# Quality of Service in Data Networks: Products

**Gautam Ray**, *ray.129@osu.edu*

## Abstract:

IP only provides best-effort service. This paper highlights the importance of quality of service (QOS) in data networks and reviews four important standards for providing quality of service; ATM, IntServ, DiffServ, and MPLS. The paper then compares ten contemporary products in terms of how they provide quality of service.

See Also : QoS in Data Networks: Protocols and Standards | QoS/Policy/Constraint based routing | QoS over Data Networks (Prof. Jain's Lecture) | Quality of Service over IP: References | Books on Quality of Service over IP
Other reports on recent advances in networking
Back to Raj Jain's Home Page

# Table of Contents

# 1.0 INTRODUCTION

Quality of Service (henceforth QOS) is described in CCITT Recommendation E.800 as: “The collective effect of service performance, which determines the degree of satisfaction of a user of the service”. QOS, which is closely related to and is sometimes referred to as Policy-based Networking, lets network managers define service policies that govern how much bandwidth goes to specific applications and end users. QOS makes it possible to implement policies across devices on LAN and IP networks. Converting best-effort IP into QOS capable network could mean huge savings for corporations. It would allow network managers to guarantee that revenue generating traffic gets more bandwidth compared to e-mail and casual web surfing, all with their current infrastructure.

Before delving into the details of QOS, an important question must be answered first: why should we be concerned with

QOS at all when most networking problems can be solved by increasing capacity? The main point is that data is inherently bursty, no matter how high the capacity, congestion will always occur for short periods. Another consideration is that routing protocols don't know about load levels, so congestion will build up on some paths while others have bandwidth to spare. There is also a speed mismatch at the LAN/WAN border. Even with faster WAN circuits, traffic from LAN is likely to congest the wide area link. Finally, bandwidth alone doesn't ensure low and predictable delay as even with huge bandwidth, there is still the danger that real-time applications will get stuck behind large file transfers. Therefore, QOS mechanisms are important because they enable networks to deliver defined levels of service with the existing network infrastructure.

The rest of this paper is organized as follows. In section two the important standards for QOS: ATM, IntServ, DiffServ, and MPLS are described. In section three, ten contemporary QOS products from Allot Communications, Alcatel, AVICI Systems, Cabletron, Cisco, Extreme Networks, IP Highway, Nortel Networks, Ukiah Software, and Xedia are compared on the basis of how they implement QOS. Section four concludes the paper.

Back to Table of Contents

# 2.0 PROTOCOLS AND STANDARDS

This section describes four protocols and standards for providing QOS. We start by defining the characteristics that are controlled to provide a specific QOS and the two generic methods for providing QOS. The two important properties that define QOS are: (i) delay, and (ii) throughput [Trillium98].

(i) Delay: Two aspects of delay have an impact on QOS: end-to-end delay, and delay variation or jitter. Interactive real time applications (e.g., voice communication) are sensitive to end-to-end delay and jitter. Long delays reduce the interactivity of the communication. Non-interactive real time applications (e.g., one way broadcast) are not sensitive to end-to-end delay but are affected by jitter. Jitter is usually accommodated by using a buffer at the receiver where the received packets are stored and then replayed at the appropriate time offset. Non real time applications are not delay sensitive.

(ii) Throughput: Throughput is the amount of bandwidth available to an application. This determines how much traffic an application can get across the network. Throughput is also influenced by the link error rate and losses (often related to buffer capacity).

The fundamental idea in QOS is that traffic can be differentiated and provided different levels of service. The granularity of differentiation can be a small set of classes or individual flows. The amount of traffic allowed into the network is also controlled based on the available resources. There are two generic methods for providing QOS: (a) Reservation-based, and (ii) Reservation-less [Trillium98].

(a) Reservation-based: In this model, resources are reserved explicitly. The network classifies incoming packets and uses the reservations made to provide a differentiated service. Typically, a dynamic resource reservation protocol is used, in conjunction with admission control, to make reservations.

(b) Reservation-less: In this model, no resources are explicitly reserved. Instead, traffic is differentiated into a set of classes, and the network provides services to these classes based on their priority. However, it is necessary to control the amount of traffic in a given class that is allowed into the network, to preserve the quality of service being provided to other packets of the same class.

There are four main standards for providing QOS: (i) ATM, (ii) IntServ, (iii) DiffServ, and (iv) MPLS [Seaman99].

## 2.1 ATM

ATM, using a reservation-based model, allows two end-points to negotiate session parameters and specify such metrics as acceptable cell loss rate, minimum throughput, and maximum jitter and latency. The ATM network, within constraints of overall capacity, can therefore adapt to the precise needs of different traffic. The following features enable ATM to provide QOS:

Virtual Circuits (VCs): VCs provide a mechanism for ATM to differentiate traffic and classify packets. The VPI/VCI field in the packet headers is an efficient mechanism to classify packets.

Service Categories: ATM provides a number of service categories such as Constant bit rate (CBR), Variable bit rate (VBR), Available bit rate (ABR), and Unspecified bit rate (UBR), which provide different levels of service with respect to delay and throughput.

QOS Signaling and Routing: ATM provides dynamic signaling and routing protocols for setting up resource reservations. Signaling messages carry the traffic and QOS characteristics, which are used by the network to perform admission control and to route the reservation request along a path that is likely to satisfy the required QOS.

IP QOS Mechanisms - Integrated Services (IntServ) and Differentiated Services (DiffServ)

There are two main QOS mechanisms in IP: IntServ and DiffServ. IntServ is a reservation-based mechanism. It reserves resources explicitly for individual flows using a dynamic signaling protocol and employs admission control, packet classification, and scheduling to achieve the desired QOS. On the other hand, DiffServ is a reservation-less mechanism. It classifies packets into a small number of service types and uses priority to provide QOS to the traffic. No explicit resource reservation or admission control is used, though the network has to use some method to differentiate traffic.

## 2.2 Integrated Services (IntServ)

IntServ is a flow oriented QOS mechanism. A flow is defined as a stream of packets that originate from a specific user activity e.g., a single application session. A flow may be identified by a variety of mechanisms: IPv4 uses source and destination IP address and the destination port number. IntServ reserves bandwidth for individual flows to provide QOS. Two main service categories are defined based on the delay and loss requirements.

(i) Guaranteed Delay service provides absolute guarantees on the delay and loss experienced by a flow. Packets are not lost, nor do they experience delay exceeding the specified bound. These firm guarantees are provided using resource reservations.

(ii) Controlled Load service provides service equivalent to that of an unloaded network. Most packets are not lost, nor do they experience queuing delay. However, no specific quantitative guarantees are provided.

Resource Reservation

IntServ requires that resources be reserved for flows in order to provide the requested QOS. This can be done via a dynamic reservation protocol, manual configuration, or by using a network management protocol. IntServ is not tied to any specific mechanism. However, RSVP is a protocol designed to provide resource reservations and it is designed to work with IntServ, though it can be used with other service models as well.

Two important characteristics of RSVP are:

- Receiver oriented: The protocol requires receivers to make reservations. Receivers request resource reservations based on senders traffic specifications and path characteristics.
- No built-in mechanisms for routing or packet scheduling: RSVP is just a signaling protocol. It relies on IP to compute the reservation route. RSVP is also not concerned with how nodes implement the reservation requests (e.g., admission control, packet classification, packet scheduling).

IntServ over ATM

Another way that IntServ can deliver QOS is through integration with ATM networks. The Guaranteed Delay service can be mapped to Constant bit rate (CBR) or real time Variable bit rate (rt-VBR) ATM service categories. Both provide real-time delivery and can provide predictable delay characteristics. Similarly, the Controlled Load service can be mapped to Constant bit rate (CBR), non-real time Variable bit rate (nrt-VBR), or Available bit rate (ABR). Available bit rate (ABR) with a positive Minimum cell rate (MCR) can also be used, but the available bandwidth may vary.

Limitations of IntServ

The most important concern about IntServ is whether it can scale to large backbones.

(a) The number of individual flows in a backbone network can be very large. The number of control messages for making resource reservation for large number of flows can be large and may require a lot of processing power. Similarly, maintaining state information for all the flows can require a lot of storage capacity. There is also a need to classify a large number of packets and schedule numerous queues making the router extremely complex.

(b) Policy issues need to be resolved to determine who can make reservations. Similarly, security issues need to be resolved to ensure that unauthorized sources do not make spurious reservations.

It is believed that IntServ is appropriate for small intranets where there are a small number of flows and where policy and security issues can be managed easily. Large backbone networks will need more scalable mechanisms for differentiating traffic and providing differentiated services to them.

## 2.3 Differentiated Services (DiffServ)

The key features of DiffServ that overcome some of the limitations of IntServ are:

- Coarse Differentiation: DiffServ does not differentiate per flow traffic. Instead there are a small number of well defined classes which are provided differentiated services.

- No Packet Classification in the Network: In RSVP, each router implements packet classification in order to provide different levels of service. To make the solution more scalable, in DiffServ, packet classification is moved to the edge of the network. Edge routers classify and mark packets appropriately. Interior routers simply process packets based on these markings. This implies that interior routers do not recognize individual flows, instead, they deal with aggregate classes.

Static Provisioning

IntServ requires dynamic resource reservation. This may result in a large number of control packets. It also requires dynamic admission control in each router. DiffServ moves admission control to the edge of the network. It also does not require dynamic reservations. Instead, it uses long-term static provisioning to establish service agreements with the users of the network, whose traffic it can police at the ingress to the network.

No Absolute Guarantees

The guaranteed delay service in IntServ provides hard bounds on the delay experienced by packets by explicitly reserving resources along the path. DiffServ, in general, does not provide hard guarantees. The goal in DiffServ is to monitor the traffic that enters the network at the ingress node and check for compliance against some predefined service profiles. Based on this, packets can be marked as being "in" or "out" of their profiles. Inside the network, routers preferentially drop packets that are tagged as being “out”.

In addition to drop precedence, there is additional information in the packet headers that communicates the type of service (TOS) desired by the packet. For instance, a premium service can offer the equivalent of a Constant bit rate (CBR) connection. Similarly, another example is that of an assured service that is characterized by bursty behavior and is provisioned using expected capacity; hence its bandwidth is allocated statistically.

Service Profiles

A service profile indicates which traffic is to be treated differentially and what type of service is requested. The former may be indicated by setting a packet filter based on packet header fields such as IP addresses. The latter can be specified using a token bucket filter. The service profiles are set up at the edge nodes based on customer subscriptions. They are therefore relatively static. The decision to accept a new subscription can be made centrally based on knowledge of network topology and capacity. Thus a network provider can provision its network according to the expected demand and subscriptions.

Packet Classification and Marking

The ingress router must check all received packets against service profiles to check if a packet should receive differential treatment. Packets that do not meet service profiles can either be discarded or sent into the network with higher drop precedence. The source can also police and shape the traffic it is offering to the network in order to maximize the probability that the traffic will meet the service profile and receive the desired quality of service. The ingress router marks the packets as they enter the network so that interior routers can handle the packets differentially. The marking use header fields, for example for IPv4 packets, the TOS octet is used.

Differential Queuing

Differentiated packets have to be handled differently. In order to do do so, the interiors router employ multiple queues with Class Based Queuing (CBQ). Generally, delay-sensitive traffic needs to be serviced sooner, and loss-sensitive traffic needs to be given larger buffers. The loss behavior can also be controlled using various forms of Random Early Detection (RED). These methods use probabilistic methods to start dropping packets when certain queue thresholds are reached, in order to increase the probability that higher priority packets can be buffered at the expense of more dispensable packets. For example, packets of different service types are put into different queues, and within a given service type, packets with higher drop precedence are discarded earlier than those with lower drop precedence.

DiffServ over ATM

Another way that DiffServ can deliver QOS is through integration with ATM networks. A number of Virtual Circuits (VCs) can be set up between adjacent routers, say, one VC per traffic class. The VCs can be provisioned as Constant bit rate (CBR) with bandwidth proportional to the amount of traffic of each class that is expected to be carried. In this manner, the ATM network will treat the traffic classes differentially. It is also possible to use Variable bit rate (VBR) VCs in order to provide some statistical multiplexing gains. If Variable bit rate (VBR) VCs are used to connect routers and two levels of precedence are applied, cells in excess of their profiles could be tagged. Tagged cells can be handled in a best-effort manner.

## 2.4 Multi-Protocol Label Switching (MPLS)

The MPLS approach to IP QOS is different from DiffServ. MPLS uses fields in the 32-bit (4-byte) label it adds to the IP packet. This label is intended to improve efficiency of the router network and allow routers to forward packets using predetermined paths according to, among other things, specified QOS levels. At the edge of the MPLS network, a label is added to each packet containing information that alerts the next hop MPLS router to the packet's predefined path. As the packet traverses the network, it may be relabeled to travel a more efficient path. Upon leaving the MPLS network the packet is stripped of its label and restored to its original size.

MPLS attempts to set up paths in a network along which packets that carry appropriate labels can be forwarded very efficiently since the forwarding engine would not look at the entire packet header, rather only at the label and use that to forward the packet. This not only allows packets to be forwarded more quickly, it also allows the paths to be set up in a variety of ways: the path could represent the normal destination-based routing path, a policy-based explicit route, or a reservation-based flow path. Ingress routers classify incoming packets and wrap them in an MPLS header that carries the appropriate label for forwarding by the interior routers. The labels are distributed by a dynamic Label Distribution Protocol (LDP), which effectively sets up a Label Switched Path (LSP) along the Label Switched Routers (LSR).

Weue a packet enters the MPLS network, a Label Switched Router (LSR) may analyze the IP header to determine its desired service level. As with addressing, the MPLS network will read this information only once. The label also contains 3 bits, called experimental bits, that may be used for specifying QOS. These bits will permit the Label Switched Routers (LSRs) to examine a packet's required service level and handle it accordingly. MPLS also permits explicit routing, where the hops a packet will take are specified in advance and the label is used to indicate this route. Explicit routing is a useful capability for allowing QOS and enabling network managers to set up defined paths through the MPLS network that apply to certain traffic streams.

MPLS can also deliver QOS through tight integration with ATM. In MPLS-over-ATM networks, an MPLS tunnel could be mapped directly to an ATM VC (virtual circuit) with an appropriate service level for the traffic. In other words, QOS information in the MPLS label could specify whether the traffic requires a Constant bit rate (CBR) or Variable bit rate (VBR) VC, and the Label Switched Routers (LSR) could send the packet to the appropriate VC.

In this section we reviewed four important standards for providing QOS: (i) ATM, (ii) IntServ, (iii) DiffServ, and (iv) MPLS. In the next section we compare how ten different vendors implement QOS.

Back to Table of Contents

# 3.0 PRODUCTS AND SERVICES

There are three core components to a QOS system: the policy console, the policy server, and policy-enabled switch or router (see Figure 1). Network managers use the policy console to set QOS policies. The policy console typically provides a graphical interface to enable the network manager to assign traffic to predefined service classes (bronze, silver, gold, and so on). The policy server is in turn responsible for telling switches and routers how to handle different types of traffic in order to meet the specified QOS levels. Policy servers define policies that are then enforced by switches and routers.

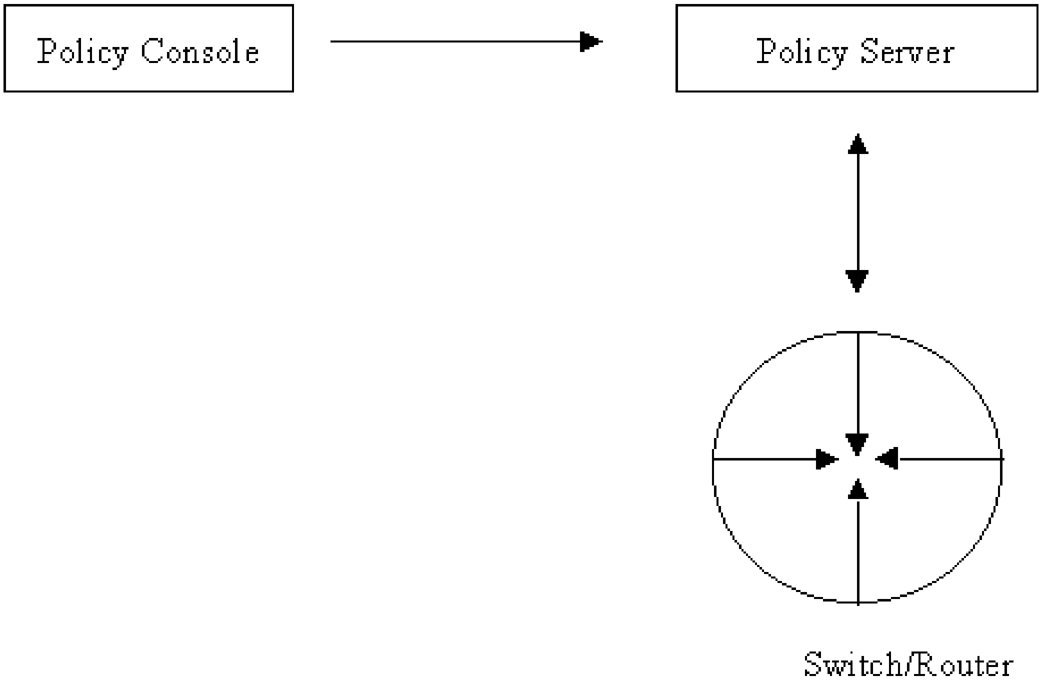

**Figure 1: Basic Components of a quality of service system.**

In this section we compare products from 10 different vendors: Alcatel, Allot Communications, AVICI Systems, Cabletron, Cisco, Extreme Networks, IP Highway, Nortel Networks, Ukiah Software, and Xedia. The products are compared on the basis of how they classify traffic, whether they implement QOS using a standalone or embedded server, what are the policy propagation mechanisms, where do they implement policy – at the edge or end-to-end, and how do they queue, shape, and police traffic [Saunders99].

## 3.1 Standalone or Embedded Server

Where the policy server is positioned can affect scalability, cost, and performance. In the embedded approach each switch or router has its own policy server. Embedding the server into each switch means that products will scale better. With standalone policy servers, far more devices to have to be managed and far more devices that could fail. Another scalability concern is that standalone servers can only distribute policy to a finite number of network devices before they become a bottleneck. This raises the question, how many devices can a standalone policy server serve? Some believe a safe ratio is five switches to one server while others thinks 500 may be the maximum. Some vendors contend that scalability problems prompted them to deploy a two-tier architecture, in which a central policy server interacts with distributed policy agents. Embedded servers also have the advantage that they avoid slowdowns since their switches don't have to communicate with standalone servers as embedding the policy server reduces the overhead associated with the switch-to-server exchanges. Alcatel, AVICI Systems, Extreme Networks, Xedia implement QOS with embedded policy server.

The advantage of standalone servers is that they reduce the number of policy profiles. In the embedded approach, each switch maintains its own profile, and all those polices must be synchronized. Standalone servers also enable network managers to support QOS on networks built with older, cheaper networking devices like aging routers and LAN switches that don't have the memory or processing power to run policy software themselves. Allot Communication, Cabletron, Cisco, IP Highway, Nortel, and Ukiah Software implement QOS with a standalone policy server.

## 3.2 Classification

Classification enables network managers to pick out a particular application and give it differentiated service - more bandwidth, lower latency, and so forth. The simplest classification mechanism is IEEE 802.1p, which sets the "P" bit in the Virtual LAN (VLAN) packet header. Since it works at Layer 2, the scheme is cheap to implement, and it doesn't matter what protocols are running on the network. The downside is that 802.1p only allows rudimentary prioritization. It also can't run over a wide area network. Switches and routers on the other hand classify traffic based on source and destination IP address, TOS (type of service) bit in the IP header, and Layer 4 criteria such as Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) port number, as virtually all popular IP applications use distinct and well known port numbers. Products that tag the TOS bit can identify priority traffic from hop to hop without adding anything new to the IP packet header. Some routers also filter on such higher-layer criteria as the DNS (domain name system), the application, or the URL. This is important because port number by itself doesn't distinguish between a customer order transmitted over HTTP and employees surfing the Web.

Alcatel, Allot Communications, Cabletron, Cisco, Extreme Networks, IP Highway, Nortel Networks, and Xedia classify traffic on the basis of source and destination IP address as well as TCP/UDP port numbers. Alcatel, Cabletron, Cisco, Extreme Networks, Ukiah Software also classify traffic on the basis of 802.1p. Allot Communications, Nortel Networks, Ukiah Software, and Xedia also classify traffic on the basis of application layer criterion such as DNS or URL. Alcatel, Allot Communications, IP Highway, Nortel Networks, and Ukiah Software also classify traffic on the basis of time of the day. Allot Communication, AVICI Systems, Cabletron, Cisco, Extreme Networks, Nortel Networks, Ukiah Software, and Xedia classify traffic on the basis of DiffServ (TOS) byte.

## 3.3 Where do we implement QOS Policy?

Where should QOS policy be enabled: at the LAN/WAN border or all the way to the end-station? The devices policies are enforced on, affects how much of a network can be brought under the rule of QOS. The decision to target the LAN/WAN demarcation is reasoned by pointing out that they make the most of scarce, expensive wide-area bandwidth. When faced with the decision to implement QOS in the LAN, as opposed to investing in more bandwidth, the second option is often preferred. QOS assumes significance at the LAN/WAN boundary, where bandwidth is more expensive. Allot Communications, Cabletron, Cisco, AVICI systems, Nortel, and Xedia implement QOS on edge devices.

The advantage of end-to-end QOS is that policies are assigned to applications before they get onto the network, thus delivering much more granular control over traffic. If applications are not policy enabled at the end-node, the traffic must reach the network device before policies can be enforced. Enabling policy in the end-station also offers the ability to identify applications that flow through dynamically assigned ports. Most policy-enabled switches and routers employ two items to identify traffic: source and destination IP addresses and an application's Layer 4 TCP/UDP port number. This approach works very well for applications like FTP that use known port numbers, but does not work very well for applications like H.323 video that don't have known port numbers. Enabling policies in the end-station solves this problem. Alcatel, Extreme Networks, IP Highway, and Ukiah Software provide end to end QOS.

## 3.4 How QOS policies are sent over the network?

The next question is what are the propagation mechanisms used to distribute policies or keep them in synchronization. Commonly used protocols are Simple Network Management Protocol (SNMP) and Lightweight Directory Access Protocol (LDAP). SNMP is good for doing simple things, but it's unreliable, it’s not secure, and it’s not fast enough to do bulk data transfers. LDAP could also exhibit similar drawbacks, since it was designed as a simple query mechanism for accessing X.500 directories. Some vendors employ Common Open Policy Service (COPS), a new IETF specification developed specifically for QOS. The limitation of Common Open Policy Service (COPS) is that it’s a per-flow protocol. This makes it difficult to deploy it in high speed networks because the switch has to ask the policy server every time it sees a new flow.

COPS is used by Allot Communications, Cisco, Extreme Networks, IP Highway and Nortel Networks. AVICI Systems, Cabletron, and Ukiah software also use SNMP. Xedia uses only SNMP. Allot Communications also uses LDAP. Alcatel only uses LDAP. AVICI Systems, Cabletron, Extreme Networks, Nortel Networks, and Ukiah Software also use Command Line Interface (CLI).

## 3.5 Queuing, Scheduling, and Shaping

Switches and routers mark traffic and use internal queuing mechanisms to enforce policy so that their counterparts on the network can also enforce policies. Most switches and routers enforce priorities by assigning packets from different streams to different queues (see Figure 2). Different types of traffic are held in different queues and high priority traffic is moved on to the network before low priority traffic. Queuing is also sometimes combined with explicit TCP/IP rate control, which involves the adjustment of window sizes and insertion of acknowledgments to control transmission rates from the source.

There are basically three types of queuing: priority, weighted, and class-based queuing. Products using priority queuing (PQ) classify traffic and set policies for high and low priority data. The high priority queues have to be emptied before lower priority traffic is transmitted. In other words, traffic is sent according to importance only and there are no guaranteed amounts of bandwidth. This approach works well for bursty traffic, but if policies aren't properly set then low priority traffic can be starved of bandwidth. This could lead to dropped packets and retransmission making the congestion problems worse.

Another queuing approach is weighted fair queuing (WFQ). It categorizes traffic into high and low bandwidth flows. WFQ first ensures that there is enough capacity available for the low-bandwidth flows, and then splits the rest among the large-bandwidth flows. A primary goal of WFQ is to avoid bandwidth starvation for low-priority traffic. Thus in WFQ, traffic is not only assigned to specific priority queues, it is also apportioned a share of bandwidth. For instance, a device can be informed that one type of traffic is twice as important as another and thus it should send two packets from the high-priority queue for every one packet it sends from the lower priority queue. Finally, there is class-based queuing (CBQ). In this approach, the queue itself is guaranteed a certain transmission rate. If the queue doesn't use all of its bandwidth, traffic from other classes can borrow as needed.

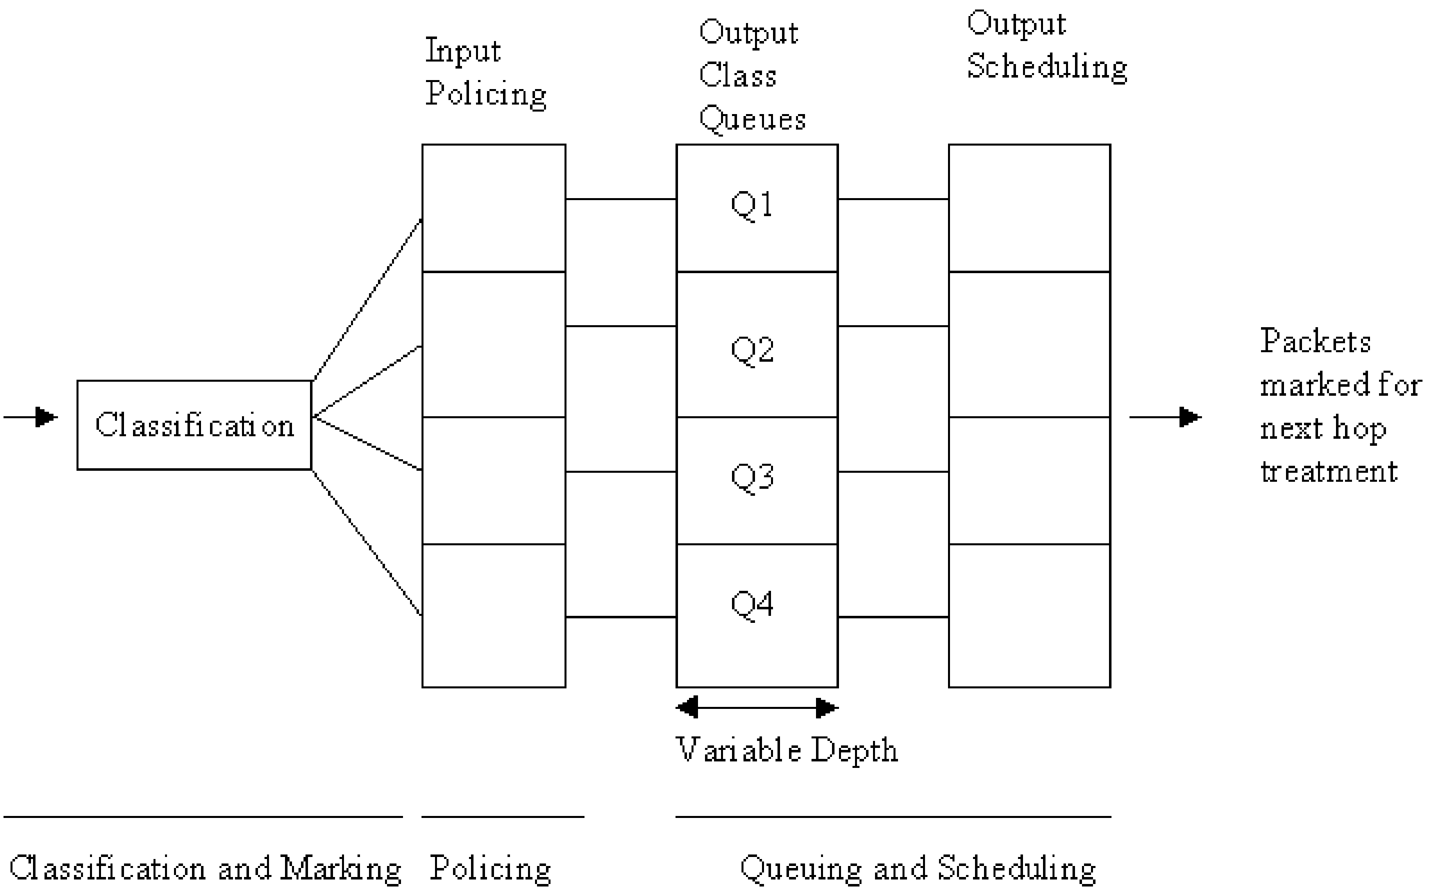

**Figure 2: Quality of service functional model (figure courtesy [Nortel Networks 98])**

QOS products can also make use of TCP/IP rate control. Ordinarily, TCP automatically informs the source to reduce or increase the size of the TCP window, depending on the amount of congestion; reducing it slows the traffic flow, while increasing it speeds things up. But with rate control, the intervals at which sources get this information are continually adjusted; instead of waiting for congestion to trigger the window-size request, the size is always under control and this regulates the rate at which traffic is sent. Essentially, rate control looks to smooth out spikes in transmission and avoid bottlenecks altogether. Products using TCP/IP rate control either define policies by establishing rate guarantees or prioritizing traffic. Since rate control can't shape anything other than TCP/IP, products that deploy it also deploy queuing.

The arguments in support of queuing are that rate control wastes bandwidth and time by sending acknowledgments back and forth, particularly when large numbers of flows and short-lived applications like HTTP sessions are involved. The proponents of rate control claim queuing induces packet loss and latency. The argument is that even though rate control may take a second or two to adjust flows, it provides end-to-end control when traffic patterns don't fit the way the queues have been set up.

Weighted round robin and round robin (WRR, RR) are scheduling mechanisms switches use to service the queues. For example, it might forward four packets from the high priority queue before checking whether there are any packets in the medium priority queue, forward two of those, and then move on to the low priority queue and forward one packet. In this example, the switch weights the priorities in a 4:2:1 ratio.

When traffic builds up in a queue, incoming packets can be lost because there's no room in the buffers. This in turn can lead to time-outs for multiple TCP senders, each of which then attempt to retransmit simultaneously, causing even more packets to pile up at the switch. It is this vicious circle that random early drop and weighted random early drop (RED, WRED) attempt to avoid by randomly discarding packets well before the buffers begin to tail-drop. It works at random to prevent any one TCP session from timing out altogether. Instead, only a few sessions slow down while the buffers empty out enough to avoid dropping more packets. WRED weighting is possible because different thresholds can be defined for various priorities.

Allot Communications, AVICI Systems, Cabletron, Cisco, Extreme Networks, and Ukiah Software use Weighted Fair Queuing (WFQ). Cabletron, IP Highway, Nortel Networks, and Xedia use Class-Based-Queuing (CBQ). Allot Communication also uses TCP rate control in addition to WFQ. AVICI Systems also uses WRR, random early drop RED, and TCP rate control in addition to WFQ. Cabletron also uses RED, WRED, and committed access rate (CAR) bandwidth shaping methods in addition to WFQ and CBQ. Cisco also uses FIFO, WRED, and committed access rate (CAR) to shape traffic. Extreme Networks also uses Round Robin (RR) and Weighted Round Robin (WRR) in addition to WFQ. Ukiah Software uses WRED, CAR and TCP rate control to shape traffic.

## 3.6 Policing and Feedback

In order to know whether QOS is effective or not, one needs to be able to see how the network is behaving before it’s applied, as well as how it responds afterwards. Most vendors support RMON (remote monitoring) using embedded agents on their switches or routers or use standalone probes. This allows network managers to monitor performance at Layer 2 by collecting statistics from different segments. Some also implement RMON2 on their hardware. This adds end-to-end monitoring and application layer reporting.

Alcatel, AVICI Systems, Cabletron, Xedia use Simple Network Management Protocol (SNMP) based interface to collect statistics such as packet counts, octet counts, and borrowing attempts that can be viewed via SNMP enabled trend analysis packages. These products collect and summarize the raw statistics into simple charts, graphs, and tables that reveal actual traffic patterns without creating data overload. Cabletron, Cisco, and Ukiah Software use RMON/RMON2 to measure bandwidth utilization for each layer-4 flow, allowing real time network baselining and fine-grained accounting.

## 3.7 Architecture

The architecture describes how a vendor implements QOS. The architecture - hardware, software, or a hybrid; indicates the type of performance that can be expected. Hardware-based products offer the best performance, since they rely on Application Specific Integrated Circuits (ASICs) to handle the bandwidth management tasks. Putting QOS functions into an ASIC chip speeds things up enormously. Hardware products on the other hand, are harder to upgrade because adding new features often means changing silicon.

IP Highway, Nortel Networks, and Ukiah take the software approach. While this generally makes for slower performance and longer setup time, upgrading is easier because the only thing it involves is a software download. Alcatel, Allot Communications, Cabletron, Cisco, Extreme Networks, and Xedia combine hardware with software. These hybrids perform

QOS functions using software embedded in dedicated hardware devices.

**Table 1: Comparison of Products**

| Company/Product | Server | Classification | Implement Policy-Where? | Propagation Mechanism | Queuing, Scheduling | Architecture |
|---|---|---|---|---|---|---|
| Alcatel (Switched Network Services) | Embedded | Source/Des IP, TCP/UDP Port, 802.1p, Time of Day. | End-to-end | LDAP |  | Hybrid |
| Allot Communications (AC Policy Manager and AC Enforcer) | Standalone | Source/Des IP, TCP/UDP Port, DNS, Time of Day, DS (TOS) | Edge Devices | COPS, LDAP | WFQ, TCP Rate Control | Hybrid |
| AVICI Systems (Terabit Switch Router) | Embedded | DS (TOS) | Edge Devices | SNMP, CLI | WFQ, WRR, RED, TCP rate Control | Hybrid |
| Cabletron (SmartSwitches and SmartSwitch Router) | Standalone | Source/Des IP, TCP/UDP Port, 802.1p, DS (TOS) | Edge Devices | SNMP, CLI | WFQ, CBQ, RED, CAR, WRED | Hybrid |
| Cisco (Internetworking Operating System) | Standalone | Source/Des IP, TCP/UDP Port, 802.1p, DS (TOS) | Edge Devices | COPS | WFQ, FIFO, WRED, CAR | Hybrid |
| Extreme Networks (Summit, BlackDiamond) | Embedded | Source/Des IP, TCP/UDP Port, 802.1p, DS (TOS) | End-to-end | COPS, CLI | WFQ, RR, WRR | Hybrid |
| IP Highway (ThruQos) | Standalone | Source/Des IP, TCP/UDP Port, Time of Day | End-to-end | COPS | CBQ | Software |
| Nortel (Preside, Optivity Policy Services) | Standalone | Source/Des IP, TCP/UDP Port, DNS, Time of Day, DS (TOS) | Edge Devices | COPS, CLI | CBQ | Software |
| Ukiah Software (NetRoad Active Policy System) | Standalone | 802.1p, DNS, Time of Day, DS (TOS) | End-to-end | SNMP, CLI | WFQ, CAR, WRED, TCP Rate Control | Software |
| Xedia (Access Point Router) | Embedded | Source/Des IP, TCP/UDP Port, DNS, DS (TOS) | Edge Devices | SNMP | CBQ | Hybrid |

Back to Table of Contents

# 4.0 SUMMARY

The fundamental idea of QOS is that traffic can be classified and provided differentiated service. Traffic classification makes two things possible. Packets can be prioritized according to the needs of specific applications and provided with the required QOS. QOS ensures that high priority, delay and jitter sensitive traffic gets the nod over lower priority traffic. The 802.1p standard gives switches the intelligence to recognize eight levels of priority. These bits operate at the MAC (media access control) OSI Layer 2. While 802.1p is effective for implementing QOS in the LAN, it typically doesn't go beyond the LAN/WAN boundary. For traffic that's headed out over the wide area, DiffServ is a promising technology. DiffServ makes use of the eight-bit TOS field of IPv4. The TOS field (also known as the DS byte) is set by DiffServ capable NICs, routers, and switches. Since DiffServ works at Layer 3, the tagging scheme and QOS parameters work across the WAN. DiffServ QOS specifications are recognizable by routers and switches or any device that reads the IP header and DS byte. The DS byte specifies the per-hop forwarding behavior (PHB) for that packet. A PHB might simply specify precedence, or it might include other performance characteristics.

When packets hit the WAN, they pass through a DiffServ boundary node, and then to DiffServ interior nodes. The boundary node also performs important conditioning functions (including class-based queuing) to keep PHBs "in profile". Conditioning functions also include metering (measuring each traffic stream's rate of flow); remarking the DiffServ byte to downgrade any traffic running in excess of the agreed maximum; shaping a traffic stream's packets to maintain conformity with the traffic profile; and policing/discarding packets within a traffic stream in order to enforce the correct traffic profile.

IEEE 802.1p provides a standard method for specifying delay or prioritization requirements over Ethernet and token ring LANs, while DiffServ defines ways of assigning specific service levels and priorities to IP traffic. Putting them together with ATM QOS mechanisms and the modified version of RSVP the IETF is now developing, could finally give network managers the ability to mark and maintain application traffic priorities and QOS from end to end, with their current infrastructure. Ten contemporary products in terms of how they provide quality of service were also compared.

Back to Table of Contents

# References

[Trillium98] Trillium, "Trillium IP Quality of Service White paper", April 1998, 27 pages.
An excellent review of QOS theory and protocols. URL: http://www.trillium.com/whats-new/wp_ipqos.html

[Saunders99] Stephen Saunders, David Newman and Erica Roberts, "The Policy Makers: A dozen vendors have gone public with their lofty ambitions for policy-based networks. But the quest for QOS is just beginning", Data Communications, May 1999, 20 pages.
A good review of QOS protocols and products. URL: http://www.data.com/issue/990507/policy.html

[Seaman99] Mick Seaman and Bob Klessig, "Going the Distance With QOS: Specify delivery requirements end to end? Using the right mix of protocols and initiatives makes it possible", Data Communications, February 1999, 8 pages.
A good review of QOS protocols. URL: http://www.data.com/issue/990207/distance.html

[Stephenson98] Ashley Stephenson, "DiffServ and MPLS: A Quality Choice", Data Communications, November 21, 1998, 12 pages.
A good review of how DiffServ and MPLS can provide QOS. URL: http://www.data.com/issue/981121/quality.html

[Allot Communication99a] Allot Communication, "Products - Policy Based Networking Solutions", September 1999, 2 pages. An overview of Allot Communication's QOS. URL: http://www.allot.com/products/

[Allot Communications99b] Allot Communications, "AC Policy Manager: Directory Enabled, Policy Management System",
September 1999, 4 pages. A data sheet on AC Policy Manager.

[Allot Communications99c] Allot Communications, "AC Enforcer - Bandwidth Management Solution", September 1999, 7 pages. A data sheet on AC Enforcer. URL: http://www.allot.com/products/ACfamily_DS.htm

[Xedia99] Xedia, "Access Point QOS Routers", February 1999, 4 pages.
Data sheet on Access Point QOS Routers.

[Stephenson99] Ashley Stephenson, "QoS: The IP Solution", September 1999, 7 pages.
A good review of the standards for IP QoS. URL: http://www.xedia.com/products/ipsolution.html

[Extreme Networks98] Extreme Networks, "Extreme Ware Enterprise Manager Data Sheet", 7 pages.
A data sheet describing Enterprise manager.
URL: http://www.extremenetworks.com/extreme/products/datasheets/entmngr.asp

[Extreme Networks99] Extreme Networks, "Protecting the Delivery of Mission-Critical Application - The Enterprise Policy System", 9 pages.
A data sheet on Enterprise Policy System from URL: http://www.extremenetworks.com/extreme/products/eem_eps.htm

[AVICI Systems98] AVICI Systems, "Delivering Internet Quality of Service", 1998. 9 pages.
A data sheet on AVICI System's Terabit Switch router from http://www.avici.com/products/index.html

[Cabletron99] Cabletron, "SmartSwitch Routers", September 1999, 15 pages.
A data sheet on Cabletron's SmartSwitch routers from http://www.cabletron.com/smartswitch-router/

[IP Highway98] IP Highway, "The Road to ThruQoS", 8 pages.
A white paper on IP Highway's ThruQoS.

[IP Highway99] IP Highway, "Product Evolution".
The first page with links to IP Highway's complete QOS Suite from http://www.iphighway.com/products.html

[Alcatel99] Alcatel, "Alcatel Internetworking Product Features - Quality of Service",1999, 4 pages.
"Alcatel Internetworking Product Features - Policy Based Management", 5 pages.
Descriptions on Alcatels QOS products.
http://rogets.ind.alcatel.com/enterprise/products/features/sns/sns06.html.
http://rogets.ind.alcatel.com/enterprise/products/features/sns/sns07.html.

[Alcatel99] Alcatel, "Switched Network Services", 25 pages. White paper on Alcatel's QOS offerings.

[Cisco99] Cisco, "Quality of Service Overview", June 1999, 13 pages.
An overview of Cisco's QOS offering.

[Nortel Networks99a] Nortel Networks, "Preside Quality of Service Position Paper", September 1999, 11 pages.

[Nortel Networks98] Nortel Networks, "IP QoS - A Bold New Network", September 1998, 24 pages.
An excellent white paper on QOS from Nortel Networks.

[Nortel Networks99b] Nortel Networks, "Network Management Products - Optivity Policy Services 1.0 for the Enterprise", May 1999, 9 pages.
Overview, features and Benefits of Optivity Policy Services from Nortel Networks.
URL: http://www.nortelnetworks.com/products/02/datasheets/3511.html

[Ukiah Software99] Ukiah Software, "Ukiah Software's NetRoad Active Policy System", February 1999, 5 pages.
An overview of Ukiah's QOS software.

Back to Table of Contents

# List Of Acronyms

QOS: Quality of Service
IntServ: Integrated Services
DiffServ: Differentiated Services
MPLS: Multi-Protocol Label Switching
VC: Virtual Circuit
VPI: Virtual path Identifier
VCI: Virtual Channel Identifier
ABR: Available bit rate

UBR: Unspecified bit rate
CBR: Constant bit rate
VBR: Variable bit rate
rt-VBR: real-time VBR
nrt-VBR: non real-time VBR
GFR: Guaranteed Frame rate
MCR: Minimum Cell rate
RSVP: Resource Reservation Protocol
TOS: Type of Service
CBQ: Class Based Queuing
RED: Random early Detection
WRED: Weighted Random Early Detection
LDP: Label Distribution Protocol
LSP: Label Switched Path
LSR: Label Switched Router
LDAP: Lightweight Directory Access Protocol
SNMP: Simple Network Management Protocol
COPS: Common Open Policy Service
WFQ: Weighted Fair Queuing
RR: Round Robin
WRR: Weighted Round Robin
CAR: Committed Access Rate
RMON: Remote Monitoring
PHB: Per Hop Behavior

Back to Table of Contents

---

*Last Modified: November 22, 1999.*

Note: This paper is available on-line at
http://www.cis.ohio-state.edu/~jain/cis788-99/