

HIGHLY AVAILABLE NETWORKS

Resiliency, reliability and security are the key concerns for maintaining network resources.

Table of Contents

-
- 2 Executive Summary

 - 2 Optimize the Network

 - 3 Update Network Management Tools

 - 4 Boost Resiliency with Wireless

 - 5 Redouble Security Efforts

 - 6 Prepare for Disaster

 - 7 Gain Control of Power and Cooling

Executive Summary

Government is moving toward a digital future. As the public has increasingly embraced online interaction with its governmental organizations, these agencies have in turn responded by building up their digital service offerings.

To meet the demands for such online services, government organizations are now strengthening and updating their network infrastructures. In addition to allowing for incremental efficiencies, productivity and revenue, such initiatives result in a more satisfied constituency.

The agency network is at the heart of this digital future. As such, its resources are expected to be available to staff and the public at all times, with as little tolerable downtime as possible. Unplanned network disruptions can have wide-ranging consequences: disruption of public services, lost revenue, paralyzed internal operations and unattractive publicity.

Giving the tremendous negative downsides of such a disruption, keeping the network highly available and operating at peak levels has become a key concern for government IT managers. To facilitate this availability, organizations need a comprehensive approach to their networks that builds on six core pillars of resiliency and reliability:

- Network optimization
- Network management tools
- Wireless LAN best practices
- Security policies
- Continuity of operations (COOP) planning
- Power and cooling strategies

By examining and updating each of these areas, IT managers can create a wide-ranging strategy for achieving highly available networks while also addressing today's performance, policy and compliance realities.

Optimize the Network

Developing a comprehensive plan for highly available networks begins by finding ways to better manage network traffic and the organization's essential data. That means utilizing infrastructure resources so data can find the most efficient routes, move as quickly as possible from sender to receiver and keep information readily accessible and safe from corruption and loss.

Balance the Load

To improve traffic-routing efficiency, look to load-balancing technology. It provides administrators with one of their most important tools for prioritizing network traffic and making sure it reaches its destination, even if a critical component, such as an application server, malfunctions.

Load balancing may be accomplished with hardware appliances or software. Both methods direct network traffic

using predefined prioritizations based on what kind of data is being moved, where it's coming from and where it's going. Together, these characteristics help IT managers determine the importance of the information and manage it accordingly. The tools then find the most efficient routes to send the data and scout for trouble spots to avoid.

For extra peace of mind, many IT managers pair load balancers in redundant configurations. If the primary implementation malfunctions, the backup is ready to step in before network performance degrades.

Modern load balancers offer advanced features to help them work effectively. Look for capabilities such as Secure Sockets Layer (SSL) termination, which frees application servers from expending processing power to decrypt traffic, and compression, a feature that helps IT infrastructures send smaller data volumes over the wire for improved overall performance.

Accelerate the WAN

A second essential for network optimization is a WAN accelerator, which works to keep latencies to a minimum and circumvent problems that may interrupt traffic flows. The best WAN accelerators come with a central management console to make it easy for administrators to monitor the collective performance of the entire accelerator fleet and to display summaries of traffic characteristics to uncover developing problems.

Accelerators should also be able to maintain local storehouses of data. So instead of clogging the network with popular files, accelerators can relay streamlined instructions for a companion device to pull the frequently accessed data from its cache.

A number of enterprise applications, including e-mail and content-management systems, send data across networks using Transport Control Protocol (TCP), which can add network-clogging overhead to the infrastructure. Look for WAN accelerators that include components for speeding TCP processing to overcome this bottleneck.

Finally, to minimize efficiency-sapping retransmission of data packets, choose WAN accelerators with packet order correction (POC), which can quickly recompose scrambled packets into their proper sequencing.

Select the Right Storage

Effectively routing and accelerating data is of little use if the information itself becomes unavailable or corrupted because of technical problems. The first line of defense for data protection is centralized storage systems, which mitigate the consequences of equipment failures and bring efficiencies to backup and COOP strategies.

Centralized storage comes in two main forms: network-attached storage (NAS) and storage area networks (SANs).

SPOF-proof Networks

What's the biggest threat to achieving a highly available network? The dreaded single point of failure. A SPOF may be as simple as a defective hard drive that suddenly stops spinning or something much more devastating, such as a hurricane that brings down an entire data center. Either way, the results may be similar: staff are cut off from the data and applications they need to do their jobs.

One of the best ways for IT managers to mitigate SPOFs and ensure uptime is to introduce redundancy wherever possible throughout the network. That means multiple servers, network switches, load balancers and storage systems designed for fast and efficient failover when disaster strikes.

The problem is that in an era of tight IT budgets, CIOs work hard enough to fund new initiatives, let alone find extra dollars to spend on redundant equipment which may never be needed. So rather than blindly chasing every worst-case scenario, IT managers need to balance acceptable risk with budget realities.

That starts with classifying systems and data according to how critical they are to the organization. Then, IT departments can do everything possible to fully protect mission-critical resources and take a more economical approach to managing those areas of the network that can suffer some degree of downtime without severely affecting the organization.

NAS packs hard drives into dedicated storage servers that maintain their own unique LAN connections.

This configuration allows any networked server or PC to pull information from the NAS devices as easily as from a local hard drive. The difference is that the drives in the NAS box work together to pool capacities and maintain redundant copies of files in case of an individual drive failure.

SANs extend the concept of networked storage to cover the entire organization by creating pools of storage resources that may include disc-based storage servers, as well as tape and optical libraries. SANs typically use Fibre Channel or iSCSI interfaces for faster performance than IP-based NAS units. SANs also allow organizations to manage their entire storage investment from a central control.

NAS and SAN devices each offer unique benefits for optimized networks. NAS is a cost-effective tool for asynchronous communication that sends backup copies of files to remote disaster recovery sites, including increasingly popular infrastructure-as-a-service (IaaS) clouds for offsite replication. Even if a human-caused or natural disaster hits a primary data center, the distant data storehouse will keep important information safe.

High-speed SANs can transmit data synchronously, which, when combined with redundant array of independent disks (RAID) technology, supports strategies for high availability. Depending on the level of RAID an IT manager chooses to implement, SANs can copy portions of files or entire data volumes on duplicate hard drives or across multiple local or geographically dispersed devices.

When IT managers orchestrate load balancing, acceleration technologies and centralized storage systems, information remains readily accessible to enterprise applications, no matter what availability challenges hit the organization.

Update Network Management Tools

To enable IT managers to keep up with the new performance and reliability demands on today's network infrastructures, some tried and true network management technologies are evolving to include advanced features. These capabilities can boost availability by reducing overhead, provide additional processing power or offer detailed analyses of network health and red flags for emerging problems.

Optimize the WAN

WAN optimization controllers have long given managers a way to reduce communication latencies that crop up in connections between branch offices. To do this, optimization appliances employed a mix of enhanced transmission control protocols, data compression and caching to minimize the sheer volume of information going over the wire. Further optimizations focused on the overhead levied by SQL implementations, security systems and other enterprise applications.

Today, WAN optimization controllers play an even more active role in highly available networks. They can monitor traffic patterns as data flows across WAN links and alert IT managers to bottlenecks before they lead to significant communication breakdowns.

Traditionally offered as a hardware appliance, WAN optimization controllers are also available as software-based utilities that managers can reconfigure on the fly to accommodate the dynamic provisioning of IT resources in virtualized data centers and cloud implementations.

By combining new features with venerable capabilities, WAN optimization controllers help existing network links handle increased traffic demands and help IT organizations avoid the reliability-sapping consequences of bandwidth overload.

Control Application Delivery

Application delivery controllers, or ADCs, are on a similar evolutionary path. Springing from traditional load-balancing devices, ADCs can now analyze the workloads of application servers and automatically take over select processing duties.

These include data compression and compute-intensive SSL termination, which keeps application server CPUs from being overloaded. In addition, some of today's ADCs can also manage overhead associated with common protocols, such as TCP and HTTP.

ADCs fit comfortably in today's highly dynamic virtualized data centers by offering processing power to help manage security certificates and networking protocols, rather than leaving that burden to be performed by individual virtual machines (VMs).

An added bonus: Because ADCs draw a fraction of the power of typical enterprise servers, the compact appliances offer ongoing operational cost savings compared with full-size hardware.

To receive the full benefit of WAN optimization controllers and ADCs, IT departments need to see and understand performance across their entire landscape of physical, wireless and virtual networks and have an early-warning system that can spot outages and growing bottlenecks in real time.

Gain Visibility

Increasing visibility requires network monitoring systems that have basic and advanced analysis features for calculating network traffic performance, resource availability and uptime rates. Basic capabilities include Simple Network Management Protocol (SNMP) pings to assess data packets for key benchmarks, such as jitter, latency and packet loss.

In addition, today's network monitoring systems can unleash software agents across the network to probe for weak spots in network gear or send test messages to assess the health of web and e-mail servers.

When the system uncovers problems, it should send a barrage of alerts to the main console screen, as well as to smartphones and e-mail accounts, to ensure that everyone who needs to get involved becomes aware of the situation.

Finally, it may not be enough to just see real-time displays of network status. Network management systems should also be equipped to produce statistical summaries and historical reports to help administrators spot long-term trends and document the need for new resources.

Boost Resiliency with Wireless

Thanks to the higher speeds and technical innovations ushered in by the 802.11n standard, wireless networks are becoming firmly established as essential infrastructure. But greater reliance on wireless networks represents more than increased convenience for end users – the same criteria for high availability and performance that apply to wired networks also affect their wireless counterparts.

Fortunately, organizations no longer sacrifice data transmission speeds when they send data across wireless links. Thanks to

Network Management System Checklist

When shopping for network management tools, look for candidates that provide the following functions:

- Create visual models of hardware, interfaces and traffic patterns
- Provide a holistic view of performance for physical, wireless and virtual components
- Use agents to scour the network for performance problems
- Offer easy-to-understand graphics and dashboards
- Monitor traffic flow, web- and e-mail-server status, and Active Directory performance
- Encompass both communications and data systems, such as Voice over IP phones
- Send alerts to desktops and mobile devices
- Allow administrators to take action by adjusting network configurations
- Develop trend analyses and statistical reports

802.11n, these networks can match or, in some cases, surpass the transmission rates of standard 100BASE-T wiring.

But there are a few caveats. First, organizations must closely align 11n clients with similarly configured access points. Physical proximity isn't the only consideration.

Signals from existing 802.11a/b/g devices operate at slower speeds and will be intercepted by backward-compatible 11n networks. This will degrade performance for the newer and faster access points and switches.

Organizations can minimize the negative impact on 11n performance by taking advantage of the latest standard's ability to use both 2.4GHz and 5GHz channels. The latter remains free of most legacy traffic and provides a clear pathway for high-speed communications.

To fully integrate 11n wireless into wired network infrastructures, IT managers should upgrade wired-to-wireless links so they support Gigabit Ethernet. Administrators can help justify these upgrades by citing high-availability benefits: If a wired switch suddenly crashes, the wireless access point can step in to keep traffic flowing at acceptable rates.

Central Control

Ramped-up data transmission rates aren't the only news in wireless networks. Also of interest to organizations seeking high availability are centralized wireless LAN (WLAN) controllers. They put an end to the old chore of having to

physically interact with each access point whenever it's necessary to adjust configurations.

Instead, IT managers can take advantage of a central database that stores configuration settings for each device. In addition, WLAN controllers can overcome some transmission glitches by automatically tweaking power settings and channel designations.

Further aiding availability, WLAN controllers can graphically represent wireless layouts and include information about each device's specific location and performance level. The graphical consoles include utilities to make necessary changes when impending problems surface.

WLANs also provide a host of failover mechanisms. If an access point falls off the grid, a different unit can temporarily take up the slack. Good architectural design practices call for redundant wireless switches for similar failover capabilities when implementing these essential components.

Redouble Security Efforts

A long series of regulations have made auditable security practices and compliance reporting mainstays of public sector operations over the past decade. And with increasingly sophisticated hackers, a more mobile workforce and Web 2.0 computing models, organizations need to find new ways of thwarting attacks that can bring networks to a standstill.

One of the biggest changes IT organizations face is that they can no longer look to perimeter security for

complete protection. Instead, they're increasingly turning to technologies and techniques that emphasize guarding information rather than infrastructures. To do this effectively, network administrators need to orchestrate a mix of tools.

VIRTUAL PRIVATE NETWORKS: Whether implemented through hardware or software, VPNs use protocols such as SSL or Internet Protocol Security (IPsec) to establish encrypted tunnels. Because VPNs establish secure links directly between users and enterprise resources, they're particularly effective for mobile workers who need to safely access protected data and applications while away from the main office.

APPLICATION-LAYER FIREWALLS: Traditional firewalls examine Transmission Control Protocol (TCP) information to create a barrier and regulate traffic between internal networks and the public Internet. Application-layer firewalls, by contrast, are sophisticated enough to read application protocols to actively understand when an attack may be taking place. IT managers can also choose to enable utilities that allow firewalls to inspect incoming content for any suspicious characteristics that may indicate malware or spam.

NETWORK ACCESS CONTROL: NAC appliances or software act like a virtual security guard that inspects the credentials of devices trying to log into the network. NACs check whether the new devices may be infected with malware or, if granted network access, whether the devices should be allowed to interact with all or just a subset of databases, applications and other resources.

NACs come in a variety of configurations to match the size of the organization they're designed to protect. Inline NACs integrate with virtual LANs to monitor incoming IP and media access control (MAC) headers to authorize or deny outside devices trying to establish network connections. Deployment ease makes this option a common choice for smaller organizations.

Midsize agencies often opt for out-of-band NACs, which are appliances that assess traffic for safety while using fast processors to avoid some of the potential bottlenecks associated with inline varieties. Large enterprises typically turn to DHCP-registration NACs, which use DHCP and IP subnets, rather than virtual LANs (VLANs), to create secure connections between networks and outside users.

INTRUSION PREVENTION SYSTEMS: IPS appliances provide an added layer of security by inspecting incoming data packets for known threat characteristics or blocking in-progress denial-of-service attacks. The appliances use dedicated, high-performance processors to conduct inspections and keep performance delays to a minimum.

IPS solutions may also embed tools designed to identify attackers trying to use protocol violations and other advanced techniques. When unusual traffic patterns surface, IPS devices can send alerts to administrators. IPS appliances may also

Wireless Purchasing Pointers

Wireless networking has come of age, with transmission rates and resiliency characteristics on par with traditional wired networks. Here are some guidelines to follow when trying to choose the right wireless gear.

Wireless switches: 802.11n devices should support Layer 2 and Layer 3 switching and offer Gigabit Ethernet (GbE) switching for high-speed integration with wired networks. Also look for automatic failover capabilities when associated devices fail, as well as "mesh backhauling" that takes over during failures in wired networking segments.

Wireless switches should also be able to handle a variety of voice applications and provide expansion options to ease upgrades to emerging WAN technologies, such as the evolving WiMax standard. Finally, confirm and evaluate the effectiveness of an embedded intrusion-prevention system.

Access points: Choose 802.11n devices that include Power over Ethernet Plus (PoE+), which allows IT managers to place the hardware wherever networking demands warrant, even if an electrical outlet isn't available.

accept data feeds from security firms and other sources for updates on the latest threat profiles and automatically adjust their filtering operations.

How to Secure Wireless Networks

Traditional networks benefit from a wide range of proven tools to control access to internal resources. Most of these solutions can't be directly bolted onto wireless networks, which require a unique approach to security. Fortunately, security practices are maturing to address wireless networks and their connections to the primary wired infrastructure.

When devising wireless security strategies, network administrators must remain wary of "spoofing" assaults, a long-time practice where hackers hijack the communications of users who believe they're sending sensitive information on a secure pipeline.

Defending against the vulnerability is complicated by the fact that wireless radio signals can travel through walls, leaving networks open to intrusions outside an organization's building.

Start with enabling the encryption and authentication capabilities that come standard with switches and access points. IEEE-standard Wi-Fi Protected Access 2 (WPA2) uses Advanced Encryption Standard (AES) algorithms for powerful data encryption protections.

Then take it a step further with measures for securing the areas where wired and wireless networks intersect. Special intrusion prevention systems for wireless environments can help network administrators quickly identify unauthorized devices trying to break through the security defenses.

Wireless-savvy IPS devices can also beat back denial-of-service attacks designed to crash networks. Geofencing (a virtual perimeter around a geographic site) and other techniques allow IT managers to grant access only to devices running at known and trusted physical locations.

Administrators can create virtual LANS (VLANs) and regulate traffic using access control lists (ACLs) to guard against vulnerabilities that arise when guest users need to connect to the Internet over a wireless link. Alternately, IT managers may dedicate a wireless LAN controller for this purpose and use it to divert guest-user traffic to a secure location outside the organization's firewall.

The good news is that wireless LANs not only achieve the performance rates of wired networks, but also are equally secure when IT managers implement the right security measures.

Prepare for Disaster

Even the most carefully architected high-availability networking strategy – complete with redundant equipment and the latest optimization technologies – can't guarantee that disaster will never strike.

The right approach may mitigate the consequences of component failures, human-caused and natural disasters, malicious attacks, and other threats. But the risk will always remain that a system may fail sometime, somewhere in the infrastructure. That's why all comprehensive high-availability efforts require well-crafted COOP initiatives.

A COOP plan outlines the policies and procedures an organization will launch to recover data and bring operations back to life as quickly and safely as possible. Creating detailed plans before a crisis occurs helps assure a cool and well-coordinated response when problems strike and tensions run high.

COOP has become significantly easier to accomplish with the rise of server and storage virtualization. When they're combined with newer trends in desktop virtualization, these technologies enable IT managers to maintain continuity of operations with instant system failovers, fast recoveries thanks to automated backups, and dynamic reprovisioning of hardware, applications and storage systems.

VMs Become Ubiquitous

Virtualized servers act as a foundation for COOP. By packing the power of multiple application servers into one physical server, virtualization offers a proven way to cut IT costs and consolidate data centers during normal times. When disaster hits, virtualization delivers a number of additional benefits.

First, the technology will have helped IT managers with some of the COOP prep work – virtual servers can ensure that regular backups run as scheduled by reducing complexities associated with large numbers of physical machines in traditional environments. Administrators can designate virtual servers to continuously replicate data to a recovery site or decide on a range of less-demanding backup operations for less critical data.

Then, in the aftermath of a disaster, data or applications can be relocated onto alternative virtual machines or moved to a backup data center. IT managers make the move using a central management console.

For effective backups, IT administrators must be sure their data protection applications are updated for virtualized environments. Most virtualization-aware backup programs use dedicated software agents to help direct backups from virtual servers. Administrators also need to make certain that their backup software gives them the capabilities they need to keep the cause of the failure from being replicated on a new virtual machine.

Virtual Storage Benefits

Storage virtualization creates similar advantages for COOP by grouping several storage systems together into one large virtual storage pool. By pairing storage and server virtualization technologies, managers can quickly relocate data across the IT infrastructure to reduce downtime and maintain high-availability goals.

The latest technologies for storage arrays and appliances give IT managers new storage-management tools helpful for virtualization. They include data deduplication, which eliminates redundant copies of data that can overwhelm regularly scheduled backups and result in incomplete operations.

A key question that IT managers must answer when implementing storage virtualization is whether to virtualize data at the file or block level. File-level virtualization, typically associated with NAS and file servers, is well suited for web applications that need to scale quickly to meet file access demands. Block-level virtualization using SANs works well for database applications that access hard drives directly without interacting with file systems.

Paths to Successful Backup and Recovery

IT administrators have three main choices when it comes to backup and recovery strategies in virtualized-server environments.

Full backups: This approach essentially treats virtual-machine backups the same as backups for physical machines. Data is fully protected, but organizations grapple with high input/output overhead on their networks.

Backups with data deduplication: Administrators eliminate redundant data on virtual machines and then focus backups only on updates to data stores. Deduplication is performed on the virtual machines, and only the changes are backed up. This minimizes network traffic, but may increase recovery times in the aftermath of a failure.

Snapshot backups: Regularly scheduled snapshots of data on virtual machines are backed up on proxy servers at offsite locations. Windows virtualization technology lends itself to this option.

Because each approach differs in the completeness of backups and other factors, organizations should evaluate choices in test environments before committing them to production systems.

Client Virtualization Gains Traction

Client virtualization severs direct links between end-user hardware and the applications and data that staff access for their jobs. Rather than residing on desktop devices, the software and data is stored centrally in data centers, where IT administrators have direct control over its management and security. Running applications in a central location also makes it easier to administer updates and patches, an additional gain toward security efforts.

For end users, virtualization addresses the desire for more choice in the types of devices they'll use to do their work – desktop virtualization accommodates everything from standard PCs, notebooks and thin clients to tablets and smartphones. For high-risk security environments, organizations can increase security controls on end-user devices using smart-card readers, biometric scanners or two-factor authentication tokens.

But IT managers must avoid letting desktop virtualization become a weak link in their high-availability strategies. That means evaluating the entire IT infrastructure to be sure network bandwidth capacity and other resources can hold up to the increased demands of new virtual clients constantly communicating with data centers. Pay particular attention to network latency, which can make the performance that end users experience unacceptably slow for daily operations.

Gain Control of Power and Cooling

The best-laid plans for high availability may fail if IT managers overlook fundamental components of their total operations. This includes the ability to manage power consumption and reduce the risks associated with electricity spikes, brownouts and blackouts. Even minor disruptions can cause major damage to sensitive hardware, leading to downtime for critical systems and costly disaster recovery events.

Day-to-day costs are another concern. Technology researcher Enterprise Management Associates estimates that for every dollar an organization spends on new data center hardware, it needs to spend another 50 cents on power and cooling resources.

These challenges are spurring many organizations to make heating, ventilation and air-conditioning (HVAC) operations an integral part of their high-availability strategies. Fortunately, there are a number of new options available for managing power issues.

Many IT organizations are already taking the first step toward improved HVAC practices by expanding server virtualization across their enterprises. Virtual machines reduce the overall number of physical servers that organizations need, which relieves pressures on power and cooling demands.

For example, packing 20 virtual machines into one physical server is more energy efficient than powering 20 individual physical servers, which, according to industry estimates, typically run at only a fraction of their rated capacity.

Be Cool

Conscientious organizations also need to pay close attention to cooling systems. One reason relates to the transformation many IT infrastructures have undergone in recent years when they installed large numbers of multi-core blade servers into banks of computing racks.

IT managers are finding it tough to adequately cool these tightly packed processing engines. These dense hardware arrangements result in “hot spots” that stress cooling systems and potentially lead to system downtime from overheating.

To alleviate hot spots, administrators are looking beyond traditional cooling methods that simply use central air-conditioners to maintain a cool temperature throughout an entire room. Newer approaches mount air-conditioners at individual rows of blade servers or onto the racks themselves to direct cool air where it's needed most. Another option is to install a mix of room and dedicated AC units to maintain consistent temperatures throughout the data center.

Monitor Consumption

Gathering detailed information about consumption rates is another important step in managing and maintaining power supplies. This requires more effort than merely checking with facilities managers for each month's electricity bill.

Many enterprise-class uninterruptible power supply (UPS) units now come with embedded monitoring tools that show consumption rates for the entire data center down to workgroups and individual components. These power-savvy reporting tools highlight important data points for organizations, including whether voltage levels are running within acceptable ranges, if there are any emerging power problems and the current battery status of each UPS unit.

The Right Stuff

Finally, organizations can increase energy efficiency by purchasing the right hardware. Some servers today

How Green Is Your Data Center?

The Power Usage Effectiveness rating, or PUE, is a calculation developed by The Green Grid, an association of IT manufacturers. The purpose is to provide an easy way to benchmark data centers on how efficiently they consume power.

To determine a PUE rating, an IT manager divides the total amount of power dedicated solely for the data center by the amount of power consumed by the IT resources used for managing, processing, storing or routing data within the data center.

The lower the PUE rating, the better in terms of how efficiently a facility utilizes energy. A PUE rating of 1 means that a facility is 100 percent energy efficient, something Green Grid members acknowledge is an ideal more than a real-world reality.

Conversely, a 3.0 rating indicates that power demands are three times greater than theoretically necessary to power the IT equipment. Expressed another way, a PUE 3.0 data center would need 1,500 watts of power to run a server rated at 500 watts of electrical draw – a 33 percent premium. This information should spur administrators to analyze their operations to find inefficiencies that are squandering energy supplies.

The Green Grid says that a lack of comprehensive data currently makes it impossible to correlate an individual organization's PUE rating with industry averages. But early studies show many data centers come in at 3.0 or above. The Green Grid adds that the right designs and power management practices should put PUE 2.0 within an IT manager's reach.

automatically adjust CPU power draws to decrease consumption during off-peak periods.

In addition, federal agencies should adhere to mandates calling for them to purchase equipment certified for Energy Star compliance and use Energy Star Portfolio Manager to gauge their building's energy use. Many state and local governments are following suit by requiring their agencies to install or at least give preference to Energy Star equipment.