

# Traffic Problems?

## Making Way for Important Network Packets

Linda Kennard



In the immortal words of John Swigert during the *Apollo 13* flight to the moon, "Houston, we have a problem." Granted, the problem that we, or rather, you and other network administrators have doesn't threaten your lives—but it does threaten your network. The problem is you're running a variety of applications with disparate needs and sensitivities on multiservice networks that have limited bandwidth. In other words, you're not leaking critical oxygen into endless space; you're flooding finite space with critical applications.

Delay-sensitive, mission-critical traffic is struggling for space on the same pipe that is being devoured by bandwidth-intensive, nonessential network traffic. Consequently, when a few users fire up their RealAudio players to listen to clips from the *Apollo 13* movie soundtrack that they have downloaded from movietunes.com, all network traffic suffers. Although the users who are grooving to the tunes don't care about the network's waning performance, other users do. When database access and other business applications are drained of power, the users who are actually working complain.

Who takes the heat for this problem? You do. What are you going to do about the problem? You could simply add more bandwidth—a practice that vendors and industry analysts call *over-provisioning*. Although over-provisioning a LAN is arguably more cost-effective than managing the bandwidth you have, over-provisioning a WAN is not always a practical solution to network congestion. Adding bandwidth to a WAN is too costly to be widely accepted as practical. (For more information, see "More Versus Managed Bandwidth: A Poor Excuse for a Debate" on p. 8.)

Even if bandwidth were free, it isn't limitless. You have to continually add bandwidth to compensate for the law of diminishing space, which is this: The more space you have, the more space you use—whether that space is a closet, a house, or network

bandwidth. You can call this phenomenon whatever you like, but the point is that no matter how much bandwidth you have, "[you] keep coming up with new and creative ways to use" even more bandwidth, Gordon Smith, marketing vice president at Ukiah Software Inc., points out. "The golden future where bandwidth is limitless and free," Smith adds, "is still a good long way off."

### THE RIGHT (TRAFFIC-MANAGEMENT) STUFF

In the absence of this mythical era of limitless, free bandwidth, what can you do to improve the throughput and reliability of WAN traffic? You can deploy a router that touts quality of service (QoS) capabilities, or you can deploy a relatively new type of product called a *traffic shaper*. (See "Routers and Switches With Quality of Service Capabilities" on p. 14 and "Traffic Shapers" on p. 20.)

Although you undoubtedly know what routers are, you may not have heard about traffic shapers. Traffic shapers can be hardware-based products (with management software) or software-based products that serve primarily to control congestion by applying prioritization rules and allocating prespecified percentages of bandwidth to different types of traffic. To control congestion on a WAN, how do you decide which device to deploy—a traffic shaper or a router with QoS capabilities? Which device is better? And what does "better" mean?

### SO MANY CHOICES, SO LITTLE TIME

Assume for a moment that "better" refers to a device's ability to ensure a consistent level of service for various types of traffic from one end of a WAN to the other. In this case, neither device is better than the other because neither device can meet this criterion. End-to-end QoS to every office or person with whom you communicate—including those with whom you communicate over the Internet—is

## Inside This Article

*More Versus Managed Bandwidth* . . . . . p. 8

*Unfair and Partial Queuing: The Value of Queuing Algorithms* . . . . . p. 8

*All Things Are Relative: Which is Better—Queuing or TCP Window Sizing* . . . . . p. 10

*Queuing With a Splash of RED (and WRED): Techniques to Eliminate Packet Dropping* . . . . . p. 12

*Shape Up: The Value of TCP Window Sizing* . . . . . p. 14

*Routers and Switches With Quality of Service Capabilities* . . . . . p. 14

*Directory Services Simplify the Prioritization Process* . . . . . p. 16

*Discrimination—In a Good Way: How Devices Prioritize Traffic* . . . p. 18

*The New Frontier: The Future of Traffic Prioritization* . . . . . p. 19

*Traffic Shapers* . . . . . p. 20

## More Versus Managed Bandwidth: A Poor Excuse for a Debate

When a network is congested, is it better to add more bandwidth or to manage the existing bandwidth? The question of more versus managed bandwidth is addressed at length in an article titled "The Big vs. Managed Bandwidth Debate," by L. David Passmore, founder and research director of NetReference Inc., a network consulting firm. (You can download this article from <http://www.netreference.com>.)

Passmore believes that switch and router vendors have added exciting new service quality and control features, but he says network administrators shy away from providing multiple service quality levels because of the potential management complexity associated with doing so. "As a result," Passmore says, "there is an alternative view that it is much easier (and perhaps cheaper) to just 'throw bandwidth at the problem' via Gigabit Ethernet or other high-speed networking technologies."

### WHAT DEBATE?

After setting up the basis for his discussion, Passmore presents arguments to both sides of the "big bandwidth debate." The article is an excellent resource for understanding when you might need to add more bandwidth or to manage your existing bandwidth.

However, the premise of Passmore's article is slightly misleading. Although Passmore describes the issue through the eyes of supposed "proponents" and "advocates" of more bandwidth and managed bandwidth, we found no vendors that strongly supported one approach at the expense of the other approach. Instead, all of the vendors we spoke with believe that solving network congestion on a long-term basis requires both adding and managing bandwidth.

### WHEN TO ADD AND WHEN TO MANAGE

How do you know when to add more bandwidth and when to

manage the bandwidth you have? All of the vendors we spoke with agree with David Robbins, product line manager at Xedia Corp., when he says, "There is no simple answer to this question." Robbins adds that whether it is better to add more bandwidth or to manage what you have "depends on the needs of the organization."

If you want an answer that is a bit more concrete, most vendors will oblige you. Neil Gehani, senior product manager of the traffic control product line at Check Point Software Technologies Inc., says, "Where capacity is inexpensive, add capacity to avoid congestion. Where capacity is expensive and limited, manage, control, and increase efficiency to optimize and improve performance." In other words, if a LAN is congested, add more bandwidth; if a WAN is congested, manage the bandwidth you have.

Most vendors also recognize that every link has its limits and further agree that managing bandwidth within those limits will take you only so far. After all, bandwidth-management features, no matter how sophisticated, are not bandwidth-miraculous features. Not surprisingly then, most vendors agree with this succinct statement from Joel Feraud, product manager at Sun Microsystems Inc.: "A permanently saturated link simply must be upgraded."

Tim Szigeti, lab administrator at Cisco Systems Enterprise Management Business Unit, sums up the issue when he draws upon a common analogy: "Managing network traffic," Szigeti explains, "is like adding a commuter lane to a highway. Doing so makes more efficient use of the already available infrastructure. If you meet the [traffic] conditions [by adding a commuter lane], you get better commuting time. Eventually, however, you will need to build bigger highways." But Szigeti does not end there—nor would any vendor of products that provide Quality of Service (QoS) capabilities. If you build bigger highways with commuter lanes, Szigeti says, "all of your new highways will be more efficient." ●

not possible today. Barring that possibility, the best you can hope to do is to control traffic at common points of congestion, such as at the LAN/WAN border.

Is it better to use a traffic shaper or a router with QoS capabilities to control traffic at the LAN/WAN border? The strongest argument for routers is that you have to use a router anyway, so why not use one with traffic-prioritization capabilities? Unfortunately, you probably already have routers, and those routers most likely don't have traffic-prioritization capabilities.

Rather than waste the investment your company has already made by scrapping the routers you have and buying others, you may instead consider preserving that investment by buying a traffic shaper. Investment preservation represents the strongest argument for deploying a traffic shaper on a LAN/WAN border.

Beyond that, whether it is better to deploy a router or a traffic shaper is a stupid question or, at best, the wrong question. A

better question centers around the methods these devices use to control traffic.

That is, all of these devices aim to control the rate of traffic, particularly TCP traffic. To control traffic, all devices use one of two methods: queuing or TCP window sizing. The better question is this: To control congestion, is it better to use queuing, as provided by routers and a few traffic shapers, or TCP window sizing, as provided by most traffic shapers? As with all questions posed in this industry, no one can give you a straight answer. However, you can learn more about each method for controlling congestion and make an informed choice based on the needs of your company.

### UNFAIR AND PARTIAL QUEUING

Queuing algorithms that traffic-prioritization devices use to provide different levels of service to different traffic types are unfair and partial—as they should be. A fair method for forwarding traffic would be one that views all traffic as equal. But

all traffic is not equal. What you need is a method that recognizes—and responds appropriately to—different traffic types, giving priority to the most critical traffic.

The standard first-in-first-out (FIFO) queue that even the most basic routers use (not to mention banks and drive-through windows) treats all traffic equally. Treating traffic equally is a QoS nightmare. In FIFO queuing, a device has one queue per port and simply forwards packets when they reach the front of the line. As a result, Structured Query Language (SQL) packets can get stuck behind long lines of Network News-Transfer Protocol (NNTP) packets, and when the FIFO queue's buffer fills up, SQL packets may be dropped.

Fortunately, the specialized queuing algorithms that routers, switches, and traffic shapers use to prioritize traffic exercise obvious biases for traffic you designate as high priority. The following are the most common queuing algorithms.

## All Things Are Relative

Which is better—queuing or TCP window sizing? Ask one dozen vendors, and you'll get one dozen answers. Ask one dozen network consultants, and you'll get one dozen more answers. The common arguments for and against both queuing and TCP window sizing are presented below. By reading these arguments, you will be equipped to decide the answer to that question for yourself.

### QUEUING IN QUESTION

Generally, TCP window sizing disciples agree with Jeff Barker, senior product manager at Packeteer Inc., who claims that "TCP window sizing is a more proactive approach" to managing TCP traffic than queuing. According to TCP window sizing proponents, including Packeteer and Ukiah Software Inc., even the most sophisticated queuing algorithms do nothing to correct the source of the congestion they're reacting to—unless you consider dropping packets a corrective practice. Queuing, Barker points out, relies on "packet loss to infer congestion, prompting sending servers to slow down only after you've lost packets."

Although TCP window sizing proponents use packet-dropping to denigrate queuing as a method for controlling traffic, packet-dropping is not quite as dire as you may think. For one thing, not all queuing algorithms rely on packet-dropping to manage TCP traffic.

For example, devices that use Class-Based Queuing (CBQ) do not have to drop TCP packets to make TCP slow down its sending rate, according to Joel Feraud, product manager at Sun Microsystems Inc.: "What usually happens with CBQ," Feraud claims, "is that TCP packets are delayed and the round-trip time increases." This increase in round-trip time prompts the TCP sender to slow down. Generally, Feraud says, TCP slows down before the queue is saturated, and consequently, no packets are dropped.

Router vendors and traffic shaper vendors that rely primarily on queuing methods other than CBQ to control traffic are fully aware of the potential problem of packet-dropping. Consequently, these vendors use other methods to minimize or eliminate packet-dropping.

For example, Check Point Software Technologies Inc. uses proprietary capabilities along with Weighted Fair Queuing (WFQ) and claims that, as a result, FloodGate-1 does not drop TCP packets. In addition, many router vendors such as Cabletron Systems Inc. and Cisco Systems Inc. use Random Early Discard (RED) and Weighted Random Early Discard (WRED) to minimize packet-dropping. (For more information about RED and WRED, see "Queuing With a Splash of RED (and WRED)" on p. 12.)

TCP window sizing proponents also commonly cite a limited number of static queues (generally fewer than ten) as an inherent disadvantage to queuing. This alleged disadvantage, however, does not apply to CBQ, used by Xedia Corp.'s Access Point, Sun Microsystems' Solaris Bandwidth Manager, and IPHighway's QoS Master solutions, nor to the proprietary queuing algorithm that Check Point's FloodGate-1 uses.

For example, Nortel Networks' routers (and switches) support up to eight priority levels mapped into two to eight queues. In contrast, Xedia claims that Access Point can enforce explicit bandwidth control for hundreds of queues, and Check Point maintains that its FloodGate-1 supports an unlimited number of "virtual" queues.

### RATE COMPLAINTS

Queuing disciples, in contrast, commonly point out that TCP window sizing intervenes with TCP sessions. TCP window sizing, says David Robbins, product line manager at Xedia, "is using TCP in a way that is not defined by the IETF [Internet Engineering Task Force]." TCP window sizing, Xedia states in an FAQ posted on the company's web site, is "a proprietary bandwidth management mechanism that requires explicit protocol intervention, analogous to protocol spoofing." (You can download this FAQ from [http://www.xedia.com/products/cbq\\_faq.htm](http://www.xedia.com/products/cbq_faq.htm).)

Words like proprietary and spoofing send shock waves through the collective networking industry. Look past the words, however, and consider the real point. Robbins explains this point best when he says that with queuing, "TCP will rate control itself," which he believes is a better approach to rate control than "playing with the [TCP] window mechanisms and forcing" TCP to slow down.

Vendors of queuing devices also focus on the fact that TCP window sizing controls only TCP traffic. Devices that use TCP window sizing must use another method to control all other traffic, including User Datagram Protocol (UDP) traffic such as voice over IP (VoIP) traffic. Queuing proponents claim that devices that use TCP window sizing use crude queuing algorithms for all other traffic. Barker flatly disputes that accusation, saying that Packeteer understands "the importance of UDP traffic and has worked hard to develop sophisticated capabilities to address that traffic."

Using words like crude to describe the queuing algorithms that TCP window sizing devices use to handle non-TCP traffic undermines the value of these queuing algorithms and arguably distorts the truth. The truth is that TCP window sizing devices generally use strict Priority Queuing or a leaky-bucket algorithm to support other types of traffic. (For more information about Priority Queuing, see the "Get Your Priorities Straight" section on p. 12.)

A leaky bucket algorithm is a buffering algorithm that delays packets when necessary to slow the data rate to a rate that you prespecify for a particular traffic flow. In this respect, the router or buffering device is like a bucket with a leak in the bottom. That is, water can trickle or burst suddenly into a bucket, but the leak or drip out the bottom of that bucket remains constant. Similarly, packets can trickle or burst suddenly into a queue, but the algorithm that the queue uses ensures that those packets leave that queue at a constant rate. If the queue (like the bucket) becomes too full, packets (like water) drop over the edge.

TCP window sizing proponents also claim that queuing solutions (by which they probably mean router-based queuing) cannot adequately handle inbound traffic. Barker states matter-of-factly that "queuing-based solutions cannot directly control inbound traffic." Controlling inbound traffic, Barker adds, "requires signaling the sender to slow down," which TCP rate control, not queuing, can do, according to Barker.

Interestingly, queuing advocates and TCP window sizing advocates accuse each other of being unable to adequately handle inbound traffic. "TCP window sizing," Check Point claims, "cannot manage inbound traffic because there is no inbound control since the decision making process for window sizing is located too far from the sender." Of course, Check Point claims its approach is entirely unique and therefore represents neither a queuing nor a TCP window sizing advocate. Check Point's complaint against TCP window sizing, however, is not unique. Queuing proponents state similar complaints. ●

- Priority Queuing
- Weighted Fair Queuing (WFQ)
- Class-Based Queuing (CBQ)

### Get Your Priorities Straight

The Priority Queuing algorithm creates multiple queues per port and assigns each queue a relative priority value. Devices that use Priority Queuing decide which queue to place traffic in by checking preconfigured traffic-prioritization rules. (For more information about traffic-prioritization rules, see “Discrimination—In a Good Way” on pp. 18–19.) Priority Queuing then transmits traffic in high-priority queues before transmitting packets in lower-priority queues.

Priority Queuing is obviously better at ensuring different levels of service for different types of traffic than FIFO queuing is. However, Priority Queuing gives the queue being serviced all available bandwidth. The potential result is that when used alone, Priority Queuing can starve lower-priority traffic of bandwidth by always servicing a high-priority queue despite traffic piling up in a low-priority queue.

Not surprisingly, all of the vendors that use Priority Queuing compensate for this inherent Achilles’ heel by using other methods, such as TCP window sizing or WFQ. Of the products mentioned in this article, routers and switches from 3Com Corp., Cabletron Systems Inc., Lucent Technologies, and Nortel Networks Corp. support Priority Queuing, as do traffic shapers from Netscreen Technology Inc., Packeteer Inc., and Ukiyah Software Inc. (See “Routers and Switches With Quality of Service Capabilities” on p. 14 and “Traffic Shapers” on p. 20.)

### Queuing With a Splash of RED (and WRED)

TCP controls its own transmission rate by slowing down when it notices packets are being dropped. To minimize packet-dropping, many router vendors use Random Early Discard (RED) and Weighted Random Early Discard (WRED) techniques.

Unlike TCP, which slows traffic only after the network is congested, RED and WRED slow traffic before the network is congested. RED and WRED monitor queues and, when the network reaches a specified traffic threshold, randomly discard packets to slow TCP traffic before the queues they are monitoring reach capacity. By discarding

### Queuing With Some Weight to It

Like Priority Queuing, WFQ creates multiple queues for different traffic classes and assigns each queue a relative priority value. Devices that support WFQ (as with devices that support Priority Queuing) decide which queue to place traffic in by checking preconfigured traffic-prioritization rules. But the similarities between Priority Queuing and WFQ end there.

### *In contrast to Priority*

### *Queuing and WFQ, CBQ*

*enables you to allocate a guaranteed bandwidth rate to each traffic class.*

With WFQ, you assign a weight value to each queue in proportion to its level of priority. You weight queues to ensure that higher priority queues get a larger percentage of available bandwidth than lower priority queues get. The exact amount of bandwidth each queue actually gets depends on the number of queues sharing that bandwidth at a given moment.

For example, suppose that you have eight queues. You assign three queues a weight value that ensures that when the queues are busy, they get the largest amount of bandwidth. You then assign the remaining five queues lower but equal

packets before the network is congested, RED and WRED techniques slow TCP traffic and minimize the risk of larger numbers of packets being delivered to already full queues. Delivering packets to full queues could result in equally large numbers of packets being dropped.

WRED is a little less random than RED in its selection of packets. WRED discards packets in order of their drop preference, as marked in the IP Precedence portion of the IP Type of Service (TOS) field. (For an explanation of IP Precedence, see “Type of Service” in the Glossary on the NetWare Connection web site at <http://www.nwconnection.com>.)

values, ensuring they always get some bandwidth, but a smaller amount than high-priority queues get.

When one of the high-priority queues and two of the other queues are busy, the high-priority queue may get 50 percent of the available bandwidth, and the other queues may get 25 percent each. When the high-priority queues are all busy at the same time, each high-priority queue gets only a little more than 30 percent of the available bandwidth. In other words, although WFQ ensures that traffic classes always get some bandwidth, the specific rate is variable—not guaranteed.

Of the products mentioned in this article, routers from Lucent Technologies, Cisco Systems Inc., and Nortel Networks support WFQ, as do routers and switches from Cabletron Systems. (See “Routers and Switches With Quality of Service Capabilities” on p. 14.) NetGuard Inc.’s traffic shaper also supports WFQ. Check Point Software Technologies Inc.’s traffic shaper uses WFQ with other proprietary capabilities to guarantee bandwidth rates and limits. (See “Traffic Shapers” on p. 20.)

### Queuing With Class

CBQ is an open packet-scheduling algorithm that is arguably the most sophisticated of the queuing algorithms vendors commonly use today. In contrast to Priority Queuing and WFQ, CBQ enables you to allocate a guaranteed bandwidth rate to each traffic class.

Like Priority Queuing and WFQ, CBQ creates multiple queues for different traffic classes and decides which queue to place traffic in by checking preconfigured traffic-prioritization rules. Unlike Priority Queuing and WFQ, however, CBQ enables you to assign traffic classes guaranteed data rates. For example, if you allocate 56 kbit/s to a high-priority traffic class, CBQ ensures that that traffic class always gets 56 kbit/s.

With CBQ you can also define parameters that enable devices to distribute additional bandwidth to traffic classes as bandwidth is needed. In short, with CBQ you can ensure that traffic classes always get their guaranteed bandwidth rates. You can also ensure that traffic classes can burst above those guaranteed rates when necessary, depending on how you have configured the CBQ device.

Of the products mentioned in this article, traffic shapers from IPHighway and Sun Microsystems Inc. support CBQ, as do routers from Xedia Corp.

## Routers and Switches With Quality of Service Capabilities

The following are a handful of vendors that offer routers and switches with Quality of Service (QoS) capabilities, including the traffic-prioritization capabilities discussed in this article. As shown below, a growing number of vendors are recognizing the importance (and certain future) of directory-enabled networking. Accordingly, many vendors are integrating their routers and switches with Novell Directory Services (NDS) and other Lightweight Directory Access Protocol (LDAP) 3-compliant directories.

Vendor	Product	Management Software	LDAP 3-Compliant?	NDS-Enabled?	Prioritization Methods	Queuing Methods
3Com Corp. http://www.3com.com 1-408-727-7021 1-800-638-3266	Core Builder 9000 Enterprise Switch Core Builder 9400 Gigabit Ethernet Switch PathBuilder S310, S330, and S700 WAN switches SuperStack II switches	Transcend Enterprise Manager	Yes	No	Rules-based VLAN prioritization	Priority IP TOS
Cabletron Systems Inc. http://www.cabletron.com 1-603-332-9400	SmartSwitch 2000/8000/8600 routers	SPECTRUM Enterprise Manager	Yes	Yes	Rules-based VLAN prioritization IP TOS DiffServ	Priority WFQ CBQ
	SmartSwitch 2000/6000/9000 switches	SPECTRUM Enterprise Manager	Yes	Yes	Rules-based VLAN prioritization IP TOS DiffServ	Priority WFQ
Cisco Systems Inc. http://www.cisco.com 1-408-526-4000 1-800-326-1941	Cisco routers, including 2500, 3600, 4000, 4500, 4700, 7200, and 7500 Cisco 5000 and 6000 family switches	Complete QoS Policy Manager Cisco IOS (runs on router hardware)	Yes	Yes	Rules-based VLAN prioritization IP TOS DiffServ (all routers and soon in 5000 and 6000 switches)	Custom Queuing WFQ
Lucent Technologies http://www.lucent.com 1-978-318-6300 1-800-237-0016	Cajun P550 Gigabit Switch Cajun P550 Routing Switch	RealNet Rules	Yes	Yes	Rules-based	Priority Custom-designed
Northern Telecom Ltd. (Nortel Networks Corp.) http://www.nortel.com 1-506-674-5471 1-800-466-7835	Accelar 1000 routing switches Accelar 8000 enterprise switches BayRS routers	Optivity Network Management Software	Yes	Yes	Rules-based (switches/routers) VLAN prioritization IP TOS (routers) DiffServ (routers)	Priority WFQ
Xedia Corp. http://www.xedia.com 1-978-952-6000	Access Point 10 Access Point 100	Access View	No	No	Rules-based IP TOS DiffServ	CBQ

(See "Routers and Switches With Quality of Service Capabilities" above and "Traffic Shapers" on p. 20.)

### SHAPE UP!

Although some traffic shapers use CBQ to control congestion, most traffic shapers, including those from Allot Communications Inc., Elron Software Inc., NetGuard, Packeteer, and Ukiyah Software, use a method called *TCP window sizing*. To understand TCP window sizing, you must understand a few things about TCP: Devices communicating via TCP send acknowl-

edgment (ACK) packets to indicate that messages have been received. These ACK packets also contain the *TCP Receiver Window Size*, a value that indicates to the sending station how much data can be sent during the next packet transfer.

The receiving server (or an intercepting device) can enlarge or reduce the window size at any point during the communication exchange. A sending server waits for an ACK packet from the receiving server before sending the next series of packets. When that ACK packet arrives, the sending server sends as many packets

as possible to fill up the available window. Because TCP aggressively fills up the window, TCP often creates sudden bursts of large chunks of traffic. As you know, bursty, chunky traffic is not a good thing.

Devices that use queuing to manage bandwidth respond to TCP traffic as they respond to all traffic: They place the traffic in a queue. The TCP traffic then awaits its turn like every other traffic flow. When bursts of TCP traffic hit queuing buffers that are already full, packets get dropped, and this packet-dropping signals TCP sending devices to slow down.

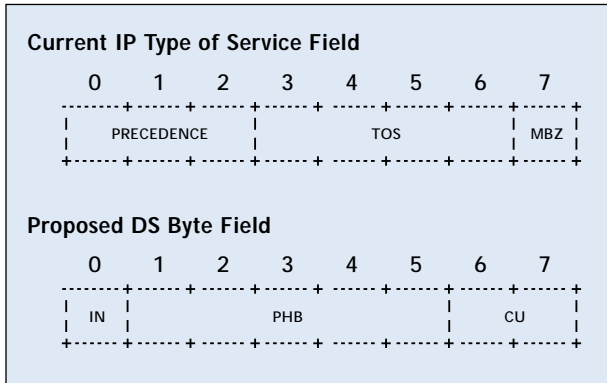


Figure 1. The DiffServ standard changes the name of the Type of Service (TOS) field to the DS Byte field.

Waiting until TCP traffic overflows queues to signal sending servers to slow down is a potential problem. If all TCP streams simultaneously slow down and then simultaneously ramp up again, the resulting cycle wreaks havoc on network performance: Queuing devices drop packets, leading servers to slow down. Later, servers ramp up, leading queuing devices to again drop packets.

Devices that use TCP window sizing intercept TCP ACK packets and adjust the window size when necessary to keep TCP traffic flowing. Rather than allowing TCP to function as it would normally, devices that use TCP window sizing intervene and reset the TCP window size to tell the sending server to slow down—before it sends a sudden burst of traffic. TCP window sizing thus combats the traffic bursts that are the defining characteristic of TCP. (For more information about TCP window sizing, visit <http://www.packeteer.com/technology/tcp.htm>.)

Of course, vendors offer different types of TCP window sizing. In fact, Packeteer, a company that produces bandwidth-management devices, called into question our use of the phrase *TCP window sizing* rather than *TCP rate control* in reference to the method its PacketShaper products use to control traffic. Using the phrase *TCP window sizing*, Jeff Barker, senior product manager at Packeteer, suggests, “is a way to trivialize the TCP rate control method that Packeteer invented.” TCP rate control, Barker adds, “involves more than simply adjusting the TCP window size.”

Chris Belthoff, Elron Software product manager, agrees. Elron’s CommandView Bandwidth Optimizer also uses a custom version of TCP rate control that Elron calls *Dynamic Traffic Control*. As Belthoff explains, products that use TCP rate con-

trol, rather than what Belthoff and Barker both refer to as *simple window sizing*, address issues such as when to apply rate control and by how much and for how long to adjust the TCP window size. “Applying TCP rate control inappropriately,” Belthoff adds, presumably implying the trouble TCP window sizing alone can cause, “can actually increase retransmits.”

### ALL THINGS CONSIDERED

Vendors of products that use queuing or TCP window sizing (or rate control as the case may be) to control traffic could probably spend days debating the advantages and disadvantages of one versus the other. (For more information about the pros and cons of both queuing and TCP window sizing, see “All Things Are Relative” on p. 10.) The bottom line is that neither queuing nor TCP window sizing is inherently better or worse but simply different. You should learn about the specific methods that individual traffic-prioritization devices

*“Traffic identification and classification is crucial,” Barker explains. “If you can’t identify the traffic, you certainly can’t control it.”*

take. Then you can purchase a traffic-prioritization device based on what is best suited for your company’s goals and network environment.

Of course, you should base your decision on more than just whether that device uses queuing or rate control. For example, you should learn how the device classifies and prioritizes traffic—and at what level the device classifies and prioritizes traffic. When a packet hits a device that provides traffic-shaping capabilities, the first thing that device needs to determine is what type of traffic this packet is flowing with and what level of priority this

type of traffic gets. Determining what type of traffic the packet is flowing with can mean determining anything from what user or application sent the packet, to which host sent or will receive the packet.

The specifics make a big difference. “Traffic identification and classification is crucial,” Barker explains. “If you can’t identify the traffic, you certainly can’t control it.” (For more information about classification and prioritization methods, see “Discrimination—In a Good Way” on pp. 18–19.)

Most traffic-prioritization devices support classification and prioritization standards, such as 802.1p (Ethernet prioritization), IP Precedence (TOS), and DiffServ. However, using these standards and even adding classification and prioritization based on IP addresses or port numbers may not be enough. According to Barker, you need a device that can “look into application content, allow dynamically negotiable port assignments, and look at URLs.” Vendors describe this level of prioritization as *highly granular prioritization*.

When you’re in the market for a traffic-prioritization device, you should also get answers to the following questions:

- Can the device manage both inbound and outbound traffic?
- Can the device analyze traffic to determine what sorts of traffic-prioritization rules might benefit the network most?
- Can the device monitor traffic after you’ve created rules to determine whether or not the rules were effective?
- Does the device integrate with a Virtual Private Network (VPN) solution? (Can the device handle encrypted traffic?)
- Does the device have an easy-to-use interface?
- Does the device support Lightweight Directory Access Protocol (LDAP)?

### DIRECTORY SERVICES SIMPLIFY THE PRIORITIZATION PROCESS

An increasing number of routers and traffic shapers support LDAP 3. Support for LDAP implies an ability to integrate with LDAP-compliant directories, such as Novell Directory Services (NDS). Directory integration implies an ability for the device to access an enterprise-wide directory to store traffic-prioritization policies that could prioritize traffic based on User and Group objects. Support for LDAP also implies the potential to create a traffic-prioritization rule one time only and store

## Discrimination—In a Good Way

When a packet hits a traffic-prioritization device that provides traffic-shaping capabilities, the device must first determine what type of traffic this packet is flowing with and what level of priority this traffic type gets. Determining what traffic type this packet is flowing with can mean determining anything from what user or application sent the packet, to which host sent the packet or will receive it.

The specific information that devices check to determine traffic type depends both on how you configure that device and on the traffic-shaping rules you create. Routers and traffic shapers usually determine relative priority based on the following information:

- A priority value specified in some Ethernet frames.
- Priority values specified in what is usually called the Type of Service (TOS) field in the IPv4 packet header.
- Other IP packet header information that you configure the device to check. (This method for determining priority is called a rules-based approach in “Routers and Switches With Quality of Service Capabilities” on p. 14 and “Traffic Shapers” on p. 20.)

### SETTING PRIORITIES

Many routers, but very few traffic shapers, comply with the Institute of Electrical and Electronic Engineers (IEEE) 802.1p specification. This specification enables devices to detect or mark a packet priority level using the priority value in 802.1Q-compliant Ethernet frames. (For more information, see “802.1p Specification” in the Glossary on the NetWare Connection web site at <http://www.nwconnection.com>.) This priority value can indicate one to eight possible priority levels, ranging from the lowest priority (which is interpreted as “apply best-effort service only”) to the highest priority (which is interpreted as “allocate reserved bandwidth”).

Although only some routers and traffic shapers use the priority value in the Ethernet frame, all routers and traffic shapers can detect (and some can mark) a packet’s priority level using information in the IP packet header’s TOS field. What this means specifically, however, varies widely.

The TOS field is actually comprised of three subfields: Precedence, TOS, and Must Be Zero (MBZ). (MBZ is not being used to indicate a packet’s priority level.) (For more information, see “Type of Service” in the Glossary on the NetWare Connection web site at <http://www.nwconnection.com>.) The designers of IPv4 intended devices to use the three-bit Precedence subfield to indicate a packet’s priority level. In addition, the IPv4 designers intended devices to use the four-bit TOS subfield to indicate the tradeoffs devices should make between throughput, delay, reliability, and cost to provide an appropriate level of service.

The trouble is although many devices can interpret (or mark) the values in the TOS field, there is no standard for interpreting (or marking) the TOS field. Members of the Internet Engineering Task Force (IETF) Differential Services (DiffServ) working group have suggested various ways to interpret Precedence and TOS values. For example, Paul Ferguson of Cisco Systems Inc. suggested using the Precedence subfield to indicate a drop preference and the TOS subfield to indicate delay requirements. (For more information, visit <http://diffserv.lcs.mit.edu/Drafts/draft-ferguson-delay-drop-00.txt>.)

In Ferguson’s proposal, devices would interpret packets marked with the lowest priority value (000) in the Precedence subfield as having the highest drop preference, and devices would interpret packets marked with the highest priority value (111) as having the lowest drop preference. In other words, low-priority packets get dropped first. Devices using Ferguson’s suggested semantics would interpret values in the TOS field as indicating a particular packet’s delay sensitivities, ranging from no delay sensitivity (0000) to highest delay sensitivity (1000).

### WHAT’S THE DIFF?

Because no standard for interpreting IP TOS and Precedence exists, devices that claim to support IP TOS or IP Precedence are not necessarily interoperable. The good news is that a new standard is in the works. The IETF DiffServ working group has a proposed standard called DiffServ. DiffServ changes the name of the TOS field to the DS Byte field and restructures that field. (See Figure 1 on p. 16.) The restructured field has the following subfields:

—continued on next page

it in an LDAP-compliant directory. All of your company’s LDAP-compliant traffic-shaping devices can then access that rule.

Of the six router vendors mentioned in this article, five router vendors have made their products LDAP-compliant. (See “Routers and Switches With Quality of Service Capabilities” on p. 14.) In addition, four router vendors either have integrated or will be integrating their products with NDS. The following routers and switches are integrated with NDS:

- Cabletron Systems’ SmartSwitch 2000/8000/8600 routers and SmartSwitch 2000/6000/9000 switches
- Cisco Systems’ 2500/3600/4000/4500/4700/7200/7500 routers and 5000/6000 switches

- Lucent Technologies’ Cajun P550 switch
- Nortel Networks’ Accelar 1000/8000 switches and BayRS routers

These vendors integrate, or will integrate, their routers or switches with NDS through their policy-management software. This integration enables you to create traffic-prioritization rules based on a User or Group object.

For example, routers and switches understand only IP addresses—they don’t understand NDS User or Group object names. Policy-management software for a particular router reads a user’s IP address from the NDS User object and sends that IP address to the router. The IP addresses themselves must be regularly updated by a

Dynamic Host Configuration Protocol (DHCP) service, such as Novell’s Domain Naming System (DNS)/DHCP service.

Prioritizing traffic based on users’ names is more useful—and more convenient—than prioritizing traffic by IP address. By prioritizing traffic based on a User object name, you ensure that the user gets the same priority level regardless of where that user logs in—whether from the accounting department, the help desk, the lab, or his or her own desk. The traffic-prioritization rules you create for users will follow those users wherever they go.

Router vendors are not alone in their pursuit to directory-enable their products. Of the nine traffic shapers mentioned in this article, six support LDAP 3. (See “Traffic Shapers” on p. 20.) In addition,

- **CU.** The 2-bit CU is "currently unused."
- **PHB.** The 5-bit PHB field marks the per-hop behavior (PHB) a particular packet requires. For example, 11100 means Expedited Forwarding (EF). Devices that support DiffServ place EF packets in short queues and service them quickly to ensure low latency and minimal packet loss and jitter.
- **IN.** The 1-bit IN field indicates whether the packet is in- or out-of-profile with respect to traffic policies at a network boundary.

Devices that support DiffServ are likely to be interoperable with other devices that support DiffServ, the implication of which is promising. For example, suppose a carrier uses routers and switches that support DiffServ, as several routers and switches now do. Also suppose that your company uses a traffic shaper that supports DiffServ, as all of them now do. In such a case, your traffic shaper could mark the DS Byte field, and your provider's routers could read that value. This traffic prioritization means you can have an end-to-end Quality of Service (QoS) solution that would ensure a consistent level of service for different traffic types over any IP network—including the Internet.

### RULES TO PRIORITIZE BY

In addition to reading priority values in Ethernet frames or in the IP TOS or DS Byte fields, devices can be preconfigured to automatically detect and prioritize certain traffic types. For example, Packeteer Inc. preconfigures PacketShaper to detect and prioritize more than 150 traffic types based on information in layers two through seven of the Open Systems Intercommunication (OSI) model.

Of course, Packeteer automatically prioritizes these traffic types only after you manually select policies from a list of Packeteer's suggested policies. Packeteer, like most router vendors and traffic-shaper vendors, believes that you know best which traffic you want to prioritize. Allowing devices "to automatically detect and prioritize traffic without any administrative action whatsoever," says Erin Curtis, senior public relations manager for Nortel Networks Corp., "may result in out-of-control QoS services."

All routers and traffic shapers provide prioritization without human intervention after you create rules that indicate where those

devices should look to detect traffic types and after you assign priority levels to those traffic types. The specific information you use to create these rules varies from device to device. However, all devices enable you to create rules to detect and prioritize traffic based on protocol type (for example, TCP or UDP), port number, and source and destination address. For example, Cisco cites the following rule as one of its favorite examples of prioritization rules that you can create for Cisco routers using Cisco's Complete QoS Policy Manager software:

```
if (protocol is TCP AND source port is 9000) then priority = HIGH
```

Tim Szigeti, lab administrator in Cisco's Enterprise Management Business Unit, explains that this rule instructs the router to protect Oracle enterprise resource planning (ERP) traffic, which typically uses TCP port number 9000. High-priority queuing, Szigeti adds, is quite extreme "and probably would not be used in an enterprise environment." Nevertheless, Cisco uses the rule in demos frequently because "it is the most dramatic," Szigeti says.

Although Cisco's example may be dramatic, Sun Microsystems Inc. does not consider it an example of a policy. Joel Feraud, product manager at Sun Microsystems, claims this example demonstrates a configuration filter, not a prioritization policy. Feraud offers the following as an example of a policy:

```
if (user is engineer) then class_of_service is gold
```

In Feraud's example, the policy includes a reference to an object in Sun Directory Services, Sun Microsystems' Lightweight Directory Access Protocol (LDAP) compliant directory where the policy would be stored. Solaris Bandwidth Manager on the Sun network then enforces this policy whenever an engineer connects to the network.

Before purchasing a traffic-prioritization device, make sure that device can detect and prioritize at least as many traffic types as you have running on your company's network. You should also learn how to configure the device to ensure that the configuration process is as simple as you expect it to be. ●

the five traffic shapers listed below are integrated with NDS.

- Allot Communications' Systems Release 2.0
- Check Point's FloodGate-1 1.5
- NetGuard's GuidePost Bandwidth Manager for NT
- Packeteer's PacketShaper 1000/2000/4000
- Uki Software's NetRoad TrafficWARE

Furthermore, one other vendor intends to integrate its traffic shaper with NDS. According to Belthoff, Elron Software plans to integrate future versions of CommandView Bandwidth Optimizer with LDAP 3-compliant directories. NDS, Belthoff states, is "at the top of Elron's list."

A second vendor suggests that, theoretically at least, its traffic shaper already can integrate with NDS. Sun Microsystems' Solaris Bandwidth Manager 1.5 supports LDAP 3 but specifically integrates only with Sun Directory Services 3.1. However, according to Joel Feraud, product manager at Sun Microsystems, the "policy-schema used by Solaris Bandwidth Manager to abstract its configuration . . . can be stored in any LDAP 3-compliant directory server supporting CRAM-MD5 authentication of login," including NDS.

### THE NEW FRONTIER

Not everyone believes LDAP support represents a significant benefit. For example, when Dave Logan, senior consultant at Acuitive Inc., is reminded of the po-

tential positive effects of deploying LDAP-compliant directories, he replies, "That's a nice story, but LDAP hasn't been widely deployed."

Logan's claim is highly questionable. After all, 80 percent of all Fortune 500 companies have deployed NDS—a figure that clearly suggests widespread deployment among what are arguably trend-setting customers.

Regardless of the exact number of companies that have deployed enterprise-wide directories, vendors clearly believe that LDAP will play an important role in the networking future. For example, Novell and Lucent Technologies recently teamed up to create open directory-enabled networking standards, as have Microsoft and Cisco. All four companies participate in

## Traffic Shapers

The following are many (but not all) of the vendors that offer traffic shapers (also called bandwidth management products or traffic tuners). A growing number of vendors of traffic shapers are recognizing the importance (and certain future) of directory-enabled networking. Accordingly, many of these vendors are integrating their products with Novell Directory Services (NDS) and other Lightweight Directory Access Protocol (LDAP) 3-compliant directories.

Vendor	Product	Hardware/ Software	LDAP 3- Compliant?	NDS- Enabled?	Prioritization Methods	Queuing Methods	TCP Window Sizing	Platforms
Allot Communications Inc. <a href="http://www.allot.com">http://www.allot.com</a> 1-408-399-3154	Allot Communications Systems Release 2.0	HW w/ mgmt. SW	Yes	Yes	Rules-based IP TOS DiffServ	Custom- designed	Yes	
Check Point Software Technologies Ltd. <a href="http://www.checkpoint.com">http://www.checkpoint.com</a> 1-650-628-2000	FloodGate-1 1.5	SW-only	Yes	Yes	Rules-based DiffServ	WFQ with hierarchical, multipath capabilities	No	Solaris 2.5, 2.6 Windows NT
Elron Software Inc. <a href="http://www.elronsw.com">http://www.elronsw.com</a> 1-617-914-5000 1-800-767-6683	CommandView Bandwidth Opti- mizer 2.1	SW-only	Plans to	Plans to	Rules-based	Custom- designed	Yes*	Windows NT
IPHighway <a href="http://www.iphighway.com">http://www.iphighway.com</a> 1-201-585-0800 1-800-964-6965	QoSMaster	SW-only	No	No	Rules-based IP TOS DiffServ	CBQ	No	Windows NT
NetGuard Ltd. <a href="http://www.ntguard.com">http://www.ntguard.com</a> 1-972-738-6900 1-800-533-8549	GuidePost Bandwidth Manager for NT	SW-only	Yes	Yes	Rules-based IP TOS	WFQ	Yes	Windows NT
Netscreen Technologies Inc. <a href="http://www.netscreen.com">http://www.netscreen.com</a> 1-408-330-7800 1-800-638-8296	Netscreen 10 Netscreen 100 Netscreen 1000	HW w/ mgmt. SW	Plans to	No	Rules-based	Priority	Yes	
Packeteer Inc. <a href="http://www.packeteer.com">http://www.packeteer.com</a> 1-408-873-4400 1-800-697-2253	PacketShaper 1000 PacketShaper 2000 PacketShaper 4000	HW w/ mgmt. SW	Yes	Yes	Rules-based VLAN priori- tization IP TOS DiffServ	Priority	Yes**	Windows NT NetWare 5
Sun Microsystems Inc. <a href="http://www.sun.com">http://www.sun.com</a> 1-888-843-5282 outside U.S. and Canada, see web site	Solaris Bandwidth Manager 1.5	SW-only	Yes	No	Rules-based IP TOS DiffServ	CBQ	No	Solaris 2.6, 2.7
Ukiah Software Inc. <a href="http://www.ukiahsoft.com">http://www.ukiahsoft.com</a> 1-408-369-2890 1-800-988-5424	TrafficWARE Active Policy System (APS)	SW-only	Yes	Yes	Rules-based IP TOS DiffServ	Priority	Yes	Windows NT NetWare 5

\*Elron Software claims to do more than size TCP windows and calls its method for controlling traffic Dynamic Traffic Control (a method that includes but is not limited to TCP window sizing).

\*\*Like Elron Software, Packeteer claims that its products do more than what it calls "simple" TCP window sizing. PacketShaper products use a method called TCP Rate Control, a method Packeteer invented (and named).

the Desktop Management Task Force's (DMTF) Directory-Enabled Networks (DEN) initiative, as do 3Com, Ukiyah Software, and others. (DMTF is a standards organization that oversees the development of industry-standard and interoperable management tools and utilities. DEN strives to define a schema for integrating network equipment into a directory service.)

Routers, switches, and traffic shapers and their integration with enterprise-wide directory services such as NDS are forging the trek into another related frontier: policy-based networking. Vendors are currently creating policy management software that uses LDAP to communicate with directories, where traffic-prioritization rules are stored. For example, Cabletron Systems' SPECTRUM Enterprise Manager, slated to be released in the fourth quarter of 1999, will enable you to create QoS policies for routers and switches from more than 200 companies, including Cisco, 3Com, Lucent Technologies, and Nortel Networks.

Ukiyah Software recently announced the upcoming release of its distributed policy-management system called *NetRoad Active Policy System (APS)*. NetRoad APS will enable you to create, distribute, and manage QoS policies that are enforced throughout the network by routers, switches, and traffic shapers—and whatever combination of those devices you use. In its initial release, NetRoad APS will support Cisco routers and 3Com switches and, predictably, Ukiyah Software's own TrafficWARE, which Ukiyah Software is porting to NetWare 5. However, this is only a partial list. Ukiyah Software plans to support other vendors as well, ultimately striving to provide a multivendor, IP services management solution.

Among other things, policy-management systems enable you to create a traffic-prioritization policy just once. After you create that policy, a policy-based network enforces that policy from one end of a network to the other.

For example, Ukiyah Software's NetRoad APS will enable you to create one policy and then push that policy out to affected devices, including routers, switches, and traffic shapers. If you create a policy for Cisco routers indicating that they should give high priority service to Service Advertising Protocol (SAP) traffic to and from members of the Finance Group object, as specified

within NDS, APS will then push that policy out to all of the Cisco routers on your company's network or, Smith claims, to all of the Cisco routers on your service provider's network.

About one year ago, author Salvatore Salamone suggested that "the entire industry's idea of QoS needs to shift from traffic prioritization on one portion of a network to policy-based management on the entire end-to-end network." ("A Seri-

ous Look at Quality of Service," *Internet-Week*, March 9, 1998.) Whether vendors heeded Salamone's words or whether those words are a prophetic suggestion of the year-2000 odyssey toward policy-based networking, the entire industry is apparently on the verge of shifting its collective idea about traffic prioritization in just the manner Salamone prescribed.

*Linda Kennard works for Niche Associates, which is located in Sandy, Utah. ●*

### Eicon Technology 1/2 Page Island AD

4 7/8" x 7 3/8"  
(4.875" x 7.375")