Quality of Service for IP Videoconferencing

Engineering White Paper

Subha Dhesikan
Cisco Systems

June 1st, 2001

Table of Contents

Terminology

Administrative domain	A collection of networks that are under the same administrative control.
Admission Control	This is the mechanism which decides whether the network device has sufficient resources to supply the requested QoS.
COPS	Common Open Policy Service. Client/Server protocol to support policy control. It is described in [7].
Flow	A flow is a set of packets belonging to one instance of the application identified by some combination of source address, source port, destination address, destination port and protocol identifier.
Jitter	Variation in delay.
Network device	This refers to a device in the network that handles traffic. Routers and switches are examples of network devices.
Policing	This is the process of enforcing the policies, which could result in delaying or dropping packets.
Policy	This is a set of rules that define the criteria for allowing access to a network resource.
Policy control	The application of policies to make a decision whether to allow access to a resource.
Policy Decision Point (PDP)	A COPS acronym. This is the device where the policy decisions are made. The PDP has usually a global knowledge of all the network polices that pertain to one administrative domain.
Policy Enforcement Point (PEP)	This is the device where the policy decisions are enforced.
Quality of Service (QoS)	This refers to the type of service provided by the network devices.
Resource	This refers to all the factors in the network device that affect the forwarding of packets such as bandwidth on an interface, queues, processing power etc.
Traffic	Traffic refers to one or more flows that traverse through the network.

1 Introduction

In recent years there has been a tremendous increase in the capacity of the networks and an accompanying increase in the consumption of that capacity. Web traffic, voice and video over IP, and other multimedia applications have significantly contributed to the increase in traffic on our networks. Some applications, such as the interactive multimedia applications, not only require large amounts of bandwidth, but also require specialized service from the network with respect to latency and loss. The network must accommodate such traffic without unduly degrading the performance of other applications in the network, some of which may be mission-critical.

IP videoconferencing is one such application that requires special service from the network. To provide predictable performance, videoconferencing requires significant bandwidth with minimal delay, jitter and loss. The inability to provide the required network quality has been one of the major obstacles affecting the popularity and deployment of this application. This document describes a solution in providing network Quality of Service (QoS) to videoconferencing applications. It describes the different QoS building blocks, the most popular QoS architectures, and solutions for obtaining the required QoS from the network for videoconferencing traffic.

2 Why QoS?

There must be adequate bandwidth in the network to meet the demands of the offered load. Even then, there could be contention for network bandwidth for short periods of time. This occurs because traffic in the network is not evenly spread during all times. When traffic arrives at a network device, the devices processes and forwards the traffic. The amount of traffic that a device can forward is limited by the capacity of its interfaces. There are certain times when too many packets arrive at a device at the same time. At such times, the device cannot forward all the packets immediately due to the limited capacity of its interfaces. These excess packets are either queued in a buffer or dropped. This causes congestion and results in delay and loss of packets. Some ways to handle this congestion are:

- **Over-provision the network**
  One could add more bandwidth and over-provision the network to ensure that the need for bandwidth can be satisfied at all times. With over-provisioning, it is possible to prevent any conflict for resources. All incoming traffic can be serviced immediately since there is always sufficient capacity in the network. This may sound like an ideal solution. However, over-provisioning leads to enormous wastage of valuable resources because the degree of over provisioning may need to be a factor of 10 or 100 the capacity needed to service the load successfully. Consequently, over-provisioning may not be the most desirable solution.

- **Separate networks**
  Separate networks can be set up for each application type so that there is no resource conflict between traffic types. For instance, an organization can have a network separate for voice, video and data. Like over-provisioning, this results in poor utilization of resources. In addition, this only solves the problem of multimedia applications interfering with data applications. It does not solve the problem of having more voice traffic than there is bandwidth on the voice network, unless it is combined with over-provisioning. Therefore, the trend is to consolidate more and more applications over a single IP network.

- **Quality of Service (QoS)**
  Some applications, such as FTP, HTTP and email, are not sensitive to delay and jitter. On the other hand applications such as interactive voice and video are vulnerable to loss, delay and jitter. During peak times, any device might delay and/or drop some packets to relieve congestion. For this solution to be effective, the QoS network devices must be able to differentiate among classes of arriving traffic and satisfy their individual requirements. QoS mechanisms provide a set of tools that can be used to do that. It enables the network to recognize traffic belonging to certain users and applications such that preferential services may be provided to them. (This is why QoS is also known as “managed unfairness”). Some QoS mechanisms also enable applications to provide information to the network to aid the devices in performing traffic classification.

QoS is therefore the best way to handle contention for network resources when the network is intended to service widely varying types of traffic. It is important to emphasize that there should be sufficient network capacity to start with. QoS does not create any additional capacity. It just helps to manage the available resources according to policies set out by the network administrator.

3 QoS Primitives

QoS tools help to manage the utilization of network resources. QoS tools fall into the following categories:

- Tools that enable the classification of incoming traffic based on either explicit signaling or heuristic identification of traffic classes and/or individual flows.
- Tools that provide a variety of *scheduling* disciplines for traffic in order to meet desired delay, bandwidth, and jitter characteristics.
- Tools that provide admission control for traffic so that use of network resources can be explicitly denied, rather than having service degraded when overloading occurs.
- Tools that police incoming traffic to protect the network from misbehaving applications.

These tools are explained in detail below:

- **Packet Classification**
  The packet classifier in the network device separates the arriving traffic into either distinct classes or individual application *flows* such that each class of traffic or flow may be provided a different service. Both types of classifiers that this paper refers to look at information in the IP header. One looks at the explicit markings in the type of service field or the traffic class field in IPv4 and IPv6, respectively. The other looks at multiple fields in the IP header (source IP address, source port, destination IP address, destination port and protocol type) to associate it with a particular application’s traffic flow. Once the traffic is classified into multiple classes, it is possible for each class of packets to receive different treatment from the device.

- **Queuing and scheduling**
  Traffic flows/classes are assigned to different queues based on their classification. By placing packets into different queues, diverse service needs can be met. For instance, special requirements such as low delay can be provided to packets by servicing their queue more frequently. Based on their classification, the schedulers determine which queue packets enter and when packets exit the queues.

  The *queuing discipline* configured in the device implements the scheduling. It determines how the packets in each of the queues are scheduled for transmission. Priority Queuing is an example of a simple queuing model. In this model, packets are forwarded based strictly on the priority of the queue. The packets in the high priority queue are forwarded first. If there are no packets in the high priority queue then the packets in the medium priority queue are forwarded and so on. In this model high priority packets are never scheduled behind lower priority packets. While providing minimal latency service to packets in a high priority queue, strict priority queuing can cause starvation of traffic in lower priority classes. Hence it is dangerous to use priority queuing without effective policing of how much traffic can enter the high priority queue. There are several other queuing disciplines such as Round Robin (RR) and Weighted Fair Queuing (WFQ), which provide different QoS services.

- **Admission Control**

  If an airplane has 8 seats in first class it is important to make sure that there are 8 or fewer people with tickets to first class, and that those who were not granted tickets in fact know that they have been denied first class service. Admission control is based on this concept.

  Admission control consists of bandwidth control and policy control. Applications can request a particular QoS for their traffic. The devices in the network through which this traffic will pass can either grant or deny the request depending on various factors such as capacity, load, policies etc. If the request is granted, then the application has a contract for that service, which will be honored in the absence of disruptive events, such network outages or pre-emption policy. If the request is denied then the denial is communicated back to the application. The application can then either stop sending traffic into the network or request a different service.

  In the absence of admission control, the application may be unaware of the resource availability in the network and may attempt to utilize the service even if they are over-subscribed. The application may not get the service that it expects from the network and, in the absence of effective policing, will also cause the deterioration of the service that is being provided to other traffic. In our airplane example, if there are 16 people with first class tickets and all of them are seated into the available 8 seats, then the service to all the 16 passengers deteriorates and does not remain first class.

- **Policy Control**

  The network administrator manages and controls some of the devices in his administrative domain. This typically includes all of the routers, switches, and policy servers, but not all of the hosts. Those devices, which only the administrator can control, are considered to be *trusted* and lie within the trust boundary. They are configured to support common administrative objective for network usage.

  Since the devices outside the trust boundary use network resources, the network administrators have to monitor, control and enforce policies with respect to the usage of services in their network. Therefore, they set up policies derived from criteria such as the identity of the users and applications, traffic requirements, security issues, etc., to manage the network usage. Policies may be configured locally in the device or they may reside in external policy servers. These external policy servers are called Policy Decision Points (PDP) and are described in [7]. The PDPs make the policy decisions. Policies in a PDP usually take precedence over the policies configured locally. The policies are enforced in network devices that forward traffic, such as the routers and switches. When devices do so they are acting as Policy Enforcement Points (PEP).

Policies are used for policing as well as for making admission control decisions. Therefore, policy control is an important part of any QoS solution.

- **Policing**
  Policing is the function that ensures that traffic conforms to its contract. The policing function in a device must monitor, control and enforce the use of network resources with respect to its installed profile and configured policies. Non-conforming packets are either dropped outright or shaped into the profile. Such policing helps the device maintain its side of the agreed upon contract.

These primitives provide the functionality to implement QoS in the network.

4 QoS Architectures

The above QoS building primitives are used in the network to create different QoS architectures. Two of the most popular QoS architectures are Differentiated Services (DiffServ) and Integrated Services (IntServ). They can be used either separately, or more often together, to achieve most practical QoS objectives. Both of these originate from the Internet Engineering Task Force (IETF).

4.1 Differentiated Services

Differentiated Services (DiffServ) [1] [2] defines a set of QoS mechanisms for implementing scalable service differentiation in the Internet. DiffServ is a multiple service model, which means that it offers more than one kind of service. It uses in-band or implicit signaling to differentiate different classes of traffic. This means that the signaling or the marking is carried in the packet itself.

DiffServ defines a field in the IP header called the DiffServ Codepoint (DSCP[1]). Hosts sending traffic requiring QoS into a network supporting DiffServ mark each packet with a DSCP value. Devices within the DiffServ network use these values to classify the traffic into distinct service classes. Based on the configuration of the service class, the packets in the stream are queued and scheduled.

One of the most popular DSCP values is Expedited Forwarding (EF), which indicates that the packet requires low latency, low loss service. If a packet arrives at a router with an EF marking and passes the policing and shaping functions, then the router schedule it into the highest priority queue that it supports. This allows the packet to be forwarded with minimal queuing delay thus incurring very little delay and jitter. EF marking is commonly used by latency and loss sensitive applications, such as Voice over IP (VOIP), as they usually require a very stringent service from the network. DSCP values for videoconferencing traffic are discussed in section 5.2.

The advantages of DiffServ are:

- **Simplicity**
  It is comparatively simple to adopt. For any application, it is usually less effort to mark the packets than it is to introduce any additional signaling.

[1] The DSCP is a six-bit field, spanning the fields formerly known as the type-of-service (TOS) fields and the IP precedence fields.

- **Scalability**
  In this architecture, the network device does not process the traffic flow-by-flow. Therefore, it is not required to maintain any elaborate state information to identify the flows. Instead the arriving traffic is aggregated into a few traffic classes based on the DSCP marking. This does not result in a huge processing overhead. The device can therefore handle a large number of flows without experiencing any deterioration in its performance. Hence, it can be deployed without the scalability concerns that flow-based schemes incur.

Some of the disadvantages are:

- **Lack of admission control**
  This mechanism does not have a facility for admission control. It is possible to ask the network device for resources with the DSCP marking in the packet but there is no mechanism either for the device to ascertain whether this traffic would cause overloads or for the device to communicate its response back to the application. The application is unaware of the service it is given by the device and therefore must depend on observing its own end-to-end performance in order to adapt to the service it is actually receiving. In addition, by attempting to utilize the resources that are unavailable, it could disrupt the service that is given to other traffic in that class.

4.2 Integrated Services

Integrated Services (IntServ) [3] is also a multiple service model, but uses explicit signaling to request QoS. The QoS signaling message carries information that allows the network device identify the traffic and informs it of the exact service requirements for that traffic. In this model, network devices that participate maintain flow state and perform flow-based classification, policing, shaping, and scheduling. IntServ uses the Resource Reservation Protocol (RSVP) to describe traffic and perform admission control. RSVP, in addition to its use as the QoS signaling protocol for IntServ, also has a number of other uses, including as a traffic engineering protocol for MPLS networks, and as a signaling protocol to communicate DiffServ marking information to endpoints that do not know a priori how to mark packets for certain traffic classes.

4.2.1 Resource Reservation Protocol (RSVP)

RSVP [4] is a signaling protocol that is used to request resources from the network. It can request resources for either unicast or multicast flows. RSVP enables applications to communicate their traffic profile and to request a specific quality of service from the network based on the applications' bandwidth, loss and delay requirements. RSVP messages are constrained to travel the same path through the network as that of the traffic being described and therefore RSVP can reserve resources in the same devices that will process the traffic flow.

Application hosts usually generate these RSVP messages. RSVP messages contain the source and destination address and port and the protocol. They also contain the exact profile of the traffic and a request for a particular class of service for that profile. When a network device receives the request for reservation, it compares the request against the available resources. The device can also base its decision on policies such as the application type, the user, etc. If the device does not have the resources to grant the request, then it denies the request and sends a reservation failure message to the requesting application. If the device decides to grant the request, then the information in the reservation is installed in the device, and the RSVP message is passed along to the next device on the path. Once all the devices on the path have granted the reservation, the application can be assured that the resources it asked for are reserved for its use.

Once the reservation is granted the device ensures that the traffic receives the service it has been promised. The classifiers in the device use the stored information to identify and classify the traffic. The traffic-control elements use the information to perform the scheduling and policing functions. The scheduler places the traffic into a queue that can provide the requested service. The policer function ensures that the application's traffic conforms to the profile it has specified and either drops or lowers the priority of the non-conforming portion of the application's traffic.

The advantages of RSVP are:

- **Explicit admission control**
  Explicit admission control ensures that first class is really first class by not oversubscribing it. If the request is denied, then the denial is communicated to the requesting application as well as the upstream devices so that corrective action may be taken or the attempt to communicate abandoned.

- **Application coordination**
  This mechanism provides for the response to the reservation request to be communicated back to the requesting devices and other entities in the path. Applications are made aware of the network's decision with respect to its request for reservation. If the request is granted, then the application remains confident of getting the service it has requested. If the request is denied, then the application can take corrective actions such as attempting to reserve lower bandwidth, increasing its buffering, etc. By contrast, in a DiffServ model, an application can ask for priority but has no way of knowing whether its packets are being remarked to a lower priority or whether the priority it has asked for is oversubscribed.

The disadvantages are:

- **Scalability**
  One major concern that is often mentioned with respect to RSVP is scalability. To support RSVP, network devices have to perform additional processing to identify the flows based on the state it maintains. If the number of flows through the router is very large, then the performance of that router is adversely affected. Efforts are underway to address the scaling concerns in such high-speed backbone routers.

It is not sufficient for the application to merely send RSVP messages. It is also necessary for the application to react based on the RSVP response. To achieve this, RSVP signaling must be synchronized with the application's signaling.

4.3 IntServ together with DiffServ

The main advantage of DiffServ is its scalability and its main disadvantage is the lack of explicit admission control. In contrast, the main advantage of IntServ with RSVP is the ability to erform explicit admission control but it causes severe scaling concerns. RFC 2998: "A Framework for Integrated Services operation over DiffServ Networks" [10] suggests combining the two to gain the advantages of both the architectures. It describes how end-to-end IntServ QoS may be supported over DiffServ networks. The following is a short explanation derived from the RFC:

There could be DiffServ regions as well as IntServ regions in the network. RSVP signaling is initiated by the endpoints and travels end to end. RSVP messages are processed as usual in the IntServ region. The routers in the IntServ regions of the network are referred to as edge routers, because IntServ is effective at the scale one finds at the edges of the network. Edge routers at the boundary between the IntServ and DiffServ parts of the network perform admission control based on their own resource availability and policies of the adjacent DiffServ region. The RSVP message travels transparently through the RSVP-unaware routers in the DiffServ region. In the DiffServ region, routers do not maintain flow state nor do flow-based processing, hence, scalability is not a concern. The traffic obtains its QoS in the DiffServ region by its DSCP marking. This marking may have originated at the host (and communicated to it through RSVP) or at the edge routers. Alternatively, the edge routers use the information in the RSVP message to mark the packets appropriately. Thus, the traffic gets end-to-end QoS with admission control without requiring the entire network to support the IntServ architecture. This solution also addresses the scaling concerns in the network.

4.4 QoS Solutions for IP Videoconferencing

The goal of IP videoconferencing is to provide a quality audio and video experience to the end user. One of the challenges to achieving this goal is getting acceptable service from the network. The best tool that is available to meet this challenge is QoS. Currently, the H.323 gatekeeper is the one videoconferencing component for providing some QoS in the case of H.323-based videoconferencing. In this section we will discuss the admission control capabilities of a gatekeeper. Then we will explore how having the gatekeeper interact with the QoS policy server can enhance the solution. Finally, we will discuss the DSCP marking and RSVP signaling required to complete the solution.

4.5 Gatekeeper

One of the most important components in an H.323-based videoconferencing network is the gatekeeper. Even though the gatekeeper is an optional component it has become an integral part of videoconferencing deployments. Terminal and infrastructure components register with the gatekeeper. In the area of QoS, the gatekeeper's main function is admission control.

The logical grouping of H.323 components managed by a single gatekeeper is called a zone. A zone may span several IP subnets. Most gatekeepers perform admission control by managing a pool of available bandwidth. This bandwidth may or may not actually exist, since the gatekeeper is not aware of the actual resource usage in the network. Each call's bandwidth is subtracted from this pool when a call is placed, and added back in when it is hung up. If there is not enough bandwidth in the pool when a call is attempted then that call is rejected. Some gatekeepers have separate pools for calls within a zone and calls between zones. Gatekeepers are also free to accept and reject calls for policy reasons as well as bandwidth.

The admission control capabilities of the gatekeeper depend upon the implementation. Theoretically, gatekeepers can implement more sophisticated mechanisms for admission control. However, currently most gatekeepers just have a variation of the above model. Though useful, the admission control facility in gatekeepers as described above is inadequate for call admission control for the following reasons:

- **Gatekeepers' bandwidth control is static**
  Gatekeepers currently base their bandwidth decisions on statically configured bandwidth pools. The gatekeeper is not aware of how a particular call flows through the network. It

does not know which network devices will service the call and whether those devices have sufficient capacity.

- **Gatekeepers are only aware of H.323 traffic**
  Gatekeepers are only aware of H.323 traffic and have no knowledge of other types of traffic traveling through the network. Hence, the decisions made by the gatekeeper regarding bandwidth availability might not coincide with the actual bandwidth availability in the network.

- **Gatekeepers cannot perform policing**
  Gatekeepers cannot control the amount of bandwidth used by the endpoints. An endpoint can make a request for a call of 384 kbps and then send 400 or more kbps of data. If the data does not flow through the gatekeeper, which is most often the case, then it cannot enforce its decisions. Such endpoints make the subsequent admission control decisions less useful.

4.5.1 Gatekeeper and PDP Interaction

The admission control capability of existing gatekeepers therefore appears insufficient. The reason is that gatekeeper makes its admission control decisions in isolation and with static information. Admission control decisions should be made based on the capacity and current load on the devices as well as the policies of the network.

To make better policy control decision, the gatekeeper can interact with the policy system in the network. As explained in section 2 above, the policy servers (PDP) in the network contain global policies that apply to all traffic in the network. Before admitting a call, the gatekeeper should ensure that the call is in accordance with the network policies. It can do that by sending a message to the PDP and asking for its approval. If the PDP approves, then a message is sent back to the gatekeeper. The protocol that is typically used to communicate with PDPs is Common Open Policy Service (COPS) [7].

The other part of the admission control problem is the capacity and load on the devices in the network. In some networks, the PDP might also have complete knowledge of network topology. In such cases, with the help of the source and destination address, the PDP can guess the route as well as the devices that will service the call. The PDP could then ensure that the individual devices have the necessary capacity before permitting the call.

The PDP may be able to go a step further. In some cases, once it decides to permit the call, it might also be able to install the parameters of the call in those devices to enable it to identify the call traffic for prioritization. The devices can also use these parameters to police other traffic and deny other requests that conflict for the same resources. This way, the devices can guarantee the service to the call traffic. Such capabilities are present in PDPs, which support the COPS-PR [13] method of pushing policy state to network devices acting as PEPs.

The above is a solution that could be deployed for providing QoS for IP videoconferences. However, there are some practical problems. It is possible for the PDPs to have sufficient topology knowledge only in very small, simple and homogenous networks, and then they must participate in the routing algorithms themselves to do so. Therefore, the above is not a practical solution in most networks. In addition, the call may span multiple networks with different QoS domains. The PDP can only check on policies for the call within its own QoS domain and usually does not know the route or the individual devices. Hence, other mechanisms may also have to be utilized.

The capacity in the network devices and the policies set up for the network will ultimately determine the service that the call traffic will get. As mentioned above, it is difficult for the PDP to identify and interact with the individual devices that will service the call in all networks. Having QoS signaling along the route of the call that will reach all the necessary devices can solve this problem. Such a QoS solution must include mechanisms to describe the traffic and provide sufficient information so that the network devices can perform classification, policing, shaping, and scheduling. Such signaling solutions are described below:

4.6 DSCP marking

One way to reach the network devices in the path of the call is by using the DSCP marking. The devices must prioritize videoconferencing traffic over best effort traffic to minimize delay, jitter and loss. The endpoint must help these devices identify low latency, low loss traffic by marking it an appropriate DSCP value.

At this date there are no standards describing which value is most appropriate for IP videoconferencing. Packets marked with Expedited Forwarding (DSCP value: 101110[2]) [6] receive the highest level of service from the network device. These packets are usually put it into a high priority queue. The packets in this queue are serviced ahead of other packets and therefore get the least delay and jitter. The high priority queues are usually kept very small, in order to control delay and to prevent starvation of lower priority traffic. This may result in dropped packets if the queue is full. EF is commonly considered most appropriate for VOIP.

Assured Forwarding (AF) [5] is another means of providing better than best effort handling for an IP packet. There are 4 Assured Forwarding classes. AF41 (DSCP value: 100010[3]) seems to provide a good match for video traffic with respect to delay, jitter and loss. AF41 is therefore considered the most suitable for video traffic. As there is no advantage in treating voice packets better than the video packets in an IP videoconferencing application, AF41 should be used as the DSCP value for both voice and video media in a videoconference.

The EF marking is not a good choice for use with video conferencing for these reasons:

- Video packets are much larger than voice packets, usually as large as the maximum link MTU size. If video packets are marked as EF then they will be allowed into the same priority queue as voice. If a small VOIP packet enters the queue behind a large video packet (or worse, a whole bunch of such packets), then the delay in the VOIP packet may increase depending upon the traffic engineering in the network. Such delay will adversely affect the performance of VOIP applications.
- Given that most EF queues are very small, using them for video traffic may lead to dropped packets.
- Video coders tend to have much longer coding delays than voice coders, and hence giving the audio stream(s) of a videoconference absolute priority only causes them to arrive early and be held in order to achieve lip sync. Therefore, it does not help to put voice packets associated with a videoconference in a queue with better service than that given to the video packets.

[2] The DSCP Value EF coincides with the IP Precedence value 5.

[3] The DSCP value AF41 coincides with the IP Precedence value 4.

Signaling traffic, such as H.225.0 and H.245 traffic, may also need better than best effort treatment to ensure the continuity of videoconference signaling on an overloaded network. Hence, it is desirable to mark such signaling packets. The suggested DSCP value is Class 3 (DSCP value: 011000[4]).

DSCP marking is simple to implement. Its main advantage is that it enables the network devices to easily identify and prioritize certain traffic. If all the devices along the path support this mechanism, then it results in an end-to-end QoS solution.

While DiffServ alone solves the problem of prioritizing video traffic over best effort traffic, it does not solve the problem of video traffic oversubscribing its queue. To take a concrete example, imagine a remote office connected by a T1 to a central office. Two users are engaged in 384 kbps videoconferences and another user starts copying large files from the remote office. Since the video traffic is prioritized over the data traffic, the videoconferences are not disrupted. In addition, the file copies are able to take place at a reasonable pace. But imagine someone now tries to place a 768 kbps call. Each call has some protocol overhead as well. Hence the bandwidth consumed now by the three calls is greater than the T1 link. The first thing that happens is that the file copies fail because the high priority video traffic has taken up all of the available bandwidth leaving best effort data (if a strict priority scheduler is used to implement the AF/EF DiffServ classes). The second thing that happens is the video calls start interfering with each other because they have oversubscribed the available bandwidth. Thus, prioritizing traffic only solves part of the problem. As noted numerous times above, a complete solution requires admission control as well.

4.7 Admission Control with RSVP Signaling

Endpoints can ensure explicit admission control by implementing RSVP and properly synchronizing it with the application signaling protocol. An application uses RSVP messages to describe its traffic profile and request the service that it requires. The RSVP message travels to all of the network devices along the path of the traffic flow. Each RSVP-capable device sees this message and decides whether to grant or deny the request. Any decision to deny the request is communicated to the requesting application in an RSVP message. This communication enables the application to adapt to the available network conditions. An application can either decide to discontinue the conference or reduce their requirements.

It is desirable that applications either wait to send traffic or send it as best effort until its request has been granted. While transmitting traffic, the application must also not exceed the traffic profile that it had advertised earlier. Any excess traffic will be policed and will not be handled as the application desires. The traffic could be dropped, or handled as best effort. In any event, this will perturb arrival order, jitter, and loss at the receiver and may result in significantly worse performance than if the application had simply used best effort transmission. Likewise, if the network device has denied the requested service for a traffic flow, then that flow will be policed and not given any priority. Such policing enables the network device to maintain the guarantees offered to other traffic flows.

IntServ QoS currently offers two kinds of service: *Guaranteed* and *Controlled load*. Guaranteed service means that the network device can carry that traffic with a measurable or bounded delay. Controlled Load service means that the device can carry that traffic as if the device is

[4] The DSCP value AF31 coincides with the IP Precedence value 3.

lightly loaded. Guaranteed service is useful primarily for applications which have no ability to adapt their reception to any delay variation in the network (e.g. CBR with no jitter buffer). Since all known videoconferencing systems are capable of adapting to substantial delay variation, guaranteed service is not needed. Therefore, it is recommended that IP videoconferencing request *Controlled Load* for its traffic.

The main advantage of RSVP signaling is that it not only provides prioritization but also provides explicit admission control and service guarantees. Thus, it satisfies all the requirements mentioned in section 5.2. Endpoints can request and get the service that they require from the network. The call traffic is guaranteed the service that it has requested or the endpoint receives a message indicating that the service requested is unavailable. RSVP messages also allow for policy information to be carried to the devices. These devices can then consult with the PDP for policy control using the COPS for RSVP protocol [9]. These devices maintain state from the RSVP signaling and use that information for scheduling and policing. Such policing prevents oversubscribing any service class and enables the devices to provide guarantees.

All of the above benefits come with a cost. Since this is flow-based model, it requires a flow state to be maintained and flow-based processing to happen. This consumes more resources than the aggregate model in the DiffServ architecture. Hence, the network administrators are sometimes reluctant to enable RSVP in all parts of the network, especially in the core of large network where devices handle millions of flows. Aggregation of RSVP [11] and other work in the IETF are expected to substantially alleviate the problem.

4.8 Putting it all together

In the above sections we have examined some network architectures for QoS. We have also discussed the applicability of these QoS solutions to IP videoconferencing and their pros and cons.

There are some parts of the network where explicit admission control and flow-based processing are required, such as the edge of the trust boundary, edge of WAN links etc. There are other areas of that network such as at the high-speed backbones, where there is no necessity for flow-based differentiation, and QoS can be provided to aggregated traffic. DiffServ might also be chosen for other areas in the network where scaling is an issue. Hence, network administrators might want to combine the IntServ and the DiffServ architectures in the single network. Any QoS solution for IP videoconferencing should be capable of providing end-to-end QoS in any combination of IntServ and DiffServ networks. This has been described in RFC 2998 and discussed in section 4.3 above.

Using DiffServ and IntServ together does avoid the limitations of relying on gatekeepers for admission control since the network devices along the path consult the PDP if required. However, the original gatekeeper’s QoS capabilities can be enabled to provide a coarser granularity or control to the QoS solution, especially when policy is based on aggregate call rates between different administrative domains. This may take into account capabilities not known to the network devices such as gatekeeper zones, endpoints capabilities etc.

An end-to-end QoS solution for IP videoconferencing needs endpoints to support RSVP. The endpoints or network devices can do the marking of the packets. Without RSVP there is inadequate admission control.

5 Synchronization of RSVP with Call Signaling

It is not enough for the endpoints to just generate RSVP signaling. It is also necessary to coordinate the functioning of the application with the RSVP signaling. The application should react to the information received in the RSVP messages to reap the full benefits. For instance, if the reservation is denied and the application continues without any changes, then the user will not only experience poor quality, but might attribute that poor quality to the application and not to network congestion. Such problems have to be avoided. Therefore it is essential to coordinate RSVP with the application's signaling.

Many IP videoconferencing applications use the H.323 suite of protocols for their signaling. H.323 is globally accepted as a standard for multimedia conferences in an IP network. The International Telecommunication Union (ITU) approved the first version of the H.323 standard in 1996. The current version is 4. Appendix II of Version 4 of the H.323 outlines an approach for using RSVP to reserve resources. The main points are:

- When placing a call, an endpoint communicates its ability to reserve resources to the gatekeeper. The gatekeeper can respond with whether or not it would like the endpoint to attempt resource reservation.
- In the H.245 phase the endpoints communicate whether they are capable of resource reservation. Using this knowledge they can than make a decision as to whether to proceed with the call.
- Once the logical channels have been opened (but before sending traffic on them) RSVP reservation messages can be sent.

The problem with the above is that by the time the reservation has been established, the far endpoint has alerted the user that there is an incoming call. If the reservation fails, releasing the call after alerting the far end creates a poor usability experience. Appendix II provides a method for reordering the protocol messages so that the application can know whether the reservation succeeded before alerting the far end. Since the far end has not been alerted, the calling application can fail the call outright by informing the user that "all circuits are busy at this time" thus protecting the other calls on the network from disruption.

A document providing details that will help with implementation will soon be available from Cisco Systems.

Moreover, Annex N of the H.323 spec will cover all aspects of providing QoS for H.323-based IP videoconferencing. This Annex is currently being worked on.

6 Summary

Network QoS is an important component for successful IP videoconferencing. With QoS, the network administrator is able to control the resource consumption of this application as well as provide acceptable network service to the end users. It is important that any QoS solution for IP videoconferencing should support all of the elements of network QoS, including classification, policing, shaping, scheduling, and admission control as well as mechanisms to support the admission control decision such as traffic control.

7 References

1. Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers (RFC 2474),
   http://www.ietf.org/rfc/rfc2474.txt

2. An Architecture for Differentiated Services (RFC 2475)
   http://www.ietf.org/rfc/rfc2475.txt

3. Integrated Services in the Internet Architecture: an Overview [RFC 1633]
   http://www.ietf.org/rfc/rfc1633.txt

4. Resource Reservation Protocol (RSVP): [RFC2205]
   http://www.ietf.org/rfc/rfc2205

5. An Assured Forwarding PHB [RFC 2597]
   http://www.ietf.org/rfc/rfc2597.txt

6. An Expedited Forwarding PHB [RFC 2598]
   http://www.ietf.org/rfc/rfc2598.txt

7. The COPS (Common Open Policy Service) Protocol (RFC 2748)
   http://www.ietf.org/rfc/rfc2748.txt

8. COPS usage for RSVP (RFC 2749)
   http://www.ietf.org/rfc/rfc2749.txt

9. RSVP Extensions for Policy Control
   http://www.ietf.org/rfc/rfc2750.txt

10. A Framework for Integrated Services Operation over Diffserv Networks (RFC 2998)
    http://www.ietf.org/rfc/rfc2998.txt

11. Aggregation of RSVP for IPV4 and IPV6 reservations
    Draft-ietf-rsvp-aggr-01.txt, Fred baker et.al.

12. A Framework for Policy-based Admission Control
    Draft-ietf-rap-framework-03.txt
    R. Yavatkar, D. Pendarakis, R. Guerin.
    April 1999

13. COPS Usage for Policy Provisioning,
    F. Reichmeyer, S. Herzog, K.H. Chan, J. Seligson, D. Durham, R. Yavatkar, S. Gai,
    K. McCloghrie, A. Smith.
    June 2000