

Configuring RSVP

This chapter describes how to configure Resource Reservation Protocol (RSVP), which is an IP service that allows end systems or hosts on either side of a router network to establish a reserved-bandwidth path between them to predetermine and ensure QoS for their data transmission.

For a complete description of the RSVP commands in this chapter, refer to the *Quality of Service Solutions Command Reference*. To locate documentation of other commands that appear in this chapter, use the command reference master index, or search online.

RSVP allows end systems to request QoS guarantees from the network. The need for network resource reservations differs for data traffic versus for real-time traffic, as follows:

- Data traffic seldom needs reserved bandwidth because internetworks provide datagram services for data traffic. This asynchronous packet switching may not need guarantees of service quality. End-to-end controls between data traffic senders and receivers help ensure adequate transmission of bursts of information.
- Real-time traffic (that is, voice or video information) experiences problems when operating over datagram services. Because real-time traffic sends an almost constant flow of information, the network “pipes” must be consistent. Some guarantee must be provided that service between real-time hosts will not vary. Routers operating on a first-in, first-out (FIFO) basis risk unrecoverable disruption of the real-time information that is being transmitted.

Data applications, with little need for resource guarantees, frequently demand relatively lower bandwidth than real-time traffic. The almost constant high bit-rate demands of a video conference application and the bursty low bit-rate demands of an interactive data application share available network resources.

RSVP prevents the demands of traffic such as large file transfers from impairing the bandwidth resources necessary for bursty data traffic. When RSVP is used, the routers sort and prioritize packets much like a statistical time-division multiplexer would sort and prioritize several signal sources that share a single channel.

RSVP mechanisms enable real-time traffic to reserve resources necessary for consistent latency. A video conferencing application can use settings in the router to propagate a request for a path with the required bandwidth and delay for video conferencing destinations. RSVP will check and repeat reservations at regular intervals. By this process, RSVP can adjust and alter the path between RSVP end systems to recover from route changes.

Real-time traffic (unlike data traffic) requires a guaranteed network consistency. Without consistent QoS, real-time traffic faces the following problems:

- Jitter—A slight time or phase movement in a transmission signal can introduce loss of synchronization or other errors.
- Insufficient bandwidth—Voice calls use a digital signal level 0 (DS0 at 64 kbps); video conferencing uses T1/E1 (1.544 Mbps or 2.048 Mbps); and higher-fidelity video uses much more.
- Delay variations—If the wait time between when signal elements are sent and when they arrive varies, the real-time traffic will no longer be synchronized and may fail.
- Information loss—When signal elements drop or arrive too late, lost audio causes distortions with noise or crackle sounds. The lost video causes image blurring, distortions, or blackouts.

RSVP works in conjunction with weighted fair queuing (WFQ) or Random Early Detection (RED). This conjunction of reservation setting with packet queuing uses two key concepts: end-to-end flows with RSVP and router-to-router conversations with WFQ:

- RSVP flow—This is a stream that operates “multidestination simplex,” because data travels across it in only one direction: from the origin to the targets. Flows travel from a set of senders to a set of receivers. The flows can be merged or left unmerged, and the method of merging them varies according to the attributes of the application using the flow.
- WFQ Conversation—This is the traffic for a single transport layer session or network layer flow that crosses a given interface. This conversation is identified from the source and destination address, protocol type, port number, or other attributes in the relevant communications layer.

RSVP allows for hosts to send packets to a subset of all hosts (multicasting). RSVP assumes that resource reservation applies primarily to multicast applications (such as video conferencing). Although the primary target for RSVP is multimedia traffic, a clear interest exists for the reservation of bandwidth for unicast traffic (such as Network File System (NFS) and virtual private network management). A unicast transmission involves a host sending packets to a single host.

RSVP Reservation Types

These are the two types of multicast flows:

- Distinct Reservation. This constitutes a flow that originates from exactly one sender.
- Shared Reservation. This constitutes a flow that originates from one or more senders.

RSVP describes these reservations as having certain algorithmic attributes.

Distinct Reservation

An example of a distinct reservation is a video application in which each sender emits a distinct data stream that requires admission and management in a queue. Such a flow, therefore, requires a separate reservation per sender on each transmission facility it crosses (such as Ethernet, an HDLC line, a Frame Relay data-link connection identifier, or an ATM virtual channel). RSVP refers to this distinct reservation as explicit and installs it using a Fixed Filter style of reservation.

Use of RSVP for unicast applications is generally a degenerate case of a distinct flow.

Shared Reservation

An example of a shared reservation is an audio application in which each sender also emits a distinct data stream that requires admission and management in a queue. However, because of the nature of the application, a limited number of senders are transmitting data at any given time. Such a flow, therefore, does not require a separate reservation per sender. Instead, it uses a single reservation that can be applied to any sender within a set as needed.

RSVP installs a shared reservation using a Wild Card or Shared Explicit style of reservation, with the difference between the two determined by the scope of application (which is either wild or explicit):

- The Wild Card Filter reserves bandwidth and delay characteristics for any sender and is limited by the list of source addresses carried in the reservation message.
- The Shared Explicit reservation style identifies the flows for specific network resources.

Planning for RSVP Configuration

You must plan carefully to successfully configure and use RSVP on your network. At a minimum, RSVP must reflect your assessment of bandwidth needs on router interfaces. Consider the following questions as you plan for RSVP configuration:

- How much bandwidth should RSVP allow per end-user application flow? You must understand the “feeds and speeds” of your applications. By default, the amount reservable by a single flow can be the entire reservable bandwidth. You can, however, limit individual reservations to smaller amounts using the single flow bandwidth parameter. The reserved bandwidth value may not exceed the interface reservable amount, and no one flow may reserve more than the amount specified.
- How much bandwidth is available for RSVP? By default, 75 percent of the bandwidth available on an interface is reservable. If you are using a tunnel interface, RSVP can make a reservation for the tunnel whose bandwidth is the sum of the bandwidths reserved within the tunnel.
- How much bandwidth must be excluded from RSVP so that it can fairly provide the timely service required by low-volume data conversations? End-to-end controls for data traffic assume that all sessions will behave so as to avoid congestion dynamically. Real-time demands do not follow this behavior. Determine the bandwidth to set aside so bursty data traffic will not be deprived as a side effect of the RSVP QoS configuration.

Plan for RSVP before entering the details needed as RSVP configuration parameters.

RSVP Implementation Considerations

You should be aware of RSVP implementation considerations as you design your reservation system. RSVP does not model all data links likely to be present on the internetwork. RSVP models an interface as having a queuing system that completely determines the mix of traffic on the interface; bandwidth or delay characteristics are only deterministic to the extent that this model holds. Unfortunately, data links are often imperfectly modeled this way. Use the following guidelines:

- Serial line interfaces—Point-to-Point Protocol, High-Level Data Link Control, Link Access Procedure, Balanced, High-Speed Serial Interface, and similar serial line interfaces are well modeled by RSVP. The device can, therefore, make guarantees on these interfaces. Nonbroadcast multiaccess (NBMA) interfaces are also most in need of reservations.

- Multiaccess LANs—These data links are not modeled well by RSVP interfaces because the LAN itself represents a queuing system that is not under the control of the device making the guarantees. The device guarantees what load it will offer, but cannot guarantee the competing loads or timings of loads that neighboring LAN systems will offer. The network administrator can use admission controls to control how much traffic is placed on the LAN. The network administrator, however, should focus on the use of admission in network design in order to use RSVP effectively.
- Public X.25 networks—It is not clear that rate or delay reservations can be usefully made on public X.25 networks.

You must use a specialized configuration on Frame Relay and ATM networks, as discussed in the next sections.

Frame Relay Internetwork Considerations

The following RSVP implementation considerations apply as you design your reservation system for a Frame Relay internetwork:

- Reservations are made for an interface or subinterface. If subinterfaces contain more than one data-link control (DLC), the bandwidth required and the bandwidth reserved may differ. Therefore, RSVP subinterfaces of Frame Relay interfaces must contain exactly one DLC to operate correctly.
- In addition, Frame Relay DLCs have committed information rates (CIR) and burst controls (Bc and Be) that may not be reflected in the configuration and may differ markedly from the interface speed (either adding up to exceed it or being significantly smaller). Therefore, the **ip rsvp bandwidth** interface configuration command must be entered for both the interface and the subinterface. Both bandwidths are used as admission criteria.

For example, suppose that a Frame Relay interface runs at a T1 rate (1.544 Mbps) and supports several DLCs to remote offices served by 128-kbps and 56-kbps lines. You must configure the amount of the total interface (75 percent of which being 1.158 Mbps) and the amount of each receiving interface (75 percent of which would be 96 and 42 kbps, respectively) that may be reserved. Admission succeeds if and only if enough bandwidth is available on the DLC (the subinterface) and on the aggregate interface.

ATM Internetwork Considerations

The following RSVP implementation considerations apply as you design your reservation system for an ATM internetwork:

- When ATM is configured, it most likely uses a usable bit rate (UBR) or an available bit rate (ABR) virtual channel (VC) connecting individual routers. With these classes of service, the ATM network makes a “best effort” to meet the traffic’s bit-rate requirements and assumes that the end stations are responsible for information that does not get through the network.
- This ATM service can open separate channels for reserved traffic having the necessary characteristics. RSVP should open these VCs and adjust the cache to make effective use of the VC for this purpose.

Resource Reservation Protocol Configuration Task List

After you have planned your RSVP configuration, enter the Cisco IOS commands that implement your configuration plan. To configure RSVP, perform the tasks in the following sections. You must enable RSVP on an interface in order to use it; the other tasks are optional.

- Enable RSVP
- Enter Senders in the RSVP Database
- Enter Receivers in the RSVP Database
- Enter Multicast Addresses
- Control Which RSVP Neighbor Can Offer a Reservation
- Monitor RSVP

Enable RSVP

By default, RSVP is disabled so that it is backward compatible with systems that do not implement RSVP. To enable RSVP on an interface, use the following command in interface configuration mode:

Command	Purpose
ip rsvp bandwidth <i>[interface-kbps] [single-flow-kbps]</i>	Enable RSVP for IP on an interface.

This command starts RSVP and sets the bandwidth and single-flow limits. The default maximum bandwidth is up to 75 percent of the bandwidth available on the interface. By default, the amount reservable by a flow can be up to the entire reservable bandwidth.

On subinterfaces, this command applies the more restrictive of the available bandwidths of the physical interface and the subinterface. For example, a Frame Relay interface might have a T1 connector nominally capable of 1.536 Mbps, and 64-kbps subinterfaces on 128-kbps circuits (64-KB CIR), with 1200 and 100 kbps, respectively.

Reservations on individual circuits that do not exceed 100 kbps normally succeed. If, however, reservations have been made on other circuits adding up to 1.2 Mbps, and a reservation is made on a subinterface that itself has enough remaining bandwidth, the reservation will still be refused because the physical interface lacks supporting bandwidth.

Enter Senders in the RSVP Database

You can configure the router to behave as though it is periodically receiving an RSVP PATH message from the sender or previous hop routes containing the indicated attributes. To enter senders in the RSVP database, use the following command in global configuration mode:

Command	Purpose
ip rsvp sender <i>session-ip-address sender-ip-address [tcp udp ip-protocol] session-dport sender-sport previous-hop-ip-address previous-hop-interface bandwidth burst-size</i>	Enter the senders in the RSVP database.

Enter Receivers in the RSVP Database

You can configure the router to behave as though it is continuously receiving an RSVP RESV message from the originator containing the indicated attributes. To enter receivers in the RSVP database, use the following command in global configuration mode:

Command	Purpose
ip rsvp reservation <i>session-ip-address sender-ip-address</i> [tcp udp <i>ip-protocol</i>] <i>session-dport sender-sport next-hop-ip-address next-hop-interface</i> { ff se wf } { rate load } <i>bandwidth burst-size</i>	Enter the receivers in the RSVP database.

Enter Multicast Addresses

If RSVP neighbors are discovered to be using UDP encapsulation, the router will automatically generate UDP-encapsulated messages for consumption by the neighbors.

To enter multicast addresses, use the following command in global configuration mode:

Command	Purpose
ip rsvp udp-multicast [<i>multicast-address</i>]	Enter any multicast addresses necessary if you use UDP.

However, in some cases, a host will not originate such a message until it has first heard from the router, which it can only do via UDP. You must instruct the router to generate UDP-encapsulated RSVP multicasts whenever it generates an IP-encapsulated multicast.

Control Which RSVP Neighbor Can Offer a Reservation

By default, any RSVP neighbor may offer a reservation. To control which RSVP neighbors can offer a reservation, use the following command in global configuration mode:

Command	Purpose
ip rsvp neighbors <i>access-list-number</i>	Limit which routers may offer reservations.

When this command is configured, only neighbors conforming to the access list are accepted. The access list is applied to the IP header.

Monitor RSVP

After you configure the RSVP reservations that reflect your network resource policy, you can verify the resulting RSVP operations. To do so, use the following commands in EXEC mode:

Command	Purpose
show ip rsvp interface [<i>type number</i>]	Display RSVP-related interface information.
show ip rsvp installed [<i>type number</i>]	Display RSVP-related filters and bandwidth information.
show ip rsvp neighbor [<i>type number</i>]	Display current RSVP neighbors.
show ip rsvp sender [<i>type number</i>]	Display RSVP sender information.
show ip rsvp request [<i>type number</i>]	Display RSVP request information.
show ip rsvp reservation [<i>type number</i>]	Display RSVP receiver information.

RSVP Configuration for a Multicast Session Example

This section describes configuration of RSVP on three Cisco 4500 routers for a multicast session. The three routers form the router network between an RSVP sender application running on an upstream (end-system) host and an RSVP receiver application running on a downstream (end-system) host—neither host is shown in this example.

The router network includes three routers: Router A, Router B, and Router C. The example presumes that Router A's upstream interface Hssi0 links to the upstream host. Router A and Router B are connected by Router A's downstream interface Ethernet1, which links to Router B's upstream interface Ethernet 1. Router B and Router C are connected by Router B's downstream interface Hssi0, which links to Router C's upstream interface Hssi0. The example presumes that Router C's downstream interface Ethernet2 links to the downstream host.

Typically, an RSVP-capable application running on an end system host on one side of a router network sends out either unicast or multicast RSVP PATH (Set Up) messages to the destination end system or host on the other side of the router network with which it wishes to communicate. The initiating application is referred to as the sender; the target or destination application is called the receiver. In this example, the sender runs on the host upstream from Router A and the receiver runs on the host downstream from Router C. The router network delivers the RSVP PATH messages from the sender to the receiver. The receiver replies with RSVP RESV messages in an attempt to reserve across the network the requested resources that are required between itself and the sender. The RSVP RESV messages specify the parameters for the requisite QoS that the router network connecting the systems should attempt to offer.

This example does not show the host that would run the sender application and the host that would run the receiver application. Normally, the first router downstream from the sender in the router network—in this case, Router A—would receive the RSVP PATH message from the sender. Normally, the last router in the router network—that is, the next hop upstream from the host running the receiver application, in this case, Router C—would receive an RSVP RESV message from the receiver.

Because this example does not explicitly include the hosts on which the sender and receiver applications run, the routers have been configured to act as if they were receiving PATH messages from a sender and RESV messages from a receiver. The commands used for this purpose, allowing RSVP to be more fully illustrated in the example, are the **ip rsvp sender** command and the **ip rsvp reservation** command. On Router A, the following command has been issued:

```
ip rsvp sender 225.1.1.1 12.1.2.1 UDP 7001 7000 12.1.2.1 Hs0 20 1
```

This command causes the router to act as if it were receiving PATH messages destined to multicast address 225.1.1.1 from a source 12.1.2.1. The previous hop of the PATH message is 12.1.2.1, and the message was received on interface Hssi0.

On Router C, the following command has been issued:

```
ip rsvp reservation 225.1.1.1 12.1.2.1 UDP 7001 7000 9.1.2.1 Et2 FF LOAD 8 1
```

This command causes the router to act as if it were receiving RESV messages for the session with multicast destination 225.1.1.1. The messages request a Fixed Filter reservation to source 12.1.1.1, and act as if they had arrived from a receiver on interface Ethernet2 with address 9.1.2.1.

In the example, the RSVP PATH messages flow in one direction: downstream from the sender, which in this example is Router A. (If the host were to initiate the RSVP PATH message, the message would flow from the host to Router A.) Router A sends the message downstream to Router B, and Router B sends it downstream to Router C. (If the downstream host were the actual receiver, Router C would send the RSVP PATH message downstream to the receiver host.) Each router in the router network must process the RSVP PATH message and route it to the next downstream hop.

The RSVP RESV messages flow in one direction: upstream from the receiver—in this example, Router C—upstream from Router C to Router B, and upstream from Router B to Router A. If the downstream host were the receiver, the message would originate with the host, which would send it to Router C. If the upstream host were the sender, the final destination of the RSVP RESV message would be the upstream host. At each hop, the router receiving the RSVP RESV message must determine whether it can honor the reservation request.

The **ip rsvp bandwidth** command both enables RSVP on an interface and specifies the amount of bandwidth on the interface that can be reserved (as well as the amount of bandwidth that can be allocated to a single flow). To ensure QoS for the RSVP reservation, weighted fair queuing is configured on the interfaces enabled for the reservation.

If the router network is capable of offering the specified (QoS) level of service, then an end-to-end reserved path is established. If not, the reservation attempt is rejected and a reservation error message is sent to the receiver. The ability of each router in the network to honor the requested level of service is verified, link by link, as the RSVP RESV messages are sent across the router network to the sender. However, the data itself for which the bandwidth is reserved travels one way only: from the sender to receiver across an established PATH. Therefore, the QoS is effective in only one direction. This is the common case for one-to-many multicast data flows.

After the three routers in the example are configured, the **show ip rsvp sender** and **show ip rsvp reservation** commands will make visible the PATH and RESV state.

On Router A, RSVP is enabled on Ethernet1 with 10 kbps to be reserved for the data transmission. A weighted fair queue is reserved on this interface to ensure RSVP QoS. (On Router A, RSVP is also enabled on interface Hssi0 with 1 kbps reserved, but this bandwidth is used simply for passing messages.)

Router A

```
!  
version 12.0  
service config  
service timestamps debug uptime  
service timestamps log uptime  
no service password-encryption  
service udp-small-servers  
service tcp-small-servers  
!  
hostname routerA  
!  
ip subnet-zero  
no ip domain-lookup  
ip multicast-routing  
ip dvmrp route-limit 20000  
!  
!  
interface Ethernet0  
ip address 2.0.0.193 255.0.0.0  
no ip directed-broadcast  
no ip route-cache  
no ip mroute-cache  
media-type 10BaseT  
!  
interface Ethernet1  
ip address 11.1.1.2 255.0.0.0  
no ip directed-broadcast  
ip pim dense-mode  
ip rsvp bandwidth 10 10  
fair-queue 64 256 1000  
media-type 10BaseT
```

```

!
interface Hssi0
 ip address 12.1.1.1 255.0.0.0
 no ip directed-broadcast
 ip pim dense-mode
 ip rsvp bandwidth 1 1
!
interface ATM0
 no ip address
 no ip directed-broadcast
 shutdown
!
router ospf 100
 network 11.0.0.0 0.255.255.255 area 10
 network 12.0.0.0 0.255.255.255 area 10
!
ip classless
ip rsvp sender 225.1.1.1 12.1.2.1 UDP 7001 7000 12.1.2.1 Hs0 20 1
!
!
!
line con 0
 exec-timeout 0 0
 length 0
 transport input none
line aux 0
line vty 0 4
 login
!
end

```

On Router B, RSVP is enabled on interface Hssi0 with 20 kbps to be reserved for the data transmission. A weighted fair queue is reserved on this interface to ensure RSVP QoS. (On Router B, RSVP is also enabled on interface Ethernet1 with 1 kbps reserved, but this bandwidth is used simply for passing messages.)

Router B

```

!
!
version 12.0
service config
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
service udp-small-servers
service tcp-small-servers
!
hostname routerB
!
ip subnet-zero
no ip domain-lookup
ip multicast-routing
ip dvmrp route-limit 20000
clock calendar-valid
!
!

```

RSVP Configuration for a Multicast Session Example

```
interface Ethernet0
 ip address 2.0.0.194 255.0.0.0
 no ip directed-broadcast
 no ip route-cache
 no ip mroute-cache
 media-type 10BaseT
!
interface Ethernet1
 ip address 11.1.1.1 255.0.0.0
 no ip directed-broadcast
 ip pim dense-mode
 ip rsvp bandwidth 1 1
 media-type 10BaseT
!
interface Hssi0
 ip address 10.1.1.2 255.0.0.0
 no ip directed-broadcast
 ip pim dense-mode
 ip rsvp bandwidth 20 20
 fair-queue 64 256 1000
 hssi internal-clock
!
interface ATM0
 no ip address
 no ip directed-broadcast
 shutdown
!
router ospf 100
 network 10.0.0.0 0.255.255.255 area 10
 network 11.0.0.0 0.255.255.255 area 10
!
ip classless
!
!
!
line con 0
 exec-timeout 0 0
 length 0
 transport input none
line aux 0
line vty 0 4
 login
!
end
```

On Router C, RSVP is enabled on interface Ethernet2 with 20 kbps to be reserved for the data transmission. A weighted fair queue is reserved on this interface to ensure RSVP QoS. (On Router C, RSVP is also enabled on interface Hssi0 with 1 kbps reserved, but this bandwidth is used simply for passing messages.)

Router C

```
!  
version 12.0  
service config  
service timestamps debug uptime  
service timestamps log uptime  
no service password-encryption  
service udp-small-servers  
service tcp-small-servers  
!  
hostname routerC  
!  
ip subnet-zero  
no ip domain-lookup  
ip multicast-routing  
ip dvmrp route-limit 20000  
!  
!  
interface Ethernet0  
 ip address 2.0.0.195 255.0.0.0  
 no ip directed-broadcast  
 no ip route-cache  
 no ip mroute-cache  
 media-type 10BaseT  
!  
interface Ethernet1  
 no ip address  
 no ip directed-broadcast  
 shutdown  
 media-type 10BaseT  
!  
interface Ethernet2  
 ip address 9.1.1.2 255.0.0.0  
 no ip directed-broadcast  
 ip pim dense-mode  
 ip rsvp bandwidth 20 20  
 fair-queue 64 256 1000  
 media-type 10BaseT  
!  
interface Ethernet3  
 no ip address  
 no ip directed-broadcast  
 shutdown  
 media-type 10BaseT  
!  
interface Ethernet4  
 no ip address  
 no ip directed-broadcast  
 shutdown  
 media-type 10BaseT  
!  
interface Ethernet5  
 no ip address  
 no ip directed-broadcast  
 shutdown  
 media-type 10BaseT  
!  
interface Hssi0  
 ip address 10.1.1.1 255.0.0.0  
 no ip directed-broadcast  
 ip pim dense-mode  
 ip rsvp bandwidth 1 1  
 hssi internal-clock  
!
```

RSVP Configuration for a Multicast Session Example

```
interface ATM0
  no ip address
  no ip directed-broadcast
  shutdown
!
router ospf 100
  network 9.0.0.0 0.255.255.255 area 10
  network 10.0.0.0 0.255.255.255 area 10
  network 11.0.0.0 0.255.255.255 area 10
!
ip classless
ip rsvp reservation 225.1.1.1 12.1.2.1 UDP 7001 7000 9.1.2.1 Et2 FF LOAD 8 1
!
!
!
line con 0
  exec-timeout 0 0
  length 0
  transport input none
line aux 0
line vty 0 4
  login
!
end
```