


Realtime
publishers

"Leading the Conversation"

The Shortcut Guide[™] To



Optimized
WAN Application
Delivery

sponsored by

Blue  **Coat**[®]

Ed Tittel

Chapter 4: Challenges and Best Practices for WAN Optimization	69
Benefits of an Improved Application Infrastructure	69
Holistic View of Entire Application Delivery Environment	70
Observation of Application and User Pairings	70
Organizational Priority and Policy Compliance	71
SSL Interception, Split Video Streams, and Local Object Caching	71
Mobile and Roving User Management	72
Application Awareness and Real-Time Results	73
Symbolic Dictionaries and Sent References	74
LAN-Side Optimization Impacts WAN Delivery	74
Proxies and Tunneling	75
Special Considerations for Disaster Recovery	75
Fixing Broken Protocols	76
How Transactions and Conversations Can Be Short-Circuited (and Fast-Tracked)	77
Enhanced Compression and Transaction Optimization	77
Achieving Efficiencies: WAN Optimization Usage Cases	78
Extending Inmate Tracking Systems for the State of Nevada	78
Improve Inter-Campus Communications at a Multinational University	79
Networx Australia Accelerates Application and Content Delivery	80
Conclusion and Recommendations	81

Copyright Statement

© 2008 Realtime Publishers, Inc. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtime Publishers, Inc. (the “Materials”) and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtime Publishers, Inc or its web site sponsors. In no event shall Realtime Publishers, Inc. or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtime Publishers and the Realtime Publishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtime Publishers, please contact us via e-mail at info@realtimepublishers.com.

[**Editor's Note:** This eBook was downloaded from Realtime Nexus—The Digital Library. All leading technology guides from Realtime Publishers can be found at <http://nexus.realtimepublishers.com>.]

Chapter 4: Challenges and Best Practices for WAN Optimization

WAN acceleration is a key enabler of strategic IT initiatives and enterprise goals, including branch office networking, central storage repositories, and business continuity planning. WAN connections and delivery may be established using dedicated leased lines or cloud services that are owned and operated by providers and shared by multiple subscribers. Furthermore, the diversity among protocols, platforms, and performance rates adds layers of complexity to traffic optimization for network engineers and infrastructure implementers.

Ultimately, the real meat of the WAN optimization discussion hinges on application delivery. This topic is probably best understood as a form of optimization that incorporates a deep and detailed understanding of application traffic patterns, network messages, and overall behavior. This approach provides a foundation for optimization techniques that includes smart use of proxies, protocol optimization, application behavior optimization, and more. It makes effective use of the techniques and technologies described in the previous chapter, particularly compression, various levels of caching, and streamlining of network traffic behavior to make the most of the WAN bandwidth available for application use. Above all, with a serious focus on delivering the best overall end-use experience, optimization and its metrics can improve productivity and usability. In turn, this depends on creating and maintaining an optimal application infrastructure, designed to improve response time and to deliver a positive application access experience.

Benefits of an Improved Application Infrastructure

The benefits of proxies and WAN optimization devices in the enterprise network landscape are numerous. A centralized outlook on network transactions creates a holistic view of the application delivery environment and lets IT observe application-user pairings including applicable content. A proxy coordinates with service priority and security policy requirements to allow or deny transit, then to prioritize all allowed traffic. WAN optimization incorporates such functionality and makes use of protocol optimization (where applicable) and various levels of caching to minimize the size and frequency of networked communications.

These are only a few of the immediately appreciable benefits of utilizing a WAN optimization architecture that includes improved application delivery, proxy accelerator, and security appliance capabilities at the network boundary. Let's revisit a few of these and other key points for an individual perspective on WAN optimization benefits.

Holistic View of Entire Application Delivery Environment

The primary approach to acquiring a complete knowledge of the application delivery environment is to ascertain all traffic types and content that passes through some central or shared access point. A centralized WAN optimization device provides comprehensive visibility into and control over information about network protocols and payloads of interest. Such a device serves as a single gateway through which all monitored traffic passes, and offers a single source viewpoint from which to observe end-user and system-to-system network interactions.

WAN optimization devices are designed to peer into network traffic and—in some cases—perform granular byte-level optimization on application and protocol traffic. Traffic acceleration is facilitated by predefined proxies, protocol optimizations, and traffic filtering/priority schemes—each with individual parameter tweaks—for many types of common traffic including instant messaging, SOCKS servers, Telnet, and DNS or through customized HTTP and CIFS proxies. SSL/TLS traffic—which some optimization devices will not touch—can also be selectively decrypted, optimized, and re-encrypted according to administratively defined policy and procedure.

Because of its more holistic perspective on network traffic and activity, this comprehensive perspective into network application and protocol payloads provides significant performance enhancements, along with greater granularity in control. Contrast this with individual end-point solutions (such as firewalls, Quality of Service—QoS—appliances, and routers), which provide varying levels of granularity and inconsistency or incompatibility because their perspective on network traffic is invariably obtained from what flows between pairs of endpoints, not from end-to-end across an entire network infrastructure.

Observation of Application and User Pairings

WAN optimization device presence permits IT staff to monitor and observe interactions between end users complete with contextual content and application data. Network staff can analyze traffic trends through information collected, analyze report data, then tune performance and security parameters accordingly.

Where unwanted or unauthorized protocols are discovered, it becomes easy to block them completely or to throttle back available bandwidth to a mere trickle (5 to 10Kbps). This latter approach has the advantage of keeping such connections alive so that network administrators can identify the parties (or at least the computers) involved in such activity. This provides an opportunity to identify the individuals who may be violating acceptable use policies and to offer counseling or take appropriate remedial action as such policies require.

Organizational Priority and Policy Compliance

Preservation of confidentiality, privacy, and security are reasonable expectations of any WAN accelerator that is likely to pass and process sensitive or proprietary data and documents. Accordingly, WAN optimization solutions should strictly adhere to organizational priorities and policies when handling data.

By way of example, MS SharePoint services marks select content as private and non-cacheable, which defeats some of the optimization strategies in a WAN optimization device. Newer document formats utilize Open XML specifications that cannot be further compressed or differenced by conventional means. The Open XML format uses compression, and subsequently inhibits any further WAN optimization and document differencing strategies.

Bottom line: organizations must determine the applicability and serviceability of such applications over the WAN and assess related performance issues. For SharePoint services and Open XML specifications, some WAN optimization devices have taken to blueprinting application protocol behaviors to overcome inherently prohibitive factors. In the case of Open XML, certain WAN optimization solutions can and will decompress, decompose, and determine differences within document formats for data reduction and caching purposes, thereby ensuring optimal delivery. This is a similar approach to that which sometimes applies to encrypted traffic described in the next section, where a WAN optimization device is granted the ability to peer into and operate on content inside already compressed and quantized document representations, enabling it to optimize the data it finds based on caching, protocol optimization, and so forth.

SSL Interception, Split Video Streams, and Local Object Caching

Communication via WAN often involves cryptographic protocols necessary to protect sensitive information in transit. If encryption occurs upstream from the WAN optimization device, special actions are required to handle such data at the appliance: it must decrypt that traffic, optimize it by all possible means, then encrypt that traffic (again). Otherwise, the WAN optimization device has no visibility into this traffic, and cannot provide optimization services. The same is true for compressed data streams as well. That explains why in most cases encryption or compression is best leveraged on the WAN optimization device itself.

Because it is guaranteed to be party to all relevant network traffic, a WAN optimization device is the only practical platform suitable for intercepting encrypted SSL transactions, split video streams, and for caching local objects. On a distributed scale, WAN optimization device intervention is also the best possible way to intercept, optimize, and continue network transmissions for multiple clients in a controlled and unified manner. Using such processing elements on the WAN can apply hardware-based SSL acceleration to increase network-driven encryption response time and offload end-point processing burdens related to cryptographic processing.

Some appliances take a single video stream and divide it locally into a sufficient number of streams to service all participating viewers. They can also record live video for future playback and even identify individual end users to track employee viewership trends.

Mobile and Roving User Management

Additionally, a WAN optimization device provides the operational oversight needed to manage roving or mobile users. Roaming employees operate outside the WAN perimeter, which is normally beyond the reach of organizationally defined security policies. Managing and securing remote users' access to organizational resources thus turns into a special-case situation.

In a typical WAN topology, stationary acceleration devices operate at corresponding ends of a WAN link. The problem with a mobile computing workforce is that these appliances cannot accompany those who roam beyond the boundaries of the WAN. Instead, some WAN optimization products utilize mobile client software to approximate LAN-like performance across public telecommunication media, as depicted in Figure 4.1. Client software reproduces WAN accelerator functionality on the mobile platform to reduce application and protocol latency, accelerate file sharing, and expedite email exchange. Where specific applications must be supported and software accelerator components are available, they too can be incorporated into this kind of “virtual WAN optimization end-point” architecture.

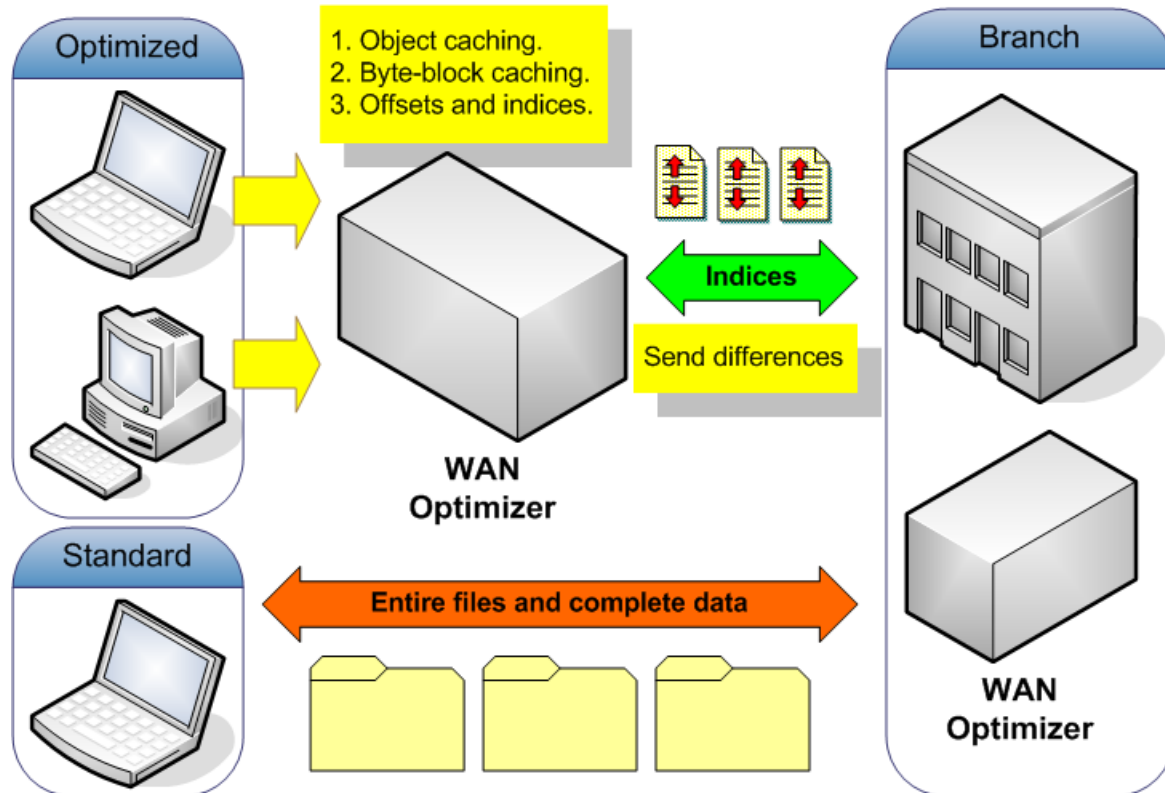


Figure 4.1: What WAN optimization devices do for WAN links, client software does for individual remote access.

Application Awareness and Real-Time Results

WAN accelerators capable of handling applications and protocols directly can deliver better optimization than those that look only at traffic at the packet, byte, or bit level. Of course, this is a primary motivation for building and maintaining application optimization and delivery mechanisms. Furthermore, you can configure and manage optimization appliances in much the same way as routers, using either command-line interfaces (CLI) or Graphical User Interfaces (GUIs).

Certain implementations require enterprises to apply configuration changes to routing devices to redirect traffic into the optimization appliance. Thus, all relevant routing protocols, load-balancing acts, and asymmetric traffic flow must be considered during the design phase, prior to deploying WAN optimization hardware. However, unlike a router, a WAN optimization appliance operates at up to the application layer (which is what endows it with a quality we call “application awareness”) and therefore involves both network engineers (who usually focus their activities at OSI layers 1 through 3) and IT administrators (who work primarily at OSI layers 4 through 7).

Altering client-server behavior can produce complex, sometimes unintended and unforeseen consequences that can be either harmless or disastrous, depending on the nature of the network transactions affected. Incremental performance benefits may be earned through application-specific latency optimization techniques that include acceleration of HTTP, SSL, and SQL traffic, and using pre-fetching techniques to anticipate upcoming content requests.

WAN accelerators can mitigate and mend broken protocols and interrupted communications on the fly without requiring manual intervention. As an example, forward error correction permits recipient WAN accelerators to detect and correct erroneous transmissions without having to reissue original data transmissions. Packet order correction handles packets delivered out of sequence and transparently rearranges them in proper order. Likewise, WAN accelerators can deliver and manage application data in a reliable, repeatable manner, then store current application data values locally to permit repeated access to such information without requiring repeated WAN traffic to match.

In terms of WAN optimization, application awareness is essential to accelerating robust business platforms. To illustrate, a WAN accelerator utilizes built-in blueprinting mechanisms for commonly utilized business applications (including Oracle eBusiness Suite, Microsoft Office System, and Microsoft SharePoint). Therefore, the appliance must learn only local transaction patterns, and when the time to transmit across the WAN comes, it needs to push only changes and differences instead of entire documents and datasets.

Fast differencing algorithms utilized to accelerate WAN traffic also exercise application awareness. These functions store individual historical records for each object an application uses. Each history can be separately aggregated and versioned. This way, differencing occurs only against correlated content.

Security is an issue primarily in two situations: data at rest and data in motion. Data at rest is any information stored on the WAN accelerator and must therefore comply with any applicable organizational and federal regulations governing the storage of private and confidential data. At the same time, data in motion—anything sent across the wire—must also be securely encrypted where applicable. For these reasons, encryption must necessarily occur for drive partitions and data pathways whenever sensitive information is involved. This data security must also be backed up by proper access control and user authentication systems.

Symbolic Dictionaries and Sent References

WAN optimization devices can use local symbol dictionaries and effect quick, efficient pattern lookups to eliminate redundant transmission of objects or data sequences (see Figure 4.2). This way, individual end users never need to send or receive complete communications from end to end; instead, transmissions are sent with placeholder references to defined symbolic “keywords” contained in each end-point symbolic dictionary. This is a direct consequence of the dual data caching and data reduction strategy built-in to basic WAN optimization device behavior.

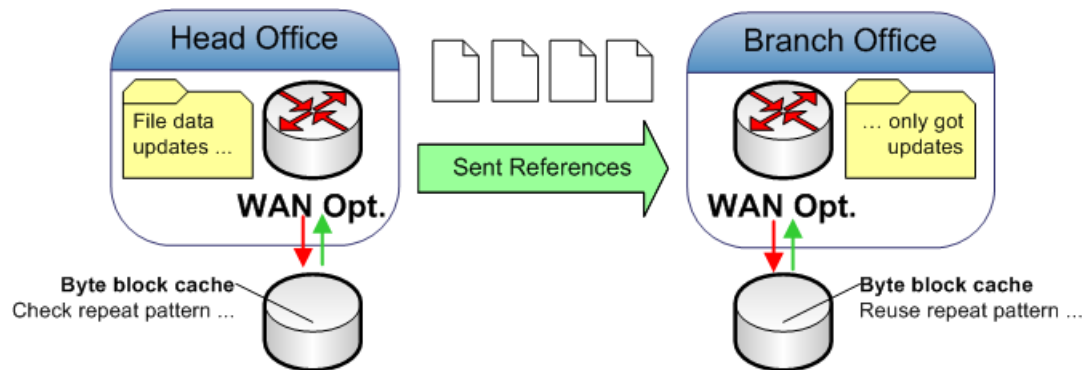


Figure 4.2: Moving symbol dictionary references instead of the data referenced thereby can achieve data reductions of three orders or magnitude or better.

LAN-Side Optimization Impacts WAN Delivery

LAN-side optimization handling accounts for network acknowledgements, maintains short response times, and uses shared symbol or content dictionaries to exchange large amounts of information by reference (rather than by outright exchange). This approach essentially short-circuits local application delivery to bring rapid-fire resumption and succession of remote activities across the LAN. The presence of WAN optimization devices at the LAN to WAN transition point also maintains consistency and helps synchronize both ends of a conversation quickly and efficiently.

Proxies and Tunneling

The entire problem with managing and maintaining globally, regionally, or territorially disparate networks is in finding a cohesive, comprehensive manner to unify diverse and mostly incompatible applications and protocols. When a WAN optimization device controls the conversation between tunneled endpoint pairings, greater and more granular control is exercised over the domain of operation between them.

WAN accelerators are designed to have functional counterparts on each end of a conversation to establish mechanisms for controlling application protocol delivery. Therefore, the WAN optimization appliance platform is an ideal checkpoint through which sub-optimal traffic passes to be prioritized and optimized for distribution across distant networks. It lets LAN conditions persist right up to the network periphery, then imposes transparent optimization onto traffic aimed at the WAN to improve the end-user experience, speed delivery, and shorten overall response times.

Special Considerations for Disaster Recovery

WAN links also service disaster recovery procedures and processes, so there are special requirements for deploying WAN optimization in such environments. Limited bandwidth, high latency, and lost and out-of-order packets put data replication and backup initiatives in jeopardy. Consequently, Recovery Point Objectives (RPO) and Recovery Time Objectives (RTO) aren't met in a timely fashion and the entire disaster recovery process suffers. This places heightened demand on WAN optimization tools capable of improving replication time while obtaining maximal bandwidth efficiency during such processes.

At the forefront of the evaluation process reside a few key principles:

- Know your applications and protocols—Observe how much traffic generates from replication processes, what and when peak loads occur, and the duration of data in transit between points.
- Know your network—Determine how much effective throughput is actually achieved across the WAN topology, taking into account excessive retransmissions from MPLS and VPN usage.
- Know your limits—Quantify traffic flow across the network to gauge how WAN optimization devices appropriately handle the volume.

Slight drops in throughput can have significant impact on particularly large backup and replication processes, which may consequently fail to complete. Latency also adversely affects disaster recovery operations, which is unavoidable for geographically disparate WAN endpoints. In certain instances, TCP acceleration techniques (including selective acknowledgements—SACK, adjustable window sizing, and round-trip time monitoring) can address this issue. All individual network appliances and intermediary network devices account for the many bottlenecking points that may limit effective throughput. Routers and firewalls may impose restrictions against the delivery of certain forms of traffic that adversely affects high-volume flow. Such issues need to be addressed directly through individual configuration on each network element or indirectly by way of WAN optimization techniques such as packet striping.

Packet striping overcomes quota restrictions emplaced on bandwidth for TCP streams and enforced by firewalls or routers. The intention is to prevent overutilization of available bandwidth. Packet striping divides the aggregate throughput for any given data stream among multiple flows. In this way, multiple smaller streams can still play by the rules without a large payload subject to checkpoint restrictions in passage. A single, bulky 100Mbps stream transiting a router restriction of 10Mbps per flow easily divides into 10 separate streams for optimal delivery.

Though beneficial, these routing restriction enforcements may unintentionally inhibit important traffic (such as scheduled online backups) with equal prejudice among several competing but less significant flows (for example, client HTTP traffic, Simple Network Management Protocol—SNMP—interactions, and routine network status checks). Striping breaks high-bandwidth traffic into several discrete flows for optimal transmission and later reassembly at the receiving end. It may be prudent to synthesize your approximate WAN conditions using a simulator or emulator as part of your WAN optimization evaluation process. Switching or upgrading existing network infrastructure to MPLS and VPN technology necessitates this discovery process and greatly benefits from its results. Good WAN emulators effectively reproduce real-world conditions specific to your network properties to include effective bandwidth, inherent latency, and non-sequential packet delivery.

Fixing Broken Protocols

Even when fixes are short-lived and properly addressed through application update or redesign, WAN accelerators can address notoriously WAN-unfriendly protocols. Known offenders on the itemized list of “broken WAN protocols” include both CIFS (file sharing) and MAPI (mail exchange). Vendors overcome the limitations of these protocols across the WAN using content prefetching or preemptive read-ahead and write-behind techniques.



Data *pre-fetching* builds a cache repository based on read and requested file segments for maximal efficiency. Read requests are served from partially cached files if requested elements are present.

Data *read-ahead* takes a predictive approach to accelerating WAN traffic by preemptively requesting file data ahead of the current cached portion to increase cache hits and performance.

Data *write-behind* techniques accelerate file transfers by deferring write requests until sufficient data accumulates to warrant issuing an all-at-once write.

For example, previous MAPI issues associated with Outlook 2000 are addressed in Microsoft Exchange 2003, which includes a new cached mode of operation to improve WAN performance. Although CIFS is designed for remote file sharing access across the Internet and other IP-based networks, it’s a fairly chatty protocol issuing hundreds to thousands of round-trip packets for a single file transfer (see Figure 4.3). CIFS performance is strictly LAN-bound and its chatty nature directly impinges WAN performance. Across the WAN, file shares accessed from a centralized data center undergo bandwidth and latency constraints that negatively impact performance.

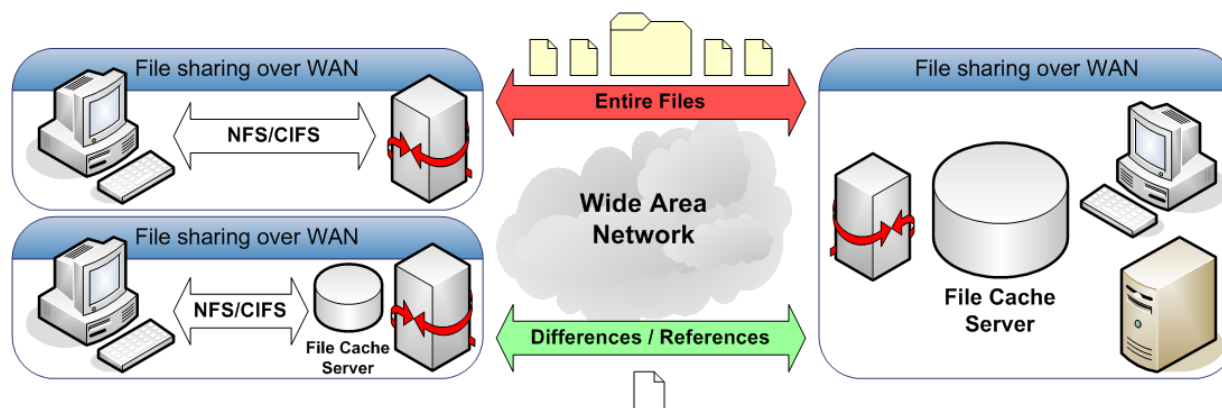


Figure 4.3: Replacing NFS or CIFS with WAN-capable file services and caching enables greatly improved communication efficiencies.

How Transactions and Conversations Can Be Short-Circuited (and Fast-Track)

Caching and reduction strategies facilitate faster end-user conversations and endpoint transactions by eliminating duplicate or wasteful transmissions. However, LAN-side acceleration appliances challenge how WAN optimizers implement QoS in branch offices and data centers. Traffic flowing internal (LAN) to external (WAN) has its protocol headers and payloads obscured by these appliances, thus preventing downstream devices (such as WAN routers) from applying optimization logic.

QoS mechanisms maximize WAN utilization particularly on oversubscribed WAN links or unmanaged sources of over-saturating data. Less important data usurps bandwidth from significantly more important data where demand exceeds capacity on the WAN. Unmanaged traffic in tight contention for the same limited resource deprives critical application and protocol data from timely and efficient delivery.

QoS serves to classify traffic based on application or protocol characteristics and prioritize delivery of critical conversations over insignificant traffic. Policy-based routing decisions ensure that proper and timely delivery is handled with utmost efficiency, including weighted queuing mechanisms that correspond to varying delay and usage constraints. Important data gets fast-tracked over lesser important data.

Enhanced Compression and Transaction Optimization

Compression strategies are present virtually anywhere storage volumes thrive in anticipation of eventual over-saturation. Many types of storage device perform basic compression techniques that will not prevent WAN optimization device operation but will significantly hamper or hinder its overall effectiveness. Preemptive data compression techniques restrict the visibility into application and protocol payloads for WAN optimization solutions.

In many cases, WAN optimization techniques perform similarly or superiorly compared with native compression strategies in storage arrays and should therefore be disabled for such volumes. Typically, this yields overall net performance improvement gains and alleviates the burden of processing compression algorithms and data on the storage arrays themselves. An enhanced level of compression and transaction optimization can occur—in some cases factoring up to 100 times the improvement for end-user response time.

Achieving Efficiencies: WAN Optimization Usage Cases

As the following use cases illustrate, specific WAN topologies and circumstances are amenable to various types of performance improvements using WAN optimization technologies. In the headings that follow, we examine a few different but specific cases in point.

Extending Inmate Tracking Systems for the State of Nevada

The Department of Corrections in Nevada operates numerous facilities in extremely remote locations. Some of them are more than 8 hours' travel time away from the department's Tahoe headquarters location. Though that remoteness serves as an additional barrier against escape, the distances involved posed interesting challenges to the department's key homegrown inmate tracking system. A batch-oriented system with no real-time data interface, its developer is no longer available, which makes internal code changes impractical.

Thanks to its age and old-fashioned design, this inmate tracking system presented numerous deficiencies. Thus, the department chose to replace it using a centralized real-time management system purchased from Syscon Justice Systems that features a Web-based interface accessible from any Internet connection. The new software handles all basic inmate operations, such as calculating sentences, tracking funds held in trust for inmates, and monitoring inmate location at all times.

Unfortunately, the Syscon program was not developed to take WAN latency issues into account. Though it performs adequately on a LAN, across slower WAN links, individual pages took nearly a full minute to load and display data. For sites not yet equipped with T1 lines, this could involve substantial added expense and even then, T1 connections were simply not available to the most remote of these sites.

Application-specific WAN optimization came to the rescue here. An appliance was installed near the WAN gateway at each location to manage traffic on each link. This led to substantial improvements in throughput, while caching enabled load times to decrease to around 5 seconds in most cases (after initial page loads had occurred). In addition, the Nevada Department of Corrections found itself able to manage all Internet access at all of its sites, and to control and groom Internet access. As a nice side effect, WAN optimization eliminated unnecessary WAN traffic, thereby freeing additional bandwidth for legitimate users and applications.

With 20 WAN optimization appliances at HQ and remote locations, implementation took some time and effort. The department had to develop access and security policies for the appliances to enforce, and they had to be tuned to provide an optimal mix of functions to meet user needs. As a result, database applications have manifested a 600% performance improvement, and page loads have improved by as much as 1000%.

Improve Inter-Campus Communications at a Multinational University

Though Heriot-Watt University is based in the United Kingdom, it operates a campus in Dubai, UAE, as well as the main campus in Edinburgh, Scotland. Though this involves a distance halfway around the globe, communications between both campuses are critical: the majority of the teaching staff and instructional resources are housed in Edinburgh, but students and staff on the Dubai campus need ready access to them. To meet its instructional needs, Heriot-Watt has implemented a Virtual Learning Environment (VLE) that permits inter-campus sharing of educational materials and online interactive tools that students share with instructors and each other. Also, Heriot-Watt uses the inter-campus WAN link for voice and video conferencing facilities.

Unfortunately, the WAN link in use provides only limited bandwidth so that data volume is limited and is subject to high latency, which poses severe problems for streaming or near real-time applications such as voice and video. This combination of hurdles made VLE problematic between the two campuses, where page load times might easily run as high as 20 seconds, and where a load of more than half a dozen active users would bog down the link sufficiently to interrupt communications.

Because of the distance between the two campuses, Heriot-Watt decided that a bandwidth upgrade was too expensive to afford and too difficult to implement. Even then, overall latency would still have remained fairly high owing to use of long-haul satellite links. Careful implementation of a pair of WAN optimization appliances with application proxy capabilities enabled the university to implement its VLE across that link. A combination of WAN optimization techniques that included compression, object and byte caching, and traffic prioritization, along with application protocol optimization, allowed Heriot-Watt to make use of video-based e-learning technology across the WAN link as well. Page load times declined from nearly 20 seconds to around 1 second, and delivered significant performance boosts to VLE and email communications. The ability for students and staff on both campuses to access the same learning materials at nearly the same instant also means the university can offer virtual classes to students on both campuses at the same time.

Because of application and file-level caching, on-demand video e-learning modules can be transported across the WAN during non-peak hours, and then stored locally on both campuses. The university can also offer live video feeds across the WAN by sending a single video stream, then splitting that feed into as many individual streams as are needed on each campus. This has enabled numerous e-learning pilots, which the university had hitherto considered infeasible.

The WAN optimization appliance also supports acceleration of SSL-encrypted traffic so that Heriot-Watt can move optimized, secure traffic related to administrative applications across the WAN link. This lets the university make much better use of available bandwidth, yet still make use of critical applications as they're needed. Also, only traffic that needs to move between Dubai and Edinburgh traverses the WAN link, as both WAN optimization devices can provide Web security for local Internet connections as well as block against malware and other malicious threats locally. Likewise, these appliances can impose security and allowable use policies on the Internet link to block or limit unwanted or inappropriate applications, such as peer-to-peer file sharing.

Networx Australia Accelerates Application and Content Delivery

Networx Australia is a managed service provider that works with Australian enterprises and organizations. It offers its clients Internet access, WAN connectivity, and security solutions for all kinds of business needs. Today's managed service environment in Australia is highly competitive so that service providers are continually on the lookout for ways and means to differentiate themselves and to inspire customer loyalty and retention.

To that end, Networx Australia chose to provide WAN optimization services to its customers to accelerate application and content delivery without increasing—and in many cases, decreasing—bandwidth consumption. Because WAN bandwidth is often limited and/or metered, reductions in bandwidth consumption can translate into cost savings or make room for rich media applications including teleconferencing, telepresence, and VoIP.

For many of its customers, Networx Australia provides the WAN infrastructure between their own headquarters' centralized data centers and branch offices. End users in the branch locations need access to data, applications, and services housed in the data center but want to obtain LAN-grade performance and response time when doing so. Prior to deployment of its WAN optimization devices, many branch users experienced excessive network latencies to the point where applications ran painfully slowly or did not even work at all. Adding bandwidth can't completely address such problems, but WAN optimization helped bring latency levels under control and enabled branch users to increase their productivity while speeding overall response time and usability.

Networx Australia's use of WAN optimization devices also enabled them to improve the security as well as the performance of their WAN infrastructure. In addition to accelerating application access and reducing network latency, these devices provided a means whereby companies could look to Networx Australia to provide malware protection, impose URL filters to block access to inappropriate or disallowed sites and content, and prohibit use of unwanted or unsafe applications that might be illegal, unauthorized, unlicensed, or not related to normal workaday use (such as peer-to-peer music or video access, BitTorrent downloads, and other "personal use" protocols or services not needed on business networks).

By deploying WAN optimization devices at customer branch offices and at their data centers, Networx Australia was able to accelerate application and content delivery at the same time it improved Web security and network controls. Customers report improved behavior and usability of approved distributed applications and file services, more overall available bandwidth, and improved application experiences for end users.

Behind the scenes, Networx Australia also was able to reduce its own operational costs, thanks to more efficient use of its WAN infrastructure, and to improve customer loyalty and retention. Caching technology and protocol optimization have enabled customers to speed CIFS file transfers by factors of up to 200 times. Web-based applications now run seconds faster per click, on average, and login/startup times for customers have improved dramatically.

The same security protections that customers enjoy in their data centers and branch offices also benefit Networx Australia, with increased protection against malware and unauthorized or unwanted applications. Traffic management (and reductions from elimination or throttling of unauthorized programs and services) has also let Networx Australia make better use of its WAN infrastructure, and improve its overall profitability without increasing operation costs. Across the board, introducing WAN optimization has been a win-win for both the provider and its customers, and has helped them to develop and retain a growing and satisfied customer base.

Conclusion and Recommendations

Companies and organizations seeking to optimize use of existing WAN links, and to maximize their ROI on new WAN links, will find that in addition to increasing the utility of bandwidth consumed, WAN optimization devices offer improved security, increased control, and more effective use of key applications through appropriate proxy deployments. Many buyers discover to their delight that the payback period for investments in WAN optimization are shorter than originally projected because increased usability, improved response time, and better application control often leads to higher-than-projected growth in WAN usage and unexpected productivity gains. Any company or organization that uses WAN links for regular, ongoing communications will find it worthwhile to contemplate, and probably to implement, some kind of WAN optimization strategy.

Download Additional eBooks from Realtime Nexus!

Realtime Nexus—The Digital Library provides world-class expert resources that IT professionals depend on to learn about the newest technologies. If you found this eBook to be informative, we encourage you to download more of our industry-leading technology eBooks and video guides at Realtime Nexus. Please visit <http://nexus.realtimepublishers.com>.