WHITE PAPER

WAN OPTIMIZATION BEST PRACTICES FOR GOVERNMENT AGENCIES

Juniper Networks Innovative WAN Application Acceleration Solutions Enable Government Agencies to Make the Most of their Existing WAN Resources

Table of Contents

Table of Figures

Executive Summary

Government organizations at all levels are employing an increasing variety of electronic applications being used to help facilitate effective communications and enhance department productivity. These applications may range from simple email systems to more complicated Web-based servers that help enhance an e-government's reach to workers and citizens. At the same time, government IT departments are working to ensure acceptable application performance for all users both local and remote. However, WAN bandwidth constraints pose a serious challenge to meeting this goal.

This document will examine ways government IT managers within all levels of government can efficiently utilize scarce WAN resources to effectively accommodate operations traffic and the growing number of electronic applications. We will discuss how to implement a reliable, comprehensive WAN optimization solution that enables the delivery of optimized application performance without compromising day-to-day operations with upgrades, redesigns or incremental investments.

Introduction

The ultimate goal of all government network infrastructures is to deliver applications effectively to users. Continuous performance improvements in PCs and LAN infrastructure have made it easy for IT to enable effective application delivery within a building or campus. But maintaining an acceptable level of application performance across the WAN proves a consistent challenge.

To cope with poor performance across the WAN, government organizations have had to make a number of accommodations. Government IT organizations have been forced to proliferate data centers and server hardware around the world, install applications locally within branch offices, and endure the high cost of increased WAN bandwidth— typically second only to staffing as IT's highest expense—just because applications run slowly or simply cannot operate across the typically poor performing WAN link.

As government functions demand a more tightly integrated globalization that pulls even the most remote branch-office workers into core business processes and applications, application performance improvement through WAN optimization has become paramount. To meet this requirement, government IT needs to address the following WAN attributes that have slowed application performance for decades:

- Bandwidth limitations
- Latency
- Application contention

The staggering drop in bandwidth as traffic moves from the LAN to the WAN is obvious and well understood. The effects of latency are sometimes less obvious, but latency often slows application performance even when ample bandwidth is available. Finally, application contention becomes far more prevalent on bandwidth-restricted WAN links, and sometimes gets worse as a result of addressing the bandwidth limitation.

To improve application delivery, government IT must look for tools that increase available WAN bandwidth, accelerate applications despite the presence of latency, and resolve contention. In their search for these tools, government IT is best served by looking for options that incorporate support for both general TCP acceleration and application-specific acceleration, as they work to improve the performance of key government applications and processes.

Ultimately, government IT needs a broad-based approach to WAN optimization that will accelerate the widest cross section of their business applications, improve the flexibility of their delivery options, and provide the necessary monitoring and reporting to track application performance.

The Changing Application Acceleration Market

As the application-acceleration market has evolved, several single-function devices have emerged that offer compression, caching, acceleration, bandwidth management and reporting. Government IT has increasingly recognized, however, the impracticality of deploying multiple discrete devices and has instead sought solutions that integrate these capabilities into a single platform. Even more valuable than simple platform reduction, this integration—when designed properly— provides better overall functionality as features tune each other based on dynamic feedback. This multiplicative environment results in greater improvement in application acceleration and WAN optimization. The inclusion of integrated measurement and monitoring also provides a clearer, more complete picture of WAN and application performance.

Juniper Networks delivers innovative technologies designed to accelerate the performance of applications over government WANs. A number of these technologies have been integrated into the company's WAN Acceleration WXC Series Framework, which offers a broad range of capabilities for accelerating applications and optimizing WANs. The WXC Series Framework is incorporated into the Juniper Networks® WXC Series Application Acceleration Platforms, two members of a larger family of solutions that improve application response times within central sites, to branch offices, and for remote users. The technologies delivered in the WXC Series Framework include compression and caching, acceleration, application control and visibility. Ease-of-use and scalability features include routing protocol snooping for topology learning, continuous communications among WXC Series to dynamically update the entire network, and deployment flexibility to install WXC Series platforms without changing the existing network. This integrated approach allows IT to successfully provide LAN-quality application delivery across distributed enterprises.

As a result of this integration, Juniper Networks delivers industry-leading capabilities in:

- Increasing application performance across WANs
- Quality of Service (QoS) support and configuration
- Deployment flexibility
- Ease of use
- Multipath support
- Monitoring
- Scalability
- Security

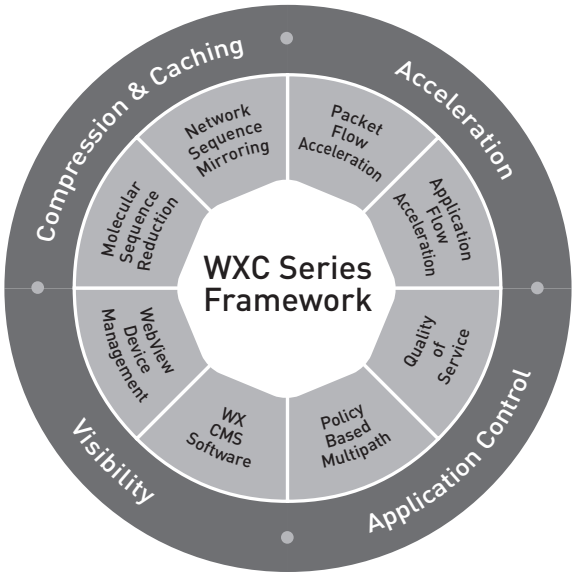

Integrated WXC Series framework

Increasing Application Performance

For application performance across the WAN to improve, the WAN must fundamentally behave more like a LAN. Making that improvement requires a combination of increasing WAN capacity through compression and caching, as well as removing the limitations of distance by speeding up transport and application protocols

Increasing WAN Capacity

The classic option for increasing the size of the WAN link is to upgrade the capacity of that constrained link. The upgrade option, however, either comes at a high price in both dollars and time, or sometimes isn't even available. Juniper Networks provides a far more cost-effective and timely solution, using market-leading compression and caching to gain instant WAN capacity on the existing network.

Juniper Networks next-generation compression techniques begin with the patented Molecular Sequence Reduction (MSR) technology, which uses a highly efficient memory-based pattern dictionary to detect and eliminate data repetitions that occur across a broad range of applications. Depending on the application mix, government organizations typically gain a two- to four-fold increase in capacity on their existing WAN links by eliminating this repetitious data; some departments have seen as much as a 10-fold increase.

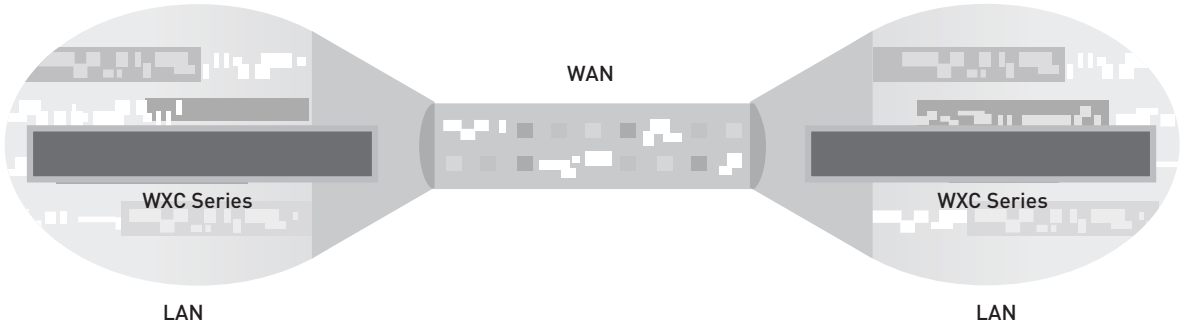

The groundbreaking contribution of MSR technology is its efficiency. Even though its memory of repeated patterns is very large, the compression technique adds very little latency—typically around 2 milliseconds (ms). Another key attribute of MSR technology is its ability to maintain its compression capabilities even on very large amounts of bandwidth, scaling to support OC-3 links. These attributes are traditionally mutually exclusive when using compression techniques such as Lempel-Ziv, its derivatives, or Predictor.

The WXC Series Framework complements the award-winning MSR compression technology with an innovative technique called Network Sequence Caching, which enables the WXC Series to dramatically increase WAN capacity by recognizing much larger data patterns than MSR compression. The Sequence Caching technology relies on embedded hard disks to store longer data patterns for longer periods of time, replacing them with a label for transmission over the WAN.

On the surface, because it operates on large patterns of data, the Sequence Caching technology appears to be similar to file caching. There are, however, significant differences. File caching is designed to eliminate large redundant file transmissions, but it often fails to deliver because of two key limitations. First, file caching works only on a single application, and since enterprises have a heterogeneous mix of applications, the overall impact that file caching has on reducing WAN transmissions is limited. Second, file caching operates only on exactly repeated files.

In contrast, Juniper's Sequence Caching technique works on any IP-based application. Also, it recognizes repeated data patterns and eliminates them, even when a file has been modified. Since most large files transmitted over the WAN are simply modified versions of previously sent files, Sequence Caching is far more effective than file caching. For example, a couple of bullets in a 30-page PowerPoint file may be changed and the file resent over the WAN. Sequence Caching will see the repetition and eliminate 99 percent of the WAN transmission, but file caching will miss it.

In addition, the use of on-board hard disks enables Sequence Caching to provide tens to hundreds of gigabytes of persistent storage, so even sequences seen several days earlier can be eliminated.

Sequence Caching complements MSR compression, as they both recognize and eliminate different types of repeated data patterns. IT will enjoy maximum traffic reduction by combining Sequence Caching and MSR technologies, and they provide very compelling benefits in contrast to traditional compression techniques.

MSR and Sequence Caching techniques reduce traffic for any IP traffic—not just TCP or UDP—so the WXC Series benefit a broader set of applications than many other compression techniques. In addition, with traditional compression approaches, the WAN endpoints store data-replacement labels on a per-tunnel basis, so a hub site with connections to multiple remote locations cannot transfer knowledge of repeated data patterns on one tunnel to other tunnels connecting to other locations. The WXC Series platforms build and maintain a dictionary of repeated patterns across the aggregate of tunnels to remote devices to improve compression results overall.

Ultimately, the efficiency of traditional compression techniques is limited, since they can buffer only a limited number of repeated patterns. In contrast, MSR and Sequence Caching techniques store many more and longer repeated data patterns. The WXC Series remove another limitation that traditional compression tools create, which is the introduction of additional latency as the CPU works to recognize data patterns and replace them with a label. MSR and Sequence Caching dramatically reduce traffic flows while adding only a negligible amount of latency.

Bottom Line: The compression capabilities in the WXC Series Application Acceleration Platforms work across the broad spectrum of repeated data pattern sizes and dramatically increase the capacity of government WAN links.

Speeding Transmissions

To speed transmissions across the broadest range of government business applications, WAN optimization and application acceleration platforms need to overcome the impact of latency. Latency affects application performance because the underlying protocols rely on a back-and-forth sequence of data transmissions and acknowledgements. Each "send and acknowledge" operation is subject to the full delay of a WAN round trip time (RTT), so an application that uses an inefficient protocol may require hundreds or even thousands of RTTs to complete transmissions or downloads.

To overcome these delays, an application acceleration platform requires both TCP acceleration to benefit a broad range of applications, and application-specific acceleration to speed applications whose Layer 7 protocols are less efficient than TCP.

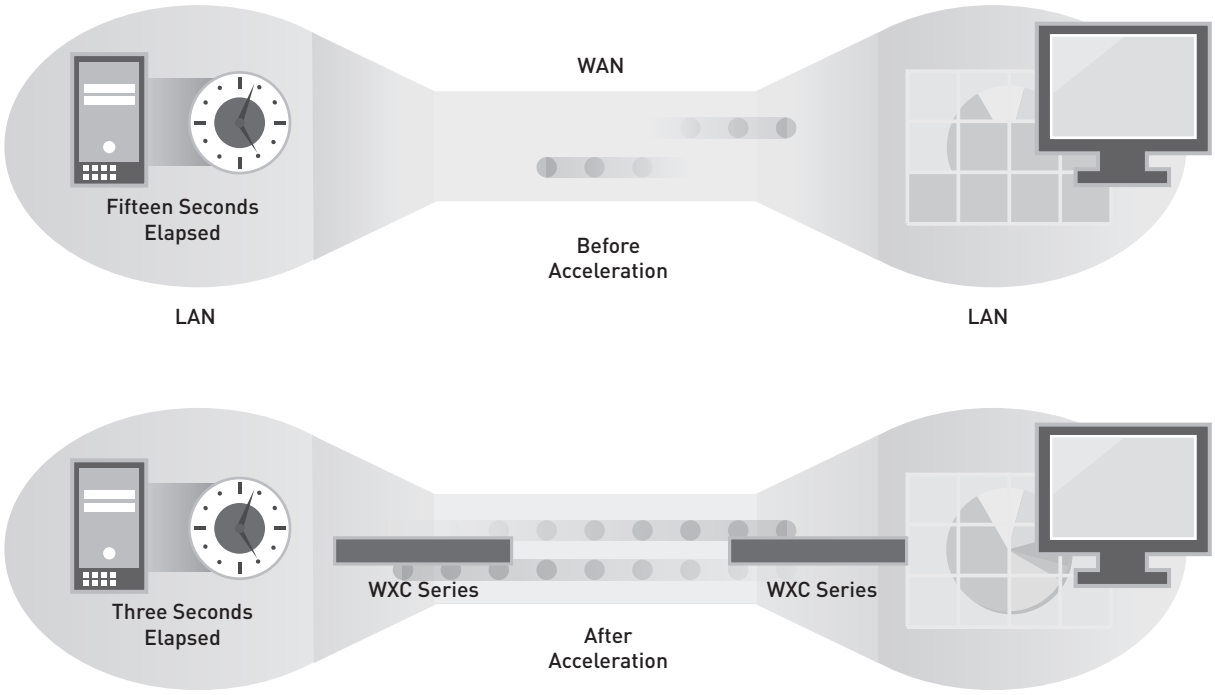

Several TCP acceleration techniques can be implemented to benefit applications based on either short-lived or long-lived TCP connections. The WXC Series Framework's Packet Flow Acceleration (PFA) technology includes a series of techniques that accelerate TCP to boost application performance:

- The Fast Connection Setup technique improves the performance of short-lived connections by eliminating one RTT from the TCP connection setup, speeding up applications that use short connections and have chatty protocols.
- The Active Flow Pipelining technique accelerates TCP performance by terminating the TCP connection local to the sender and using a more efficient transport protocol between WXC Series platforms. This feature significantly benefits application performance on high bandwidth, high-latency connections.
- The Forward Error Correction technique limits the need for retransmissions on lossy networks. It makes use of recovery packets sent alongside data packets that index those data packets, allowing for reconstruction of lost packets.

The protocols of some applications limit throughput even more than TCP and so are subject to delay on WAN links with even modest latency. Microsoft Exchange, Microsoft File Services, and Web-based applications are three broadly used business applications that require application-specific acceleration. These applications can't benefit from TCP acceleration until they are accelerated at the application layer, at which point application-specific acceleration and TCP acceleration provide a compound performance improvement.

The WXC Series Framework's Application Flow Acceleration (AppFlow) technology accelerates the underlying protocols of these applications: the Messaging Application Programming Interface (MAPI) used by Microsoft Exchange; the Common Internet File System (CIFS) at the heart of Microsoft file services; and HTTP for Web applications.

In the case of Exchange and file services, their underlying protocols send data in small blocks and require an acknowledgement for each, resulting in hundreds or even thousands of RTTs to complete a single transaction. As a result, performance drops dramatically when used across a WAN link with even modest latency—20 ms or 30 ms—resulting in user frustration and lower productivity. For Web applications, HTTP requests objects one at a time, so dozens of RTTs are needed to load a single Web page, again impacting user productivity.

The AppFlow technology accelerates these applications by pipelining the data blocks and Web objects, sending as many in quick succession as needed to fill the available WAN capacity and deliver up to a 50-fold improvement in application performance.

Bottom Line: *WXC Series Application Acceleration Platforms enhance performance across a broad range of government application types and WAN link characteristics, improve user productivity associated with existing applications, and enable deployment of new applications that previously could not run across a WAN.*

Security

Government organizations no longer have the luxury of securing just their private WANs and avoiding the Internet for business transactions. Instead, they need to make use of both transports. Thus, government IT needs tools to make all transports secure enough for these business transmissions. For WAN optimization platforms, two aspects of security are critical: securing the device itself and securing the data that traverses the device.

Juniper Networks hasn't left anything to chance within the WXC Series. All methods of access for the WXC Series are secure, using HTTPS and SSH. IT can also define Access Control Lists (ACLs) to allow or disallow access to the platforms, and IT can deploy Authentication, Authorization and Accounting (AAA)-based access to the platforms via RADIUS.

Government IT also has the option to disable all network access to the WXC Series and support only console access. In addition, Juniper Networks actively monitors the security warnings from industry security watchdog groups to make sure that all vulnerabilities are removed.

Juniper Networks WXC Series IPsec Implementation
• Advanced Encryption Standard (AES) and triple Digit Encryption Standard (3DES) for encryption • HMAC-SHA-1 and HMAC-MD5 for packet authentication • Dynamic key exchange (IKE) • "Retail" export approval from the Department of Commerce (–15 Mbps)

To secure the platforms from a physical standpoint, the WXC Series use no exposed flash memory cards that can be stolen and/or compromised, and IT can choose to deactivate the front panel configuration feature.

The WXC Series also ensure the security of their data transmissions. IT can optionally deploy a standards-based IPsec encryption feature in sites without a virtual private network (VPN) deployment, securing data sent over unsecure links such as the Internet or satellite, and also securing device-to-device communications.

Bottom Line: *The WXC Series provide a wide range of security features.*

QoS Support and Configuration

Since speeds between the LAN and the WAN differ by orders of magnitude, no amount of compression or acceleration will solve all problems. This dramatic discontinuity in bandwidth means that contention for WAN real estate is a very real problem that needs to be addressed with an effective and realistic QoS and bandwidth allocation model that enforces business priorities.

Historically, configuring QoS has been incredibly complex, limiting its effective use and consuming too much government IT staff time.

Many WAN optimization platforms rely on per-flow manipulation of traffic to apply QoS policies. This granular micromanagement of applications, however, simply takes too much of IT's precious time and, worse, quickly becomes out of step with true requirements as new sites or new applications appear. This kind of granular approach overly complicates the goal—to ensure that important business traffic makes it through the WAN by de-prioritizing, when needed, the less critical traffic.

Juniper Networks changes all this. The WXC Series Framework approaches QoS and bandwidth allocation with a pragmatic, solutions-based orientation, making sure that mission-critical application traffic arrives where it needs to within the required service-level agreement (SLA). The WXC Operating System (WXOS) software uses wizard-based QoS templates based on business priorities to simplify QoS configuration and wizard-based configurations. This approach strikes the needed balance between performing QoS in the most appropriate device—the WAN optimization platform—and keeping the implementation simple.

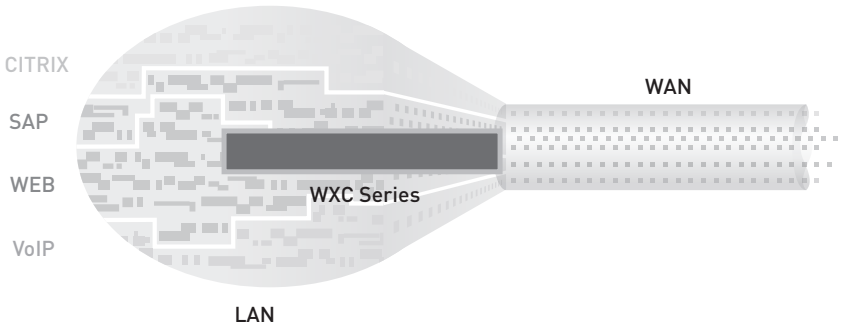

Why is the WAN optimization and application acceleration platform the best point in the network to perform QoS and bandwidth allocation? Because it's the last point in the network that sees uncompressed traffic. If an edge WAN router performs QoS rather than the WAN optimization device, for instance, and it enforces QoS on already compressed traffic, the QoS policy will be incorrectly applied since compressed traffic will have a disproportionate share of the bandwidth.

For example, if the QoS policy specifies that voice traffic should receive 10 percent of the available bandwidth and text traffic should receive five percent, but those parameters are applied after compression, text will constitute far more of the total bandwidth since it's highly compressible, while voice traffic is not. The QoS policy will be accurate over the WAN but will not provide the appropriate user experience on the far end of the WAN link. Consequently, a WAN optimization platform that lacks QoS will ultimately break the QoS policies as defined in a separate platform.

Juniper Networks has also integrated the QoS function with the visibility features of the WXC Framework. Getting constant feedback on WAN dynamics is essential to implementing QoS successfully. For example, QoS must understand how much compression is happening on the traffic; knowing the capacity of the WAN is a prerequisite for knowing when to invoke prioritization techniques.

The alternative, required by many WAN optimization platforms lacking this insight, is for government IT to manually adjust the QoS policies after investigating and analyzing the compression results. Other platforms simply enforce QoS without regard to capacity at all, throttling back all traffic all the time to avoid contention entirely. These approaches, however, make the fundamental goal— better utilization of the WAN link—unattainable.

In addition, for QoS to operate effectively throughout the enterprise, the WAN optimization platform requires a holistic view of the WAN. This broad perspective includes seeing both those sites outfitted with an optimization device and those without them. Many WAN optimization platforms have no understanding of the destination location. That kind of "one sided" QoS deployment is important for including sites without WAN optimization platforms but in the enterprise QoS strategy, that mode should not be the only QoS technique supported.

Understanding both ends of the WAN link provides extensive advantages. This kind of "dual sided" deployment allows the WXC Series, for example, to automatically map traffic according to the QoS template assigned to that destination device. The dynamic link knowledge also enables IT to add additional WXC Series devices to the network very simply, with just a few mouse clicks, and have them adopt the appropriate QoS policies.

While the dual-sided approach provides for the greatest link understanding and dynamic behavior, the WXC Series also support "one-sided" deployments of QoS as well, to enable a consistent QoS policy across the distributed enterprise.

Insight into the traffic type is essential for applying QoS appropriately. Many WAN optimization platforms lack the ability to identify business applications at all levels of the network stack, from simple addresses to deep inspection within the payload. All Citrix applications, for example, look the same at Layer 3 but by looking inside the payload, a WAN optimization platform can distinguish critical ERP traffic from simple print jobs.

Government organizations need the flexibility to augment the QoS embedded within a WAN optimization platform with other techniques for marking traffic, and the optimization device should not do anything that disables the markings set by other network devices. Rather than overwrite any needed information, Juniper Networks has designed the QoS feature of the WXC Series Framework such that type of service (ToS) / differentiated services (DiffServ) settings can be mapped to other network devices, tunneled traffic is still identifiable by application, and

MPLS class of service (CoS) information can be communicated to the edge devices in service provider networks without permanently overwriting anything in the original packet.

Bottom Line: *WXC Series Application Acceleration Platforms provide the adaptive and dynamic QoS functionality enterprises need, without the complexity typically associated with defining granular bandwidth control.*

Deployment Flexibility

Delivering applications throughout a distributed government organization requires great flexibility in the options for network placement and configuration, transport support and capacity ranges. It's critical that the WAN optimization and application acceleration platform conform to the existing network design rather than forcing the network design to conform to the platform's needs.

For network placement, government IT must have the choice to deploy WAN optimization equipment either on the network between a LAN switch and WAN router, or attached to a switch and router in a one-armed fashion. The WXC Series support both these modes, with both Inline Mode and Off-Path Mode options.

A basic inline configuration is a common feature among WAN optimizers, but the WXC Series are unique in their awareness of 802.1Q and their ability to compress traffic within the full 4095 number of supported VLANs. The platforms can optionally preserve the VLAN tags as packets are transported to other destinations through the tunnel.

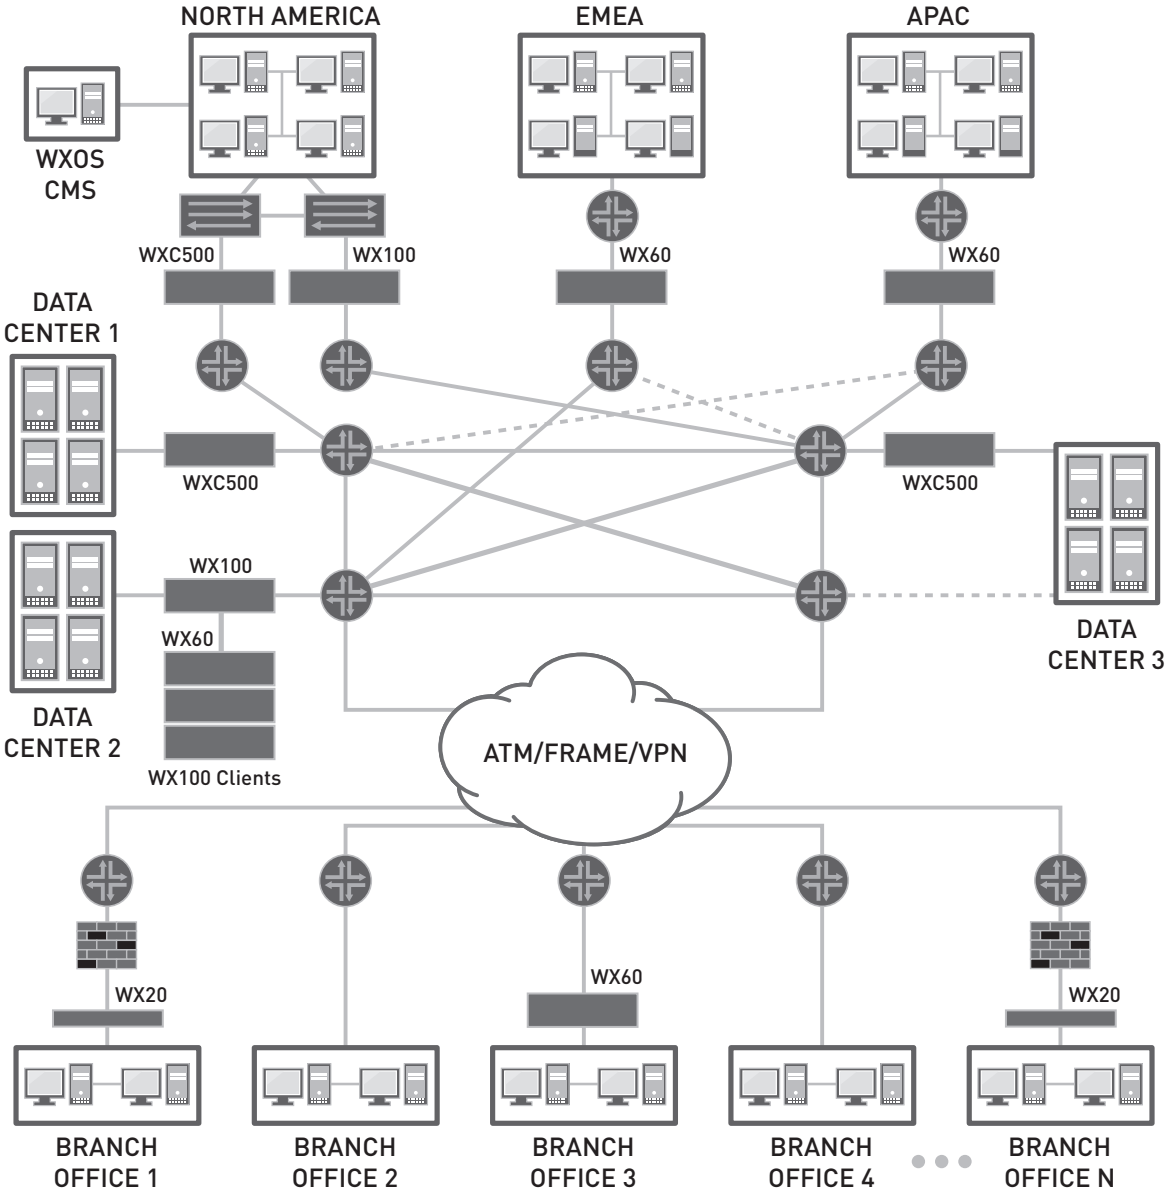

Other WAN optimization and application acceleration devices, in contrast, cannot operate in an off-path mode, and off-path deployments are essential for interoperability with some WAN architecture. For example, if the WAN router acts as a collapsed backbone, serving both local LANs and remote networks, government IT needs to attach the WAN optimization platform directly to a port on the router. The WXC Series are also unique in that they allow IT staff using off-path mode to selectively choose what traffic is redirected to a WXC Series device, and what traffic is left untouched.

The WXC Series include a number of redundancy features unique in the marketplace. The devices support an active fail-to-wire mechanism so that in case of a device failure, all traffic will pass through the box at wire speed untouched. For additional levels of redundancy, WXC Series platforms support dual-active redundancy with no need for extra configuration of surrounding network devices, as well as an n+1 backup mechanism. The platforms also work with routers configured with redundancy protocols, and the devices can load-balance tunnel traffic to redundant WAN routers or load-balance to redundant destination WXC Series platforms.

Another critical feature to enabling flexible deployment options is supporting both point-to-point and point-to-multipoint configurations. Many government networks are built in a hub-and-spoke design, with several branch offices feeding into regional or centralized sites. That kind of traffic aggregation must be supported by the WAN optimization platforms so that government IT maintains the efficiencies of that hierarchy. To further support network hierarchy, the application delivery platforms should support the appropriate tunnel configurations to enable traffic aggregation from regional offices to centralized locations. The WXC Series, for example, feature Tunnel Switching, which allows IT to create multiple layers of tunnel aggregation, as is often found in frame relay networks using the hub/regional/spoke topology.

Deployment flexibility also must incorporate a variety of transport types. Government organizations use a variety of WAN services, including private WANs, frame relay, public VPN services and MPLS.

The traffic marking used by the WAN optimization platforms should not interfere with any markings needed to support the full gamut of transport types. To support MPLS, WXC Series platforms implement transparent QoS communications to place traffic in the right label switched path (LSP) without ever changing the packet. To support satellite links, the WXC Series Framework supports a series of techniques that reduce the impact that latency has on application performance.

Bottom Line: *WXC Series Application Acceleration Platforms provide the deployment flexibility needed for government IT to meet the varying configuration and transport needs of different locations of the distributed government organization.*

Ease of Use

Given the increasingly pervasive deployment of application acceleration platforms throughout a government's locations, having a system that's easy to use is vital. A number of factors determine a device's usability—an intuitive interface, automated configuration and deployment capabilities, and synchronized communications among platforms are a few of the necessary elements.

The WXC Series are easy to use and manage, offering both Graphical User Interface (GUI)- and Command-Line Interface (CLI)-based views that allow government IT to manage the devices individually or as a collection. A common complaint about many WAN optimization and application acceleration platforms is the complexity associated with defining policies for them. Many optimization devices fail to meet the "human factor" needs in this regard and are well known for bogging down IT with a slow GUI. They require many separate screens to complete a configuration and take a long time for each interaction. As part of their equipment evaluation, IT staff should be sure to step through the process for configuring a QoS policy.

Automating key functions is another critical element of ease of use, and the pursuit of increased automation makes the WXC Series easy to deploy. The most significant achievement in this area is the auto-deployment feature for branch offices. Through a combination of the WXOS software and the WX Central Management System, IT can pre-stage configurations centrally via templates and then have remote WXC Series platforms download them automatically.

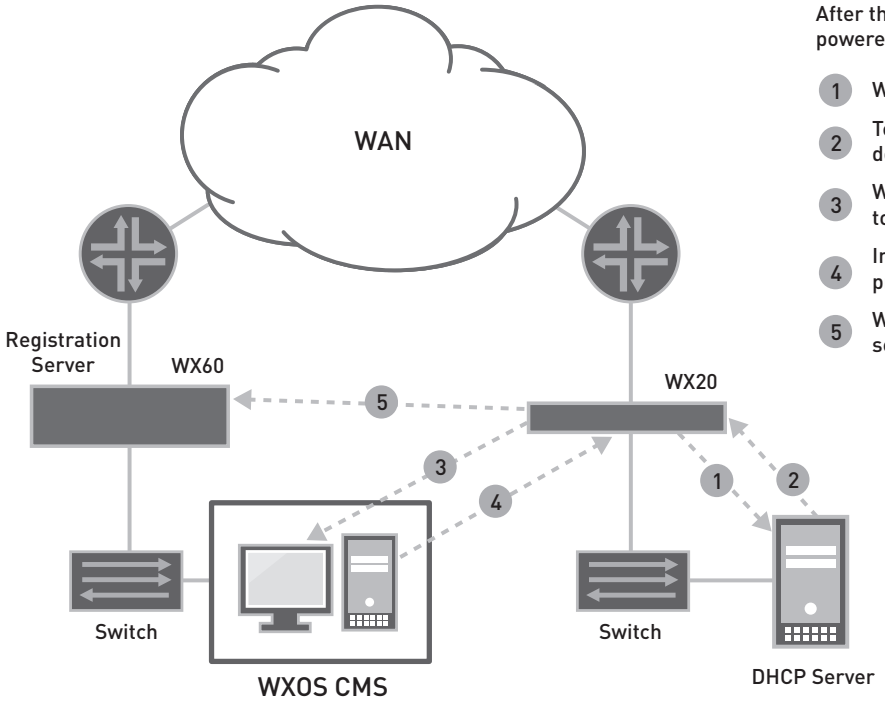

After the WXC Series device is powered on, configuration is automatic:

1. WXC Series issues DHCP request
2. Temporary IP address and domain returned
3. WXC Series performs DNS lookup to locate CMS
4. Image file (if needed) and pre-staged config file downloaded
5. WXC Series connects to registration server and begins operation

Branch-office staff needs only to plug in the device and connect it to the network. Straight out of the box, the device will automatically procure a network address, locate the centralized management software via the domain name service (DNS), request a configuration, download it, and begin operation. The lack of government IT involvement in deploying devices to branch offices produces a significant savings in staff time and money, and it enables rapid, pervasive deployment.

The WXC Series automate several other tasks that are typically cumbersome to perform on WAN optimization and application acceleration devices. After a WXC Series device is up and running, it connects to a registration server to learn about other remote WXC Series devices and how they're configured—for example, which devices are hubs and spokes; whether key compression techniques are enabled, whether IPsec is enabled, and whether a device is sending out traffic over two WAN links and using path optimization.

This registration information is essential to the synchronized communications among the WXC Series devices. These communications provide government IT with distributed stateful intelligence about the network state and enable devices to act on changes such as link loss or increased congestion on a path.

These automated synchronizations and communications dramatically simplify both setup and ongoing operations of the WXC Series. In contrast to the manual intervention needed to support other WAN optimization and application acceleration devices, Juniper Networks simplifies tasks such as applying QoS policies to remote devices, upgrading those policies, upgrading the system software, and accommodating topology changes.

Bottom Line: *A sophisticated but intuitive GUI, significant automation capabilities, and synchronized communications among WXC Series Application Acceleration Platforms result in a very simple-to-use product set.*

Using Multiple WAN Links

Government organizations increasingly seek to take advantage of hybrid public/private WAN transports, but maintain the assurance that key performance criteria will still be met. To make effective use of both paths, government IT needs to apply business policies to each link and monitor its performance. Most WAN optimization platforms overlook this WAN deployment scenario and fail to help IT make full use of these dual links.

The WXC Series Framework includes a unique feature called Policy-Based Multipath (Multipath) that enables government IT to define which applications traverse which link and under what conditions. For example, IT can designate that latency-sensitive traffic such as VoIP will run over the private link, while delay-tolerant applications such as email and bulk file transfers will use the Internet/ VPN link. But in addition to enabling this simple allocation, the WXC Series software also allows IT to set latency and loss thresholds for each link. IT can use the Multipath function to define, by class of application, how to treat traffic when a performance threshold is exceeded.

Government IT can designate, for example, the application classes that will switch to the other link when performance suffers.

The Multipath feature demonstrates the integrated nature of the WXC Series Framework. For example, when traffic is diverted from one link to another, QoS policies ensure that applications already flowing over the second link are not negatively impacted.

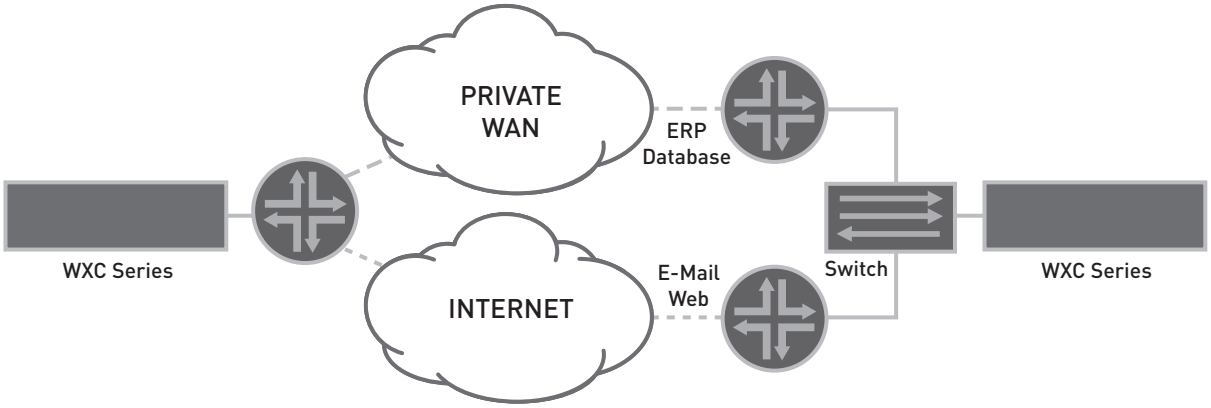

Bottom Line: *WXC Series Application Acceleration Platforms offer path selection flexibility that allows government IT to maximize the value of dual WAN links while adhering to essential QoS and security policies.*

Complete Monitoring

Defining policies for optimizing traffic flows over the WAN requires that government IT understands the actual traffic flows. Monitoring tools that provide unified insight into distributed applications and networks are essential to effective application delivery.

The WXOS software provides the broadest and deepest set of analysis tools for understanding WAN traffic characteristics and performance levels for applications flowing through the WXC Series. Government IT can choose to view information per WXC Series device or in aggregated form, and they can gain insight into such aspects as packet size distribution, error rates, throughput statistics, and TCP and application acceleration data.

What truly sets the WXC Series monitoring capabilities apart, however, is their combination of data reduction and capacity improvement statistics with an understanding of the QoS statistics. It's crucial for IT to see what's happening to traffic when QoS is being invoked. If the system is applying a QoS policy, then that means congestion is present. Too many government applications are contending for too little bandwidth, so some traffic has to be restricted to enable priority applications to transmit. IT needs to understand that impact to better tune QoS policies.

In addition to these statistics, the WXOS software provides other unique views, including a sophisticated pass-through monitor to communicate what traffic is not compressed and why; a link SLA monitor showing path latency and path packet loss; packet size histograms that display incoming and tunneled traffic; the results of TCP and application acceleration; and a tunnel summary showing the status of all the tunnels on the device.

The WXOS software goes beyond simply displaying information, allowing for many methods of exporting information. Supported methods include SNMP, NetFlow, Comma Separated Values (CSV) and Excel graphs. The software can also facilitate in-depth troubleshooting at remote locations by taking a remote sniffer trace.

Complementing the WXOS software is the WX Central Management System, which allows government IT to configure and manage multiple WXC Series platforms from one central location. The WXC CMS provides government IT with a unified view into system capabilities throughout the distributed enterprise. Government IT can look at metrics about compression performance, application acceleration, WAN utilization, and QoS and bandwidth allocation.

The WXC CMS provides government IT with information about what traffic is traversing the WAN, which applications are consuming most of the valuable WAN capacity, and which traffic is being impacted by the application of QoS. Government IT can also use the WXC CMS to schedule system upgrades, apply new configurations, update QoS policies, and automate license management.

Bottom Line: *WXC Series Application Acceleration Platforms provide comprehensive insight into WAN traffic patterns but in a manner that's easy for government IT to understand, quickly enabling policy adjustments that maximize the application acceleration capabilities of the WXC Series devices.*

Scalability

Because government organizations have a variety of locations that range in size, applications and link types, departments need a wide range of WAN optimization and application acceleration platforms.

Essentially, government IT needs an architecture that can scale from small branch offices to regional deployments to centralized hub locations. Other aspects of scalability include the scope and breadth of network traffic that a WAN optimization and application acceleration platform can benefit, the number of remote locations a single-hub WAN optimization and application acceleration platform can handle, the ultimate throughput of any single optimizer, and the ability to support a worldwide implementation.

Government IT must be able to optimize traffic delivery across a broad range of applications. Designing a platform that's highly specific to just one application may achieve superlative benefits for that application, but the goal of improving overall traffic delivery will suffer. Rather than focusing on application-specific approaches, IT should look for platforms that scale the performance of a wide cross-section of government applications.

Government IT also needs scalability in the throughput levels the application acceleration products enable. Many organizations have WAN capacities that range from 64 Kbps links in small offices all the way to OC-3 connections in headquarter locations. Government IT should not only be able to buy a range of equipment types to deliver various bandwidth amounts, but should also be able to scale a single platform to higher capacity levels via clustering. IT also needs these optimization platforms to support large numbers of connections to other sites in these centralized hub locations.

The WXC Series meet these demands, easily scaling to support compressed output speeds up to 155 Mbps and 2,000 connections to other application acceleration platforms. In addition to a high tunnel count, application acceleration platforms should also enable tunnel configurations that support large, hierarchical networks. The architecture should provide automatic tunnel aggregation to support these complex topologies as networks increase in size. The Tunnel Switching feature on the WXC Series provides exactly that capability, enabling any-to-any communications between any two sites without requiring a full mesh setup of tunnels linking the sites

Monitoring capabilities must also be scalable. Government IT must be able to look at multiple parameters at the same time, such as application, destination and flow direction. Without seeing this information tied together, IT will not be able to scale the WAN optimization to meet the broad needs of the enterprise.

Bottom Line: *WXC Series Application Acceleration Platforms provide a highly scalable product set and software suite that meet the wide range of application, bandwidth and monitoring capabilities demanded in a distributed government environment.*

Conclusion

Government network users are demanding an acceptable level of application performance across the LAN and WAN. They require near real-time access to core government business processes and applications in order to perform their functions effectively. This is a challenge for government IT departments due to WAN bandwidth constraints.

Juniper Networks WAN application acceleration solutions address this challenge by enabling government organizations to make the most of their existing WAN resources. Not only does the WXC Series Framework multiply the effective capacity of existing WAN connections, but it also accelerates applications and ensures that mission-critical applications receive the bandwidth they need.

The Juniper solution delivers the resources needed to provide LAN-quality application delivery without requiring any changes or additions to the existing WAN. With a single platform, government agencies gain WAN capacity, improved application support, and optimized throughput. The end results are satisfied users with a cost-effective solution for government IT.

About Juniper Networks

Juniper Networks, Inc. is the leader in high-performance networking. Juniper offers a high-performance network infrastructure that creates a responsive and trusted environment for accelerating the deployment of services and applications over a single network. This fuels high-performance businesses. Additional information can be found at **www.juniper.net.**

Corporate and Sales Headquarters
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

APAC Headquarters
Juniper Networks (Hong Kong)
26/F, Cityplaza One
1111 King's Road
Taikoo Shing, Hong Kong
Phone: 852.2332.3636
Fax: 852.2574.7803

EMEA Headquarters
Juniper Networks Ireland
Airside Business Park
Swords, County Dublin, Ireland
Phone: 35.31.8903.600
EMEA Sales: 00800.4586.4737
Fax: 35.31.8903.601

To purchase Juniper Networks solutions, please contact your Juniper Networks representative at 1-866-298-6428 or authorized reseller.

2000248-002-EN Oct 2010