

Sponsored by:



# NETWORKWORLD

This story appeared on Network World at <http://www.networkworld.com/reviews/2007/081307-test-wan.html>

## WAN acceleration offers huge payoff

### Riverbed edges Cisco, Silver Peak and Blue Coat in test of devices that speed traffic, cut bandwidth and save money

[Clear Choice Tests](#) By [David Newman](#), Network World, 08/13/07

#### CLEAR CHOICE TEST

Imagine walking into the CIO's office tomorrow and saying, "I can cut our [WAN](#) consumption by as much as 80 times, speed file transfers as much as 45 times and make our [Windows](#) users a whole lot happier." Think you'd get the CIO's attention?

Those aren't just idle claims. Seven months of rigorous testing showed us why application [acceleration](#) is such a hot area: These devices really work.

We tested boxes from Blue Coat Systems, [Cisco](#), Riverbed Technology and Silver Peak Systems in a true enterprise context, with a massive test bed pushing data over multiple T-3 and T-1 links (see "[How we did it](#)"). After pounding the systems with the most popular enterprise [applications](#), we're inclined to believe the hype.

Even if average speedups are "only" around 5 to 10 times, that's still a big improvement. With 31% of IT budgets eaten up by recurring monthly WAN costs, according to a recent Nemertes Research study, application acceleration promises potentially huge cost savings.

Riverbed's Steelhead appliances outperformed the field in most tests, and won our Clear Choice award.

- [A buyer's checklist for application acceleration](#)
- [What About SSL?](#)
- [How we tested application-acceleration devices](#)

But all these devices deserve serious consideration: Blue Coat's SG appliances for solid HTTP optimization; Cisco's Wide Area Application System (WAAS) for excellent compression, traffic transparency and interoperability with other devices; and Silver Peak's NX appliances for strong scalability and intuitive traffic reporting tools.

Sponsored by:



## Why is Windows so bad?

The problem statement for application acceleration is simple: Windows performance in the WAN is lousy. To begin with, Windows' two workhorse protocols -- TCP and NetBIOS -- were never intended for use in low-bandwidth or high-delay networks. Windows XP Service Pack 2 compounds these problems with some spectacularly suboptimal configuration defaults. (Windows Vista is better, but it isn't widely implemented yet.)

By default, XP's TCP stack advertises a receive window -- the maximum amount of data allowed in flight without acknowledgment -- of 64KB. That's fine as far as it goes, but XP isn't very responsive about resizing that window in response to loss or delay. A large, static receive window contributes to retransmissions, possible packet loss and poor response time.

To make matters worse, XP doesn't use a common TCP option called window scaling that can expand a 64KB receive window by a factor of four or more. Even when network conditions let XP go much faster, it won't. (There is a registry hack to enable window-scaling, but even then, it isn't used by the Windows file-handling protocol.)

WAN performance is always limited by the so-called bandwidth-delay product, but the constraints with Windows clients are especially severe. For example, if a link between Boston and Los Angeles has a 100-msec round-trip delay and the Windows TCP receive window is 64KB, the highest transmission rate possible is only around 5.6Mbit/s, regardless of link speed. Ordering up a T-3 or OC-3 won't help, at least not for any given Windows TCP connection; 5.6Mbit/s is as good as it gets.

WAN-acceleration devices compensate for these shortcomings with a variety of tricks, including block caching, compression, connection multiplexing and application-layer optimization. While not all devices implement every method, all sharply reduce response time and bandwidth for Windows applications across the WAN.

## Faster file service

As part of our research for this test, we asked [vendors](#) and several corporate IT shops to name their top five candidates for application acceleration, and every respondent named Common Internet File System ([CIFS](#)) as its top pick. This is understandable, given that Microsoft's notoriously chatty file-handling protocol originally was intended for LAN-only operations. Given its popularity and performance issues, we made CIFS the highlight of our performance testing.

We tested application acceleration the way enterprises use it -- with multiple WAN links and round-trip times. Our test bed modeled a hub-and-spoke WAN linking with a headquarters office plus four remote sites, two apiece on T-1 and T-3 links. The remote sites represented every permutation of high and low bandwidth and delay.

At each of the remote sites, we configured XP clients to upload and download directories containing Word documents from a Windows Server 2003 machine at headquarters.

To measure the effects of block and/or file caching, we ran the CIFS tests three times. First was a "cold run" with all caches empty. Second was a "warm run" that repeated the same transfer as the cold run, this time with the files already in cache. Finally, we changed the contents of 10% of the files; this "10% run" forced devices to serve some but not all content from the origin server.


The two most important application-acceleration metrics are bandwidth reduction and response-time improvement. While we measured both in this test, our results show there's not necessarily a strong correlation between the two. A device with a powerful compression engine might do well at reducing bandwidth consumption, but the time spent putting the squeeze on data might increase response time or, at best, yield only modest improvements. Conversely, some devices might willingly trade off a bit more bandwidth consumption if the net result is faster overall data delivery.

Looking first at bandwidth-reduction results, all products substantially lightened the WAN load, but big differences exist across devices depending on cache contents. See graphic:

[Click to see: CIFS WAN bandwidth](#)

## CIFS WAN bandwidth reduction [ FIGURE 1

Application-acceleration devices reduce WAN bandwidth consumption the most when they handle data they've previously seen. We tested each device three times: "cold run" with caches empty, a "warm run" with caches populated and a "10% run" with 10% of files changed. In the warm run, Riverbed's Steelhead used nearly 84 times less bandwidth than a baseline test with no device.



For example, in the cold run (caches empty), Cisco's Wide Area Engine (WAE) appliances were by far the most effective at compression, using nearly 28 times less bandwidth than was used in our baseline, no-device test. In contrast, the bandwidth savings for other devices seeing data for the first time was usually less than a two-times reduction in bandwidth, according to measurements taken by a ClearSight Networks Network Analyzer.

Note that we're presenting all results in terms of relative improvement rather than absolute numbers. For example, in the CIFS cold run, Cisco's devices consumed 130MB of WAN bandwidth, compared with 3.6GB with no acceleration device inline, which translates into using 27.82 times less bandwidth. (The absolute numbers from all tests are [available online](#) (Excel download).

Given that enterprise data patterns are repetitive and subject to change, bandwidth reduction in the warm and 10% test cases can be more meaningful -- and this is where these devices really shine.

Riverbed's Steelhead appliances topped these tests, reducing bandwidth by a factor of 84 in the warm run and a factor of 32 in the 10% run. While the other devices reduced bandwidth by a lesser degree, the improvements were still dramatic. Any device that reduces bandwidth use by 20 or 30 times must be considered a boon to IT budgets.

We also used the ClearSight analyzer to measure LAN bandwidth consumption (see graphic "[CIFS LAN bandwidth reduction](#)" and other [online-only performance results](#)). LAN differences among products were not as dramatic as WAN differences. The Blue Coat and Cisco devices reduced LAN bandwidth consumption by factors of 1.5 to 2 in our warm run and 10% run, because these vendors' headquarters devices served objects out of cache instead of from the origin servers. In contrast, the Riverbed and Silver Peak devices increased LAN use by 2% to

10%, probably because of appliance-control traffic. Changes in bandwidth use don't always correlate with changes in response time, however.

### Measuring CIFS response time

We used a common enterprise task to gauge CIFS response time, measuring how long it took for a client to upload or download a set of Word files to or from a server. We measured transfer times at each of our four remote sites -- each representing a different permutation of high and low bandwidth and delay. We're presenting the results for each site because users' requirements differ depending on where they work. As our results suggest, some appliances do a better job at accelerating CIFS in low-bandwidth settings; others are better for high-delay settings.

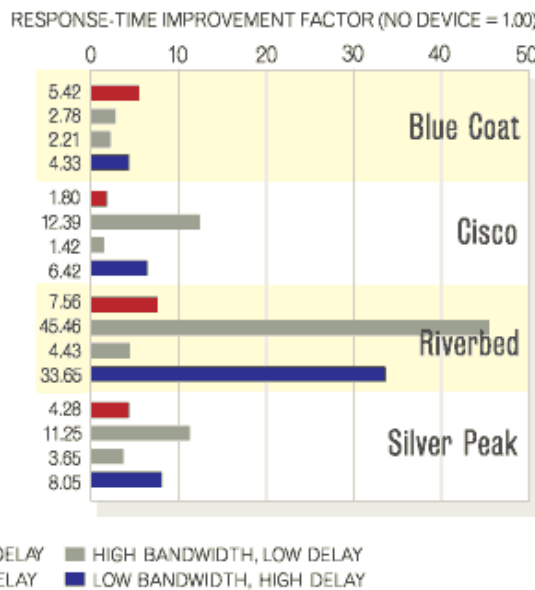
Arguably, the most important results for enterprises are from the 10% runs, where we offered 10% new content and 90% existing content to each set of appliances. This represents an enterprise where many users might see the same documents repeatedly but where there also would be some new documents added to the mix.

In the download tests, low-bandwidth sites tended to see the biggest improvements in response time, regardless of the amount of delay present. See graphic:

[Click to see: Downloading CIFS](#)

### Downloading CIFS [ FIGURE 2

File downloads go much faster with application acceleration. In situations where 10% of user data changed, downloads moved as much as 45 times faster with acceleration than without it. The biggest boosts generally came on lower-speed T-1 links, regardless of delay.



Riverbed's Steelhead appliances sped up file transfers 45 times to a low-bandwidth, low-delay site and 34 times to a low-bandwidth, high-delay site. The Steelhead appliances were also tops for the high-bandwidth sites, but to a lesser degree, with speed increases of four to seven times.

The Silver Peak NX appliances were next most efficient overall, with speedups of three to 16 times (again, with the most improvement shown for low-bandwidth sites), followed by the Cisco and Blue Coat appliances.

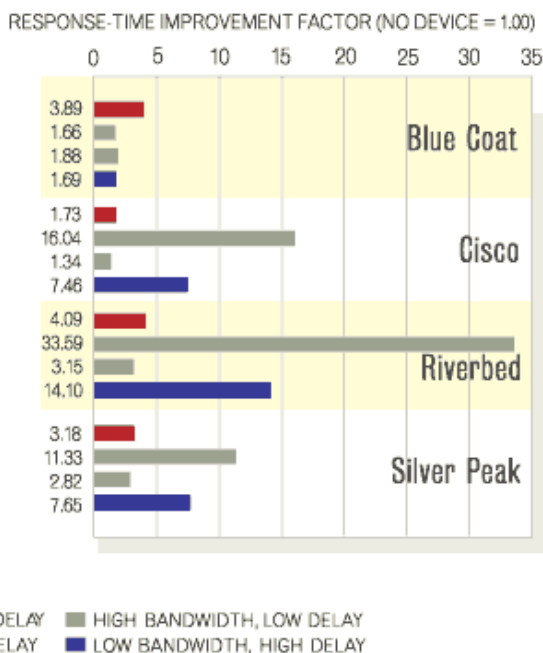
File uploads generally don't benefit from application acceleration as much as downloads do. When handling client downloads, acceleration devices either serve content from a client-side cache, pipeline data using read-ahead operations or employ some combination of the two approaches. That's not possible with write operations, because an acceleration device can't predict in advance what data the client will send.

Even so, big improvements in upload performance are still possible. See graphic:

[Click to see: Uploading CIFS](#)

## Uploading CIFS [ FIGURE 3

While file uploads don't benefit from acceleration as much as downloads, the benefits still can be substantial. Uploads from clients to servers moved as much as 33 times faster with acceleration than without it. As with downloads, the greatest benefits came on lower-speed T-1 links, although all the speedups were significant.



Riverbed's Steelhead appliance again led the pack, with speedups of three to 34 times compared with no acceleration. Accelerations from the Silver Peak, Cisco and Blue Coat devices were less dramatic but still significant, moving traffic 1.3 to 16 times faster than our baseline test. Most devices sped up data the most from low-bandwidth sites. Blue Coat's SG was an exception; it delivered the greatest upload benefit to the high-bandwidth, high-delay site.

Note that response-time improvements do not track linearly with bandwidth-reduction results. For example, Cisco's devices were more efficient, relative to their competitors, at reducing WAN bandwidth consumption than at speeding CIFS transfer times.

In reviewing the CIFS results, Riverbed commented that it achieved even greater improvement over no-acceleration baselines by using many small files. Our tests used a mix of random file sizes of 25KB to 1MB. Both approaches have their merits: Riverbed's short-file methodology is more stressful on devices' CIFS processing engines (stress is a good thing in device benchmarking), while a mix of larger files may offer a more meaningful prediction of device performance in production settings.

### Mail call

After CIFS, the next most popular candidate for acceleration is Messaging API (MAPI) traffic. MAPI is the e-mail protocol used by the Microsoft Exchange server and Outlook clients. All devices tested can speed up MAPI traffic, but in our tests the improvements were far less significant than in the CIFS tests.

In our MAPI tests, all clients sent messages -- some with Word attachments, some without -- to all other clients through an Exchange 2003 server. As with the CIFS tests, the number of messages was proportional to each site's link speed -- fewer messages for clients at T-1 sites, more for those at T-3 sites.

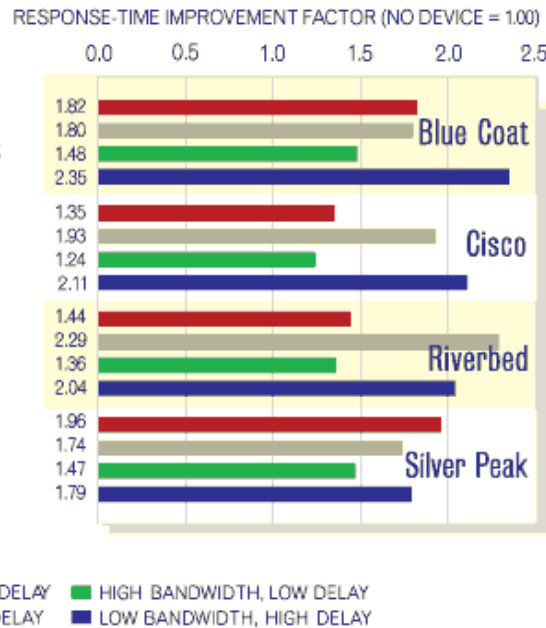
There was significantly less differentiation among products when accelerating MAPI traffic, compared to CIFS

traffic. See graphic:

[Click to see: MAPI acceleration](#)

## MAPI acceleration [ FIGURE 4

The MAPI e-mail protocol showed lower but still significant improvement in our tests compared with CIFS file traffic. Differences between products also were much smaller. This may be tied to our use of Outlook 2007 clients, which by default enable encryption of e-mail contents, giving appliances fewer opportunities for optimization.



All products sped mail delivery, but only by factors of 1.24 to 2.39 compared with a no-device baseline. Averaging results across all sites, the Blue Coat devices provided the biggest boost for mail traffic, but by a relatively small margin over the Riverbed, Silver Peak and Cisco devices.

Doubling e-mail performance is nothing to sneeze at, but we also wanted to understand why MAPI performance didn't match CIFS performance. A few minutes with the ClearSight analyzer gave us the answer: The Outlook 2007 clients we used in this test encrypt e-mail traffic by default.

To the acceleration appliances, most of the MAPI data structures weren't visible to be optimized. Some acceleration was still possible, through TCP optimizations or because some MAPI traffic was visible. After reviewing the results, Riverbed said it encourages Outlook 2007 users to disable encryption for highest performance. That said, network managers using the new version of Outlook should consider whether the [security/performance](#) tradeoff is worthwhile.

### A faster Web

We measured acceleration of HTTP traffic in two tests, one with 248 and one with 2,480 concurrent users. The results were a bit surprising: While the products delivered Web traffic as much as seven times faster than a baseline test without acceleration, performance didn't necessarily improve as we added more users.

To avoid overloading the sites on slower links, we put proportionately fewer users at the T-1 sites than at the T-3 sites. For example, our 2,480-user test involved 1,200 clients at each of two sites on a T-3, and 40 clients at each of two sites on a T-1. We used Spirent Communications' Avalanche/Reflector tool to emulate Web clients and servers. Because previous studies of Web objects place the average size at 8K to 13KB, we configured the clients to request an 11KB object from the servers.

As in the CIFS and MAPI tests, the Riverbed Steelhead appliances delivered Web traffic the fastest (see graphic [Web acceleration](#)). In all three ways we measured -- transactions per second, traffic rates and response time -- the

Steelhead appliances delivered Web traffic seven times faster than tests with no device inline. We observed the same seven-times improvement with 248 and 2,480 users; because LAN and WAN bandwidth use was almost identical in each test, it's likely that WAN bandwidth was the bottleneck.

Blue Coat's SG appliances were second fastest, but that result must be stated with a caveat: The Blue Coat boxes worked better with fewer Web users, not more. Compared with no acceleration, the Blue Coat appliances boosted Web performance by around seven times for 248 users, but by around six times for 2,480 users (and that's just for transactions per second and data rate; the response time improved by only a factor of three).

We noticed some erratic Address Resolution Protocol ([ARP](#)) behavior in tests involving 2,480 users when Blue Coat forwarded either Web or SSL traffic (see "[What about SSL?](#)"). Although Blue Coat replicated our issue in-house and produced a software fix (now available to customers), we still observed sluggish behavior in the 2,480-user tests after applying the update.

Silver Peak's NX appliances were third-fastest, tripling transaction and data rates and reducing response time by around 2.5 times when handling 248 users. With 2,480 users, performance dipped slightly (by about the same margin as Blue Coat's appliances), though traffic still moved substantially faster than in our no-device baseline test. Silver Peak says these results are roughly in line with its in-house testing.

Cisco's WAE appliances better than doubled performance with 248 users, and more than tripled performance with 2,480 users. Cisco's WAE devices don't proxy Web traffic as they do with CIFS, so the performance improvements here are largely attributable to TCP optimizations.

## QoS testing

[QoS](#) testing revealed some of the most interesting -- and in some ways most problematic -- results of all our performance testing. While three of four products did a virtually perfect job of prioritizing traffic, the path there was anything but straightforward, involving much tuning -- and in some cases external devices to protect key flows during congestion.

To measure QoS capabilities, we offered a small amount of high-priority traffic -- in this case, a single [VoIP](#) call, which is sensitive to delay and jitter -- while walloping the WAN with huge amounts of background traffic. We used User Datagram Protocol ([UDP](#)) for both high- and low-priority flows; VoIP uses UDP by default, and TCP was not suitable as background traffic, because of its built-in congestion control.

We also determined whether devices could "re-mark" [Diff-Serv](#) code points ([DSCP](#)), a good practice in guarding against rogue users or applications marking their flows with an inappropriate priority.

Blue Coat's SG appliances couldn't participate in this test, because they don't optimize UDP traffic. The other vendors turned in excellent results but used different paths to get there.

Cisco recommends using WAN routers (in this case, the Cisco 3845 and ISR 2800 Series devices it supplied) rather than application accelerators for shaping traffic. Cisco's WAAS-acceleration devices and [routers](#) work together using network-based application recognition (NBAR). We verified in testing that flows the acceleration devices classified using NBAR will be prioritized by the routers during congestion. The routers turned in great results; the ClearSight analyzer measured R-value, an audio-quality metric, as 92.03 out of a possible 93, and they correctly re-marked DSCPs.

Note that ultimately Cisco's entry performed prioritization on its routers, not on the application-acceleration devices, though the latter did play a role in classifying traffic. This differs from the Riverbed and Silver Peak devices, which performed prioritization on board. Many network managers already run QoS on WAN routers, and for them handing off this function to a router isn't a big deal. For users just getting started with QoS, it may be simpler to set it up on application-acceleration devices, and leave routers alone, at least for now.

The Riverbed and Silver Peak appliances also protected voice traffic, with R-value scores of 91.80 and 90.07, respectively, and both correctly re-marked DSCPs.

Of the two, the Silver Peak NX appliances were easier to configure. They correctly classified VoIP streams and shaped traffic according to the parameters we defined. Riverbed's Steelhead appliances don't classify real-time protocol streams automatically, and a bug in the software version we tested wouldn't let us manually define port ranges. Instead, we used other criteria, such as source address, to classify VoIP streams.

## Concurrent connections

Our final performance test determined the maximum number of TCP connections each system could optimize. This is an important metric for enterprises with many remote offices and hub-and-spoke network designs, where connection counts for data-center devices can run into the tens of thousands. All the devices we tested get into that tens-of-thousands range, but there was more than a fourfold difference between the highest and lowest capacities.

To measure connection concurrency, we configured Spirent's Avalanche to issue a Web request once a minute, letting us establish many connections and keep them alive. We kept adding connections until transactions began to fail or the devices stopped optimizing new connections.

Cisco's new WAE-7371 came out tops in this test, accelerating more than 50,000 TCP connections (see graphic [Maximum accelerated TCP connections](#)). Silver Peak's NX appliances were next, optimizing 43,306 concurrent connections. This is well short of the NX 7500's rated capacity of 128,000 optimized connections, a level that Silver Peak achieved in internal testing. We were unable to reproduce that result in our lab, and, despite extensive troubleshooting, neither we nor Silver Peak's engineers were able to explain the difference. The Blue Coat SG appliances were next, handling about 19,500 optimized connections.

Riverbed's Steelhead 5520 optimized more than 12,200 connections, but that result reflects the limits of the two Steelhead 3520 units through which we set up connections. Riverbed says the higher-end 5520 model can optimize 15,000 connections. We were unable to confirm that result, but our tests did show that each 3520 slightly outperformed its rated limit of 6,000 connections to get to the 12,200 total mentioned previously.

## Features and functions

Most testing focused on performance, but we also assessed devices for functionality, manageability and usability. Each of these areas turned up at least as many differences as the performance tests did.

All acceleration devices reduce the number of bits sent across the WAN, but they do this in very different ways. The Blue Coat and Cisco devices act as proxies, terminating connections between clients and servers and setting up new sessions on their behalf. Riverbed's devices can proxy traffic, though the vendor did not enable that feature for this test. Silver Peak's NX appliances don't proxy traffic.

Transparency is another architectural difference. Blue Coat and Silver Peak engineers respectively configured SSL or generic routing-encapsulated tunnels between appliances, and Riverbed can use SSL tunneling. Tunneling may pose a problem if other inline devices, such as [firewalls](#) or bandwidth managers, need to inspect traffic.

Cisco claims this is a major differentiator for its WAAS offering, which doesn't hide traffic from other devices and automatically learns about new traffic types from other Cisco devices using NBAR. A powerful classification engine, NBAR in our tests classified even applications using ephemeral port numbers, such as for such as those used for H.323 and [Session Initiation Protocol](#). Silver Peak's appliances also classified such traffic. Then again, transparency isn't an issue for users who don't need application visibility among acceleration devices.

Application support also varies, but it's less important a differentiator than performance, manageability and usability. It's tempting -- but also a bit misleading -- to compare the number of predefined application types each

vendor claims to optimize. First, the applications involved are important only if they're running in your enterprise. Second, acceleration devices still may boost performance even if a given application isn't predefined, thanks to compression and TCP optimization. Finally, all devices we tested allow manual definition of new application classes based on addresses and port numbers (though these may not be subject to the same speedups as some predefined types).

## Manageability

To look after all the devices in our test bed's enterprise, we asked each vendor to supply a central management system.

We assessed centralized management in terms of functions and reporting features. On the functions side, all vendors but Blue Coat offer a centralized method of pushing out configuration changes or software upgrades to all appliances. Blue Coat indeed can push changes and upgrades but only by manually defining a job to push out the change. All vendors allow appliances to be defined into groups (though Blue Coat's Director appliance requires a manually defined job to perform an action on a given group).

All devices use a dashboard display to show application distribution and volume during predefined periods. These displays can be enormously helpful in managing application traffic even before acceleration is enabled. It's pretty common to find during installation that enterprises are running applications they didn't know about.

Once acceleration is enabled, these devices use pie charts and bar graphs to report on compression, percentage of optimized vs. pass-through traffic and data reduction.

The Cisco, Riverbed and Silver Peak appliances aggregate displays across multiple devices, a useful feature for capacity planning. There were differences in terms of the application data and time periods supported; for example, Silver Peak's display was useful in troubleshooting because -- uniquely among the products tested -- it reported on packet loss and did so in per-minute intervals.

## Usability

There are significant usability differences among the accelerators, but we'll be the first to admit this is a highly subjective area. If we had to rank the systems in terms of ease of use, the lineup would be Riverbed, Silver Peak, Cisco and Blue Coat.

Riverbed's Steelhead appliances came closest to the goal of "just working." Setup took less than half a day. Once we were up and running, we found the user interface to be simple and well designed. It was easy to make changes and view reports, even without delving into the company's well-written documentation.

Silver Peak's NX appliances also feature a simple user interface with excellent reporting on current and historical statistics. The central management display wasn't as polished or fully featured as Riverbed's, although unlike Riverbed's, it includes a topology map of all appliances.

Cisco's display bristles with features and commands -- perhaps too many. Cisco's redesigned dashboard offers whizzy graphics, useful pie charts on CIFS application performance and (like Riverbed and Silver Peak devices) real-time connection monitoring and per-device reporting on connection statistics. Getting to specific commands or opening logs often took more steps than with other devices, however; further, not all the commands available from the device command line were available from the GUI, and vice versa.

Blue Coat's management software, while powerful, was the most difficult to use. Individual appliances used a Web-based Java application that was sluggish; further, it worked with Internet Explorer but not Firefox. And some predefined tasks in other vendors' devices, such as updating configuration or images, required manual definition in the Blue Coat devices, or touching each appliance individually.

*Newman is president of Network Test, an independent test lab in Westlake Village, Calif. He can be reached at [dnewman@networktest.com](mailto:dnewman@networktest.com).*



Newman is also a member of the Network World Lab Alliance, a cooperative of the premier reviewers in the network industry, each bringing to bear years of practical experience on every review. For more Lab Alliance information, including what it takes to become a member, go to [www.networkworld.com/alliance](http://www.networkworld.com/alliance).

All contents copyright 1995-2007 Network World, Inc. <http://www.networkworld.com>