

White Paper

# Essential IT Guide: Ensuring Highly Available Access to Business-critical Applications

---

Supporting Branch-office and Remote Users in a  
Centralized World



Juniper Networks, Inc.  
1194 North Mathilda Avenue  
Sunnyvale, CA 94089 USA  
408 745 2000 or 888 JUNIPER  
[www.juniper.net](http://www.juniper.net)

Part Number: 200197-001 July 2006

Centralized Servers and Distributed Users.....	3
Risks and Requirements for Application Availability.....	4
The WAN Optimization Controller is Critical.....	4
Device Reliability .....	5
Device Redundancy .....	5
Path Redundancy .....	6
Ease of Installation .....	6
Creating a Total Solution.....	6
Juniper Networks: Ensuring Access to Applications.....	7
Reliable Devices .....	7
Flexible Deployment.....	7
Intelligent Path Selection.....	9
Achieving Operational Efficiency.....	9
Highly Available Secure Connections.....	10
Highly Available Routing Architecture .....	11
Maintaining Availability at Large Scales.....	11
Juniper Networks: Creating Complete Solutions .....	11
Additional Reading .....	12

## Centralized Servers and Distributed Users

It's often said that the problems of the past come back to haunt us. This is the situation many organizations are facing with the initiative to move application servers out of branch offices and centralize them in the corporate data center. Organizations find that they are running into familiar problems from the past, as well as some new ones.

The centralized server model is similar to the old mainframe computing model from years ago, where applications were served to distributed users on terminals in the branch office. In the mainframe model, link speed was slow and if the link went down, users were cut off.

Everything changed with the introduction of the PC, which in turn spawned local area networks (LANs) and the client-server model, in which local servers deployed in the branch office delivered applications to users over a fast LAN connection. Since the servers were local, there was no long-distance link connecting the users to the mainframe, virtually eliminating the threat of a downed connection cutting remote users off from centralized applications and data.

However, while the LAN model provided tremendous improvements in performance and reliability, it did have its downside: every branch office essentially became its own data center, resulting in server sprawl and an accompanying growth in ownership costs.


To reduce those costs, gain better control and improve operational efficiencies, organizations have now embarked on an industry-wide initiative to move servers back to a centralized data center, where a single IT group can make upgrades, roll-out new applications, and perform routine maintenance and management tasks. Not only does this centralization initiative save on hardware, staffing, software and facilities, it also helps companies maintain the agility they need to deliver better service levels as they move towards a real-time infrastructure. Centralization also makes regulatory compliance easier to achieve since back-up, administration and recovery of data and applications is centralized as well.

Yet for all its benefits, centralization has its drawbacks. Deploying application servers in a single central location creates a single point of access for all users, local and remote, as well as a single point of failure. Because today's corporate enterprise is widely distributed, with dozens of branch offices located around the world, IT is faced with a challenge: how to maintain high availability of business-critical applications for branch-office users in a distributed enterprise.

The greatest threat to availability of any centralized application is the state of the WAN that connects distributed branch offices and remote and mobile users to the centralized business applications. Under the decentralized model, where application servers were deployed in each branch office, work could continue if the WAN link to the data center went down; data would be replicated to the main servers in the data center once the link was restored. In today's centralized environment, a WAN link failure means branch offices and other remote users are completely cut off from the servers in the data center and the business-critical applications that they host.

In searching for a solution to the availability problem, IT managers have evaluated storage devices to provide proxy file services in the branch office in the event of a WAN link failure. Unfortunately, this approach does not support transactional

applications, or real-time applications like VoIP and video conferencing, which are very critical to a business. To overcome the access challenge, IT leaders must develop a strategy for building a highly available wide-area network infrastructure designed to provide continuous access to all centralized business applications in the data center. This strategy must utilize reliable networking devices and a resilient network design – including WAN optimization controllers, routers, security devices and WAN services – to ensure branch office, remote and mobile users always have uninterrupted access to their business-critical applications. Choosing equipment with the right capabilities – for both the branch office and the data center – is critical.



**Figure 1:** As applications are centralized, the WAN link becomes a lifeline. High availability of the WAN link is essential to maintaining access to business-critical applications, most of which are transactional or real time.

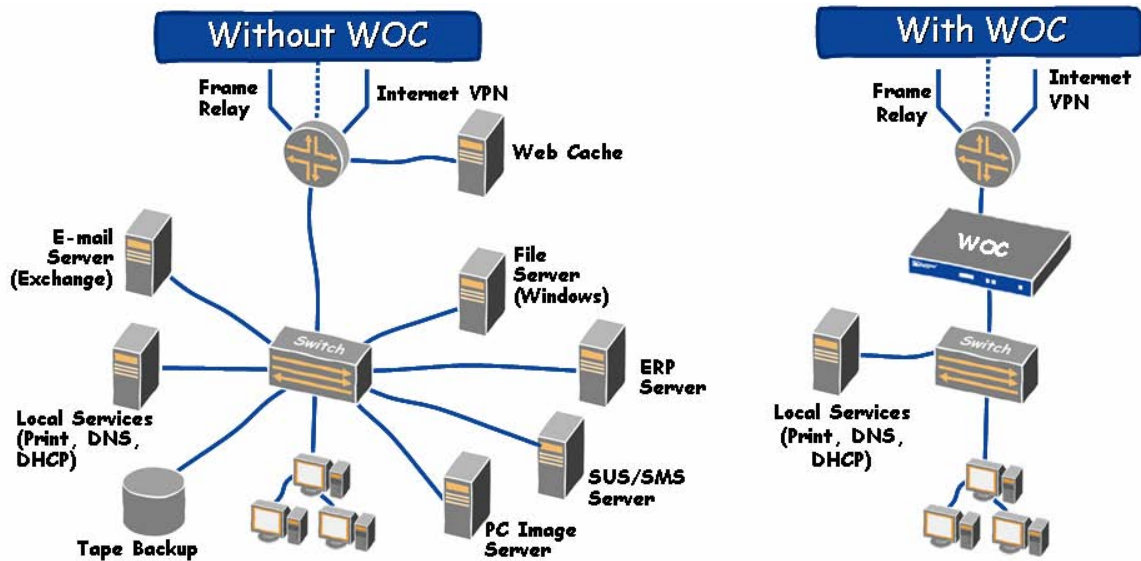
## Risks and Requirements for Application Availability

Ensuring LAN-like response times for centralized applications over the WAN introduces a host of challenges due to limited capacity, the effects of latency on round-trip times, and contention for fixed amounts of bandwidth. This situation is amplified by the move to Web-enabled applications that typically require more than 10 times the bandwidth and create thousands of additional TCP sessions, each requiring a request and acknowledgement for each block of data that is transmitted.

### The WAN Optimization Controller is Critical

The WAN optimization controller (WOC) delivers essential functionality for application performance in wide-area environments. As a result, WAN optimization is a key component of a highly available network. If the WOC fails, then bandwidth capacity it added will be lost. Contention between applications will increase, causing

business-critical applications to be slowed by non-critical traffic. TCP sessions will be dropped and throughput will plummet as connections are slowly reestablished. TCP and application protocols will no longer be accelerated, and the impact can cause the performance of business-critical applications to become intolerable. As a result, ensuring performance and availability to applications requires utilizing a WOC that fully incorporates high availability.



**Figure 2:** Building out infrastructure in the branch has created an unmanageable server sprawl. As organizations consolidate servers in the data center they are deploying WAN optimization to ensure high performance for applications running over the WAN.

### Device Reliability

To deliver high availability, the WOC must be able to survive internal failures such as the loss of power, cooling or storage. If a failure does occur, the WOC must not impede traffic flow; it must include a pass-through feature that allows traffic to flow even if the WAN optimization features are disabled.

### Device Redundancy

To ensure continuing traffic optimization in the event of a device failure or the failure of a network link, the WAN optimization solution must support redundant deployment options including active/active, active/passive, or N+1 redundancy configurations. It must support flexible deployment schemes, including a single device installed inline on the link between a LAN switch and WAN router; multiple devices in a mesh to support highly redundant configurations in data centers; and a one-armed deployment to support smaller branches with a collapsed backbone switch/router or where multiple devices are not required.

### **Path Redundancy**

A high-availability design should incorporate redundant links for the data center and from each of the branch-office locations. Dual links, also known as multi-homing, provide the ability to choose the best performing path during times of congestion and, in the event of a primary link failure, to quickly divert traffic to an available secondary link. Multi-homing can be provisioned using two ports on one router or, for full redundancy, using two routers. To ensure against network failure, each link should be provisioned to a different provider. Routers should be configured with a routing failover mechanism and a router redundancy protocol so that they respond to both the failure of the link and the failure of the router. The WOC and routing equipment must work together, providing the ability to direct specific application flows to a specific link.

Not only do redundant links ensure that high-priority applications always have sufficient WAN capacity, they also allow applications to recover rapidly from device or link failures by automatically diverting traffic to an alternate path across available equipment.

### **Ease of Installation**

To enable large scale deployments, the ability to manage communities of remote devices from a central location is essential. To cut costs, organizations have reduced, or even eliminated, branch-office IT staff as they shrink the branch infrastructure; therefore, any equipment required for a small branch must be easy to install by local, non-technical personnel. Also, a significant percentage of network failures are due to equipment configuration errors; in order to minimize these errors, equipment installation procedures should be optimized for operational simplicity. Easy installation should include features such as templates for applying standard configurations and the ability to automatically download configuration information and software updates from a “master” device.

### **Creating a Total Solution**

High availability is a network-wide issue. All of the components in the network play a critical role in delivering applications with high availability from end-to-end. A complete solution for delivering applications over the WAN requires interoperability between the WOC, routers, and security devices. When designing networks for high availability, organizations should choose hardware that is designed to deliver critical functionality across devices.

Building a total solution for high availability hinges on choosing a trusted supplier — one whose products deliver the essential capabilities required to ensure high availability in today’s distributed enterprise where services and applications are centralized. Organizations need a vendor whose products have a track record of performance, reliability and scalability, as well as one that offers a complete range of products that work together to deliver a total solution for application delivery that ensures data security, device reliability and network resiliency.

## Juniper Networks: Ensuring Access to Applications

A critical component of a high-availability solution for delivering mission-critical business applications in the distributed enterprise is the WOC. The Juniper WX and WXC application acceleration platforms meet this mandate, delivering the required features that make them a perfect fit for ensuring high availability of business-critical applications.

### Reliable Devices


Device reliability is the first line of defense in creating a highly available network. The WX/WXC platforms can be deployed with redundant power supply units (PSUs), and WXC platforms contain redundant hard drives so that, in the event of a drive failure, the device can continue to operate.

Since the WX/WXC platforms boot from persistent memory, not from a hard drive, the WXC platforms can operate even if both drives fail, providing memory-based compression similar to that of the WX platform. All WX/WXC platforms support fail-safe operation; in the event of a failure of any kind, including total loss of power, the WX/WXC interfaces automatically convert to a bypass mode in which all traffic simply passes through the device untouched, ensuring an uninterrupted flow of data.

### Flexible Deployment

Designing redundancy into a network is key to ensuring high availability, and the WX/WXC devices support multiple deployment options to provide redundancy for the data center, regional locations and small branch offices. For instance, in standard or switch/router mesh environments, the WX/WXC platforms can be deployed inline between switches and routers, either individually or in tandem, in an active-standby, active-active or, for maximum redundancy, an N:1 scheme.

Conversely, an off-path mode feature enables a WX/WXC platform to attach to a port on a Layer 3 switch or router, providing the flexibility to support collapsed backbone switch/router environments serving both local LANs and the WAN. The off-path mode also allows WX/WXC platforms to be deployed in a mesh in situations where inline deployments would be impractical. The off-path mode feature makes the WX/WXC platforms unique in that they allow the IT staff to choose which traffic is redirected to the WX/WXC device.




**Figure 3: The WX/WXC platforms support redundant inline deployments in both mesh and standard environments.**

The WX/WXC platforms operate in communities that dynamically exchange information such as topology, reachability and path-performance metrics, providing IT with distributed stateful intelligence about local- and wide-area network conditions to assist with management and troubleshooting. For large-scale deployments, IT has the option to partition WX/WXC platforms into separate domains, increasing operational scalability and reducing the potential for errors.

### Intelligent Path Selection

While dual links protect against loss of connectivity, link congestion is a much more common problem. The WX/WXC platforms' Policy-based Multipath feature allows IT managers to set policies for congestion thresholds that direct traffic over specific links when multiple links are available. These links can be provisioned on one router using one WX/WXC platform or on a pair of redundant routers using one or two WX/WXC platforms. The multipath feature works independently of the WAN transport and can work with any connection type so that IT departments with legacy Frame Relay connections can complement them with a DSL connection to the Internet and then direct traffic to either link based on predetermined policies. The WX/WXC platforms can application load-balance traffic to redundant WAN routers and to redundant destination WX and WXC devices.




**Figure 4:** The multipath feature enables IT to direct traffic over specific links to ensure mission-critical and delay-sensitive application flows take advantage of more reliable links.

With the multipath feature, IT can provision transactional applications such as SAP and Oracle, or delay-sensitive applications such as VoIP, over the most reliable link and setting a policy that automatically switches the traffic to the backup link if the primary link fails or performance falls below a set threshold. By leveraging both legacy links and Internet connections, the multipath feature can significantly increase the enterprise WAN's overall availability and ease the migration from legacy links to Internet connections, as organizations can easily swap out one link while the other carries traffic.

### Achieving Operational Efficiency

Automating key functions is a critical element of ease-of-use and error-free deployment. Using the branch office auto-deployment feature, enabled through a combination of the WX operating system (WXOS™) and WX Central Management System™ (CMS™) software, IT creates configuration templates that remote WX and WXC platforms automatically download from a central device. When deploying a WX or WXC platform in a branch office, the staff simply needs to power up the device and connect it to the network. The device takes it from there, automatically procuring a network address, locating the centralized management server via the domain name service (DNS), requesting a configuration, downloading it, and then beginning operation.

The WX CMS software manages software upgrades and creates, manages and compares WX/WXC device configurations to ensure consistency. The WX CMS software also manages license keys for all WX/WXC devices. With the remote packet capture feature, the WX CMS software also enables analysis for troubleshooting.



**Figure 5:** Juniper WX/WXC WAN optimization platforms can be deployed in configurations to suit all requirements.

## Highly Available Secure Connections

For secure connectivity across public networks, the WX/WXC platforms can be deployed in conjunction with the Juniper Firewall / VPN devices, providing a highly scalable and resilient solution for delivering centralized applications to branch office locations using IPsec. In addition to standalone firewalls, Juniper also offers the Secure Services Gateway (SSG), a high-performance firewall with integrated routing. Organizations deploying the WX/WXC and Juniper SSG platforms can use the virtual routing capabilities of the firewall to provide a solution that delivers intrusion detection and prevention, WAN optimization and IPsec encryption for branch offices and central sites. This solution requires only a single SSG and WX/WXC platform at each site, minimizing the number of devices in the network.

To ensure high availability of VPN connections, Juniper security devices employ a dynamic VPN routing capability that enables the device to choose an alternate VPN tunnel when a connection fails. For security device availability, Juniper provides Netscreen Redundancy Protocol (NSRP), which enables a redundant pair of security systems to be integrated into the network architecture for high availability, eliminating many common causes of system failures, such as a physical port failure or a faulty cable, from bringing down the connection.

## Highly Available Routing Architecture

For organizations that choose to deploy separate routers and firewalls, the WX/WXC platforms can be paired with the Juniper SSG or J-Series Enterprise router products at the branch and with the M-Series routers in the data center. Juniper routing platforms ensure resource availability through an architecture that cleanly separates the routing engine, forwarding engine and services engine. Juniper routers can be configured with redundant power supply units, packet forwarding engine, and routing engine hardware; a hitless routing engine failover capability protects against single node hardware failure.

To further enhance reliability, the Virtual Router Redundancy Protocol (VRRP) prevents a router failure from becoming a single point of failure by allowing seamless failover to a backup router. Separation of routing, forwarding and services ensures that VRRP will work even under severe congestion and other corner cases. The Juniper Networks Graceful Restart feature supports non-stop forwarding by allowing a restarting router and its neighbors to continue forwarding packets without disrupting the network.


## Maintaining Availability at Large Scales

Juniper routers provide market-leading multiprotocol label switching (MPLS) capabilities for organizations that are designing their networks to take advantage of the end-to-end performance and availability capabilities, such as path engineering and fast reroute, that MPLS provides. In MPLS environments, organizations can use the ability of the WX/WXC platform to mark ToS/DSCP to ensure prioritization for applications across the wide-area network over the MPLS core.

Many organizations seek to take advantage of the high availability and performance capabilities of MPLS in the core while also taking advantage of the cost savings of connecting to branches over a public network. In this situation, organizations can scale the deployment with the IPSec-to-MPLS core interoperability capability of the Juniper routers and security devices, thus avoiding the scaling limitations of this point-to-point tunneling technology when it is overlaid in a mesh environment.

## Juniper Networks: Creating Complete Solutions

With the adoption of a centralized server model, organizations face challenges delivering applications to distributed users across the WAN — challenges such as low or limited bandwidth, high latency and congestion. While many organizations have addressed these issues with a WOC to maintain high performance over long-distance links, they are now faced with the greater challenge of maintaining those WAN links to prevent users from being cut-off from centralized applications. To this end, organizations are adopting a networking model that incorporates high availability as a key component of the networking equipment and the network topology.



**Figure 6: Juniper provides a full portfolio of highly available, scalable networking platforms and solutions.**

Juniper Networks provides a portfolio of highly available, scalable networking platforms and solutions that meet the demanding requirements for delivering highly available applications. With market-leading integration features, the Juniper WX/WXC platforms, firewalls, VPN devices and routers deliver a complete solution with the performance, security, reliability and resiliency required to build a highly available enterprise network. As a global company with 24/7 service and support and a long history of product leadership, Juniper Networks is in a unique position to meet the networking and application delivery requirements of the global enterprise.

### Additional Reading

- “Best Practices for WAN Acceleration” (Juniper Networks white paper): [http://www.juniper.net/solutions/literature/white\\_papers/200136.pdf](http://www.juniper.net/solutions/literature/white_papers/200136.pdf)
- “Implementing Data Center Consolidation: How Juniper Networks Enables DCC Through Application Acceleration” (Juniper Networks solution brief): <http://www.juniper.net/solutions/literature/solutionbriefs/351106.pdf>
- “Policy Based Multipath: Avoiding the High Cost of Private Leased-line WANS” (Juniper Networks white paper): [http://www.juniper.net/solutions/literature/white\\_papers/200153.pdf](http://www.juniper.net/solutions/literature/white_papers/200153.pdf)
- “High Availability at the Central Site Edge” (Juniper Networks application note): [http://www.juniper.net/solutions/literature/app\\_note/350062.pdf](http://www.juniper.net/solutions/literature/app_note/350062.pdf)
- “High Availability for Business IP Telephony” (Juniper Networks white paper): [http://www.juniper.net/solutions/literature/white\\_papers/200127.pdf](http://www.juniper.net/solutions/literature/white_papers/200127.pdf)
- “IP Multihoming: Reducing Internet Access Downtime” – The Burton Group (login required): [http://www.burtongroup.com/research\\_consulting/doc.aspx?cid=286](http://www.burtongroup.com/research_consulting/doc.aspx?cid=286)

- “MPLS in Private Networks: Is It a Good Idea?” white paper by Jim Metzler, vice president, Ashton, Metzler & Associates:  
[http://www.juniper.net/solutions/literature/white\\_papers/mpls\\_private.pdf](http://www.juniper.net/solutions/literature/white_papers/mpls_private.pdf)

---

Copyright © 2006, Juniper Networks, Inc. All rights reserved. Juniper Networks and the Juniper Networks logo are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered trademarks, or registered service marks in this document are the property of Juniper Networks or their respective owners. All specifications are subject to change without notice. Juniper Networks assumes no responsibility for any inaccuracies in this document or for any obligation to update information in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.