Application
Acceleration Challenge

Application
Acceleration Challenge

Introduction

Setting the Stage

Network managers face a complex problem: their users and customers demand ever-decreasing response time, but at the same time, both the complexity and the volume of data are increasing. New applications are being built on web platforms using HTTP and XML, and older client/server applications are migrating to HTTP. Users expect network managers to apply some sort of "network magic" and make response time better than in the past, or at the very least, they expect response time not to increase. Fortunately, there is some magic available that addresses this problem: Application Acceleration.

Application acceleration is a key component in an emerging Application Delivery Architecture. Application acceleration is ensured a place as a practical tool in the network manager's arsenal because it doesn't require any changes to the applications. Moreover, application acceleration can significantly reduce response time by a factor of 2, 10, or maybe even 100 times.

Application acceleration addresses several trends that are currently undermining response time. Many of these trends are just now or will soon be appearing in most networks. The trends include:

- The migration of applications away from proprietary applications protocols to HTTP and XML. The move to HTTP and XML allows applications to use richer content, thus increasing the amount of data being sent.

- Cost and regulatory issues are driving servers to be consolidated in the data center. Previously, it was common for remote sites to have their own file and email servers. The problems with this older system are that few of these servers are effectively utilized, the cost of remotely maintaining the servers is high, and new regulations are driving the cost even higher. A more effective solution is to consolidate the servers in the data center, allowing each remote site to have its own virtual servers. Using this method, server resources are used more effectively, maintenance and backups are

cheaper and easier, and compliance with new regulations is simplified.

- Emails with large attachments are today's norm. The files can be Word documents, PowerPoint presentations, PDFs, or a host of other types of files. While email is not considered a "real time" application, this has not stopped people from expecting their immediate delivery. It is not uncommon two network users who are talking on the phone to be sending each other emails with attachments and expect the attachments to arrive instantaneously.

- Most workers are accessing applications, emails, and files from the road, home, and other locations. They expect to have the same interactions with their network from wherever they are, even though they are considerably further from the server when they are outside the office.

This is the main problem for network managers: remote users expect the same response time when connected remotely that they received when connected to the server over a LAN. Data consolidation, using the Web and HTTP and remote access should not result in lower productivity due to slower response time. It is the responsibility of the network manager to ensure that users have the same experience as if they were in headquarters. Adding to the problem is that the move to HTTP and the Web makes it easy for developers to include more objects, graphics, pictures, and other items that create significantly larger transactions. Consolidating servers in the data center means sending large attachments over the WAN.

Bandwidth is Not the Solution

The traditional answer to problems with response time has been to increase bandwidth. This approach works in many cases. Doubling the speed of the line does reduce the amount of time it takes to transmit the packets and is generally a good solution when the link is being heavily used. But

what is a network manager to do when link utilization is not the problem, and users are still unhappy with their response times? If the users are experiencing slow response time and the link utilization is low – under 50% – then adding more bandwidth will only marginally improve their response time.

Even if adding bandwidth is the chosen solution, there is no guarantee that it will solve today's response time problems because of the way newer protocols work. A Web page is composed of many objects. Each object is sent over its own connection. Additionally, many objects have to be sent in a sequence, decreasing the benefit of additional bandwidth. Given the start up time, the transmission times, and the sequencing problem, adding bandwidth may only have a small impact. Files and attachments pose their own problems. The protocols used, such as CFIS and MAPI, are designed for use over a LAN. The problem is that these protocols are "chatty" – they exchange many messages – and use a window size of one, which restricts the number of packets that can be sent before an acknowledgement is received. Network protocols such as TCP overcome this problem by using a variable sliding window protocol, but they cannot speed up response time if the application layer protocol still uses a window size of one, thus negating their efficiency.

Additionally, network managers cannot attempt to solve every problem by increasing bandwidth. Workers or business partners accessing applications over the Internet must accept the bandwidth available to them from their ISPs.

The Solution

The solution to these and other more general response time problems can be found in the application acceleration function of an application delivery architecture. Application acceleration is not a single technique. It is a collection of techniques and tools are applied to network traffic that can dramatically improve response time.

The techniques include:
- **Compression:** This can be as simple as applying standard compression found in every browser, GZIP, to the data. It can also go further to include advance pattern recognition or dictionary compression techniques that can reduce the amount of data being sent by a factor of 10 to 100 times.

- **Caching:** The accelerator can cache objects, saving the time it takes to retrieve the object from the server. The accelerator can also direct the user's browser to cache the object and automatically tell the browser to use the version it has already downloaded. Caching can be applied to both static and dynamic objects; it can also be rule-based for dynamic objects.

- **Connection pooling:** The accelerator can reuse TCP connections to the server, saving the time and resources it takes to set up a separate connection for each object.

- **SSL off-load:** The accelerator can have a high-speed processor to more efficiently encrypt the data.

- **TCP:** TCP's flow control is well designed for long-lived connections. With the increasing use of short-lived connections, TCP's slow start and general flow control can slow down the transfer. The accelerators, with their understanding of what is happening in the network, can improve the process, resulting in a faster transfer. Techniques can also be applied to TCP error correction to increase its efficiency.

- **Traffic shaping:** Traditional shaping techniques can be applied at both the TCP and application layer to ensure that mission critical traffic is not negatively impacted by less important traffic.

- **File and attachments:** Inefficient file protocols can be spoofed and turned into efficient file protocols. Additionally, the accelerator can apply caching to decrease delivery times.

- **Message Flow Reduction:** Accelerators can recognize what is happening in a message flow and automatically redirect the flow to the right server, saving the back and forth communication that would normally be necessary.

This is just a partial list, as accelerator vendors are constantly developing new methods and techniques to improve user response time. In the future, vendors will develop techniques to help accelerate XML and to address problems with specific applications such as Oracle or SAP.

Which Vendor?

The question is not whether an accelerator is in your network's future, but whose accelerator you will use. Vendors differ in their architecture, features, and techniques. The Application Accelerator Challenge goal is to help you, the network manager, understand acceleration and the vendors who are offering it. I have gathered together the leading vendors - Riverbed, Packeteer, Expand Networks, Citrix and Juniper - and asked them to give an overview of the solutions they offer. Specifically, I have asked each of them to answer the following questions:

- What type of application acceleration do you provide?

- What makes your solution the best solution?

- What is different about your solution?

I did not ask the vendors to spend time explaining the importance of application acceleration, so you will not see the normal "marketing" position on the importance of application acceleration. In the brief space, I asked them to concentrate only on the key advantages of their product to allow you to gain a flavor for the type of solution each offers. At the end of each section, the vendor's contact information is provided. I encourage you to talk with them to learn more about application acceleration and their unique solutions.

If you would like to learn more about the technology and issues involved in application acceleration, the second portion of the challenge provides more background. In these presentations, vendors and I talk about key technologies used in providing their solutions. The presentations are modeled after a conference session in which each vendor takes 10 to 15 minutes to present the technology using slides with voice over. I encourage you to view the presentation after you have read their answers to gain an even better understanding of their solutions and the technology of application acceleration.

If you have any comments about this Challenge you can contact me at Robin@Layland.com.

Wide-area Data Services

is the comprehensive approach to accelerating applications

riverbed™

Alan Saldich
VP of Product Marketing
Riverbed Technology

Background

Riverbed™, the performance leader in Wide-area Data Services (WDS), builds application acceleration appliances that transparently drop into any enterprise IT infrastructure. Riverbed's award-winning Steelhead™ appliances were introduced in 2004, and enable enterprises to simplify, consolidate and accelerate their IT infrastructure anywhere in the world by improving the performance of any application running over a WAN. In many cases, Steelhead appliances can speed up the throughput of applications by up to 100 times.

By offering the highest performance for the range of applications enterprises care the most about, combined with the easiest product to deploy on the market, Riverbed's products offer a multi-faceted ROI with a payback well under a year.

Applications

Steelhead appliances help simplify and accelerate IT in many specific and important ways.

- **Application Acceleration.** Steelhead appliances can be used simply to speed up applications like Lotus Notes, Windows or NFS (UNIX) file sharing, Microsoft Exchange, FTP, MS-SQL, SMS, document management, web-based applications, many client-server applications as well as backup and replication processes. In many cases, the degree of acceleration can be five to fifty times, or even up to 100 times or more under some conditions.

- **Bandwidth Optimization.** Steelhead appliances can be used to reduce WAN traffic, typically by 60% to 95%, which allows network managers to defer or completely avoid WAN bandwidth upgrades. In several instances customers have actually de-commissioned existing WAN links after deploying Steelhead appliances. With Riverbed's high-speed TCP optimizations, Steelhead appliances can also be used to optimize very high bandwidth WAN links to help "fill the pipe".

- **Site Consolidation.** Steelhead appliances also enable successful consolidation of remote office IT infrastructure from remote offices to data centers. This consolidated infrastructure could be Exchange servers (5.5, 2000 or 2003), file servers, filers or tape backup equipment like auto-loaders.

## The Riverbed Difference – Proven Speed and Simplicity

There are many products available to attack the problems facing IT with respect to remote offices, but no one but Riverbed provides a comprehensive solution to the problem. With over a thousand customers, including many Fortune 500 and Global 200 enterprises, having deployed our products across the world, our solution has been tested in some of the most demanding, complex network environments. In addition, Steelhead appliances have won tremendous industry recognition, including the *InfoWorld* Technology of the Year (2005 & 2006) for Best WAN Accelerator, the 2006 *eWeek* "Excellence Award" for Best Network Infrastructure product, the 2006 *Network Computing* "Well-Connected" award for Best Remote Office Solution, and the 2005 *Wall Street Journal* "Innovation of the Year" in networking, broadband and Internet.

Riverbed's success is based on an integrated approach to application acceleration. While other approaches started with a limited architecture and attempt to "add-on" additional optimizations, Riverbed built the Riverbed Optimization System (RiOS) from the ground-up to simultaneously address storage, networking, and application-level constraints to wide-area application performance.

Data Streamlining

Riverbed's Data Streamlining algorithms typically eliminate 65% to 99% of all data moving across a WAN. Data Streamlining stores all TCP traffic as small application-independent segments of data that can be re-used in the future by any file, web page, email or application. Unlike a WAFS or cache-based architecture, the data stored on the disk is not in the form of a file or other identifiable object, it's not related to file names, nor is it application specific. In this way, all redundant traffic that flows on your WAN, regardless of the application that generates it, is removed – in most companies, this turns out to be a significant fraction of the traffic.

Data streamlining overcomes bandwidth limitations
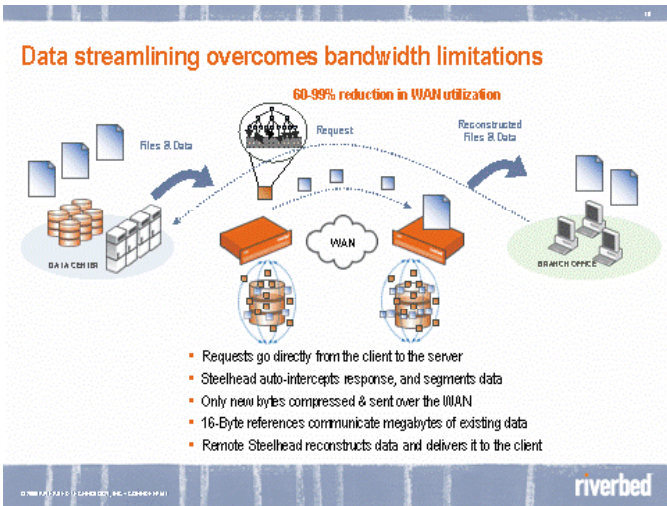
Both the segments and the references are stored on the disk on the appliance. If the data has never been seen before, then the segments and accompanying references are sent across the WAN and stored on the Steelhead appliance on the far side of the network. Subsequent client requests for the same data are fulfilled by the server normally, but the response from the server is intercepted and optimized – no repetitive data is sent again so long as those segments are stored on the disk.

Riverbed's patented technology has another key feature that allows an arbitrarily large amount of data to be represented by one reference of a few bytes, so that a small amount of data (a reference) can be sent over the WAN, yet represent a virtually unlimited amount of data. These powerful techniques enable unparalleled data reduction across all TCP-based applications.

As shown in the figure above, Data Streamlining algorithms take bulky responses from servers and then send only a thin envelope of information that represents the data without interfering with the application protocol semantic – in other words, the application continues to work just as it did prior to the Steelhead appliances being deployed, just much faster.
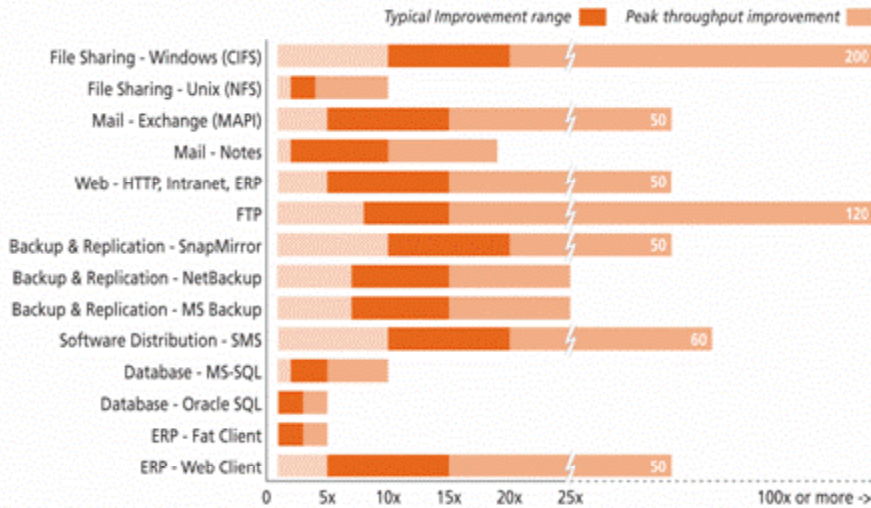
Transport Streamlining

Steelhead appliances also use Transport Streamlining in order to minimize the "chattiness" of TCP across the WAN. Transport Streamlining consists of half a dozen optimizations that streamline the performance of TCP. For example, Virtual Window Expansion (VWE) uses the references to data that are created by Data Streamlining to "re-pack" the TCP frame, and instead of just loading it with data, loads it with references to data that is already on the far side of the network. Thus, Transport Streamlining "virtually" expands the TCP payload far beyond what is possible using typical window scaling. Riverbed has also introduced its high-speed TCP algorithms that overcome the congestion control algorithms in TCP, and allow data to flow at rates over 700 Mbps, even at round-trip latencies well over 100 ms. for disaster recovery, large scale database replication and other highly data intensive workloads over LFNs.

Application Streamlining

Steelhead appliances also perform application level latency optimizations for particularly problematic applications. Many applications have their own protocols that ride on top of TCP, and in many cases those protocols are even more "chatty" (they require many more round trips to accomplish a task) than TCP. Riverbed's Transaction Prediction algorithms are application-specific optimizations built that eliminate most of the round trips generated by some of the most important and widely used applications within enterprises, including web traffic (HTTP), Windows (CIFS), UNIX (NFS), MS-SQL, and Microsoft Exchange (MAPI). In addition, Riverbed provides latency optimization for common backup and replication tasks. Riverbed will be introducing additional application level optimizations very soon.

What kind of results can you expect?

Typical Improvement range ▮   Peak throughput improvement ▮

Application	Value
File Sharing - Windows (CIFS)	200
File Sharing - Unix (NFS)	
Mail - Exchange (MAPI)	50
Mail - Notes	
Web - HTTP, Intranet, ERP	50
FTP	120
Backup & Replication - SnapMirror	50
Backup & Replication - NetBackup	
Backup & Replication - MS Backup	
Software Distribution - SMS	60
Database - MS-SQL	
Database - Oracle SQL	
ERP - Fat Client	
ERP - Web Client	50

0   5x   10x   15x   20x   25x   100x or more ->

riverbed

Integrated File Sharing and Disconnected Operations

An exciting feature of Steelhead appliances, Proxy File Service (PFS) enables integrated file sharing, and some disconnected operations into Riverbed's existing product line. There are situations where a WAN link fails, yet users in a branch office need access to files in order to complete their work. PFS provides access to files stored on the disk in the Steelhead appliance.

With PFS enabled, even when a WAN link fails, certain files will be accessible locally, and globally shared files will be available in a read-only mode. Unlike caching products, which try to make the choice for the administrator as to what files are available, PFS gives the administrator transparency and flexibility to decide what files should be available to users.

Management Streamlining

A critical benefit of Riverbed's approach is transparency. Riverbed's approach combines unparalleled application acceleration with ease of deployment, management, and monitoring. Unlike caching-based products that require client and server changes (for example mapping remote file shares), and WAN Optimization products that require administrators

to configure tedious tunnels among all offices, deploying Steelhead appliances requires absolutely no changes to your underlying router infrastructure, servers or clients – every part of your IT infrastructure continues to work as it does today. Deploying Steelhead appliances can typically be up and running in 10 to 15 minutes, allowing virtually immediate optimization of WAN traffic.

Summary

Riverbed offers a solution to application performance that is unparalleled by any other vendor. Unlike compression oriented solutions, Riverbed's Steelhead appliances add both TCP and application level optimizations. Unlike caching-oriented solutions that are application specific, Riverbed's products optimize all TCP traffic. The results speak for themselves – as one customer, Craig Mathews said: "McKim & Creed has seen performance increases of around 50 times the throughput we had before deploying the Steelheads. We performed extensive tests before and after deploying the Steelheads, and the performance increases are astounding."

Since Steelhead appliances accelerate all TCP traffic, applications as diverse as SMS, SnapMirror, FTP, Lotus Notes, Microsoft Windows, and many more can be significantly accelerated with Steelhead appliances.

For more information on Riverbed, visit www.riverbed.com, call +1 415 247 8800 (or, in the US call Toll Free: 1-87-RIVERBED), or email info@riverbed.com.

Packeteer –
Every Application at Every Location

Mark Urban
Director of Product Marketing
Packeteer, Inc.

Network managers have the responsibility for maintaining the health of the corporate network. The real challenge they face, however, may not be the one immediately recognized by most people. With reality of hundreds of different applications, the potential for so many network issues and, equally vast solution sets at their disposal (consisting of a myriad of tools, techniques and technologies), how do they best determine the ideal solution — or combination of solutions — for any given issue at any given time?

In the past, network managers were forced to rely on "point products," a narrowly-focused group of devices that were geared to address a select set of problems or single sets of applications. The problem was, however, that many soon found that these devices created an impenetrable thicket of overlapping technologies and protocols. Never has this more evident than with application acceleration.

Many "accelerator" products focus on a single or a few technologies. Upon installation, a network manager may find that one function works really well. Suddenly, opening or saving documents, e-mail applications or document storage archival across the network gets perceptibly faster. Which on the surface sounds great — but what if other problems exist around a business-critical application like SAP, Oracle or Citrix? What if the issue is protecting jitter-prone voice and video applications?

It is precisely issues such as this that have many network managers embracing Packeteer's complete system. Packeteer gives network managers a multifunctional solution

- **Find the problem** – Packeteer helps determine source of the problem (or problems) using deep Layer 7 Plus visibility into the applications, their utilization and the user experience, it just as importantly gives them the power to fix the issues with key optimization technologies.

- **Protect availability of key applications** – Built in QOS technologies that work on a per application basis

and per connection basis help ensure availability of key applications. These ensure availability and performance of transactional applications at any time by protecting them from congestion events, like an FTP download that steps on a voice call, or containing the impact of DDOS.

- **Accelerate and Increase capacity** – Packeteer solutions have a full suite of technologies for accelerating performance — for CIFS, TCP, HTTP and Exchange specific acceleration — removing the issues seen in protocols that weren't designed for the WAN. Using advance dictionary-based compression and fragment caching techniques, Packeteer data reduction technology increase network capacity and offload the WAN.

- **Extend virtual storage/WAFS and branch office network services to the branch** – Single point file caching, with advanced storage techniques to protect data coherency allow you to provide virtual storage in branch offices, with high speed access locally, but maintaining all information at the data center. Key IT services like System Management Server (SMS), Microsoft Domain Controller, Print, DNS & DHCP provide a true branch office box.

There are key questions you should ask when reviewing alternatives to determine if they have the right solution to your network problem.

Question # 1.
What do I want to accelerate?

Do alternatives accelerate all the important applications on my network? ERP applications running remotely? File access? Web pages? E-mail attachments? DHCP requests? Google Earth animations or Napster2? Can they prevent unimportant applications from hurting performance?

Can you help the applications I'm having trouble with? How do your technologies apply to (insert your key applications here)? Can they deal with congestion? How do I know which 'acceleration' technologies you have or which are actually useful in this situation? Any acceleration solution you invest in needs to include some kind of intelligence about the application itself and what is causing the problem; that's a more informed way to approach acceleration.

For that approach to work, the solution needs to include "application intelligence" –able to identify the application and the issues associated with it. In some cases, it must have application specific optimizations like for CIFS or HTTP. Even more specific "file-aware" optimization may be required – in order to take advantage of file-differencing technologies that achieve superior WAN bandwidth savings by sending only changes to files and to manage the data integrity through locks, read and write management. It also needs to enable you to isolate and optimize all the apps that are important to you and causing problems and choose the right tools to solve them.

For example, VoIP traffic is not helped by TCP acceleration (because VoIP uses UDP), caching or payload compression. "Second-pass" acceleration (byte caching type of technologies) won't really help if your users don't keep going back to the same file. Unless your vendor has a method for speeding up DNS requests and optimizing SSL, no amount of HTTP acceleration or caching will make a difference to opening a Web page or reaching a secure server for the first time. Unless your vendor can offer some way to spool prints locally, they won't ever speed up.

Packeteer can optimize the full range of apps. Beyond simply offering Microsoft file and TCP acceleration, compression and caching, Packeteer can optimize UDP (VoIP, ping, etc.), transactional (Citrix, Telnet, SAP, etc.) and secure (SSL, SSH, etc.) applications. In reality, those types of apps make up a significant percentage of the traffic out there — and a significant source of performance problems.

One additional item is the key network services required for applications at the branch. Other vendors just don't have a solution for Microsoft services – they simply have no way of providing Domain Controller, System Management Server (SMS) software distribution and Microsoft print services, DNS & DHCP at the branch location, relying.

As well as having specific application-intelligent technologies to apply to the full range of traffic that's causing performance problems, make sure you can differentiate between apps and apply business rules to acceleration. It's often called "QoS" for applications; the ability to apply QoS at the application layer was invented by Packeteer ten years ago.

The premise behind QoS or "shaping" as we now call it is simple: do you really want to accelerate things that aren't critical? What if the protocol is fine, and it's just a iTunes download that is causing problems? How do you control that? It might not make sense to accelerate or spend cycles trying to compress "bursty" FTP traffic that isn't time-sensitive if it's at the expense of time-critical, latency-sensitive apps. It makes no sense at all to accelerate recreational traffic, and give Napster music downloads equal billing with something business-critical that's being delivered by Marimba or SMS.

In fact, speed is often less about 'acceleration technologies' than ensuring there is the base amount of resource/ bandwidth to make the application work in the first place – so 'availability' is a key foundation to performance. To protect the availability of your key applications, requiring the ability to create policies for key applications (and identify them properly), for key 'sub application processes' (to keep, for example, Citrix print from stepping on Citrix real-time traffic), to provision per flow or per call QOS for important sessions for all application types (for example, voice over IP calls). Solutions should consider that acceleration only makes sense with control and should not try to argue against the case for shaping. Only Packeteer has the level of capabilities for QOS that extend beyond your router infrastructure to ensure availability of key applications.

Remember, you can't control what you can't see. Only Packeteer really knows the applications — with application intelligence that helps all the way through the solution to any WAN performance problems. Only Packeteer has hierarchical auto-discovery; only Packeteer can validate that applications are what they say they are; and only Packeteer allows for **sub**-classification of many applications like Citrix, Oracle, Postgress, HTTP, Web Services, FTP, Exchange and Director Services.

Question # 2.
Acceleration – Is It The Only Consideration?

Your vendor should not try to make you see acceleration in isolation. Packeteer believes that you should never be asked to sacrifice important operational aims like -- availability of key applications, data integrity, scalability, disaster recovery or security — for performance.

If you accelerate with out intelligence, for example, the aggressive TCP applications or CIFS files that you accelerate can very well interfere with sensitive real time traffic in SAP, Oracle, Voice and video. Do you want that?

For data integrity, Packeteer ensures that you maintain the latest copy and that data is written back properly and can be accessed locally even if the WAN is not. Packeteer's *write* back caching, which allows the cache to acknowledge the write updates before updating the master. This is much more sophisticated than the write *through* caching offered by certain vendors. Because many alternatives have to wait for the server back at the data center, your data integrity is very exposed during saves — and that's a huge risk in the event of WAN disruption. The kind of file-awareness and file management you need for an enterprise-class approach to acceleration is not present in any other WAFS solutions.

Scalability is another key consideration. All other WAFS solutions work by maintaining a matching cache for each of the remote sites back at the data center. As you can imagine, this sets a ceiling on how much you can scale. Packeteer WAFS design is fundamentally different and does not require a bloated matching cache in the server, enabling Packeteer solutions to scale in support of many more users. Also, look at the scalability of the equipment itself. The math's not hard: if a data center side box has maximum 4500 TCP connections, and the average user takes up 10 connections, you have a maximum of 450 users — world-wide!

Finally, consider security issues. It's a risk to accelerate everything blindly. Packeteer has implemented WAFS using native CIFS, providing compatibility with security architectures and maintaining a very close relationship with Microsoft. Some other vendors require network managers to turn off SMB "signing", a security construct where Microsoft encrypts the headers of the data blocks used by CIFS, thereby rendering spoofing approaches to acceleration useless.

In summary, Packeteer's comprehensive system solution provides the key technologies to accelerate application performance, but also the tools to ensure their availability and to determine the source of issues, so that you can deploy the right technologies at the right place – simply.

More info and general contact info:
info@packeteer.com for contact information, and
http://www.packeteer.com/company/moreinfo/

Enabling Productivity
through Application Acceleration

Nat Smith
Director of Product Marketing
Expand Networks
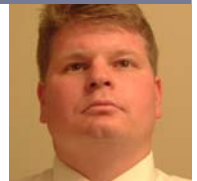
The Mandate

The face of IT organizations is changing. No longer is technical savvy sufficient to successfully run the IT organization; now business acumen is mandatory. Everyone in the enterprise is aware of IT and in almost every case relies on IT for some or all of the tools they use to conduct business. Rather than considering IT to be a small cost center of insignificance, every line of business is aware of the services that IT delivers to them, and is much more likely to scrutinize their funds expensed in the name of IT. IT can no longer afford to simply purchase technology that strikes their fancy, but must account for all expenditures in terms of the mission and goals of the enterprise. In fact, most IT management is compensated on their ability to further the success of enterprise-wide goals.

- Simplicity

- Regulatory Compliance

- Productivity

- Globalization

- Cost Reduction

- Continuity

- Consolidation

In such an environment, the success of IT and its effect on the business is much greater when every IT project or initiative is steeply grounded in the commercial goals of the enterprise.

Needing to account for recurring expenses, as well as the expense of new technology initiatives, forces most IT organizations to build business cases on a regular basis. Business cases steeped in the enterprise commercial drivers are much more likely to be adopted. In particular, the horizontal nature of productivity cuts across all IT technology and even

across all lines of business. In fact, the strength of building a case for productivity improvements can even be related to (and enhanced by) the other commercial drivers.

In order to build a business case based on productivity (and the other commercial drivers), IT must understand a fundamental equation that relates IT performance to business productivity.

Application Performance → **Application Response Time** → **User Experience** → **Productivity**

Improvements in application performance will likely mean faster application response times, which will improve the user experience and make that user more productive. In the same manner, reduction in performance and slower response times degrade the user experience and generally reduce (often dramatically) user productivity.

The enterprise mandates every line of business comply with their goals that drive commercial success. For IT, among other things, this often means enhancing the productivity of the enterprise through superior application performance and response time.

The Challenge

What contributes to poor application response times?

As the enterprise continues to supply the global economy, the workforce is increasingly located outside of headquarters. While there are many factors that contribute to the quality of application response time, there are two large issues that are faced in almost all distributed enterprise environments:

- Congestion

- Latency

Congestion occurs when there is more traffic for the network than the network has capacity to handle. When this happens, packets are delayed or even discarded until sufficient network capacity is available. While most transport can deal with the delay and even discarded packets and still eventually get all of its traffic to the destination, performance (response time) is not so resilient. When delays or retransmissions occur, performance suffers and response time goes up.

Latency has always been the cause of performance degradation and there is really no way to get rid of latency. Latency, in a purely theoretical sense, is simply the time it takes for light to travel between two points. Consider a clear 64KB fiber link between New York and Tokyo, a distance of roughly 10,000 kilometers. Under the best possible conditions (light traveling in a vacuum), a 1 KB packet would take just over 33 milliseconds to make the one-way trip. What if you increased the bandwidth to an OC3 (approximately 155 MB) – how long would that 1 KB packet take to make the trip? The same 33 milliseconds. So what? A 33 millisecond delay for every transaction doesn't seem like that big of a deal to make. Well, the thing is that a transaction is not just a one-way trip from client to server – there are many, many roundtrips required just to transfer the content in one transaction if everything goes perfectly – many more if there are complications. Now consider the fact that the real latency between New York and Tokyo is closer to 170 milliseconds and most application "clicks" involves multiple transactions. This can quickly be the kind of poor response time that not just degrades productivity, but grinds it to a halt.

The Solution

The solution that helps IT reach the goals of the corporation and ensure performance-based productivity for remote users is not based on quick fix or point solution technology. In contrast, the real solution considers the complete IT environment and appreciates the intricacies and integral roles the networking technologies play in the response time equation. The complete solution understands that "1 + 1 = 3" and that the whole is greater than the sum of its parts. The solution is the standards' based integrated Compass platform.

The Compass platform was designed from the start with the end in mind. Rather than a single technology meeting the needs of the market today, but requiring compromises in the future when the needs of the market change, the

Compass platform was designed for the technology today, as well as with the ability to easily integrate the technology of tomorrow. One single platform that offers all of the best technologies individually, but more importantly, intelligently integrates them.

- WAN Optimization

- Acceleration

- Management & Visibility

- WAFS

- Security

So how does the Compass platform resolve congestion issues that degrade response time?

There are two areas where the Compass platform attacks congestion. The first is the stable and low latency compression technology that put Expand on the map. Although not new to the market, this industry proven technology still has relevance today, consistently improving bandwidth by more than 400% the first time it sees the traffic. This is an easy return on investment, often solving complex performance problems even today. The second method of attacking congestion stems from the reality that networks do not natively care about the business. For example, business critical CRM packets are much more important than packets for web traffic to www.espn.com. Without this insight, when

congestion occurs, the network is just as likely to forward the ESPN traffic as it is to forward the CRM traffic – commercially, the wrong decision. The Compass platform's QoS (Quality of Service) solves this problem handily. In fact, without QoS, any other optimization that accelerates traffic is worthless once congestion hits – again, the Compass platform considers the big picture.

But what can the Compass platform do to combat response time issues caused by latency?

The Compass platform cannot eliminate latency – no application acceleration solution can really – but the Compass platform does know how to mitigate the effects of latency, the effects that degrade response time and ultimately decrease productivity. How is this done? With an understanding of how application protocols and transactions work, the Compass platform can optimize those flows. Instead of crossing the WAN several hundred times to complete a full transaction, the Compass platform's application intelligence allows it to seamlessly combine serial round trips into single, parallel round trips. The effect is reducing (or even eliminating) the number of round trips across the latency intensive WAN. This reduction in round trips significantly improves the response time without threatening the integrity of the data. The result is a LAN-like user experience – positive for productivity.

Returning to the beginning, success for IT comes when initiatives and expenditures are tightly aligned with corporate goals and drivers. Guaranteeing productivity through superior application response time is certainly one way to make

that happen. However, much like the Compass platform considers the whole picture, picking products and solutions for your networks that only serve one project, will not strategically position IT for success. The multifaceted Compass platform has been shown effective across multiple IT initiatives, including some of the more popular ones today:

- Server Consolidation
- Disaster Recovery
- Web Services
- Server-based Computing
- Voice and Data Convergence
- Satellite Communications
- Enabling MPLS

Only an integrated, single box solution will meet the strategic and cost management goals of the corporation at the remote office – but for Expand, while still meeting these real needs, the passion is in application acceleration that enables productivity.

Expand Networks
info@expand.com
+1.888.892.1250
+1.973.618.9000
www.expand.com

CITRIX

Citrix
WANScaler™

Gordon Smith
Application Networking Group
Citrix

Introduction

Citrix WANScaler accelerates applications that are most problematic for remote workers by up to 50 times, including applications such as file sharing and bulk data transfer. WANScaler products, which can be deployed either as hardware appliances or as software-based remote clients, use a *symmetric* approach to increasing application performance over the WAN. That is, these products control both ends of the WAN link to provide the degree of control over traffic that is required to achieve the highest throughput and the fastest application response time.

Designed for speeding business-critical applications utilized over WANs, WANScaler can turbo-charge a broad variety of applications. A few examples: WANScaler has been shown to improve the performance of Citrix Presentation Server traffic by as much as 5x for end users, especially for remote printing and file access. Backup operations with Network Appliance's SnapMirror or with Double-Take from Double-Take Software that may have previously taken half a day can now be completed within a half hour. Microsoft Word files hosted at a U.S. corporate datacenter that used to take three minutes to open by a branch employee in Beijing can now be accessed in seconds.

WANScaler enables virtually any TCP-based application that sends data between geographically dispersed offices to run faster, increasing WAN throughput and improving application response times by as much as 750%. In this paper, we first take a look at the internals of Citrix's breakthrough WAN acceleration technology. Next, we discuss an important (and easily overlooked) aspect of WANScaler, its network transparency, as well as some of its other key benefits. We conclude by highlighting WANScaler performance and scalability testing results validated by Miercom.

Technology

The cornerstone technology of WANScaler is the AutoOptimizer™ Engine, the core software which automatically and dynamically applies to each data flow the best combination of performance boosting techniques – depending on the application, the data and the network conditions.

Here is a quick overview of the key performance enhancement techniques used by WANScaler:

- **TCP Flow Control**. Instead of TCP's general-purpose link congestion control and packet retransmission algorithms, WANScaler appliances incorporate adaptive TCP algorithms optimized for WAN links. This allows data to be transferred at rates close to the theoretical throughput limit of the link, regardless of the distance between the sender and the receiver, even over links with high latency and high packet loss.

- **Multi-Level Compression**. WANScaler multi-level compression features algorithms that incorporate both first-pass and long-history schemes, which work at the intra-session and inter-session level. The compression system can attain compression ratios as high as 3500:1, providing extreme improvements in transfer time.

- **CIFS Acceleration**. WANScaler CIFS acceleration improves Windows file sharing by reducing the client-server messaging overhead that is a by-product of using CIFS for Windows-based workflows. In addition to simple read-ahead and write-behind operations that are more commonly found in other CIFS-acceleration implementations, WANScaler CIFS also delivers metadata caching, directory browsing and specific Microsoft Office application optimizations that work together to improve all aspects of an end-user's file sharing and browsing experience.

- **Service Class Policies**. WANScaler appliances allow granular management of enterprise traffic with service class policies. Traffic segmentation can be achieved by

identifying applications, business groups, application servers or remote sites.

Products and Solutions

WANScaler can benefit the delivery of a wide variety of applications over the WAN, providing improved performance for remote users throughout the enterprise. The diagram below shows some of the typical solution environments.

Typical Solution Environments
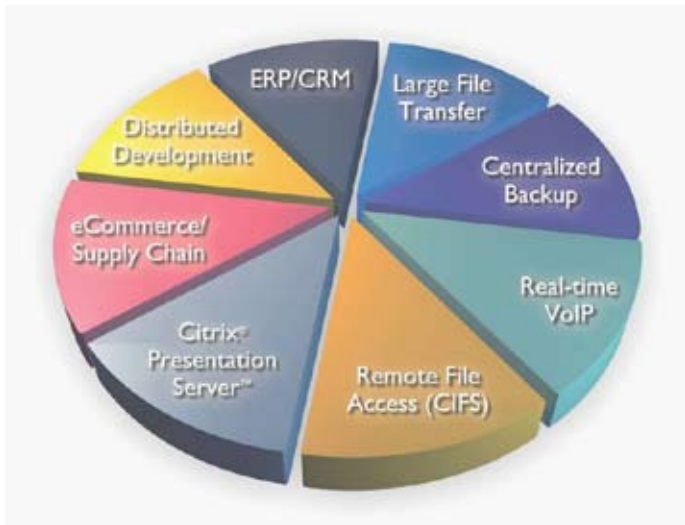
Citrix WANScaler accelerates a wide variety of application environments.

A couple of common ones include:

- **Citrix Presentation Server:** Citrix Presentation Server is the #1 application virtualization solution for providing any user with secure access to client/server applications from anywhere using any device or connection. Presentation Server enables IT to centrally deploy and manage line of business applications while providing secure, on-demand access to these resources, and delivering the lowest total cost of ownership (TCO), highest security, best performance and greatest scalability. Although ICA, the protocol used for Presentation Server sessions, is highly optimized and is designed to operate over bandwidth-constrained links, because it is based on TCP it suffers from many of the limitations imposed on other interactive TCP-based application. WANScaler solutions provide the end-to-end control over the WAN link that guarantees reduced end-user wait times regardless of WAN latency, packet loss or bandwidth contention

caused by multiple competing applications. In internal tests, WANScaler has been shown to improve the performance of Presentation Server by as much as five times for end users, especially for file printing and local file save tasks that require the transfer of large volumes of data over the WAN.

- **Centralized backup / disaster recovery:** Backing up huge, multi-gigabyte files over an inefficient WAN link can mean that the operation won't complete in a timeframe that meets the business' needs. If a backup operation that should ideally be completed overnight spills over into regular working hours, remote users of the WAN may find that their applications have become very slow. A typical scenario might be one where the backup of a 20GB file over a 20Mbps coast-to-coast WAN link requires 16 hours to complete over a regular WAN link, but only 15 minutes over a link optimized by WANScaler appliances on each end.

The Importance of Transparency

All WAN acceleration vendors advertise their products as being "transparent": neither the client nor the server is aware of the WAN acceleration device, and applications do not have to be modified to benefit. We call this "application transparency". When WAN acceleration vendors state that their products are "transparent", that's typically all that they mean.

Almost all WAN acceleration vendors, however, have a dirty little secret: they break another equally important form of transparency – **network** transparency.

Network transparency is defined simply as preserving the key characteristics of traffic flows on a network. These characteristics include IP addresses, port numbers, and the protocol (TCP or UDP). A network-transparent product does not change these fields. An HTTP GET request destined for port 80 on web server 10.0.0.1 crosses the WAN as a TCP-based HTTP request with the destination address 10.0.0.1 and port 80, even with WANScaler appliances deployed at each end of the WAN link.

Why is network transparency important? Only network-transparent products interoperate with other networking products such as firewall, monitoring and QoS devices, and ensure that no configuration or deployment changes are required for continued operation of these devices.

By default, all WAN acceleration and optimization vendors, other than Citrix, use TCP or UDP tunnels, whether they are explicitly configured by the administrator or automatically configured by the system. All application traffic is multiplexed over a single TCP or UDP connection between WAN acceleration devices on each end of the WAN link, and the addresses of packets flowing between WAN accelerators are replaced with the addresses of the devices themselves. These proprietary traffic streams are not understood by any device operating at layer 4 and above.

In typical enterprise networks, with data centers serving multiple branch offices, tunneling does not scale. Being application and network transparent, WANScaler solutions are ideal for enterprise deployments of any size. The network-transparent nature of WANScaler appliances, combined with their ease of configuration, minimizes the time it takes for initial deployment and reduces the effort required to manage the appliances on an on-going basis.

Technical benefits of WANScaler

- **Ease of Deployment and On-going Management**. The ability to deploy a WAN acceleration system *without* breaking existing network monitoring tools (such as NetFlow), firewall security rules, or client-server applications is a key distinctive benefit of the transparency of Citrix's WANScaler. Auto-optimization is another key benefit, dramatically simplifying on-going management. The **AutoOptimizer Engine** dramatically increases the efficiency of WAN optimization and application acceleration, even under dynamically changing conditions across a wide variety of types of networks.

- **Extended Reach to the Entire Enterprise**. The WANScaler architecture allows a flexible choice of acceleration endpoints. At the datacenter, the endpoint will always be a WANScaler appliance; at the remote end of the link, that endpoint can be either an WANScaler branch office appliance OR a WANScaler software-based client. With this flexibility, ALL remote users in the entire enterprise can cost-effectively benefit from the Citrix WAN acceleration solution -- whether these users are in large branch offices where an appliance would be deployed, or in small branch offices, home offices or on the road, where a software client is the right answer.

- **Higher Performance**. WANScaler increases transactional performance through features such as automatic

scheduling of interactive traffic ahead of bulk traffic and superior retransmit logic. It also speeds bulk transfer with extremely efficient compression (up to 3500:1). And virtually all TCP applications benefit from flow control, with link utilization frequently an order of magnitude higher than with stock TCP.

Performance and scalability results

Applications can be as variable as snowflakes, due to such factors as first-pass data compressibility, repetitive data patterns, the "chattiness" of the application, network link conditions such as latency and loss, and so on. WANScaler will automatically use the appropriate mix of acceleration techniques to drive the best performance for that application and that link.
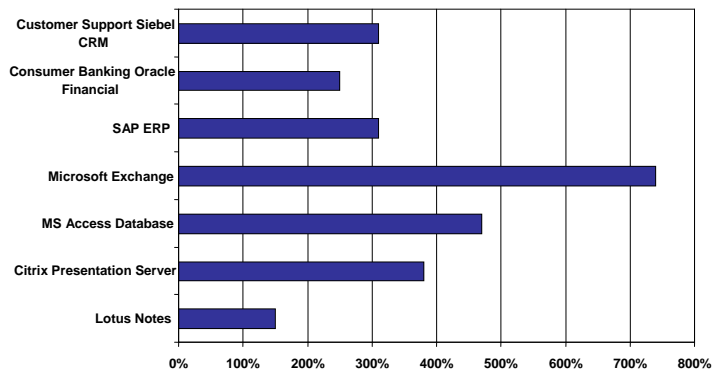
Citrix WANScaler accelerates response time by as much as 750% for common application types.

Conclusion

WAN acceleration technology has quickly become a "must have" for the distributed enterprise. Improved WAN performance can reduce costs while dramatically improving remote user productivity – including remote users in branch offices, at home or on the road. Key criteria in vendor selection must include factors such as transparency, ease of deployment and management, high performance, and the flexibility to be able to deploy either a hardware appliance or a software client solution to extend the benefits of improved application delivery to every user throughout the extended enterprise.

For more information, see www.citrix.com or contact us at: 1.800.393.1888 toll free 1.954.267.3000 main

Meeting the Challenges
of Today's Distributed Enterprise

by Michael Banic
Director, Product Marketing
Juniper Network

Today's corporate enterprise is at a crossroads.

Multiple conflicting trends are currently underway, leaving businesses struggling to meet the challenges of a constantly shifting landscape. The ability to adapt to these changes while remaining flexible enough to meet similar challenges in the future will determine who survives and thrives — and who falls hopelessly behind.

These seemingly contradictory trends can be distilled down to three interrelated issues: dependence on applications to support business processes; globalization; and server centralization and data center consolidation. These trends are on a collision course, threatening to drive up costs, lower productivity and create serious compliance liabilities.

The first issue reflects a truism of our time: business runs on applications. From ERP to CRM to collaboration tools, businesses are dependent on application software to get the job done. By extension, this makes the corporate network critical to business processes; if the network isn't enabling optimal application performance, the business suffers.

Globalization compounds the application reliance issue by scattering employees — all of whom depend on uncompromised application performance — around the globe. The corporate network must expand apace to maintain connectivity, forcing communication and collaboration to take place over long-distance wide-area networks (WANs).

Server centralization and data center consolidation deliver the final blow. In an effort to reduce costs, simplify backup and recovery, and improve regulatory compliance, many businesses are redeploying critical file and application servers from branch offices to central data centers. As a result, employees must access these files and applications — the lifeblood of the business — over long-distance, fixed-capacity WAN links.

These trends — which lead to a growing reliance on the WAN — bring to light four primary impediments to effective application performance over wide-area links: a lack of sufficient bandwidth; latency imposed by distances between distributed users and centralized resources; contention among applications for space on the network; and the absence of systemwide management and reporting tools. Predictably, these four conditions causes performance to suffer, productivity to plummet, and the business to slow to a crawl.

Meeting the Challenge

To keep pace with these evolving needs, IT managers around the world are working feverishly to evolve and optimize their WAN infrastructures to provide secure and assured application delivery that keeps workers working, servers serving, and business productivity growing.

WAN application acceleration and optimization solutions are playing a critical role in this effort. Such solutions help businesses make more efficient use of their WAN resources and deliver LAN-like response times to globally distributed users of centralized applications.

However, all WAN optimization and application acceleration solutions are not created equal. According to a report by the Burton Group*, the basic WAN optimization technologies are compression, which squeezes out repeated sequences in a dataflow; caching, which stores frequently requested files; and protocol modification, which improves the performance of protocols over high-latency WAN links. These techniques, says the report, are often coupled with quality of service (QoS) capabilities to provide a level of control for prioritizing application traffic.

Juniper Networks, a premier provider of secured and assured networking solutions, takes this definition a step further. Juniper contends that a truly complete WAN opti-

mization and application acceleration solution must integrate these capabilities — and more — on a single platform to address the bandwidth, latency, contention and manageability issues posed by the WAN environment. Working together, these attributes — compression and caching, acceleration, application control, and visibility — not only address the most pressing needs of today's distributed enterprise, but also help position customers to meet continuously evolving needs in the future.

In addition to these integrated WAN application acceleration solutions, providers must also have an eye to the future, providing a foundation that considers future developments and provides a path for meeting those challenges as they arise.

The Juniper WAN Optimization and Acceleration Platforms

This is precisely the approach Juniper Networks has followed in its WAN optimization and application acceleration solutions. The products — the Juniper WX and WXC WAN application acceleration platforms — are based on a unique WAN acceleration (WX) Framework that outlines the attributes defined above and describes how specific features of the WX and WXC platforms meet those requirements.

Compression and Caching: To satisfy the compression and caching component, the WX Framework integrates memory-based Molecular Sequence Reduction™ (MSR™) technology, which increases WAN throughput up to 10 times by eliminating repeated data patterns from traffic flows traversing the WAN. MSR compression is complemented by Network Sequence Caching, which uses hard disks to store larger repeated patterns last seen days or even weeks earlier to increase throughput up to 50 times.

Acceleration: Acceleration is delivered in the form of Packet Flow Acceleration™ (PFA™) and Application Flow Acceleration™ (AppFlow™) technologies, two Juniper-specific features of the WX Framework that distinguish the WX and WXC platforms. The PFA techniques — including Fast Connection Setup and Active Flow Pipelining — combat the effects of TCP latency by accelerating connection setup and substituting a more efficient transport protocol across the WAN to dramatically improve performance. For lossy networks, an additional PFA feature — Forward Error Correction — makes use of recovery packets to reconstruct lost data, eliminating the need for retransmissions.

Application-specific acceleration is delivered via the AppFlow technology, which augments PFA by accelerating applications that are constrained by their underlying protocols — specifically MAPI (used by Microsoft Exchange), CIFS (used by Windows file services) and HTTP (used by Web-based applications). The AppFlow feature compensates for these chatty protocols by pipelining data blocks and web objects across the WAN simultaneously, rather than sequentially, providing users accessing these centralized applications a more LAN-like experience that dramatically reduces wait times and improves productivity.

Application Control: The WX Framework delivers application control capabilities via QoS and Policy-based Multipath capabilities. QoS is combined with bandwidth-management tools that allow users to prioritize flows to ensure sufficient bandwidth is always available for critical traffic or delay-sensitive applications such as voice over IP (VoIP). Policy-based Multipath complements QoS by enabling users to assign designated application flows to a specific WAN link when more than one path is available, ensuring critical traffic is sent over more reliable private links while less sensitive traffic is relegated to the Internet.

Visibility: Finally, the WX Framework integrates visibility and reporting functions that arm IT with tools for monitoring, understanding and controlling application performance over the WAN. The WX Central Management System™ (WX CMS™) software provides unprecedented system-wide visibility into application performance while WebView device management enables IT to configure and manage individual WX and WXC platforms from a central location.
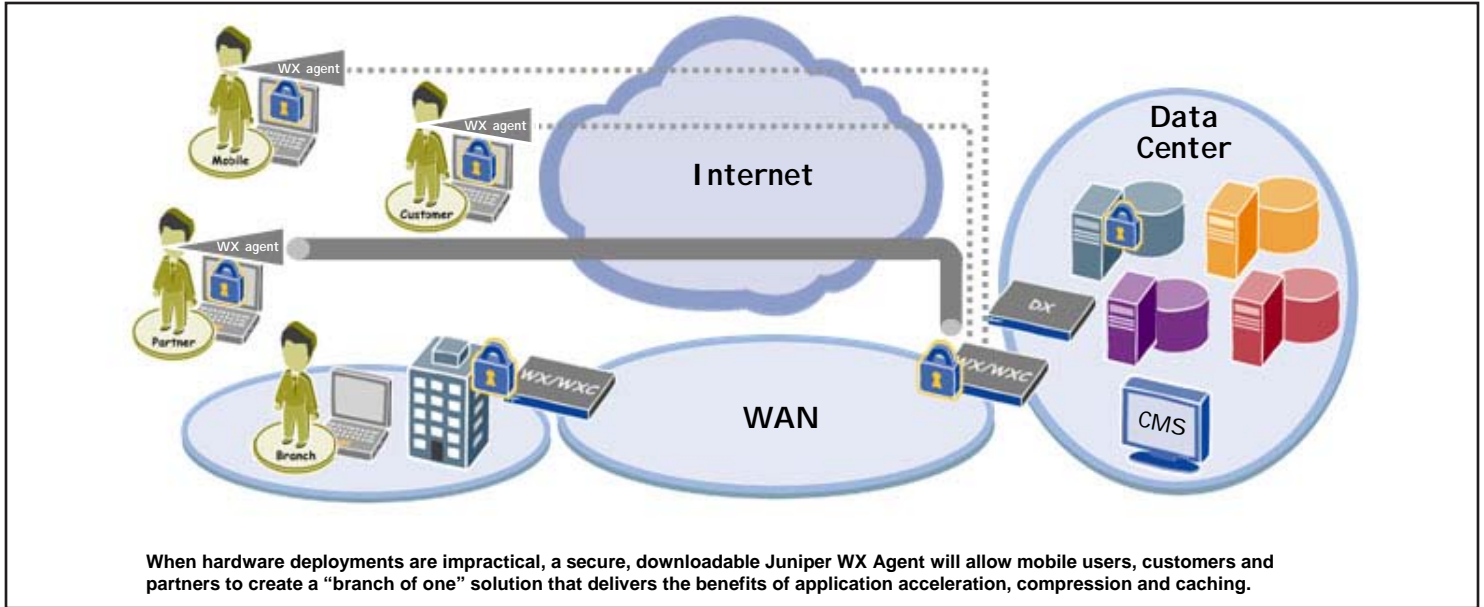
Working together in the WX and WXC platforms, each of these elements interact to dynamically adjust and improve each other's capabilities and deliver the industry's most complete WAN optimization and application acceleration solution.

Looking to the Future

While the WX Framework provides IT with the greatest degree of application acceleration and WAN optimization available today, it also represents a sound foundation for delivering even more powerful capabilities in the future.

Over the years, Juniper has established itself as an industry visionary by consistently delivering increasingly powerful WAN optimization and application acceleration solutions that

When hardware deployments are impractical, a secure, downloadable Juniper WX Agent will allow mobile users, customers and partners to create a "branch of one" solution that delivers the benefits of application acceleration, compression and caching.

build upon earlier capabilities to provide a flexible and evolutionary platform that meets current and future business needs. That tradition continues with the recent announcement of an application acceleration strategy that defines new capabilities designed to address the challenges associated with limited bandwidth, latency, application contention and system-wide management in the face of globalization, server centralization and data center consolidation.

The strategy recognizes that today's enterprise supports a wide range of local, remote and mobile users — including employees, partners and customers — using a variety of applications and web services to streamline business processes, improve productivity, and create true customer value and competitive differentiation. To assist in this effort, the strategy calls for extending the WX Framework's AppFlow capabilities to a broad range of IP- and Web-based business applications, including SAP and SQL, in addition to Microsoft Exchange and file services.

Likewise, the vision also defines a method for extending these capabilities to mobile users, partners and customers by securely deploying a downloadable WX agent to user desktops. This agent will work in tandem with WX/WXC platforms in the data center to speed delivery of business applications for remote and mobile users.

The strategy also addresses the acceleration of SSL-encrypted flows with web-based applications by including WX/WXC platforms in the "trust model" — making them

"friends in the middle" — to accelerate encrypted web applications rather than allowing them to pass through untouched like other WAN optimization solutions. And finally, to meet demands from remote and mobile users for a more LAN-like experience, the strategy also includes plans to deliver web content caching functionality and adaptive content processing features at branch locations.

Juniper Networks:
Solutions for Today and Tomorrow

As a leading provider of the industry's most complete WAN optimization and application acceleration solution, Juniper Networks has shown that it understands the problems businesses are facing. And Juniper has consistently delivered solutions that address those problems with unique and innovative technologies designed to maintain employee productivity, regardless of their location around the globe.

Juniper delivers those solutions today — powerful application acceleration platforms that will provide a foundation for delivering LAN-like performance to local, remote and mobile users for years to come.

**For more information about the
Juniper application acceleration platforms, visit
http://www.juniper.net/products/appaccel/wan/
or call Juniper at 866.298.6428 (inside the U.S.) or
978.589.0500 (outside the U.S.).**