

Signalling Overview

In the most general sense, QoS signalling is a form of network communication that provides a way for 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 all of these entities must 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 the Internet Protocol (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 Resource Reservation Protocol (RSVP) and IP Precedence.


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 IP Precedence and RSVP 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 12, IP Precedence utilizes the three precedence bits in the IPv4 header's type of service (ToS) field 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 12 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 queueing mechanisms, such as WFQ and 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, network managers 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 queueing mechanisms. RSVP requests the particular QoS, but it is up to the particular interface queueing mechanism, such as weighted fair queueing (WFQ) or Weighted Random Early Detection (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 “Introduction: Quality of Service Overview.”

A primary feature of RSVP is its scalability. RSVP scales well using multicast’s inherent scalability. 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 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 queuing 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.

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. If both attempts succeed, the RSVP daemon sets parameters in a packet classifier and packet scheduler to obtain the desired QoS. The packet classifier determines the QoS class for each packet and the scheduler orders packet transmission to achieve the promised QoS for each stream.

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.

