NORTEL
NORTHERN TELECOM

Bay Networks

Carrier Packet Networks

IP QoS—A Bold New Network

An IP Quality of Service backgrounder for service providers

White Paper

Abstract

Businesses use the Internet for remote access, information searches, e-mail, and other applications, but do not yet rely on it for all networking needs. Service providers see potential revenue growth in corporate networking services—*if* the security and performance issues of the current Internet can be resolved.

IP quality of service (IP QoS) refers to the performance of IP packet flow through networks. Its purpose is to deliver end-to-end QoS to user traffic. It is is characterized by a small set of metrics, including service availability, delay, delay variation, throughput, and packet loss rate. IP QoS is predicted to lead the way to high-margin business customers, higher-priced service levels, more efficient bandwidth use, and more. It will be a critical enabling technology for the growth of IP networks.

Corporate services are the primary focus of IP QoS, with Service Level Agreements (SLAs) defining the guarantees and responsibilities between subscribers and providers. To forge an agreement that customers can trust, a service provider needs a network with QoS capabilities and a policy management system to configure, control, and maintain performance levels.

Although some work has been done to research, define, and develop IP QoS systems, it is generally agreed that a mature architectural framework, the required supporting hardware, and the appropriate operational techniques are not yet in place.

The evolution of the IP network toward guaranteed QoS promises to be rapid, exciting, and rewarding.

Table of Contents

Executive Summary

From the user side, the Internet has become a powerful consumer and business tool despite its well-publicized shortcomings. Businesses are using the Internet for remote access, information searches, e-mail, and other applications, but do not yet rely on the Internet for all their networking needs.

From the service provider side, corporate networking services constitute a large and profitable revenue opportunity for providers who can solve the security and performance drawbacks of the current Internet.

A NEW LEVEL OF QUALITY

The cornerstone of future IP network growth will be *IP quality of service* (IP QoS). With IP QoS, service providers can achieve greater profitability through high-margin business customers, higher-priced service levels, more efficient bandwidth use, and more.

They can also be more competitive through enhanced service differentiation, better-than-best-effort service, and customized solutions.

IP QOS DEFINED

IP QoS refers to the performance of IP packet flow through one or more networks. The aim is to deliver end-to-end QoS to user traffic. IP QoS is characterized by a small set of metrics, including service availability, delay, delay variation (jitter), throughput, and packet loss rate.

Corporate services are the primary focus of IP QoS, with Service Level Agreements (SLAs) defining the guarantees and responsibilities between subscribers and providers.

ARCHITECTURE

To make a contractual agreement that customers can trust, a service provider needs a network with QoS capabilities and a policy management system to configure, control, and maintain performance levels.

Two IP QoS architectures—Integrated Services Architecture (Int-Serv) and Differentiated Services Framework (Diff-Serv) are currently defined by the Internet Engineering Task Force (IETF). Each has a role and they must be able to interwork.

Int-Serv is implemented at the edge of enterprise networks where user flows can be managed at the desktop user level. More scalable than Int-Serv, Diff-Serv is used in enterprise WANs and plays a key role in the service provider network, based on its ability to prioritize by application or traffic path.

NETWORK SOLUTIONS

In addition to the architectural framework, other elements are required to build real-world IP networks that meet QoS goals.

Routers and switches must meet carrier reliability goals. The network must recover quickly from nodal or link failures. And the QoS mechanisms at each node must be configured to act in concert to deliver end-to-end QoS across the network—a goal that cannot be realistically achieved in any sizable network without a policy manager.

Despite the early stage of development of IP QoS, many components of tomorrow's high-performance, reliable, and flexible IP network have been identified, including:

- Separating traffic according to classification into queues
- A policy manager for managing QoS and SLAs and configuring routers and switches
- Traffic marking and policing mechanisms for entry traffic
- Filtering exit traffic for security and congestion control
- Active output queue management
- Packet discard algorithms
- Monitoring traffic levels at each outgoing interface
- Traffic policies to ensure the safety of premium traffic
- Leveraging of ATM switching and QoS technologies

THE FUTURE

Though IP QoS is in its infancy, it is quite clear that it will be an absolute requirement in commercial IP networks. Its evolution will be rapid, exciting, and rewarding.

Introduction

During the past twenty-five years, the Internet has evolved from a U.S.-government-sponsored research network to today's international, commercially operated network. The first grand-scale application of the Internet Protocol (IP), the Internet is driving the migration of other data traffic from voice, frame relay, asynchronous transfer mode (ATM), and other network architectures to IP networks.

IP technologies are now established as the fundamental platform for the world of webtone and are generally predicted to play a critical—and perhaps dominant—role in the evolution of the public network and private networks such as corporate intranets.

Migrating business network traffic onto public IP networks—including virtual private networks (VPNs)—presents great opportunities for business customers to reduce operating costs, investment risk, and operational complexity.

REMAINING ISSUES

Despite the Internet's rapid growth, implementation issues remain.

For example, the emergence of multimedia traffic over IP networks places great demands on *quality of service* (QoS) in the IP environment. Through the efforts of companies such as Intel and Microsoft, multimedia applications have become an integral part of PC architecture, driving both public and private networks even more rapidly toward a diverse and challenging traffic mix.

Voice and fax over the Internet also provide convincing cost savings and threaten to revolutionize the communications industry. All of these real-time multimedia applications demand better than the current *best-effort* Internet QoS.

The fact is that today's Internet falls far short of delivering the kind of reliability and performance guarantees that enterprises are demanding and are accustomed to in their private networks. Businesses will not place their mission-critical data, voice, and multimedia applications onto public IP networks until they receive secure, predictable, measurable, and guaranteed service.

Furthermore, during the period that the Internet was enjoying such rapid growth, intense competition was pushing margins extremely low in the traditional IP services market.

It is very difficult, if not impossible, to create a successful business model based on a $9.95 per month (with per-hour charges) or a $19.95 per month (unlimited hours) pricing structure. To improve this picture, service providers are now striving to find new sources of revenue and service differentiation that can improve their margins.

QOS OPPORTUNITIES

Moving business traffic—primarily data, but some IP—based voice traffic as well-onto public IP networks is one of the huge opportunities identified by providers in recent years.

A major prerequisite for attracting business customers with this type of mission-critical traffic is to offer alternative IP-based services with guaranteed QoS. By implementing IP QoS solutions, service providers can achieve:

- **Profitability**—improving top-line revenue by attracting high-margin business customers and offering higher-priced levels of services while reducing bottom-line cost by using bandwidth more efficiently.
- **Competitiveness**—enhancing service differentiation by offering multiple classes of *better-than-best-effort* service and by offering customized solutions based on individual requirements.

However, the path to profitability and competitiveness is not straightforward at this time. IP QoS is still a relatively new concept, with vendors offering different proprietary solutions while standards are still being developed.

In this currently uncertain environment, service providers should ask themselves these questions when implementing an IP QoS solution:

- What set of service levels should I offer my customers?
- How can I simplify my IP QoS offerings to communicate easily with my customers?
- How can I offer and cost-effectively manage IP QoS on an end-to-end basis?
- How can I take advantage of my existing IP or ATM infrastructure?

- How can I prepare for future growth and emerging IP QoS standards?
- How can I offer IP QoS in conjunction with Corporate Virtual Private Intranet services?

Service providers who weigh these questions carefully before planning and building IP networks will have a distinct advantage over their competition.

IP QoS Defined

Most industry experts agree that QoS can be a critical differentiator among service providers. However, general agreement on key concepts and terminology relating to service attributes—an important prerequisite for building standardized service offerings—still lags behind.

For example, the term *IP QoS* itself is frequently misused, even by people in the industry. What is advertised as IP QoS is often a set of features for implementing a *class of service* (CoS).

In general communications parlance, *CoS* is a broad term describing a more or less standardized set of features and other characteristics available with a specific service or service package.

QoS is a more precise term, chiefly used to measure a specified set of *performance attributes* typically associated with a service. In the IP network environment, *IP QoS* refers to the performance of IP packets flowing through one or more networks.

Given the current drive toward greater performance and reliability on the Internet, the ultimate aim of service providers is to deliver end-to-end, guaranteed IP QoS to user traffic on IP networks—including data, video, multimedia, and voice.

As a first step toward meeting this goal, a clear definition of QoS, within the context of a definable administrative authority (such as the network defined by a service provider's demarcation points), is a critical prerequisite.

With this aim in mind, QoS can be characterized by a small set of measurable parameters:

- **Service availability**—the reliability of the user's connection to the Internet service.
- **Delay**—also known as *latency*; refers to the interval between transmitting and receiving packets between two reference points.
- **Delay variation**—also called *jitter*, refers to the variation in time duration between all packets in a stream taking the same route.
- **Throughput**—the rate at which packets are transmitted in a network; can be expressed as an average or peak rate.
- **Packet loss rate**—the *maximum* rate at which packets can be discarded during transfer through a network; packet loss typically results from congestion.

With these definitions and parameters in mind, it is now time to look at a key mechanism that can help to ensure QoS in the IP network of the future.

Service Level Agreement

Service Level Agreements (SLAs), although usually thought of in conjunction with VPNs, can apply to all customers of a service provider, including dial-up, corporate, wholesale, or peer network users. An SLA could be a simple standard contract for mass consumers or customized and multidimensional for business customers.

An SLA defines end-to-end service specifications and may consist of the following:

- **Availability**—guaranteed uptime, service latency (where relevant, this is the delay accessing the network)
- **Services offered**—specification of the service levels offered
- **Service guarantees**—for each class; for throughput, loss rate, delay, delay variation, and class over-subscription handling
- **Responsibilities**—consequences for breaking the contract rules; location of the demarcation point; *24 x 7* support and customer service
- **Auditing the service**
- **Pricing**

TABLE 1. QUALITY OF SERVICE PARAMETERS

Service Level	Application	Priority Mapping
1	• Non-critical data • Similar to Internet today (see UBR on ATM) • No minimum information rate guaranteed	• Best-effort delivery • Unmanaged performance
2	• Mission-critical data • VPN outsourcing, e-commerce • Similar to frame relay CIR, ATM VBR	• Low loss rate • Controlled delay and delay variation
3	• Real time applications • Video streaming, voice, videoconferencing	• Low delay and delay variation • Low loss Rate

Central to the service level agreement are the service levels or classes that are available to the user's traffic. *Level of service* (LoS) and CoS are often used interchangeably. Traffic traveling under different service classes receives different levels of quality. An important function of the SLA is, therefore, to assign responsibility for mapping traffic to the different service classes offered.

Developing IP service levels is going to require a phased approach. In the first phases, very simple schemes will be implemented such as the two-bit differentiated services architecture (see reference 1) or the Assured Service (see reference 2), where only two to four service levels are defined. Subsequent phases will be evolutionary based on experience with early deployments and development of the market.

Another factor in favor of simplicity and a limited number of service levels is the user's perception of quality. Even when users can detect variations between the service classes through measurement and monitoring, they have not indicated the willingness to pay an incremental amount for the differences between highly granular performance variations. Early services will likely identify a premium service for mission-critical applications with guaranteed delivery and well-controlled delay, jitter, and throughput.

The next step may be to allow integrated services, with a low-delay, real-time service for voice applications. The natural environment to offer these services is within VPNs for intranet traffic. Table 1 shows an example of a simple set of IP QoS levels and their associated applications.

Note that the example in Table 1 represents current industry thinking about a simple move beyond the best-efforts-only Internet services that users are familiar with today. As the technologies, techniques, and service offerings mature, more sophisticated services will almost certainly be developed and marketed.

A final broad point should also be made about SLA. Because a legal contract is in place between the two parties, each desires to monitor the service performance and usage for different purposes.

The customer monitors to ensure the service provider is meeting the terms of the contract and to track utilization for it's own purposes, one of which may be internal accounting. The service provider monitors to verify any complaints made by the customer and for early detection of any potential violation in order to take preventative measures.

There is also monitoring to ensure that the customer is not over-subscribing services—although this is usually part of traffic conditioning at the trusted boundary point of the service provider's network (discussed later in detail).

IP QoS Architecture

A number of QoS architectures have been defined by various organizations in the communications industries (see reference 3). For IP QoS, the researchers are now focusing on two architectures developed by the

MORE ABOUT INT-SERV

The Integrated Services (*Int-Serv*) model for IP QoS architecture defines three classes of service:

- **Guaranteed**—with bandwidth, bounded delay, and no-loss guarantees.
- **Controlled load**—approximating best-effort service in a lightly loaded network.
- **Best-effort**—similar to what the Internet currently provides under a variety of load conditions, from light to heavy

Using a method similar to ATM's SVCs, Int-Serv uses RSVP between senders and receivers for per-flow signaling. RSVP messages traverse the network to request/reserve resources. Routers along the path—including core routers—must maintain soft states for RSVP flows.

Note: A soft state is a temporary state governed by the periodic expiration of resource reservations, so that no explicit path teardown request is required. Soft states are refreshed by periodic RSVP messages.

Internet Engineering Task Force (IETF)—the Integrated Services architecture (often referred to as *Int-Serv*), and the Differentiated Services architecture (often referred to as *Diff-Serv*).

INT-SERV

Int-Serv was defined in Request for Comments (RFC) 1633, which proposed the Resource Reservation Protocol (RSVP) as a working protocol for signaling in the Int-Serv architecture. This protocol assumes that resources are reserved for every flow requiring QoS at every router hop in the path between receiver and transmitter, using end-to-end signaling.

Scalability is a key architectural concern, since Int-Serv requires end-to-end signaling and must maintain a per-flow soft state at every router along the path. Other concerns are (1) how to authorize and prioritize reservation requests and (2) what happens when signaling is not deployed end-to-end.

It seems likely to current analysts that Int-Serv will be implemented at the edge of enterprise networks where user flows can be managed at the desktop user level. An important driver for Int-Serv in the vicinity of the desktop is Microsoft's implementation of RSVP and QoS capabilities in Windows 98 and NT 5.0.

MORE ABOUT DIFF-SERV

The Differentiated Services (Diff-Serv) model for IP QoS architecture uses a new implementation of the IP Version 4 type of service (ToS) header field. This field can now be marked, so that downstream nodes receive the information required to handle packets arriving at their entry ports and forward them appropriately to the next hop routers. Diff-Serv also renames the eight-bit ToS field as the DS field, with six bits available for current use and two reserved for future use.

Within the six available bits, only one mapping has currently been defined:

- **DE**—(Default), a best-effort class of service.

Another draft is proposing a second code point:

- **EF**—(Expedited Forwarding), not quantitatively defined at present; however, it is described as a forwarding treatment where the departure rate of the traffic from any Diff-Serv node must equal or exceed a configurable rate independent of the intensity of any other traffic attempting to transit the node; there are several implementation schemes that have been proposed but none is standardized yet.

DS Field

DSCP = Diff-Serv code point (6 bits)
DSCP = 000000 indicates DE
CU = currently unused (2 bits)
DSCP = 101100 indicates EF

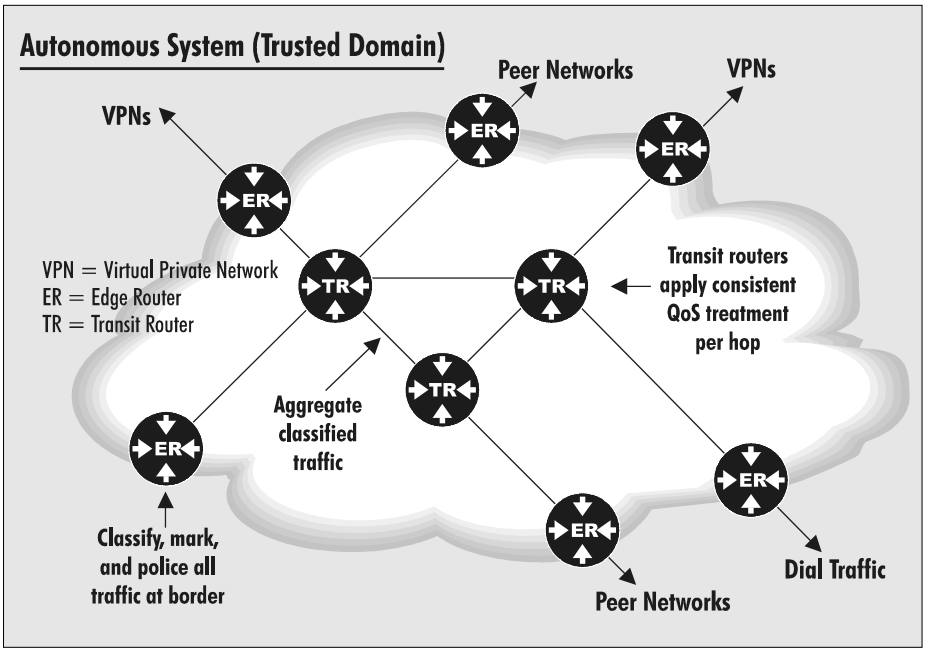

Figure 1. Diff-Serv framework.

DIFF-SERV

Diff-Serv is a relatively new IETF working group that has defined a more scalable way to apply IP QoS. It has particular relevance to the service provider and carrier networks.

Diff-Serv minimizes signaling and concentrates on aggregated flows and per hop behavior applied to a network-wide set of traffic classes. Flows are classified according to predetermined rules, such that many application flows are aggregated to a limited set of class flows.

Traffic entering the network domain at the edge router (ER) is first classified for consistent treatment at each transit router (TR) inside the network (see Figure 1). Treatment will usually be applied by separating the traffic into queues according to the class of traffic.

The eight-bit IP Version 4 type of service (ToS) field is used as a marker to notify downstream routers which treatment to apply to each arriving packet. Diff-Serv has renamed this field the DS (Differentiated Services) field.

Diff-Serv takes control of the ToS field and gives it a simple role in a flexible framework, so that equipment providers can develop configurable QoS capabilities that can interpret bit patterns (code points) in this field as sophisticated per hop behaviors.

Diff-Serv also outlines an initial architectural philosophy intended to provide a framework for interprovider agreements and make it possible to extend QoS beyond a single network domain (see Figure 2).

The Diff-serv framework is more scalable than Int-Serv because it handles flow aggregates and minimizes signaling, thus avoiding the complexity of per-flow soft state at each node. It will likely be applied most commonly in enterprise backbones and in service provider networks.

However, there will probably be domains where Int-Serv and Diff-Serv co-exist, so there is a need to interwork them at boundaries. This interworking will require a set of rules governing the aggregation of individual flows into class flows suitable for transport through a Diff-Serv domain. Several interworking schemes have been posited (see references 4 and 5).

The responsibility for mapping traffic to classes rests most logically with the customer. However, demarcation points can vary, so in some situations the service provider can manage this role on behalf of the customer. VPN services are particularly affected by such considerations, as will be discussed later in this paper.

REMAINING ISSUES

Diff-Serv lays a valuable foundation for IP QoS, but it cannot provide an end-to-end QoS architecture by itself. Effectively, Diff-Serv markings behave as a lightweight signaling mechanism between domain borders and network nodes, carrying information about each packet's service quality requirements.

Another set of requirements must be addressed before a workable implementation can be built. The principle requirements are:

1. A set of DS field code points in lieu of standards

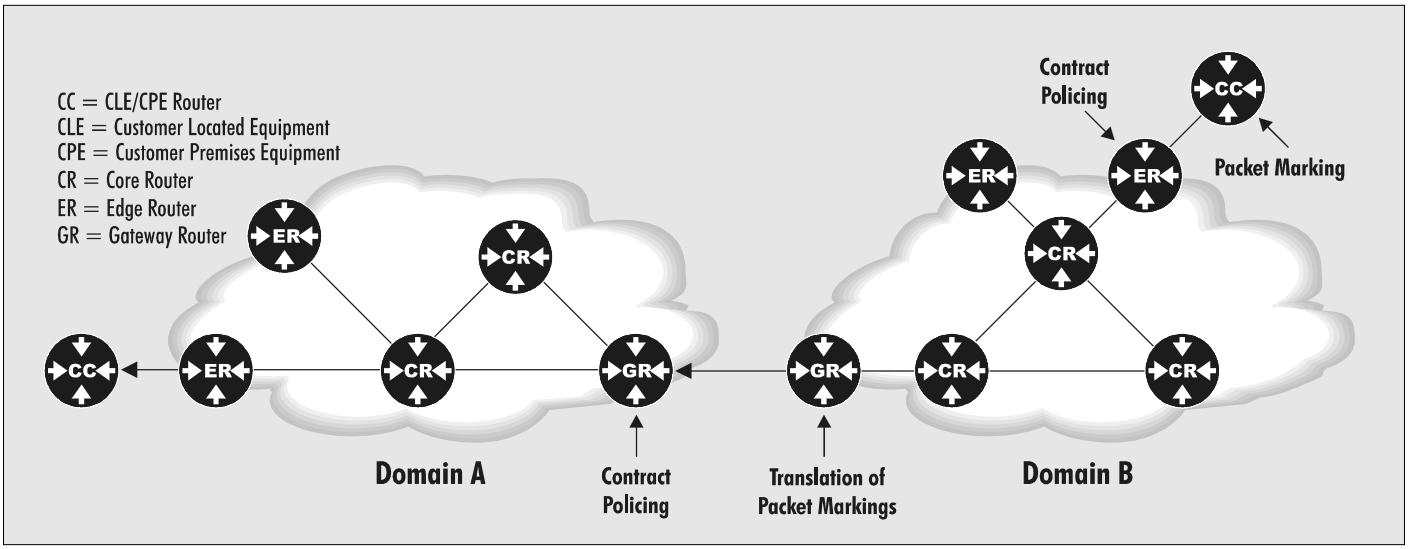

Figure 2. Diff-Serv inter-domain operation.

2. Quantitative descriptions of class performance attributes
3. A mechanism for efficiently aggregating the many sources of premium class traffic that can converge at transit routers
4. A solution to the single-ended SLA problem
5. An interworking solution for mapping IP CoS to ATM QoS
6. Management tools to facilitate deployment and operation

The first two points—standardized DS field code points and quantification of performance attributes—may not be as critical as some of the others in terms of developing standardized implementations. In fact, leaving these two issues unresolved will allow the service provider to develop proprietary solutions and achieve a competitive advantage.

However, lack of resolution in these areas is likely to slow down multi-domain service interworking. Moreover, providers may be able to negotiate agreements and service mappings at borders despite the lack of standardization.

Point 3—aggregation at transit routers—seems much more serious at this juncture of the evolution of IP QoS (see "Traffic management for IP QoS" later in this paper for potential solutions to this problem). It should be noted, however, that aggregation at transit routers is an issue that the communications industries have much to learn about. It will take some experimentation to find which levels of premium traffic can be handled safely. Initially, premium traffic may represent less than five percent of total traffic, but it may increase as confidence rises and new techniques emerge.

Point 4—the single-ended SLA problem—is also a serious challenge. Diff-Serv only manages traffic at the network entry points and does not provide a way to ensure appropriate exit capacity. This is particularly problematic in VPNs, where even high priority traffic might not terminate at a site if the access link is blocked by traffic from other sites. One solution is to over-dimension the access link. Another is to implement filtering (see "Traffic filtering" later in this paper).

Point 5—IP/ATM QoS interworking—is also challenging. Although ATM has excellent and well-defined QoS capabilities, they are path-based. Unfortunately, techniques for mapping IP packets to paths are still at an early stage of development, much like the Int-Serv and Diff-Serv QoS architectures. In addition, ToS-based routing has largely been unimplemented in routing protocols, since the IP ToS field has not been used by applications until recently.

Other ATM solutions are either scale limited—such as Multiprotocol Over ATM (MPOA)—or are proprietary and unlikely to be standardized. A scheme that many industry experts see as more promising in terms of standardization and scalability is Multiprotocol Label Switching

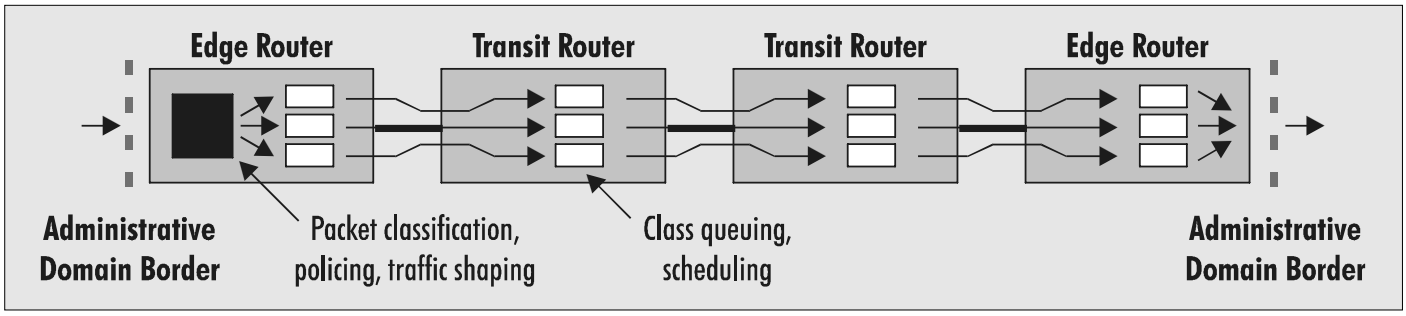

Figure 3. Traffic flow across a domain.

(MPLS). See "Leveraging ATM Infrastructure" later in this paper for a discussion of MPLS.

Resolving the final point—the need for management tools—should also prove to be a formidable task. Note that IP QoS is a framework around which service quality can be designed and engineered. It requires a large number of other mechanisms and network elements to operate in har mony before end-to-end service quality can be delivered to users.

Because of the highly distributed nature of these components and the need to manage them centrally, a set of management tools is a critical requirement. The policy manager is the delivery vehicle for this tool set (see "Policy-based management" later in this paper for a discussion of this topic).

NETWORK DELAY

Four different types of delay have been identified in IP networks:

Propagation delay: An inherent delay associated with signals traveling on any physical medium. In the case of fiber optics, propagation delay is somewhat more than the speed of light delay (the theoretical minimum).

Link speed delay: Data transfer rate is determined by the bit rate of the link. A fast link will obviously transfer a packet much faster than a slower link, so the slower link introduces a relative delay. Link speed delay is independent of propagation delay and is by far the greater of the two components.

If traffic is allocated some share of a very fast link (such that its capacity is the same as if it fully occupied the capacity of a slower link), link delay can be reduced—provided that interleaving is at the packet level.

Queuing delay: Every switch and router employs queues, where packets can be stored until capacity is available to transfer them out to the link. Time spent in queues constitutes queuing delay, which accumulates with each device traversed.

Hop Count: Each switch or router traversed by a packet is considered a *hop*. Queuing delay grows as hop count increases, so hop count is an important metric to control.

Implementing IP QoS

Figure 3 shows how traffic flows across an IP network through queues at each node. Queues are provided at each outgoing interface, and, when appropriate, there is a dedicated queue for each traffic class.

The transit routers implement queuing at their output interfaces. Policing is not needed because traffic arrives only from reliable sources.

Based solely on a packet's DS marking, it is inserted into the associated class queue at the appropriate outgoing interface. The traffic in output queues is conditioned by traffic management mechanisms acting on each queue to create a well-defined class behavior. Key functions are allocation of the output bandwidth and establishing rules for how to drop packets when congestion occurs.

Edge routers have the same capabilities as transit routers, but use policing to monitor the customer contract and a classifier to classify and mark the traffic at the incoming interface. The packet arrival rate can be measured for each class to ensure compliance with the SLA. In most cases the average rate over a defined period is checked to minimize the effects of

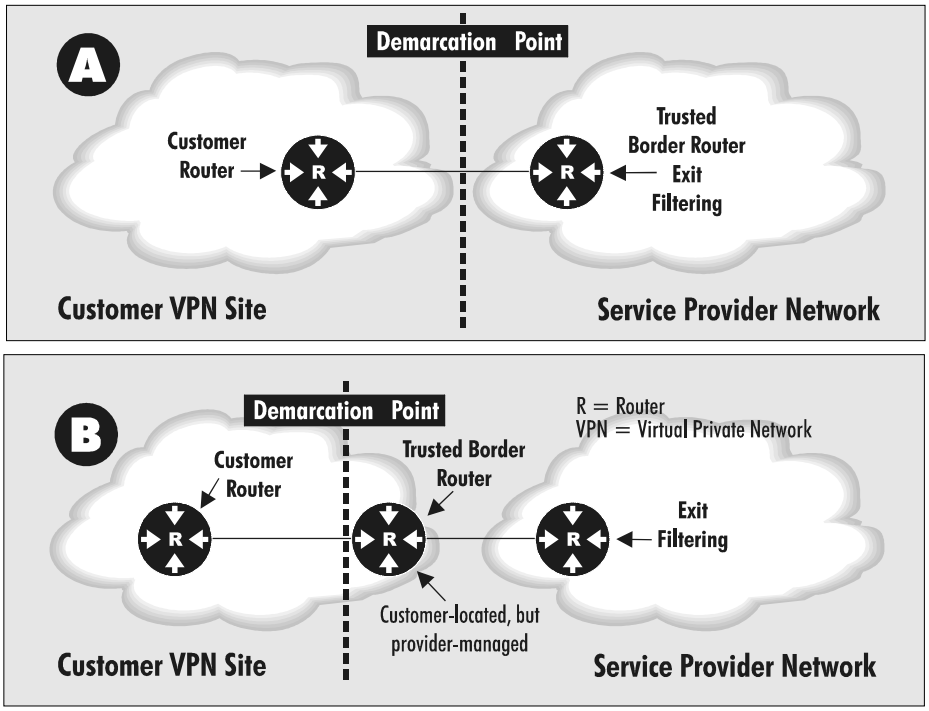

Figure 4. Traffic flow across a domain.

bursty traffic. Traffic can be classified in a number of ways, which are discussed later in this paper.

SLA AND NETWORK DESIGN

Earlier in this paper, the "Service Level Agreement" section discussed the various specifications of an SLA. The following sections discuss two specifications that relate directly to network attributes—availability and service guarantees.

Availability

Availability requires a network robust enough to survive failures such as a fiber cuts, port failures, or switch failures. Today, transport equipment often provides survivability of physical media failure that is almost transparent to higher network services.

Thus, in the case of a fiber failure, IP traffic may be totally unaffected. However, for equal service availability in the case of a non-transport related failure, the network must maintain services—particularly premium services—while minimizing service degradation overall.

One important part of managing service availability is ensuring that the traffic mix is composed of sufficient amounts of drop-tolerant traffic to prevent service degradation from affecting SLA traffic.

Service guarantee factors

The following paragraphs describe the challenges confronting the industry in the evolution toward reliable service guarantees.

Nodal Delay—such as propagation and link speed delay, which are relatively constant, and queuing delay are introduced into the network at each node (see the "Network Delay" sidebar on this page). Network design and planning can control link speed and minimize hop count.

Nodal delay can also be controlled in the queuing stages, where some traffic can be segregated by characteristics scheduling factors into queues, so that a share of the output link is allocated according to traffic engineering rules.

Delay variation—(or *jitter*) can be introduced by path variation, especially when poor network design is a factor.

However, most delay variation results from variations in queuing duration and packets getting stuck behind other long packets. Class-based queuing and output scheduling can be used to reduce jitter for premium types of traffic.

Loss Rate—defines the probability that a packet will be dropped before delivery to the destination. The transient nature of IP traffic patterns makes it difficult to eliminate packet loss.

Over-engineering link capacity is one solution but this may not be cost effective. It is virtually impossible to over-dimension links to the point that no traffic is ever lost. In addition, when failure conditions are taken into account, average utilization would be very low.

A reasonable solution is to implement *some* over-dimensioning and maintain a mixture of high- and low-value traffic, so that low-value traffic is potentially over-subscribed, but insensitive to loss in the event of failure or traffic surges.

QUEUING AND SCHEDULING MECHANISMS

First-In First-Out (FIFO)—the most straightforward approach and very simple to implement. However, with FIFO, a high-priority packet could be stuck behind thousands of best-effort packets.

Strict priority scheduling—where a class is served only if there are no queued packets belonging to any higher-priority classes. This is simple to implement but suffers from the problem that all but one class (highest-priority) could starve.

Fair Queuing or Round Robin (RR)—simple round robin scheduling from multiple queues. This helps in making the bandwidth availability fair to the different queues. One of the problems with fair queuing is that streams with large packets require a bigger share of the available bandwidth.

Weighted-Fair Queuing (WFQ)—an improvement to Fair Queuing. In this scheme, each queue is given a weight that determines the share of that queue to the link bandwidth.

Class-based queuing—uses several queues, each corresponding to a different traffic class (probably as defined by the PHB). Different methods for servicing or scheduling the queues can be used.

Hierarchical Class Based Queuing (CBQ)—Traffic is divided into classes and each class can have sub-classes. This hierarchy forms a tree. If a sub-class exceeds its share of link throughput, it will first try to borrow bandwidth from its sister sub-classes. This tree can be used to distinguish between types of traffic at many hierarchical levels.

Finally, this discussion points to two critical prerequisites for making service guarantees possible:

- *Good network design* is a pre-requisite for QoS delivery.
- *Queuing and scheduling mechanisms* in routers and switches play a vital role, which must be examined.

IP QoS Traffic Management

There are three distinct phases in the flow of every packet through a demarcated network (or *domain*). They are the *entry* phase, the *forwarding* phase, and the *exit* phase.

The network operator, whether a business customer or service provider, must first be concerned with how traffic enters its domain—typically via a trusted border router. This router applies appropriate traffic management processes (or *mechanisms*) to the traffic by agreement between the network operators on each side of the border. The agreed-upon mechanisms that control traffic entry and exit are the basis for the term *trusted border router*.

The network operator must also be concerned with how traffic flows within the domain (for example, is premium traffic handled consistently at each hop?), and with how traffic exits the domain (is exiting traffic marked appropriately for handling in the desired manner after leaving the operators domain of control?).

Thus, at each phase of the journey through a domain, packets may encounter multiple traffic management mechanisms—such as policing, security, filtering, conditioning, or classification mechanisms—that influence the quality of service during the journey.

ENTRY ARCHITECTURE

Upon entering a domain, a packet can be examined in a number of ways, not all of which are necessary for a particular type of traffic. Figure 4 shows two contrasting examples of traffic flowing between a business customer's VPN and a service provider network.

In Figure 4A, the customer owns and administers a WAN access router, typically shaping traffic into the link. Packets are classified by marking the DS field according to agreed-upon policies.

TABLE 2. TRAFFIC CLASSIFICATION

Network Layer	Application	Priority Mapping
4	Port Number	n/a
3	Type of transport protocol	ToS/DS field
2	n/a	Ethernet 802.1p, ATM, frame relay

PACKET DISCARD MECHANISMS

Tail drop—drops arriving packets only when the allocated buffer space is fully occupied. While being easy to implement, it is well known that this approach can lead to network collapse because it triggers the TCP global synchronization.

Random early detection (RED)—very effective at breaking TCP global synchronization. The idea is to try to maintain a small average queue size by randomly dropping arriving packets as the queue occupancy starts building up (but long before real congestion occurs). This causes only a few TCP sources to slow down and reduces the potential for congestion. The probability that an arriving packet will be discarded increases as the average queue size increases. Weighted RED (WRED) is a variant of RED that attempts to influence the selection of packets to be discarded. There are many other variants of RED.

Marking allows the network operators to aggregate individual flows from Int-Serv domains as discussed earlier in this paper. In this case, the trusted border router in the service provider's network *polices* the contract for compliance.

In Figure 4B, the service provider owns and administers the router collocated at the customer's site, so the demarcation and the policing points shift.

In this case, the service provider can shape traffic across the access link according to both the customer's policies and its own. Of course, this type of cooperative arrangement would depend on the level of trust between the parties.

In addition, the connection from the customer's network to the service provider's collocated router can be over Ethernet, instead of a WAN interface, as required in the architecture shown in Figure 4A. This gives the customer the option of classifying traffic using the Ethernet 802.1p priority scheme and letting the service provider map the priority to the packet's DS field according instructions in the SLA.

TRAFFIC FILTERING

Filtering is typically applied to traffic exiting a domain. Exit requirements may simply be filtering for security purposes and to prevent the access link from becoming blocked by low-value traffic.

For example, an exit-filtering policy might be used to dimension traffic termination capacity from other sites so that mission-critical traffic has priority to terminate over low-value traffic. This mitigates some of the problems of single-ended contracts alluded to earlier.

Security filtering might also be needed to prevent unauthorized traffic from entering a private domain. Filtering must be done at the service provider's end of the access link. Otherwise, malicious users could flood the link, causing denial of service for legitimate users. Thus, in Figure 4A and 4B, filtering is implemented on the edge router located at the service provider's premises.

Forwarding behavior in this case is different from classical IP forwarding in the sense that traffic is intentionally treated unequally so that packets marked for better treatment can be isolated and handled consistently at each hop. Forwarding treatment is applied at every stage including entry and exit.

TRAFFIC CLASSIFICATION

A network implementing Diff-Serv defines a standard set of classes throughout the domain. The number of classes may grow over time, but is relatively static and independent of the number of customer SLAs supported.

All traffic inside the network is treated as a standardized set of class flows. Customer service differentia-

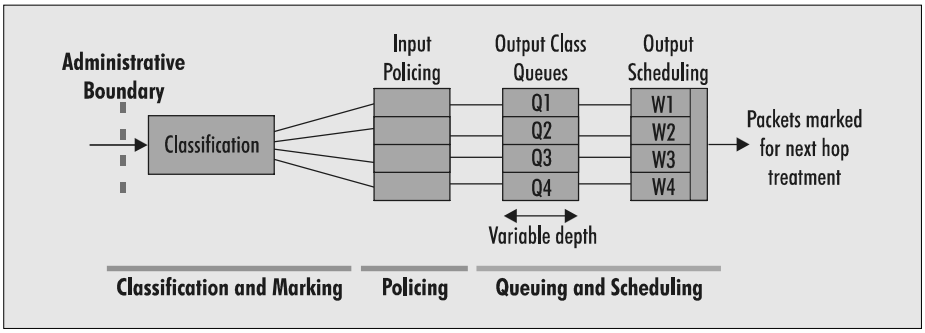

Figure 5. QoS functional model.

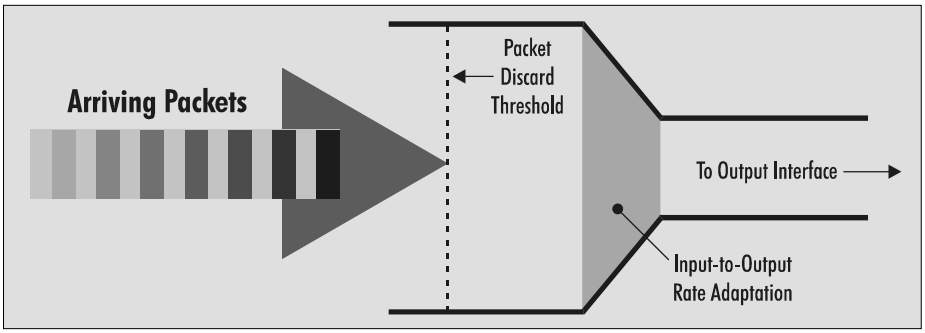

Figure 6. Functional model of output queuing.

tion is achieved entirely through contract negotiation and shaping at the point of entry. Typical customer-specific parameters might be price, penalty clauses, capacity per class, filtering or others.

Traffic entering a Diff-Serv domain must be classified for treatment inside the network. It must either be pre-marked by the customer or marked at the first router on the service provider's side of the demarcation point (see Figure 4).

Customer traffic classified by the service provider's edge router can be based on multiple criteria, ranging from the interworking of various priority schemes to application level analysis of traffic within the IP packet. Table 2 summarizes the options. It should be pointed out that security mechanisms, such as encryption and IPSec, will in some cases prevent application level analysis and classification of the traffic.

TRAFFIC POLICING

Traffic policing is implemented using a classifier (for classifying traffic), a token bucket or similar mechanism (for monitoring entry traffic levels at each class), and markers (for identifying or downgrading non-compliant traffic). Figure 5 shows the QoS functional model, including the policing segment.

Note that downgrading non-compliant traffic on a per-packet basis is not generally considered useful. Diff-Serv deliberately does not look at flows, so downgrading some packets from a premium flow would cause packet re-ordering—which defeats the purpose of enhanced service quality.

TRAFFIC CONDITIONING

Traffic at output interfaces is first classified and inserted into the correct output queues. Each queue will have selectable drop algorithms such as Random Early Detection (RED) or tail-drop, configurable by the requirements of the class. Each queue will also have programmable schedulers that implement algorithms such as Weighted Fair Queuing (WFQ), Round Robin (RR), and strict priority. These algorithms are also configurable by class requirements.

Figure 6 shows how queues adapt arrival rates to the output interface rate.

In addition, to accommodate different throughput and delay requirements of a class, queue depth is also a configurable parameter. However, there is a tradeoff to be aware of. Short queues can overflow quickly, but offer low delay. Longer queues are better at handling bursty traffic and provide enhanced throughput, but delay is correspondingly worsened. Queue depth must therefore be configured in conjunction with link scheduling and dimensioning in mind, as well as the characteristics of the traffic that will utilize the class.

Network Implementation

Network implementation can be just as complex as issues such as architecture, network design, standardization, and service levels. After all the industry standardization, planning, and development is done, networks must be built in a huge variety of environments, with complex hardware and software configurations, legacy devices, mixed technologies, and many other practical hurdles to overcome. This section provides some practical guidelines for network implementation.

TCP GLOBAL SYNCHRONIZATION

TCP Global Synchronization occurs when a large number of TCP sources lose packets at approximately the same time. This phenomenon leads to cycles of underload (when the involved TCP sources cut their rates simultaneously) and severe congestion (when the involved TCP sources ramp-up their rates simultaneously).

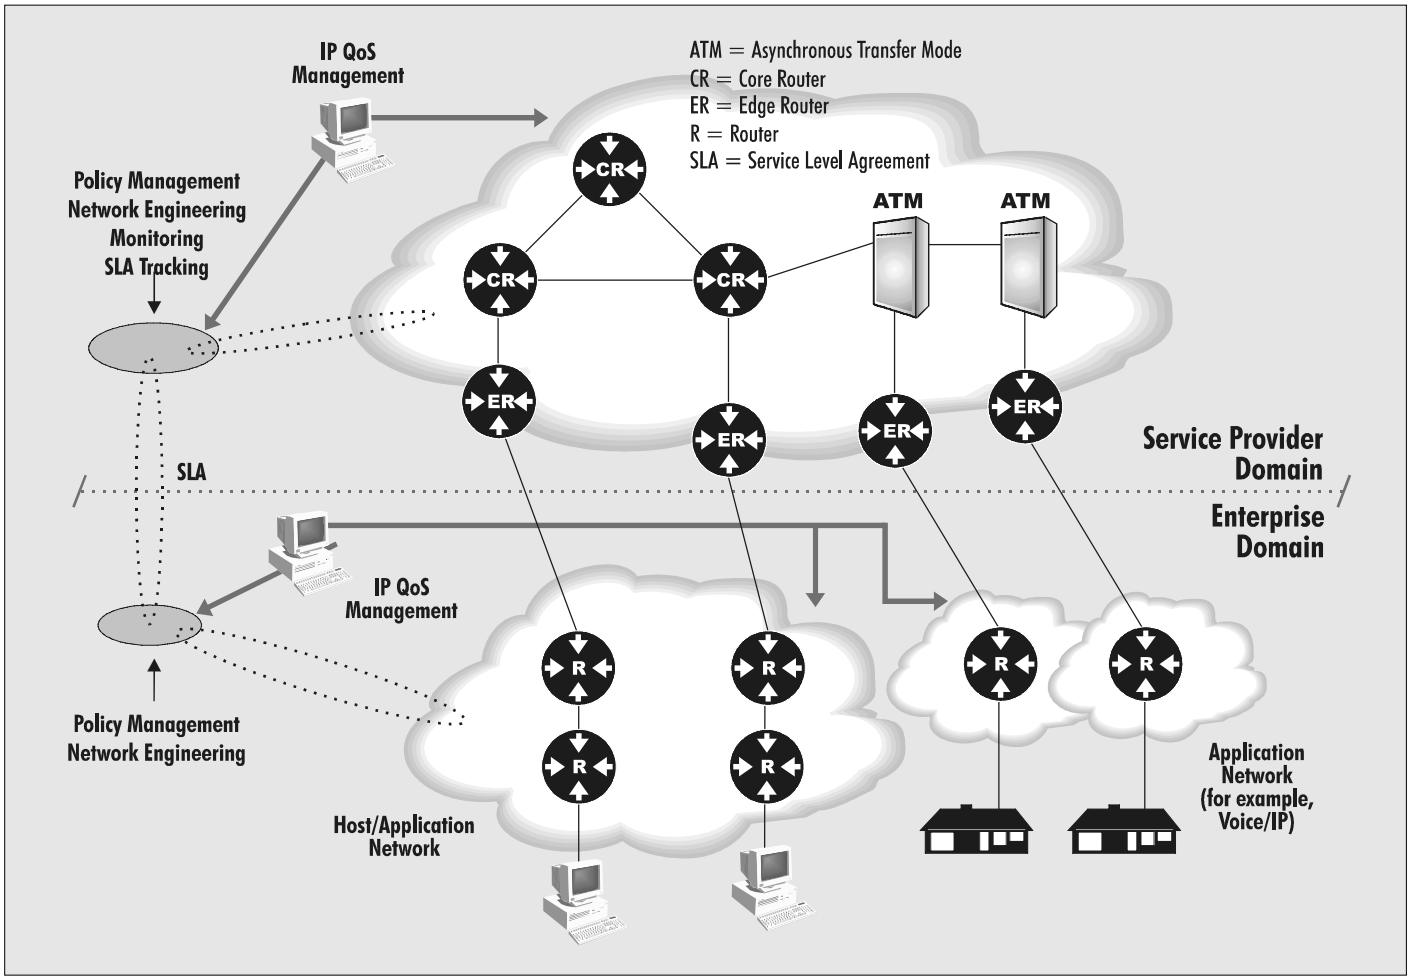

Figure 7. IP QoS architecture.

IP QOS ROUTER CHECKLIST

Router switches that can forward packets and apply traffic conditioning at wire speeds are going to be essential for IP QoS delivery. However, there are other important QoS-related factors to be aware of when selecting router products:

- **Carrier-class fault tolerance and reliability.** True carrier-class reliability will reduce routing instability and both support and improve availability guarantees to customers. The aim is to achieve the so-called *five nines* (99.999%) reliability.
- **Highly flexible QoS mechanisms.** QoS products should offer upwards of four queues (service classes) per interface with configurable discard and scheduling algorithms that can be selected independently for each queue. Look for a choice of mechanisms such as RED, WFQ, and strict priority, so that a rich set of service classes can be constructed.
- **Highly configurable QoS mechanisms.** QoS products should also be able to configure DS field code mappings flexibly to classifications that are user defined. Fixed or limited configuration capability could very quickly prevent service development and differentiation in both the current and future market environments, given the rapidly evolving standards that are predicted. Expect new mechanisms to emerge, such as the ability to create constant bit rate services by metering traffic onto the line.
- **Contract policing.** As service contracts become more complex, they should be rigorously checked for compliance. Token buckets or similar packet-counting mechanisms can be critical IP QoS components, since they allow traffic arrival rate to be verified for each class of service. This information

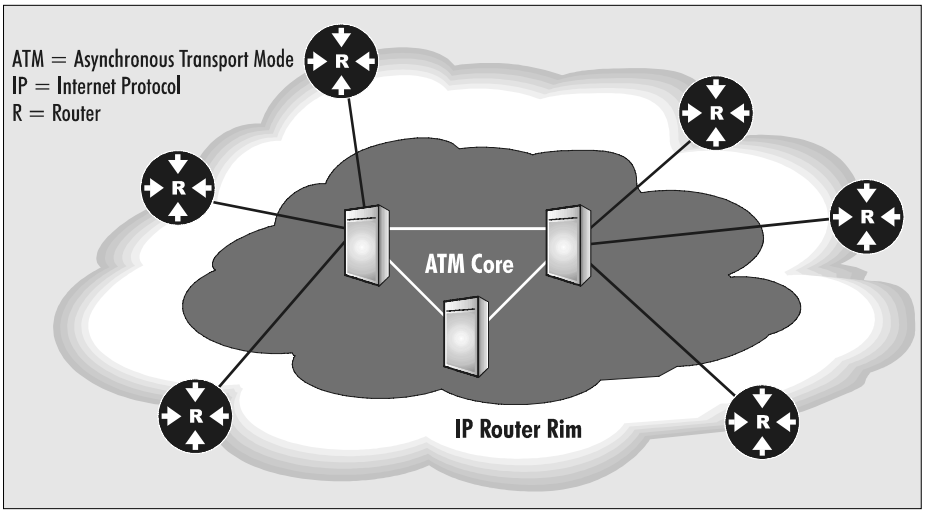

Figure 8. IP traffic channeled to an ATM core network.

can also be invaluable for billing and providing audit trails to customers.

- **Statistics gathering.** QoS products should offer a rich set of counters that can be configured to collect interface statistics on congestion and throughput by class. This information will be vital for traffic engineering and service monitoring.
- **Policy management.** Vendors who offer management tools that allow QoS to be configured and managed in multivendor installations will add value to their products and to customer and provider networks. IP QoS will be difficult to deploy in any reasonable-sized network without these tools.

COPING WITH LEGACY ROUTERS

As QoS services develop, routers will need to be able to process large numbers of packets at full wire-speed, which in the worst cases could be only 40 bytes long. However, legacy routers will be present in some networks, potentially limiting or complicating service offerings. Because of the per-hop behavior model of DiffServ, a network-wide set of QoS classes would have to default to the capability set of the lowest performing router.

A possible solution would be to confine legacy routers to best-effort only roles with policy- or QoS-sensitive routing techniques keeping valuable traffic away from them. For example, in a case where ATM is available alongside a legacy IP router network, the identified premium traffic can be groomed onto ATM virtual circuits with appropriate QoS attributes. Another option might be to redeploy the legacy devices to Internet traffic collector roles, feeding to new generation aggregation routers.

Hop-level packet re-marking is another potential limitation of legacy routers. Some routers assign hop behaviors to a non-user-configurable bit pattern in the IP precedence segment of the old ToS field. This would require packet remarking at the entry and exit of legacy environments within a network. Again, the impact could be limited by re-assignment to a best-effort role until scalability considerations allow these routers to be retired from the network economically.

LEVERAGING ATM INFRASTRUCTURE

Before delving into the technical issues of implementation, it is important to briefly consider the roles of ATM switches and IP routers and determine where and when they can be most effectively deployed.

Network solutions for adding QoS to IP traffic vary according to the needs of each service provider. When analyzed in detail, each proposed network has its own complex and subtle requirements, so a generalized approach can fail to find the optimum solution. With this caveat in mind, it is still useful to consider some general criteria involved in the decision process.

MPLS FOR IP AND ATM

Multiprotocol Label Switching (MPLS) labels are assigned at the network's edge router. Information from the routing protocols is used to assign and distribute labels to MPLS peers. In general, an MPLS node receives an outgoing label mapping from the peer that is the next hop for a stream, and allocates and distributes incoming labels to upstream peers for a given stream. The labels are extended into a switched path through the network (in a given service provider's domain) as each MPLS node *splices* the incoming to outgoing labels.

Figure 7 shows an architecture that includes IP routers and ATM switches at the core of an IP network, showing that either technology—and in some cases a mixed solution—is valid.

Considering the following factors can help a network planner decide the right implementation choice:

- Existing infrastructure
- Level of risk involved versus what is considered acceptable
- Time scale for maturity of products
- Amount of IP traffic and growth rate as a percentage of the total traffic mix in the network

As discussed earlier in this paper, there are areas independent of the technology where development and standardization are ongoing. There is, therefore, a choice to be made about how proven the technology is and whether it is standards-based or proprietary.

All of these factors affect risk. For example, choosing a new technique and an unproven product for the same network implementation raises the risk level—but could be acceptable for a new player seeking to steal market share from incumbent providers.

Another important factor is the percentage of IP traffic in the network. If the percentage is low and other types of traffic must be consolidated onto one network, ATM is a solid choice. Some of the more complex decisions arise for IP networks aiming to serve business markets. In this case, there is a particularly delicate trade off to be made between risk levels, time frame for network deployment, and the startup revenue needs of the business case.

Returning to technical consolidations, there are two primary *functional* roles to consider for ATM and IP router technologies—border traffic treatment and class handling inside the network. The model presented here allows Diff-Serv to be implemented over either an ATM- or IP-based core network.

IP can easily make use of the speed and performance of ATM at the core. Variable-length packet data can be adapted to the fixed-length cell transport using ATM adaptation layers (AALs). Both the adaptation to ATM and the switching of cells from one virtual circuit to another commonly take place in hardware. Figure 8 shows a rim of routers channeling IP traffic across ATM output interfaces towards an ATM core transport network.

MPLS can be implemented on ATM switches without modifying the hardware. Supporting MPLS on an ATM switch means that switch operation is controlled by the label switching component by running protocols such as OSPF, BGP, and PIM rather than protocols such as UNI and PNNI. RSVP is one of the methods for allocating QoS resources in IP networks.

More coarse-grained QoS capabilities can be supported by the Label Distribution Protocol (LDP). Such support would be more along the lines of *differentiated services*. The LDP provides the upstream node with Virtual Channel Identifier/Virtual Path Identifier (VCI/VPI) along with the CoS value. The VPI/VCI is used as a label, and QoS is signaled through LDP, based on the previously obtained CoS value in the IP header. The IP QoS Service Level is mapped into ATM as described in Table 1.

Handling of IP and ATM traffic will be based on common traffic management architecture. Some of the issues being investigated include MPLS/-ATM support for loop prevention. The interoperability between MPLS and the overlay ATM subnet require further investigation to eliminate the IP forwarding hop between the network boundary.

Traffic Engineering

For Diff-Serv to function, a traffic policy is required that allows relatively large amounts of traffic tolerant to packet loss to be dropped to ensure the safety of mission-critical and other highly valued traffic.

From the discussion of network issues in the previous section, it can be seen that network design and planning are an essential part of delivering service quality to users. Techniques such as policy or QoS-based routing can have tremendous value in networks with a diverse set of link media (such as wireless and satellite) such that application- and destination-based decisions allow traffic to be routed optimally.

However, path-based decisions have much less relevance to high-scale fiber networks, where delay and

bandwidth are much less of a limitation. Path engineering in this type of network is more relevant for route diversity—independent of the routing layer.

Managing Quality of Service

So far, we have considered how SLAs can be implemented from IP QoS structures within a service provider's network—independent of provisioning or maintenance of device configurations. In fact, configuration is not a trivial task, especially when one considers the number of queues that must be configured at each interface and the translation of SLAs into policing contracts at customer interfaces.

Policy management is the solution to this administrative challenge.

POLICY-BASED MANAGEMENT

In fact, policy management, in solving QoS administration issues, enhances the service provider's ability to manage network resources efficiently and offer subscribers new service features. With policy-based management, it is suddenly possible to control bandwidth utilization based on dynamic factors—such as time of day, application priority, and conditions in the network—according to defined policies.

Policies are used to define and dynamically control traffic behavior within a network domain. The alternative to policies is nodal configuration, where intended network-level behavior must be manually translated down to device-level instruction sets.

It may be helpful to think of policies as analogous to high-level programming language statements. Extending this analogy, a device configuration is analogous to a set of machine code instructions. Thus, the relationship of policies to device configurations is high-level to low-level.

In practice, a set of policies effectively creates a device independent program for the network. The program is verified for errors, such as policy conflicts (for example, a local policy might contradict a global policy), and compiled into device specific instructions.

One departure from the programming language analogy is that a program compiler generates machine code for a particular processor, while the policy generator has to create sets of device-level instructions for potentially many different types of network devices.

Different network devices might have equivalent sets of traffic management capabilities but different configuration requirements, a configuration which is reasonably straightforward to manage.

However, complexity arises when the devices have very different capability levels. In some cases, it may be satisfactory to restrict policies to the lowest common set of capabilities. However, in others, some level of manual intervention might be required to address this issue.

In time, these compromises will be eliminated with equipment and network evolution, but for now, they are key issues.

Thus, a policy-based manager (PBM) acts globally across the network domain, supervising device configurations that pertain to traffic management of user SLAs.

The PBM consists of five functions:

- Policy editing
- Policy verification and conflict resolution
- Policy generation
- Policy distribution
- Policy evolution

The policy editor is used by a network administrator to create the network and subscriber policies. Subscriber services (SLAs in particular) need to be interpreted into policy statements, a process that can be performed manually or automated by using service templates from a service management system. The entered policies must be checked for errors and potential conflicts before the device-level instruction sets are created for all the network nodes.

PBMs work with network management to distribute the configurations to the network elements.

Some policies may have dependencies (for example, a dependency on the network state or the time of day might exist), which are sensed by the PBM and result in updates to the device configuration of some nodes. The policy evolution stage looks after these activities.

The PBM system must be robust to failure, so it should use a distributed architecture. Of course, the administrator control console can be central-

ized to a few locations or even driven from a Web-based terminal that can be accessed from almost anywhere. While security imposes some restrictions and authorization requirements, such an architecture permits a high degree of flexibility.

The distributed architecture also mirrors the nature of global and local policies. *Globally* refers to policies that affect traffic at a network level. *Locally* refers to policies that affect a sub-set of the traffic, as is the case with customer SLAs.

Global policies may pertain to traffic dimensioning rules, nodal QoS requirements for the network service classes, and response actions in the presence of fault conditions. Local policies may include time-of-day and day-of-week dependencies, filtering policies for security, and SLA policies.

MONITORING AND TRACKING

Earlier in this paper, it was mentioned that enterprise subscribers and service providers need to monitor and track service quality to confirm that it is contract-compliant. To facilitate this requirement, the network nodes can collect and store statistics from each node about the traffic flowing through each of the queues.

A large amount of valuable information is thus available from each output and customer interface. Statistics can reflect average and peak throughput and packet discard levels for each traffic class. The statistics can be periodically collected from each node via Simple Network Management Protocol (SNMP) for storage and later processing.

Measuring delay is much more difficult, since it needs to be calculated between end points across the network for a particular packet. It therefore needs to be calculated periodically for each customer's traffic classes with delay variation being discovered over time from the minimum and maximum measurements observed.

Functionality for delay measurements could be integrated into edge nodes or implemented as separate monitoring equipment. The customer could either trust the service provider and request tracking reports prepared by the service provider or implement its own monitoring and tracking solutions at its premises. In the latter case, equipment at customer end points can communicate periodically and sample network performance.

Some of the collected statistics have the more valuable role of charging. Billing might be at least partially flat rate rental and independent of usage. However, SLAs could be written to allow some or all of the service to be usage based. For example, in the case where a fractional service that only partially uses the available capacity is deployed, the contract may allow flat rate up to a certain level but per-packet or per-megabit billing thereafter.

A View into the Future

IP QoS will be the cornerstone of carrier-class IP networking solutions that can be trusted to carry business-critical applications alongside public Internet traffic. Many processes have already been set in motion that will, in turn, trigger other processes and accelerate the evolution toward carrier-class IP networks.

The engine of change is competition for lucrative markets opened by telecom liberalization and the wealth of opportunities afforded by the technology change to connectionless IP networks. The key areas that will feed each other are discussed in the following paragraphs.

IP TRAFFIC PATTERNS

Analysts generally agree that over the next two years IP traffic will grow rapidly and will be the dominant form of traffic in the majority of service provider networks—not just ISPs. The industry structure will change as IP services continue to commercialize and the drive to create profitable businesses intensifies.

Problems associated with *hot potato* routing affect public Internet traffic quality in particular, since the path taken by user traffic is determined by how the user is connected to the service provider, how the service provider is connected to regional, national, and international networks, and the entire network path from user to destination point.

Current commercial pressure seems to be leading to the development of a three-tier hierarchy of providers—from small, local ISPs through larger regional ISPs up to national scale providers. Local ISPs will need to connect to national networks via regional networks.

The rule of markets should eventually limit the players at each level to

around three major providers (plus a number of niche providers), simplifying and improving interconnectivity between networks. The result of the simplified industry structure will be that traffic traversing multiple networks will be expedited by far better network performance.

As the industry structure changes, new content and service offerings with local and regional scope will emerge to take advantage of the improved connectivity between users and services.

In addition, new traffic patterns will result, based on communities of interest, so that most traffic will remain local—the reverse of the situation today. Improvements in caching techniques will also help to localize traffic patterns.

DEVELOPMENT OF IP QOS

IP QoS standards will be created both by standards organizations and by the process of *de facto* standards arising and gaining industry-wide adoption. Major areas in need of standardization will be traffic conditioning methodology, CoS definition, policy management protocols, and policy definition language.

Richer sets of QoS traffic conditioners will emerge to become the standard for routing and switching and to facilitate more advanced services and finer control. For example, traffic metering—where packets are transmitted at a fixed rate to break up packet trains and bursts—will help downstream aggregation and lead to more controllable traffic patterns.

Growing numbers of queues will be offered to facilitate finer service granularity. In the future, it could well be viable to define and allocate queues for specific flows.

Policy management algorithms will become fully tuned to network topology and other environmental factors to allow sophisticated high-level policies to be applied to the network, and more effectively govern SLAs, network and traffic engineering, and service restoration. For example, under certain failure conditions, some users may have the option of paying extra to receive the highest priority for early and preferential restoration.

The result of more well-defined traffic patterns and an enhanced ability to control IP traffic is that service quality levels will evolve rapidly and become available to subscribers in increasing numbers and richness.

References

1 "A Two-Bit Differentiated Services Architecture for the Internet," L. Zhang, V. Jacobson, K. Nichols, December, 1997.

2 "IP Precedence in Differentiated Services Using the Assured Service," S. Brim, F. Kastenholz, F. Baker, J. Renwick, T. Li, S. Jagannath, April, 1998.

3 "A Survey of QoS Architectures," C. Aurrecoechea, A. T. Campbell, L. Hauw, Center for Telecommunication Research, Columbia University, New York, NY 10027, USA.

4 "A Framework for Use of RSVP with Diff-serv Networks," Y. Bernet, R. Yavatkar, P. Ford, F. Baker, L. Zhang, K. Nichols, M. Speer, Internet Draft, June, 1998.

5 "A Framework for End-to-End QoS Combining RSVP/Intserv and Differentiated Services," Y. Bernet, R. Yavatkar, P. Ford, F. Baker, L. Zhang, Internet Draft, March 1998.

Glossary

24 x 7	24 hours, 7 days a week
AAL	ATM Adaptation Layer
Administrative Domain	An administrative partition of a network
ATM	Asynchronous Transfer Mode
BGP	Border Gateway Protocol
CBQ	Class-Based Queuing
CIR	Committed Information Rate
CLE	Customer Located Equipment
CoS	Class of Service
CPE	Customer Premises Equipment
CU	Currently Unused
DE	Default
Diff-Serv	Differentiated Services
Domain	Architectural partition of a network
Downstream	Network element or other network component that follows an upstream element
DS	Differentiated Services
DSCP	Differentiated Services Code Point
EF	Expedited Forwarding
ER	Edge Router
FIFO	First In, First Out
IETF	Internet Engineering Task Force
Gatekeeper	Element that manages addressing, admission, and bandwidth
IETF	Internet Engineering Task Force
Int-Serv	Integrated Services
IP	Internet Protocol
IPSec	Internet Protocol Security protocols
ISP	Internet Service Provider
LDP	Label Distribution Protocol
LoS	Level of Service
LSP	Label-Switched Path
MPLS	Multiprotocol Label Switching
MPOA	Multiprotocol Over ATM
n/a	Not Applicable
OSPF	Open Shortest Path First
PBM	Policy-Based Manager
PC	Personal Computer
PHB	Per-Hop Behavior
PIM	Protocol-Independent Multicast (both Sparse and Dense modes)
PNNI	Private Network-Network Interface
QoS	Quality of Service
RED	Random Early Detection
RFC	Request for Comments
RR	Round Robin
RSVP	Resource Reservation Protocol
SLA	Service Level Agreement
SNMP	Simple Network Management Protocol
SVC	Switched Virtual Circuit
TCP	Transmission Control Protocol
ToS	Type of Service
TR	Transit Router
UBR	Unspecified Bit Rate
UNI	User-to-Network Interface
Upstream	Network element or other network component that precedes a downstream element
VBR	Variable Bit Rate
VCI/VPI	Virtual Channel Identifier/Virtual Path Identifier
VPN	Virtual Private Network
WAN	Wide Area Network
WFQ	Weighted Fair Queuing
WRED	Weighted Random Early Detection

Information subject to change since Nortel reserves the right to make changes, without notice, in equipment design or components as engineering or manufacturing methods may warrant. Product capabilities and availability dates described in this document pertain solely to Nortel's marketing activities in the United States and Canada. Availability in other markets may vary.

For more information, or to order more copies of this document, contact your Nortel sales representative or call 1-800-4 NORTEL (1-800-466-7835) from anywhere in North America. Product and service information is also available on the Internet at Nortel's World Wide Web home page (http://www.nortel.com).

Address correspondence to:

In the United States:

Nortel
P.O. Box 13010
Research Triangle Park, NC 27709
USA

In Canada:

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

Published by:

Nortel
Marketing Publications, Department 4262
PO Box 13010
Research Triangle Park, NC 27709

74003.25/09-98 Issue 1
Printed in USA September 1998