show commands in user EXEC mode.
Step 2	show ip rsvp signalling fast-local-repair [statistics [detail]] Example: <pre>Router# show ip rsvp signalling fast-local-repair statistics detail</pre>	Displays FLR-specific information that RSVP maintains. The optional statistics and detail keywords display additional information about the FLR parameters.
Step 3	show ip rsvp interface [detail] [interface-type interface-number] Example: <pre>Router# show ip rsvp interface gigabitEthernet 0/0/0</pre>	Displays RSVP-related information. The optional detail keyword displays additional information including FLR parameters.
Step 4	show ip rsvp Example: <pre>Router# show ip rsvp</pre>	Displays general RSVP-related information.
Step 5	show ip rsvp sender [detail] [filter [destination ip-address hostname] [dst-port port-number] [source ip-address hostname] [src-port port-number]] Example: <pre>Router# show ip rsvp sender detail</pre>	Displays RSVP PATH-related sender information currently in the database. The optional detail keyword displays additional output including the FLR parameters.
Step 6	exit Example: <pre>Router# exit</pre>	(Optional) Exits privileged EXEC mode and returns to user EXEC mode.

Configuration Examples for RSVP FLR

Example Configuring RSVP FLR

The configuration options for RSVP FLR are the following:

- Wait time
- Number of notifications
- Repair rate



Note You can configure these options in any order.

Configuring the Wait Time

The following example configures gigabitEthernet interface 0/0/0 with a bandwidth of 200 kbps and a wait time of 1000 ms:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface gigabitEthernet 0/0/0
Router(config-if)# ip rsvp bandwidth 200
Router(config-if)# ip rsvp signalling fast-local-repair wait-time 1000
Router(config-if)# end
```

Configuring the Number of Notifications

The following example configures the number of flows that are repaired before suspending to 100:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# ip rsvp signalling fast-local-repair notifications 100
Router(config)# exit
```

Configuring the Repair Rate

The following example configures a repair rate of 100 messages per second:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# ip rsvp signalling fast-local-repair rate 100
Router(config)# exit
```

Example Verifying the RSVP FLR Configuration

Verifying the Details for FLR Procedures

The following example displays detailed information about FLR procedures:

```
Router# show ip rsvp signalling fast-local-repair statistics detail
Fast Local Repair: enabled
  Max repair rate (paths/sec): 10
  Max processed (paths/run): 10
FLR Statistics:
  FLR 1: DONE
    Start Time: 05:18:54 IST Mon Nov 5 2007
    Number of PSBs repaired: 2
```

```

Used Repair Rate (msgs/sec):      10
RIB notification processing time: 0(us).
Time of last PSB refresh:         5025(ms).
Time of last Resv received:       6086(ms).
Time of last Perr received:       0(us).
Suspend count: 0
FLR Pacing Unit: 100 msec.
Affected neighbors:
  Nbr Address      Interface      Relative Delay Values (msec)      VRF
  10.1.2.12        Et0/3          [5000 ,..., 5000 ]              vrf1
  10.1.2.12        Et1/3          [5000 ,..., 5000 ]              vrf2

```

Verifying Configuration Details for a Specific Interface

The following example from the **show ip rsvp interface detail** command displays detailed information, including FLR, for the gigabitEthernet 0/0/0 interface:

```

Router# show ip rsvp interface detail gigabitethernet 0/0/0
Et1/0:
  RSVP: Enabled
  Interface State: Up
  Bandwidth:
    Curr allocated: 9K bits/sec
    Max. allowed (total): 300K bits/sec
    Max. allowed (per flow): 300K bits/sec
    Max. allowed for LSP tunnels using sub-pools (pool 1): 0 bits/sec
    Set aside by policy (total): 0 bits/sec
  Traffic Control:
    RSVP Data Packet Classification is ON via CEF callbacks
  Signalling:
    DSCP value used in RSVP msgs: 0x30
    Number of refresh intervals to enforce blockade state: 4
  FLR Wait Time (IPv4 flows):
    Repair is delayed by 1000 msec.
  Authentication: disabled
    Key chain: <none>
    Type:      md5
    Window size: 1
    Challenge: disabled
  Hello Extension:
    State: Disabled

```

Verifying Configuration Details Before During and After an FLR Procedure

The following is sample output from the **showiprsvpsenderdetail** command before an FLR procedure has occurred:

```

Router# show ip rsvp sender detail
PATH:
  Destination 192.168.101.21, Protocol_Id 17, Don't Police , DstPort 1
  Sender address: 10.10.10.10, port: 1
  Path refreshes:
    arriving: from PHOP 172.3.31.34 on Et0/0 every 30000 msecs
  Traffic params - Rate: 9K bits/sec, Max. burst: 9K bytes
    Min Policed Unit: 0 bytes, Max Pkt Size 2147483647 bytes
  Path ID handle: 01000401.
  Incoming policy: Accepted. Policy source(s): Default
  Status:
  Output on gigabitEthernet 0/0/0. Policy status: Forwarding. Handle: 02000400
    Policy source(s): Default
  Path FLR: Never repaired

```

The following is sample output from the **showiprvpsenderdetail** command at the PLR during an FLR procedure:

```
Router# show ip rsvp sender detail
PATH:
  Destination 192.168.101.21, Protocol_Id 17, Don't Police , DstPort 1
  Sender address: 10.10.10.10, port: 1
  Path refreshes:
    arriving: from PHOP 172.16.31.34 on Et0/0 every 30000 msec
  Traffic params - Rate: 9K bits/sec, Max. burst: 9K bytes
    Min Policed Unit: 0 bytes, Max Pkt Size 2147483647 bytes
  Path ID handle: 01000401.
  Incoming policy: Accepted. Policy source(s): Default
  Status:
  Path FLR: PSB is currently being repaired...try later
  PLR - Old Segments: 1
    Output on Ethernet1/0, nhop 172.5.36.34
  Time before expiry: 2 refreshes
  Policy status: Forwarding. Handle: 02000400
    Policy source(s): Default
```

The following is sample output from the **showiprvpsenderdetail** command at the MP during an FLR procedure:

```
Router# show ip rsvp sender detail
PATH:
  Destination 192.168.101.21, Protocol_Id 17, Don't Police , DstPort 1
  Sender address: 10.10.10.10, port: 1
  Path refreshes:
    arriving: from PHOP 172.16.37.35 on Et1/0 every 30000 msec
  Traffic params - Rate: 9K bits/sec, Max. burst: 9K bytes
    Min Policed Unit: 0 bytes, Max Pkt Size 2147483647 bytes
  Path ID handle: 09000406.
  Incoming policy: Accepted. Policy source(s): Default
  Status: Proxy-terminated
  Path FLR: Never repaired
  MP - Old Segments: 1
    Input on Serial2/0, phop 172.16.36.35
  Time before expiry: 9 refreshes
```

The following is sample output from the **showiprvpsenderdetail** command at the PLR after an FLR procedure:

```
Router# show ip rsvp sender detail
PATH:
  Destination 192.168.101.21, Protocol_Id 17, Don't Police , DstPort 1
  Sender address: 10.10.10.10, port: 1
  Path refreshes:
    arriving: from PHOP 172.16.31.34 on Et0/0 every 30000 msec
  Traffic params - Rate: 9K bits/sec, Max. burst: 9K bytes
    Min Policed Unit: 0 bytes, Max Pkt Size 2147483647 bytes
  Path ID handle: 05000401.
  Incoming policy: Accepted. Policy source(s): Default
  Status:
  Output on Serial3/0. Policy status: Forwarding. Handle: 3B000406
    Policy source(s): Default
  Path FLR: Started 12:56:16 EST Thu Nov 16 2006, PSB repaired 532(ms) after.
    Resv/Perr: Received 992(ms) after.
```

Additional References

The following sections provide references related to the Control Plane DSCP Support for RSVP feature.

Related Documents

Related Topic	Document Title
Cisco IOS commands	Cisco IOS Master Commands List, All Releases
RSVP Commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>Cisco IOS Quality of Service Solutions Command Reference</i>
Quality of service overview	"Quality of Service Overview" module

Standards

Standard	Title
None	--

MIBs

MIB	MIBs Link
RFC 2206 (RSVP Management Information Base using SMIv2)	To locate and download MIBs for selected platforms, software releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

RFCs

RFC	Title
RFC 2205	<i>Resource Reservation Protocol</i>

Technical Assistance

Description	Link
The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html

Feature Information for RSVP FLR

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 1: Feature Information for RSVP FLR

Feature Name	Releases	Feature Information
RSVP Fast Local Repair	Cisco IOS XE Release 2.6 Cisco IOS XE Release 3.8S	The RSVP Fast Local Repair feature provides quick adaptation to routing changes without the overhead of the refresh period to guarantee QoS for data flows. With FLR, RSVP speeds up its response to routing changes from 30 seconds to a few seconds. The following commands were introduced or modified: clear ip rsvp signalling fast-local-repair statistics, ip rsvp signalling fast-local-repair notifications, ip rsvp signalling fast-local-repair rate, ip rsvp signalling fast-local-repair wait-time, show ip rsvp, show ip rsvp interface, show ip rsvp sender, show ip rsvp signalling fast-local-repair. In Cisco IOS XE Release 3.8S, support was added for the Cisco ASR 903 Router.

Glossary

admission control --The process by which an RSVP reservation is accepted or rejected on the basis of end-to-end available network resources.

bandwidth --The difference between the highest and lowest frequencies available for network signals. The term is also used to describe the rated throughput capacity of a given network medium or protocol.

message pacing-- A system for managing volume and timing that permits messages from multiple sources to be spaced apart over time. RSVP message pacing maintains, on an outgoing basis, a count of the messages that it has been forced to drop because the output queue for the interface used for the message pacing was full.

MP --merge point. The node where the new and old FLR segments meet.

PLR --point of local repair. The node that initiates an FLR procedure.

QoS --quality of service. A measure of performance for a transmission system that reflects its transmission quality and service availability.

RSVP --Resource Reservation Protocol. A protocol that supports the reservation of resources across an IP network. Applications running on IP end systems can use RSVP to indicate to other nodes the nature (bandwidth, jitter, maximum burst, and so on) of the packet streams that they want to receive.

VRF--virtual routing and forwarding. VRF is a VPN routing and forwarding instance. A VRF consists of an IP routing table, a derived forwarding table, a set of interfaces that use the forwarding table, and a set of rules and routing protocols that determine what goes into the forwarding table. In general, a VRF includes the routing information that defines a customer VPN site that is attached to a PE router.

