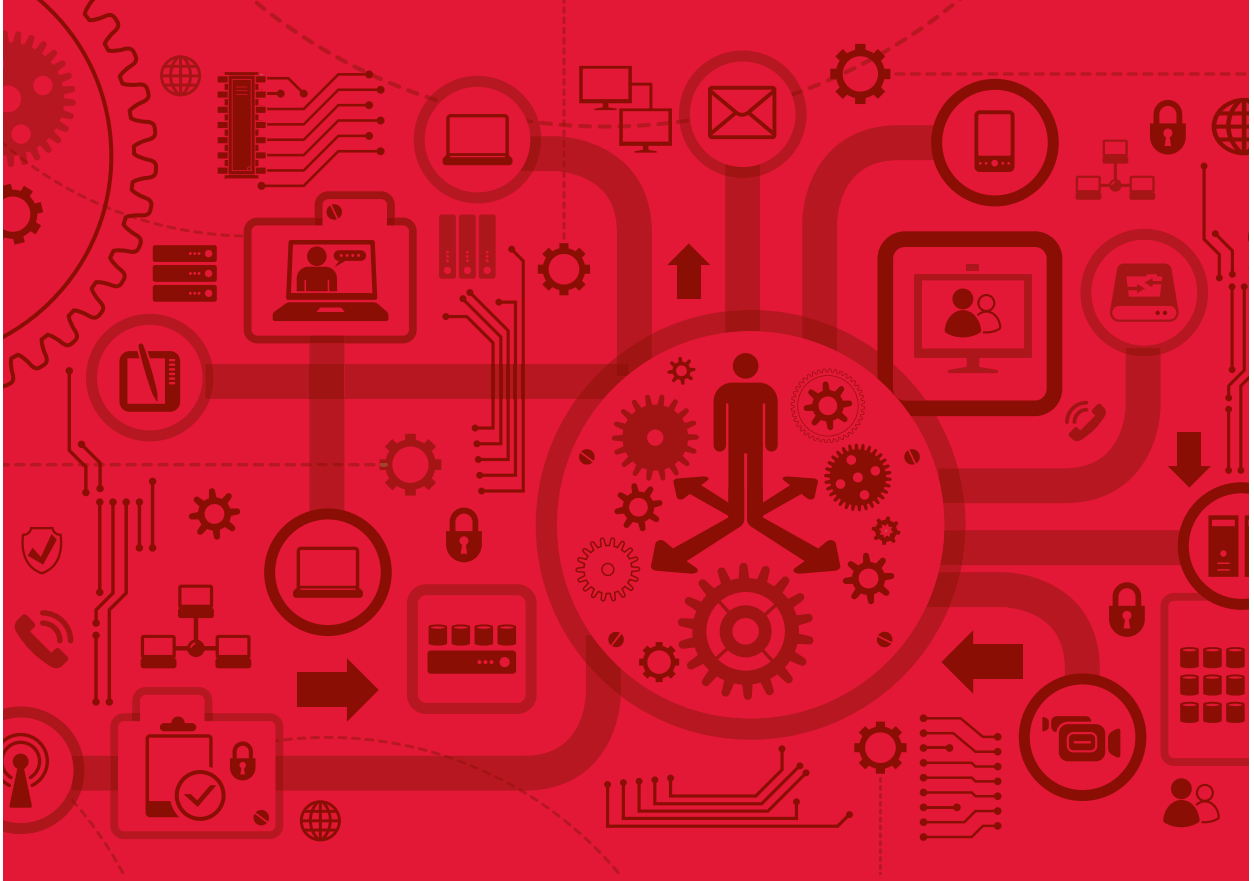

NETWORK INFRASTRUCTURE

Adapting networking strategies and techniques for today's challenges

CDW REFERENCE GUIDE
MAY 2013 | 800.800.4239 | CDW.COM/NETWORKGUIDE

CDW PEOPLE WHO GET IT™

WHAT'S INSIDE:

Making it easy to find out what's new >>>

800.800.4239
CDW.COM/NETWORKGUIDE

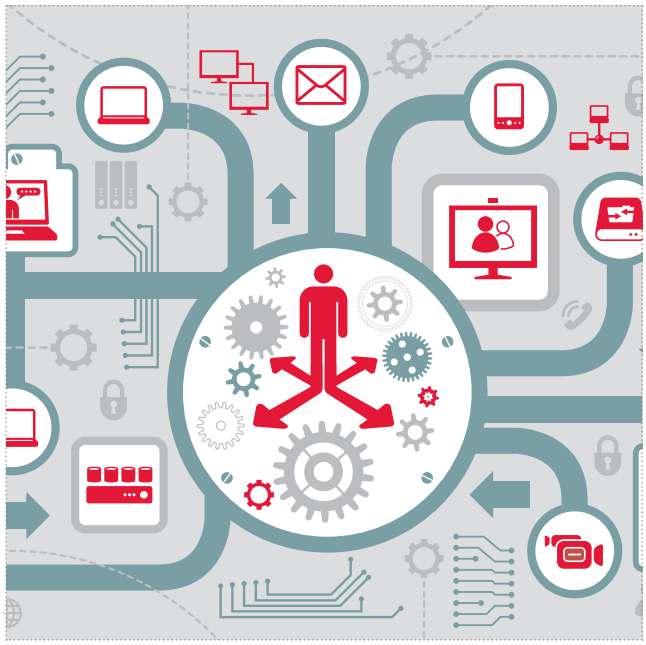

6 ENTERPRISE CENTRAL COMMAND

For more information on network infrastructure optimization, visit **CDW.com/networking-solutions**

WHAT IS A
CDW REFERENCE GUIDE?

At CDW, 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 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.

SCAN THIS!

Find out how team captain Charles Barkley and the CDW squad win at home *and* on the road.

GET M.CDW.COM ON THE GO

m.cdw.com is now available anywhere with our new mobile-friendly website or download the CDW app for your iPhone from the App Store.

GET IT at **m.cdw.com**

THE NETWORK: ACCOMMODATING CHANGE

Today's networking strategies must take into account four important technology trends.

Without a reliable network, many (probably most) organizations would have to simply shut their doors. They could not operate on a daily basis. Today's IT leaders are serious about their network investments because they know that networks are serious business. March 2013 results from the Duke University/*CFO* magazine Global Business Outlook Survey show that technology spending by organizations will increase by 8 percent this year, a big jump from 2012. The question is not whether to invest in the network, but how to make the wisest and most economical investments.

Every year, the landscape of organizational computing changes. Sometimes these shifts are long lasting. At other times, they are discarded quickly when they don't yield sufficient benefit or when the timing is not quite right for this particular evolution.

Currently, there are four significant trends influencing IT operations in all types of organizations. Each of these developments brings a new focus on the network that IT managers should consider carefully as they set technology maintenance, upgrade and expansion strategies for their organizations.

Development 1: Cloud Computing

C-level executives everywhere are urging their IT teams to move services to the cloud. For many network managers, who've long represented networks on their infrastructure plans as fluffy clouds, this latest trend is often rather puzzling.

In the realm of enterprise networking, organizations essentially moved to cloud-based application deployment years ago. Applications are already running in these enterprise private clouds: Users connect to the data center via the cloud — the organizational network, often supplemented with the Internet and virtual private network (VPN) concentrators. Data transmission is transparent to the end user; the data center location is unimportant.

Fortunately, network and data center

>

managers don't need to wade into the cloud definition hysteria and can avoid debating software/platform/infrastructure as a service (SaaS, PaaS, IaaS) with their colleagues. From a network point of view, cloud computing simply implies that the traditional geographic congruence of application and user no longer holds true.

In addition to cloud computing, organizations everywhere are making greater use of interoffice and Internet bandwidth as part of their workday – whether because they've moved from sharing text documents to sharing multimedia presentations and videos, or because software apps have become more resource-intensive. It all adds up to greater dependence on organizational wide area network (WAN) and Internet connectivity.

So what is the implication for network and data center managers? They must plan for high-reliability and high-quality networking over the complete path from the user to the application, no matter where the user or the app happens to be.

Traditional metrics, such as bandwidth, latency and uptime guarantees must be extended to places where the organization has little or no control. Thus, although the notion of cloud computing may not be new, this reliability challenge is a new task that many network managers are struggling with as IT shops commit themselves to the cloud.

TODAY'S MAJOR DEVELOPMENTS IN NETWORKING

Trend	Effect on the Network
Cloud computing	Requires high-reliability external bandwidth; increases overall consumption of bandwidth offsite
Consumerization and bring your own device (BYOD)	Requires re-evaluation of LAN security; may require creation of multiple connection strategies to securely accommodate users' personal devices
Mobility	Necessitates that wireless capabilities for both networks and devices achieve production-level quality without compromising security; requires organizationwide availability of collaboration tools – such as Voice over IP, email, video conferencing, presence awareness and instant messaging – anytime, anywhere and for any device
Virtualization	Requires rearchitecting data center LANs for high-density, high-use bandwidth per port

Development 2: Consumerization and BYOD

Consumerization of IT and bring-your-own-device (BYOD) programs are two sides of the same coin, a radical shift in the way organizations view technology generally. Originally, IT departments were insular groups concerned with deep techno arcana, providing technological guidance to a docile crowd of end users.

Now that everyone in the organization has access to inexpensive desktop and notebook computers, handheld devices, Internet connectivity and world-class Internet-based apps, IT departments must support a very different breed of user and executive. Workers now are often well educated about technology; experienced with many types of systems, IT gear and apps; and have high expectations about how the IT group can enhance personal productivity and organizational efficiency.

The term *consumerization* describes the movement of IT products and services to the consumer level. Email, Internet access, Voice over IP (VoIP), multimedia, real-time collaboration and online access to information are no longer the domain of well-heeled organizations, research groups and universities.

The entry price for many IT devices and tools is low enough that almost everyone in an organization likely has some type of personal device or computing system. What only a couple of decades ago was a resource for rarified environments has become as mundane and common as the telephone.

What's more, consumerization has changed the way that workers view enterprise apps and IT groups. When workers at all levels manage their own home networks, sometimes with dozens of devices, their expectations of what enterprise IT groups should be doing changes dramatically.

Every IT choice, from upgrades and antimalware to wireless and device choice, suddenly comes into question. When workers have access to world-class products on home computers and via the Internet – resources such as web-based email and collaboration tools, multimedia communications products and fast-responding Asynchronous JavaScript and XML (AJAX) applications – senior IT managers can find themselves defending and justifying enterprise application architectures, network designs and device choices in situations far beyond the boardroom.

BYOD programs can take many forms:

- allowing staff to read work email on their smartphones and tablets;
- exporting enterprise data via mobile apps and web-based query tools;
- taking over some (or all) device management tasks;
- allowing enterprise apps to be installed on personal devices.

No matter where an organization falls on the BYOD deployment spectrum, the IT group must address two key questions: What will support for user devices cost? How will the organization ensure data and network security?

Development 3: Mobility

Mobility isn't so much a trend as it is a new way of life. What started before the turn of the last century as a series of pressures — decreasing hardware costs, increasing real estate expenses, evolving work habits and attention to work-life balance — has now become a mainstream way of thinking about teams, productivity and daily work.

With notebook purchases representing a substantial portion of new computers sold, it's clear that mobility within the organization is not a fad; it's here to stay. The growing adoption of business-use tablets and the release of the tablet-optimized Microsoft Windows 8 operating system (OS) only serve to emphasize the mobile nature of today's workforce.

IT and network managers must carefully consider the impact of mobility initiatives within their organizations to keep ahead of the power curve in supporting mobility. In some organizations, mobility may simply mean notebooks are dragged into conference rooms and occasionally used from home. In others, it may mean that users don't have a fixed location anymore, moving from workgroup to workgroup, city to city, or even country to country — with a requirement for total productivity wherever they are in the world.

Generally though, mobility means that the traditional detritus of the office, including a desktop computer and telephone, as well as a "known location," are nonexistent or minimized. Although the burden of ensuring the collaboration and communications functions necessary to support a mobile workforce doesn't fall entirely on the IT department, there are certainly technologies available to make up for the move away from a standard, fixed office.

IT and network managers will want to focus on beefing up wireless capabilities, deploying broader access to tools such as VoIP, email, video conferencing, presence and IM, all while keeping a close eye on security.

Development 4: Virtualization

If mobility is a near-universal trend for users, virtualization fits the same niche when it comes to system managers. The flexibility, reliability and expandability provided by virtualized systems such as VMware vSphere and Microsoft Hyper-V, have taken hold in data centers everywhere and are dramatically changing the way that organizations provision and deploy servers and apps.

Highly desirable features, such as hardware redundancy and storage replication, are now much simpler because of the abstraction layer provided by virtualization. The few downsides of virtualization, such as the potential for virtual machine sprawl, are barely noticeable compared with the technology's enormous advantages.

Virtualization has been a quiet revolution because system managers can migrate slowly and methodically with few abrupt impacts. Intel continues to deliver on Moore's Law, with faster and faster systems at affordable prices. And attached subsystems, including memory and disk, have increased in capacity without increasing in price. These supply-side economic realities mean that the added cost and overhead of virtualization is easily absorbed into existing budgets.

But as system managers slowly grow their virtualization farms, they must address other limitations in infrastructure brought about by the new technology, such as ceilings on LAN bandwidth within the data center, oversubscription at the aggregation or core, or even troubleshooting down to the virtual machine level.

There are also physical infrastructure issues, such as power and cooling distribution. A fairly low-priority problem before virtualization, power and cooling have come front and center given today's broader green initiatives and the effect of increased equipment density brought about by virtualization, which influences power and cooling use within a facility.

No one wants to reduce their commitment to virtualization, but system managers must realize that its use has far-reaching effects, often unanticipated, and that rebuilding data centers may be required to provide the infrastructure needed for effective use of virtualization technologies. ■

WAN acceleration is not a high-profile trend like cloud computing or BYOD. But this development is having major effects on the network. **To learn more about it, see Chapter 3, starting on Page 20.**

ENTERPRISE CENTRAL COMMAND

The data center is the starting point for network improvements.

The outlook for data center managers is both scary and exciting. The scary news first: Everything in the data center is in motion, being built differently from the way it was 10 years ago. And now the good news: Even though the infrastructure has changed radically, the same core requirements and many of the same technologies apply. Previous experience will help manage the change and deliver strong return on investment.

All of the technology has evolved, from the servers and the cable plant to the air-conditioning and management tools. But the goals of increasing density, improving reliability and speeding communications all still apply.

Server virtualization and consolidation, the key drivers of today's data center changes, have been building for more than a decade. Back in 1983, data centers typically had a small number of mainframes running a large number of applications. For the next 20 years, the shift to Windows-based computing caused an explosion in the number of computing elements, supported by the shrinking size and cost of servers. With many apps designed to require 10 to 15 servers, typical data centers jumped from two servers to 200.

Keeping these burgeoning server loads manageable has made virtualization a must-have technology. In addition to controlling the sheer number of physical devices, virtualization has also made the process of deploying new servers for production, test, development, disaster recovery and quality assurance environments as simple as a few clicks on a console.

High-density and High-speed Networks

The use of space-efficient servers is one of the first signs that a data center is growing its virtualization capabilities. With storage moved to a storage area network (SAN), and modern designs offering efficient in-server cooling and packaging, most data center managers prefer to buy

the smallest servers available. By standardizing on easily replaced and conveniently sized computing elements, organizations can increase capacity in large enough chunks to keep from having to add new servers every day.

Because of this, buying patterns have shifted from physically large servers with built-in disks and redundant array of independent disks (RAID) to 1U and 2U servers. In fact, few manufacturers even make general-purpose servers larger than 2U, with the exception of blade server chassis, which are themselves ultra-high-density devices that pack more than one server per 1U of rack space.

The biggest issues resulting from these high-density server designs are their power and cooling. Design goals for data centers built in 2003 called for power densities of 2 kilowatts to 4 kilowatts per filled rack. But increasing density has brought that number closer to 10 kW when racks are filled with modern 1U and 2U servers or blade chassis.

Power going into a rack means heat flows out of it. So data centers have had to rebuild their heating, ventilation and air-conditioning (HVAC) infrastructures and revise room layouts using hot and cold aisles to keep racks from melting down.

Fortunately for network managers, Ethernet speeds keep increasing, especially in space-constrained environments such as data centers. Although 1 Gigabit Ethernet connections are now considered the standard minimum for any production device in a data center, most network managers are preparing for 10 Gig-E deployment in their next upgrade cycle.

10 Gig-E is commonly used for interswitch links; for example, between top-of-rack switches, distribution layer switches and core switches. Servers are also shipping with 10 Gig-E capability using inexpensive, off-the-shelf adapters. Network managers should keep in close communication with server managers to be sure that all production servers being installed either have 10 Gig-E already installed or have open slots for an upgrade.

In the meantime, manufacturers are working on 100 Gig-E for future high-capacity data center networks. An interim standard, 40 Gig-E, can be deployed in areas where 10 Gig-E or multiple 10 Gig-E connections don't meet requirements. Generally, organizations should carefully consider 40 Gig-E because it's an interim standard and won't reach the adoption levels or economies of scale that 100 Gig-E eventually will.

VIRTUALIZATION'S **UNIQUE CHALLENGES**

Common Requirements	Challenges
High density	Virtualized environments pack many more servers per rack than ever before, as traditional space hogs such as disk arrays are moved to the SAN, and servers shrink to 1U (1.75 inches) or 2U (3.50 inches).
High reliability	Virtualized environments spread their virtual machines (VMs) across physical hosts in an unpredictable way to manage loads, which means that any breakdown in the data center network can affect many apps.
High speed	Virtualized environments with multiple VMs per physical host do not follow traditional guidelines for estimating bandwidth per port, which requires rethinking speeds and feeds, and may require flattening the network from routed to bridged in some areas.

Top-of-rack Switches

One of the challenges that virtualization creates is the difficulty of bringing sufficient network connectivity to a rack filled with virtualized systems. Most network managers are turning to top-of-rack switches and building high-reliability stacks in each rack to handle the connections.

Top-of-rack switches (the term used no matter where they're placed in a rack) are specifically designed for server connections at the edge of a network. Each switch has 48 network ports for servers or other in-rack devices and several high-speed uplinks to either distribution layer or core switches in the data center.

Top-of-rack switches may be based with 1 Gig-E or 10 Gig-E. But their uplinks are typically 10 Gig-E, with at least two ports dedicated to uplinks. Top-of-rack switches usually also have dual power supplies, or at least dual power cords, so that they are compatible with the A/B power systems installed in most high-availability data centers.

In a 42U rack (the standard size for data centers), most designs will aim for about 30 1U servers. This leaves about 10U for other infrastructure, including patch panels, network switches, power distribution, network and power cable management, and other infrastructure devices, such as environmental sensors, local KVM switches, serial console devices and the like.

>

FLATTENING DATA CENTER **NETWORKS**

Data center networks have traditionally followed a three-tier model: core, distribution and edge. Top-of-rack switches cement this arrangement into place at the server edge of the network.

To simplify system complexity and add speed, some network managers try to flatten networks down to two tiers when 10 Gigabit Ethernet or speedier connections are in place.

Here are some of the benefits and challenges of flattening networks:

Goal	Three-tier Network	Two-tier Network
Simplicity	The more devices on the network, the more management it will require.	With fewer devices, this topology requires less management.
Low latency	More hops lead to higher latency.	Fewer hops result in lower latency.
High bandwidth and reduced oversubscription ratios	Interswitch links can be a bottleneck because of oversubscription. Often, lower-cost links (one or more 1 Gig-E) are used.	Less oversubscription occurs because there are fewer devices (and links between them). But higher-speed links are required for increased bandwidth, which increases costs.
Efficient use of power and space	More devices drive up power consumption and footprint. Cabling is simplified across multiple devices.	Fewer devices need less power and less space. Cabling requires careful planning to manage densities.
Scalability	It is easy to scale up – just add edge switches.	It's not very scalable. When the distribution is full, adding one more device requires a major redesign.

Each of the 30 1U servers typically has a minimum of two to four Ethernet ports for data communications, one port for a baseboard management controller or lights-out management system, and two ports for some type of SAN connection (generally Fibre Channel or iSCSI).

The network complexity multiplies quickly: five to seven ports per server times 30 servers equals 150 to 210 ports per rack. For this reason, network managers are moving away from a huge core switch (or switches) in the center of the facility and turning instead to top-of-rack or end-of-row devices.

No one wants to backhaul 200 or more network ports per rack, some of which may be fiber optic (in the case of Fibre Channel or 10 Gig-E) to a central switch. Patch panels, cabling and patching costs, and cable management make this a support nightmare.

Five ports per server may be a conservative estimate in some data centers. Current virtualization best practices call for binding several high-speed network ports into a single, highly reliable and fast network channel with multiple virtual LANs, requiring two to four ports per server.

Some virtualization architects don't agree with that approach and have elected to use separate physical channels for the data plane, virtualization system management and virtual machine (VM) migration. This can boost per-rack loads to 300 network ports or more.

All of this adds up to a logical design based on three to five top-of-rack switches. In a network with separate data and storage networking, four 48-port switches are commonly used, with a fifth switch used for out-of-band management. If data and storage networking are converged, fewer ports are required (minimally, three ports per server). But because converged ports are always 10 Gig-E ports, switch densities can be lower, with fewer 48-port switches available today.

Storage Area Networks

There are differing opinions on how SANs should best be handled in data centers. Although iSCSI developers have shown tremendous performance, especially using 10 Gig-E connections, data center managers generally prefer Fibre Channel networking for storage networks. For some data center managers, keeping data and storage networking separate is a desirable goal, while for others convergence eases complexity and reduces the management burden.

Standard Fibre Channel switch prices (per port) and host adapters are more expensive than Ethernet. In addition, Fibre Channel has improved speed (from 1 gigabit per second to

16Gbps) across multiple generations of devices in a short period of time. During Fibre Channel's effective lifetime, Ethernet has changed equipment and generations once, from 1Gbps in 1999 to 10Gbps in 2002.

Compared to Ethernet, Fibre Channel's many small speed increases have been inconvenient. As virtualization growth has brought ever-larger numbers of servers and SAN devices together, many storage managers have upgraded through five generations of Fibre Channel switches, from 1Gbps (200 megabytes-per-second throughput) to 2Gbps, 4Gbps, 8Gbps and finally to 16Gbps (3,200MBps). The 16Gbps Fibre Channel devices have been available since 2011, although adoption of 16Gbps Fibre Channel has been slow because of interoperability issues with existing SANs and the difficulty of upgrading existing installed storage networks.

The inconvenience and cost of maintaining a separate network for storage has led many network managers to look toward Fibre Channel over Ethernet (FCoE) to maintain the performance of Fibre Channel while gaining the benefits and simplicity of Ethernet. But Fibre Channel can't simply be transmitted over existing Ethernet hardware.

It requires more of the network than Ethernet does. Fibre Channel's protocols expect a lossless network, and Fibre Channel switches are more active in routing and access control (security) than their Ethernet equivalent.

For example, Fibre Channel fabrics actively try to spread load across multiple paths to increase performance, while standards-based Ethernet keeps all traffic for the same virtual LAN (VLAN) on the same path. Therefore, FCoE isn't just a buzzword, but an extension of Ethernet switching technology that supports the additional requirements of Fibre Channel networks.

FCoE also requires 10 Gig-E because the Ethernet network must appear lossless to the Fibre Channel devices. Even the slowest Fibre Channel network connection cannot reside in a standard 1 Gig-E pipe. FCoE is currently standardized for 10 Gig-E, 40 Gig-E and 100 Gig-E. Most experts believe that FCoE connections between servers and switches will use 10 Gig-E for the foreseeable future, with the faster 40 Gig-E and 100 Gig-E speeds reserved for interswitch links.

Data center managers interested in FCoE should also specify multihop FCoE when selecting equipment and designing topologies. Initial FCoE products were mainly aimed at converged adapters on the server side, cutting port counts

>

SELECTING **TOP-OF-RACK SWITCHES**

Feature	Data Center Switching Requirement	Benefit
Stackability	Stacking of switches reduces management overhead. All switches in a rack should be managed as a single unit, interconnected with 10 Gig-E or higher speed links.	Management is hard enough without having to worry about which switch a server is plugged into in a single rack.
Aggregation	Switches must support large numbers of link aggregation groups, and the groups should be spread across multiple switches.	Each server will need its own group. Groups should be spread across switches to eliminate single points of failure.
Uplink	At least two 10 Gig-E ports should be available for uplink. Four may be required if standard switch ports are used for stack communications, rather than a dedicated stacking port.	Virtualization pushes hardware much closer to its limits, which means that normal network utilization metrics no longer apply. Server densities have increased tenfold, so bandwidth requirements for a single rack increase as well. High availability requires port pairs.
Reliability and topology healing	All switches should support (and be configured for) 802.1s Multiple Spanning Tree Protocol (MSTP), which also includes 802.1w Rapid Spanning Tree Protocol (RSTP).	MSTP can increase network utilization by spreading the load across redundant links, which is important when using vMotion, disk-to-disk backups, Network File System or iSCSI. RSTP improves high availability by reducing convergence time for the spanning tree during errors.

in half in the rack by combining data port and Fibre Channel traffic on the same Ethernet wire. The FCoE connection was made to the edge switch only, and edge switches had to have both Fibre Channel and Ethernet uplinks to the rest of the network.

With multihop FCoE, Fibre Channel is carried across multiple Ethernet switches end to end between SAN disk subsystems and client servers, or it is only converted back to pure Fibre Channel just before connecting to the SAN disk subsystem. The use of multihop drives up performance and speeds routing within the data center.

Network Uptime for Virtual Servers

Virtualization represents the opportunity to rethink how servers are connected in the data center. Best practices call for redundancy in every connection, which requires identifying — and eliminating — single points of failure between every server and the network core (typically, already redundant). When server-to-network redundancy is done right, it also increases performance by doubling or quadrupling total bandwidth from server to network.

From a high-availability point of view, it's sufficient to just double every connection: two ports from each server feeding to different switches, two connections from each switch (or switch stack) going to a different distribution layer, and then a redundant connection from the data center distribution layer up to the core of the network.

Many network managers have already moved to redundant network connections at the server level, but often without understanding and implementing the basic technology at work. Redundant connections without proper link aggregation tend to cause more problems than they solve. Similarly, redundant backbone linkages made without paying attention to spanning tree protocol (STP) choice and configuration can create unexpected results and may not offer the desired benefits.

Link aggregation is the standards-based term commonly used for the practice of combining multiple Ethernet connections into a single larger connection. Other terms used to describe this technique are "bonding," "teaming," "port trunking" and "EtherChannel."

The IEEE 802.1AX standard (previously published as 802.3ad in 2000) defines link aggregation. Although all modern managed switches support standards-

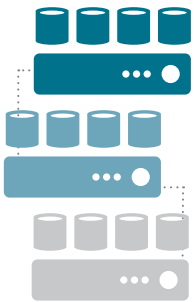

VIRTUALIZATION REPRESENTS THE OPPORTUNITY TO RETHINK HOW SERVERS ARE CONNECTED IN THE DATA CENTER.

IS **INFINIBAND** ***REPLACING*** FIBRE CHANNEL?

Infiniband is a high-speed network connection with speeds ranging from 20 gigabits per second to 56Gbps. It relies on switching fabric that is commonly used in high-performance computing and some types of storage subsystems.

Originally conceived of as a universal input/output bus, Infiniband was used to simplify and speed communications within computing systems and in the data center, building on Ethernet, Peripheral Component Interconnect (PCI) and Fibre Channel topologies, among others. Infiniband has become popular among builders of supercomputers because it offers higher speeds than typical Ethernet devices and can meet the needs for very high-speed internal interconnects.

Makers of Infiniband equipment have taken the lessons learned from supercomputer deployments and are gradually releasing commercial SAN products based on Infiniband for enterprise use. Typical Infiniband products start with a capacity of 1 petabyte (1,000 terabytes, or 1 million gigabytes) and are designed to handle Big Data applications, such as scientific computing, governmentwide databases and huge commercial applications (think Amazon, Facebook and Wal-Mart).

Although Infiniband outperforms Fibre Channel, the enormous previously installed base of Fibre Channel equipment and the slow-to-shift nature of most enterprise computing environments suggest that Infiniband will remain a niche technology for at least the near future.

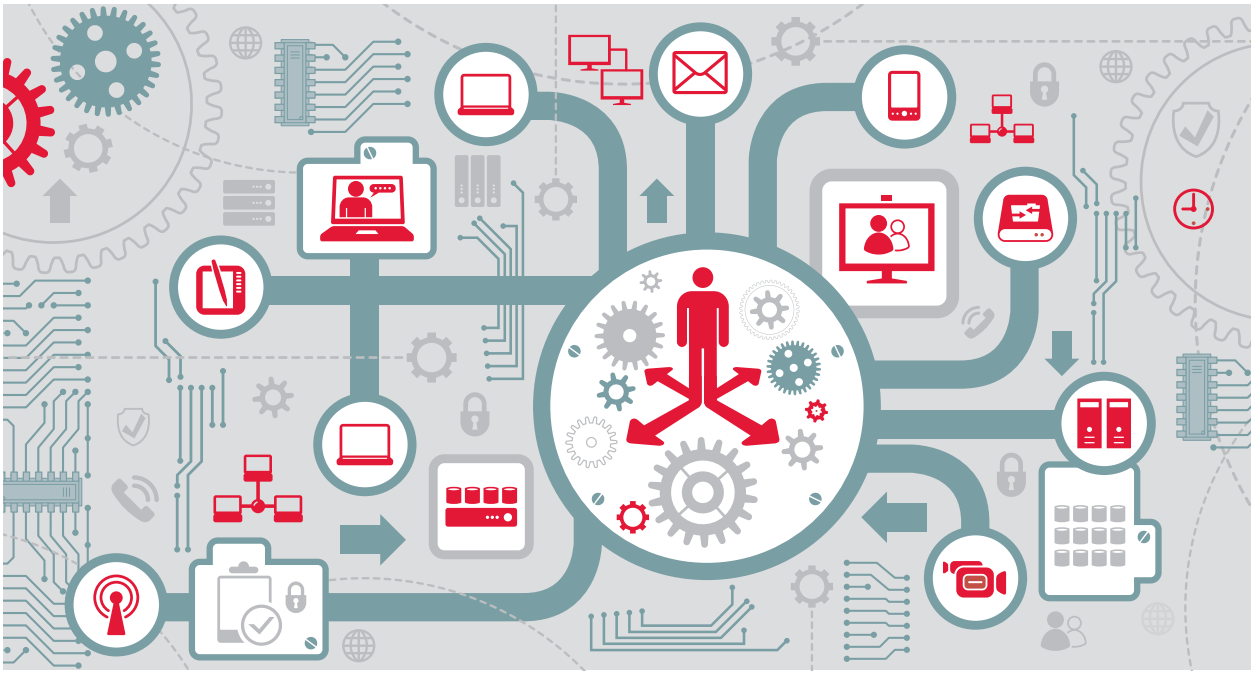

based link aggregation, many system managers have not deployed proper link aggregation technology.

One reason that system managers and network managers do not see eye-to-eye on this topic is the complexity of managing link aggregation within the operating system. Although network managers are generally comfortable with the terminology and technology used for link aggregation, system managers find the added confusion of proper configuration a barrier. Additionally, because many OS and LAN switch combinations work without any such configuration, system managers have been misled into thinking that no configuration is fine.

But link aggregation provides worthwhile benefits. Generally, two links from the same device connected to the same switch will operate in an active/passive mode, where one link randomly becomes the active link and the other link is used only when the first link fails.

Software and app makers have come up with a variety of nonstandard mechanisms to manage the links and even to perform some types of load balancing. But in the absence of a protocol-based communications channel between the host and the LAN switches, all of these nonstandard mechanisms have significant failure modes that can interrupt or degrade service.

Because one of the main reasons to have multiple links is to increase overall reliability, the best strategy is for network and system managers to make the additional effort required to properly configure link aggregation between critical servers and the LAN infrastructure. All enterprise-level managed switches already support standards-based link aggregation, so this is one area where network managers can help train system managers on the benefits, terminology and compatibility of link aggregation technology.

A subtle issue that can be easily overlooked when implementing link aggregation of any type for higher availability is multidevice link port groups. Best practices for deploying virtualized servers call for each server to have two data ports connected, each to separate switches. (In the case of Fibre Channel, the same two-switch rule is also present, but Fibre Channel handles the two ports within the intelligent fabric of the switches differently than Ethernet does.)

By connecting Ethernet links from the same server to two different switches, a failure of a switch, disconnection of a switch uplink or even periodic maintenance of a single switch will not cause any interruption in services. But for this to work properly, network managers have to be able to offer link aggregation port groups that span two switches serving the same rack.

Although all enterprise switches support link aggregation, not all switches allow ports to span different switches. Network managers, therefore, should verify the multichassis port group capabilities of their switches to be sure that link aggregation won't be blocked by limits in the LAN switches. ■

WE GET WIRELESS NETWORKS

WIRELESS THAT DELIVERS PRODUCTIVITY AND RELIABLITY

WITH 450 MILLION

ESTIMATED MOBILE WORKERS AROUND THE WORLD, BUSINESSES NOW OPERATE EVERYWHERE, ALL OF THE TIME.

Source: Riverbed Technologies; 2012 white paper: *The CIO's New Guide to Design of Global IT Infrastructure*

As mobile devices become more common in the workplace, IT managers must prepare their infrastructures for wireless access and networking. A key part of their strategies will have to address network performance, which is critical for providing a wireless user experience that is as good as a wired one. Consequently, many IT managers are upgrading their infrastructures to better support mobile devices and provide ubiquitous connectivity.

Advancements in wireless networking have allowed for more agile and accessible infrastructure. So much so that challenges such as bandwidth and application performance are no longer an issue. But before implementing a new wireless infrastructure, there are many things you must consider:

Layout | The number of locations, the layout of the buildings and the users are all important factors in the design of a wireless architecture.

Environment | Understanding the environment is important to the architecture of a wireless infrastructure. Vibrations, as well as dust and other particulates, might require a different configuration in order to achieve optimum performance.

Applications | It's critical to know which applications need transmission. Voice, video and data applications have different demands, so identifying the applications will help determine the required bandwidth.

Management | Wireless network management software can report valuable usage statistics, detect rogue access points and more. The right management software will depend on the monitoring wants and requirements of the organization.

Security | Wireless networks can be more secure than wired networks because of encryption of the data when transmitted, but there needs to be a special emphasis on threat prevention solutions when possible. A detailed assessment can determine which security precautions must be implemented.

With technologies constantly changing and evolving, it is important to select a partner that has a pulse on the industry. CDW has account managers ready to assist you and experienced solution architects with the latest industry training and partner certifications, including Aruba, Cisco and HP. Our solution architects will help you select the best products for your infrastructure needs and IT environment — whether they are from one vendor or many — so you have a solution made of the best the industry has to offer.

Beyond product selection, our wireless infrastructure experts are prepared to support you through planning, implementation and maintenance and to answer your questions along the way. Let our decades of experience and thousands of implementations guide your solution selection and implementation so you can move forward with confidence.

meraki

Meraki provides powerful and intuitive centralized management without the cost and complexity of traditional wireless controller hardware. Seamlessly manage organizationwide Wi-Fi deployments and distributed multisite networks with zero-touch AP provisioning, networkwide visibility and control, cloud-based RF optimization, seamless firmware updates and more — without training or dedicated staff.

ARUBA networks CDW.com/aruba

Aruba WLANs are amazingly versatile, giving you the freedom to make the best network design decision. Aruba WLANs can be deployed with or without mobility controllers. Bring on the mobile devices. Aruba WLANs are BYOD-ready and use contextual data — user identity, device type, applications and location — to enforce security and QoS policies. Powerful RF technology enables Aruba WLANs to deliver a predictable Wi-Fi experience in high-density mobile device environments. Personal mobile apps no longer adversely impact the performance of corporate mobile apps that are vital to your business.

D-Link CDW.com/dlink

A secure wireless connection everywhere.

D-Link wireless solutions bring robust, stable and secure access for business. The new generation of wireless products offers seamless connectivity, self-healing mechanisms, traffic segmentation and centralized management to achieve a wireless environment as productive and secure as a wired network.

D-Link's range of wireless N products provide stable connectivity, which is robust enough to be deployed at the very core of your network. They give greatly enhanced reliability and coverage, and include advanced security features to keep your network completely safe from intrusion. Products under this category include:

- Unified wireless switches and controllers
- Stand-alone and unified managed access points, in both single- and dual-band designs

---

cisco Gold Partner CDW.com/cisco

Wireless local area networks (wireless LANs, or WLANs) are changing the landscape of computer networking. Organizations around the world are installing overlay WLANs and freestanding, all-wireless networks to increase employee productivity.

Cisco Mobile Office Net software allows end users to manage email, schedule meetings, and access files and applications on the corporate network from conference rooms, break rooms, coworkers' desks, and virtually anywhere in the building. With wireless networking, end users are just a mouse-click away from key information and applications, regardless of where they are in a facility.

27%

THE PERCENTAGE OF **INCREASED PRODUCTIVITY** IN ORGANIZATIONS THAT HAVE **IMPLEMENTED A WIRELESS LAN.**

Source: Wireless Networking Basics IV: Planning and Deploying a Wireless LAN; squidoo.com

WIRELESS INFRASTRUCTURE SERVICES

When it comes to upgrading your wireless infrastructure, you need to make sure the solution fits your organization's needs. We can help you do just that from planning and implementation to support and security. Our pre-sales services include a logical wireless LAN design and comprehensive site surveys. And our experts can tailor your network offerings around:

- Wireless LAN architecture
- Indoor and outdoor coverage
- Voice over wireless solutions
- Active RFID and asset tracking
- Centralized wireless client management

LEARN MORE AT CDW.COM/WIRELESSNETWORK

WE GET NETWORK SECURITY

DEFEND YOUR NETWORK

A DATA BREACH IS ESTIMATED TO COST $214 PER RECORD REACHED, TOTALING **7.2 MILLION** PER INCIDENT.

Source: Ponemon Institute, *The Business Case for Data protection*

Your network is constantly under attack. Not only from external threats, but internally as well. A well-meaning staff member is almost as likely as an external attack to compromise your network. Whether it's through a gap in your security or an unwitting employee, threats like viruses, worms, rootkits, denial of service (DoS) attacks or malicious bots, are always looking for a way in. If you and your staff are not prepared, they're likely to get in. And if they do get in, the costs incurred can be substantial.

Securing your network requires more than just new hardware and software. It's about a change in your organization's culture to a more security-conscious environment. Your staff needs strong and diverse passwords. They have to be aware of every link they click on and every attachment they download. Even the most sincere worker can cause a small vulnerability that leads to a large data breach.

Defense-in-depth Approach

Defense in depth is the coordinated use of multiple security countermeasures to protect the information assets of an organization. It is based on the principle that it is more difficult for a threat to defeat a multilayered defense system than to breach a single barrier. It protects key network areas most susceptible to threats. These include: network gateway, server, client and application.

While there is no such thing as a fully secure network, a defense-in-depth strategy can help you keep your guard up and minimize risk. Plus, you'll get other benefits like:

Risk Mitigation

You can easily limit exposure to threats and safeguard against data loss by implementing foundational security tools like:

- Antivirus
- Antimalware
- Antispam
- Firewalls
- Content filtering
- Access controls
- Encryption

Enabling Productivity

A secure network can actually boost productivity. When social networking sites and remote access protocols are secured, workers can instantly access the data they need without worry or delay. You can give your staff the tools they need to be productive and still maintain a high level of security.

Forward-moving Organization

In most organizations, IT is the foundation of almost every activity. When the IT foundation is secured properly, the organization can move ahead smoothly and pursue its mission and goals. A secure IT environment ensures business and organizational success.

CDW.com/fluke

For nearly two decades, Fluke Networks has provided innovative solutions used by enterprises to provide their network installers, owners and maintenance staff with exceptional vision: combining speed, accuracy and ease of use to optimize network and application performance. Network optimization starts with knowing your network — who is on the network, where are they connected, how are links performing, measuring for over- or under-utilized resources, and more. Fluke Networks' OptiView XG provides networkwide discovery and path analysis, to determine the connection paths throughout your network and the health of every device and interface on the path.

CDW.com/juniper

Your security ecosystem starts here. Juniper Networks connectivity solutions enable organizations to use the same enterprise network infrastructure to securely connect remote and local users and devices. This security is essential technology for enabling initiatives, such as BYOD, that increase employee satisfaction and productivity. Enterprise network security solutions help IT administrators guarantee consistent security and policy applications across the entire data center. Juniper's connectivity solutions work to protect critical web-based applications against attacks as well as help maintain user-centric, policy-driven and application-aware visibility and control with extreme granularity.

CDW.com/symantec

Symantec offers products to help you improve threat monitoring, manage web traffic, prevent data loss and reduce the IT burden of protecting critical endpoints such as desktops, servers, notebooks and mobile devices. You'll be able to maximize the accessibility, availability and security of your IT infrastructures while protecting confidential data.

CDW.com/mcafee

Safeguard critical data and help ensure regulatory compliance with McAfee Data Protection solutions. Available individually or in suites, McAfee Endpoint Encryption and McAfee Data Loss Prevention solutions provide multilayered protection for your data regardless of where it resides — on the network, in storage systems, or at the endpoint.

ACCORDING TO THE 2012 VERIZON DATA BREACH INVESTIGATION REPORT, THE TOP METHODS OF ATTACK INCLUDED THE FOLLOWING:

55% Exploitation of default or guessable credentials

40% Use of stolen credentials

29% Brute-force attacks

25% Exploitation of backdoor or command-and-control channels

CDW'S COMPREHENSIVE **SECURITY ASSESSMENTS**

Our security assessments are customized to reflect your specific needs. We make a point of understanding what assets are important to your organization, and we don't run tests that deliver results in a vacuum. We understand that every network is different, and we make sure our report reflects your concerns and goals. We go beyond just testing for obvious problems and look for ways that an attacker could exploit design decisions, configuration choices, administrative practices and other factors to target your critical information. We typically take the role of a would-be intruder to determine whether your security measures can be defeated by custom attacks.

LEARN MORE AT CDW.COM/SECURITY

WE GET
WAN OPTIMIZATION

A WAN THAT WORKS

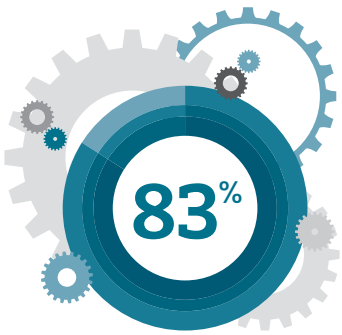

SOLVE PROBLEMS UP TO 83% FASTER WITH A WAN OPTIMIZATION SOLUTION.

Source: Riverbed Technology

Now more than ever, your network is the backbone of your organization. Workers need to accomplish their tasks from virtually anywhere. When out of the building, your staff needs applications to respond fast, and they need to access data quickly. Nothing slows your organization down like an overburdened WAN.

Your WAN keeps everyone in your organization moving. It's as important to your productivity as your workers are. When it's running slow, your organization runs slow too.

Data use is growing exponentially. Organizations like yours have become more reliant on bandwidth-intensive activities like:

VoIP | Requires a constant stream of data, up to 87.2Kbps.

Video Streaming | A typical web video can consume 2.5Mbps.

Video Conferencing | Usually consumes double the amount of video streaming for long periods of time.

File Sharing | Can cannibalize large amounts of bandwidth in heavy spurts.

Running without an optimized WAN can leave your organization exposed to:

Reduced Productivity | Without sufficient bandwidth VoIP, calls get interrupted and video conferencing becomes impossible.

Downtime | The only thing worse than a slow network is no network. Losses in revenue system damage and legal issues are all effects of a down network.

Three steps to consider when building your WAN:

1. Understand your network's strengths and weaknesses with application-level awareness tools.
2. Conduct end-to-end testing to find weaknesses before they become security issues.
3. Implement WAN controllers and ADCs to improve bandwidth management and increase application performance.

At CDW, we understand that network optimization can seem like a daunting task. That's why we offer you more than just products. We offer you the people and the plan to turn them into real solutions. Our certified networking solution architects can run an assessment of your current infrastructure and design a network that operates properly and securely. They can work with your account manager to help recommend products and services based on the needs of your organization. We can help you plan properly for deployment using a phased approach. And with Managed Network Services and our cloud offerings, we can help manage and monitor your improved WAN for you, reducing the workload on your IT staff.

riverbed Think fast.™ CDW.com/riverbed

A customer's Wide Area Network (WAN) is the foundation of their globally connected enterprise, which enables collaboration, communication, business productivity and risk mitigation. With Riverbed WAN Optimization solutions, networks run faster and more efficiently, delivering consistent service levels and cutting the costs of IT infrastructure.

Since WANs are fundamental to many businesses, optimization can deliver improvements to a wide range of top IT initiatives, including:

- Accelerating application performance, increasing business productivity
- Consolidating in the branch and data center
- Optimizing file sharing, web, email and even voice and video – all at the same time
- Leveraging cloud economics without compromise
- Protecting data more completely, with less cost and effort

Riverbed excels in WAN optimization, offering comprehensive innovations and strong performance for many applications.

f5 IT agility. Your way. CDW.com/f5

The performance of your website and applications sets the pace at which you do business. Improve performance and you can increase employee productivity, boost business operations and drive e-commerce revenue.

F5 Application Delivery Optimization solutions help you achieve top performance by making your network and applications work faster, use fewer resources, and operate more cost-effectively. With F5, you can:

- Improve web application performance – accelerate application response time, minimize latency and delays, and reduce the number of data round-trips necessary to complete a web request
- Speed replication between sites – eliminate latency and minimize packet loss and congestion so you can back up your critical data and applications in record time
- Keep mobile users happy – ensure applications load fast and look good on any combination of device, operating system and browser your employees and customers may use
- Scale your data center – contain costly infrastructure and bandwidth upgrades with solutions that offload processing-intensive tasks from servers, manage network connections, and reduce overall network traffic
- Secure access to applications and data from anywhere – F5 delivers an intelligent services platform that integrates application delivery, monitoring and context-based policy enforcement. You get a highly scalable, extensible and simplified approach to maximizing security

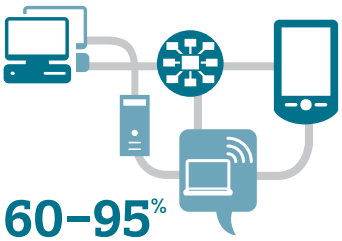

60–95%
OF TRAFFIC TO ANY ONE LINK CAN BE **ELIMINATED WITH WAN OPTIMIZATION.**

Source: *Adding Wow to the Wan*, CDW.com

CALL OUT **A NETWORK WITHOUT LIMITS**

It's probably no exaggeration to stay that today, your organization is only as strong as its network. Because the network stands in this critical central location, you're constantly on the lookout for ways to improve its performance and speed. We get it and we can help. Our network solution architects can show you how implementing WAN optimization solutions can help you achieve the performance levels you need.

LEARN MORE AT CDW.COM/NETWORK

WE GET
RESILIENT NETWORKS

MEETING EVER-GROWING NETWORK DEMANDS

THE AMOUNT OF DATA MANAGED IN ENTERPRISE DATA CENTERS WILL BE

50 TIMES LARGER IN 2020

THAN IT IS TODAY.

Source: *EMC Digital Universe Study*, 2012

The technologies that power organizations today are more powerful and flexible than ever. For example, virtualization helps data centers run at peak efficiency while keeping costs down. Software as a service (SaaS) simplifies the delivery and management of applications by keeping them in the cloud. Tablets, notebooks, smartphones and video collaboration tools keep workforces productive anywhere and everywhere they go.

However, all of these technologies live or die on one thing: your network. Many organizations took the leap into virtualization, SaaS, mobile computing and unified communication without making sure their network was ready. The network operations team was often left out of the planning conversations and now the network is buckling under the increased traffic.

The result? Sluggish performance across the board.

- Virtualization causes throughput instability and delays in the data center
- SaaS applications run at a crawl
- Tablets are frustratingly slow and not used
- Consumer-grade video apps like Skype work better than high-end video conferencing rooms
- Branch offices suffer slow connections to their data

Now, the network team is under pressure to resolve these issues, quickly and affordably, or else the benefits of these new technologies will never be realized. The organization may even end up being worse off because of the impact on network performance.

The key problem here is latency, which can't be solved just by adding more bandwidth — something that can be too costly, anyway. What's needed is a comprehensive strategy for optimizing your entire network, ensuring that it is ready to meet your needs both today and in the years ahead.

Network optimization is a cost-effective way to improve performance and the user experience, but can be a complicated process due to so many moving parts. That's why many organizations rely on CDW to get the most out of their networks. We understand the challenges you're facing and can help overcome them with expert advice and vendor recommendations.

CDW.com/hp

Enable tools to configure the network for efficient transport of information, bandwidth optimization, and enhanced application performance or automation of operations. Applications in this category enable efficient transport from LAN to service provider, across WAN links between sites and applications hosted in the data center.

CDW.com/cisco

The Cisco 7600 Series is a carrier-class edge router that offers integrated, high-density Ethernet switching, carrier-class IP/MPLS routing, and 10Gbps interfaces, benefiting enterprises and helping enable service providers to deliver both consumer and business services over a single converged Carrier Ethernet network.

CDW.com/barracuda

Free up your bandwidth, stay connected, and save money with Barracuda Networks. Your critical data always gets through thanks to traffic prioritization over multiple connections and 3G failover, even if your wired connections go down. The aggregation of disparate links of varying speeds – e.g., MPLS, T1, DSL, Cable, 3G – eliminates the need for costly high-capacity backup links. And, a Barracuda Networks solution means affordable, all-inclusive pricing with no per-user fees.

CDW.com/belkin

Belkin offers an extensive and complete series of networking cables. From Cat5e, Cat6, Cat6a to fiber, Belkin has you covered in every length, connection style and colors. Belkin will also take customized orders for specialized applications and network configurations.

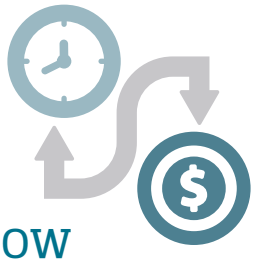

SLOW NETWORKS CAN BE COSTLY

A report by Aberdeen Group found a one-second increase in response time reduced conversion rates by **7%**, page views by **11%**, and customer satisfaction rates by **17%**.

Source: *When Application Tuning Becomes Network Optimization*, Dan Sullivan, February 2013

GENERAL **NETWORK ASSESSMENT**

Our engineers can provide a tailored assessment of your current network. They can help determine how it's performing and what areas need improvement. You will receive a written report that clearly details our recommendations to improve your network's speed, scalability and security.

GET STARTED AT CDW.COM/ NETWORKING-SOLUTIONS

IN TRANSIT: BUILDING RELIABILITY AND BANDWIDTH

Strategies to keep growing network traffic moving

Cloud services mean two critical things to network managers: more bandwidth and an increased emphasis on reliability. Apps are migrating from decentralized branches back to the central data center, and eventually to a hosting center or to a third-party cloud provider. Users need to get to those apps, and they need to get to them reliably.

Cloud services aren't the only processes demanding additional bandwidth. As organizations embrace BYOD programs, one direct effect is more devices on the network.

For example, a branch office worker who started with a desktop computer might now also carry a smartphone and a tablet, both of which use enterprise network resources. All three devices require Internet and WAN bandwidth. And when those BYOD users go home or are on the road, they need access to enterprise apps – especially collaboration tools such as email and VoIP – which requires more bandwidth.

Increasing Bandwidth

The choke point for many networks is not the LAN in and around the data center, but the WAN between the data center and distributed users. The WAN, whether it's a private Multiprotocol Label Switching (MPLS) network or the public Internet, becomes especially important when cloud apps are being deployed.

This is because users expect to access those services farther away from the data center. Simple answers, such as bumping a link speed between buildings or doubling up links, won't work if offices are widely geographically dispersed or if the app is being served from a provider via the Internet.

Obviously, one option is to simply call up WAN service providers and ask them to increase bandwidth between locations. But that can get pricey – and worse, it's a continuing expense, not a one-time cost. It might not initially appear to be very costly, but that expense, for a hundred offices, for 12 months – it adds up to a heavy expenditure.

Network managers have technology they can apply to optimize bandwidth use. Several product niches have come together, including WAN optimization controllers (WOCs), application delivery controllers and firewalls, to help improve the overall user experience without directly increasing speeds.

The solutions can be divided into two broad categories: private and public.

Private solutions require that users and apps all run on the same enterprise network, and that the organization maintains control of the network. This might involve running VPN tunnels over a public network such the Internet, having a traditional private WAN based on technologies such as MPLS or private Ethernet, or contracting for a service provided by a private network operator.

Private networking solutions improve the user experience by:

- ☑ Compressing data streams by reducing overall bandwidth used
- ☑ Caching data by maintaining a store of recently requested data objects, such as files or email attachments, at the remote end of the WAN for quick transmission when a cached object is requested
- ☑ Optimizing protocols (such as TCP or UDP) to make more efficient use of bandwidth-limited or congested circuits

Public solutions are broader in scope and don't require equipment or software at both ends of the connection. These solutions can perform application performance management and bandwidth allocation or quality of service. They can also give visibility into network flows and control application usage to provide a better user experience.

Public solutions can also modify applications on the fly by:

- ☑ Making smarter use of objects such as JavaScript, which reduces retransmission by enabling the user's own browser to cache the object
- ☑ Compressing and optimizing content and images, which reduces bandwidth used for handheld devices by reformatting images, and otherwise compressing any other data streams — taking advantage of the built-in compression features of most browsers
- ☑ Interleaving HTTP connections and enabling HTTP extensions

>

MAKING THE MOST EFFICIENT USE OF BANDWIDTH

Technique	Most Common Solution Category	Additional Solution Categories
Data compression and reduction	WAN optimization controllers (WOCs)	Some functionality may be available in web security gateways, but pure data compression and reduction are not often found in other product categories.
Application optimization	Application delivery controllers (ADCs) and load balancers	WOCs often include some application optimization features. Web application firewalls are a separate niche unto themselves.
Traffic prioritization and bandwidth management	Quality of service and visibility products	WOCs often include traffic prioritization. Unified threat management (UTM) and next-generation firewalls often include basic bandwidth management and prioritization.
Visibility (based on IP Flow Information Export/ NetFlow protocols)	Stand-alone IPFIX/ NetFlow exporters	High-end WOCs, firewalls and routers all generally act as IPFIX/NetFlow exporters. In most cases, no stand-alone exporter will be needed.
Routing and link balancing; using multiple links	Branch and edge firewalls or combination router/VPN devices	Stand-alone edge routers tend to have this capability, but the location of the router outside of the firewall (where it's unable to see into encrypted VPN tunnels) pushes this feature to whatever device is handling the VPN for the branch.
Load balancing to servers	ADCs	Some firewalls offer this feature, although rarely with enterprise-class capabilities.
Security	Unified threat management and next-generation firewalls	Web security gateways and proxy servers may include limited web-focused features. Stand-alone intrusion prevention systems can be used, but are rare in a branch setting if UTM or next-gen firewalls are available.

*Additional information about optimizing WANs can be found in the CDW white paper, *The High-Performance WAN*, CDW.com/wan-whitepaper

to increase parallel processing of HTTP streams

Ensuring Reliability

Building out more bandwidth in enterprise networks is not simple, but it's not rocket science either. It requires engineering and technology work in tandem to increase performance.

In challenging environments, network managers may have to turn to expensive technologies such as very small aperture terminal (VSAT) satellite networks, but adding speed to a network is fairly straightforward. Making sure that a network is reliable, on the other hand, often requires some inspired brainstorming with engineering. Network managers have to bring a variety of clever configurations, advanced technologies and clear thinking to the table to enhance overall network reliability.

One effect of cloud computing is the requirement to increase overall network reliability in WANs and LANs. The use of multihoming Internet service and tools that make use of the Border Gateway Protocol (BGP) can improve Internet reliability in large locations.

But to increase overall network reliability, network managers should consider implementing Ethernet delivery for WAN services, building resilient LANs and using dynamic routing between WAN sites.

A NETWORK WITHOUT SPANNING TREE IS MORE SUSCEPTIBLE TO FAILURE AND SERVICE INTERRUPTION THAN ONE THAT HAS IT ENABLED.

Moving to WAN Ethernet

Some developments in bandwidth provisioning favor reliability. Older WAN technologies, such as T1/E1 circuits (1.5Mbps to 2Mbps point-to-point lines) are being abandoned as quickly as possible. Because most organizations have more than two locations, carriers are using MPLS service to deliver a replacement VPN service over inexpensive Ethernet links.

Whatever the technology on the carrier side, organizations are increasingly using EtherLoop (essentially, high-speed Internet service over Ethernet) as their point of connection to public and private network services, rather than serial ports at T1 or T3 speeds.

Ethernet delivery of WAN services has significant advantages over serial point-to-point technologies, including:

Decreased cost | The hardware cost for an Ethernet endpoint is much lower. A T3/E3 card for an edge router can run $5,000, while Ethernet ports might cost only a few hundred dollars.

Failover support | Tools that use the high-availability Virtual Router Redundancy Protocol (VRRP) can operate over Ethernet and make it easy to cluster edge routers for failover.

Easier upgrades | Ethernet line rates of 10Mbps, 100Mbps, 1Gbps and even 10Gbps offer an easy growth path. Bandwidth upgrades from carriers are usually just a phone call away and don't require reengineering.

Security support | Security appliances such as intrusion prevention systems, data loss prevention and antimalware are easily installed in-line or as taps in Ethernet circuits if needed.

Decreased footprint | Carrier equipment is often smaller for Ethernet, especially in high-speed services, saving space and power in remote branches.

High-reliability LANs and Spanning Tree

Many network managers might brush aside the notion of using the Spanning Tree Protocol as a reliability builder. In fact, more than a few organizations have disabled STP based on negative experiences.

As true as these experiences may be, a network without Spanning Tree is more susceptible to failure and service interruption than one that has it enabled. Without it, there can be only one path through a network. And if that path is interrupted in any way, the network breaks in two.

Before diving into newer versions of the protocol, it's worth defining some basic configuration rules that apply to all Spanning Tree versions. Because it doesn't have any real authentication, Spanning Tree configuration within the network must protect it against intentional or unintentional misconfiguration from outside the network.

STP Rule No. 1 | In a traditional multilayer network, edge ports where user devices or servers connect should block Spanning Tree. This action alone will solve almost all common STP problems. (This approach is commonly referred to as using a Bridge Protocol Data Unit (BPDU) guard and filter.)

STP Rule No. 2 | Spanning Tree generates a valid loop-free configuration for every network, but that might not be the optimal configuration. A proper STP implementation will identify the desired root and backup root switches in the network core, set the priorities properly (to 0 and 1, generally) and then block other switches from acting as roots. (This technique is known as using "root guard.")

STP Rule No. 3 | Various hardware, configuration, wiring and

software problems can cause STP to misbehave. Some of these are rare, such as bugs in the software, while others are more common, such as hardware misconfiguration. Switch features that guard against common problems, such as one-way links or overburdened or buggy CPUs, should be used when they are available.

These guidelines don't substitute for a solid Layer 2 network design, but network managers should make sure that they are following at least these suggestions no matter how simple their network might be. Turning STP on throughout a Layer 2 network is the first step toward achieving a solid high-reliability design to support virtualization.

STP has evolved considerably over the past 30 years, resulting in the current Multiple Spanning Tree Protocol (MSTP), which incorporates the older Rapid Spanning Tree Protocol (RSTP) that is also important for high-reliability configurations. Network managers should require 802.1s MSTP in all LAN switches. 802.1s is backward-compatible with both standard 802.1D STP, the version that most network staff think of as Spanning Tree, and the newer 802.1w RSTP, which speeds recovery in the event of a failure.

The key benefit of 802.1s MSTP is that it supports multiple Spanning Trees at once across the same infrastructure. Why is that good? Because multiple Spanning Trees on multiple VLANs support true multipath functionality.

For example, imagine a fully redundant set of links from edge 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 just waiting for a failure that might never occur. With the multiple Spanning Trees allowed in 802.1s MSTP, some VLANs can use one data path, through one set of switches and links, and other VLANs can use a different path, through different switches and across different links.

MSTP is not true active/active load balancing, and it's not dynamic routing. But it nonetheless spreads the load in a static way that provides better accommodation for peaks in traffic.

Organizations that are using VLANs for iSCSI SAN-based disk service,

>

NEXT-GENERATION SPANNING TREE

The Institute of Electrical and Electronics Engineers (IEEE) and the Internet Engineering Task Force (IETF) are both designing next-generation versions of the Spanning Tree protocols.

The IEEE approach, called Shortest Path Bridging (SPB), was standardized in 2012 as IEEE 802.1aq. The IETF approach, called Transparent Interconnection of Lots of Links (TRILL), was standardized in 2011 in a series of IETF requests for comment.

Both technologies build on older Spanning Tree technology, as well as lessons learned from MPLS, to create a new Layer 2 network interconnection technology that can solve some of the problems of STP. These include enabling multipath load sharing, supporting larger networks (those with more than 4,094 virtual LANs), becoming more resilient to configuration and operation problems, and speeding convergence times.

A few enterprise-class network manufacturers have begun implementing SPB, and several interoperability tests have shown successful interconnection of multiple vendors' network equipment. Although no vendors have implemented TRILL, both Brocade and Cisco Systems have proprietary versions.

As new technologies, SPB and TRILL are still being tested for applicability in enterprise environments. Because they are based on a dynamic routing protocol, they are potentially more complicated to configure than traditional STP.

It's worth noting, however, that the lack of commercial implementations doesn't really tell users how hard these new versions will be to configure once they're commonly available. Both have been designed to be more resilient to human error, such as plugging the wrong switch into the wrong port — the type of problem that has plagued STP implementations.

The jury is still out on whether SPB or TRILL will gain ground in enterprise networks. Network managers shouldn't be demanding them now, but they should keep an eye out for developments on both fronts.

virtual machine migration or backup-to-disk applications can push that bursty, high-volume data across another path in the network. This ensures that the iSCSI service (or vMotion or backup-to-disk) gets as much bandwidth as it needs, without affecting the rest of the network.

Using Dynamic Routing for Branch Offices

Branch offices can benefit from the same STP lessons that apply to primary locations. But WAN reliability in a branch office can't be accomplished using the same tools as in a large data center. It's usually not economical to double-up both carriers and connectivity in every branch office to compensate for the occasional failure.

Instead, network designers can turn to hybrid networks that combine private and public circuits, or to entirely public Internet services, to boost reliability. Hybrid networks that include Internet circuits are particularly appropriate when cloud-based apps are being delivered over the Internet by a service provider.

The key to increasing reliability in hybrid networks is proper use of internal dynamic-routing protocols, usually Open Shortest Path First or (in an all-Cisco network) Enhanced Interior Gateway Routing Protocol (EIGRP). By establishing dynamic-routing adjacencies between each branch office and the enterprise WAN — across multiple carriers — network managers can ensure survivability of a branch office's connection should any branch office circuit fail.

Firewalls used by small organizations often include some type of load-balancing or failover tool to move traffic from one Internet circuit to another in the case of circuit or connectivity failure. This is useful when all services are on the Internet, as in a single-office organization. But this approach doesn't work well in a case where services are both Internet-based and WAN-based, especially in an organization with multiple data centers or which is operating a private cloud.

If an organization has only one or two data centers, the complexity of dynamic routing to branch offices can often be avoided by using a small firewall's load-balancing capability to detect and recover from failures. But these small firewalls are generally very limited in their capabilities, so network managers should be careful not to grow dependent on a product feature that won't grow with the organization.

For instance, use of such a limited firewall cannot ensure reliability when enterprise apps are distributed across multiple sites or if branch offices employ tools such as VoIP, real-time collaboration or video conferencing directly between offices.

Instead, network managers should consider two common techniques for running dynamic routing over organizational WANs when both public and private connections are in use.

In one approach, enterprise dynamic routing integrates with dynamic-routing services provided by the organization's carrier — assuming that the carrier uses tools such as MPLS to keep its customers' networks separate. For Internet services, VPN tunnels establish dynamic routing and encapsulate traffic. In the other approach, enterprise dynamic routing is kept entirely separate from each carrier by encapsulating interoffice traffic in a VPN tunnel, even when private network circuits are involved.

The first approach (integrating with the carrier) maximizes available bandwidth (as nothing is wasted for VPN encapsulation over private networks) and leverages use of the carrier's own WAN. Effectively, the enterprise and carrier networks are merged, allowing the carrier to choose the best path between sites and outsourcing the job of dynamic-routing management and configuration over private circuits. Internet circuits still must be managed, tunneled and routed, but the enterprise network manager's workload is reduced and less complex.

The second approach (using tunnels even when not required) benefits the organization by allowing it to be completely carrier-independent. For example, if multiple carriers are in use, integrating with each carrier makes engineering and debugging troublesome.

Similarly, if the organization changes carriers every three to five years, integration turns into a problem that must be solved frequently, costing more than it saves in the long run. Tight integration between the enterprise and carrier networks also has the potential to create vendor lock-in, which reduces flexibility and negotiating power in an environment where long-haul telecommunications costs are constantly dropping. ■

HYBRID NETWORKS THAT INCLUDE INTERNET CIRCUITS ARE PARTICULARLY APPROPRIATE WHEN CLOUD-BASED APPS ARE BEING DELIVERED OVER THE INTERNET BY A SERVICE PROVIDER.

On The Road With CDW

When I'm on the court, I'm in the zone. I have to know where the ball is, at all times. In other words, I have to be totally awesome. Achieving total awesomeness is a lot to live up to. As an IT pro, you feel my pain. You feel it because you have a lot riding on your shoulders – data centers to run, networks to manage, glitches to find, mobile devices to protect, security issues to thwart. CDW can help you with all of that. It's like having one player who can do it all – so basically it's like having me.

Home or away, CDW and its partners can help you win. Visit CDW.com/barkley

CDW® PEOPLE WHO GET IT™

IMPROVING THE WIRELESS NETWORK

Increasing enterprise reliance on mobility requires a re-evaluation of wireless strategies.

Mobility is a beneficial, long-term technology trend that comes with some challenges for network management. Although mobility does not change the goals of an organization, it does affect how IT and network teams support the achievement of those goals. And some de facto rules about how to build networks, deliver computing and secure the environment must change when organizations embrace mobility.

Best practices of the past don't necessarily apply anymore. It used to be that the IT team operated under the premise that the organization owned, controlled and secured all the devices on the network. But initiatives such as BYOD turn some basic assumptions on their heads.

The final shape of a truly mobile workforce isn't fully defined. There are many issues left to work out over the next few years. Nonetheless, one of the core requirements for mobility is clear: Organizations will be expected to support secure and highly reliable wireless network access.

Today, for anywhere, anytime computing, standards and business models are competing to gain position as the dominant mobile approach. Inside office spaces and many other controlled environments, such as indoor public spaces, wireless adheres to one standard almost exclusively: IEEE 802.11.

Therefore, network managers need to understand how to deploy enterprise-class 802.11 networks and how to make intelligent management decisions about 802.11 availability, security and performance.

Planning Out Enterprise-class 802.11 Networks

With the increased performance and capabilities of these newest updates to 802.11, network managers now view wireless as a production network, rather than as a convenience tool for staff and guests roaming around the organization. Network managers now must help their IT teams understand that managing an enterprise-class wireless

AN 802.11 DECODER RING

Wireless LANs in organizations are all based on the IEEE 802.11 family of standards.

When a standard document name ends in a lower-case letter, such as 802.11b, it's an indication that the document is really an update or amendment and will be rolled up into the next version of the main standard. The latest version of 802.11, called IEEE 802.11-2012, is almost 3,000 pages long and incorporates the nearly 20 updates and amendments that came before it.

Despite the many changes, most network engineers are accustomed to talking about parts of 802.11 using the name of the update, even if it has been rolled up into a current document. For example, people will refer to 802.11a and 802.11b, neither of which technically exist anymore, instead of calling those elements the 5GHz band and the 2.4GHz band.

Network managers should memorize the six most important amendments (all but 802.11ac are rolled into 802.11-2012), which are listed in this chart.

802.11a 1999	802.11 operating in the 5GHz band, which greatly increased the number of channels; now 21 nonoverlapping in the United States, although channels have been added or deleted; documents written before October 2010 may refer to 23 channels
802.11b 1999	The first usable version of 802.11, an 11Mbps wireless network on the 2.4GHz band; three nonoverlapping channels in 802.11b in the United States and Europe
802.11g 2003	A modulation change for 802.11b, which increased the speed from 11Mbps to 54Mbps, retaining the 2.4GHz band
802.11i 2004	Fixed security problems in the original 802.11 and brought 802.11 into alignment with enterprise network security requirements; Wi-Fi Protected Access (WPA) and WPA2 security profiles, taken from 802.11i, were made industry standards by the Wi-Fi Alliance
802.11n 2009	Another modulation change for 802.11; applies to both 2.4GHz and 5GHz bands; brings speeds of up to 600Mbps for wireless, although notebooks typically see 150Mbps to 300Mbps because of antenna restrictions
802.11ac 2013 (expected)	A speed bump for the 5GHz band, including 2x-, 4x- and 8x-wide channels; also a modulation change, more spatial streams (up to eight); speeds top out at 1.3Gbps, with notebooks typically maxing out at 450Mbps

LAN is a growing priority. As wireless networks become operational tools, the IT team must make sure that the quality and availability of the wireless signal meets enterprise standards.

Consulting on a Design

Designing a buildingwide wireless deployment requires a specific set of radio-frequency engineering skills that the typical network manager may not have. That means either allotting resources for training or using third-party help to ensure proper access point placement, which is a requirement for achieving a high-performance, high-reliability wireless network.

When working with a consulting team, the organization's network managers need to be prepared to hand over a set of requirements and building maps showing offices, meeting rooms and dead space. The consulting team, in turn, will craft a set of maps illustrating access point (AP) locations, along with details on managing radio frequency (RF) channels and levels. The network plan will also lay out wireless coverage for each floor of each building.

A site survey done by a consultant need not be overly complex, and the requirements don't have to stretch for pages and pages. For example, they might be as simple as "Wireless

>

NETWORK MANAGERS NOW VIEW WIRELESS AS A PRODUCTION ENVIRONMENT, RATHER THAN AS A CONVENIENCE TOOL FOR STAFF AND GUESTS.

8 COMMON 802.11 CONFIGURATION ERRORS *AND HOW TO FIX THEM*

With updates to 802.11, organizations do not need to scrap wireless networks that grew over time, access point by access point, to cover an entire building. But they do need to re-evaluate those networks to be sure that they meet organizational goals.

Once any gaps or performance issues are identified, the network team can set an action plan. Here is a list of common 802.11 configuration errors and quick fixes that can improve performance.

Error	Fix
Incorrect channel assignments in the 2.4GHz band	Use only channels 1, 6 and 11 (even in Europe), and position APs to minimize adjacency of identical channels.
Not addressing wireless's 3D nature	Wireless signals propagate through surfaces, so stagger APs on adjacent floors to maximize coverage and minimize overlap.
Too much or too little authentication	Integrate wireless authentication (based on 802.11i/WPA2) with enterprise directory and wireless guest services to ensure that users aren't frustrated trying to use the wireless network.
Incorrect power levels	Generally, turn down power and increase AP density (counterintuitive, but true).
Ignoring 5GHz band	Use the 5GHz band because it offers higher speeds and is much less crowded than the 2.4GHz band (802.11n and 802.11ac being the most effective in the 5GHz band).
Accepting the defaults for "minimum allowed speed"	Increase minimum allowed speeds. Beacon frames and minimum allowed connection speeds should be raised above 11Mbps to ensure edge devices don't degrade services networkwide.
Not taking mobile devices into account	Support physical mobility for tablets and smartphones, which users may tote around in the environment more than notebooks.

coverage should be set at a minimum RSSI of –67 dB in 90 percent of the offices and no worse than –75 dB anywhere in the coverage area. Density should be no more than 25 users per access point." As a reminder, site surveys for enterprise networks should be designed around the 5GHz band.

Active Management

Every major wireless network manufacturer now offers centrally managed and coordinated networks. In some cases, this means that 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. There is a global management system, controlling and tuning the access points, but traffic doesn't flow through a central point. In both instances, though, the central system manages RF settings, which is necessary to ensure quality performance for any large buildingwide wireless network.

From the moment a wireless network is installed, the physical environment of the building starts to change. People move around more frequently to carry out their work, and their physical bodies will affect network performance because bodies absorb wireless energy — quite well.

Over time, organizations also move desks, filing cabinets, bookshelves and other equipment (even walls). The wireless network must be tuned to handle these changes in the RF environment. And the only way to effectively keep performance high is to use a central management system to continuously monitor and automatically tweak the network in response to environmental changes.

Wi-Fi networks are amazingly fault-tolerant. They can often survive the loss of access points and the addition of interference without

significant effects. This means that small blips might go unnoticed without active monitoring. But as small changes occur, performance degradation will mount. Good network management practices, including regular scanning of logs and active monitoring of devices and usage, will help identify problems before they affect user performance.

Prioritizing and Segregating Traffic

Wireless networks share a finite amount of RF spectrum. An organization can increase performance by adding APs (and reducing power levels), but multiple users still must share the spectrum. Therefore, tuning to prioritize and segregate traffic will increase performance more efficiently and at less cost.

Mission-critical services, such as VoIP or transaction processing applications, should take priority over nonwork and casual usage. Firewalls and access points themselves can throttle bandwidth based on service, while tools such as Wi-Fi Multimedia help ensure performance of VoIP traffic.

Organizations typically use multiple service set identifiers (SSIDs), essentially network names, to segregate traffic by defining different service levels. For example, most networks will have a "guest" SSID for unauthenticated or lightly authenticated guest users.

This level of network service provides convenient Internet access to visitors and other third-party users on organizational campuses. Networks also typically have a "production" SSID that provides staff members with wireless network access that is similar to their wired service.

The security profiles of these two SSIDs will differ dramatically. The guest network will have almost no wireless security (although a firewall is expected). The production network will have WPA2-Enterprise security, authenticating users against the organization's central directory (such as Active Directory) and ensuring encryption of all traffic.

If the organization allows BYOD, additional SSIDs might be appropriate. For example, some organizations set up a special network to authenticate and provide encrypted access for staff members' personal mobile devices but don't provide end-to-end access to the organization's network assets.

Picking the Right Hardware

Today, 802.11n is the best, highest-performance wireless solution for organizations. Because 802.11n is a multiple-antenna, multiple-stream standard, user devices and access points can utilize from one to four streams.

Most devices supporting 802.11n have two streams and two antennas (2x2:2) or two streams and three antennas (2x3:2). These are limited to a top speed of about 130Mbps (or 270Mbps if double-wide channels are used in the 5GHz band). Newer devices are becoming available with three streams across three antennas (3x3:3) and a top speed of 195Mbps, or 405Mbps with doublewide channels.

Because access points need to be long-lived, it's worthwhile purchasing APs that support three antennas and

>

THICK *VERSUS* THIN

There is some general confusion over the relative value of "thick" versus "thin" access points.

Network managers may even have heard that thin APs are less expensive. It's a myth: Thin APs are about the same price (or may even be more expensive) after factoring in the cost of the controller. It's also a myth that thin APs are a prerequisite for a fully managed solution.

From an enterprise perspective, an organization needs a fully managed solution that supports mobility services and optimizes wireless RF networking continuously. Whether the APs are thin, thick or something in between is really irrelevant.

There are important differences among different manufacturers' fully managed wireless solutions, and different fully managed thick and thin architectures are appropriate for different types of network topologies. One major difference between thick and thin AP products is in the amount of traffic backhauled to the management controller. For some networks, bandwidth between wiring closets and data centers is essentially unlimited, but in other environments there may be bottlenecks that require a certain architectural solution – even though all are fully managed.

Therefore, the decision between thick versus thin hinges on the site analysis and an evaluation of the specific dynamics of the organization's environment.

CASE STUDY

WIRELESS CATCH-UP

Read about how several companies upgraded their wireless networks to meet growing user expectations:

CDW.com/networkguide1

three data streams to avoid an upgrade. When installing new devices, or upgrading an existing network, use 3x3:3 access points whenever possible.

The new 802.11ac standard will increase data rates (in the 5GHz band only) above 1Gbps by using even wider channels (four-wide 80MHz and eight-wide 160MHz channels will be available) and up to eight antennas per access point. The new standard is only available in consumer-level devices at this time because one of the big targets for 802.11ac bandwidth is multiple high-definition video streams, typically a home application. Organizations will be able to enjoy the benefits of 802.11ac as well, even in lower-speed apps, because lower-speed (subgigabit) devices will use less power for the same throughput level.

The upcoming 802.11ac upgrade shouldn't be a reason to put off a planned wireless project. Devices can be rolled in gradually as availability increases and pricing levels off. New 802.11ac devices will slowly filter into the product lines of enterprise wireless makers in 2013 and 2014. Fortunately, AP densities for 802.11ac should be about the same as for 802.11n in the 5GHz channel.

Organizations planning to upgrade their wireless networks in 2014 should design around 802.11n topologies but be prepared to swap out chosen AP models for 802.11ac devices at the time of purchase.

One aspect of wireless deployment that 802.11ac emphasizes is picking the right switches to connect APs. Wireless users should be given the same (or higher) speeds as wired users: 1 Gig-E with 10 Gig-E uplinks to the network core. Plus, the network should use Power over Ethernet (PoE) everywhere to ensure maximum flexibility. ■

NETWORK SECURITY: DEFENSIVE ADAPTATION

Flexible strategies are key for protecting against evolving threat vectors.

Network security requires constant re-evaluation of the security posture as the threat, risk and regulatory landscape changes. For most medium to large organizations, security is a nonstop cycle of product evaluation, purchases and replacements; setting and refining policies; patching and reconfigurations; and the adoption of new procedures to accommodate the evolution of the environment and technology.

The main constant is change. Therefore, the best approach is one of adaptability.

BYOD, consumerization and mobility all require significant adjustments in enterprise security. These three trends mean that network and security managers should rethink some basic assumptions and ask themselves difficult questions:

- Can the organization control the security posture of every device on the network?
- What are the real assets of the organization that need to be protected?
- Where are the right places for security controls in the network?
- Which of the security policies and practices have become outdated and need to be jettisoned, and which continue to be valuable and should be retained?
- How does the organization effectively communicate changes in security up the management chain and across the organization?

Strategies for Threat Mitigation

Network and security managers should take the opportunities presented by these new technology developments to revisit inconsistencies in their security deployment. For example, an organization may have very separate approaches to securing the main office LAN, branch office network and wireless network, leading to dangerous vulnerabilities. As BYOD and mobility are spread through the network, consistency should be a goal for security reengineering.

>

Although there are many new security tools, techniques and issues specific to BYOD and mobility that will be addressed here, organizations should not let the routines of patching, antimalware, intrusion protection and antispam technologies fall by the wayside. New zero-day and advanced persistent threats receive a lot of press and are real concerns for organizations. But research has shown that it's old threats, rather than new ones, that are most often to blame when systems are compromised by malware.

4 MYTHS OF THREAT MANAGEMENT

Myth	Reality
1. Intrusion prevention and detection systems block most intrusions.	IPS and IDS provide important visibility into the security posture of a network and can block occasional intrusions. IPS improves visibility and reduces the size of the patch window.
2. Antimalware at the edge of the network is most effective.	Edge antimalware provides secondary and tertiary protections. Desktop antimalware and Layer 7–aware antimalware (example: a typical antispam gateway with included antimalware) are the first lines of defense for most intrusions.
3. Malware poses the greatest threat.	Viruses come in last in the statistics for actual data breaches. What's No. 1? Lost, stolen or improperly disposed equipment. Phishing, misaddressed email and leakage of important information are all more significant threats to end users. Application-layer vulnerabilities (such as poorly written apps) that attackers can easily exploit are the biggest problems.
4. Zero-day threats are a huge problem.	Known threats wreak the most havoc: According to Microsoft research, 82 percent of malware found on desktop PCs results from threats reported six or more months earlier.

Mobile Device Management

BYOD initiatives take different forms. But generally, they involve allowing workers to use their own personal mobile devices — notebooks, tablets and smartphones — to carry out their job duties in some capacity.

As BYOD programs have grown and taken hold in large organizations, network staffs have adopted mobile device management (MDM) strategies both to monitor and track the devices, as well as provide security to enterprise systems and data. MDM doesn't prevent the potential for data loss, but it does give an organization tools to help enforce security policies aimed at reducing the risk of loss.

For example, MDM can ensure the encryption of data on the BYOD devices and enforcement of strong passwords. MDM also can mitigate the repercussions of a device being lost or stolen by allowing the organization to remotely wipe all or some of the content on BYOD systems.

The market for MDM is quickly becoming a stand-alone business, separate from endpoint protection products (although both McAfee and Symantec compete with solid offerings). The unique requirements of devices running Android, Apple iOS, BlackBerry OS, Symbian and Windows 8 have segmented device management tools into smartphones and tablets on one hand, and desktops and notebooks on another. Now that Windows 8 is becoming widely available, the Windows tablet and notebook markets will begin to merge from an MDM point of view.

But Windows 8 represents just a tiny fraction of devices in 2013, with BYOD demand really coming from users with Apple iOS and Android devices. Although the same issues of security are present on all platforms, network and security managers, at least for now, will generally need separate products. They will need one to cover Windows 8 (and earlier versions) and another to handle mobile operating systems such as iOS, Android and BlackBerry OS.

Network and security managers can choose either cloud-based or on-premises MDM platforms, depending on their preferences and organization size. Smaller organizations or those with only a few devices should definitely lean toward SaaS for a cost-effective solution, while larger organizations might prefer MDM appliances (both virtual and physical) and standard software.

The main goal of MDM is to implement policies for mobile devices, with a secondary interest in reducing support costs and simplifying deployment (especially for organizations that supply their users with devices). In some cases, MDM and endpoint protection are merged into a single tool.

Network Access Control and BYOD

Like many technologies, network access control tools came to market ahead of their time. They have been slowly gaining acceptance as organizations discover what NAC can (and can't) do for them.

NAC was first designed in 2004 to solve a particular problem: Enterprise security managers were often led to believe that their desktop systems were fully patched and up to date with antimalware – but they weren't. A few high-profile viruses swept across the Internet, causing many organizations to suffer extended network downtime. The problem wasn't that the malware was zero-day; rather, the systems were not sufficiently patched and updated.

Using NAC to check systems for security posture works well for systems completely under enterprise control, but it's a poor approach for nonenterprise devices that might be connected to the network. NAC vendors have worked hard to create universal browser plug-ins to enforce policies such as "if any known antimalware is installed and updated, then this system is considered compliant," but the success rate for these tools has not made NAC a popular tool.

But NAC is useful for much more than checking security posture. These tools include two other key components: end-user authentication (typically, using 802.1X through a wireless WPA2 authentication) and per-user access control. In the world of BYOD, posture checking has become much less important than the other two NAC capabilities that can identify who the user is and what access they should have to the network.

Most organizations considering BYOD are not ready to let foreign devices, even those owned by their employees, onto the most sensitive parts of their networks. Barriers can be created in a variety of ways. For example, tools such as virtual desktops and remote terminals can isolate specified enterprise systems and applications from mobile devices.

APPLYING MDM TO BYOD

MDM Tool	BYOD Applicability
Enforces device configuration, such as patch levels, OS versions, application whitelists and blacklists; manages device software, including installation, upgrades and removal; provides alerts or blocks usage when devices are noncompliant	BYOD users typically wouldn't expect the IT team to manage software, but application blacklists may be an exception. Minimum OS and patch levels could be used to block out-of-date user access.
Backup and restore	BYOD users and the IT group will usually agree that backup of personal data is not the organization's responsibility.
Control of device security policies, such as configuration of endpoint protection tools, device locks and password rules	BYOD policies may require that devices conform to organizational policies at all times; MDM can both configure and enforce these policies even when devices are not on the network.
Configuration of enterprise applications, including email, VPNs and hotspot access; control of wireless communications, such as whitelisting and blacklisting, SSIDs and insecure Wi-Fi usage	BYOD support has the potential to be a huge black hole for organizations, no matter what the policy states. By autoconfiguring applications, MDM can reduce costs and help mitigate the threat of insecure usage, such unencrypted wireless.
Remote wipe	Remote wipe of personal devices can pose a legal liability issue, yet it represents one of the best tools for dealing with the most frequent mobile device security risk: loss or theft.
Data security policy enforcement, including local encryption of enterprise data and encryption of all network communications	As with device security, BYOD policies should require conformance to enterprise rules regarding data security. MDM is a key tool to configure and enforce data security.

Some organizations allow access, but treat mobile devices as only slightly more trusted than general Internet traffic. In these cases, the network needs a way to distinguish enterprise-owned devices from staff personal devices (and from notebooks, tablets and smartphones brought in by visitors, which might be completely untrusted). NAC solutions can make such distinctions, based on authentication and other information, and can also create access controls in the network to keep multiple groups of users separate from one another.

Most NAC products fall into a category called "preadmission," meaning that a device is authenticated and access controls are applied at the moment the device connects to the network. These NAC products are the most secure and provide the

>

most foolproof methods of access control. Preadmission NAC products are generally proactive, limiting or permitting traffic based on policy.

"Postadmission" NAC products kick in after a device user tries to perform a prohibited action. In such a case, the product may intercept a browser request and display a captive portal or block traffic entirely.

The advantage of postadmission NAC products is that they are designed to overlay existing networks and, if they're compatible with a network, usually require few intrusive changes. Postadmission NAC products are generally reactive, stepping in only when certain boundaries are crossed.

There's considerable debate in the security industry over the two NAC types. Preadmission NAC in the wireless world has been successful because the 802.1X authentication built into WPA2 is generally highly compatible across different devices, access points and authentication servers.

Whether NAC is right for an organization (and if so, which type of NAC) is a difficult question to answer and will depend on both management and budget resources. But the pressure of BYOD has caused NAC to reemerge as a key technology for organizations grappling with mobile security.

Next-generation Firewalls

Next-generation firewalls shift the focus away from protecting servers to protecting users. They work by widening the access control rules. Instead of simply defining access as source and destination IP addresses and ports, next-gen firewalls let security managers control access based on apps and users. These features raise the firewall up from TCP/IP network concepts (such as port number and IP address) so that the emphasis is on meeting the goals of the security team.

In some ways, next-gen firewalls and NAC are similar technologies because they both have a user-focused security profile. In fact, some security providers have elected to enforce their NAC policies using their next-gen firewall products, so the overlap is a natural one.

IT groups undertaking a BYOD deployment can make good use of next-gen firewalls to provide fine-grained access control for untrusted devices. Although next-gen firewalls are primarily designed to protect users against attack, they can also effectively become application whitelisting and blacklisting tools to enforce policies on semitrusted subnets (wired and wireless) inhabited by BYOD users. ■

WILL **WINDOWS 8** MAKE *ANTIMALWARE OBSOLETE?*

Microsoft Windows 8 continues the antimalware security improvements that Microsoft has been gradually integrating into Windows for the past decade.

Address space layout randomization is used to randomize the layout of portions of an application, which is an effective defense against many buffer overflow attacks. Similarly, data execution prevention works with features of the system hardware to prevent an app from suddenly executing the data portion of its memory. This reduces app instability (or at least catches it sooner) and also eliminates a trick used by many attackers.

Windows 8 also incorporates other security features, including application sandboxing (creating what's referred to as "app containers") and enforcement of a fine-grained mandatory access control system for files, services, registry keys and other operating system objects. Microsoft, which calls this the Windows integrity mechanism, first introduced this approach in Vista. All of these features are designed to increase overall security, but many have been specifically aimed at problems associated with untrustworthy software and malware.

Whether these changes finally free desktop managers from installing antimalware products depends on whether the IT staff believes that antimalware products are obsolete. Certainly, the case can be made that signature-based antimalware has lost the battle against the flood of threats to Windows OSs.

Although security vendors and analysts offer different estimates, all agree that the number of threats (in the millions) are growing dramatically and that pure signature-based antimalware can't keep up. Obviously, there aren't millions of pieces of malware, which means that heuristic (rather than signature-based) detection is the path forward for these tools. Even so, Windows 8 comes with Windows Defender, an antispyware/antimalware tool, which indicates that Microsoft believes that antimalware still provides a layer of protection.

Index

ABOUT THE CONTRIBUTORS

NEAL CZAPLEWSKI is the director of CDW's Network Solutions practice. He has more than 15 years of experience in the IT industry with various roles in consulting, management and technical sales. Neal's primary responsibilities include leading the strategy and execution for CDW's national networking practice and data center solution architects.

JOEL SNYDER, Ph.D., is a senior IT consultant with 30 years of practice. An internationally recognized expert in the areas of security, messaging and networks, Dr. Snyder is a popular speaker and author and is known for his unbiased and comprehensive tests of security and networking products. His clients include major organizations on six continents.

DAN VARGAS is the Lead Solutions Architect for CDW's Network Solutions practice. He is a Cisco Triple CCIE, holding an IE in Route/Switch, Security and Voice. He has more than 15 years of experience in IT, with various roles in consulting, engineering and technical sales. Dan's primary responsibilities include setting the technical direction for a national practice of network field solution architects, as well as determining the networking equipment and demo execution for technology labs at CDW.

LOOK INSIDE FOR MORE INFORMATION ON:

- Accommodating increased mobile traffic on the network
- Optimizing network performance
- Deploying a comprehensive network security strategy
- Updating wireless network bandwidth and reliability