

Introduction to **Quality of Service (QoS)**

Introduction

Quality of Service (QoS) is a broad term used to describe the overall experience a user or application will receive over a network. QoS involves a broad range of technologies, architecture, and protocols. Network operators achieve end-to-end QoS by ensuring that network elements apply consistent treatment to traffic flows as they traverse the network.

The goal of this paper is to put all of this into perspective and provide a holistic view of the need for and the value of creating a QoS-enabled network. This paper will describe the relationship between the numerous technologies and acronyms used to describe QoS.

This paper assumes that the reader is familiar with different networking technologies and architectures. This paper will provide answers to the following:

- Why is a QoS-enabled network important?
- What considerations should be made when deploying a QoS-enabled network?
- Where should QoS be applied in the network?

Why is QoS important?

Today, network traffic is highly diverse and each traffic type has unique requirements in terms of bandwidth, delay, loss, and availability. With the explosive growth of the Internet, most network traffic today is IP-based. Having a single end-to-end transport protocol is beneficial because networking equipment becomes less complex to maintain, resulting in lower operational costs. This benefit, however, is countered by the fact that IP is a connectionless protocol, i.e., IP packets do not take a specific path as they traverse the network. This results in unpredictable QoS in a best-effort network.

The IP protocol was originally designed to reliably get a packet to its destination with less consideration to the amount of time it takes to get there. IP networks must now support many different types of applications. Many of these applications require low latency. Otherwise, the end-user quality may be significantly affected or in some cases, the application simply does not function at all.

Consider a voice application. Voice applications originated on public telephone networks using TDM (Time Division Multiplexing) technology which has a very deterministic behavior. On TDM networks, the voice traffic experienced a low and fixed amount of delay with essentially no loss. Voice applications require this type of behavior to function properly. Voice applications also require this same level of “TDM voice” quality to meet user expectations.

Take this “TDM voice” application and now transport it over a best-effort IP network. The best-effort IP network introduces a variable and unpredictable

amount of delay to the voice packets and also drops voice packets when the network is congested. As you can see, the best-effort IP network does not provide the behavior that the voice application requires. QoS techniques can be applied to the best-effort IP network to make it capable of supporting VoIP with acceptable, consistent, and predictable voice quality.

QoS and network convergence

Since the early 1990s, there has been a movement towards network convergence, i.e., transport all services over the same network infrastructure. Traditionally, there were separate, dedicated networks for different types of applications. However, many of these networks are being consolidated to reduce operational costs or improve profit margins.

Not too long ago, an enterprise may have had a private TDM-based voice network, an IP network to the Internet, an ISDN video conferencing network, an SNA network, and a multi-protocol (IPX, AppleTalk, etc.) LAN. Similarly, a service provider may have had a TDM-based voice network, an ATM or SONET backbone network, and a Frame Relay or ISDN access network.

Today, all data networks are converging on IP transport because the applications have migrated towards being IP-based. The TDM-based voice networks have also begun moving towards IP. Video conferencing is also moving towards IP, albeit at a lesser pace. When the different applications had dedicated networks, QoS technologies played a smaller role because the traffic was similar in behavior and the dedicated networks were fine-tuned to meet the required behavior of the particular application.

The converged network mixes different types of traffic, each with very different requirements. These different traffic types often react unfavorably together. For example, a voice application expects to experience essentially no packet loss and a minimal but fixed amount of packet delay. The voice application operates in a steady-state fashion with voice channels (or packets) being transmitted at fixed time intervals. The voice application receives this performance level when it operates over a TDM network. Now take the voice application and run it over a best-effort IP network as VoIP (Voice over IP). The best-effort IP network has varying amounts of packet loss and potentially large amounts of variable delay (typically caused by network congestion points). The best-effort IP network provides almost exactly the opposite performance required by the voice application. Therefore, QoS technologies play a crucial role to ensure that diverse applications can be properly supported in a multiservice IP network.

“QoS technologies play a crucial role... in a multiservice IP network.”

QoS versus bandwidth

Some believe that QoS is not needed and that increasing bandwidth will suffice and provide good QoS for all applications. They argue that implementing QoS is complicated and adding bandwidth is simple. While there is some truth to these statements, one first must look closer at the QoS problems to be solved and whether adding bandwidth will solve them.

If all network connections had infinite bandwidth such that networks never became congested, then one would not need to apply QoS technologies. Indeed, there are parts of some carrier networks that have tremendous amounts of bandwidth and have been carefully engineered to minimize congestion. However, high-bandwidth connections are not available throughout the network, from the traffic's source to the traffic's destination. This is especially true for access networks where the most commonly available bandwidth is typically only hundreds of kbps. Furthermore, bandwidth discontinuities in the network are potential congestion points resulting in variable and unpredictable QoS that a user or application experiences.

Carriers have been aggressively adding bandwidth to their IP networks to meet the exploding demand of the Internet. Some carriers can offer low latency connections across their metropolitan area networks (MANs) or cross-country or continental long-haul networks. Traffic must be treated consistently to achieve a subscribed QoS level. If the access network aggregation point is a congestion point in the network, the resulting end-to-end QoS can be poor even though the long-haul network may offer excellent QoS performance.

Bandwidth owner versus renter

If the network operator owns the network connection's copper, fiber, or frequencies (wireless), adding bandwidth to provide QoS may be an attractive choice in lieu of implementing more complicated QoS traffic management mechanisms. Some interconnect technologies, such as DWDM (Dense Wavelength Division Multiplexing) over fiber optic connections allow bandwidth to be added simply and cost-effectively over the existing cable infrastructure. Other

interconnect technologies, such as fixed or mobile wireless, are much more constrained due to frequency spectrum limitations regulated by government agencies.

Those who own their network connections have more choices than those who lease their bandwidth.

For example, a network operator has dark fiber in their backbone network and DWDM interfaces on their backbone switches. The operator has determined that sufficient bandwidth is no longer available to provide the network users with the required performance objectives. In this scenario, it is simpler and more cost-effective to add additional wavelengths to provide QoS by increasing bandwidth than add complicated QoS traffic management mechanisms.

Those who lease their network connections are more constrained with their bandwidth choices.

For example, a network operator provides best-effort services and rents (leases) bandwidth from another service provider. This network operator wants to offer premium data services with performance guarantees. Through the application of QoS mechanisms, this network operator can offer service differentiation between the premium data service and best-effort subscribers. For a lightly to moderately loaded network, this may be accomplished without increasing bandwidth. Hence, for the same fixed recurring cost of this leased bandwidth, the network operator can offer additional, higher-priced services and increase profit margins.

QoS performance dimensions

A number of QoS parameters can be measured and monitored to determine whether a service level offered or received is being achieved. These parameters consist of the following:

- Network availability
- Bandwidth
- Delay
- Jitter
- Loss

There are also QoS performance-affecting parameters that cannot be measured but provide the traffic management mechanisms for the network routers and switches. These consist of:

- Emission priority
- Discard priority

Each of these QoS parameters affects the application's performance or end-user's experience.

Network availability

Network availability can have a significant effect on QoS. Simply put, if the network is unavailable, even during brief periods of time, the user or application may achieve unpredictable or undesirable performance (QoS).

Network availability is the summation of the availability of many items that are used to create a network. These include networking device redundancy, e.g., redundant interfaces, processor cards or power supplies in routers and switches, resilient networking protocols, multiple physical connections, e.g., fiber or copper, backup power sources, etc. Network operators can increase their network's availability by implementing varying degrees of each of these items.

The greatest challenge for network operators today is to provide highly available IP networks.

Bandwidth

Bandwidth is probably the second most significant parameter that affects QoS. Bandwidth allocation can be subdivided into two types:

- Available bandwidth
- Guaranteed bandwidth

Available bandwidth

Many network operators oversubscribe the bandwidth on their network to maximize the return on investment of their network infrastructure or leased bandwidth. Oversubscribing bandwidth means the bandwidth a user subscribed to is not always available to them. This allows all users to compete for available bandwidth. They get more or less bandwidth depending upon the amount of traffic from other users on the network at any given time.

Available bandwidth is a technique commonly used over consumer ADSL networks, e.g., a customer signs up for a 384-kbps service that provides no QoS (bandwidth) guarantee in the SLA. The SLA points out that the 384-kbps is "typical" but does not make any guarantees. Under lightly loaded conditions, the user may achieve 384-kbps but upon network loading, this bandwidth will not be achieved consistently. This is most noticeable during certain times of the day when more users access the network.

Guaranteed bandwidth

Network operators offer a service that provides a guaranteed minimum bandwidth and burst bandwidth in the SLA. Because the bandwidth is guaranteed, the service is priced higher than the

available bandwidth service. The network operator must ensure that those who subscribe to this guaranteed bandwidth service get preferential treatment (QoS bandwidth guarantee) over the available bandwidth subscribers.

In some cases, the network operator separates the subscribers by different physical or logical networks, e.g., VLANs, Virtual Circuits, etc. In some cases, the guaranteed bandwidth service traffic may share the same network infrastructure with the available bandwidth service traffic. This is often the case at locations where network connections are expensive or the bandwidth is leased from another service provider. When subscribers share the same network infrastructure, the network operator must prioritize the guaranteed bandwidth subscribers' traffic over the available bandwidth subscribers' traffic so that in times of network congestion, the guaranteed bandwidth subscribers' SLAs are met.

Burst bandwidth can be specified in terms of amount and duration of excess bandwidth (burst) above the guaranteed minimum. QoS mechanisms may be activated to discard traffic that is consistently above the guaranteed minimum bandwidth that the subscriber agreed to in the SLA.

Delay

Network delay is the transit time an application experiences from the ingress point to the egress point of the network. Delay can cause significant QoS issues with applications such as voice and video, and applications such as SNA and fax transmission that simply time-out and fail under excessive delay conditions. Some applications can compensate for small amounts of delay but once a certain amount is exceeded, the QoS becomes compromised.

For example, some networking equipment can "spoof" an SNA session on a host by providing local acknowledgments when the network delay would cause the SNA session to time-out. Similarly, VoIP gateways and phones provide some local buffering to compensate for network delay.

Finally, delay can be both fixed and variable. Examples of fixed delay are:

- Application-based delay, e.g., voice codec processing time and IP packet creation time by the TCP/IP software stack
- Data transmission (queuing delay) over the physical network media at each network hop
- Propagation delay across the network based on transmission distance

Examples of variable delays are:

- Ingress queuing delay for traffic entering a network node
- Contention with other traffic at each network node
- Egress queuing delay for traffic exiting a network node

"Some applications can compensate for small amounts of delay but once a certain amount is exceeded, the QoS becomes compromised."

Jitter

Jitter is the measure of delay variation between consecutive packets for a given traffic flow. Jitter has a pronounced effect on real-time, delay-sensitive applications such as voice and video. These real-time applications expect to receive packets at a fairly constant rate with fixed delay between consecutive packets. As the arrival rate varies, the jitter impacts the

application's performance. A minimal amount of jitter may be acceptable but as jitter increases, the application may become unusable.

Some applications, such as voice gateways and IP phones, can compensate for small amounts of jitter. Since a voice application requires the audio to play out at a constant rate, if the next packet does not arrive within the playback time, the application will replay the previous voice packet until the next voice packet arrives. However, if the next packet is delayed too long, it is simply discarded when it arrives, resulting in a small amount of distorted audio.

All networks introduce some jitter because of variability in delay introduced by each network node as packets are queued. However, as long as the jitter is bounded, QoS can be maintained.

“...as long as the jitter is bounded, QoS can be maintained.”

Loss

Loss can occur due to errors introduced by the physical transmission medium. For example, most landline connections have very low loss as measured in the Bit Error Rate (BER). However, wireless connections such as satellite, mobile, or fixed wireless networks have a high BER that varies due to environment or geographical conditions such as fog, rain, RF interference, cell handoff during roaming, and physical obstacles such as trees, buildings, and mountains. Wireless technologies often transmit redundant information since packets will inherently get dropped some of the time due to the nature of the transmission medium.

Loss can also occur when congested network nodes drop packets. Some networking protocols such as TCP (Transmission Control Protocol) offer packet loss protection by retransmitting packets that may have been dropped or corrupted by the network. When a network becomes increasingly congested, more packets are dropped and hence more TCP retransmissions. If congestion continues, the network performance will significantly decrease because much of the bandwidth is being used to retransmit dropped packets. TCP will eventually reduce its transmission window size, resulting in smaller and smaller packets being transmitted. This eventually will reduce congestion, resulting in fewer packets being dropped.

Because congestion has a direct impact on packet loss, congestion avoidance mechanisms are often deployed. One such mechanism is called Random Early Discard (RED). RED algorithms randomly and intentionally drop packets once the traffic reaches one or more configured thresholds. RED takes advantage of the TCP protocol's window size throttling feature and provides more efficient congestion management for TCP-based flows. Note that RED only provides effective congestion control for applications or protocols with “TCP-like” throttling mechanisms.

Emission priorities

Emission priorities determine the order in which traffic is forwarded as it exits a network node. Traffic with a higher emission priority is forwarded ahead of traffic with a lower emission priority. Emission priorities also determine the amount of latency introduced to the traffic by the network node's queuing mechanism.

For example, delay-tolerant applications such as e-mail would be configured to have a lower emission priority than delay-sensitive real-time applications such as voice or video. These delay-tolerant applications may be buffered while the delay-sensitive applications are being transmitted.

In its simplest of forms, emission priorities use a simple transmit priority scheme whereby higher emission priority traffic is always forwarded ahead of lower emission priority traffic. This is typically accomplished using strict priority scheduling (queuing). The downside of this approach is that low emission priority queues may never get serviced (starved) if there is always higher emission priority traffic with no bandwidth rate limiting.

A more elaborate scheme provides a weighted scheduling approach to the transmission of traffic to improve fairness, i.e., the lower emission priority traffic doesn't always have to wait until the higher emission priority traffic is transmitted. Finally, some emission priority schemes provide a mixture of both priority and weighted schedulers.

Discard priorities

Discard priorities are used to determine the order in which traffic gets discarded. The traffic may get dropped due to network node congestion or when the traffic is out-of-profile, i.e., the traffic exceeds its prescribed amount of bandwidth for some period of time.

Under congestion, traffic with a higher discard priority gets dropped before traffic with a lower discard priority. Traffic with similar QoS performance requirements can be subdivided using discard priorities. This allows the traffic to receive the same performance when the network node is not congested.

However, when the network node is congested, the discard priority is used to drop the “more eligible” traffic first.

Discard priorities also allow traffic with the same emission priority to be discarded when the traffic is out-of-profile. Without discard priorities, traffic would need to be separated into different queues in a network node to provide service differentiation. This can be expensive since only a limited number of hardware queues (typically eight or less) are available on networking devices. Some devices may have software-based queues but as these are increasingly used, network node performance is typically reduced.

With discard priorities, traffic can be placed in the same queue but in effect, the queue is subdivided into virtual queues, each with a different discard priority. For example, if a product supports three discard priorities, then one hardware queue in effect provides three QoS levels.

Application requirements

Table 1 illustrates the QoS performance dimensions required by some common applications. As you can see from this table, applications can have very different QoS requirements. As these are mixed over a common IP transport network, without applying QoS technologies, the traffic will experience unpredictable behavior.

“...without applying QoS technologies, the traffic will experience unpredictable behavior.”

Table 1: Application performance dimensions

Performance dimensions				
Application	Bandwidth	Sensitivity to:		
		Delay	Jitter	Loss
VoIP	Low	High	High	Med
Video Conferencing	High	High	High	Med
Streaming Video on Demand	High	Med	Med	Med
Streaming Audio	Low	Med	Med	Med
Client/Server Transactions	Med	Med	Low	High
E-mail	Low	Low	Low	High
File transfer	Med	Low	Low	High

Table 2: Application traffic categories

Traffic category	Example application
Network control	Critical alarms, routing, billing, Critical OAM
Interactive	VoIP, interactive gaming, video conferencing
Responsive	Streaming audio/video, client/server transactions
Timely	E-mail, non-critical OAM

Categorizing applications

Networked applications can be categorized based on end-user expectations or application requirements. Some applications are between people while other applications are between a person and a networked device’s application, e.g., a PC and a Web server. Finally, some applications are between networking devices, e.g., router-to-router.

Table 2 categorizes applications into four different traffic categories—namely Interactive, Responsive, Timely, and Network Control. The table also includes example applications that fall into the different categories.

Interactive applications

Some applications are “interactive” whereby two or more people actively participate. The participants expect the networked application to respond in “real-time”. In this context, “real-time” means that there is minimal delay (latency) and delay variation (jitter) between the sender and receiver. Some interactive applications, such as a telephone call, have operated in real-time over the telephone companies’ circuit switched networks for over 100 years. The QoS expectations for voice applications have been set and therefore must also be achieved for packetized voice such as VoIP.

Other interactive applications include video conferencing and interactive gaming. Since the interactive applications operate in real-time, packet loss must be minimized. Imagine if you were speaking over the telephone and lost parts of a word every so often during the conversation. This QoS level would be unsatisfactory.

Interactive applications typically are UDP-based (Universal Datagram Protocol) and hence cannot retransmit lost or dropped packets as with TCP-based (Transport Control Protocol) applications. However, packet retransmission would not be beneficial because interactive applications are time-based. For example, if a voice packet was lost, it doesn't make sense for the sender to retransmit it because the conversation has already progressed and the lost packet might be from part of the conversation that has already passed in time.

Interactive applications expect the network QoS to provide packets with the lowest possible delay, jitter, and loss.

Responsive applications

Some applications are between a person and a networked device's application. End users require these applications to be "responsive" so a request sent to the networking device requires a relatively quick response back to the sender. These applications are sometimes referred to as being "near real-time". These applications require relatively low packet delay, jitter, and loss. However, QoS requirements for responsive applications are not as stringent as real-time, interactive application requirements. This category includes streaming media and client/server Web-based applications.

Streaming media applications include Internet radio (talk radio and music) and audio/video broadcasts (news, training, education, and motion pictures). Streaming applications require the network to be responsive when they are initiated so the user doesn't wait too long before the media begins playing. These applications also require the network to be responsive for certain types of signaling. For example, with movies on demand, when one changes channels or "forwards", "rewinds", or "pauses" the media, one expects the application to react similarly to the response time of their VCR controls.

Client/server Web-based applications typically involve the user selecting a hyperlink to jump to a new page or submit information (place an order, submit a request, etc.). These applications also require the network to be responsive such that once the hyperlink is selected, a response—e.g., a new page begins loading—occurs typically within one to two seconds. With broadband Internet connections, this is often achieved over a best-effort network, albeit inconsistently. These types of applications may include a financial transaction, e.g., place credit card order and quickly provide feedback to the user indicating the transaction has completed. Otherwise, the user may be unsure the transaction completed and attempt to initiate a duplicate order (unknowingly). Alternatively, the user may assume that the order was placed correctly but it may not have. In either case, the user will be dissatisfied with the network or application's performance.

Responsive applications can use either UDP or TCP-based transport. Streaming media applications typically use UDP (but can also use TCP). Web-based applications are based on the HyperText Transport Protocol (HTTP) and always use TCP. For Web-based applications,

packet loss is managed by TCP which retransmits lost packets. Retransmission of lost streaming media packets is typically handled by application-level protocols as long as the media is sufficiently buffered. If not, then the lost packets are discarded, resulting in some distortions in the media (audio or video).

Timely applications

Some applications between a person and networked device's application do not require "near real-time" performance but do require the information to be delivered in a timely manner. Such examples include store and forward e-mail applications and file transfers. The relative importance of these applications is based on their business priorities.

These applications require that the packets arrive with a bounded amount of delay. For example, if an e-mail takes a few minutes to arrive at its destination, this is acceptable. However, in a business environment, if an e-mail took 10 minutes to arrive at its destination, this is often unacceptable. The same bounded delay applies to file transfers. Once a file transfer is initiated, delay and jitter are inconsequential because file transfers often take minutes to complete. Note that timely applications use TCP-based transport and therefore, packet loss is managed by TCP which retransmits any lost packets resulting in no packet loss.

In summary, timely applications expect the network QoS to provide packets with a bounded amount of delay. Jitter has a negligible effect on these types of applications. Loss is reduced to zero due to TCP's loss recovery mechanisms.

Network control applications

Some applications are used to control the operation and administration of the network. Such applications include network routing protocols, billing applications, and

“Timely applications expect the network QoS to provide packets with a bounded amount of delay.”

QoS monitoring and measuring for SLAs. These applications can be subdivided into those required for critical and standard network operating conditions. To create high-availability networks, network control applications require “priority” over end-user applications because if the network is not operating properly, end-user application performance will suffer.

In summary, network control applications require a relatively low amount of delay. Jitter has a negligible effect on these types of applications. Loss needs to be minimized because some of the applications are not transported via TCP and hence do not have lost packet recovery mechanisms.

“Network control applications require a relatively low amount of delay. Loss needs to be minimized.”

QoS technologies— a layered approach

Building a QoS-enabled network requires a number of different QoS technologies. *Figure 1* provides a quasi-OSI model of the different QoS technology layers. Some QoS technologies used to measure and monitor QoS services span all layers. The intent here is not to delve into each of the technologies. The intent is to provide an understanding of where each of these QoS technologies can be used and the benefit they provide.

Figure 1: QoS technologies

QoS measurement and monitoring	IP QoS	IP Differentiated Services (DiffServ)
	Network-signaled QoS	ATM PNNI, MPLS RSVP-TE, or MPLS CR-LDP
	Traffic-engineered paths	ATM PVCs, MPLS Label Switched Paths (LSPs)
	Link layer QoS	Ethernet 802.1Q VLANs, 802.1p, ATM, MPLS, PPP, UMTS, DOCSIS, Frame Relay
	Physical layer QoS	Wavelengths, Virtual Circuits (VCs), ports, frequencies

“Building a QoS-enabled network requires a number of different QoS technologies.”

Physical layer QoS

These technologies allow for the separation of traffic. The separation and prioritization may take the form of wavelengths (lambdas), Virtual Circuits (VCs), ports on a device, or frequencies over the air. This is the simplest form of QoS whereby different levels of QoS are provided through traffic separation at the physical layer. For example, the blue wavelength may provide a priority service and the red wavelength may provide a best-effort service. In some cases, this type of QoS can be inexpensive, e.g., adding additional wavelengths over an existing fiber optic cable. However, this approach can also be very expensive if the resources are leased or limited, e.g., frequency spectrum.

Link layer QoS

Each link layer has a different type of QoS technology that can be applied. The most common link layers are Ethernet, ATM, PPP, MPLS, DOCSIS (HFC cable), Frame Relay, and Mobile wireless technologies (only UMTS/3GPP will be discussed here). The following sections will provide a brief overview of the different link layer QoS technologies. For additional details, references are provided at the end of this document.

Ethernet

Ethernet (IEEE 802.3) provides two different QoS mechanisms. One mechanism is via 802.1p which provides eight classes of service. The other mechanism is via VLANs whereby traffic can be separated, isolated, and prioritized by the VLAN ID. VLANs allow for the logical grouping of users or devices with similar QoS or security requirements. Although VLANs are Layer 2-based, users belonging to the same VLAN do not need to be physically connected to the same Ethernet subnet. VLANs most commonly allow for traffic separation and prioritization based on the particular Ethernet switch port to which a user is connected (called port-based VLANs). VLANs can also be created based on Ethernet MAC addresses, protocol types, or other user-defined information for the Ethernet switches to classify.

ATM

The ATM Forum created ATM service categories, each with different QoS traffic management parameters and performance levels. The most widely available ATM service categories are CBR (Constant Bit Rate), rt-VBR (real-time Variable Bit Rate), nrt-VBR (non real-time VBR), and UBR (Unspecified Bit Rate). In general, CBR is used for circuit emulation services (including circuit-based voice or video transport), rt-VBR is used for real-time packet-based voice

or video services, nrt-VBR is used for priority data services, and UBR is used for best-effort data services.

Other ATM service categories defined that are less widely available are ABR (Available Bit Rate) and GFR (Guaranteed Frame Rate), both of which are enhancements over the UBR service and provide additional service guarantees that are not provided by UBR.

ATM also provides a number of traffic management parameters for each of the ATM service categories such as Peak Cell Rate (PCR), Sustained Cell Rate (SCR), and Cell Loss Priority (CLP). These parameters define the traffic's performance level in the particular ATM service category.

PPP

When using PPP over a low-bandwidth connection, the IP packets are typically fragmented to reduce queuing delay. Each packet fragment can be assigned a PPP Class Number to use for service differentiation. When the packets arrive at a PPP session termination point for reassembly, the Class Numbers are used to determine the service class to which the particular packet fragments belong.

Two PPP multi-class formats are supported. The short sequence number format provides four classes of service and the long sequence number format provides 16 classes of service.

MPLS

MPLS provides for two different forms of QoS determined by the EXP (Experimental) bits in the MPLS shim header. When using E-LSPs (EXP-inferred Label Switched Paths), the EXP bits provide 8 service classes that support both emission and discard priorities and Differentiated Services (DiffServ) traffic class behaviors. When using L-LSPs (Label-inferred LSPs), the EXP bits provide eight discard priorities.

MPLS also supports a number of traffic management parameters to define the behavior the traffic will receive as it traverses a particular LSP.

DOCSIS

The DOCSIS protocol is used over HFC cable networks and has been proposed for use over fixed wireless networks. The DOCSIS protocol supports QoS via traffic separation through the use of Service IDs (SIDs). There are four SIDs. They are rtPS (real-time Polling Service), nrtPS (non real-time Polling Service), UGS (Unsolicited Grant Service), UGS-AD (Unsolicited Grant with Activity Detection), and BE (Best Effort Services). rtPS is used for bursty real-time services such as video and VoIP with silence suppression. nrtPS is used for bursty, non real-time flows such as Web browsing. UGS is used for periodic real-time flows such as VoIP. UGS-AD is used for intermittent yet periodic real-time flows such as VoIP with Voice Activity Detection (VAD).

Frame Relay

Frame Relay has a QoS mechanism called Discard Eligibility (DE) that can be set by the networking device for traffic that can be discarded under network congestion. Frame Relay also provides a Committed Information Rate (CIR) and Excess Information Rate (EIR) that define the minimum and burst bandwidth guarantees, respectively.

UMTS Wireless Networks

The 3rd-Generation Partnership Project (3GPP) has defined wireless networks that are the first to offer QoS through service differentiation. UMTS (Universal Mobile Telecommunications System) standards have defined four wireless traffic classes used for traffic separation. The four classes are Conversational, Streaming, Interactive, and Background.

UMTS networks provide traffic management parameters such as Allocation/Retention priorities used to determine the forwarding behavior of the traffic.

Traffic-engineered paths

As traffic traverses a network, it can often take different paths depending upon the networking technology. For example, routed IP networks are connectionless, i.e., a packet can take different paths. Because there is not a dedicated path for the traffic to traverse, QoS can become less predictable under network congestion. Because network operators want to offer service level guarantees, network paths can be engineered to provide guaranteed QoS performance. Traffic that is within the service level criteria can be steered to these traffic engineered paths and obtain a predictable QoS level.

Network-signaled QoS

MPLS and ATM use signaling protocols to request a desired QoS level from other network nodes prior to connection establishment (known as connection admission control or CAC). ATM uses a protocol called PNNI (Private Network-Network Interface) to accomplish this. MPLS uses a protocol called LDP (Label Distribution Protocol) to set up the Label-Switched Paths (LSPs). To signal QoS for these traffic-engineered paths, MPLS uses RSVP-TE (RSVP for Traffic Engineered paths) or CR-LDP (Constraint-based Routed LDP). RSVP-TE provides extensions to and uses the basic RSVP protocol to set up stateful connections between MPLS Label Edge Routers (LER) and Label Switch Routers (LSR). CR-LDP uses the TCP protocol to set up connections between the LERs and LSRs. Both protocols achieve similar goals through different approaches.

IP QoS via DiffServ

DiffServ (Differentiated Services) defines a number of service classes and QoS mechanisms that are applied to packets in those service classes (called Per-Hop Behaviors or PHBs). The DiffServ Code Point (DSCP) is located in the IP packet header and is used to determine the PHB. Note that the DSCP occupies the TOS (Type of Service) field yet is not compatible with it. (DiffServ has replaced the TOS field definition.) The standard PHBs each have a unique DSCP associated with them. The DSCP is used to determine the respective DiffServ behavior the packet is to receive. Different types of applications have different traffic characteristics and require different types of QoS behaviors to be applied to them. Note that custom PHBs can also be created using a unique DSCP to identify them.

Expedited Forwarding DiffServ class

The Expedited Forwarding (EF) DiffServ behavior provides a low-latency, high-priority service that is ideally suited for VoIP. The EF behavior is implemented with a high (or highest) emission priority and lowest discard priority. Basically, in order to achieve the required behavior, each network node must ensure that the EF traffic has the lowest possible delay, jitter, and loss since the service is attempting to emulate a leased line over an IP network.

Assured Forwarding DiffServ class

The Assured Forwarding (AF) DiffServ behavior consists of four different service classes, each with three different discard priority levels.

Table 3: Assured Forwarding DiffServ classes

Discard priority	AF class			
	Class 1	Class 2	Class 3	Class 4
Low	AF11	AF21	AF31	AF41
Medium	AF12	AF22	AF32	AF42
High	AF13	AF23	AF33	AF43

Each AF class also has three different discard priorities (drop precedence levels), resulting in twelve different DSCP values. Routers use these drop precedence values to determine the discard priority of packets under network congestion. *Table 3* lists the different AF classes by discard priorities.

Class Selector DiffServ class

The Class Selector (CS) DiffServ behavior may be represented by eight priority classes and uses the same bit positions as the IP Precedence field in the TOS definition. In this usage, the CS7 DSCP has the highest emission (forwarding) priority and the CS0 DSCP has the lowest forwarding priority. CS0 is equivalent to a “best effort” service. Note that the CS behavior does not support discard priorities.

Class Selector DSCPs may also be used to indicate other standardized DiffServ behaviors. For example, a router may be configured such that packets marked with the CS5 DSCP receive Expedited Forwarding (EF) behavior. In this example, packets marked with either CS5 and EF will receive Expedited Forwarding behavior by the router. This is useful where different applications require the same DiffServ behavior but their bandwidth is managed separately using the different DSCP value.

Default DiffServ class

The Default (DE) DiffServ behavior is used to transport best-effort traffic. Any traffic that is not classified into another standard or custom DiffServ PHB must be transported using the DE PHB. The DE DSCP value is 0 and typically has the highest discard priority and lowest emission priority.

QoS measurement and monitoring

Service providers need to ensure that their networks are properly engineered to support new services such as IP telephony, video on demand (movies), interactive learning, and video conferencing services. These services cannot be offered without certain QoS-related network capabilities.

In order to offer such services to subscribers, the network QoS must be measured and monitored to ensure that the service is being adequately supported. Furthermore, since the QoS performance may be specified in a Service Level Agreement (SLA), the service provider needs to ensure that the network is providing the performance as specified. The SLA may consist of parameters such as maximum packet loss (over some time interval) and maximum packet delay.

The SLA specifies the terms and conditions of the service being offered. Once a service provider can accurately measure the network capabilities and provide a guaranteed performance level, he can confidently offer a billable service to his subscribers.

Subscribers also want to monitor network performance to ensure that they receive the services to which they subscribe. Therefore, the service provider should also provide the subscriber with tools to monitor network QoS statistics that are relevant for the particular SLA.

Making QoS simple

After reading through this document, you will quickly conclude that QoS can be quite a complicated topic and there are many different technologies, standards, and network architectures to consider. Furthermore, there is no single QoS technology or standard that can be used across your network. How can QoS be simplified?

In 1999, Nortel Networks started simplifying QoS by creating standardized, default QoS configurations and behaviors for its products in the form of end-to-end network service classes. These are called Nortel Networks Service Classes (NNSCs) and the NNSCs are a superset of the QoS Classes defined in the ITU-T Y.1541 standard. The NNSCs have been defined based upon the most common types of applications as illustrated in *Table 4*. The NNSCs provide default settings and behaviors for the different QoS technologies.

The NNSCs have been designed to provide the appropriate behavior required for different types of applications. A service provider or enterprise network manager determines the services to be offered and the application to be supported. Based on this information, the network elements are configured to place the traffic into the NNSC that provides the closest performance behavior required by the application or service offering. Services can now be quickly added using the NNSCs' default QoS behaviors while not having to deal with the underlying QoS technologies required to create the service behaviors.

NNSCs are being built into Nortel Networks products but can also be created in other vendors' products through QoS policy management systems. The default behaviors would be created in the policy management systems and the device-specific configuration rules (commands) would then be transmitted to the device.

Table 4. Nortel Networks Service Classes (NNSC)

Traffic category	Example application	Nortel Networks Service Class
Network control	Critical alarms	Critical
	Routing, billing, critical OAM	Network
Interactive	IP telephony	Premium
	Video conferencing, interactive gaming	Platinum
Responsive	Streaming audio/video	Gold
	Client/server transactions	Silver
Timely	E-mail, non-critical OAM	Bronze
	Best effort	Standard

QoS performance consistency

In order to maintain QoS performance for a given service offering, QoS technologies must be implemented consistently across the network. Since IP is the predominant networking protocol, IP QoS is offered across the network using DiffServ. However, implementing DiffServ to provide end-to-end QoS is not sufficient because packets traverse different link layers, each with their unique QoS technologies. Therefore, in order to achieve consistent end-to-end QoS performance, DiffServ must be mapped to the different link layer QoS technologies.

Mapping DiffServ to link layer QoS

DiffServ provides a standard set of service classes known as Per Hop Behaviors (PHBs). These PHBs also provide a standardized DiffServ Code Point (DSCP) value associated with the PHB. Therefore, the DSCP must be mapped to the different link layer QoS to provide the closest behavior possible to create an end-to-end service.

There are many possible approaches to mapping IP QoS to link layer QoS. The following sections will provide sample default link layer mapping that can be used to provide consistent behavior between the IP and link layer QoS technologies.

Mapping DiffServ over Ethernet

Ethernet provides eight classes of service via the three 802.1p bits. These eight classes traditionally have been used to provide eight priority levels. DiffServ can be mapped to the Ethernet 802.1p user priorities as illustrated in *Table 5*.

Note that in this example, 802.1p user priority 1 is not used. 802.1p user priority 0 must be the default for best-effort traffic which is why it is mapped to the DE (Default) DiffServ PHB.

Table 5. Mapping DiffServ to Ethernet 802.1p

DiffServ Code Point (DSCP)	Ethernet 802.1p User Priority
CS7, CS6	7
EF, CS5	6
AF4x ¹ , CS4	5
AF3x ¹ , CS3	4
AF2x ¹ , CS2	3
AF1x ¹ , CS1	2
DE, CS0	0

¹x=1, 2 or 3

Mapping DiffServ over ATM

The most popular ATM service categories are CBR, rt-VBR, nrt-VBR, and UBR. DiffServ can be mapped to these as illustrated in *Table 6*. By mapping DiffServ to the ATM service categories, IP QoS behavior is preserved over the ATM link layer.

Table 6. Mapping DiffServ to ATM service categories

DiffServ Code Point (DSCP)	ATM service category
CS7, CS6, CS5, EF	CBR or rt-VBR
AF4x ¹ , CS4	rt-VBR
AF3x ¹ , CS3	
AF2x ¹ , CS2	nrt-VBR
AF1x ¹ , CS1	
DE, CS0	UBR

¹x=1, 2 or 3

Mapping DiffServ over PPP

When using PPP over a low-bandwidth connection, the fragmented IP packets are assigned a PPP Class Number. By mapping DiffServ-marked packets to the PPP Class Numbers as illustrated in *Table 7*, IP QoS is preserved across the PPP connection.

Note that EF-marked traffic is marked with the highest PPP Class Number. This is done because the queuing delay over a low-bandwidth connection is high and therefore the EF-marked traffic must be transmitted ahead all other traffic to meet its strict QoS delay requirements.

Table 7. Mapping DiffServ to PPP Class Numbers

DiffServ Code Point (DSCP)	PPP Class Number (long sequence number format)
EF	7
CS7, CS6, CS5	6
AF4x ¹ , CS4	5
AF3x ¹ , CS3	4
AF2x ¹ , CS2	3
AF1x ¹ , CS1	2
DE, CS0	1

¹x=1, 2 or 3

Nortel Networks Service Class mapping

Nortel Networks Service Classes (NNSCs) allow network operators to construct end-to-end services. The NNSCs provide the default mapping and QoS behaviors as illustrated in *Table 8*. Traffic is classified and placed into a NNSC which provides the default QoS behaviors in each network router and switch. Products that do not natively support NNSCs can have the NNSC behavior constructed using a QoS policy management system such as Nortel Networks Opitivity Policy Services.

Table 8. NNSC QoS technology mapping

NNSC	DiffServ Code Point (DSCP)	ATM Service Category	PPP Class Numbers	802.1p User Priority
Critical Network	CS7 CS6	rt-VBR	6	7
Premium	EF, CS5	CBR or rt-VBR	7	6
Platinum	AF4x ¹ , CS4	rt-VBR	5	5
Gold	AF3x ¹ , CS3		4	4
Silver	AF2x ¹ , CS2	nrt-VBR	3	3
Bronze	AF1x ¹ , CS1		2	2
Standard	DE, CS0	UBR	1	0

¹x=1, 2 or 3

Summary

QoS technologies are required for every type of network. Low-bandwidth and leased bandwidth connections often require more complex QoS technologies to be implemented to provide adequate performance levels while maximizing bandwidth efficiency. The degree to which these are applied depends upon the services being offered to the subscriber. Finally, end-to-end services require QoS technologies to be implemented consistently across the network in order to achieve the performance required by certain applications or specified in an SLA to a subscriber.

References

- ATM Forum AF-TM-0121.000 Version 4.1 “Traffic Management Specification,” <ftp://ftp.atmforum.com/pub/approved-specs/af-tm-0121.000.pdf>
- IEEE 802.1D, “Media Access Control (MAC) Bridges,” <http://standards.ieee.org/reading/ieee/std/lanman/802.1D1998.pdf>
- IEEE 802.1Q, “Virtual Bridged Local Area Networks,” <http://standards.ieee.org/reading/ieee/std/lanman/802.1Q-1998.pdf>
- RFC 3270, “MPLS Support of Differentiated Services,” <http://www.ietf.org/rfc/rfc3270.txt>
- Optivity Policy Services, <http://www.nortelnetworks.com/products/01/optivity/policy/>
- PacketCable Standards, <http://www.packetcable.com/specifications.html>
- RFC 2205, “Resource ReSerVation Protocol (RSVP),” <http://www.ietf.org/rfc/rfc2205.txt>
- RFC 2474, “Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers,” <http://www.ietf.org/rfc/rfc2474.txt>
- RFC 2475, “An Architecture for Differentiated Services,” <http://www.ietf.org/rfc/rfc2475.txt>
- RFC 3246, “An Expedited Forwarding PHB,” <http://www.ietf.org/rfc/rfc3246.txt>
- RFC 2597, “Assured Forwarding PHB Group,” <http://www.ietf.org/rfc/rfc2597.txt>
- RFC 2686, “The Multi-Class Extension to Multi-Link PPP,” <http://www.ietf.org/rfc/rfc2686.txt>
- RFC 2748, “The COPS (Common Open Policy Service) Protocol,” <http://www.ietf.org/rfc/rfc2748.txt>
- RFC 2749, “COPS Usage for RSVP,” <http://www.ietf.org/rfc/rfc2749.txt>
- RFC 3084 “COPS Usage for Policy Provisioning,” <http://www.ietf.org/rfc/rfc3084.txt>
- UMTS Wireless standards—3GPP, <http://www.3gpp.org/>
- Y.1541 “Network performance objectives for IP-based services,” <http://www.itu.int/recl/recommendation.asp?type=items&lang=en&parent=T-REC-Y.1541-200205-I>

In the United States:

Nortel Networks
35 Davis Drive
Research Triangle Park, NC 27709
USA

In Canada:

Nortel Networks
8200 Dixie Road,
Suite 100
Brampton, Ontario L6T 5P6
Canada

In Caribbean and Latin America:

Nortel Networks
1500 Concorde Terrace
Sunrise, FL 33323
USA

In Europe:

Nortel Networks
Maidenhead Office Park
Westacott Way
Maidenhead Berkshire SL6 3QH
UK

In Asia:

Nortel Networks
6/F Cityplaza 4,
Taikooshing,
12 Taikoo Wan Road,
Hong Kong



Nortel Networks is an industry leader and innovator focused on transforming how the world communicates and exchanges information. The company is supplying its service provider and enterprise customers with communications technology and infrastructure to enable value-added IP data, voice and multimedia services spanning Wireless Networks, Wireline Networks, Enterprise Networks, and Optical Networks. As a global company, Nortel Networks does business in more than 150 countries. More information about Nortel Networks can be found on the Web at:

www.nortelnetworks.com

For more information, contact your Nortel Networks representative, or call 1-800-4 NORTEL or 1-800-466-7835 from anywhere in North America.

*Nortel Networks, the Nortel Networks logo, and the globemark design, BayStack and Passport are trademarks of Nortel Networks. All other trademarks are the property of their owners.

Copyright © 2003 Nortel Networks. All rights reserved. Information in this document is subject to change without notice. Nortel Networks assumes no responsibility for any errors that may appear in this document.

GSA Schedule GS-35F-0140L
1-888-GSA-NTEL

56058.25-021803