

WAN Acceleration Solutions

A collection of educational articles and technical resources on WAN acceleration and its impact on disaster recovery, network backup and WAN performance.

INSIDE

Top 10 Considerations for Scaling a WAN Acceleration Solution	2
Push Your WAN to Peak Performance	7
Building a Better WAN	9
Silver Peak Squeezes Into the WAN Game	13
Data Security as a Service	14
Boosting Network Performance: 5 Tips	17
Six Steps Toward Disaster Recovery	19
Effective Long-Distance Data Protection	21

Compliments of:



Silver Peak

Top 10 Considerations for Scaling a WAN Acceleration Solution

SECURITY AND COMPLIANCE CONCERNS ARE CREATING increased pressure to remove data out of branch offices and to implement robust business continuity plans. At the same time, new WAN acceleration technologies have emerged that deliver order of magnitude performance improvements across a wide breadth of applications. These two factors are creating a “perfect storm” in the IT industry, which is elevating WAN acceleration from a tactical fix to a strategic enabler of key business initiatives.

As WAN acceleration becomes more critical to business operations, the average size and scope of deployment is steadily increasing. This creates a variety of questions pertaining to product scalability. What happens when new applications are deployed? How easy and cost effective is it to add new offices? How is performance affected by network load? Do higher capacity WAN links behave the same as lower capacity links?

There are numerous factors that affect the scalability of a WAN acceleration solution, from hardware architecture to ease of use. This paper identifies the top 10, and provides guidelines for overcoming scalability challenges in real-world enterprise environments.

WAN Throughput

Approximately one third of all enterprises have data centers with 45 Mbps or higher WAN connections, fifteen percent of which are over 100 Mbps. Large WAN links are required to connect many remote offices to centralized resources (i.e. to support a large “hub and spoke” architecture), and to support data center-to-data center applications like replication, backup, and disaster recovery.

To accommodate large offices, WAN acceleration appliances require adequate WAN throughput capabilities. A scalable

solution can achieve several hundred Mbps of throughput in a single appliance. This is much more cost effective and easier to manage than solutions that require multiple appliances to be clustered within a data center to deliver adequate WAN capacity. In addition, clustering often leads to reduced performance in large networks. When only one appliance is deployed, it is guaranteed that successive requests for data will always go through that device, which makes it easy to detect and eliminate duplicate data across the WAN. When multiple appliances are deployed, successive requests for data are sometimes handled by different devices, limiting the effectiveness of data reduction and compression technologies.

WAN throughput can vary with acceleration techniques. For example, when only latency mitigation is employed, such as TCP window sizing and selective acknowledgement, it is fairly easy to exceed 600 Mbps of throughput in a single appliance. It is more challenging, however, to achieve high throughput when more advanced techniques are employed, such as data reduction. Enterprises should be aware of this distinction, making their vendors highlight the maximum throughput that can be achieved (with a subset of features) separately from the throughput that is expected when all features are enabled.

It is often useful to perform a “sanity check” on vendor’s throughput claims by dividing WAN capacity by LAN capacity. If, for example, an appliance supports only 1 GB of LAN traffic and the vendor is claiming 600 Mbps of WAN throughput, then that product is delivering less than 2x improvement when fully loaded. That is fine when latency mitigation is all that is required, as discussed above. But, that ratio is not enough to support the massive performance gains that are achieved when doing data reduction across the WAN, which can easily average in the 10-20x range.

Copyright © 2006 InfoWorld Media Group. All rights reserved.

See the full selection of *InfoWorld* “IT Strategy Guide” reports at <http://www.infoworld.com/store/>.

Flow Limits

WAN acceleration appliances may never reach their stated WAN capacity if the maximum number of TCP sessions supported is reached well before the WAN pipe is filled.

For example, a mid-sized enterprise with 45 offices and 50 users in each office would generate approximately 22,500 TCP flows (assuming the average user has 10 TCP flows open at any given time.) If all of these offices are connected to a main data center via a 45 Mbps link, one might assume that a single WAN acceleration appliance with 45 Mbps of WAN capacity could be deployed at the head-end to support this entire network. However, this is not always the case. If the appliance has a low TCP flow limit, then WAN capacity is irrelevant when it comes to sizing this network. For example, if only 4,500 TCP sessions are supported simultaneously per appliance, then 5 appliances would actually be required in this data center instead of the 1 previously assumed. That is a significant difference in terms of cost, management, and support – which becomes even more significant as the network grows even larger.

It is important to determine the TCP flow limits of a WAN acceleration solution prior to sizing a deployment. In addition, one must determine what scenarios may affect this number. For example, do flow limits actually decrease when specific features are enabled, such as disconnected file services? In addition, what happens when TCP flow limits are exceeded - is WAN traffic blocked or sent un-accelerated across the WAN? This last question is particularly important when dealing with applications that have long-lived connections, such as CIFS and email. If a user logs on during a busy period and their connection is sent across the WAN un-accelerated, will the session remain un-accelerated for the lifetime of the connection – even if TCP flow counts drop to acceptable levels?

Not surprisingly, a product may demo great in a lab trial where flows are kept to a minimum, but fall over when deployed in a large network. To ensure that there are no surprises during full-scale deployment, one must take measures to test TCP flow limits from the onset of a WAN acceleration project.

Performance Under Load

Performance can vary based upon the number of users and/or applications deployed in a network. If a WAN acceleration appliance does not have enough memory or disk space to support large environments, performance will start to degrade as the network grows. This often starts to occur well before stated throughput and flow limits are reached, and worsens as the network continues to increase in size.

The right hardware platform can help to ensure better performance under load. A 64 bit architecture, for example, will have more memory and disk space than older platforms, making it better suited for large environment. By using multi processor, multi-core processors, appliances can achieve higher performance with increased memory addressability. In addition, while native 64-bit architectures enable physical and virtual memory to be directly addressed, 32 bit architectures can only address 4 GB of physical memory directly. This requires more complex memory de-referencing, which is substantially slower under heavy load.

Hardware Acceleration for Encryption

The last thing that enterprises want to do is sacrifice security for the sake of application performance when they replace branch office servers with new WAN acceleration appliances. To overcome this issue, many vendors will employ encryption on their WAN acceleration devices. Encryption of local drives can protect data at rest, while technologies like IPsec can use encryption to protect data sent across the WAN. To perform data reduction and compression techniques on SSL traffic, the WAN appliance must securely become part of the trusted security domain, decrypt the SSL streams, optimize the traffic, then re-encrypt the traffic. This requires fast authentication, and high speed encryption and decryption. Often these techniques are used together. For instance it is very dangerous to store decrypted SSL content in-the-clear on a disk. So encrypting data at rest is a co-requisite for secure SSL optimization.

Encryption is very computationally intensive, which can have an adverse impact on the performance and scalability of a WAN acceleration appliance. Even if the

WAN link is only a few Mbps (and within reach of software encryption), peak LAN bandwidth can be greater than 100 Mbps. In addition, WAN acceleration appliances often require several hundred Mbps of bandwidth to perform disk read and writes. As all of this traffic needs to be encrypted and decrypted, dedicated, multi-Gbps security hardware co-processors are required to offload encryption functionality. This is the only way to ensure maximum throughput when encrypting network traffic, enabling a WAN acceleration solution to grow in a secure and reliable fashion. Even small branch office appliances can benefit from hardware acceleration for encryption to handle SSL acceleration and encryption of local data stores.

It is generally good practice to see how performance changes when all encryption services are enabled. This is the only way to ensure that there is adequate processing capacity with a WAN acceleration solution to perform security with sacrificing scalability and performance.

Breadth of Application Supported

The business case for WAN acceleration is often predicated on the number of applications being accelerated. When more applications are supported by a WAN acceleration solution, enterprises experience larger productivity gains. This makes it easier to justify investments in new WAN acceleration hardware and software.

The most scalable WAN acceleration solutions are designed to support new applications and future revisions of existing applications with minimal configuration. To achieve this, the following is required:

- Transparency. When modifications are required to clients, servers, routers, or the application itself, it becomes increasingly harder to support future revisions or new types of traffic. The most scalable WAN acceleration solutions are completely transparent to existing infrastructure.
- Multiprotocol support. A WAN acceleration solution that supports only TCP-based applications will have less applicability across the enterprise than a solution that supports both TCP and UDP. The latter, for example, can be used to improve the performance of VoIP, video,

disaster recovery, and other applications.

- Low latency. If a WAN acceleration solution imposes high latency on WAN traffic, it cannot be used for time sensitive traffic. This limits its focus to bulk applications, such as file transfers and email.

While initial deployments sometimes focus on a few key applications, such as file and email, most enterprises ultimately want to scale their WAN acceleration investments to support all business critical WAN traffic, including Voice over IP (VoIP), video, ERP, CRM, and custom applications. To satisfy this scalability requirement, application breadth is an important consideration when evaluating WAN acceleration solutions.

Quality of Service (QoS)

QoS is a key feature for scalability as it enables more applications to co-exist with one another on a single network. In addition, it ensures that bandwidth is adequately distributed as more applications are vying for a shared limited resource.

Policing and traffic shaping are an indispensable part of any complete WAN acceleration solution. Even with drastic reductions in aggregate bandwidth consumption, there are still bursts of uncompressible data, so it is critical to guarantee bandwidth for enterprise critical applications. In a well-designed network, QoS will be managed at every potential bottleneck point. It is particularly important to implement QoS in the WAN acceleration appliance as it is the only element that has both a pre- and post-optimization view of the traffic.

RAM vs. Disk-Based Data Reduction

The amount of storage that is used for data reduction techniques will impact the overall performance in larger networks. That is because the efficiency of a compression or data reduction technique is dependent on the likelihood that data streams that have traversed the network in the past can be recognized on subsequent passes and eliminated across the WAN.

The most effective data reduction solutions leverage hard drives on appliances to store several weeks or months' worth of traffic patterns, a concept called "disk based data

reduction”. This efficiently eliminates the transfer of duplicate data by leveraging information collected over extended periods of time. In contrast, some products perform data reduction in RAM (as opposed to using disk drives.). While the two solutions may provide similar benchmark results for short data runs, (e.g., in simple lab tests), significant differences will be seen over time and under increased load. Better “long term” memory and higher storage capacity will perform better in real world environments as resources fill up and more information is learned about the particular network environment.

A fully utilized 155 Mbps WAN link, for example, will send 60 Gigabytes (GB) of data per hour in a single direction (120 GB per hour if operating in full duplex.) At 12% utilization, that same link will send nearly a Terabyte (TB) of traffic over a 24 hour period. A WAN acceleration appliance that has a large, rapidly referenceable effective data store of multiple TBs will be extremely efficient in this type of environment. In contrast, an appliance that operates only out of physical RAM will be limited to 4GB or 8 GB of active data store, resulting in less than 10 minutes of traffic retention. While the differences may not be noticeable when transferring small file sets (e.g., a few GBs), they are apparent when large volumes of information are sent over extended periods of time, as is the case with data center replication and disaster recovery.

Effective Storage Capacity vs. Stated Capacity

When performing disk based data reduction, vendors use different methods to store information. Some techniques are more efficient than others, resulting in better usage of available storage space, which leads to improved scalability.

A scalable WAN acceleration appliance will store a single local instance of information that is independent of all peer appliances. The alternative – storing a separate instance for each individual WAN link (i.e. peer appliance) – consumes significantly more disk space, particularly when many offices are involved. In addition, detecting and storing repetitive patterns in a bidirectional fashion can dramatically improve scalability by eliminating the

need to store a separate instance of information for each direction of a WAN link. For example, if a 10 MB file is sent back and forth between 100 offices, a head-end appliance that works bi-directionally and stores a single instance of information would use up approximately 10 MB of local data store. If the same appliance stored a separate instance for each direction of each WAN link, it would require 2 GB of local hard drive capacity to support this scenario. The first solution is getting 200x better effective storage capacity than the second solution!

Management

Remote provisioning, management, and monitoring are important to scaling WAN operations. This is particularly true in large networks, especially where servers are removed from remote offices and there is limited IT staff on-site to handle day-to-day tasks. Features such as zero-touch provisioning, which allow a device to be sent to a remote office and then automatically provisioned from a central location, are critical for growing a large deployment. The ability to centrally generate network traffic reports and export them to commonly used formats, such as Cisco Netflow, are also important for scalability. Other capabilities that can help ease the burden of managing a large WAN acceleration deployment include centralized alarm aggregation, scheduled tasks, auditing, trending reports, and centralized software updates.

High Availability and Clustering

As WAN acceleration deployments become larger, they often become more strategic to the enterprise. As a result, availability often goes hand in hand with scalability. This is especially true when WAN acceleration is used for data center-to-data center applications, such as disaster recovery, where downtime can mean a significant loss of revenue.

To address the need for high availability, most vendors will support a wide range of network configurations. These often use common techniques to ensure redundancy, such as Virtual Router Redundancy Protocol (VRRP) and Web Cache Coordination Protocol (WCCP). These same protocols can also be used to cluster multiple appliances togeth-

er for better scalability. This enables traffic to be balanced across multiple devices for better performance under heavy load.

Resiliency can also be built into the WAN acceleration appliance itself. Common features include redundant load-sharing power supplies, fail-to-wire network interfaces, ECC memory and hot-plug RAID drive configurations.

Additional resources are available at:

application-delivery.org ↗

—report compiled by *ApplicationDelivery.org*

Push Your WAN to Peak Performance

YOU CAN ALWAYS PICK OUT THE IT HEROES: THEY are the ones that keep their users happy — and productive. Sometimes that's hard to do when the users have to access applications and data on the other side of an over-subscribed WAN link. It isn't just a bandwidth issue; a big pipe isn't always a fast pipe. Users need a reduction in latency and protocol chattiness coupled with TCP optimization and intelligent caching. They need a WAN acceleration and optimization solution.

Silver Peak, a new player in the WAN optimization and acceleration game, is shipping the NX family of appliances that not only scale well and provide terrific performance increases but optimize UDP (User Datagram Protocol)-based applications along with TCP-based traffic. The NX-3500 includes a flexible QoS engine and can define ACLs (access control lists) on a per-tunnel basis. To monitor the overall health of the WAN circuit, reporting information (in the form of real-time and historical charts) is available via the browser-based UI.

A 2U appliance, the NX-3500 comes with dual redundant power supplies, fail-to-wire Gigabit WAN ports, and 500GB of local hard drive space. I tested the NX-3500 in my lab for a few weeks and found the appliance capable of handling a wide range of applications efficiently, although it suffers a bit from lackluster reporting capabilities.

Birds of a Feather

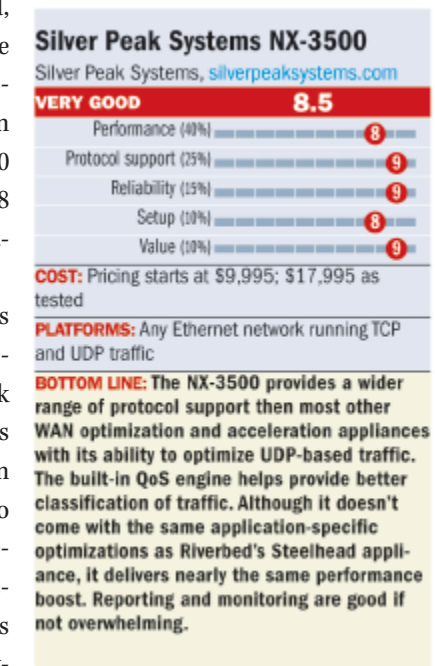
I ran the NX-3500 against the same suite of tests I used to evaluate the Riverbed Steelhead 3010 WAN accelerator (infoworld.com/3661). These tests involved FTP, CIFS (Common Internet File System), and Exchange traffic generated by MJ Net's Macro Scheduler. The only difference



in my test bed was the upgrade to the latest Shunra Virtual Enterprise WAN simulation appliance. Overall, Silver Peak's performance was on par with Riverbed's, with most results within a few percentage points, but generally slightly lower. Depending on the traffic type, performance increases ranged from double to nearly 30 times that of nonoptimized traffic.

Setting up and installing the NX-3500 was not overly complex; my test setup was operational in less than 45 minutes, with much of the time spent on tunnel definition. A tunnel is made up of an ACL and the various optimization settings such as compression, application acceleration, traffic classification, and QoS. The ACL allows admins to specify which traffic and applications to optimize, as well as which IP addresses should, or should not, be optimized. To protect data while in flight, the NX-3500 supports AES128 hardware-encrypted IPsec tunnels.

Silver Peak's caching technology, called Network Memory, stores byte segments on the appliances to help speed up repetitive file transfers. It inspects traffic flows look-



ing for pattern matches down to the byte level. As I saw with Steelhead, if a file was passed through via FTP one time and then sent out as a renamed e-mail attachment the next, the NX-3500 would recognize it and only send the changed data.

The NX-3500 includes CIFS and TCP acceleration, much like Riverbed's Steelhead, but it doesn't offer any application-specific optimization for MAPI, HTTP, or SQL traffic as Riverbed does. Moreover, unlike Steelhead, it does not perform read-ahead of MAPI or SQL requests, but it will still apply its overall TCP optimization to the traffic to help reduce network chatter. Also missing is any type of WAFS (wide area file services). Silver Peak doesn't plan to build WAFS into this platform.

UDP for You and Me

Silver Peak does include a feature that separates it from Riverbed. The NX-3500 can accelerate and cache bulk-UDP traffic, such as NFS — commonly found in Unix and Linux systems — and Veritas Volume Replicator.

Also benefiting from UDP support are real-time voice and video, as well as interactive traffic, such as that of Citrix Presentation Server. UDP traffic is much harder to optimize than TCP because it is usually very sensitive to latency and has no built-in "resiliency" like TCP does. The NX-3500 includes FEC (forward error correction) to help protect sensitive UDP flows by supplying its own error correction to the data streams.

The NX-3500's reporting is solid, providing tunnel and application graphs and charts showing real-time and historical trends. The Application reports display WAN usage for the top 10 applications as they pass through the appliance. I was able to change the time frame for viewing the reports although I found it limiting. For example, although I could choose a period of "last week," I could not pick a specific date range. Also, in this release, reports cannot be exported to any other format, such as Excel, PDF, or CSV. Silver Peak did tell me about a global management system coming later this year that will help with collating reports and centrally managing device configurations.

Even though the Silver Peak solution is one of the newest on the market, it performs like a veteran and doesn't

leave out any must-have features. The UDP support is new to this space and allows for even more flexibility in defining an overall optimization scheme. The reporting engine provides enough information to be helpful but could use some improvements. Overall, if you are evaluating WAN optimization solutions, here's one more to add to your list. 🦹

— Keith Schultz

Building a Better WAN

TEN YEARS AGO, THE WAN WAS THE EXCLUSIVE domain of frame-relay communication and leased lines. Today, a WAN may use anything from IPsec connections and cable modems to MPLS (multiprotocol label switching) tunneled over multimegabit networks. The methods may have changed, but the challenge remains the same: How do you make a WAN seem like one big LAN?

Simply throwing more bandwidth at the problem won't solve it. MPLS, as described in Paul Venezia's "Supercharge Your WAN" (infoworld.com/2787), can go a long way toward improving WAN performance, but the root cause of the problem lies well below the MPLS level.

Other forces are at work conspiring to rob your WAN's performance and response time; latency, congestion, chaty applications, and traffic contention all affect in how the WAN may respond at any given time. These are the dirty secrets of WAN performance that are usually swept under the rug — if they're even detected at all. Most of the time, the focus is on the size of the pipe, not on how the pipe is being used.

Size Doesn't Matter

In the world of the WAN, the size (that is, bandwidth) of the link often makes little difference in overall performance, particularly when the link is a long one ("long" being more than a few hundred miles). Part of the problem is that TCP and other protocols weren't intended to function beyond the local-network edge. "The reason why long-distance networks don't work is that the protocols weren't designed to do that," explains Dick Pierce, CEO of Orbital Data, which sells WAN-optimization appliances. "They work pretty well on

a local basis, and in some cases even short distances. But wide-area networks don't. The whole history of how this market segment [WAN optimization] developed was on that basis."

The problem is that the protocols' efficiency suffers as latency increases. Latency is based on the speed of light and the overall length of the WAN link, something we have little control over. Don't think speed of light is a factor? Just experience the latency in a satellite link. (A few years back, one could have argued that routers and switches added significant latency to WAN links, but most backbone equipment today works in the sub-millisecond range.)

Latency affects network protocols in various ways. TCP, for example, uses ACK (acknowledgement) packets to help provide reliability. By receiving an ACK from the receiving endpoint, the sending system knows the packet made it

WAN Glossary

New to WANs? This grab bag of terms will help avoid confusion.

ACK Acknowledgment packet used by TCP/IP to let the sending system know the packet arrived intact

Bandwidth-Delay Product A number expressed in either bytes or packets that is the bandwidth multiplied by the latency; used to determine proper router queue sizing, with $BDP = RTT * C$, where RTT is the average round-trip time, and C is the bandwidth of the link in kilobits per second

CIFS Common Internet File System, the file-sharing protocol for Windows networks

Latency The time, usually measured in milliseconds, that a packet takes to travel from one end of the WAN to the other

MPLS MultiProtocol Label Switching, a standard for routing traffic over an IP network so that all packets follow the same path

MTU Maximum transmission unit, the largest size of a packet or frame

QoS Quality of Service, a method of defining a certain level of performance for distinct IP traffic flows

RTT Average round-trip time, expressed in milliseconds

without any errors. But on high-latency links, waiting for ACKs chokes throughput.

Thus, latency is one of the biggest — if not the biggest — killer of WAN performance, both in response time and overall throughput. Long fat networks (LFNs) run at T1 speeds and higher, but suffer greatly from the inherent latency of the link. For most U.S. terrestrial links, the average round-trip time is approximately 150 ms, with satellite links averaging approximately 800 ms. Global links vary greatly, but it isn't uncommon to see 200 ms to 400 ms or higher RTTs (round-trip times). And increasing the bandwidth doesn't help.

In fact, due to latency, LFNs are largely underutilized. “The reason people built long-distance pipes that turned out to be empty was they were trying to get predictable application performance by overprovisioning,” Orbital's Pierce says. “Yet the inherent design of the networks — that they weren't designed for long distance — was the problem.”

Rush Hour

Congestion also affects WAN performance, of course. Congestion occurs when no bandwidth-allocation policy has been applied to traffic on the WAN. Traffic flows can be bursty, such as when one user tries to retrieve a large e-mail attachment while another user logs in to a CRM portal. With no bandwidth management, the download can bring the smaller link to a grinding halt.

P.G. Narayanan, CEO of Allot Communications, believes that much of the congestion problem can be solved by applying QoS to the traffic. “The problem most of these networks have, though, is temporary ... that second, or that minute it's congested, you can get away with just prioritizing applications. So what you can do is put a gigabit box at the central site to prioritize those applications, the critical applications, on a temporary basis, and you can avoid the congestion, and all other times you're OK anyway,” says Narayanan.

Prioritizing application flows is an important part of managing your WAN traffic, but it isn't going to solve TCP's inherent limitations when latency creeps in. On shorter links where latency isn't an issue, simply preallocating your bandwidth will help keep important packets moving, re-

gardless of what else is in the pipe. But on LFNs, latency, not congestion, is the culprit.

Talk, Talk, Talk

From the end-user point of view, latency gets less tolerable as the back-and-forth communication required for some action increases. And layer 7 protocols — where applications live — are chatty, requiring an absurd number of round-trips to complete a single task. Much like TCP, protocols such as CIFS and MAPI (mail application programming interface) were designed to run inside the LAN, not over the WAN.

The chattiness reaches a crescendo when users map drive letters over the WAN using CIFS (used in Windows networks). Any user that has had to open, edit, and save a Microsoft Word or Excel document from a remote file server knows how long this simple task can take, even over a fat WAN connection.

By the same token, users of Microsoft Outlook and Exchange 2000 suffer when they open an e-mail with an attachment over a WAN link. The message appeared to be in their inbox, but in reality it was still on the server waiting to be retrieved.

Microsoft Exchange Server 2003 was designed to mask this problem by downloading messages and attachments in the background (cached Exchange mode). Although this is great for the end-user, it adds additional traffic on the WAN. For example, Outlook now downloads all attachments to your inbox, regardless if you were going to open them in the first place. This places an additional load on the WAN link, which should never happen.

Out With the Old...

Traditionally, WAN performance was attacked at the packet level. Back in 1998, Expand Networks was one of the leaders in WAN compression. Liad Ofek, vice president of technical services at Expand Networks, says that, at the time, the goal was to “squeeze as much data as possible” into existing links.

Expand used a series of compression algorithms to reduce the number of packets on the wire. Other vendors, most notably Packeteer, also used highly advanced com-

pression schemes and began adding QoS to further allocate and manage WAN traffic flows.

File-caching provides yet another way to reduce traffic by storing a copy of recently accessed files on an appliance near requesting users. As with a browser cache, files and objects are kept closer to the remote user, helping to overcome latency and prevent excessive, redundant requests over the WAN.

This is typically a “full file” cache and not made up of smaller data segments. Full-file caching isn’t nearly as effective as newer segment-caching methods, because the chance of a second or third user requesting the same file is slim. Also, if the file on the file server is renamed or changed, then it won’t match the file already in cache and must be transferred again anyway.

...In With the New

In recent years, TCP acceleration has taken center stage as one way to improve performance by reducing ACKs and

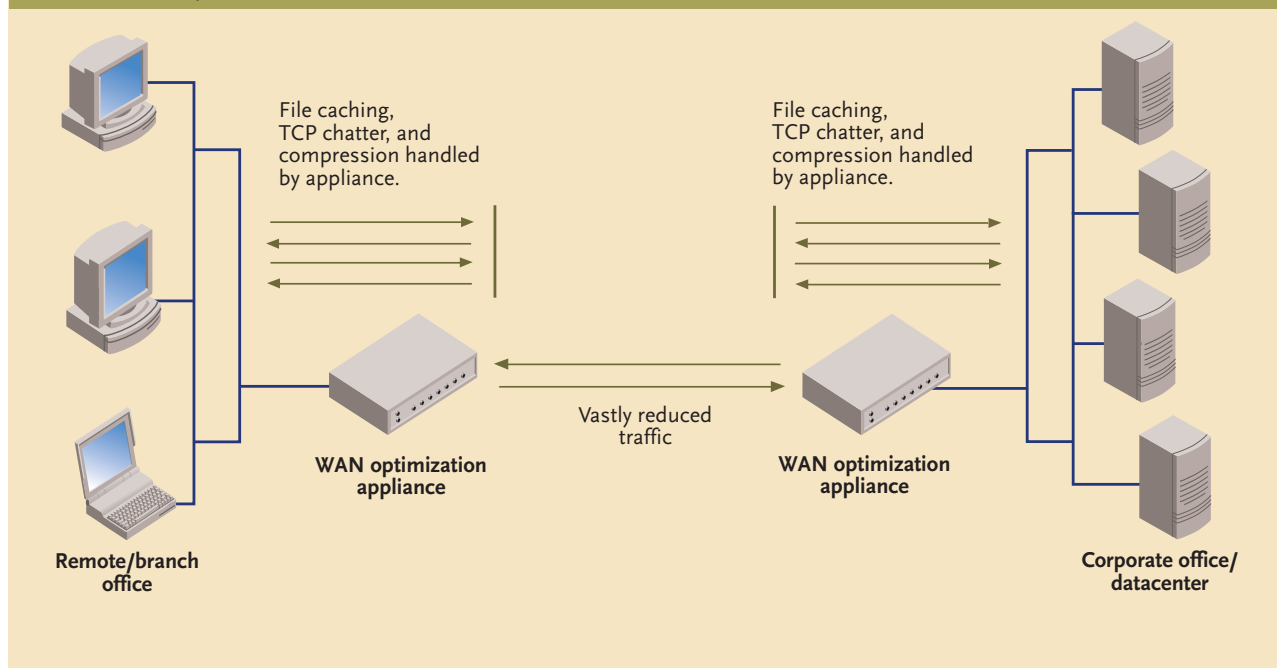
playing games with the TCP window size. Vendors such as Swan Labs, Peribit (now owned by Juniper Networks), Expand Networks, and Riverbed Technology have all developed solutions based on improving TCP’s performance.

One of the most effective methods is to handle TCP ACKs locally, using an appliance. The appliance bundles multiple ACKs into a single request, thereby reducing the delays caused by high latency. To the application requesting the data, it receives an ACK just as it expects to, except the ACK comes from the local WAN appliance and not from the far side of the WAN.

The next step beyond TCP tricks is application-specific acceleration. Some WAN optimization vendors use plug-ins in their appliances to help improve application response. Applications such as DNS, Exchange, FTP, Citrix, Notes, and CIFS/NFS can all benefit from reduced chatter on the wire. The plug-ins work much like the TCP ACK optimization in that they handle redundant requests locally instead of sending each one.

Feels Like Home

WAN optimization appliances use caching to localize file access and compression to reduce congestion. They also simulate TCP acknowledgments locally and reduce chatter for specific applications. Appliances may specialize in one type of optimization or consolidate multiple functions.



There Is No Quilt

The WAN optimization and acceleration space is heading toward a convergence of sorts. In the past some vendors specialized in a single technology solution, but now they're adding technologies to solve other pieces of the WAN problem. Orbital's Pierce sees the multiple approaches to solving WAN issues as "patches, in the context of patches and a quilt. In the end, it's about the quilt; it's not about the patches themselves. Customers buy patches today because there is no quilt." The trend is for vendors to move from "point" solutions to a more comprehensive managed system.

Several WAN appliances include compression and TCP acceleration along with file-caching and application-specific acceleration. But not all vendors agree that such consolidation is wise. "I think more customers are more worried about just the visibility into the network," says Allot's Narayanan. "They want a good traffic-management company with the ability to decode any application layer properly, not falsify it."

Other vendors, such as Swan Labs, Riverbed, Disk sites, and Juniper Networks, are banking on single-box solutions. Tom Tansy, vice president of Marketing at Swan Labs, sees a further consolidation of technology. He believes many customers are suffering from a "box proliferation problem" and will want to roll out a single appliance instead of many disparate solutions.

Either way, when it comes to speeding up WANs, everyone agrees that more bandwidth alone is not the answer. As long as TCP remains unchanged (and for now it has to) and the speed of light governs latency, boosting WAN performance will require tricks at the protocol level, combined with traffic-flow prioritization and application-specific packet reduction. WAN acceleration solutions will continue to evolve to include multiple techniques for getting most out of your link, at least until we find a way to send data faster than the speed of light. 🚀

— Keith Schultz

Shaking Off the WAN Blues

The WAN optimization techniques you choose depend on the challenges. Today, many vendors bundle multiple techniques in one appliance, which may be stronger in some areas than in others — so it pays to evaluate products carefully.

Scenario	Solution
Short distance link (low latency)	QoS to provide expected response times
Limited bandwidth	Compression between sites
Voice over IP	QoS to prioritize voice traffic
Long link (high latency)	TCP acceleration to reduce ACKs
Layer 7 applications	Application-specific acceleration to reduce program chatter; file and segment caching to keep recently used data near remote users

Silver Peak Squeezes Into the WAN Game

AS THE WAN OPTIMIZATION MARKET CONTINUES to grow and evolve, new and established vendors are developing systems that accelerate previously ignored varieties of network traffic. One promising newcomer is Silver Peak Systems, whose breadth of application support impressed me in a recent demo.

Silver Peak showed me how its NX-3500 WAN acceleration appliance can improve response time over a wide range of enterprise applications that utilize bulk TCP (CIFS, MAPI) and UDP (NFS, TFTP) as the transports. Other applications supported are real-time and streaming data (such as VoIP and video), Citrix, Microsoft Remote Desktop Connection, and SQL. Silver Peak offers three models suitable for branch offices and large datacenters. WAN capacities range from 2Mbps to 155Mbps of bandwidth capacity and as much as 2TB of raw local storage.

Much like Riverbed's Steelhead (infoworld.com/3661), the Silver Peak solution fingerprints byte segments and stores them in each appliance's local data store. When a user requests a segment that has already traveled the WAN, the data is fetched from the local cache — only original content must be sent over the link. During the demo, I watched the NX-3500 reduce a PowerPoint file copy from a few minutes to seconds, regardless of the underlying transport.

Silver Peak has implemented hardware-based AES encryption to protect the data store, and secures communications between NX appliances using IPSec and 128-bit AES. Also impressive is Silver Peak's reporting, which allows traffic data to be sliced and diced many different ways.

The WAN acceleration short list just got a little longer. 
— Keith Schultz

Data Security as a Service

FOR ALL THEIR DIFFERENCES, SMBs AND ROBOs (remote offices/branch offices) have one unavoidable headache in common: designing a robust backup and recovery system at a justifiable cost.

When backing up network assets, larger, centralized organizations typically employ expansive — and expensive — automated tape systems. Although such backups may go to disk first for performance reasons, almost all end up on tape. An off-site vault provider then maintains copies of backup tapes in case of disaster. To meet recovery requirements for important applications, some large-scale enterprises tap more advanced methods, such as replication or CDP (continuous data protection).

SMBs and ROBOs rarely have the luxury, however, of duplicating big-time backup schemes on a small scale. Typically, they lack the administrative and operational expertise, the capital for tape hardware, or the money to pay an off-site vault company month after month.

The unfortunate result is that many small offices do not back up their data at all — or they use an inexpensive system fraught with design flaws and operational challenges, such as a single tape drive that performs a full backup every night. Those tapes typically stay on site and in many cases sit inside the backup server, allowing a single break-in or fire to destroy everything. Worse, lack of oversight may mean that backups are routinely falling under the radar — until a failed attempt at restoring them gets somebody fired (infoworld.com/2594).

SMBs and ROBOs know they need backups that work. They just can't perform them affordably and reliably. What's needed is the equivalent of an AOL for backups — click OK, pick a screen name, and make backups happen. But the relatively slow connections typical of SMBs and ROBOs mean that conventional backup

schemes, in which one change to a huge file results in that entire file being backed up, must be replaced by more intelligent, incremental schemes.

Backup on a Human Scale

Vendors such as Asigra, Avamar, Connected, EVault, and LiveVault offer products and services that enable administrators to perform advanced incremental backups with point-and-click ease. All allow you to load their software onto your environment, which you then back up to a remote vaulting service via the Internet. And they all encrypt the data for security reasons.


Administrators can select individual drives and directories as well as certain file types to include or exclude. Most offerings support auto-discovery, allowing you to back up all drives on the system automatically, without having to update the software every time you add a new drive or file system. LiveVault and Connected enable you to manage their products via the Web, whereas the other products are managed by software loaded onto your environment, such as a Windows workstation. Some also have Java consoles that can be installed on other platforms.

In most cases, you need to install an agent on each machine that is to be backed up. With Asigra's software, however, you select one system in your environment to be the "ds-client," which then communicates automatically with all systems in your environment using a variety of protocols, including SSH, CIFS, or NFS. It even performs hot backups of databases using this approach. Asigra doesn't charge for its ds-client or database agents; it bills only for the amount of data you're protecting.

Asigra also provides the broadest platform support, as its agentless model not only supports major Unix platforms but any platform that can export an NFS or CIFS share.

Long-Distance Backups

A remote backup architecture can be laid out with or without a local recovery server, depending on budget and the performance and reliability of the WAN connection.



Second in terms of platform support is EVault, followed by LiveVault and Connected. Most products and services provide flexible backup scheduling, allowing customers to perform backups every hour, every minute, and so on. CDP — in which a file is backed up automatically as soon as it is created or changed — is currently offered only by LiveVault, although Asigra says it has plans to support CDP in the future.

A New Backup Paradigm

Backup services add a remote wrinkle to a familiar architecture: There are clients to be backed up, a remote recovery server, an optional tape archive, and optional local recovery server. Client software is installed on the systems to be backed up, allowing backups to either the local or remote recovery server. If stored locally, backups are automatically replicated to the remote recovery server, which may be owned by a BSP (backup server provider) or by a large enterprise that wants to maintain the process.

Companies usually start out by backing up directly to a BSP, minimizing the capital outlay — no servers to buy or maintain. The charges are based solely on the number of gigabytes stored per month at the BSP. The downside, however, is that all data is remote. Small files can be restored remotely; large restores, however, require the BSP to cut a CD, tape, or portable disk and ship back it to the customer.

A more sophisticated backup methodology involves

obtaining a local recovery server — which can provide quick restores of large systems — and then replicating backed-up data to a BSP for disaster purposes. This model gives the customer exactly the same level of data protection that an enterprise datacenter does, but for a fraction of the cost. The local-recovery-server option is available from Asigra, Avamar, EVault, and LiveVault and ranges in cost from free (Asigra) to tens of thousands of dollars,

depending on vendor and data volume.

Companies with considerable backup volumes may eventually grow disenchanted with per-gigabyte monthly fees. Such companies should consider purchasing a remote recovery server from their BSP and managing it themselves. Even a small business can do this by putting the server at a collocation facility. All of the vendors covered here offer this option.

Make or Break Decisions

Remote backup providers have gone to great lengths to develop features that minimize bandwidth and capacity demands. Take LiveVault's delta restore feature. The software knows which blocks of a file have changed since the time you asked to restore it, so it only needs to send those blocks back to the client to reassemble the file. That can save a lot of bandwidth when your file is corrupted and not deleted.

Because you're paying for what you're storing, it's also important to consider what each backup offering does to eliminate redundant data. The solutions from Asigra, Avamar, and Connected eliminate redundant files in the vault. If you have the same spreadsheet on three different systems, for example, these products ensure it is stored in the vault only once. If you have a high amount of redundancy in your environment, this can save a lot of money.

Another significant data point is the number of protected

terabytes — that is, the total size of the customer data protected by the solution. In this area, Asigra is the clear winner; it claims its BSPs protect more than 3 petabytes. That makes Asigra the best-kept secret in data protection, probably because its software is usually rebranded.

The final word on these products comes from those who use them. When he first looked at his Avamar system, Steve Merkel of Data393 says he was certain he was “seeing things.” He had been performing virtual, full backups every night on 65GB of data but noticed only 0.05 percent of the data was going across the wire. In the end, however, his “six-month testing cycle” proved what he was seeing was correct.

“Everyone would be doing backups like this” if they knew how easy and cost-effective it is, adds Tim Hannibal, who works at VaultLogix, an Asigra customer and service provider.

Such offerings show that you don’t need to buy a \$20,000 tape library and sign a large contract with an off-site vault vendor to have automated backups. You just need to install some software, pay a monthly fee to a BSP, and go worry about something else for a change.

About 5 petabytes of data are being backed up by electronic vaulting services today. Although that number just scratches the surface of all the data out there, it also represents millions of happy customers and untold thousands of successful restores. Based on the evidence, a sizable portion of low-volume remote backups are very likely working better than those in big datacenters. ↻

— *Curtis Preston is vice president of the Glasshouse Technologies consultancy*

Boosting Network Performance: 5 Tips

THE TROUBLE WITH PERFORMANCE BOTTLENECKS is that they can be tough to identify. And when performance riddles remain unsolved, IT management may find itself faced with a Hobson's choice between admitting ignorance and making up excuses. Drawing on our years of sleuthing and experimentation, we've collected a few of the most likely ailments – and suggested remedies.

The Woe of the WAN

Think you need to reclaim WAN bandwidth? You can easily spend a bundle on traffic-shaping appliances or caching engines in an attempt to rein in WAN bandwidth utilization. But what if it's not the pipe?

First things first: Before you buy anything, get a solid idea of what traffic is crossing the WAN. Network analysis tools such as Ethereal, ntop, Network Instrument's Observer, or WildPacket's EtherPeek NX can give you a fresh look at what's really on the wire.

You may find that replication times for your Active Directory are set far too low and simply configuring longer replication intervals can buy you breathing room during the workday. Are some users in remote locations mapping shares to the wrong servers and pulling large files across the WAN without realizing it? Are the vestiges of a long-disabled IPX network still floating around? Some WAN problems boil down to application misconfiguration, where traffic is directed across the WAN when it should have stayed local. Regular reports on WAN traffic patterns will save money and headaches.

Let's Play Nice

All too often, applications, Web services, and Web sites from multiple departments across the enterprise compete for server resources. Although each of these components may be well-tuned in its own right, an application from another department that is also using the same production

clusters may have a poorly tuned query or some other issue, which in turn affects your users or customers.

In the near term, all you can do is work with your system administrators and the department that is having the performance problem to obtain resolution for your users or customers. Longer term, create a community across all of the departments that use the production clusters where your objects are deployed. Work across teams to ensure that there is adequate funding for a staging environment that is truly representative of the mixed workload production environment. Ultimately, you'll want to develop a series of benchmarks that can be used to validate mixed workload performance in the staging environment.

Caching, Shaping, Limiting, Oh My!

If your WAN is truly undersized – and you can't afford a long-haul frame-relay network – traffic shaping and caching can help unclog the pipe.

Traffic-shaping configurations are more art than science. Prioritizing apps is often more political than technical but may have tremendous effects on perceived network performance.

Caching is a different beast altogether. It requires less work than traffic shaping, but the impact will likely be smaller. Caching engines store and serve up local copies of commonly accessed data to reduce WAN traffic. The downside is that dynamic content can't truly be cached, so e-mail won't enjoy the same performance bump.

Reining in Mirroring and Replication

Slowdowns often plague enterprises that use mirroring or replication for high availability or disaster recovery between locations. If you have many locations or many database tables – or a lot of transactions or journaling that needs to stay in sync between multiple locations – watch out, because the performance loss can be dramatic.

If possible, run your mirroring and replication activity

across separate WAN “pipes” to keep it isolated from production traffic. Your network design team can assist with a viable topology. At the same time, go over the configuration (star topology, for example) you intend to use to support mirroring and replication. Vendor representatives and the network design team can provide useful input on constructing a configuration that will prevent network saturation.

Configuration aside, mirroring and replication products – such as XOssoft’s WANSyncHA Oracle or High Availability Linux Project’s Heartbeat – usually provide options to control timing and traffic flow between sites. Some products enable you to schedule syncing activity during off hours, whereas others let you activate syncing tasks only if a particular threshold is reached. Use those controls; that’s what they’re there for.

Estimating Future Speed

One of the trickier tasks is projecting infrastructure requirements in the face of constantly evolving business demands. Most people start with capacity planning tools, such as TeamQuest’s View and Modeling products. No tool alone, however, can accurately predict the nature of tomorrow’s workload.

A little sleuthing is needed. Put on that detective hat in a sandbox, staging, or fail-over environment and do some proof-of-concept work.

Create a virtualized, grid-type environment and replicate a representative sample of workload. From this exercise you’ll obtain and use numbers to extrapolate the overall change in infrastructure performance if the topology were to undergo a wholesale change. Execute baseline performance tests that detail today’s workload. Then execute load tests that mimic the addition of the expected number of transactions in six months or an increase in the number of users over time.

Even if testing resources are not as plentiful as those found in production, you can still run tests that mimic the type of infrastructure changes required to meet business demands and get reasonably good numbers. With this information you should be able to project when to add resources to keep performance on track. 🦋

— Maggie Biggs and Paul Venezia

Six Steps Toward Disaster Recovery

I RECENTLY GOT TO WRITE A FUN PIECE FOR *InfoWorld* called “Stupid User Tricks” about protecting your network from human error (infoworld.com/4534). Researching the article revealed to me how many variables folks tend to miss when running a network, as well as when planning to protect and recover that network.

I suppose some of the errors I encountered during my research are more surprising to us consultant types because we live and breathe best practices. We live it, we breathe it, we get to install and bill for it, and then we get to walk away and do it all someplace else. Day-to-day systems administrators live and breathe a just-get-it-done philosophy, and they can't walk away.

So in that spirit, I've condensed some of the disaster-recovery best practices into a top six list. Make sure you've got these six points covered, and you're much more likely to survive not only stupid human tricks but any kind of network disaster curveball Lady Fate may decide to pitch your way.

Test Your Backups

This is first because it was by far the most popular entry. Someone installs a tape drive setup, installs the backup software, and schedules daily, weekly, monthly backups. Something happens a year later, and it turns out nothing's actually been running.

Backups are boring, I know. Not to mention mind-crushingly tedious. But if you don't have them when you need them, you're done.

So do a test backup and restore after installing a new backup system. Then – and this is critical, not optional – do a test restore every week. That's right: every week. Not the whole tape, just a specific subset of folders. Shouldn't take more than 15 minutes, and it can save your professional career in a crisis. Just do it.

Spend a Little Money on Your Backup Software

Don't just buy Bob's Basic Backup package because it's cheap or came with the tape drive. Spend some bucks here. Make sure the thing can support dynamic backups; also ensure that it can support individual folder and file restores. Take a step back and think about investing in a disk-heavy server to act as a disk-based backup between the tape drive and the network. Many of the better packages, including those from CA, IBM/Tivoli, and Veritas can manage this NAS-type device as well as the backup, which means not only safer data but much faster restore times. And the cost really isn't that huge.

Store a Weekly Copy Off-Site

This was the next most popular entry for the article, even though it didn't get much play in the finished article. If you're worried about recovering data should the office building burn down, then keeping all the data in the office building isn't all that bright.

Explain all this to your tightfisted boss using small words. Get a safety deposit box or a secure business storage locker and bring your tapes there. One tape a week'll do you. Likely this is a quick 30 minutes out of your day door-to-door. Look at it like this: It's less desk work.

Block off Access to Servers

If your business runs on its server applications, they shouldn't be accessible to just anyone – including cleaning people. Put them in a room. Get ventilation. Think about things like sprinklers (bye-bye servers) versus Halon systems (servers live), UPS protection, a building-based power generator, and maybe even a Webcam-based monitoring system, such as the NetBotz system from APC. Know that room is safe, and know what's going on inside it.

Then add this new thing to the door called a lock. Make sure only you, the IT staff, and a responsible member or two on the executive team have the clearance to open this lock. If the cleaning people need to get in there, open the door for them and show them where they can plug in their gear.

Map Out a Plan for What Happens If the Office Building Burns Down

We're talking worst-case scenario day here. There are oodles of options in this department, so I'm not going to try and list them all here, but do decide if your business can shut down due to one of these occurrences or if it needs to recover somewhere else right away. And how quickly it needs to recover. Then figure out what it needs in order to recover. You should also make sure you can deliver all these requirements in time. Yes, this is a lot of work.

Write All of the Above Down and Title It "Disaster Recovery Plan"

Put gold and red star stickers on the cover, then put your name, your IT staffer's names, and some executive manager names on it, and make sure it's distributed to everyone who needs to see it. Then show it to everyone who doesn't need to see it. Make sure the section on what employees should do if the building blows up gets to the worker bees.

My snide tone and I are making all this sound obvious, but both the recent stupid-user article and my wide IT travels have continually shown me throngs of people that keep putting these things off or simply ignoring them altogether. Yes, getting all this done is a month or more of real work. But having it in place when Godzilla steps on your server room: priceless. 🐞

— *Oliver Rist*

Effective Long-Distance Data Protection

PROTECTING DATA PROPERLY IS CHALLENGING IN ANY circumstances but can be even more difficult to do at a remote office. It's easy to understand why: Most data-protection tasks require both human labor and the computing power necessary to move large amounts of data, digging into two resources that are typically in short supply at a remote office.

To minimize the effects of WAN latency, you may choose to store branch data locally, a strategy that ensures easy, reliable access for branch workers but makes life harder for admins. Nobody likes to worry about critical backup jobs at a remote office getting stuck in the middle of the night.

A centralized data structure — where all the major databases and file systems are stored and protected in the glass house, while branches access the data remotely — consolidates control and reduces the operational burden. The shortcomings of this approach may include unacceptably slow application response times; further, connection problems will have a noticeable impact on business activities.

How do you find a balance between prompt, reliable user access and foolproof data protection for remote offices? Luckily there are numerous tools and technologies to choose from. Your company must find a balance between centralized and distributed data that meets the needs of the business.

For example, leaving data locally at the branch office and contracting an online, hands-off backup service could be an acceptable solution if remote access to branch data is an occasional or infrequent event. Some service providers offer a consolidated view of all remote-office backups, which simplifies monitoring and enforces centralized control.

By contrast, if interoffice data access is a frequent occurrence, you may be better off choosing a solution, such as a set of WAFS (Wide Area File Services) appliances, that mirrors remote data at a central location and keeps it carefully synced to minimize access delays.

Naturally, the solution you choose should support the range of applications and the business process flow inside your company. For a Windows environment that does not require infinite restore points, Microsoft DPM (Data Protection Manager) offers good automated data protection of files at remote offices.

However, companies with a decentralized IT structure — think of multiple datacenters serving remote users at various locations — may have requirements for fail-over, performance, and multiple mirroring instances that neither a WAFS solution nor a Microsoft DPM is capable of supporting.

For those demanding environments, a solution based on YottaYotta NetStorage control nodes can provide seamless data sharing at the block level — as opposed to the file level — and create what's essentially a resilient, distributed SAN that users and applications can access quickly and transparently from any location.

Implementing a foolproof data-protection scheme for your remote offices is not a luxury but a necessity. And the goal is within reach. If you understand how and when the data and files flow at your branches, you can find a solution that is safe and sane for both users and admins. It might even give your company that additional competitive advantage that makes all the difference with customers. ↪

— *Mario Apicella*

For more information, contact:

Silver Peak Systems
471 El Camino Real
Santa Clara, CA 95050
www.silver-peak.com
US: Phone: 408-935-1800