



Essential Requirements for WAN Acceleration

WAN Acceleration Evaluation

Revision 1.0
25 July 2007

ESSENTIAL REQUIREMENTS FOR WAN ACCELERATION

Introduction

This document proposes a list of suggested requirements and evaluation criteria for selection of a WAN Acceleration product that can be deployed robustly to support a large-scale enterprise network. While many WAN acceleration products seem to perform adequately in simple lab tests in isolated network environments, very few have proven to be able to sustain full-scale production operation across a large number of different applications, under heavy traffic loads, and in larger-scale networks having from dozens to hundreds or thousands of remote locations. WAN optimization products that have proven to meet the stringent requirements of the larger, more complex networks will deliver greatest returns through stable trouble-free operation, including for deployments at more normal-sized networks. This document lists evaluation criteria used to identify products that are capable of meeting robust scaling and performance requirements.

How to use this document

In this document each suggested requirement is listed in **bold**, followed by a discussion about why the requirement is relevant and important. The reader should formulate and tailor their own list of applicable requirements to fit their specific needs. This document focuses on the essential requirements (e.g., "the product must support CIFS acceleration"), as opposed to trivial requirements (e.g., the product must be weigh less than 10 pounds).

Not all requirements are necessarily applicable in every network environment; each environment has its own unique characteristics, and the reader should tailor their requirement to fit the needs of their specific network and applications.

Table of Contents

Configuration and Network Integration.....	4
Must be able to provide references of other similar-sized successful customer deployments.....	4
Must support multiple bypass port pairs for in-path deployment.....	4
Must be deployable out-of-path with WCCP	4
Must support autodiscovery of remote peer devices.....	4
Must support manual configuration of remote peer devices where necessary	4
Appliance-based solutions must not require additional software plugins or agents on client or server hosts.....	4
Optimized traffic must preserve IETF-specified (RFC 2581) TCP congestion control behavior over the WAN.....	4
Must detect asymmetrically-routed network traffic	5
Must optimize asymmetrically-routed network traffic	5
“Transparency” Mode should be optional, not mandatory (alternatively a requirement for “accurate addressing”).....	5
Scaling and High Availability.....	6
Must be able to simultaneously communicate with at least ___ peer devices (fill in your specific requirement)	6
Must support clustering without WCCP or PBR.....	6
Must not block additional application sessions when capacity is reached.....	6
Must support ___ TCP connections and ___ application sessions per device (fill in your specific requirement)	7
Must support ___ TCP connections and ___ application sessions through clustering (fill in your specific requirement)	7
Data Reduction Technology.....	7
Must have disk-based data store for byte-level data.....	7
Data store must be persistent through reboots and WAN disconnections.....	7
Must be able to synchronize Data Stores to backup WAN optimization device.....	7
Must have at least __GB of storage capacity for byte-level data (fill in your specific requirement)	7
Must use a universal single-instance data store architecture	8
Application Acceleration Technology.....	8
Must address protocol chattiness and latency issues for CIFS (Windows File Sharing)	8
Must not interfere with the operation of CIFS file locks between client and server.....	8
Must address protocol chattiness and latency issues for Exchange/MAPI	9
Must address protocol chattiness and latency issues for Exchange 2003/MAPI2003.....	9
Must support cross-protocol data recognition and data reduction for Exchange/MAPI	9
Must address protocol chattiness and latency issues for NFS version 3.....	9
Must address protocol chattiness and latency issues for enterprise web applications	9
Must address protocol chattiness and latency issues for SSL encrypted applications	9
Must be available as a software client installable on Windows PC’s and Laptops	10
Management.....	10
Must have a central manager capability available	10
Central manager must be optional, not required.....	10

Configuration and Network Integration

Must be able to provide references of other similar-sized successful customer deployments

With a large number of WAN acceleration vendors offering products in this relatively-new technology area, one of the best ways to identify the most suitable product is to look at other success stories. Many WAN acceleration products look good on paper, but only a small number of vendors have had their products successfully rolled out in large-scale production network environments. The solution must have consistent and demonstrated success in a real-life production network, and not just in a simple two-box lab demonstration in an isolated test lab. Some vendors encourage customers to limit their evaluations to a lab environment, where they can avoid unfavorable conditions (e.g., unknown applications, heavy traffic loads, etc...). Often they introduce contrived conditions (e.g., artificially-high packet loss levels) that unfairly benefit their own products. The best way to see if a product will work—and work well—is to look for concrete examples of full-scale successful deployments in other customer networks that are similar to your own. The vendor should be able to furnish these customer references if indeed they exist. Be sure to check that the references have actually deployed at a scale similar to your own network, not just that the customer has purchased a comparable number of devices.

Must support multiple bypass port pairs for in-path deployment

Larger networks, particularly those designed for high availability and redundancy will have redundant Ethernet LAN connections. To efficiently support these environments, the WAN optimization product should support multiple Ethernet bypass port pairs so that a single appliance can be simultaneously connected to multiple Ethernet LAN connections. This avoids having to use a separate appliance for each Ethernet connection.

Must be deployable out-of-path with WCCP

WCCP should be an available deployment option, as it allows the WAN acceleration device to physically reside outside of the network path. This option allows traffic that cannot or should not be optimized to avoid any contact with the WAN acceleration device, should that be desirable.

Must support autodiscovery of remote peer devices

The WAN acceleration device must be able to dynamically detect the presence and IP address of any other WAN acceleration devices, on demand and without advance administrator configuration. True autodiscovery significantly reduces the amount of tedious configuration work and potential for configuration error, especially in networks with large numbers of WAN acceleration devices or sites with WAN acceleration devices. Furthermore, autodiscovery avoids the need to enter host subnets in the configuration process, another source of headaches and errors.

In contrast, tunneling-based products require that all host subnets be identified. This is an administrative burden, particularly for networks with large numbers of subnets and those that use formal subnet tracking tools management processes. Each and every time an IP address subnet assignment is created, deleted, or changed, the change must be reflected not only in your existing address tracking tools, but also must be reflected in the WAN optimization solution configuration. The long-term task of keeping these two configuration databases synchronized, up-to-date, and accurate is a potential source of conflict and error in the foreseeable future.

Must support manual configuration of remote peer devices where necessary

In some environments, firewalls and intrusion protection devices may block use of the TCP options field, and this can interfere with the operation of the autodiscovery mechanism. In these environments, one can either change the configuration of the firewall, or alternatively disable the autodiscovery mechanism and instead manually configure IP address information of each remote peer device. Because the former approach can compromise security, the latter option of manually configuring peering relationships is the better choice in many cases. However, some WAN optimization products lack this ability to disable autodiscovery and instead use fixed rules that specify the IP address of the peer device. These WAN optimization products should be avoided as they lack the flexibility to be deployed in highly-secure environments where firewalls disallow use of the TCP options field.

Appliance-based solutions must not require additional software plugins or agents on client or server hosts

A truly transparent appliance-based solution should not require software agents on client or server hosts. This defeats the purpose of having an appliance-based solution.

Optimized traffic must preserve IETF-specified (RFC 2581) TCP congestion control behavior over the WAN

Some WAN acceleration products do not honor IETF-specified TCP congestion control semantics. Rather, they use TCP

windowing tricks to boost performance and throughput to levels that normally wouldn't be allowed by standard TCP congestion control safeguards specified in RFC2581. These products attempt to boost performance by stealing bandwidth from other applications sharing the WAN infrastructure.

If you are planning to deploy WAN acceleration devices in an environment with shared WAN connections, then non-TCP compliant products should be avoided. The danger for these products is when multiple non-TCP traffic flows collide in a shared WAN environment, significant levels of network congestion and packet loss will occur, affecting all traffic sharing the WAN. Products that do not honor slow start and congestion avoidance principles outlined in RFC2581 will not know how to adequately handle such situations. Rather than fall back and reduce their transmission throughput under heavy network congestion, these products attempt to "power-through" the congestion, by maintaining their transmission rate in the face of packet loss. The result is severely degraded network conditions for all traffic sharing the WAN.

Must detect asymmetrically-routed network traffic

Asymmetric routing is unavoidable in most large and highly-available IP networks. Whether your network already has redundancy features such as multiple network paths, or at some future point you believe you may add additional connections to provide the redundancy, the WAN acceleration product that you select should have the capability of detecting and not harming or blocking traffic that is asymmetrically-routed.

Some WAN optimization products block traffic that is routed through a different WAN optimization device for each direction of traffic flow. These products are incapable of detecting asymmetric routing conditions, particularly when a given WAN optimization device only sees the client's request, but not the server's response. When this condition occurs, the WAN optimization device must be able to realize that asymmetric routing is probably going on, and that it should not attempt to intercept and optimize the traffic, but rather let it pass-through without disruption. Products that cannot do this should be avoided, because they block traffic and break applications when asymmetric routing exists in a network.

This may be an important requirement even if you don't believe you have asymmetric routing. Often, in the initial state of the network asymmetric routing will not exist at all until a link failure or some other event occurs, which triggers activation of an alternate network path. In this case, a WAN acceleration deployment that initially works fine may suddenly start blocking traffic and breaking applications as a result of normal IP re-routing operations through the dynamic routing protocol (e.g., OSPF, BGP, etc.).

Note that detecting asymmetrically-routed traffic only means that the traffic will be passed-through without optimization, and does not necessarily mean that the traffic will be optimized. The next requirement deals with the ability to actually optimize asymmetrically-routed network traffic.

Must optimize asymmetrically-routed network traffic

The WAN acceleration product should not only detect, but given adequate configuration, also be able to optimize asymmetrically-routed network traffic. This requires that the product have an ability to forward and re-direct traffic appropriately so that the asymmetry is eliminated as far as the WAN optimization device is concerned.

Note that WAN acceleration devices that address protocol chattiness issues must be able to intercept and optimize traffic in both directions of data flow. If one direction of data flow is intercepted and optimized by one WAN optimization device, and the return path for that traffic follows a path supported by a different device, then problems can occur because each of the two different devices only see one duplex of the traffic. This requirement states that the WAN optimization devices must communicate, coordinate, and appropriately forward the traffic flow so that the device performing the optimization service can see both directions of traffic flow.

"Transparency" Mode should be optional, not mandatory (alternatively a requirement for "accurate addressing")

Some WAN acceleration vendors offer a "transparency" feature that uses the original IP address of the client and server to address packets containing optimized data for delivery over the WAN. This feature allows network monitoring devices to observe the original source and destination of the network traffic, but it hides the real sender and receiver of the processed, compressed, and transformed traffic, which are the WAN optimization devices, not the original client and server. The alternative approach to "transparency" is "accurate addressing" the processed, compressed, and transformed traffic. This approach accurately reflects on the compressed traffic that the WAN optimization device is sending the traffic to another WAN optimization device, so that there is no possibility of misrouting of the traffic.

While a "transparency" feature can be useful in some specialized network environments, it also has a number of significant impacts to the overall integrity of the functioning IP network, and as a result many customers may need to disable this feature to

avoid routing problems and application disruptions. Unfortunately, “Transparency” mode is not only the default mode for some WAN acceleration products, but for some products it cannot be disabled even if the customer does not want to use this feature.

This raises an issue for customers with large networks having redundant WAN routers and asymmetrically-routed traffic. Because the optimized WAN traffic contains the IP address of the end-host, and not the WAN optimization device, traffic can potentially be routed to the wrong WAN optimization device, particularly in networks with asymmetric routing and multiple WAN optimization devices deployed at the data center site. Furthermore, routing loops can develop because network routers cannot distinguish between traffic that has already been processed by the WAN acceleration device, and traffic that has yet to be processed, and as a result they may continuously route already-processed traffic back to the WAN acceleration device in a continuous loop, resulting in black-holing of the traffic. Finally, a very-real possibility exists for the wrongly-addressed packets to actually be delivered to the end-host, resulting in confusion and application disruptions when that traffic contains random byte-level data compressed by the WAN acceleration device.

In order to avoid these issues, the WAN optimization solution must have a capability for accurate addressing. This capability can be optionally enabled, but it must exist, and when it is utilized, the optimized, compressed, and transformed data sent over the WAN segment of the network will accurately utilize the real IP addresses of the WAN optimization devices. This prevents any chance of misrouting of the traffic from occurring.

While the “Transparency” feature might be useful in a limited number of cases, nevertheless in many enterprise networks, it can cause serious problems and challenges, and it may need to be disabled in order for the WAN acceleration solution to function without disrupting the network and various applications.

Scaling and High Availability

Must be able to simultaneously communicate with at least ___ peer devices (fill in your specific requirement)

This requirement is important for customers having (or potentially having) large numbers of sites where WAN acceleration devices will be deployed. It becomes an even more important requirement if your network includes an MPLS-based WAN, which provides for any-to-any connectivity. Many WAN acceleration devices can only connect with a small number of remote peer devices, while others can connect with thousands of peer devices, each deployed at a different site.

Be careful about products that claim to increase the number of peers that they can communicate with through clustering approaches. Load balancing and clustering through WCCP unaware of device peering relationships, and the same is true for many proprietary clustering approaches (e.g., Juniper’s WX/WXC cluster). With random traffic load balancing, there is no way to determine which clustered/load balanced device will receive a given traffic flow. As a result, every device in the cluster must form relationships with the same peer devices. For example, if a given WAN acceleration product can peer with a maximum of 50 devices using WCCP to cluster 3 of these devices together does not necessarily raise the peer limit to 150, because each of the 3 clustered devices typically must support peering relationships with the same 50 peer devices.

Furthermore, note that this architectural requirement has implications on the solution’s ability to support individual mobile software clients. If a given WAN optimization device can only support connectivity to 50 peer devices, then effectively each data center device can only support 50 mobile software clients. On the other hand, many enterprise environments expect to have far greater than 50 mobile employees using a software-based Windows client-PC solution.

Must support clustering without WCCP or PBR

WCCP does have drawbacks, one of which is increased demand on router processing resources. In some environments the traffic load is such that WCCP places too much demands on the processing resources of the router or switch, and therefore an alternate load balancing and clustering approach must be supported, without use of WCCP. Another problem with both WCCP and PBR is that they are unaware of device peering relationships or device overload: they can tell when a device is up or down, but they have no finer-grained visibility in to the state of the cluster.

Must not block additional application sessions when capacity is reached

A solution that scales properly must be able to survive overload conditions without disrupting application sessions for additional traffic beyond the device capacity. In a large massively-scaled network environment, there will be time periods where application demand spikes to higher levels. To accommodate this, when the session capacity of the WAN acceleration device is reached, additional sessions should be passed-through and not be adversely affected.

Note that some WAN acceleration devices use an application-specific caching (e.g., file cache) approach. These devices handle

session requests similar to how an application server would—if the maximum number of application sessions that can be supported by the server is reached, then additional connections are denied. Unfortunately, this behavior is not appropriate for a WAN acceleration device deployed to intercept and generically optimize all application traffic, and any such device should be eliminated from consideration as it will not scale, unless the device can be adequately oversized so as to guarantee that the connection-denial never occurs.

Must support __ TCP connections and __ application sessions per device (fill in your specific requirement)

Scaling the solution requires that each WAN optimization device be able to accelerate a significant amount of network traffic. Bandwidth metrics are often deceptive, as raw packet compression capacity does not accurately measure the solution's ability to apply layer-7 application-specific acceleration techniques designed to address latency and protocol chattiness issues. Rather, a more accurate metric of a device's scaling capability is the number of TCP connections and application sessions that can be intercepted and accelerated. Note that these are two separate and equally-important metrics that must be considered and compared across all WAN optimization solutions being considered. Note that some vendors emphasize a very high connection limit for generic TCP connections, but they have significantly lower application-level session (e.g., CIFS, MAPI, etc...) limits that they avoid disclosing. Make sure you specifically enquire on the vendor's limits for both TCP connections and application sessions.

Must support __ TCP connections and __ application sessions through clustering (fill in your specific requirement)

Many network environments are sufficiently large that multiple WAN optimization devices must be deployed at some sites, particularly the data center hub site. In these cases, multiple WAN optimization devices must be clustered in order to scale the solution as required, and the total number of TCP connections and application sessions that the clustered solution must be compared across all WAN optimization solutions being considered.

Data Reduction Technology

Must have disk-based data store for byte-level data

A sufficiently-large disk-based data store with multiple gigabytes to terabytes of data storage capacity, used to retain byte-level repetitive byte patterns is essential for effective WAN optimization. A disk-based data store is persistent across reboots, and will retain learned repetitive byte patterns for a period of weeks, if not months or years. The overall end-user experience is significantly enhanced when a file or email attachment last accessed a number of weeks ago is accessed with LAN-like speeds. Furthermore, transfer times for large files, such as a 100MB FTP file transfer, are significantly faster if the entire byte-level data can be retained in the disk-based data store.

In contrast, WAN acceleration products that only retain learned byte patterns in DRAM memory are significantly less effective due to their lack of storage capacity. Some DRAM-only products, due to storage constraints, only store repetitive byte patterns that are observed at least twice—these products only provide “warm” performance acceleration on the third transfer of a given file. Other DRAM-only products store all byte patterns that they observe, and any learned files are retained in memory for a period of a few minutes before they are flushed out by newer byte patterns that are more recently acquired.

Data store must be persistent through reboots and WAN disconnections

A data store that will be retained for months to years must be persistent through device reboots and loss of WAN connectivity. Some WAN optimization products will lose their data contents if the device is rebooted, and some will even lose their data if connectivity to the peer WAN optimization device is disrupted.

Must be able to synchronize Data Stores to backup WAN optimization device

In order to support high availability, redundancy, and failover of the WAN optimization solution, the byte-level contents of the disk-based data store must be able to be copied and synchronized to the backup WAN optimization device. In the event of failure or loss of the primary WAN optimization device, then there will be no loss of data and the backup device will be able to continue accelerating with the same data store as the previously-operational primary device. Data store synchronization between redundant WAN optimization devices must be supported regardless of the deployment approach, and this includes not only for in-path deployments, but also deployments that make use of WCCP.

Must have at least __GB of storage capacity for byte-level data (fill in your specific requirement)

Backup and data replication applications involve transfers of potentially hundreds of GB to several TB of data. The WAN optimization device must be capable of storing the entire set of byte patterns transferred for the entire backup and data replication process. Of course, this will vary depending on your specific application and environment, but the data store must

match or exceed the nominal amount of application data to be backed-up or replicated.

It is important to distinguish between a byte-level data store and an application-level cache. Disk-based storage can be utilized to store either byte-level data or application-level data. The byte-level data storage capacity required in the WAN optimization device is typically far less than the amount of file-level data being transferred over the WAN. A byte-level data store eliminates the byte-level data redundancies found among different versions of the same or similar files. It is not uncommon to be able to support three to five times as much application-level data with a given byte-level data store (e.g., a 500GB data store can often support between 1.5 to 2.5TB of files, email attachments, web objects, etc...). On the other hand an application-level cache (e.g., a file cache or web cache) must have a data storage capacity that is greater than the amount of application-level data being transferred over the WAN. Note that for some WAN optimization products, data is stored twice—first at the byte level and then again at the file level. For such products, the amount of disk storage needed is far greater than the amount of application-level (i.e., files) data that must be transferred over the WAN.

A memory-only WAN acceleration device stores much less data than disk-based devices, and will provide minimal benefit for backup and data replication applications because the stored data will quickly be overwritten.

Must use a universal single-instance data store architecture

The WAN optimization solution must use a universal data store that stores each byte-pattern once, regardless of the number of remote peer devices accessing the file. In the case of an identical file fetched by users at 10 different branch offices, the relevant data must only be stored and represented only once in the data center WAN optimization device—not 10 separate times. Each additional accelerated transfer of that file to any remote site leverages the same instance of data stored in the central WAN optimization device, and must not consume additional storage.

Note that some WAN optimization devices use a partitioned data store for each peer device, which stores each instance of data once for each tunnel or remote peer device. For a given file retrieved by users at different remote sites, the data center WAN optimization device must store the same instances of byte-level data at least once for each remote site. As a concrete example, where users at 10 different branch offices retrieve the same file, the core WAN optimization device in the data center must store the relevant data 10 times—once for each tunnel or remote peer WAN optimization device.

A partitioned data store architecture also has implications on the solution's capability to support a software client-based solution for mobile users. In this case a separate data partition must be created and maintained for each mobile software client. If 10GB were set aside for each mobile software client, then supporting 1000 mobile software users would consume 10TB of data store capacity in the central data center appliance.

The per-peer data store architecture used by some WAN optimization devices will not scale to meet the needs of larger more complex networks having many remote sites and/or individual mobile client users. As a result, the WAN optimization solution must use a single-instance universal data store architecture.

Application Acceleration Technology

Must address protocol chattiness and latency issues for CIFS (Windows File Sharing)

CIFS, also known as Windows File Sharing, is one of the more common application protocols seen over the WAN. The WAN optimization solution must not only address scarce bandwidth issues through data reduction and compression technologies, but it must also address latency and protocol chattiness issues associated with CIFS chatty protocol behavior. Furthermore, the WAN optimization must not only accelerate simple CIFS desktop drag & drop copy operations, but it must also accelerate more complex CIFS operations that are initiated from within applications, such as when an application user clicks the "save-as" button within the MS-Excel application.

Must not interfere with the operation of CIFS file locks between client and server

The CIFS acceleration mechanism must be transparent to CIFS file lock operation between client and server. In order to provide safe acceleration of CIFS traffic, end-to-end file locking mechanisms employed by the Windows operating system must be able to continue functioning. As a result, a WAN outage should be handled no differently from a PC being disconnected while a file operation is taking place—in this case the Windows operating system will save a local backup copy of the file, and the end-user can recover normally through existing Windows file recovery processes.

Note that file caches inherently intercept these file locks and use a proprietary mechanism to relay file lock information across the WAN. If a WAN disruption or other event (such as a glitch in the WAN optimization device) prevents the local and remote WAN

optimization devices from communicating over the WAN, the result can be data corruption or blocked file access because of incorrect file locking.

Must address protocol chattiness and latency issues for Exchange/MAPI

The WAN optimization solution must be able to not only reduce bandwidth consumption through compression and disk-based data reduction techniques, but also address latency and protocol chattiness issues associated with the chatty Exchange/MAPI protocol.

Exchange email is one of the more common applications used over the WAN; acceleration of Exchange/MAPI traffic allows consolidation of Exchange servers to central data centers, thus saving cost, improving backup and data protection processes, and increasing overall IT efficiency.

Must address protocol chattiness and latency issues for Exchange 2003/MAPI2003

When an Outlook 2003 client connects to an Exchange 2003 server, chatty protocol behavior will slow down transfer performance, although the slow performance can be masked if Outlook Cache Mode is used.

Exchange 2003 uses a different version of the MAPI protocol than earlier Exchange servers. WAN optimization products that claim to optimize Exchange/MAPI do not necessarily optimize Exchange 2003/MAPI2003, because the latter uses a different protocol.

Must support cross-protocol data recognition and data reduction for Exchange/MAPI

Both Exchange and Exchange 2003 servers encode their attachments before sending them to the Outlook client. Because of the Exchange special encoding, the original byte-patterns of the file are changed when they are viewed over the network. As a result, the byte-patterns for the encoded attachment are different than if the same file was sent using a different protocol such as CIFS, HTTP, or FTP. A WAN optimization device must first undo the encoding before applying data reduction algorithms, if the byte patterns of the file are to be stored and recognized in their native format. This approach allows the WAN optimization device to deliver cross-protocol data recognition and data reduction. For example, when an attachment is initially retrieved by the Outlook client from an Exchange server, it will also be accelerated when that attachment is then sent over for the first time using CIFS or FTP. If a WAN optimization device applies data reduction algorithms to the file attachment's encoded form, then the WAN optimization device cannot deliver cross-protocol data reduction because the data had initially been stored and recognized in its encoded format by that WAN optimization device.

Must address protocol chattiness and latency issues for NFS version 3

Environments with Unix hosts often use NFS to transfer files across the WAN. If your environment includes Unix-based workstations and servers, then you may want to include NFS protocol optimization to your requirements.

Must address protocol chattiness and latency issues for enterprise web applications

Simple web applications and static web pages generally perform well with WAN optimization solutions that deliver disk-based compression and data reduction mechanisms. However, more complex enterprise-class business applications often exhibit chatty behavior when accessing various web objects needed to populate a given web page. If your environment includes such applications (e.g., Siebel, Oracle, SAP, etc...), then a feature that pre-fetches web objects and addresses latency-related performance issues for enterprise web applications should be one of your requirements.

Must address protocol chattiness and latency issues for SSL encrypted applications

SSL-encrypted web traffic comprises an increasing portion of today's enterprise network traffic. The traffic must not only be encrypted for confidentiality, but also authenticated to ensure the data comes from the correct source. Unfortunately, these requirements also prevent many WAN acceleration solutions from accelerating SSL traffic, because the encrypted traffic has been scrambled and any repetitive data patterns have been hidden in the encrypted data stream.

However, there are some WAN optimization solutions available that are now able to securely terminate the security layer in a safe manner, without compromising security. Once they can access the underlying clear-text data, WAN optimization algorithms and techniques can now be applied before the traffic is re-encrypted for delivery over the WAN.

If the SSL-encrypted traffic in your network is slow due to bandwidth and latency issues in the WAN, then consider adding a requirement for the WAN optimization device to accelerate the traffic in a safe and secure manner, without disrupting the existing trust model.

Must be available as a software client installable on Windows PC's and Laptops

According to IDC's Mobile Worker Forecast 2006-2009, about 450 million workers will use mobile computers in the near future. If your corporation has a number of mobile users, then the requirements for the WAN optimization solution must be available not only as a separate hardware appliance deployed in data centers and branch offices, but also as a software client on mobile Windows laptop computers.

The software client should support an invisible mode, allowing installation and configuration to be centralized. Operation of the mobile software client should be completely transparent to the end user. The software client should be tested and work well with your VPN client and anti-virus software, and interoperate well with the various applications that are used by the end user.

Note that each of the requirements discussed earlier in this document should also apply to the mobile software client solution. Specifically, the software client must not only address bandwidth reduction, but also latency and application protocol chattiness issues. It should be integrated and interoperable with the hardware WAN optimization appliance platform, in order to avoid management of two separate products.

Management

Must have a central manager capability available

Management is an important issue when scaling a solution. As more sites deploy the WAN optimization solution, it is important to have a central manager capability, in order to monitor the health and performance of large numbers of remote appliances. The central manager must be able to store configuration files and software images. These will be handy whenever the need to replace a given hardware device comes up, as the configuration file and software image can then be pushed out to the replacement hardware.

Central manager must be optional, not required

Often, there will be a need to deploy in limited environments, where only two devices—one at each site—are involved. In these cases, use of a central manager should be optional, not required. A central manager in a limited 2-3 site deployment not only raises cost, but also increases administrative burden and complexity for what should be a very simple deployment. In very large networks, processes and custom tools built around scripting, statistics export, and SNMP may be a better fit for the organization. In such a situation a mandatory central manager is not just an unfortunate waste of money, but may become a point at which the organization's preferred management system is undercut.

Riverbed Technology, Inc.
199 Fremont Street
San Francisco, CA 94105
Tel: (415) 247-8800
www.riverbed.com

Riverbed Technology Ltd.
1, The Courtyard, Eastern Road
Bracknell
Berkshire RG12 2XB, United Kingdom
Tel: +44 118 949 7002

Riverbed Technology Pte. Ltd.
350 Orchard Road #21-01/03
Shaw House
Singapore 238868
Tel: +65 68328082

Riverbed Technology K.K.
Shiba-Koen Plaza Building 9F
3-6-9, Shiba, Minato-ku
Tokyo, Japan 105-0014
Tel: +81 3 5419 1990

© 2007 Riverbed Technology, Inc. All rights reserved. Riverbed Technology, Riverbed, RIOS, Interceptor, Steelhead and the Riverbed logo are trademarks or registered trademarks of Riverbed Technology, Inc. Portions of Riverbed's products are protected under Riverbed patents, as well as patents pending.