

# NETWORK INFRASTRUCTURE

Building a more efficient network with greater capabilities, capacity and security

---

800.808.4239 | [CDWG.com/networkguide](http://CDWG.com/networkguide)



## CDW-G REFERENCE GUIDE

A guide to the latest technology for people who get IT



# NETWORK INFRASTRUCTURE REFERENCE GUIDE

---

## IN THIS ISSUE:

<b>CHAPTER 1: Network Evolution</b> .....	3
· Where Networking Is Going	
· Network Improvements	
<b>CHAPTER 2: Data Center and Infrastructure</b> .....	6
· Server Consolidation and Virtualization	
· Data Center Virtualization	
· High Availability and Redundancy	
· Cloud Applications	
<b>CHAPTER 3: Network Optimization</b> .....	21
· Building Resiliency	
· Creating Visibility	
· Increasing Flexibility	
<b>CHAPTER 4: Wireless Mobility</b> .....	26
· Wireless Standards	
· 40MHz Channels	
· Enterprise-class 802.11 Networks	
<b>CHAPTER 5: Network Security</b> .....	30
· Network Security Concerns	
· IPS and DLP	
<b>GLOSSARY</b> .....	33
<b>INDEX</b> .....	35

## WHAT IS A CDW·G REFERENCE GUIDE?

---

At CDW·G, we're committed to getting you everything you need to make the right purchasing decisions – from products and services to information about the latest technology.

Our Reference Guides are designed to provide you with an in-depth look at topics that relate directly to the IT challenges you face. Consider them an extension of your account manager's knowledge and expertise. We hope you find this guide to be a useful resource.

# NETWORKING EVOLUTION

## Preparing Enterprise Networks for the Future

"The network is the computer." When Sun Microsystems popularized this phrase several years ago, it was helping usher in a 21st century in which networks have become indispensable. Without a secure infrastructure to connect information, systems, stakeholders and staff, a great deal of productivity would cease at many organizations today.

This move toward ubiquitous connectivity is pushing IT managers to rethink and retool their enterprise networks. Keeping an enterprise network up to date in 2011 requires more consideration than blindly swapping 10/100 switches for 10/100/1000 switches and upgrading the organization's Internet connections.

Every IT department has to make technology and design decisions to accommodate the rise of virtualization, mobile and cloud computing, and the ever-present risk of security breaches.

This reference guide discusses the tools and technologies that will help IT managers develop a strategy for future-proofing their

networks. It will focus on changes in data center infrastructure, network optimization, wireless mobility and network security – all the information required to adapt to the changing dynamics of enterprise networks.

### Where Networking Is Going

Although many of the same principles that IT professionals practiced during the rising adoption of PC LANs and the Internet still apply, much has changed. Heading into 2011, there are at least five major trends at work that will determine how IT departments will deploy enterprise networks. (See the sidebar "Change Is on the Way: 5 Networking Trends" on page 4 for more information.)

**1. User expectations:** Users now expect to interact with the organization using the Internet and other interconnected networks. While the trend has been visible for a long time, with the proliferation of mobile devices and an increasing reliance on the Internet for productivity, even progressive organizations need to re-evaluate their networking strategies.

**2. Security concerns:** The ubiquitous nature of the network puts all the organization's assets within reach. Unfortunately, that reach can also extend to people with malevolent intent.

Networks full of powerful computers, high-speed connections and organizational and personal data are prime targets for criminals. These networks are the bank vaults of today's economy. Separating good data from potential viruses is a constant battle with a changing enemy, one that requires the IT staff to compensate and adapt its tactics and technology.

**3. Data center changes:** The technological pendulum that swung from huge, central mainframes to many small, distributed servers has swung back toward centralized computing in the form of virtualization and server consolidation. The result is the same: more and more data moving through a smaller number of physical systems.

This has created a new focus on extremely redundant designs, very high bandwidth concentration and capacity, more granular security

## Change Is on the Way: 5 Networking Trends

Trend	Design Changes	Technology Changes
<b>1. User expectations</b>	<ul style="list-style-type: none"> <li>• Scaling up via multiple data centers, content distribution networks and geographic load balancing</li> <li>• Tighter integration of old green-screen apps with self-service web apps</li> <li>• Rethinking operations to move to 24x7 availability</li> </ul>	<ul style="list-style-type: none"> <li>• Federated identity management that links to other web service providers</li> <li>• Tunnels to e-commerce and fulfillment partners</li> <li>• Secure Sockets Layer (SSL) accelerators</li> <li>• SSL decryption appliances</li> <li>• Quality of Service (QoS) management for traffic</li> </ul>
<b>2. Security concerns</b>	<ul style="list-style-type: none"> <li>• Increased use of defense-in-depth with built-in layers of apparently redundant security</li> <li>• More internal access control points within all Layer 3 devices</li> <li>• More traffic inspection points designed in</li> </ul>	<ul style="list-style-type: none"> <li>• High-speed, high-density firewalls with many interfaces</li> <li>• Gigabit-speed intrusion prevention system (IPS)</li> <li>• Data loss prevention, both for outbound and inbound traffic</li> <li>• Application-layer firewalls for servers</li> <li>• Network access control (NAC) and application-control tools for users</li> </ul>
<b>3. Data center changes</b>	<ul style="list-style-type: none"> <li>• Completely redundant design everywhere</li> <li>• Collapsed backbone into chassis-type products to reduce Layer 3 routing decisions</li> <li>• Higher density in-cabinet switches</li> <li>• Green power and HVAC-aware design</li> </ul>	<ul style="list-style-type: none"> <li>• Link aggregation at switches to scale up speeds</li> <li>• 10Gbps interfaces to top-of-rack or end-of-row</li> <li>• 10Gbps directly to blade servers</li> <li>• Application delivery controllers placed in front of server farms</li> <li>• Converged data/storage networks</li> </ul>
<b>4. Reliability requirements</b>	<ul style="list-style-type: none"> <li>• Application designs that make use of multiple redundant systems and redundant storage elements distributed across data centers</li> <li>• Active-active device design</li> <li>• Change in design to permit rolling upgrades of infrastructure without downtime</li> </ul>	<ul style="list-style-type: none"> <li>• High-speed inter-data center links</li> <li>• Data deduplication</li> <li>• Continuous data protection backup</li> <li>• Use of redundant 10Gbps interfaces</li> <li>• Rapid speed spanning tree protocol (RSTP)</li> </ul>
<b>5. Mobility, the cloud and more efficient WANs</b>	<ul style="list-style-type: none"> <li>• Rethink apps to reduce traffic for WAN and mobile users</li> <li>• Push apps to cloud</li> <li>• Unified communications, linking voice, video, presence and apps</li> <li>• Threat mitigation at the perimeter for VPN users</li> </ul>	<ul style="list-style-type: none"> <li>• Wireless LAN 802.11n deployments</li> <li>• WAN acceleration for branches</li> <li>• UTM security devices</li> <li>• Mobile device management tools</li> <li>• Service level agreement (SLA) monitoring systems</li> </ul>

and monitoring tools, and green computing.

**4. Reliability requirements:** Planned network downtime is a luxury that many organizations are finding increasingly difficult to accommodate. Unexpected downtime is even less welcome. Old-school ideas such as disaster recovery are losing ground to continuity of operations (COOP). This suggests a shift in thinking away from concern for network availability only during disaster scenarios and evolving toward continuous availability – whatever the scenario.

**5. Mobility, the cloud and more efficient WANs:** People will always work together in groups. But for many organizations, being in the same room at the same time is becoming less important.

To accommodate the need for interaction across distances, network managers are building Voice over IP (VoIP) and video conferencing networks, using cloud-based services, supporting work-anywhere mobility tools such as smartphones and notebooks, and upgrading WANs to handle latency-sensitive and bandwidth-sensitive applications.

## Network Improvements

Networks are always in a state of flux, constantly evolving to meet the needs of the organization and its users. These needs often change as technology evolves and improves. Skillful IT management includes sensing when a particular technology development might align very well with a pressing network need and applying it.

This requires research and testing – knowing what technologies are on the horizon and what they might be able to add to the network. There are four technologies in particular that IT managers should start to learn more about because they will more than likely affect their network in the near future.

**Virtualization:** Whether it is the desktop, the server, storage, the switch or router, or the firewall, virtualization is a growing presence throughout the network. Virtualization offers so many benefits that organizations of all sizes are implementing the technology, especially server virtualization.

By deploying farms of virtual hosts housing tens, hundreds or even thousands of guest virtual machines (VMs), organizations are seeing higher utilization of hardware, which saves space and lowers power and cooling costs. Virtual hardware platforms reduce dependence on particular brands or generations of devices, simplifying upgrades, disaster recovery and COOP planning. Organizations also report increased agility, being able to roll out servers quickly and efficiently.

**10-Gigabit networking:** 100BASE-T Ethernet is now the base LAN device speed, while 10-gigabit-per-second (10Gbps) gear is the minimum for core interconnectivity in the LAN. Make sure to invest in equipment that's future-proof up to those speeds. Faster devices with 40Gbps speeds are coming, but they will be more useful to carriers and hosting service providers.

The new standard for distribution-to-core links is 10Gbps. These new higher speeds make high-speed wireless a possibility. Each access point requires a 1Gbps connection today, so it makes sense to have 10Gbps Ethernet in the wiring closet. Remote deployment technology and virtual desktops also put heavy burdens on the edge of the network. And without 10Gbps links to wiring closets, both wireless and wired systems can't reach their full potential.

In the data center, 10Gbps is an investment requirement. Large-scale server virtualization needs that level of bandwidth to operate. And 10Gbps Ethernet switches also support large databases and backup-to-disk initiatives.

**Wireless:** The new base level for wireless is 802.11n, in both the 2.4 gigahertz and 5GHz bands. Organizations will want to be prepared to remove old 802.11b equipment from the network. Wireless networks are not just for guests; they're being upgraded and used by staff to make collaboration more efficient.

Initiatives such as paperless meetings, which are green and save money, can work only if the organization has a strong wireless network to support them. It also takes wireless to support webcams on notebooks, which makes it possible for more people across the organization to take advantage of the communication and collaboration benefits of video conferencing.

And, finally, as organizations roll out more wireless notebooks and other devices, upgrading the wireless network ensures that they can support the added traffic and make staff more productive, maximizing the investment in wireless.

**Multilevel security:** Firewalls positioned only at the perimeter is an outdated security strategy. Organizations are now embedding firewalls throughout the network, with new requirements for speed, reliability and manageability. Be prepared to re-evaluate everything about the organization's security architecture. Cleaning up security problems is very expensive, in terms of both time and lost organizational reputation.

Many IT managers have discovered that the old three-zone segmentation of their network (inside, outside, DMZ) leaves them open to infections and insider threats. Pushing security tools into the core of the network helps contain and block attacks, no matter what the source. At the same time, increased emphasis on managed endpoint security helps protect devices from infections in the first place. ■

# DATA CENTER AND INFRASTRUCTURE

## Planning for Virtualization in the Data Center

In many ways, strategies around network infrastructure in data centers and buildings have changed little: numerous Ethernet ports (usually at higher speeds), a heavy emphasis on Transmission Control Protocol/Internet Protocol (TCP/IP) networking with more mobility and wireless, and faster security devices to protect it all.

But scratch beneath the surface and it's clear that there's a big difference between today's networks and the first PC LANs. Fundamental changes in the way that data centers are built and in how end users connect to applications are driving both design and technology changes throughout networks in many organizations.

### Server Consolidation and Virtualization

Data center networks are undergoing a refresh driven by the biggest changes in IT design in the past decade: virtualization and server consolidation. During the 1990s, what used to be called the mainframe was

broken out into hundreds, sometimes thousands, of different servers.

Now, the pendulum is swinging back, but with a twist. All of those services are being pushed into a relatively small number of high-performance virtual host hardware devices, thereby reducing the complexity of managing physical hardware but increasing configuration needs. At the same time, all these virtual hosts are having their internal storage stripped away and moved to storage area networks (SANs).

### How Virtualization Changes Data Centers

Virtualization and server consolidation have resulted in three priority requirements for data center networks: high density, high speeds and very high reliability and redundancy. Virtualization also has shifted the emphasis in data center networking to reduce reliance on Layer 3 routing in favor of Layer 2 switching.

This is because Layer 2 switching gives the virtualization manager

greater flexibility to move VMs between hosts and – most important – between different data centers in the same or different buildings and campuses. The concentration of more services into fewer devices has also brought out an intense focus on monitoring and management.

### Virtualization Best Practices

Very large virtualization environments are new to the IT industry, so best practices and recommended deployment approaches are still being worked out. However, most virtualization architects prefer to build very large clusters of servers connected to a common SAN.

This lets them scale up and down dynamically and easily move applications transparently from one piece of physical hardware to another – which may involve moving across the room, across town or across the country.

For the network manager, the result is often racks of "pizza box"

servers, each 1.75 inches high (referred to as 1U) and stacked 30 or more together into a rack. Going from four huge servers in a rack to 30 1U servers is a major challenge for the data center manager who has to deal with power and heating, ventilating and air-conditioning (HVAC) issues. But it also presents challenges for the network manager.

Each of those servers will have a minimum of two Gigabit Ethernet ports (more likely four) and often a connection for "lights out" management (which enables remote management of servers without physically attaching a keyboard, monitor and mouse). These servers may also have additional Fibre Channel ports for connections to the SAN. In a few years, those two Gigabit Ethernet ports will probably expand to two 10-Gigabit Ethernet ports.

Advice on how to configure those ports varies depending on the consultant's preferences. Most network managers are settling on building the largest trunk possible. This means combining at least two (but preferably four, and possibly even more) Gigabit Ethernet ports into a single large trunk port with virtual LANs used to separate the different subnets.

The same trunk is used for virtual machine data ports, system management, virtual machine movement or migration (vMotion, in VMware terms) and any Ethernet-based SAN, such as iSCSI. By combining multiple physical ports into the largest

aggregation possible, the highest performance is achieved.

All of this consolidation adds up: 30 to 40 servers in a rack, multiplied by three to five ports per server, is 90 to 200 ports of LAN connectivity required per rack. Network managers could backhaul all of this to a large core switch if they wanted to run 200 Ethernet drops out of each rack to the core.

But most networking professionals are choosing to aggregate servers to switches within the rack to reduce cabling and complexity. These switches, called top-of-rack switches, are typically inexpensive 1U or 2U devices with 48 ports – a less expensive approach compared to the high cost per port of an end-of-row switch.

### Top-of-rack Switches

Network managers who have not used top-of-rack (or end-of-row) switches in their data centers should be careful in selecting devices because the top-of-rack environment is quite different from the typical end-user distribution environment found in a wiring closet deployment. Some manufacturers have different versions of their normal end-user 48-port 1U switch specifically designed for top-of-rack applications, with additional redundancy features and a higher uplink capacity built in.

The "Top-of-rack Switch Requirements" sidebar

## Top-of-rack Switch Requirements

Area	Data center switching requirement	Why is this important?
<b>Stackability</b>	<ul style="list-style-type: none"> <li>• Look for switches that can stack so that all switches in a rack are managed as a single unit.</li> <li>• A high-speed inter-switch link is also desirable.</li> </ul>	<ul style="list-style-type: none"> <li>• Network management is challenging enough without worrying about which switch a server is plugged into in a single rack.</li> </ul>
<b>Aggregation</b>	<ul style="list-style-type: none"> <li>• Switches should let the networking staff create large numbers of link aggregation groups, and the groups should be able to spread across multiple switches.</li> </ul>	<ul style="list-style-type: none"> <li>• Each server will need its own group, and the network manager will want to spread across switches to eliminate single points of failure.</li> </ul>
<b>Uplink</b>	<ul style="list-style-type: none"> <li>• A minimum of two 10-Gigabit Ethernet ports should be available for uplink. Four ports are better and may be needed in some environments.</li> </ul>	<ul style="list-style-type: none"> <li>• Virtualization pushes hardware much closer to its limits, which means that Gigabit ports from each server can get quite busy. Think in terms of port pairs for high availability.</li> </ul>
<b>Spanning Tree</b>	<ul style="list-style-type: none"> <li>• 802.1w Rapid Spanning Tree Protocol (RSTP) support is needed.</li> <li>• 802.1s Multiple Spanning Tree Protocol (MSTP) is an even better choice.</li> </ul>	<ul style="list-style-type: none"> <li>• Rapid Spanning Tree helps with high availability by reducing downtime when a switch goes offline. Multiple Spanning Tree can increase network utilization by spreading the load across redundant links, which is important when using iSCSI, VMware vMotion or disk-to-disk backups.</li> </ul>

shows some of the important requirements to look for when deploying top-of-rack switches into high-density virtualization environments.

### Data Center Virtualization

One of the side benefits of virtualization is a more efficient use of computing resources. In fact, a primary reason for virtualization's popularity lies in one of the main drives behind server proliferation – the difficulty of managing multiple applications on a single server – rather than a need for more computing resources.

For example, even in a small branch office, five or more physical servers can be found sharing the load of Windows Domain Controllers, Domain Name System (DNS) and Dynamic Host Configuration Protocol (DHCP), print and file services, and e-mail service – while the actual load presented by those services could have fit easily in the resources of a single system.

Take that simple example and expand it to a central data center, with a mix of production applications, application development and test environments, internal management systems, and infrastructure services. This kind of environment houses hundreds or even thousands of servers, overconfigured and underutilized.

### Increasing Utilization: The 10Gb Ethernet Upgrade

As virtualization is used to push multiple server loads onto a smaller number of servers, the utilization of each server will go up – a plus for the data center manager in every way. But for the network manager, normal metrics on port utilization used as rules of thumb for data center network design will change dramatically.

For example, consider a previrtualization data center, where statistics have shown that the typical usage of a 1Gbps Ethernet port per server is about 1 percent. Put 10 of those servers in a rack and the uplink speed from that rack could also be 1Gbps. This is because the network manager is expecting only 10 times that 1 percent, or a 10 percent load on the 1Gbps uplink port.

Consider swapping out those 10 servers for 10 virtualized servers, each with 10 VMs on them. Suddenly the average load is 100 percent on a 1Gbps uplink port – a major bottleneck, blocking peaks and causing disruption and unhappy application users throughout the organization.

In that situation, a 10Gbps (or more) bump is required. And a rack holding 10 servers with 10 VMs on them is very modest from a virtualization standpoint – many network designers will try to get five or 10 times as many virtual machines in the same space.

The obvious answer is that the organization's rack uplinks

need to be increased to 10Gbps at a minimum, and possibly higher. Keep in mind that reducing the number of servers does not reduce the amount of data generated. If there is a concern about total bandwidth out of a single rack, one option is to subdivide the rack of servers by virtually cutting it in half, creating two subracks, each with its own set of top-of-rack switches and 10Gbps uplink ports.

That's an expensive approach to avoid turning to the extremely expensive, fiber-only 40Gbps (or higher) speed interconnects that are becoming available.

### High Availability and Redundancy

Most data center networks have a high-availability core, yet individual servers end up with only a single connection to the network. Usually this is done for all the wrong reasons – the application architect didn't bother to specify redundant network connections; the network manager didn't insist on redundancy because it's always been done with a single connection before.

Virtualization represents an opportunity to rethink how servers are connected in the data center. Best practices call for redundancy in every connection, and that means identifying – and eliminating – single points of failure between the server and the existing redundant core. Server-to-network redundancy, done correctly, also provides the opportunity for increased performance by doubling or quadrupling total bandwidth from server to network.

### Rapid Spanning Tree Protocol

From a high-availability standpoint, it's sufficient to simply double every connection. Two ports from each server feed to different switches. Two connections from each switch (or switch stack) connect to different distribution layers. And then a redundant connection joins the data center distribution layer up to the core of the network.

Data center network managers should insist on 802.1w Rapid Spanning Tree Protocol (RSTP) support, at a minimum, to ensure the fastest recovery when there is a component failure.\*

Spanning Tree Protocol (STP) is used to eliminate loops in a bridged (Layer 2) Ethernet environment. If the organization's Ethernet network has any redundancy at all, it is almost certainly running 802.1D (with updates from 802.1t) STP.

RSTP is spanning tree that recovers from changes in topology much more quickly. When configuring 802.1w RSTP, be careful to check that all of the devices in the tree have support for this fairly new standard. It's unlikely that a device being sold today doesn't support it, but older devices may not. Mixing and matching is not advised because it dilutes the benefits of RSTP.



RSTP doesn't require much more configuration beyond normal 802.1D spanning tree. Network managers will get better results if they identify for their switches which are uplink ports and downlink ports to other switches, and which are simply talking to servers. (This is important for another reason: all spanning tree protocol features must be disabled on any server or end-user port for security purposes.)

### **Multiple Spanning Tree Protocol**

An even better strategy in the data center is 802.1s, Multiple Spanning Tree Protocol.\*\*


MSTP is backward-compatible with STP and RSTP. However, MSTP allows an organization to run multiple spanning trees at once across the same infrastructure. Why is it good to run different spanning trees on the same topology? This allows network managers to run different virtual LANs (VLANs) across different parts of the network.

For example, imagine that the organization has a fully redundant set of links from edge to distribution to core. With normal STP, half of the links (and probably half of the switches) are unused and don't pass any traffic – they are in place for a failure that might never come.

With the multiple spanning trees allowed in MSTP, network managers can run some VLANs over one data path and through one set of switches and links, and other VLANs over a different path through different switches and across different links.

It's not true active-active load balancing, and it's not dynamic routing. But MSTP lets network managers spread the load out in a static way that can provide better accommodation for peaks in traffic.

If the organization is using some VLANs for iSCSI SAN-based disk service, virtualization management



or backup-to-disk applications, then pushing that data across another path is a great way to make sure that iSCSI (or VM migration or backup-to-disk) gets as much bandwidth as it needs, without impacting the rest of the network.

The downside of MSTP is that configuration is considerably more complex. So it really pays to use MSTP only when the organization's network utilization peaks frequently to full saturation of critical links.

### **Cloud Applications**

In some organizations, data center networking is changing in an entirely different way. Applications are being moved to the cloud – out of the data center, across the Internet and running on someone else's infrastructure. While cloud service providers are quick to claim that pushing applications into the cloud doesn't require changes in the rest of the network, that's an incorrect view.

Cloud computing simply moves applications from one part of the network to another, but the same security and management concerns have to be met. From the network point of view, there are five main areas to look at when applications move out of

\*\*802.1s was folded into the 802.1Q virtual LAN standard in 2005, so switch literature may refer to MSTP as 802.1Q-2005.

the data center and across the Internet: bandwidth management, encryption policy, access controls, audit and compliance, and authentication and authorization.

### **Bandwidth Management**

The most obvious effect that cloud computing will have on the organization is a bandwidth-use jump. If people are reading their e-mail via the Internet, for example, then it's important to have enough bandwidth to the Internet from the organization to support that kind of use.

Measuring and provisioning Internet bandwidth isn't as easy when the organization is running applications from the cloud, because now the quality of the Internet

---

## **APPLICATION DELIVERY CONTROLLERS**

Load balancers have been one of the most successful tools for network and application managers in their quest to add redundancy and reliability to network services. But other services such as Domain Name System (DNS) are also often protected by load balancers. Over time, the load-balancer marketplace has added so much functionality to these common appliances that a new name for them was coined. Thus, the application delivery controller (ADC) was born.

What's the difference between a load balancer and an ADC? It's all a question of added features. ADCs take the basic functions of load balancing (switching traffic at layers 4 through 7) and add SSL acceleration, compression and caching, protocol optimization, and connection multiplexing. Together, these devices give the application user a better experience by making a web service seem faster and more reliable. They're ideal for virtualized environments as well.

---

connection is critical to satisfaction with cloud applications.

Connections that were good enough for simple web browsing may not be good enough when users are accessing applications over the Internet that they used to connect to over the LAN. So before contacting the organization's provider and asking for an increase in the connection speed, make sure to develop some metrics for throughput, latency to critical sites (such as the cloud service provider) and availability.

A service-level agreement (SLA) that includes these statistics should be built into the organization's Internet contracts. It won't necessarily help the organization get a better connection. But if the ISP can't meet the organization's requirements, the SLA is a good tool to get the organization out of that contract and over to a better provider quickly.

Network managers should also look at devices in the critical path between end users and the Internet, such as firewalls, intrusion prevention systems (IPSs) and even switches. Some network managers might find it surprising how often an old 100Mbps switch is used as the transit network between the Internet and an organization's firewalls.

That's fine, until the organization increases its connection speed past 100Mbps, at which point the switch would be throttling the network unnecessarily. Firewalls and IPS devices may also need to be upgraded to handle the increased bandwidth load to the Internet.

### **Encryption Policy**

Traffic in the organization's data center probably moves unencrypted, but that approach can be very risky while running applications over the cloud. Moving to the cloud requires a new encryption policy.

Start by enforcing communication encryption on both end-user PCs and within the network infrastructure – something network managers didn't have to do before cloud-based applications. Another important part of getting the organization's network ready for the cloud is making sure that its firewalls and web proxy systems are set so that no one can accidentally make an unencrypted connection.

With this new increase in encryption comes another problem: IPSs, data loss prevention (DLP) solutions and intrusion detection systems (IDSs) won't be able to handle encrypted traffic without special help. This may push organizations to install special Secure Sockets Layer (SSL) decryption appliances that can be used to let these security monitoring and enforcement points do their job.

If the organization was doing packet capture

for forensics purposes, the forensics team needs to know that the IT staff is going to be capturing a great deal of useless data once end-to-end encryption is turned on. This will alert them either to add their own SSL decryption tools or to filter out the encrypted traffic that won't be very useful.

### **Access Controls**

Many organizational network and security teams are a little sloppy with their access controls, depending on known IP addresses to help define permissions within the network. That works well as long as the organization is in total control of all IP addresses within the network.

However, as soon as the organization pushes applications off-network into the cloud, it loses whatever control it had on IP addresses and any ability to use them to define security permissions.

A better access control system for the cloud era is moving from an IP-centric approach to a user-centric approach. It's not a question of what IP addresses the organization happens to have today, but who is sitting behind that keyboard. Network Access Control (NAC) can help organizations move to a user-centric set of access controls, and may need to be part of a cloud deployment.

NAC brings other benefits beyond facilitating cloud deployments. With a stronger interest in internal access controls, NAC gives the network manager and security manager the tools they need to build a scalable internal access control system that will also stand up to the bright glare of regulatory scrutiny.

NAC may be a larger-than-average deployment struggle, as it touches all aspects of the network (and end-user systems to boot). But the benefits in fine-grained access control are invaluable and impossible to achieve any other way.

### **Authentication and Authorization**

Cloud applications should be integrated with the organization's authentication and authorization system. Some poorly designed cloud applications don't do this well. However, the enterprise-class ones will talk to the organization's existing authentication system (often Windows Active Directory) to validate user credentials and authorizations.

When authentication didn't stretch beyond the firewall, most Windows system managers had little knowledge or interest in how external authentication worked and was secured. Now that the cloud is stretching into the organizational network, it's up to the network and security teams to ensure that proper, secure

connections are made to enterprise directories.

The opportunity to abuse the organization's authentication system is an additional concern. Once a cloud-based application accepts a username and password that will be checked against the organization's directory, an attacker has the leverage he or she needs to begin a brute-force authentication attack against the organization, or to attempt some clever social engineering.

In either case, it's the responsibility of the network and security teams to ensure that solid password management is properly communicated to the user community and is enforced by technology on the server side, along with end-to-end break-in detection to block brute-force attempts.

### **Audit and Compliance**

No matter what regulatory regime the organization is covered by, there is going to be a requirement to coordinate auditing information from both the data center and the Internet-based cloud – assuming that the cloud service provider will be sending the organization logging information.

If the organization hasn't taken a serious look at log management and built (or bought) a system to capture and archive all of its logs, it needs to with the cloud. Keep in mind that the cloud-based applications are only going to make things more challenging by bringing in new log information that may be difficult to interpret. Log management systems range in sophistication from the very simple to more complex tools with searching, alerting and archiving capabilities.

At the very high end of the log management chain are security information management (SIM) products, also known as security event management (SEM) and security information and event management (SIEM) solutions. These devices act to correlate events across the network (and the cloud) to bring a deeper understanding of what is happening.

These products sound as if they are security-only devices, but that's deceptive. These tools bring value to network, security, operating system and application managers in organizations of all sizes. ■

# NETWORK OPTIMIZATION

## Fine-tuning Enterprise Networks

Maintaining a top-performing network requires three important qualities: good design, constant visibility and flexibility to accommodate changing needs. A good network manager will build in a combination of processes, configurations and technologies to ensure that the network keeps operating reliably and at top speed.

### **Building Resiliency**

Constant visibility requires using network management and monitoring tools to always know what's happening on the network. Start with simple reachability, or the ability to contact a remote server or device using a command line tool such as the ping utility. Second, make service health checks, which keep track of whether the expected services on a server are running and responding with the correct information.

Finally, identify link-loading levels and top talkers/listeners on all points of the network. This is even more important if the organization has a WAN component because WAN congestion is hard to

diagnose without the right tools.

Networks designed to be flexible can accommodate changing requirements without project delays or expensive upgrades. Network managers should always be one step ahead of changing requirements. For example, if the IT department is asked to install 200 servers within a few days, it can be done without numerous last-minute changes.

### **Network Design**

Most networks follow a traditional tiered design, whether in the data center or in buildings where end users are served. The names of the tiers may change over time, but the basic approach is the same: Run connections out to several Ethernet ports where users and servers are located. Then gradually aggregate to higher-speed links until reaching a high-speed central backbone where switching and routing both take place.

Network managers typically use three basic terms:

- **Edge:** Where the users and servers are located

- **Distribution:** The aggregation layer that builds up toward the backbone
- **Core:** The backbone where routing occurs

These terms may be unfamiliar to security managers in this context. For them, the edge of the network is where the network stops and the firewall begins – so be careful to define this terminology when appropriate. There's no special magic about the layers or about the way they are laid out. However, these terms are so ingrained into training classes and with equipment manufacturers that network managers may have a difficult time convincing others to use different terminology.

### *Optimizing Networks*


The way to design better LANs is to think seriously about making the network resilient. The best way to do that is to assess every piece of equipment and every patch cable and determine what would happen if a device or cable were to fail. Consider some examples.

Starting at the edge of the network with the user device, there's a clear single point of failure – the user's workstation and the cable to the edge switch. But what about the edge switch? When it fails, how long is it going to take to find the switch, configure a replacement and get it into place?

Some network managers use large-chassis switches with individual blades of 24 or 48 ports. That saves a lot of time, assuming three things. First, that there's a blade on the shelf ready to drop into that switch. Second, that the cables are labeled so that when 48 ports are disconnected the IT staff can reconnect them in the right order. Third, that the IT staff can physically get the old blade out of the switch without disconnecting every cable going to every other blade.

Buying spare parts is easy, but making sure that the organization's cable management allows the group to swap out a blade and reconnect cables without disrupting operations everywhere is a more difficult task. Designing a LAN with this in mind will ensure a more reliable network.

Chassis are more expensive per port than individual switches, plus they represent a single point of failure. Many network managers prefer the scalability and flexibility of using stacks



## THE NEED TO OPTIMIZE

Why don't most organizations develop capacity plans? Because it's been easier to overengineer the network than to figure out exactly what is needed. Deploying Gigabit Ethernet for everyone, for example, takes less time than figuring out who needs it and then managing a multispeed network.

Unfortunately, the networking group's ability to overengineer the network is creating a real challenge. Network managers are discovering that they need to gain a deeper understanding of what is happening on their networks to optimize them for best performance.

Gaining visibility into network flows is the critical first step. Once network managers know what's happening, more expensive techniques, such as upgrading some links to 10Gbps or installing acceleration appliances on WAN links, can be used to balance performance increases with budget.

of 1U switches rather than a chassis.

Network managers still must have a spare switch on the shelf, good cable management and, most important, a valid copy of the configuration for the failed switch. Capturing copies of the configurations on all devices every night means avoiding having to rebuild a network from scratch.

It's also important to make estimates of how often devices fail, how many users the failure will affect, how long it will take the network to recover, and how much it will cost for the IT staff to mitigate the potential failure. Such estimates will give the organization the information it needs to build in reliability – and will free up money and time for focusing on other critical parts of the network.

### ***Building Redundant Networks***

Move up one level to the distribution switches and consider the implications of a failure here. Most networks with more than 200 devices have a distribution layer. What happens when any of the switches fail?

Network managers need to have hot spares and online redundancy. Each edge switch should connect to two different distribution switches. And the distribution switches themselves should have redundant routes to the core.

Many network managers quickly realize that redundancy means doubling up on equipment. In a traditional spanning tree configuration, the budget-minded network manager may not think it's necessary to buy two distribution layers and two cores in case a device breaks. That's a great reason to learn about 802.1s, the Multiple Spanning Tree Protocol.

MSTP is covered in more detail in Chapter 2, but to recap: 802.1s lets network managers actually use the redundant paths in the network. The configuration is manual, offering the ability to decide which VLANs run on which sets of gear. On the other

hand, 802.1s can be used to increase performance at the same time that it increases reliability.

If the decision is made not to double up the organization's distribution layer, at least double up on the uplink connections. Network equipment made by reputable manufacturers is generally very reliable, but components can fail, especially fans and power supplies.

One of the least reliable components in any machine room or wiring closet is the cabling. It doesn't matter how much an IT department spends on its network patch cables, some fail every year. Patch cables fail more frequently as they age, as the cable insulation becomes more brittle and environmental factors (such as blowing cold air) take their toll.

If the uplink between edge and distribution or distribution and core is 1Gbps Ethernet, then doubling uplinks is an inexpensive way to achieve redundancy and higher performance. There's no reason not to double up those connections.

If the organization jumped to 10Gbps in wiring closets, then doubling those connections is harder to justify, because 10Gbps Ethernet modules are still fairly expensive. In that case, consider specifying switches that are capable of supporting double uplinks.

But only populate one uplink during the initial installation. Have a firm plan in place to revisit each switch, add the second 10Gbps module or transceiver and light up a second uplink in 18 to 24 months when costs come down.

### ***Creating Visibility***

Visibility is about seeing and understanding what is going on with the network. Consider it a critical ability that's needed both to verify network design and to identify bottlenecks and problem spots before they begin to affect operations.

### ***Network Management Tools***

The keys to gaining good visibility are proper tools and instrumentation. Remember that this is a common challenge: Every network manager is looking to increase their visibility and control, and this need has inspired a rich set of tools and best practices.

Many software products have tried to take on the entire spectrum of network management, but the different requirements for each area make that an impossible task. Instead, set up a portfolio of network management tools that the organization can use to gain the visibility it needs.

For example, a great reachability/availability and alerting tool requires a very different design and features

## Networking Tools Checklist

Strategic tool	Function	Why is this important?
<b>Inventory</b>	This type of tool finds all devices, routers, switches and systems that are on the network.	Knowing what's on the network is the first step to visibility.
<b>Reachability and alerting</b>	These tools constantly check whether critical systems are reachable, applications are responding and metrics for network devices are within limits.	Proactive monitoring and alerting when devices, systems and applications stop responding is a minimum requirement.
<b>Performance details</b>	This category gathers performance statistics from network elements, mostly switches and routers, using Simple Network Management Protocol (SNMP) or similar protocols.	Long-term statistics on traffic down to the port level are useful when replacing or upgrading equipment or debugging a problem.
<b>Traffic summary</b>	Traffic tools watch network traffic at critical points and generate summary information, usually at the IP layer and up, toward the application layer.	Broad views of traffic across the network (top talkers and listeners, protocols, applications and conversations) will help identify problems and misconfigured systems, and plan topology changes.
<b>Protocol analysis</b>	These tools capture traffic for replay, display or export.	This broad set of tools is most useful when debugging problems or as part of a design session in which IT staff try to understand where traffic is going and why.
<b>Intrusion detection and prevention/data loss prevention</b>	IDS/IPS/DLP tools watch application-layer traffic as it flows past key points in the network (LAN to servers, WAN, wireless, Internet), looking for known attacks and anomalies.	The overlap between network and security management makes these tools, typically part of the security manager's toolset, important to the network manager as well.
<b>Configuration control</b>	This category of tools polls network elements to capture configuration changes and push or roll back configurations.	Tracking network configurations and backing up configuration changes can protect organizations when devices fail.
<b>IP address management</b>	These tools handle IP address assignment and recovery, DNS synchronization and Dynamic Host Configuration Protocol (DHCP) service.	By tracking a network issue back to a known IP or media access control (MAC) address, IT staff can identify the person or device involved and resolve the problem.

compared to a tool to identify the top talkers and listeners on the network. Most organizations probably want both abilities, which means two different tools to accomplish those different tasks.

### **Designing Visibility**

A second important strategy for gaining network visibility is recognizing that network management isn't something to layer on top of the network after it's completed. Instead, network managers must design the network knowing that the organization needs visibility. One example is traffic mirroring (called Switched Port Analyzer, or SPAN, by Cisco), which offers traffic capture capabilities for statistical analysis, protocol analysis or intrusion detection.

The more highly switched a network, the more difficult it is to find places to mirror traffic. It may be necessary to deoptimize the network by inserting additional interswitch connections with different VLANs, or even inter-VLAN connections within a single switch. This ensures that the organization can tap traffic for analysis at strategic points, especially heading into the data center and out toward wireless networks, the WAN and the Internet.

There are several formal network management models available and many products to help IT staffs manage the network and gain much-needed visibility. The "Networking Tools Checklist" sidebar on page 24 offers a checklist of the most useful management and visibility tools available for large networks. Use it as a starting point for developing a network management strategy.

### **Increasing Flexibility**

Being flexible allows the network to adjust to changing user and organizational needs. Here are some best practices that will help network managers build more flexible networks.

**Know the network.** This is very nontechnical advice, but it's also the most important. The key in networking is to be slightly ahead of the power curve. Network managers have to keep their antennas tuned in two directions: industry trends and organizational needs.

Why? Because if network managers know what's coming before it impacts their network, they can think strategically about what its repercussions will be and have time to prepare and do a good job of integration and technology adoption.

**Test new technologies.** Become informed about technologies such as VLANs, Quality of Service (QoS), Power over Ethernet (PoE), Fibre Channel over Ethernet (FCoE), iSCSI and wireless. Don't just read about them – fight for some budget dollars to spend on future technologies, get some hardware or software and actually do some testing.

Even the most budget-strained organization knows that it needs to invest in technology to stay competitive. By showing that it is investing in understanding new technologies, and by providing test results to management, the IT department will find it easier to secure the time and money to move forward on projects.

Don't just read a web page or white paper – push to establish a test lab. The network manager's formal job description may not include research and development, but putting aside some time every month to try new technology in a lab is essential.

**Be open to new ideas and continue to pursue your own.** When someone comes to the IT department with an unworkable idea or a plan, don't reject it out of hand. Instead, explain the obstacles that have to be overcome.

It may be security, it may be resources or it may be a policy that has to be changed. Let whoever is making the request understand the hurdles in the way – and while they are working on them, begin planning for what the IT staff will have to do when those constraints disappear.

At the same time, when someone rejects one of the IT department's ideas, ask for information on why the project wasn't approved. It's not being impertinent – it's a reasonable question to ask. Turning the situation around may be as easy as asking for some extra time to make your point. ■



# WIRELESS MOBILITY

## Taking Advantage of Mobility

The most frequently misunderstood technology shift in the past 10 years has been the move to mobile computing. Any network manager who does not see the potential of mobile computing is missing an opportunity to make the organization's staff more productive — both in and out of the office.

One of the core components of mobile computing is wireless network access. In the field, there are competing standards and models, ranging from IEEE 802.16 WiMax to a host of mobile telephony technology families, such as 3G and 4G. However, in the office and in many controlled environments such as indoor public spaces, wireless access is gained exclusively through one standard: IEEE 802.11.

Therefore, it's important to understand how to deploy enterprise-class 802.11 networks and how to make intelligent management decisions about 802.11 availability, security and performance.

### **Wireless Standards**

IEEE 802.11 was first released in 1997 and has since undergone multiple updates. Initially, most 802.11 hardware had a top speed of 2 megabits per second, and security specifications took up a dozen pages of the standard. With the latest version, 802.11n, the top theoretical station speed is now 600Mbps, and the security specifications are 207 pages long.

What's most important to know is that 802.11n is more reliable, faster and supports more users than previous versions of 802.11.

### **MIMO Technology**

Multiple Input Multiple Output (MIMO) is a new capability in 802.11n that uses multiple streams of data and sophisticated signal-processing techniques to process more bits through the same amount of space.

Before delving into a detailed explanation of MIMO, it's important to understand that not all 802.11n devices share the same capabilities.

There are physical limitations, such as the number of antennas, their spacing and the amount of power required, that may keep some smaller devices from realizing huge performance gains on an 802.11n network.

The "multiple" aspect of MIMO is that there must be multiple antennas to get multiple data streams going. On most 802.11n products, access points are described as 2x3 or 3x3.

The notation  $a \times b$  (more properly,  $a \times b:c$ ) is used to show how many transmitting antennas (the  $a$  number) and how many receiving antennas (the  $b$  number) are used by the device. The  $c$  number, which doesn't show up often on specification sheets, is actually the number of data streams that are sent or received.

Simply adding antennas provides better immunity to noise and interference and gives users a better experience. Adding data streams actually speeds data processing, delivering better throughput. So a device advertised as 3x3 will generally offer users a better experience than a 2x2 device.

The speed will be higher only if there are more data streams, something that's difficult to determine from the specification sheet. Both 3x3:2 and 3x3:3 devices are common. The 3x3:2 device is limited to about 130Mbps performance in normal channels, while the 3x3:3 device can reach speeds up to 195Mbps in normal channels.

For planning purposes, consider 130Mbps the maximum end-user speed that will be delivered for the next five years. The theoretical maximum 802.11n performance speed of 600Mbps may never be available to normal end users because it requires a 4x4:4 configuration and doublewide 40-megahertz channels. That's more antennas, more radio power consumption and more CPU requirements than a typical notebook can meet today.

### **Benefits from 802.11n**

Network managers need to look beyond the 802.11n label on the box and consider what they need to get the real benefits from these new devices. An 802.11n device with only a single data stream – even if it is a 3x3 device – won't be much faster than a typical 802.11g device.

At a minimum, in-building access points should be 2x2:2 with dual radios (2.4GHz and 5GHz bands, usually called the "b" or "b/g" and "a" bands) to get the real advantages of 802.11n. Dual data streams and 2x2 antennas deliver the high bandwidth of 802.11n.

And dual radios will let network managers pack more users into a smaller space by making use of more radio-frequency spectrum. The 2x3:2 or 3x3:2 devices also offer end users a better overall experience.

When buying 802.11n equipment, focus on the specifications. It's not possible to tell how many antennas or data streams a device has by looking at it. Unless the building (or a specific room) has a particularly challenging radio frequency (RF) environment, specify access points without visible antennas to reduce costs and increase reliability.

Remember that for end users to see any benefit, the 802.11n clients will need to support multiple data streams as well. For the moment, most notebooks are limited to a maximum of two data streams.

This means that buying 3x3:3 devices may help future-proof a network, but they will not provide end users with better performance than 3x3:2 devices because current notebooks are able to handle two data streams at most. If the manufacturer offers both, specify 3x3:3 only if the price difference is negligible.

### **40MHz Channels**

Another enhancement of 802.11n is the use of 40MHz wide channels, which doubles the 20MHz width of pre-802.11n wireless channels. By taking two adjacent channels and using them as a single, larger channel, 802.11n can double the per-user theoretical maximum speed.

The challenge is that there aren't that many channels available. Reducing the number of channels by using two for every access point reduces the number of users who can work simultaneously at full speed on the wireless network.

There aren't many 802.11n enterprise networks deployed today, so there is little consensus on how best to configure wide channels. Most networks today make use of 20MHz channels in the 2.4GHz band, sometimes called the 802.11b/g band.



Because there are only three nonoverlapping 20MHz channels, there's only one 40MHz channel available.\* When laying out a building network with multiple access points, keep the 2.4GHz band populated as 20MHz channels 1, 6 and 11 because access points and clients will interfere with each other if they are all crowded into the same single 40MHz channel.

Besides the shortage of channels, there are other reasons to stay away from 40MHz in the 2.4GHz band. 802.11n has a variety of protection mechanisms, including an aversion to any non-802.11n devices on the network. If an older 20MHz device appears in operation, then 802.11n devices will fall back to 20MHz channels.

The 5GHz band, sometimes referred to as the 802.11a band, has a much larger allocation of channels. The channel count varies because additional channels were added to the 5GHz band in 2007. Only devices that have Dynamic Frequency

---

## TURN THE POWER DOWN

One of the counterintuitive aspects of wireless LANs is that they perform better when power levels are turned down, rather than up. In fact, when deploying production-quality wireless service, turning the power up is not recommended.

Remember that the wireless signal will bounce back from radio frequency (RF) reflective surfaces, such as metal doors and cabinetry, solid exterior walls, and even the building across the parking lot. The stronger the signal, the more it bounces.

The more it bounces, the worse the performance of the network, because individual devices will see the same signal multiple times, something called multipath interference. The more interference, the lower the network throughput.

One way to understand this is to think of the wireless signal as something one person is throwing to another. The goal is to throw the object so the receiving person gets it gently and directly.

In wireless networking, the sender wants to put exactly as much energy into the signal as it takes to get it there, and not have the signal bounce around the room like a rubber ball.

In addition, even if a stronger signal is available, it may not be desirable. Think of each access point as a meeting room in which only one person at a time is allowed to speak.

Wireless is like Ethernet before switches, when only hubs were available and only one LAN station at a time could talk. If there is one meeting room (access point) with 30 people in it and the meeting lasts 30 minutes, then each person is allocated only a very small amount of time to talk. That slows everyone down.

By breaking things up so that there are two smaller meeting rooms (two low-power access points) with only 15 people each, more people can talk at once.

---

Selection (DFS) support (required after 2007 to avoid interference with radar) are allowed to use these new channels.

In the United States, the result is that if the organization has newer equipment, it can usually allocate up to 11 802.11n 40MHz nonoverlapping channels, which is plenty of room for high-speed devices to operate without colliding into each other.

Network managers still have to worry about older gear operating in the 5GHz band and reducing overall performance by causing 802.11n protection and noninterference features to kick in. Because there are more channels to work with, and network design should call for fewer devices per access point, these effects can be minimized.

### **Enterprise-class 802.11 Networks**

With the increased performance and capabilities of these newest additions to 802.11, most network managers now view wireless as a production network, rather than as a convenience for staff and guests. To elevate 802.11 to full production status, consider the following three tips.

**Find a good consultant.** Designing a wireless deployment for a large building or a campus requires a specific set of RF engineering skills that the typical network manager may not have. This is where it makes sense to bring in a consultant who has these skills.

Remember that above all, building an optimized, high-performance, highly reliable wireless network requires proper access-point placement. The IT department needs to hand over a set of requirements and building maps showing offices, meeting rooms and dead space.

The consultant's job is to hand back a set of maps showing where to put the access points and what the coverage will be throughout the campus or buildings. The requirements turned over should be short and to the point.

**Use a smart wireless-management system.** Every major wireless manufacturer now offers centrally managed and coordinated wireless networks. In some cases, this means that where there is a wireless controller device the access points are tightly coupled and all traffic flows through the controller.

In other cases, the binding is more relaxed and there is a global management system controlling and tuning the access points, but traffic doesn't flow through a central point. In both situations, the central management system is responsible for RF management. This is a requirement for any large enterprise wireless network.

Remember that even if the building looks the same on the outside, it is constantly changing on the inside. People walk around. Desks, filing cabinets and bookshelves are moved. And someone's metal office door that was closed in the summer is now open in the winter.

The wireless network has to be tuned to handle these changes in the RF environment. The only way to effectively keep performance at a high level is via the constant monitoring and adjusting that a central management system offers.

Centralized management and coordination also lets network managers quickly detect access-point failure and automatically adjust power levels and coverage of other access points in the area to help eliminate dead spots and keep the network reliable.

**Use multiple Service Set Identifiers.** The SSID refers to the name of the network. Most organizations need multiple SSIDs because they have networks that offer different services.

Don't be afraid to propagate two or four SSIDs if that's what it takes to deliver effective wireless services. For example, many organizations use

a guest SSID for unauthenticated or lightly authenticated guest users. The IT group can offer the convenience of allowing third parties to browse the Internet, yet not provide any special access to the organization's network.

Network managers should also have a production SSID that gives staff members network access similar to what they'd have on a wired connection. Of course, the security profiles of these two SSIDs are dramatically different.

Guest users would expect almost no wireless security (although a firewall is expected). Production networks typically have Wi-Fi Protected Access version 2 (WPA2), authenticating users to the organization's central directory (such as Active Directory) and ensuring encryption of all traffic.

Some organizations may choose to have even more SSIDs. For example, many set up a special network for wireless devices that give staff members easily authenticated and encrypted access to the Internet and a small selection of internal resources without letting them into the corporate network.

The staff wireless network in this case should remove traditional wireless captive portals and web-based authentication that are typically found on guest networks. ■

# NETWORK SECURITY

## Securing the Enterprise

The technology of network security has consistently improved over the past several years. IPSs get better; firewalls get faster; unified threat management (UTM) devices get more sophisticated; and DLP gets more precise. Despite the upgrades, the fundamental approach remains much the same.

### Network Security Concerns

What is different about network security compared with five years ago is the environment. And here there are three significant changes to consider when designing security.

**Borderless networks:** Borderless networks help alleviate some of the complications around the numerous amounts and types of devices now accessing the network. This approach helps support an infrastructure capable of delivering seamless and secure access to users regardless of their physical location.

Clearly, a single firewall device at the Internet firewall cannot support all the requirements for remote access, branch

office WANs, partner virtual private networks (VPNs) and different security levels within the same organization.

**Pervasive computing:** Driven by the low cost of both personal computers and smartphones, as well as high-speed wireless communications, pervasive computing is the new reality – one in which both staff and other end users all have mobile devices that are always on. And all of these users want to be continuously connected, whether via an in-building wireless LAN or worldwide telecommunication networks.

**Cloud computing:** This untethered and flexible approach to computing often involves having an organization's data received, processed and stored in data centers and on servers entirely managed by third parties.\*

These environmental changes have a clear effect on any data network, even one not directly connected to the Internet.

All three trends point toward changes we will likely see more of in the future: cheaper devices that

interconnect, operate autonomously and constantly update large cloud databases with new information. This may include raw data points (for example, the engine temperature in a car), geographic information (a cell phone's location) and image data (foot traffic through a parking lot).

Building a 21st century network security strategy requires considering carefully the effects of these three environmental changes.

### **Network Access Control**

The single most important network security response to this new environment is network access control. With NAC, the network perimeter has shrunk to the point where every Ethernet port is a potential access control device – in effect, a firewall at the physical port level.

The strategy behind NAC calls for user-focused access control. This means that a user's access is a function of who they are and which groups they belong to in the organization's directory, as well as the state of any endpoint security software they use.

Although NAC is a good idea, many organizations have been slow to adopt it in their networks over concerns of additional cost and complexity. However, it is simple for most network managers to include some valuable NAC concepts in their network without going through a full deployment.

Mobile computing is a good example of how NAC can add security without adding complexity. One of the security complications of pervasive computing is the state of the endpoint. In other words, is that notebook, smartphone or tablet computer properly secured?

In a full NAC deployment, each device would be authenticated and checked before appropriate access is granted, a difficult task given the large number of mobile platforms. However, it is easy for network managers to create a wireless network

to bridge the gap between no access and unbarred access, specifically for local mobile devices.

This avoids the difficult problem of checking endpoint security and adds an internal access control point, a much easier task. By adding a separate SSID for staff-owned wireless devices that requires authentication to the enterprise directory (such as an Active Directory service), the network manager can create an intermediate zone that allows access to appropriate services, such as e-mail synchronization and limited intranet access.

Building a separate wireless network that is somewhere between full access and guest Internet access has other desirable effects. For example, it encourages a strong awareness of different security zones within data centers, possibly enforcing a separation of servers based on the sensitivity of the data they contain.

This is how it should be done originally. However, the organic growth of most data centers often creates a hodgepodge of services and servers that follows mostly chronological and budget lines, rather than functional and security lines. If access controls are differentiated within the network, this may result in further reorganization that helps network security overall.

### **Remote Access via VPNs**

Remote access via VPN concentrators is another security strategy with a simple deployment and great benefits. With NAC, every remote user has an approved endpoint security product on their notebook or smartphone. The VPN concentrator is responsible for authenticating users, checking endpoint security status and enforcing specific access controls.

Without going to full NAC, how can an organization reap some of these benefits? One way is by shifting to updated Internet protocol security (IPsec) technology. Another option is an SSL VPN.

IPsec VPN technology is older than SSL VPN and has a less sophisticated set of security controls. There's nothing wrong with IPsec VPN technology. In fact, it is more resistant than SSL VPN to common attack methods such as eavesdropping and man-in-the-middle attacks.

However, SSL VPN devices were introduced after IPsec and generally have a more sophisticated access control model that is easier to deploy. Enterprise-class SSL VPN devices make it easier to give remote users exactly the access they need, and no more.

So when moving to a more modern VPN deployment, it makes sense to increase the organization's ability to define and apply fine-grained access control. This offers the benefits of NAC, without having to commit to a full NAC deployment.

### IPS and DLP

Intrusion prevention systems have been underused in many networks where they could be providing valuable benefits. Many network managers have hesitated to deploy IPS units, considering the technology too new and unreliable.

However, IPS devices are effective at identifying known threats and blocking them, which means that they are good for identifying systems that are infected with malware or are participating in a botnet attack. Placing an IPS between the organization's end users and the most valuable assets in the data center will help identify the workstations that have become infected and need to be cleaned up.

If having an IPS between end users and servers is cause for concern, then put the IPS on the organization's outgoing Internet connection (always on the inside of the firewall). If workstations become infected, they'll trigger the IPS on outbound traffic almost as quickly. This strategy won't necessarily protect the servers, but it may provide an early warning that a threat is present.

IPS units can also be used on WANs, assisting with remote access and wireless links. They can help quickly identify and block systems that are misbehaving, either because malware is present or a malicious user has stumbled onto the organization's VPN.



### Data Loss Prevention

There's a clear parallel between IPS devices (which decode network traffic, look for known attacks and block them from entering the network) and DLP devices (which decode network traffic, look for known data and block it from leaving the network). DLP products have very different design criteria, but the basic concept is almost identical to an IPS.

Because of the steep cost involved in cleaning up a data breach, DLP is not a difficult investment to justify. Not all DLP technology requires an in-line network device. DLP can also be deployed on top of many web proxies, firewalls and e-mail security appliances. ■

## IS SSL VPN ACTUALLY NAC?

Network access control solutions and Secure Sockets Layer virtual private network (SSL VPN) devices have similar evaluation criteria. Is an SSL VPN product the same as a NAC device? Some manufacturers think so. For example, Juniper Networks repurposed the policy engine from its SSL VPN product into its NAC product.

Since the best SSL VPN products focus on authorizing users, controlling access and using endpoint posture assessment, this is a natural fit. Just as intrusion detection system manufacturers were ideally situated to start selling IPS products, SSL VPN manufacturers are going to have a natural advantage in the NAC market.

If you apply NAC evaluation criteria to SSL VPN devices, the similarity is even clearer. For example, good NAC solutions include a broad range of authentication methods and types, just as good SSL VPNs do: web-based authentication, installed or downloadable clients and full integration with Active Directory. NAC also generally has endpoint security assessment, a common feature of all enterprise SSL VPN products.

And NAC always includes access control enforcement, something most SSL VPNs excel at. For example, most SSL VPNs can provide access control at the URL level, while NAC solutions often use a much coarser access control model based on virtual LANS (VLANs) or subnets.

The one area where SSL VPNs don't compare well with NAC is with LAN-based access controls. Because SSL VPNs are designed for remote access, they don't really fit in a LAN-based NAC deployment. However, SSL VPNs are a good entry point for network managers looking to get some experience with NAC in securing remote access, branches or wireless.

This glossary serves as a quick reference to some of the essential terms touched on in this guide. Please note that acronyms are commonly used in the IT field and that variations exist.

# GLOSSARY

## **Active-active**

This term refers to a high-availability technique in which multiple systems operate simultaneously to provide services. Two firewalls in an active-active configuration would share the load of traffic, with one firewall handling half of the traffic and the other firewall handling the other half.

## **Active-passive**

This term refers to a high-availability technique in which two systems provide service, but only one system delivers service at a time. If two firewalls are in active-passive mode, one firewall handles all the traffic while the other watches and maintains state information so it can take over when the first firewall fails.

## **Application delivery controller (ADC)**

ADC is the new marketing term for load balancer. An ADC takes the basic functions of load balancing and adds SSL acceleration, compression and caching, protocol optimization and connection multiplexing.

## **Collapsed backbone**

A collapsed backbone is a network architecture in which the network

backbone is contained in a single device, offering simplified management and higher performance. This is in contrast to a distributed backbone where routing and core switching functions are spread across multiple devices usually located in various hubs and wiring closets.

## **Data deduplication**

Data dedup is a type of data compression in which duplicate data chunks, ranging in size from 4-kilobyte blocks up to entire files, are replaced with a single copy of the data and pointers to the copy. Deduplication can be done as part of a process, such as when making backups or storing files to disk.

## **Data loss prevention (DLP)**

DLP refers to a family of security products aimed at mitigating the threat of sensitive or critical data being taken outside of organizational control. DLP products help protect against intentional/malicious and unintentional leakage of sensitive information.

## **Dynamic Host Configuration Protocol (DHCP)**

DHCP is a critical service in networks used to assign IP addresses and other

network configuration information to otherwise unconfigured devices. DHCP services need to be carefully engineered for security purposes and to ensure continuity of operations.

## **Edge, distribution, core**

These are common terms that describe layers in an organization's building or campus LAN. Edge ports connect to users at the edge of the network. Edge switches are aggregated through distribution layer switches, such as in a wiring closet or at the building level. Sometimes the distribution layer will involve Layer 3 routing. Distribution switches are then connected to a core switch. Most core switches do some Layer 3 routing.

## **Fibre Channel over Ethernet (FCoE)**

FCoE uses Ethernet to transmit Fibre Channel SAN traffic. A special protocol is required because the physical layer for Fibre Channel has different characteristics than Ethernet. FCoE compensates for the differences. This traffic moves across the link layer and uses Ethernet to transmit the Fibre Channel protocol.



### **Lights out**

"Lights out" is a style of data center management that doesn't require anyone to actually go into the machine room. Lights-out management is supported by remote operators who have full access to device consoles over a network.

### **Network access control (NAC)**

NAC embeds access controls into network devices. These access controls are dynamically established, based on the identification of the user connecting to the network. Frequently, endpoint security posture checks are also done, which may affect the access controls, such as sending a user to quarantine if their antivirus software is not up to date.

### **Port mirroring**

Port mirroring is used in intrusion detection, protocol analysis and application debugging. It refers to a configuration of an Ethernet switch used to send copies of frames from a particular virtual LAN or switch port to a different switch port.

### **Power over Ethernet (PoE)**

PoE is a technique used to send DC power over the same cable used for Ethernet communications. The original PoE standard, IEEE 802.3af, allowed for devices to consume up to 13 watts of 48 volts of DC power. Newer standards provide higher wattage (up to 50W) for devices that require more, such as some 802.11n access points.

### **Quality of Service (QoS)**

QoS is the ability to guarantee different levels of performance to different data traffic across a network, or to prioritize different flows. QoS may include specific protocol-based resource reservations for a particular flow, or it could include more static configuration for entire classes of traffic, such as all VoIP traffic or all e-mail traffic.

### **Secure Sockets Layer Virtual Private Network (SSL VPN)**

SSL VPN is a remote access VPN technology that encrypts user traffic to a VPN gateway using the SSL/TLS encryption protocols. SSL VPNs can be used in network extension mode, competing directly with IPsec VPNs to allow end users true IP access to a remote network.

### **Service set identifier (SSID)**

SSID is an identifier for a single 802.11 wireless LAN. All access points advertising the same SSID should be offering equivalent service, allowing wireless client devices to move between access points transparently.

### **Spanning tree**

Spanning tree is a network protocol that removes topological loops from bridged Ethernet LANs. This protocol creates a tree within a meshed network, then disables redundant links to leave a single path between any two nodes in the network.

### **Storage area network (SAN)**

A SAN is a type of network-based storage that presents block-addressable devices to each host. SAN volumes appear as if they were unformatted disk drives. Typical SAN access protocols include Fibre Channel (FC), which encapsulates SCSI commands over a dedicated physical network, and iSCSI, which encapsulates SCSI over TCP/IP.

### **Switched port analyzer (SPAN)**

SPAN is a term used by Cisco to refer to the manufacturer's technology for port mirroring.

### **Uplink**

An Ethernet switch uplink port is used to connect a switch to higher layers in the organization's LANs. For example, edge is uplinked to distribution, and distribution is uplinked to core.

### **Virtual local area network (VLAN)**

A VLAN is an IEEE standard that lets multiple administratively separate networks operate over the same cable. Each frame on the LAN is tagged with a 12-bit value indicating which VLAN the frame belongs to. VLAN-capable switches can add or remove tags on a per-port basis to accommodate the devices being connected.

### **Virtualization**

This term refers to the creation of a virtual version of a device or resource that is then located on a partitioned environment, such as a SAN or a server. Virtualization increases hardware utilization levels, but has many other benefits as well, such as simplifying system deployment and creating high-availability configurations.

### **Wi-Fi Protected Access (WPA and WPA2)**

WPA and WPA2 are protocols established by the Wi-Fi Alliance based on IEEE 802 standards for wireless security. WPA and WPA2 replace WEP.

### **Wired Equivalent Privacy (WEP)**

WEP was introduced as part of the original 802.11 protocol in the 1990s. It was one of the first security algorithms for wireless networks. Now largely discredited, WEP is not recommended.

### **Worldwide Interoperability for Microwave Access (WiMAX)**

WiMAX is a family of wireless technologies designed to offer medium-speed (40Mbps) connectivity to end users, households and businesses. WiMAX is typically seen as a competitor with other broadband technologies, such as DSL and cable modems.

## Disclaimer

The terms and conditions of product sales are limited to those contained on CDW-G's website at CDWG.com. Notice of objection to and rejection of any additional or different terms in any form delivered by customer is hereby given.

For all products, services and offers, CDW-G reserves the right to make adjustments due to changing market conditions, product/service discontinuation, manufacturer price changes, errors in advertisements and other extenuating circumstances. CDW®, CDWG® and The Right Technology. Right Away® are registered trademarks of CDW LLC. PEOPLE WHO GET IT™ is a trademark of CDW LLC. All other trademarks and registered trademarks are the sole property of their respective owners. CDW and the Circle of Service logo are registered trademarks of CDW LLC. Intel Trademark Acknowledgement: Celeron, Celeron Inside, Centrino, Centrino Inside, Core Inside, Intel, Intel Logo, Intel Atom, Intel Atom Inside, Intel Core, Intel Inside, Intel Inside Logo, Intel Viiv, Intel vPro, Itanium, Itanium Inside, Pentium, Pentium Inside, Viiv Inside, vPro Inside, Xeon and Xeon Inside are trademarks of Intel Corporation in the U.S. and other countries. Intel's processor ratings are not a measure of system performance. For more information please see [www.intel.com/go/rating](http://www.intel.com/go/rating). AMD Trademark Acknowledgement: AMD, the AMD Arrow, AMD Opteron, AMD Phenom, AMD Athlon, AMD Turion, AMD Sempron, AMD Geode, Cool 'n' Quiet and PowerNow! and combinations thereof are trademarks of Advanced Micro Devices, Inc. This document may not be reproduced or distributed for any reason. Federal law provides for severe and criminal penalties for the unauthorized reproduction and distribution of copyrighted materials. Criminal copyright infringement is investigated by the Federal Bureau of Investigation (FBI) and may constitute a felony with a maximum penalty of up to five (5) years in prison and/or a \$250,000 fine. Title 17 U.S.C. Sections 501 and 506. This reference guide is designed to provide readers with information regarding network infrastructure. CDW-G makes no warranty as to the accuracy or completeness of the information contained in this reference guide nor specific application by readers in making decisions regarding network infrastructure. Furthermore, CDW-G assumes no liability for compensatory, consequential or other damages arising out of or related to the use of this publication. The content contained in this publication represents the views of the authors and not necessarily those of the publisher.

©2011 CDW Government LLC  
All rights reserved.

# INDEX

10Gb Ethernet.....	4-5, 7-8, 22, 23	Network management tools .....	23-25
40MHz channels.....	27-29	Network optimization .....	21-25
802.11n.....	4, 5, 26-29	Power over Ethernet (PoE) .....	25
802.1s .....	7, 9, 23	Quality of Service (QoS).....	4, 25
Application delivery controller (ADC).....	4, 10	Redundancy.....	7, 8, 10, 23
Bandwidth.....	5, 8, 9-10, 27	Secure Sockets Layer virtual private network (SSL VPN).....	21, 32
Cloud computing .....	4-5, 9-11, 30	Service set identifier (SSID).....	29, 31
Collapsed backbone .....	4	Spanning tree.....	4, 7-9, 23
Continuity of operations (COOP) .....	5	Storage area network (SAN) .....	6-7, 9
Data deduplication .....	4	Switched port analyzer (SPAN) .....	25
Data loss prevention (DLP).....	4, 10, 24, 30, 32	Top-of-rack switches.....	4, 7-8
Dynamic Host Configuration Protocol (DHCP).....	8, 24	Unified threat management (UTM) .....	4, 30
Edge, distribution, core.....	21-22	Uplink.....	7-8, 23
Fibre Channel over Ethernet (FCoE).....	25	Virtual local area network (VLAN) .....	9, 23, 25, 32
Lights out .....	7	Virtual private network (VPN).....	4, 30-32
Log management.....	11	Virtualization.....	3, 5, 6-9
Multiple Input Multiple Output (MIMO) .....	26-27	Wi-Fi Protected Access (WPA and WPA2) .....	29
Network access control (NAC).....	4, 11, 31-32	Worldwide Interoperability for Microwave Access (WiMAX) .....	26