

Solution Brief

Accelerating Web Applications across the WAN: How Juniper Networks Optimizes Web-Based Applications

Whether developed in house or purchased off the shelf, web-based applications have become a staple for enterprises of all sizes and are increasingly popular because of the business and IT benefits they deliver. Web-based applications are ideal for distributed and mobile workforces, giving users access to applications anywhere, anytime, via browser-based devices ranging from desktop and laptop computers to PDAs and cell phones.

Equally important, web-based applications have enabled much greater communication among enterprise employees, customers, suppliers, and other strategic partners, facilitating new business processes. For example, by web-enabling an order-entry system originally designed for in-house sales representatives, an enterprise can provide self-service order entry for customers. Similarly, many organizations have designed web portals that allow employees to perform a variety of transactions, from entering expense reports and benefits changes to checking pay slips.

Industry experts agree that virtually every class of application will eventually offer some sort of web-enabled delivery capability, providing universal access to employees, partners, and customers alike across the Internet.

Beyond the business benefits of web-enabling applications, IT benefits include the ability to bring new users online at minimal cost, without client device configuration and with little or no training.

By providing a web front end, IT can make the same data available to users whether they're using a Windows, Mac, or Unix-based system; a tablet PC; or any other web-enabled device. Web-based applications can be deployed on the Internet, on a company intranet, or both, giving users the flexibility to access them via public or private networks.

The Performance Challenge

Unlike their client-server counterparts, web-based apps tend to be graphically intensive, featuring images, streaming audio/video, and embedded files such as Java applets. This richer content requires greater bandwidth – as much as a ten-fold increase, by some accounts – putting stress on existing network links:

According to Gartner, Inc., “Depending on the application, business goals may require the network infrastructure to add what traditional network managers might consider to be ‘excess capacity.’”*

In addition, HTTP – the protocol used by web-based applications to request and gain access to services and/or resources – is excessively “chatty,” which can dramatically slow web applications running over a WAN. For example, the first time a browser loads a web page, the browser retrieves the HTML text file, which contains the formatting and text for the page. The browser can start to display the text immediately, but the HTML file will likely contain references to many additional objects

– such as jpg and gif images, javascripts, or style sheets – that need to be retrieved. HTTP requests and receives these objects individually, one at a time.

While web caching does reduce bandwidth demands, it does nothing to accelerate page display times since confirming object freshness remains subject to delays from WAN latency.

These back-and-forth exchanges happen serially, resulting in “ping-pong” behavior. Users accessing a web server via a LAN won't notice any delay from this protocol behavior, but for users crossing the WAN to a web server, the latency will degrade the application's performance. The higher the WAN latency and the greater the number of objects associated with the URL, the slower the page builds. Casual Internet surfers may find the delay tolerable, but users trying to complete specific job functions will likely find these slow response times very frustrating.

After a browser has loaded a web page, it stores any objects that are cacheable in its local cache; however, caching doesn't eliminate the latency delay. The browser is typically required to contact the server to verify – for each and every object – whether the cached object is still current before using it. So while the object itself, if current, may not have to be transferred across the WAN again, saving on WAN bandwidth, the required protocol exchange means building the web page is still subject to delay as a result of WAN latency.

Web Acceleration Options

Web caching, whether on the client or via a proxy device, is one of the key methods to improving the performance of web-based applications running over a WAN. Web caching helps reduce bandwidth consumption by cutting the number of objects that traverse the network and delivering responses from a local cache.

While web caching does help reduce bandwidth demands once a web page is cached, it does nothing to accelerate initial page loads or to refresh objects that have changed since being cached. Nor do proxy caches distinguish one type of web traffic from another, so order-entry traffic must contend with traffic from employee payroll deposit inquiries, for example. Consequently, the performance of mission-critical web applications can degrade over WAN links when non-mission-critical traffic is high – for example, at mid-month and the end of the month as employees check payroll status.

In remote offices where no IT staff is present, IT may be reluctant to deploy proxy caches due to their cost and support overhead and their narrow focus on web traffic alone. Yet remote offices are often the sites that could most benefit from web acceleration because of their relatively slow WAN links.

What remote offices need is a WAN optimization solution that combines HTTP acceleration, caching, URL-level Quality of Service (QoS), and data compression for optimum bandwidth utilization and latency reduction. The solution must also be easy to deploy and require no ongoing support so that remote offices without IT help can use the devices. Finally, web acceleration platforms should also provide benefits for overall WAN traffic rather than only web-based applications.

Juniper Networks delivers powerful application acceleration and WAN optimization technologies that increase the performance of web applications.

Ultimately, a WAN optimization solution for web acceleration must:

Reduce the negative impact of latency on web-based applications

Latency is a significant problem for web-based applications operating across a WAN. To increase the productivity of employees in remote offices, a WAN optimization solution with web acceleration must reduce the delay associated with HTTP's ping-pong behavior.

Reduce the bandwidth impact of web applications

A WAN optimization solution must address the bandwidth-intensive nature of web applications both through compression and caching. Compression reduces the volume of both static and dynamic web traffic traversing the WAN, while caching limits the traffic across the WAN to initial page downloads and object refresh messages.

Ensure sufficient bandwidth for mission-critical web applications

As more and more applications are web-enabled, the volume of traffic traversing the WAN naturally increases, and contention for bandwidth can become an issue. A WAN optimization solution should provide bandwidth-management and QoS tools that enable IT to allocate bandwidth and prioritize traffic across the WAN, ensuring that mission-critical web applications such as order entry and inventory management aren't negatively impacted by employees checking web-based HR records, for example.

Accelerate overall WAN traffic

Although web-based applications are increasingly common, they do not represent all traffic traversing WAN links to remote offices. In addition to web traffic, a WAN optimization solution must accelerate the full range of application traffic to enable maximum employee productivity and efficient WAN utilization.

Provide visibility into web application performance

To troubleshoot web application problems, understand the traffic patterns associated with these flexible platforms, and provide feedback to executives about WAN performance, the web acceleration platform should provide in-depth reporting and monitoring for web and other applications traversing the WAN. IT needs to understand key metrics such as application performance, top talkers, application mix, and bandwidth-consumption patterns.

Operate transparently, with little or no IT intervention

To remove the administrative and management burden on IT associated with remote proxy cache servers, a WAN optimization solution should accelerate web traffic transparently, without requiring changes to clients and servers or ongoing IT support. A web optimization product must operate seamlessly within an existing WAN environment, regardless of the technology or topology in use. In addition, it should not conflict with existing web-caching platforms or other server-acceleration technologies.


The Juniper Networks Solution

Juniper Networks delivers a complete family of application acceleration platforms designed to improve application response times within central sites, to branch offices, and for remote users. Two members of that family – the WX™ and the WXC™ application acceleration platforms – provide powerful compression and caching, acceleration, application control, and visibility capabilities that dramatically improve the performance of web-based applications.

The WX and WXC platforms are based on the unique WX Framework™, which integrates interdependent technologies that work together and continuously share information to provide knowledge of the WAN link from both ends. This distributed stateful intelligence enables the WX and WXC platforms to tune their settings in response to changing conditions, ensuring maximum application performance in a dynamic WAN environment.

The WX Framework's integrated technology set includes:

- The patented Molecular Sequence Reduction™ (MSR™) technology, next-generation compression capabilities based on DNA pattern matching that increase WAN capacity four-fold and more by reducing WAN transmissions.
- The patent-pending Network Sequence Caching technique, which reduces the overhead of repeated file transfers by identifying large redundant data patterns and storing them on hard drives for long periods of time, thereby improving data reduction results.



- Quality of Service (QoS) and bandwidth-management tools that let IT prioritize mission-critical applications and time-sensitive traffic relative to other traffic types.
- Packet Flow Acceleration™ (PFA™) technologies that enable IT to accelerate application performance by reducing the impact of latency on TCP-based applications.
- Application Flow Acceleration™ (AppFlow™) techniques that optimize application-specific protocols for web traffic, Exchange, and Microsoft file sharing, improving the response times of these applications.
- Policy-Based Multipath™ (Multipath™) selection technology, which allows IT to define which applications follow which paths when multiple WAN links are available and supports automatic migration to an alternate path if performance degrades or a link fails.
- Management and monitoring tools such as the WX Central Management System™ (CMS™) software, which provides a single, unified view into application and network behavior; automates device configuration and deployment; and produces executive-level reports that aggregate network performance information.

The WX Framework integrates key technologies that work together and influence each other, providing IT with distributed stateful intelligence about their WAN links and applications.

Customer Success Story:

Global Manufacturing Company

Business Benefits:

- Accelerated performance of web-based inventory management application two to four times
- Delivered LAN-like response times to remote and satellite-connected offices
- Enabled global consolidation of critical applications in U.S. data center

A global manufacturing company was having a web problem.

The company relies heavily on its corporate intranet to share information with thousands of employees at more than 80 facilities around the world. More importantly, the intranet also supports a proprietary inventory-management application, which tracks literally millions of parts across the globe.

The application, housed at the company's U.S.-based data center and accessed around the clock, was suffering significant performance problems. Remote users – particularly those in satellite-connected offices – were experiencing unacceptable delays, with latency hovering at 500 ms. The company needed a solution that could restore user productivity and deliver LAN-like performance over the WAN.

After deploying Juniper Networks WX application acceleration platforms with AppFlow technology at the U.S. data center and the 10 largest remote offices, the company noticed a dramatic improvement. Response times for the custom inventory application improved two to four times for users in the 10 test offices, including three satellite-connected facilities.

The results were so impressive the company intends to roll out WX platforms to the remaining offices in a phased deployment. Even better, the company is now confident it can proceed with a planned server consolidation project without sacrificing performance for remote users.

With the WX Framework, the WX and WXC platforms provide branch office users with WAN access to web content at near-LAN speeds through a rich set of web optimization capabilities that:

Reduce the negative impact of latency on web-based applications

Juniper Networks' key innovation for accelerating web applications is its "pre-fetch" technique. As users request web pages, the WX and WXC devices build tables of the URLs and the objects associated with them. When a user's browser initiates a request to view a web page, the device closest to the user sees this request and asks the device at the other end of the WAN link, nearest the web server, to not only fetch the object being requested but also to "pre-fetch" all objects associated with that page, in advance of the user's browser actually requesting them. The server-side device sends the objects across the WAN in quick succession to the user-side platform. By the time the user requests the objects, the local WX or WXC device can send them to the user at LAN speeds.

This "pre-fetch" technique vastly reduces the impact of latency on web applications, eliminating the delays associated with the "ping-pong" behavior of retrieving objects one at a time.

Reduce the bandwidth impact of web-based applications

In addition to accelerating page loads through "pre-fetching," the MSR and sequence caching technologies reduce the amount of web traffic traversing the WAN link, providing up to a 20-fold increase in WAN capacity. MSR compression and sequence caching technologies identify variable-sized, repeating data patterns and replace them with a short flag. The sequence caching technique has the added advantage, over traditional caching, of being able to recognize repeated patterns even when a file has changed, so it can reduce both static and dynamic content.

Ensure sufficient bandwidth for mission-critical applications

As more applications become web-enabled, mission-critical web traffic will be contending with less important traffic for WAN bandwidth. For example, order-entry transactions may be crowded out when employees view web-based payroll information. Likewise, mission-critical web apps can be impacted by other non-critical applications such as e-mail.

To ensure that key web applications get sufficient bandwidth across the WAN, the WX and WXC platforms include easy-to-use, template-based bandwidth-management tools, including QoS and bandwidth allocation. With these tools, IT staff can prioritize applications and allocate bandwidth among different applications, defining minimum and maximum throughput levels for mission-critical web-based applications and delay-sensitive applications, such as Voice over IP (VoIP).

The WX and WXC devices perform layer-7 classification and so can distinguish web traffic at the URL level, allowing IT managers to delineate traffic for each web application. Alternately, IT can classify traffic based on other factors, from IP headers and TOS/DiffServ information to the actual payload content. In addition, the WX and WXC platforms have the ability to see and preserve QoS markings applied by other devices elsewhere in the network.

The WX Framework enables WAN access to web content at near-LAN speeds.

Accelerate overall application performance

Beyond simply accelerating HTTP, the WX and WXC platforms accelerate and increase the WAN bandwidth for all application traffic. The devices provide broad-based acceleration for any TCP-based applications with the PFA technology. The platforms also support AppFlow techniques, which – like the HTTP acceleration discussed here – improve response and file-transfer times for Microsoft Exchange and the Microsoft file services Common Internet File System (CIFS) protocol. By optimizing the full range of WAN traffic, the WX and WXC devices deliver significant application performance improvements and productivity gains for users in remote offices.

Provide visibility into web application performance

The WX and WXC platforms provide a set of monitoring and management tools that enable IT to maximize web acceleration and overall WAN optimization with minimum time. These tools include WebView, which provides device-level management, and the WX CMS software, which enables system-wide control and monitoring and enables automated deployment of WX and WXC platforms. With the WX CMS software, IT can configure and manage multiple platforms from one central location and gain a unified view into metrics such as web and other application acceleration, compression and caching performance, WAN utilization, QoS, and bandwidth allocation.

These real-time WAN monitoring and reporting capabilities allow IT staff to identify top talkers, gather WAN and application statistics, and perform troubleshooting and problem-resolution activities such as packet capture. The intuitive GUI provides instant visibility into WAN performance, including both real-time and historical views. Application acceleration reports provide insight into the performance of different web-based applications as well other key business traffic.

Operate transparently, with little or no IT intervention

The WX and WXC platforms provide application acceleration and WAN optimization well beyond HTTP acceleration, delivering the benefits of web applications over existing WAN links while helping to reduce overall WAN expenses and IT operational time. The devices accelerate applications in a non-intrusive manner in any existing network environment; they're transparent to existing applications, networks, routers, and WAN interfaces, as well as to IPsec VPNs, MPLS, firewalls, and encryption technologies. No changes are required to clients at remote sites or to servers within the data center.

The WX and WXC platforms are also robust. However, should one of the devices fail, it automatically switches to bypass mode, allowing all traffic to flow across the WAN untouched. This approach ensures uninterrupted operations and eliminates the WX or WXC platform as a potential point of network failure.

Unlike many web caching products, the WX and WXC platforms can be installed and configured at centralized sites in less than 10 minutes. Remote devices are even simpler to deploy. IT defines configurations for remote devices via the WX CMS software. Personnel at remote sites simply connect the devices to electrical power and an Ethernet network and the rest of the configuration is automatic. When the remote device boots up, it automatically requests a temporary address, looks up the central WX CMS server, downloads its configuration, and begins operation. The WX CMS software also automates license management.

Juniper Networks: Accelerating Web Traffic as Part of WAN Optimization

Web-enabled applications offer a myriad of benefits, from flexible user access to ease of deployment. However, because of their rich content and HTTP's ping-pong behavior, web applications often perform poorly over wide-area connections. To realize the full benefits of web applications and get the most bang for their buck, enterprises need an easy-to-use, full-featured WAN optimization solution that includes specific acceleration for web applications.

With its suite of application acceleration and WAN optimization technologies, Juniper Networks gives remote offices LAN-like access to web servers, eliminating the need for remote caches. The WX and WXC platforms also ensure that WAN bandwidth is used efficiently and that critical applications get the needed bandwidth. With the WX and WXC application acceleration platforms, enterprises accelerate web-enabled and all other applications operating across the WAN.

* Gartner Research "New Application Architectures Will Impact Networks" by Ray Valdes, March 2004



CORPORATE HEADQUARTERS
AND SALES HEADQUARTERS
FOR NORTH AND SOUTH AMERICA

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888-JUNIPER (888-586-4737)
or 408-745-2000
Fax: 408-745-2100

www.juniper.net

EAST COAST OFFICE

Juniper Networks, Inc.
10 Technology Park Drive
Westford, MA 01886-3146 USA
Phone: 978-589-5800
Fax: 978-589-0800

ASIA PACIFIC REGIONAL
SALES HEADQUARTERS

Juniper Networks (Hong Kong) Ltd.
Suite 2507-11, Asia Pacific Finance Tower
Citibank Plaza, 3 Garden Road
Central, Hong Kong
Phone: 852-2332-3636
Fax: 852-2574-7803

EUROPE, MIDDLE EAST, AFRICA
REGIONAL SALES HEADQUARTERS

Juniper Networks (UK) Limited
Juniper House
Guildford Road
Leatherhead
Surrey, KT22 9JH, U. K.
Phone: 44(0)-1372-385500
Fax: 44(0)-1372-385501