



Signalling Overview

In the most general sense, QoS signalling is a form of network communication that allows an end station or network node to communicate with, or signal, its neighbors to request special handling of certain traffic. QoS signalling is useful for coordinating the traffic handling techniques provided by other QoS features. It plays a key role in configuring successful overall end-to-end QoS service across your network.

True end-to-end QoS requires that every element in the network path—switch, router, firewall, host, client, and so forth—deliver its part of QoS, and that all of these entities be coordinated with QoS signalling.

Many viable QoS signalling solutions provide QoS at some places in the infrastructure; however, they often have limited scope across the network. To achieve end-to-end QoS, signalling must span the entire network.


Cisco IOS QoS software takes advantage of IP to meet the challenge of finding a robust QoS signalling solution that can operate over heterogeneous network infrastructures. It overlays Layer 2 technology-specific QoS signalling solutions with Layer 3 IP QoS signalling methods of the Resource Reservation Protocol (RSVP) and IP Precedence features.

An IP network can achieve end-to-end QoS, for example, by using part of the IP packet header to request special handling of priority or time-sensitive traffic. Given the ubiquity of IP, QoS signalling that takes advantage of IP provides powerful end-to-end signalling. Both RSVP and IP Precedence fit this category.

Either in-band (IP Precedence, 802.1p) or out-of-band (RSVP) signalling is used to indicate that a particular QoS service is desired for a particular traffic classification. IP Precedence signals for differentiated QoS and RSVP for guaranteed QoS.

IP Precedence

As shown in Figure 11, IP Precedence utilizes the three precedence bits in the type of service (ToS) field of the IPv4 header to specify class of service for each packet. You can partition traffic in up to six classes of service using IP Precedence. The queueing technologies throughout the network can then use this signal to provide the appropriate expedited handling.

Figure 11 IP Precedence ToS Field

You can use features such as policy-based routing (PBR) and committed access rate (CAR) to set precedence based on extended access list classification. Use of these features allows considerable flexibility of precedence assignment, including assignment by application or user, or by destination or source subnet. Typically, you deploy these features as close to the edge of the network or the administrative domain as possible, so that each subsequent network element can provide service based on the determined policy. IP Precedence can also be set in the host or the network client; however, IP Precedence can be overridden by policy within the network.

IP Precedence enables service classes to be established using existing network queuing mechanisms, such as weighted fair queuing (WFQ) and Weighted Random Early Detection (WRED), with no changes to existing applications and with no complicated network requirements.

Resource Reservation Protocol

RSVP is the first significant industry-standard protocol for dynamically setting up end-to-end QoS across a heterogeneous network. RSVP, which runs over IP, allows an application to dynamically reserve network bandwidth. Using RSVP, applications can request a certain level of QoS for a data flow across a network.

The Cisco IOS QoS implementation allows RSVP to be initiated within the network using configured proxy RSVP. Using this capability, you can take advantage of the benefits of RSVP in the network even for non-RSVP enabled applications and hosts. RSVP is the only standard signalling protocol designed to guarantee network bandwidth from end-to-end for IP networks.

RSVP does not perform its own routing; instead it uses underlying routing protocols to determine where it should carry reservation requests. As routing changes paths to adapt to topology changes, RSVP adapts its reservation to the new paths wherever reservations are in place. This modularity does not prevent RSVP from using other routing services. RSVP provides transparent operation through router nodes that do not support RSVP.

RSVP works in conjunction with, not in place of, current queuing mechanisms. RSVP requests the particular QoS, but it is up to the particular interface queuing mechanism, such as WFQ or WRED, to implement the reservation.

You can use RSVP to make two types of dynamic reservations: controlled load and guaranteed rate services, both of which are briefly described in the chapter “Quality of Service Overview.”

A primary feature of RSVP is its scalability. RSVP scales well using the inherent scalability of multicast. RSVP scales to very large multicast groups because it uses receiver-oriented reservation requests that merge as they progress up the multicast tree. Although RSVP is designed specifically for multicast applications, it may also make unicast reservations though it does not scale as well with a large number of unicast reservations.

RSVP is an important QoS feature, but it does not solve all problems addressed by QoS, and it imposes a few hindrances such as the time required to set up end-to-end reservation.

How It Works

Hosts and routers use RSVP to deliver QoS requests to the routers along the paths of the data stream and to maintain router and host state to provide the requested service, usually bandwidth and latency. RSVP uses a mean data rate—the largest amount of data the router will keep in the queue—and minimum QoS (that is, guarantee of the requested bandwidth specified when you made the reservation using RSVP) to determine bandwidth reservation.

A host uses RSVP to request a specific QoS service from the network on behalf of an application data stream. RSVP requests the particular QoS, but it is up to the interface queueing mechanism to implement the reservation. RSVP carries the request through the network, visiting each node the network uses to carry the stream. At each node, RSVP attempts to make a resource reservation for the stream using its own admission control module, exclusive to RSVP, which determines whether the node has sufficient available resources to supply the requested QoS.

**Note**

For RSVP, an application could send traffic at a rate higher than the requested QoS, but the application is guaranteed only the minimum requested rate. If bandwidth is available, traffic surpassing the requested rate will go through if sent; if bandwidth is not available, the exceeding traffic will be dropped.

If the required resources are available and the user is granted administrative access, the RSVP daemon sets arguments in the packet classifier and packet scheduler to obtain the desired QoS. The classifier determines the QoS class for each packet and the scheduler orders packet transmission to achieve the promised QoS for each stream. If either resource is unavailable or the user is denied administrative permission, the RSVP program returns an error notification to the application process that originated the request.

WFQ or WRED sets up the packet classification and the scheduling required for the reserved flows. Using WFQ, RSVP can deliver an integrated services Guaranteed Rate Service. Using WRED, it can deliver a Controlled Load Service.

**Note**

For information on how to configure RSVP, see the chapter “Configuring RSVP” in this book.

Subnetwork Bandwidth Manager

RSVP and its service class definitions are largely independent of the underlying network technologies. This independence requires that a user define the mapping of RSVP onto subnetwork technologies.

The Subnetwork Bandwidth Manager (SBM) feature answers this requirement for RSVP in relation to IEEE 802-based networks. SBM specifies a signalling method and protocol for LAN-based admission control for RSVP flows. SBM allows RSVP-enabled routers and Layer 2 and Layer 3 devices to support reservation of LAN resources for RSVP-enabled data flows. The SBM signalling method is similar to that of RSVP itself. SBM protocol entities have the following features:

- Reside in Layer 2 or Layer 3 devices.
- Can manage resources on a segment. A segment is a Layer 2 physical segment shared by one or more senders, such as a shared Ethernet or Token Ring wire.

- Can become candidates in a dynamic election process that designates one SBM as the segment manager. The elected candidate is called the Designated Subnetwork Bandwidth Manager (DSBM). The elected DSBM is responsible for exercising admission control over requests for resource reservations on a managed segment.


A managed segment includes those interconnected parts of a shared LAN that are not separated by DSBMs. The presence of a DSBM makes the segment a managed one. One or more SBMs may exist on a managed segment, but there can be only one DSBM on each managed segment.

You can configure an interface on routers connected to the segment to participate in the DSBM election process. The contender configured with the highest priority becomes the DSBM for the managed segment.

If you do not configure a router as a DSBM candidate and RSVP is enabled, then the system interacts with the DSBM if a DSBM is present on the segment. In fact, if a DSBM, identifying itself as such, exists on the segment, the segment is considered a managed segment and all RSVP message forwarding will be based on the SBM message forwarding rules. This behavior exists to allow cases in which you might not want an RSVP-enabled interface on a router connected to a managed segment interface to become a DSBM, but you want it to interact with the DSBM if one is present managing the segment.

Figure 12 shows a managed segment in a Layer 2 domain that interconnects a set of hosts and routers.

Figure 12 DSBM Managed Segment



When a DSBM client sends or forwards an RSVP PATH message over an interface attached to a managed segment, it sends the PATH message to the DSBM of the segment instead of to the RSVP session destination address, as is done in conventional RSVP processing. As part of its message processing procedure, the DSBM builds and maintains a PATH state for the session and notes the previous Layer 2/Layer 3 hop from which it received the PATH message. After processing the PATH message, the DSBM forwards it toward its destination address.

The DSBM receives the RSVP reservation request (RSVP RESV) message and processes it in a manner similar to how RSVP itself handles reservation request processing, basing the outcome on available bandwidth. The procedure is as follows:

- If it cannot grant the request because of lack of resources, the DSBM returns a RESVERR message to the requester.
- If sufficient resources are available and the DSBM can grant the reservation request, it forwards the RESV message toward the previous hop(s) using the local PATH state for the session.



Note

For information on how to configure SBM, see the chapter “Configuring Subnetwork Bandwidth Manager” in this book.

RSVP-ATM QoS Interworking

The RSVP-ATM QoS Interworking feature provides support for Controlled Load Service using RSVP over an ATM core network. This feature requires the ability to signal for establishment of switched virtual circuits (SVCs) across the ATM cloud in response to RSVP reservation request messages. To meet this requirement, RSVP over ATM supports mapping of RSVP sessions to ATM SVCs.

The RSVP-ATM QoS Interworking feature allows you to do the following:

- Configure an interface or subinterface to dynamically create SVCs in response to RSVP reservation request messages. To ensure defined QoS, these SVCs are established having QoS profiles consistent with the mapped RSVP flow specifications (flowspecs).
- Attach Distributed Weighted Random Early Detection (DWRED) group definitions to the Enhanced ATM port adapter (PA-A3) interface to support per-VC DWRED drop policy. Use of per-VC DWRED ensures that if packets must be dropped, then best-effort packets are dropped first and not those that conform to the appropriate QoS determined by the RSVP's token bucket.
- Configure the IP Precedence and ToS values to be used for packets that conform to or exceed QoS profiles. As part of its input processing, RSVP uses the values that you specify to set the ToS and IP Precedence bits on incoming packets. If per-VC DWRED is configured, it then uses the ToS and IP Precedence bit settings on the output interface of the same router in determining which packets to drop. Also, interfaces on downstream routers use these settings in processing packets.

How It Works

Traditionally, RSVP has been coupled with WFQ. WFQ provides bandwidth guarantees to RSVP and gives RSVP visibility to all packets visible to it. This visibility allows RSVP to identify and mark packets pertinent to it.

The RSVP-ATM QoS Interworking feature allows you to decouple RSVP from WFQ, and instead associate it with ATM SVCs to handle reservation request messages (and provide bandwidth guarantees) and NetFlow to make packets visible to RSVP.

To configure an interface or subinterface to use the RSVP-ATM QoS Interworking feature, the **ip rsvp svc-required** command is used. Then, whenever a new RSVP reservation is requested, the router software establishes a new ATM SVC to service the reservation.

To ensure correspondence between RSVP and ATM SVC values, the software algorithmically maps the rate and burst size parameters in the RSVP flowspec to the ATM sustained cell rate (SCR) and maximum burst size (MBS). For the peak cell rate (PCR), it uses the value you configure or it defaults to the line rate. RSVP-ATM QoS Interworking requires an Enhanced ATM port adapter (PA-A3) with OC-3 speed.


When a packet belonging to a reserved flow arrives on the interface or subinterface, the RSVP-ATM QoS Interworking software uses a token bucket to manage bandwidth guarantees. It measures actual traffic rates against the reservation flowspec to determine if the packet conforms to or exceeds the flowspec. Using values you configure for conformant or exceeding traffic, it sets the IP Precedence and ToS bits in the ToS byte of the packet's header and delivers the packet to the appropriate virtual circuit (VC) for transmission. For the RSVP-ATM QoS Interworking feature, packets are shaped before they are sent on the ATM SVC. Shaping creates back pressure to the Versatile Interface Processor (VIP) when the offered load exceeds the rate.

The RSVP-ATM QoS Interworking software uses per-SVC DWRED to drop packets when shaping causes a queue to build up on the VIP. Use of per-SVC DWRED allows RSVP to deliver Controlled Load Service class, which requires that reserved packets experience performance equivalent to that of an unloaded network (which is one with very low loss and moderate delay). For a more detailed account of how the RSVP-ATM QoS Interworking feature works, see the following example scenario.

An Example Scenario

To understand the behavior of the RSVP-ATM QoS Interworking feature, consider the following example, which uses a Cisco 7500 router with VIP ingress and egress interfaces and RSVP ingress functionality implemented on the Route Switch Processor (RSP). Figure 13 illustrates this example; it shows a pair of routers that communicate over the ATM cloud. In this example, a single PVC is used for RSVP request messages and an ATM SVC is established to handle each new reservation request message.

Figure 13 Two Routers Connected over an ATM Core Network



Host X, which is upstream from Router A, is directly connected to Router A using FDDI. Host Y, which is downstream from Router B, is directly connected to Router B using FDDI. (In an alternative configuration, these host-router connections could use ATM VCs.)

For the RSVP-ATM QoS Interworking feature, reservations are needed primarily between routers across the ATM backbone network. To limit the number of locations where reservations are made, you can enable RSVP selectively only at subinterfaces corresponding to router-to-router connections across the ATM backbone network. Preventing reservations from being made between the host and the router both limits VC usage and reduces load on the router.

RSVP RESV messages flow from receiving host to sending host. In this example, Host Y is the sending host and Host X is the receiving host. (Host Y sends a RESV message to Host X.) Router B, which is at the edge of the ATM cloud, receives the RESV message and forwards it upstream to Router A across the PVC used for control messages. The example configuration shown in Figure 13 uses one PVC; as shown, it carries the RSVP request.

The ingress interface on Router A is configured for RSVP-ATM, which enables it to establish for each request an SVC to service any new RSVP RESV reservations made on the interface. When it receives a reservation request, the interface on Router A creates a new nonreal-time variable bit rate (nRTVBR) SVC with the appropriate QoS characteristics. The QoS characteristics used to establish the SVC result from algorithmic mapping of the flowspec in the RSVP RESV message to the appropriate set of ATM signalling parameters.

In this example, Controlled Load Service is used as the QoS class. The ATM PCR parameter is set to the line rate. If the **ip RSVP atm-peak-rate-limit** command is used on the interface to configure a rate limiter, the PCR is set to the peak rate limiter. The ATM SCR parameter is set to the RSVP flowspec rate and the ATM MBS is set to the RSVP flowspec burst size. Packets are shaped before they are sent on the ATM SVC. Shaping creates back pressure to the VIP when the offered load exceeds the rate.

When a new SVC is set up to handle a reservation request, another state is also set up including a classifier state that uses a packet's source and destination addresses and port numbers to determine which, if any, reservation the packet belongs to. Also, a token bucket is set up to ensure that if a source sends more data than the data rate and MBS parameters of its flowspec specify, the excess traffic does not interfere with other reservations.

More specifically, here is how data traverses the path. When a data packet destined for Router B arrives at Router A, before they traverse the ATM cloud, the packet's source and destination addresses and port numbers are checked against the RSVP filterspec to determine if the packet matches a reservation.

If the packet does not match a reservation, it is sent out the best-effort PVC to Router B. If a packet matches a reservation, it is further processed by RSVP. The packet is checked against the reservation's token bucket to determine whether it conforms to or exceeds the token bucket parameters. (All packets matching a reservation are sent out on the reservation's SVC to prevent misordering of packets.)

To introduce differentiation between flowspec-conformant and flowspec-exceeding packets, you can specify values for RSVP-ATM to use in setting the IP Precedence and ToS bits of the packets. To specify these values, you use the **ip rsvp precedence** and **ip rsvp tos** commands. When you set different precedence values for conformant and exceeding packets and use a preferential drop policy such as DWRED, RSVP-ATM ensures that flowspec-exceeding packets are dropped prior to flowspec-conformant packets when the VC is congested.

**Note**

For information on how to configure the RSVP-ATM QoS Interworking feature, see the chapter “Configuring RSVP-ATM QoS Interworking” in this book.
