

The ability truly to deliver quality of service will separate the winners from the losers in the packet-switched future

# The cost of quality in Internet-style networks

AMITAVA  
DUTTA-ROY  
Contributing  
Editor

**I**N BROAD TERMS, THE QUALITY OF service of a wide-area network is a measure of how well it does its job—how quickly and reliably it transfers various kinds of data, including digitized voice and video traffic, from source to destination. Back when networks dealt pretty much exclusively with voice telephony, the subject hardly ever came up. The circuit-switched telephone system was designed specifically to satisfy the human ear. It did, and it does.

Nowadays, with the advent of packet switching and the proliferation of many kinds of communications traffic (time-sensitive financial transactions, still images, large data files, voice, video, and so on), there are more than one set of criteria to satisfy. The data rate needed for satisfactory voice communication may take an intolerable time to transfer high-resolution images. Conversely, the degree of network latency acceptable in transferring some files may not be adequate for real-time voice. So quality of service (QoS) has become a hot topic, and the contracts that specify it, called service level agreements (SLAs), are becoming more and more common, at least between service providers and their largest customers.

In fact, as incumbent providers of telecommunications service are increasingly being challenged by competitive carriers, QoS has become a convenient marketing tool for both. "One way for a service provider to gain a competitive edge is to offer SLAs that guarantee QoS levels and offer rebates when those levels are not met," *IEEE Spectrum* was told by Benjamin S. Stump. He is senior product specialist in service activation and performance solutions at Telcordia Technologies, formerly known as Bellcore and still in Piscataway, N.J. The long-distance carrier AT&T Corp., New York City, for example, offers standard and gold versions of its SLAs. Rebates are credited to customer accounts when guaranteed service levels are not met.

## QoS DEFINED

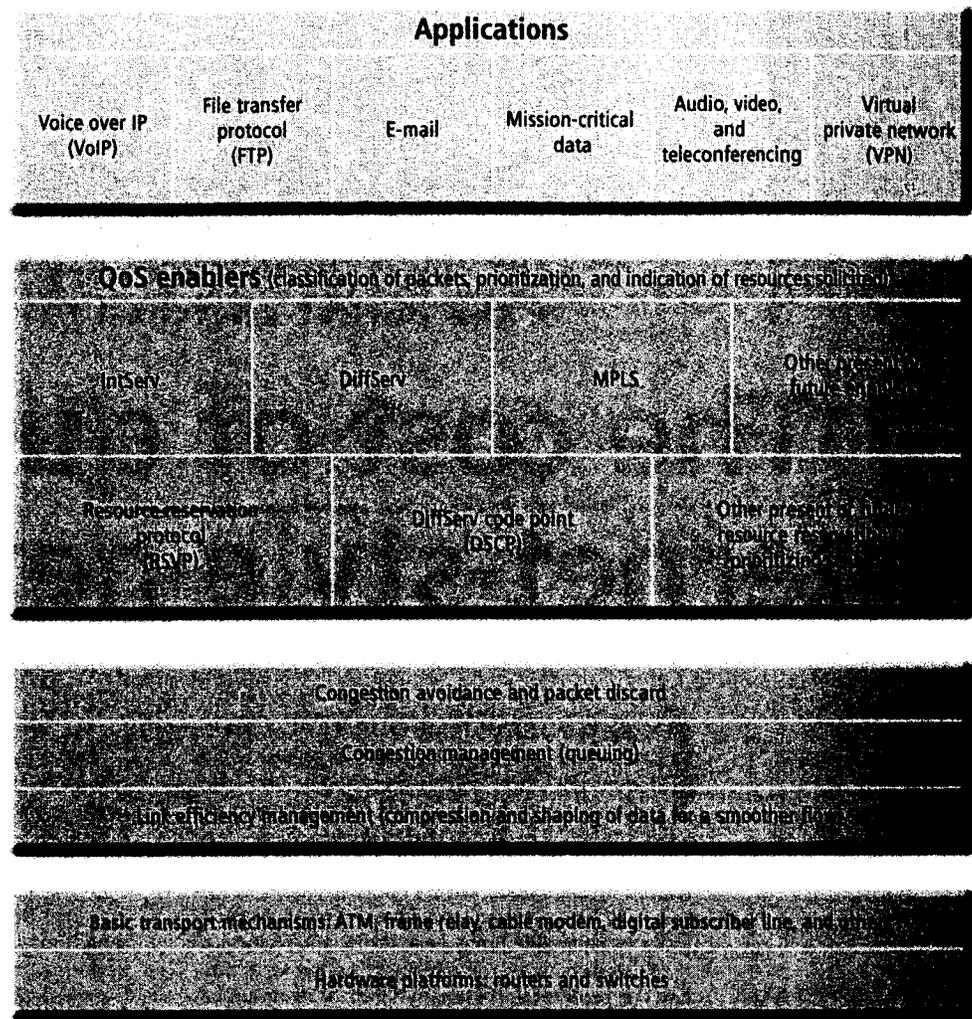
Technically, QoS refers to an aggregation of system performance metrics. The five most important of these are:

- **Availability.** Ideally, a network is available 100 percent of the time. Criteria are quite strict. Even so high-sounding a figure as 99.8 percent translates into about an hour and a half of down time per month, which may be unacceptable to a large enterprise. Serious carriers strive for 99.9999 percent availability, which they refer to as "six nines," and which translates into a downtime of 2.6 seconds a month.
- **Throughput.** This is the effective data transfer rate measured in bits per second—it is emphatically not the same as the maximum capacity, or wire speed, of the network, often erroneously called the network's bandwidth. Sharing a network lowers the throughput realizable by any user, as does the overhead imposed by the extra bits included in every packet for identification and other purposes. A minimum rate of throughput is usually guaranteed by a service provider.
- **Packet loss.** Network devices, like switches and routers, sometimes have to hold data packets in buffered queues when a link gets congested. If the link remains congested for too long, the buffered queues will overflow and data will be lost. The lost packets must be retransmitted, adding, of course, to the total transmission time. In a well-managed network, packet loss will typically be less than 1 percent averaged over, say, a month.
- **Latency.** The time taken by data to travel from the source to the destination is known as latency, or delay. Unless satellites are involved, the latency of a 5000-km voice call carried by a circuit-switched telephone network is about 25 ms. For the public Internet, a voice call may easily exceed 150 ms of latency because of delays, such as those caused by

[1] The similarities and differences among the various procedures for controlling quality of service (QoS) probably are best understood by viewing them in the same general framework as shown here.

Data from one or more applications [top] pass down through QoS enablers [green], which prioritize the data flows and indicate the resources each requires.

The data then continues through various levels of software and hardware that control packet discard mechanisms [blue] when buffered queues become too long. Finally it reaches the basic transport mechanisms and their hardware platforms [beige] that carry packets to the next node.



Source: Cisco Systems Inc.

signal processing (digitizing and compressing the analog voice input) and congestion (queuing).

- Jitter, which is another way of saying latency variation, has many causes, including: variations in queue length; variations in the processing time needed to reorder packets that arrived out of order because they traveled over different paths; and variations in the processing time needed to reassemble packets that were segmented by the source before being transmitted.

Applications vary in their QoS requirements [Table 1]. A long file transfer needs a high throughput and low packet loss, but is not very sensitive to delay and jitter. Live videoconferencing, on the other hand, also needs high throughput, plus it is sensitive to both delay and jitter. It is these differences that must be considered in writing the SLAs between service providers and their clients. The usual agreement specifies the end-to-end performance to which the client

is entitled over a specified time interval—a month or a quarter, for example.

#### A MATTER OF PRIORITIES

QoS is largely about priorities. At network aggregation points, like routers, multiplexers, and switches, data streams with different QoS needs are combined for transport over a common infrastructure. Satisfactory QoS has two main requirements: a means for labeling flows with respect to their priorities, and network mechanisms for recognizing the labels and acting on them.

Some networks—notably, those that use the asynchronous transfer mode (ATM) protocol—have extensive provisions of this kind. Unfortunately, the Internet does not, and neither do the similar IP networks based on the transmission control protocol/Internet protocol (TCP/IP) suite. So ensuring adequate QoS comes down to devising a means for labeling data flows and recognizing and acting on those labels.

IP is a best-effort protocol in that it does not guarantee delivery of data packets. Confirmation of the arrival of data packets at the destination is the responsibility of the TCP, which sits just above the IP in the well-known seven-layer open systems interconnection (OSI) reference model promulgated by the Geneva-based International Organization for Standardization (ISO), a worldwide federation of national standards bodies.

If any packet is not delivered (as determined by checking the sequence numbers of packets at the destination), TCP requests a retransmission of the missing packet, thereby ensuring that all packets eventually get to the destination. This is effective, but slow. Therefore, TCP is generally used by applications that are not time-sensitive.

Real-time applications cannot take advantage of TCP. Obviously, the time needed for keeping track of missing packets and retransmitting them is not acceptable in

such cases. So these applications rely on what is essentially a stripped-down version of TCP, known as the user datagram protocol (UDP), which runs faster than TCP by omitting some of its functionality. Applications that run over UDP must either have those missing capabilities built into them or else do without.

In the case of voice communications, where retransmitting packets takes too long to be of any value anyway, missing packets are simply lost. Internet telephony, therefore, will work only over networks that are quite reliable to begin with, like fiber-based nets with modern switches and routers.

### FIXING THE PROBLEMS

The Internet Engineering Task Force (IETF)—the protocol engineering and development arm of the Internet Society, which is headquartered in Reston, Va.—has proposed several methods for improving QoS, including IntServ, DiffServ, and MPLS [Fig. 1]. Some typical applications are indicated in the top layer of the diagram, while the second layer shows the different procedures proposed by the task force for handling them.

Integrated service (IntServ) is the earliest of these procedures. It assigns a specific flow of data to a traffic class, as it is called, which defines a certain level of service. It may, for example, require best-effort delivery or guaranteed delivery. It might even impose some limits on latency.

Once a class has been assigned to the data flow, a so-called path message is forwarded to the destination to determine whether the network has available the resources (transmission capacity, buffer space, and so on) needed to support that specific class of service. If all devices along the path are found capable of providing the required resources, the receiver generates a "resv" message and returns it to the source, indicating that the latter may start transmission of its data. The procedure, known as the resource reservation protocol (RSVP), is repeated continually to verify that the necessary resources remain available. If the required resources are not available, however, the receiver sends an RSVP error message to the transmitter.

Although IntServ has some attractive aspects, it does have its problems. One, obviously, is that it has no means of ensuring that the necessary resources will be available when wanted. Another is that it reserves network resources on a per-flow basis. If multiple flows from an aggregation point—say, a communications server in a local-area network—all require the same

resources, the flows will nevertheless all be treated individually. The resv message must be sent separately for each flow. In other words, IntServ does not scale well, and so wastes network resources.

The procedures of IntServ are improved upon in another method from the IETF, one known as differentiated service (DiffServ). With DiffServ, a short tag is appended to each packet depending on its service class. Data flows having the same resource requirements may then be aggregated on the basis of their tags when they arrive at the edge routers. The routers at the core can then forward the data flows toward their destinations on the basis of their tags without examining the individual packet headers in detail. Since most of the decision-making is in this way transferred from the core routers to the edge routers, the core network runs much faster [see "Diversifying Internet delivery," in *To Probe Further*, p. 62].

In the past, QoS planners supported both IntServ and DiffServ. At present, however, the trend is to use DiffServ supplemented by some of the resource reservation capabilities of RSVP.

### MULTIPROTOCOL LABEL SWITCHING

A newer approach to speeding the transit of data through a network is multiprotocol label switching (MPLS), also a procedure promulgated by the IETF. Normally, under IP, packet headers are examined at every transit point (multiplexer, router, or switch) in a network, which takes time and contributes to the overall data delay. A more efficient approach would be to label the

packets in such a way as to make it unnecessary for each IP packet header to be analyzed at points intermediate between the source and destination. Multiprotocol label switching does this by appropriately labeling IP packets at the input of label edge routers located at the entry points of an MPLS-enabled network [Fig. 2].

The procedure works like this: the label edge router examines the incoming packets and decides—based on the packet's source address, destination address, and priority level—where to send it for its next hop through the network. It also attaches a 32-bit tag, known as an MPLS label, to the packet. The MPLS label contains such information as whether the packet should be treated as MPLS traffic or routed as an ordinary IP packet, whether it conforms to IPv4 or IPv6, the packet's "time to live", and, of course, what its next hop should be. The edge router then forwards the packet to the router at the end of the next hop.

That router, in turn, examines the MPLS label and decides on the next hop for the packet. That second router then creates a second MPLS label. The two labels are swapped before the packet is forwarded to the second hop. The process is repeated until the packet reaches its destination.

This procedure has two advantages over normal IP routing. First of all, the routers along the path need not read and analyze a packet's complete header information, just the shorter MPLS label. This alone saves some time. Secondly, the swapping of labels leaves a trail in the registry of the routers that other packets in the same session can follow. Once the first packet establishes a

Traffic type	Sensitivities			
	Bandwidth	Loss	Delay	Jitter
Voice	Very low	Medium	Low	Low
E-commerce	Low	High	Low	Low
Transactions	Low	High	Low	Low
E-mail	Low	High	Low	Low
Telnet	Low	High	Medium	Low
Casual browsing	Low	Medium	Medium	Low
Serious browsing	Medium	High	High	Low
File transfers	High	Medium	Low	Low
Video conferencing	High	Medium	High	Low
Multicasting	High	High	High	Low

\* Complex contents may include audio and video clips and fast animations.

Source: CQOS Inc.

path, decision-making at intermediate points is eliminated to a great extent. This markedly speeds up the transfer of data.

Many network service providers have installed label edge routers and are about to roll out MPLS services. During the summer of this year, for example, Cable & Wireless PLC of London started offering MPLS for its transatlantic links, which join New York City and Washington, D.C., to London, Amsterdam, and Frankfurt, Germany. Cable & Wireless also plans to introduce MPLS in all of its OC-192 (9.953 Gb/s) fiber networks between now and the end of 2001.

### ENHANCED SERVICES COST MORE

Technologies that involve both software and hardware now exist to detect the requirements of each data flow on the fly—inferring them from, say, its source or destination IP address instead of reading them from a special label. Once a specific application in a session is detected, it can be given the priority to which it is entitled.

But until recently, a client's network administrator had to inform the service provider about each and every change in the priorities of data generated by certain applications. As this process costs time and money, many clients have been discouraged from requisitioning the enhanced services in the first place.

But a client can add advanced services much more easily, thanks to a new tool for assuring the QoS of a network. Known as the common open policy service (COPS) protocol, the tool is more adaptable to a customer's own requirements, allowing those requirements to vary with time of the day, application, or even user session. The

requirements and the rules for allocation of system resources, known as policies, are decided in advance. The objective is to specify a service in unequivocal terms and to allocate the resources required to deliver that service.

### COMMON POLICY

Policy information is stored in a policy server from where it is shared with other network devices using COPS. The rules follow an IF, WHAT, WHEN, and THEN logic. A typical sequence of events could be:

IF: The user belongs to the computer-aided design group # 003

and

WHAT: the application is the design of a rocket engine

and

WHEN: the time is between 0800 and 1400 hours on Monday through Friday

THEN: the user is entitled to: a service level, S, that gives a throughput of X kb/s with an end-to-end latency of no more than Y ms.

The service level could also specify other parameters, such as constant-bit-rate service.

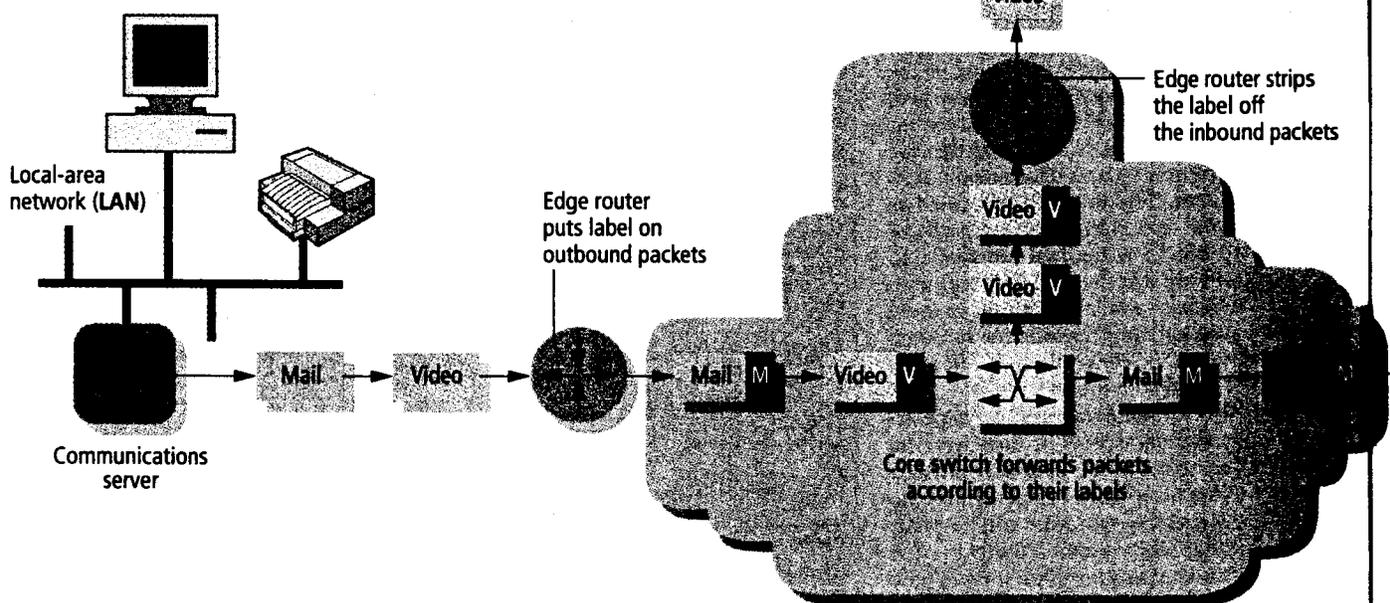
Once a user has put such a policy in place, it becomes easier for the client's network administrator to configure and adapt the system to the company's changing circumstances.

"Policy-based networking is instrumental to provide a friendly and dynamic environment of the user," said Cameron Sistanizadeh, founder of Yipes Networks Inc., San Fran-

cisco. "With proper flow-through operations support systems, network parameters can be configured to meet customer-initiated QoS requirements." Sistanizadeh's other title is vice president for network architecture for his company, a QoS-enabled service provider that offers COPS.

### CREATION, INTERPRETATION, ENFORCEMENT

The three principal elements of such a policy-based traffic management system are: policy creation and storage, interpretation, and enforcement [Fig. 3]. When a data packet arrives at the input port of the enforcement device, the device first determines the classification of the data by some predefined criteria. Then, using the COPS protocol and the well-established simple network management protocol (SNMP), it checks with the policy interpreter as to the



Source: *Managing Bandwidth*, A. Croll and E. Packman

QoS to which the packet is entitled. The policy interpreter, in its turn, verifies the status of the data by pulling the policy rules using the lightweight directory access protocol (LDAP), a protocol commonly used for exchanging information among directory databases.

With the help of the information thus retrieved, the interpreter determines what are called the rights of that particular data packet. On receipt of information on these rights, the enforcement device sends the packet, properly tagged, onward to a router. If the same type of data, such as a request to a specific Web site, is found to be repeating often, the rules could be temporarily cached in the enforcement device itself.

Vendors, such as Cisco Systems, Juniper Networks, Extreme Networks, and Nortel, are already shipping servers and routers that can run COPS. And start-up service providers like Yipes are creating dynamic QoS-on-demand environments, to provide capacity adjustable in increments that are as small as 1 Mb/s for time-sensitive applications.

But interoperability problems do exist when attempts are made to have products from different vendors work together. These problems will have to be solved before policy-based networking becomes ubiquitous.

Nevertheless, there is a growing worldwide adherence to COPS. Nicola Chiminelli, a research engineer at CSELT SpA, Telecom Italia's research institute in Turin, told *Spectrum*, "It's really important to change the paradigm used for identifying and classifying users, applications, and network resources, and the approach suggested by policy-based networking founded on COPS and LDAP standards seems to be the only

feasible manner to deal with the QoS in an efficient way."

Notwithstanding the methods used for assuring QoS in a virtual private network (VPN), measuring and displaying the parameters are clearly the bottom line. After all, if customers cannot feel assured of getting the service they are paying for, how likely are they to continue paying? Luckily, "the TCP/IP protocols are also well suited to measurement of metrics like throughput, forwarding rate, and packet loss," as David Newman, president of Network Test Inc., in Hoboken, N.J., told *Spectrum*.

Several vendors, such as Micromuse, Visual Networks, Netscout, Infovista, Sitara Networks, Netcom Systems, Lightspeed Systems, and CrossKeys Systems specialize in QoS monitoring, filtering, and reporting equipment.

Robert Mandeville, a founder and the chief executive officer at CQOS Inc., Irvine, Calif., told *Spectrum*: "Measurement of IP QoS will in future arbitrate the economics of business-grade IP services. IP QoS will be bought and sold only if it could be quantified and measured." CQOS is a start-up company dedicated to IP QoS measurement.

#### SLAS ARE KEY

The SLA is where the provider's technical competence, dedication to service, and business integrity stand revealed. To quote Joe Lardieri, managed router service offer manager, BellSouth Corp., Atlanta, Ga.: "Although the technical issues are considerable, one can never forget that SLAs are one part technical, one part contracting, and three parts negotiating. Carriers have a vested interest in minimizing their exposure to penalties; end-users have an equally

vested interest in maximizing it. In order for both sides to maximize their gains, a clear understanding of one's own business objectives is critical while drawing up a SLA."

#### A RIGHT TO KNOW

A well-crafted SLA should include the metrics of QoS and factors such as availability, maintenance scheduling, and mean time to repair. The client has the right to know about the time needed for network recovery after a power outage or equipment failure. The client also should be cognizant of the provider's ability to proactively detect and correct problems that may be looming.

Automatic generation of QoS reports, alarms, and trouble tickets, and issuance of credits for the vendor's noncompliance should be an integral part of an SLA. "[Yet] the deployment of QoS-based services is challenged by the ability to monitor and bill for such services," Azhar Sayeed, senior product manager for QoS at Cisco Systems, told *Spectrum*.

Other key questions are: what type of QoS reports should be generated? And how often should QoS metrics be taken and reported? If a service provider's report bases average availability on measurements taken over 24 hours, it may hide the problems that occur during the hours of peak usage.

Generating billing and credit records as per SLAs have not yet reached a high degree of automation. Take T-Data GmbH, Düsseldorf, which is a subsidiary of Deutsche Telekom, the German communications giant—it has signed SLAs with clients, but has yet to fully automate the billing, according to Axel Schenkel, its manager for the agreements.

Another challenge to delivering good QoS arises when a virtual private network crosses the administrative and technical domains of many providers, perhaps incumbent Baby Bells and competitive providers, who may not all adhere to the same transmission and QoS technologies. "Without uniform technologies and standards, service providers must negotiate individual operating agreements on a network-by-network basis. The associated cost can be prohibitive and limit service footprints, diluting the business promise of the Internet," said Lardieri of BellSouth.

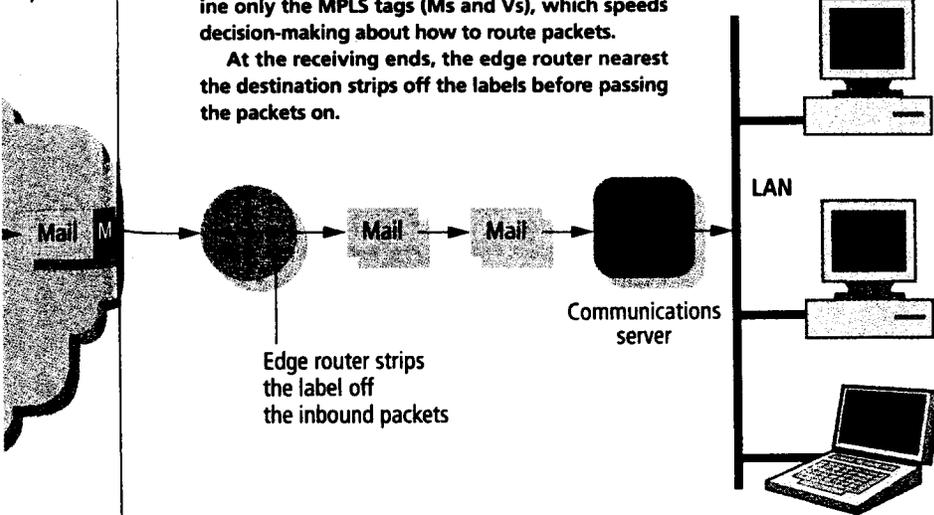
Furthermore, while designing a wide-area network for a high QoS, special attention has to be paid to the interfaces at aggregation points where there is a capacity mismatch between access links and network core links, according to Cisco System's Sayeed. Capacity mismatch occurs when a 100-Mb/s local-area network, for example, interfaces with a 1.544-Mb/s T1 wide-area network line.

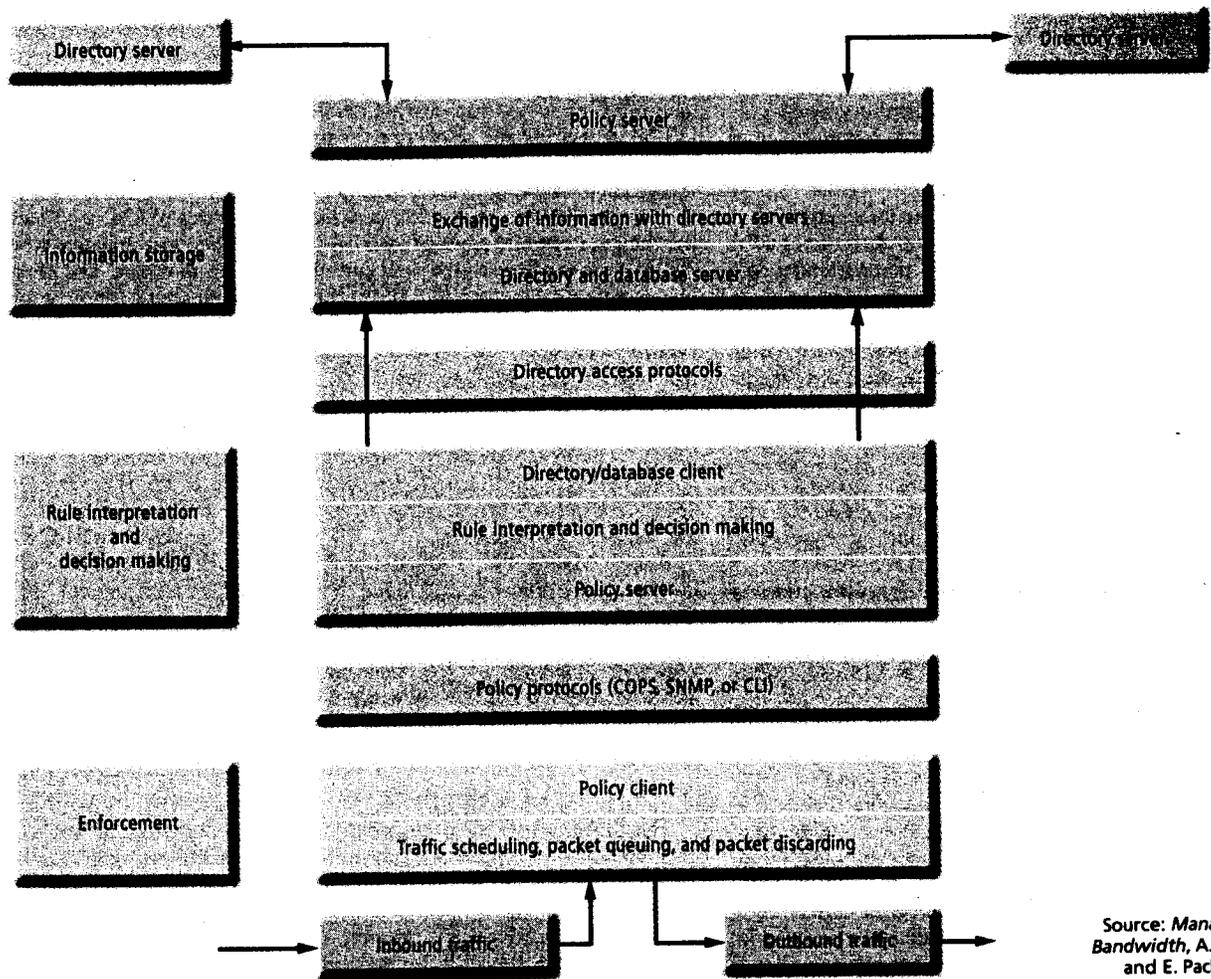
The more diverse and important a client's communications traffic becomes, the more crucial it is that the carrier main-

[2] Multiprotocol label-switching (MPLS) relies on an edge router at the network entry point [left] to label its packets with a 32-bit tag. The core switches and the routers inside the network then need to examine only the MPLS tags (Ms and Vs), which speeds decision-making about how to route packets.

At the receiving ends, the edge router nearest the destination strips off the labels before passing the packets on.

er strips off the packets





Source: *Managing Bandwidth*, A. Croll and E. Packman

[3] Adding advanced services is simplified by use of the common open policy service (COPS) protocol. The protocol makes use of a policy server [at top] that queries directory servers [beige] for the latest policy rules, which it keeps in its information storage section. A policy interpreter

makes decisions [blue] in compliance with the stored policies and the classification of the outbound data packets. The decisions are enforced [yellow] via the simple network management protocol (SNMP) and the command line interface (CLI) protocol.

tain a high QoS. Throughput, availability, packet loss, latency, and jitter must all be spelled out in SLAs, along with how each is to be measured and reported. (It is not uncommon for carriers to track QoS but not report the results to the client unless an extra fee is paid. Also, don't expect a carrier to generate credits automatically unless obliged to under the SLA.) The Latin expression *Caveat emptor* (let the buyer beware) may be old, but it remains sound advice in the world of modern telecommunications. ♦

### TO PROBE FURTHER

"Diversifying Internet delivery," *IEEE Spectrum*, November 1999, pp. 57-61, explains in detail the workings of the components of both integrated and differentiated service (IntServ and DiffServ), the two earliest standard procedures for improving quality of service (QoS).

The Internet Engineering Task Force (IETF), Reston, Va., formed in 1986, is a group of vol-

unteers who decide on the technical standards for the Internet. The task force invites papers on topics concerning the performance of the Internet and its protocol for discussion at meetings held three times a year. IP Performing Metrics (IPPM) and Benchmarking Methodologies Working Group (BMWG) are part of the task force. The IETF Web site at [www.ietf.org](http://www.ietf.org) yields a wealth of information.

Organizations involved in defining Internet metrics, including measurements of QoS parameters, have Web sites offering white papers and illustrative graphs. Check out: • Cooperative Association for Internet Data Analysis (Caida), San Diego, Calif.; [www.caida.org](http://www.caida.org). • Matrix Information Directory Services Inc. (MIDS), Austin, Texas, at [www.mids.org](http://www.mids.org). • IP Detail record (IPDR), McLean, Va., at [www.ipdr.org](http://www.ipdr.org). • Cross Country Working Team (XIWT), Reston, Va., at [www.xiwt.org](http://www.xiwt.org).

*QoS in Wide Area Networks*, by Uyles Black (Prentice Hall PTR, Upper Saddle River, N.J.,

2000), is a technical book that avoids much of the hype about quality of service issues and gets quite deeply into the technical details of the subject, down to the level of bits and bytes.

*Virtual Private Networks* by Bruce Perlmutter with Jonathan Zarkower (also from Prentice Hall PTR, 2000), gives an excellent treatment of virtual private networks (VPNs) and includes a section on QoS. Both authors write from first-hand experience with VPNs.

*Managing Bandwidth: Deploying QoS in Enterprise Networks*, by Alistair Croll and Eric Packman (again from Prentice Hall PTR, 2000), offers an easy-to-read pragmatic approach to the QoS problem. Both authors have practical experience in QoS and give real-life examples. The analogies given are good and the examples of numeric values help readers gain a firm grasp of the subject.

*Spectrum* editor: Michael J. Riezenman

Quantum transistors • Net quality of service • Issues in U.S. presidential race

SEPTEMBER 2000  
THE INSTITUTE  
OF ELECTRICAL  
AND ELECTRONICS  
ENGINEERS, INC.



# SPECTRUM

<http://www.spectrum.ieee.org>



## TRACKING WEATHER'S FLIGHT PATH

\$2.95



0 73361 64849 5

09