Taming Your WAN

Joanie Wexler

How enterprises are improving bandwidth efficiency and application performance.

By happenstance or by design, you're likely to need a WAN traffic management strategy soon, if you don't have one already. Most applications have been developed to run on LANs, which are generally congestion-free, so they behave differently when traversing the more volatile WAN environment.

"Application performance on the WAN is usually disastrous or comical," said Peter Christy, principal analyst at NetsEdge Research Group. "And independent software vendors (ISVs) can only verify performance for their customers in a predictable environment—which the WAN isn't."

Let's take a closer look at how several enterprises are solving their WAN challenges.

Increased Reliance On The WAN

Enterprises are Web-enabling their existing business applications and are deploying new applications that might affect the performance of the existing apps. Sometimes, applications that network staff don't even know about sneak in.

Such is the case at McKee Foods, (Collegedale, TN). "An increasing number of applications on the network are unplanned," explained Bo Smith, IS group manager in the company's network services group "It's a shame to buy extra bandwidth to support under-the-table network applications that weren't in the budget."

Instead, he uses a compression appliance called the Sequence Reducer from Peribit Networks. He reports an approximate 5:1 improvement on his T1 links, which has allowed him to delay investments in additional network capacity.

He said he considered compression from his existing router vendor, Nortel, but decided on an independent appliance so that "we could put it anywhere regardless of whose routers we use."

Using compression to squeeze more capacity out of an existing link is one of several steps that organizations can take to manage WAN traffic. They need this management because the wide-area network is playing a larger role than ever in enabling access to IT resources. Meta Group, for example, points to a trend toward branch-office datacenter consolidation in sites that are bandwidth constrained.

In addition, Vertical Systems Group estimates that there are 2.55 million enterprise endpoints in the U.S. connected to private lines, frame relay and ATM services. Budget-challenged companies like McKee Foods are starting to learn that it's prudent, when possible, to manage bandwidth rather than upgrading capacity at all sites.

Troublesome Traffic

Many organizations have cited Citrix traffic as prompting their initial attempts to manage WAN application behavior. Citrix provides a software framework that allows enterprises to centrally manage and provision remote access to centralized applications in a thin-computing model.

"Citrix will drop a session if it doesn't have a minimum of about 5 kbps per user," said Matt Matin, security and systems engineer at LandAmerica Financial Group in Richmond, VA. After

WAN Traffic Management Deployment Checklist

- First, deploy Layer 7 monitoring and protocol discovery capabilities in QOS appliances, routers or standalone network management systems to determine WAN traffic mix and to identify what, if any, traffic needs management.
- Choose router-based QOS software or a standalone appliance for applying QOS.
- If using an appliance, consider the pros and cons of LAN-side versus WAN-side devices. A combination of both might be in order in a large network.
- If you choose router-based QOS, determine the impact on routing resources and existing configurations. Also determine if you have to upgrade your routing software.
- If combining compression or encryption with QOS, investigate to make sure there is interoperability between the functions□

Joanie Wexler is an independent editor and writer based in Silicon Valley. She has spent 13 years analyzing and covering trends in the IT and computer-networking industries.

Just "throwing" bandwidth at the problem might not work

LandAmerica deployed Citrix, "sessions were getting dropped, and we didn't know what was going on," Matin said. "We'd add bandwidth, and two months later, the same thing would happen."

LandAmerica, with about 900 sites, 2,000 circuits and 15,000 users in the U.S., uses Packeteer PacketShaper appliances for monitoring and QOS at its primary sites. "We saved bandwidth, bringing a T1 down to a 1,024-kbps pipe," he said.

In remote sites that simply require one protocol to be prioritized over another, though, LandAmerica uses protocol discovery and associated queuing provided in its Cisco IOS router software, a function Cisco calls Network Based Application Recognition (NBAR). "Running 900 PacketShapers [one in each site] would be cost-prohibitive," said Matin.

Monitor First

Before you can adopt effective traffic management, you need to find out what applications are on your network. Most organizations report surprises when they monitor their traffic mix.

"We had a site in Brazil that was fully congested for a month," said Tony Cabe, network infrastructure analyst at Metso Corp., which makes process industry machinery and systems. "Once we began monitoring the network [using Allot Communications appliances], we learned there was a user downloading every flavor of Linux from every university around the world!"

Traditional Layer 3 protocol analyzers will tell you what protocols are running, measure your overall bandwidth utilization and track latency and jitter, but they do not provide granular detail on a per-application or per-user basis. For example, you might learn that your link is 90 percent full, but then you have to troubleshoot for yourself which users or applications are consuming the majority of the resources.

Jay Mohr, data networks team manager at the Metropolitan Water District of Southern California in Los Angeles, explained: "We used a [Network Associates] WAN Sniffer on our T1s, then tried Cisco traffic shaping a couple years ago, but it was a pretty crude setup. We had to analyze and pinpoint the trouble ourselves, based on an IP address. Then we could shape the traffic."

The district has since switched to a Packeteer platform to automate the discovery of traffic behavior by application and handle the QOS.

QOS appliances such as those from Packeteer, Allot and Sitara Networks let you see what's actually on the network, traffic volume by application and user, peak usage patterns and whether any particular applications, protocols or users are monopolizing bandwidth. These breakdowns are usually shown graphically, for non-technical people. Most QOS products are transport-agnostic; you can use them on any type of WAN service.

Sophisticated monitoring, along with QOS, can be performed in various form factors, including the appliances mentioned, router software and standalone monitoring software. In March, for example, Cisco released an NBAR Protocol Discovery management information base (MIB) that enables management applications such as InfoVista's VistaView and Concord Communications' eHealth to display pie charts and graphs of traffic, broken down by application, protocol and user for monitoring and feedback.

A recent development in application-layer monitoring is that most vendors can now track protocols that hop randomly across ports. In the past, these devices were able to discern TCP or UDP traffic, for example, but made the association based on port number. When the application moved to another port, its visibility was lost. Now, these sessions can be continually tracked.

Alan Tosi, IT manager at Delmar Learning, a Clifton Park, NY-based Packeteer customer, cites peer-to-peer traffic such as AOL Instant Messenger (IM) and KaZaa as particularly challenging in this regard. "IM has a standard port it will use, but if it can't work on that port, it will scan for other available ports to let people chat. PacketShaper knows it's IM, though—it isn't worried about ports," he explained.

This capability allows Tosi to block MP3s and other traffic. "I don't want a music company coming to me and saying someone in my organization is swapping music files," he said.

Is More Bandwidth Necessary?

Why not just "throw bandwidth at the problem?" Because this solution might not work.

Joe Cron, VP of information systems and manager of systems, networks and facilities at CHD Meridian Healthcare in Nashville, can attest to this. "For many years, we bought more and more bandwidth, hoping for better performance. Though we'd have a bigger [network service] bill every month, we didn't always have improved performance to show for it," Cron said.

So his company deployed Sitara QOS appliances, hoping to save money on bandwidth and improve the performance of core applications. "We wanted our medical applications to have far more bandwidth than Internet Explorer. So we created different priority queues using the Sitara devices," said Cron. Monitoring helped CHD Meridian discover that some 384-kbps circuits only needed to be 256 kbps; it also found that some circuits also required increased capacity.

Bandwidth isn't always the answer because "bandwidth doesn't necessarily help latency," said Peter Firstbrook, senior research analyst at Meta Group. Latency can also be caused by distance—the time it takes to retrieve information—as well as server slowdowns and bandwidth-hogging applications that need to be rate-controlled. These kinds of traffic jams are better addressed by using QOS technology to tell servers to slow down their transmission rate, limiting the bandwidth

available to certain applications and guaranteeing minimum bandwidth to other apps, Firstbrook said (see *BCR*, August 2003, pp. 56–58).

Universities often think they need more bandwidth as students adopt peer-to-peer applications. Consider Louisiana State University. In the fall of 2001, LSU increased its Internet access bandwidth from 24 Mbps to 45 Mbps. "The minute we did it, it was all consumed," said Terry Doub, network security manager.

So the following spring, LSU moved to 62 Mbps and in July upgraded again to 155 Mbps. "We presumed that our problem was peer-to-peer traffic, but we couldn't prove it. So we started looking at bandwidth-management tools," said Doub, who settled on Allot's high-end NetEnforcer AC701, because it could keep pace with the 155-Mbps speeds. The university has since added the vendor's gigabit-speed version to manage traffic on the school's gigabit-speed local Internet connection that supports faculty and staff.

"We haven't heard a peep from people who were complaining about Web browsing a year ago," Doub said.

He noted that Packeteer performs the same functions, but not at speeds as high as Allot's: "And we played with some of the built-in Cisco QOS, like policy mapping, but it didn't cut it for us. At the time, NBAR didn't ID traffic at Layer 7; it could only track Layer 4 port numbers."

Doub now can shape traffic according to three primary groupings: students, faculty and staff. "Based on IP address, different rules apply," he explained. Dormitories can use some peer-to-peer applications, but are limited in how much bandwidth is available. For the other two groups, this type of traffic is disallowed.

"Students can download music—but we prevent others from coming in and taking music from us, because it creates a liability," Doub added.

Bandwidth management saves upgrades at colleges where peer-to-peer is popular

Shaping And Prioritizing Traffic

Once you've monitored your WAN links and have settled on whether you need extra capacity, you can then control the performance of high-priority and real-time application traffic and get more efficient use out of existing bandwidth, like LSU did.

This process of "shaping" involves throttling traffic at end stations and scheduling its delivery—smoothing out bursts to avoid big chunks of network congestion, thus reducing jitter, latency and packet loss. "Shaping" may also refer to any number of QOS capabilities that can be applied to traffic. These include marking traffic with priority information, directing the marked traffic to the appropriate priority queue and allocating a

TABLE 1 WAN Optimization Summary

Function	Description	Primary Application(s)	Representative Vendor Sample
Traffic Management and Prioritization	Applying one or more of the following to WAN traffic flows: shaping, priority queuing, minimum-bandwidth allocation, rate-limiting, blocking	To limit asynchronous, bursty, or superfluous traffic monopolizing bandwidth; to prioritize critical and real-time traffic; to block disallowed traffic	Adtran Allot Cisco Expand Packeteer Peribit Sitara
Compression	Sophisticated algorithms identify repeat patterns in network content to reduce the volume of packets transmitted	Legacy traffic, ERP, Citrix, intranet traffic	Cisco Expand ITWorx Nortel Packeteer Peribit Redline
Caching	Storing content locally with periodic refreshes from a remote origin server to circumvent latency-prone WAN circuits and improve user response times	Web traffic, on-demand multimedia streaming traffic	Allot Cisco Expand Network Appliance Novell/Volera Packeteer
Route Control	Determines which ISP connection has least round-trip latency and directs traffic to that network	Dual- or multihomed Internet connections	netVmg Opnix Proficient Networks Route Science Sockeye Networks

Meta Group contributed to the content of this chart.

Where should the QOS appliance reside: On the LAN side or WAN?

minimum or maximum amount of bandwidth to a certain application, protocol or user.

Most enterprises report a six- to 18-month payback on traffic-management investments. "There have been five or six occurrences where we've saved $135 to $150 per site per month by not having to upgrade a port or PVC," said Cron of CHD Meridian. "Our average hard ROI is 18 months. The real ROI, though, is that with [the Sitara appliance], performance is great. Without it, applications don't work. It's that cut and dried."

Form Factor Options

As mentioned, there are a number of QOS deployment options, depending on your network topology, requirements, budget and general networking philosophy. Some network managers prefer standalone QOS appliances, for example, because they don't want to risk disturbing their highly tuned router configurations. They might also frown on spending the time and money to upgrade router software to handle QOS or worry about the performance impact of consuming router processing cycles with traffic management.

"If our Sitara appliance experiences a fault, it automatically takes it out of the data flow so as not to interrupt routing. QOS integrated into the router can't really do that," said Cron. His philosophy is to keep standalone functions separate so he can determine whether each is doing its job.

Similarly, St. Louis-based biotech software design company Tripos Inc. sought out appliances four years ago "because we wanted to spend less money on routers. We wanted them to just route; that's their job," said Jerry Wintrode, network architect.

And LSU's Doub points out that organizationally, a separate group of employees handles router hardware decisions and code levels, while his group oversees network security and traffic management. "I wouldn't want others to make decisions [about routers] that would affect the performance of what I'm doing," he noted.

Other enterprises, though, figure you might as well take advantage of capabilities that are already available in your routers.

"Adding devices just adds complexity," said Matthew Haynes, network manager at Pactiv, a Lake Forest, IL-based maker of specialty packaging. "If you can get the same function in a router you already have, why not just use that?"

Pactiv uses native QOS commands in its Cisco routers for packet fragmentation, priority queuing and rate limiting. "An IOS upgrade is typically cheaper than maintenance [fees] on a QOS device," he said. "And with Cisco's focus on [voice over IP], they've really beefed up their capabilities to match or surpass the appliances."

However, Pactiv has also invested in Expand Networks' Accelerator appliances, a compression/caching device, to reduce the overall volume of data running on its WAN. "I could have gotten a compression card from Cisco, but that's more hardware I didn't want to buy," Haynes explained.

Also using router-based QOS is the Department of Rehabilitation in Alabama, where Cisco 7200, 3640 and 2621 routers do priority queuing for distributed IP call centers. "The Cisco routers were brought in to replace another vendor and were purchased specifically with the prioritization capabilities in mind," said Andy Cannon, manager of network operations.

He said he is setting IP Precedence bits in about six network routers to ensure that voice-over-IP (VOIP) packets are always placed in the router's top-priority queue and to ensure that audio- and videoconferencing packets are not discarded during congestion.

LAN- vs. WAN-side Appliances

If you settle upon using a QOS appliance rather than router software, one configuration consideration is whether to install it on the LAN or WAN side of your WAN-access router. There are pros and cons to each. A LAN-side device, which shapes and prioritizes traffic before it reaches the WAN router, supports two Ethernet connections, one to the LAN switch and one to the WAN-edge router. This makes deployment simple.

Delmar Learning chose a LAN-side device because the company was doing IPSec encryption. "A device on the LAN side of the router truly sees the application, shapes and prioritizes it, then passes it to our firewall for encryption and then to the router," Alan Tosi explained. A WAN-side device wouldn't have been able to ID encrypted packets in order to apply QOS.

On the other hand, Metso specifically wanted a WAN-side device when it installed 11 NetReality WiseWAN appliances (now supplied by Allot, which acquired NetReality a year ago).

Metso runs an international frame relay network tied together by an ATM backbone supplied by Infonet, as well as some MCI IP-VPN services in the Americas. The company was looking to reduce the cost and increase the efficiency of real-time applications such as Citrix, said Tony Cabe.

"We figured using a WAN-side QOS appliance, we could eliminate the expense of international real-time PVCs, because WAN devices can actually see the CIR being used at any one time and shape traffic accordingly," Cabe said. Early LAN-side devices didn't have this dynamic visibility. He said Metso has done away with the expensive international real-time PVCs.

Cabe added that if Metso had used a LAN-side device, it would have had to reconfigure its WAN switches to connect to certain ports. "Also, we're a hub-and-spoke topology; we didn't want to get into the position where if remote sites started to communicate directly through the hub, we couldn't shape that traffic" because the traffic would never hit the LAN-side device at the hub site, he explained.

Meta Group's Peter Firstbrook observed that if you are using a managed network service, the demarcation point where the service provider's responsibility leaves off and yours begins is often in the middle of the router, where the WAN meets the LAN. Your inserting a device in the WAN path could be an issue.

In Metso's case, however, "we made an agreement in our contract with Infonet that this was OK. They are responsible for the CSU and router, but we are responsible for the WiseWAN appliance [in between]."

Tripos' Jerry Wintrode also wanted WAN-side devices when he began investigating bandwidth shapers four years ago for his international frame relay network. He was seeking a device strictly for giving preferential treatment to videoconferencing, VOIP and Citrix applications.

For example, he nails up 500 kbps for Citrix, then allocates 50 kbps per user. "You can also guarantee a certain amount of bandwidth to a certain application within Citrix, such as Microsoft Word," he noted. "This way, no one Citrix user can kill other Citrix users, and no other traffic can kill Citrix as a whole."

Wintrode said that at the time he was in the market, Packeteer and NetReality were the only prospects, and "with Packeteer, for every subnet you supported across the WAN, you needed a separate appliance. The PacketShaper didn't support virtual LANs," he said. (Packeteer notes that it added this capability in May 2001.)

Tripos now has a NetReality 600 sitting on the LAN side in the datacenter, supporting subnet tunnels, and WAN-side WiseWANs with V.35 connections in all remote sites to prevent the return path to the hub site from becoming congested, according to Wintrode.

Other vendors in this space include Expand and Peribit, whose appliances began life as monitoring and compression devices and have recently gained QOS functions. These devices can be either LAN- or WAN-side.

Compression

Another technique, compression, enables enterprises to squeeze as much capacity as possible out of their existing WAN circuits, like McKee Foods discussed above. Compression offers as much as a 10:1 throughput improvement, depending on application and compression vendor. Compression appliances are often sold with licenses so you can simply use a software key to enable them to work on larger connections, rather than having to replace existing appliances with bigger ones.

O'Reilly Auto Parts, with datacenters in Dallas and Springfield, MO, uses Expand Accelerator devices. O'Reilly was so impressed that it refused to return trial units after a 30-day trial on a disaster-recovery link that supports real-time mirroring of business-critical data. "We bought the product on the spot," said Mark Garton, business continuity team leader at the company, which runs 1,000 retail locations and 10 distribution centers in the Midwest and Southeast.

"We immediately reduced a 12-Mbps pipe running at 80 to 85 percent utilization to a 6-Mbps pipe running at 50 percent utilization," said Garton, who added that O'Reilly has reduced its WAN expenditures by $2,700 per month.

In addition, for his real-time, high-availability mirroring, "we used to create more transactions than we could send; we'd get a couple of hours behind. Now, there is no latency."

Though Pactiv uses its Cisco routers for QOS, it has opted for standalone Expand Accelerator devices to squeeze more capacity out of certain European WAN links after rolling out an enterprise resource planning (ERP) system that threatened to require an international WAN capacity upgrade. "Instead, we increased our application count in those locations but maintained or reduced the size of our WAN links," said Matthew Haynes.

Compression cards are available for Cisco routers as well, though they generally provide only a 2:1 maximum throughput improvement. "And Cisco compression cards don't work very well beyond speeds of 200 kbps," said Meta Group's Firstbrook.

At press time, LandAmerica was poised to begin using the new compression capabilities in its Packeteer appliances to reduce the volume of data replicated between mirrored datacenters in Dallas and Richmond. "We did a lab test and saved 70 percent on bandwidth," Matt Matin said. "We're going to schedule downtime for a real test, and if we save 40 to 50 percent, we will immediately purchase the capability."

With Packeteer, the customer invests in an initial appliance platform configured for monitoring only; a device supporting 128 kbps to 200 Mbps of throughput lists for $2,250, according to Packeteer. You can layer on traffic shaping and QOS using a software key; the same holds true for adding compression. Prices range from $2,250 for monitoring only at 128-kbps speeds to $58,000 for running all features at 200 Mbps. Matin noted, though, that compression is available only on the lower-end PacketShaper models; the higher-end appliances support too many classes of traffic that, when combined with compression, would bog the appliance down.

Another enterprise, Santa Clara University in Santa Clara, CA, is using a Packeteer appliance to shape and prioritize traffic, but is also leveraging Redline Networks appliances to compress network traffic. Redline gear offloads Web processing and connection-management functions, but also compresses network traffic.

"Now, we are pushing access to administrative data to our students, faculty and staff via the Web. We're seeing about a 70 percent bandwidth reduction due to compression," said Ron Danielson, chief information officer.

Compression is another way to free up bandwidth

Internet route control could someday play a larger QOS role

FIGURE 1 Worldwide Compression Unit Shipment Forecast

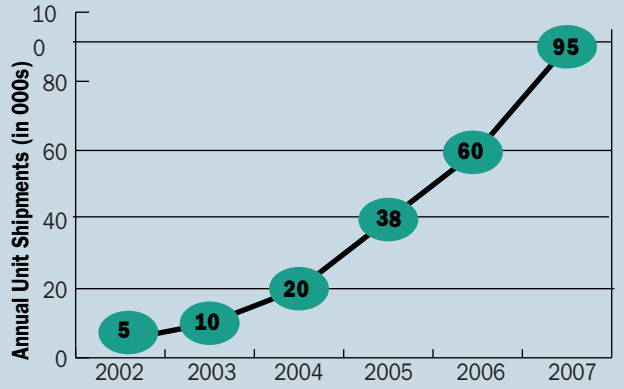

Source: netsEdge Research Group

For better or for worse, the best compression appliances use proprietary algorithms that require like appliances at either end, noted Meta Group's Firstbrook. "There are compression standards but they don't work as well as the proprietary ones," he said. "In addition, for compression and prioritization to work together, they usually have to be developed in an integrated fashion by the same vendor. Otherwise, QOS devices can't read the priority markings on compressed traffic."

Combining QOS With Other Functions

Other emerging QOS issues involve attempts to combine encryption and compression technologies with QOS using different vendors' wares. It's possible that if one vendor's device encrypts or compresses traffic and another's attempts to shape and prioritize it, the function won't work. This is because the shaping devices can't identify encrypted or compressed data.

To this end, industry players are starting to integrate these functions in their devices. Packeteer, known for its monitoring and shaping strengths, added compression to its PacketShapers in March. Likewise, Expand and Peribit, which got their start with sophisticated compression devices, have added QOS.

In Cisco routers, an IOS capability called "V3PN" (voice and video VPN) overcomes the encryption challenge by copying QOS priority markings to the IPSec tunnel header before transmission over the WAN.

Another option for devices that don't merge the two capabilities is to put all IPSec traffic in a single class with a certain priority, suggested Joe Cron of CHD Meridian. He said he asked his vendor, Sitara, in June whether the company would be including compression in its products, and "they indicated they might be investigating it."

Route Control

Another potential component of WAN traffic management is route control, or route optimization, for sites that dual-home their Internet connections to separate Internet service providers (ISPs). Route control, which usually requires full Border Gateway Protocol (BGP) peering with your ISPs, involves installing an appliance that sends test data across each network to determine which one is currently delivering the fastest round-trip time. The device then directs traffic to that network (see *BCR*, September 2002, pp. 20–24 and March 2003, pp. 18–21).

Route control may not yet be cost-effective for branch sites. "I haven't seen this take off at the edge of the network yet, because the solutions are still expensive," said Lawrence Orans, principal analyst at Gartner Group. Rather, he said, this capability is being applied mostly in enterprise datacenters.

Indeed, Wes Groves, network engineer at the Chicago Board of Options Exchange is using RouteScience PathControl appliances in dual-homed datacenter. Groves intends to configure his devices to also police the utilization of both links, and shift traffic to keep each link within its minimum contracted range as often as possible. For example, he could contract for 5 Mbps minimum, burstable to 9 Mbps, with each provider. When bursting, he pays a higher rate, so it would pay to have an automated appliance keep each link at 5 Mbps or under, when feasible.

SLA Verification

Finally, some enterprises would like to check their network availability and performance numbers against those generated by their network service providers. Generally, the issue isn't one of nickel-and-diming a provider for a credit if a service-level agreement (SLA) hasn't been met, but one of making sure that carriers have a stake in protecting the integrity of application performance.

"We use the NetReality box to draft SLAs and tell our carriers that the SLAs will be measured against our numbers, because we trust our numbers, not theirs," said Tripos' Wintrode. He said that, for example, when a carrier monitors traffic on 15-minute averages, SLAs look like they are being met. "But we found that on a second-by-second average, they were far off the mark."

The reason Tripos is so fussy is that scientists at his company run applications that require calculations that might take weeks or even months to process. "A 10-second outage could cost a scientist two weeks," Wintrode explained.

Tripos negotiates contracts in which it waives money-back guarantees. Instead, it stipulates that if SLAs are missed at all, to any degree, two months in a row, the company has the right to deem the contract null and void.

Conclusion

As organizations grow more decentralized, more traffic traverses the WAN to reach remote sites and users. In this respect, corporate reliance on the WAN for access to IT resources is skyrocketing. Because user WAN links are generally more constrained and prone to latency, jitter and packet loss than are LANs, controlling the behavior of WAN traffic is a mounting requirement.

The tools now available combine monitoring on a per-application, per-protocol and per-user basis to help determine what's actually on the network and how it is behaving and affecting other traffic. Based on this information, automated tools can be used to appropriately mark, classify, queue, shape and rate-control traffic with policies that can be propagated across many subnetworks.

Compression can also be used to boost efficiency and performance. Finally, route control can be considered an ancillary part of the picture.

Many enterprises say they like the plug-and-play simplicity of standalone traffic management appliances and the utilization reports they generate, while others prefer to leverage the QOS capabilities embedded in their edge routers. Both have their advantages and drawbacks; what's important is that enterprises use whatever method with which they are comfortable to stay in control□

Companies Mentioned In This Article

Allot Communications (www.allot.com)
AOL (www.aol.com)
CHD Meridian Healthcare (www.chdmeridian.com)
Cisco (www.cisco.com)
Citrix (www.citrix.com)
Concord Communications (www.concord.com)
Delmar Learning (www.delmar.com)
Expand Networks (www.expand.com)
Infonet (www.infonet.com)
InfoVista (www.infovista.com)
KaZaa (www.kazaa.com)
LandAmerica Financial Group (www.landam.com)
MCI (www.mci.com)
McKee Foods (www.mckeefoods.com)
Metso Corp (www.metso.com)
Microsoft (www.microsoft.com)
Nortel Networks (www.nortelnetworks.com)
O'Reilly Auto Parts (www.oreillyautoparts.com)
Packeteer (www.packeteer.com)
Pactiv (www.pactiv.com)
Peribit Networks (www.peribit.com)
Redline (www.redlinenetworks.com
RouteScience (www.routescience.com)
Sitara Networks (www.sitaranetworks.com)
Tripos Inc. (www.tripos.com)

Which approach is best? Whichever one gives you best control