VIRTUALIZATION AND INFRASTRUCTURE OPTIMIZATION

REFERENCE GUIDE

Innovating for greater efficiency and flexibility, with lower cost and complexity.

CDWG.com/virtualizationguide | **888.693.4239**

CDW·G®

The Right Technology. Right Away.®

VIRTUALIZE

VIRTUALIZATION AND INFRASTRUCTURE OPTIMIZATION

REFERENCE GUIDE

TABLE OF CONTENTS **CHAPTER**

WHAT IS A CDW•G REFERENCE GUIDE?

At CDW•G, we're committed to getting you everything you need to make the right purchasing decisions — from products and services to information about the latest technology. Our Reference Guides are designed to provide you with an in-depth look at topics that relate directly to the IT challenges you face. Consider them an extension of your account manager's knowledge and expertise. We hope you find this guide to be a useful resource.

» We are pleased to also offer the **CDW•G IT Investment Guide.** It includes products and information to help you meet your virtualization and infrastructure optimization objectives.

NETWORK INNOVATIONS

IT managers have always been on the lookout for ways to make their operations more efficient, secure and flexible. This desire for improved performance has been a strong motivator among IT manufacturers to develop pertinent technologies. New, innovative technologies are released at a constant pace, giving IT managers and organizations many options for how to go about improving their systems.

One area of strong growth in particular was the explosion of data centers in the 1990s. This expansion helped lay the groundwork for organizations' embrace of and reliance on digital communications and operations today. But this expansion was not without its drawbacks. Supporting a sprawling physical infrastructure entailed high costs, not only in terms of hardware, power and cooling, but also with regard to management and maintenance.

In addition, many organizations heavily underutilized their servers — on average, as low as 15 percent of their total capacity, according to the International Data Corporation (IDC). These machines often lacked sufficient disaster or fail-over protection. And they frequently lacked security.

VIRTUALIZATION'S ARRIVAL

By 1999, these shortcomings became untenable for many organizations. This server sprawl began to be addressed through virtualization, a method for dividing up a computer's resources into multiple environments to create a flexible, easy-to-manage, secure computing system.

Enterprising IT managers began to see the potential for using server virtualization technology across other areas of the network. While server virtualization remains the most common type of virtualization, this technology increasingly has applications across many other facets of IT.

Virtualization can improve many of the components that comprise today's computer environment. Two of the most important areas where virtualization can make a significant impact are servers and clients (desktops, notebooks, etc.).

THE VIRTUES OF VIRTUALIZATION

IT managers have embraced virtualization with enthusiasm primarily because it reduces costs and complexity. With newer focal areas such as client virtualization now reaching maturity, IT managers are finding ways to employ virtualization throughout their organizations.

Virtualization additionally allows for the separation of resources, or dependencies. As a result, it enables the control and movement of other resources to different locations in order to enhance the versatility and performance of the overall system. Different types of virtualization can be used in a variety of scenarios to deploy resources where they're needed within the organization.

This guide will provide information to assist you in making educated decisions as you move forward in helping your organization to achieve marked systems improvements through virtualization.

1A. VIRTUALIZATION: HARDWARE TO SOFTWARE

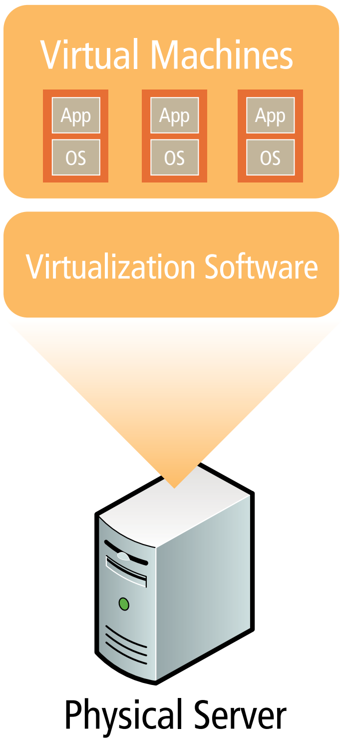

SIMPLIFY VIA OPTIMIZATION

Another recent innovation that has helped organizations improve the performance of their systems is optimization. For many years, systems standardization had IT managers stuck in a rut. Because the overall cost of support rises with each type of device added to a network, the idea of standardizing organizations to between one and three models of PCs, servers and networking equipment became quite popular.

This approach certainly has its merits in running an IT department; it helps keep costs down. However, this approach also began to present certain challenges. Every user benefits from having an individual PC, but the number and types of applications used on those machines begins to quickly proliferate.

Furthermore, every network application gets its own server, regardless of how much server utilization is required. The resulting complexity of cabling, networking, management tools, space, power and cooling begins to spiral out of control. It ultimately becomes a drain on the energies of IT personnel and an organization's finances.

Enter two of the more successful solutions in IT: virtualization and optimization. Optimization refers to the overall goal of doing more with less. (Virtualization is one means of achieving that goal.) Using a much smaller set of hardware components (via virtualization), optimization can be applied to PCs, servers, storage or networking traffic.

OPTIMIZATION'S BENEFITS

Many organizations are now realizing that the benefits of optimization and virtualization far outweigh keeping systems separate. An optimized IT environment provides a simpler backup and disaster recovery (DR) methodology to keep your organization safe, a much lower power consumption rate to keep your organization in the black, a simpler suite of data center management tools to help an organization stay on top of all of its IT assets, and a better experience for end users.

Fifteen years ago, IT managers faced a revolutionary change in moving from mainframes with attached terminals to single application servers with networked PCs. Today, IT confronts a change of similar magnitude: abandoning the old way of doing things and embracing the "data center 3.0" methodology of a highly optimized, highly available and highly flexible network.

WHILE SERVER VIRTUALIZATION REMAINS the most common type of virtualization, this technology increasingly has applications across many other facets of IT.

Now is the time to look at every application, piece of hardware and IT process to determine what can be done differently, more efficiently and ultimately more effectively.

This guide will introduce you to best practices for optimizing your data center, server and storage infrastructure, and network using mainstream techniques that have a tremendous return on investment (ROI). However, as new optimization techniques are constantly being advanced, keep in mind that the information provided here represents only a small glimpse of the solutions available. ◊

SERVER
VIRTUALIZATION

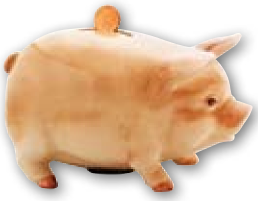

Among the different types of virtualization, server virtualization is the most widely adopted of these technologies. It allows for the virtualization of multiple application servers onto a single physical server system.

What makes this approach such a boon for IT managers is that it reduces the number of actual servers in the data center while at the same time increasing their utilization. In doing so, it presents solutions to many of the challenges that IT departments face today.

SERVER PROLIFERATION PROBLEMS

The IT industry has evolved dramatically over the past decade, and with it the way that organizations operate. Organizations gained access to greater technological capabilities through inexpensive x86 server systems, as well as the applications and operating systems that run on this platform.

However, the rapid adoption rates of these server systems have resulted in myriad difficulties for many organizations, each of which virtualization can potentially remedy. These challenges include:

- Increased total cost of ownership (TCO)
- Server sprawl
- Low server utilization
- Inefficient server hardware migration
- Inefficient application server deployment
- High-availability complexity
- Disaster recovery complexity
- Licensing complexity
- Power and cooling inefficiency

INCREASED TCO

Perhaps the most common requests heard by IT departments are to lower the TCO and achieve a faster ROI for all new IT purchases. These demands, coupled with stringent high-availability and disaster-preparedness requirements, have made cost reduction a major challenge for CIOs and IT managers everywhere.

SERVER SPRAWL

Data centers have come to rely on x86 servers to support organization growth and today's more widely distributed enterprises. This proliferation owes a great deal to the concept of distributed computing: the ability to spread server application loads across multiple tiers, different platforms and locations.

Another reason for this increased use has been the requirement by many of the application manufacturers to have applications deployed on dedicated hardware, and in some cases, dedicated database and operating systems.

LOW SERVER UTILIZATION

Industry-standard x86 systems continue to be the best-selling server platform. Processor, memory, network and disk speeds continue increasing while technologies such as multicore processor architecture continue to improve hardware performance. However, most operating systems and applications are not able to make use of these performance increases.

This is why the average server utilization rate today is under 9 percent, according to industry analysts. Low utilization on an endless group of servers has made many organizations rethink their IT strategy and implement a control structure to avoid server sprawl.

INEFFICIENT SERVER HARDWARE MIGRATION

The typical data center replaces servers every three years. Although replacing servers may seem like a straightforward process, it can be quite time-consuming, difficult and expensive. With each operating system tied directly to the hardware, it gets progressively more difficult to migrate to newer systems.

Also, applications can be tied to a particular named instance of the operating system, as well as the hardware. This makes each infrastructure refresh cycle an unattractive proposition for the data center, operating systems and applications management teams. Even though newer physical-to-physical technologies exist today, they are often not cost-effective and offer unpredictable results.

INEFFICIENT APPLICATION SERVER DEPLOYMENT

Application servers are continually added to enhance an organization's operations. However, lead times for procuring the hardware and software, performing testing and development, and conducting proof-of-concept modeling, implementation and end-user training can sometimes take months to complete.

Although a number of server deployment technologies are available, they can be very expensive to purchase and implement. Consequently, they are not in widespread use today.

HIGH-AVAILABILITY COMPLEXITY

High-availability is a series of measures undertaken to implement minimal to near-real-time fail-over for a particular application. Due to the variety of application architectures and operating systems available, high-availability can be challenging to implement.

As the size of the data center grows, it becomes increasingly more complex to make the infrastructure highly available, and thus more costly to implement and maintain.

Most organizations do not need every system to be highly available. Systems that serve the network backbone, such as directory services, file and print sharing, e-mail, and enterprise applications, fall into the high-availability category.

Determining the criticality of each application is the first step in creating a highly available infrastructure. This determination should be made by upper management (not IT) and incorporated into the organization's continuity of operations plan (COOP).

SERVER VIRTUALIZATION COST-SAVING EXAMPLE

PHYSICAL VERSUS VIRTUAL

In order to illustrate the cost benefits of implementing a virtual infrastructure, the following example compares a physical server versus a virtual server implementation.

The two solutions developed for this example are presented in detail. A number of assumptions were made while developing this model.

VMware Infrastructure 3 Enterprise Edition and VirtualCenter are used to develop the cost model for the virtual solution. Proper software licensing is still required for all operating systems and applications, whether implemented on a physical or virtual machine.

Because of the functionality and flexibility of VMware ESX Server, this product is used as a basis to illustrate cost benefits. However, note that Microsoft Virtual Server and VMware Server could potentially yield a lower TCO because of their lower price points.

These assumptions are explained in detail in CDW•G's server virtualization white paper, which is available at: http://webobjects.cdw.com/webobjects/docs/pdfs/Server-Virtualization-Small-Business.pdf

SOLUTION A: PHYSICAL SERVER

The physical server model presented includes:

- 32 two-socket servers
- 2GB of memory attached to a redundant storage area network (SAN) and network core
- 64 Ethernet cables and 64 Fibre Channel cables (required for the redundant network and Fibre Channel connectivity)

SOLUTION B: VIRTUAL SERVER

The virtual server solution consists of the following:

- Three four-socket servers (32GB of memory, VMware Infrastructure 3 Enterprise Edition)
- One VMware VirtualCenter Management Server
- 20 Ethernet cables and eight Fibre Channel cables (required for redundant network and Fibre Channel connectivity)

CH 02

HARDWARE/SOFTWARE COSTS

When comparing the cost to purchase 32 individual servers versus four physical servers with virtualization software, the total savings realized over a three-year period (the average life span of a server) amounts to $194,437 — as shown in Table 1.

OPERATIONAL COSTS

Table 2 shows the reduction in operational costs over a three-year period because of the reduced power and cooling footprint. Other excluded costs vary from the cost per-rack-unit of data center space, rack space, loaded weight, uninterruptible power supply (UPS) load, power distribution units, switches, etc. The total cost savings realized over a three-year period when implementing a virtualized solution is $20,995.

CONSULTING SERVICES COSTS

The services costs in order to implement the proposed solution are shown in Table 3.

REDUCED TCO THROUGH VIRTUALIZATION

The three-year savings when implementing a virtualized solution for this example are shown in Table 4.

After comparing the costs of implementing the virtual solution versus the physical, we arrive at a total cost savings of $203,432 over three years when implementing the virtualization solution.

Looking at this from a different perspective, the TCO from $390,320 is reduced to $186,888 with the virtual solution.

	Physical Solution (32 Servers)	Virtual Solution (4 Servers)	Cost Savings
Hardware/Software (Hw/Sw)	$300,768	$159,501	$141,267
KVM Cabling	$3,136	$98	$3,038
Fibre Channel Cabling	$1,920	$240	$1,680
Ethernet Cabling	$320	$200	$120
Backup Software (TSM/vRanger Pro)	$56,000	$7,668	$48,332
Total 3-Year Cost	$362,144	$167,707	$194,437

Table 1 – Comparison of Hardware/Software Costs

	Physical Solution (32 Servers)	Virtual Solution (4 Servers)	Cost Savings
Power Costs (3 Years)	$25,691	$6,548	$19,143
Cooling Costs (3 Years)	$2,485	$633	$1,852
Total 3-Year Cost	$28,176	$7,181	$20,995

Table 2 – Comparison of Operational Costs

	Physical Solution (32 Servers)	Virtual Solution (4 Servers)
Vi3 Jumpstart Services	$0	$12,000

Table 3 – Consulting Services Costs

	Physical Solution (32 Servers)	Virtual Solution (4 Servers)	Cost Savings
Total 3-Year Hw/Sw Costs	$362,144	$167,707	$194,437
Total 3-Year Power/ Cooling Costs	$28,176	$7,181	$20,995
Total Services Costs	$0	$12,000	-$12,000
Total Cost	$390,320	$186,888	$203,432

Table 4 – Combined Hardware and Operational Savings

DISASTER RECOVERY COMPLEXITY

Disaster recovery is a series of measures undertaken to implement minimal to near-real-time fail-over for a particular application outside the data center involving either a hot or cold site.

Disaster recovery has been a major concern for organizations in recent years. Still, very few organizations have implemented disaster-recovery plans coupled with regular testing.

Similar to a highly available infrastructure, creating a disaster recovery site increases in complexity as the size of the data center grows, making it expensive to implement and maintain.

REDUCED LICENSING COMPLEXITY

Managing the number of server and application licenses to stay in compliance has become difficult to track and maintain. As a result, many organizations today don't have a good handle on whether they are in compliance or not, and struggle to maintain different maintenance contracts that expire at different intervals. Server virtualization greatly reduces the time to manage the environment.

IMPROVED POWER AND COOLING EFFICIENCY

Data centers were originally built at a 1KW to 2KW per rack average. Today, newer systems like blade centers, multiprocessor systems and storage arrays themselves consume 4KW of power. Analysts are now suggesting that new data centers should be built with a minimum of 9KW to 15KW per rack average.

This forces most data center designs to explore larger heating, ventilation and air-conditioning (HVAC) and water-based cooling systems. Additionally, most data centers do not receive volume discounts from utility companies due to the increase in power consumption, causing operational costs to rise continuously.

Because server virtualization reduces the number of physical boxes, it decreases the required power and cooling needs and can help in prolonging the lifecycle of the data center.

Some utility companies in the United States are now issuing credits to organizations who implement virtualization that allows them to reduce their power consumption. Similar to this, there is a flurry of industry activity that is actively trying to find new ways to reduce power consumption across all data center components.

In fact, over the next five years the trend in data center evolution will focus on building and maintaining a truly green data center. Although this will most likely involve replacing all the existing data center components, it will be necessary to stabilize power consumption and create an efficient data center model.

COST BENEFITS OF SERVER VIRTUALIZATION

An operational benefit must be gained from any technology investment. The primary reason so many organizations have adopted virtualization strategies over the past few years is because of the previously mentioned reduced TCO realized by implementing a virtual infrastructure.

Major factors contributing to a lowered TCO include:

- Hardware/software costs
- Operational costs
- Consulting services costs

VIRTUALIZATION MANUFACTURERS

A server virtualization solution consists of three main components: software, servers and storage. VMware stands at the forefront of virtualization software today. Its VMware Infrastructure suite of products enables consolidation, high availability and a more efficient use of resources.

However, newer products such as Citrix XenServer and Microsoft Hyper-V have begun to compete in this space. These alternatives merit an evaluation and comparison to VMware Infrastructure.

Although most servers are compatible with server virtualization software, HP, Sun and IBM solutions are valued because of their virtualization heritage. Also, their established product lines have key enterprise features and functionality.

The last component, storage, has a number of solid contenders to choose from including IBM, NetApp, EMC, HP and LeftHand Networks. ◊

CH 02

CLIENT VIRTUALIZATION

CHAPTER 3:

Desktops, notebooks and other client devices are in use throughout the workplace. Because of their popularity, organizations are struggling to find ways to manage and maintain these devices while simultaneously reducing costs and increasing end-user productivity. Overall, client virtualization can greatly simplify and lower both IT management and administration costs.

On the heels of wide adoption of server virtualization in data centers, a renewed interest in client virtualization has emerged. However, unlike server virtualization, a number of different technologies and devices comprise the client virtualization options available.

CLIENT VIRTUALIZATION COMPONENTS

In many situations, multiple virtualization components are combined to achieve the requirements that satisfy organizational needs.

DEVICES

When building a client virtualization solution, the most important consideration is the end-user experience. Therefore you need to determine the audience, not only in terms of what devices end users will employ (such as notebooks, desktops, tablet PCs, mobile devices, etc.), but also the locations from which they will connect, available bandwidth, the peripherals (such as printers) that they will need to connect to and whether they have a single or multiple display.

Finally, for those users that require stringent security, solutions such as smart-card readers, biometric scanners and two-factor authentication tokens also require consideration. Determining the types of devices can, in fact, eliminate some potential technologies that might otherwise have been considered.

OPERATING SYSTEMS

If the end-user experience is to include an actual operating system as part of the interface, then the choices usually narrow down to Windows XP or Windows Vista. There are many ways to present an operating system to your devices.

In a hosted model, Windows can run on blade PC/workstations in the data center or on virtual machines on a hypervisor. Either way, a display protocol, such as remote desktop protocol (RDP), independent computing architecture (ICA) or remote graphics software (RGS) usually "presents" the operating system while the processing actually occurs on the hosted platform.

In other words, your client devices can be thin or thick since connecting to the remote systems requires minimal hardware. VMware View and Citrix XenDesktop both deliver powerful solutions for hypervisor-based virtual desktop connectivity.

However, hosting the Windows images on a server and then streaming the operating system to existing desktops presents another possible solution. This proves a great option for organizations looking to extend the lifecycle of existing desktops, as it doesn't involve any blade workstation or hypervisor investment in the data center. Citrix Provisioning Server for Desktops delivers this "streaming" solution.

3a. CLIENT VIRTUALIZATION OPTIONS

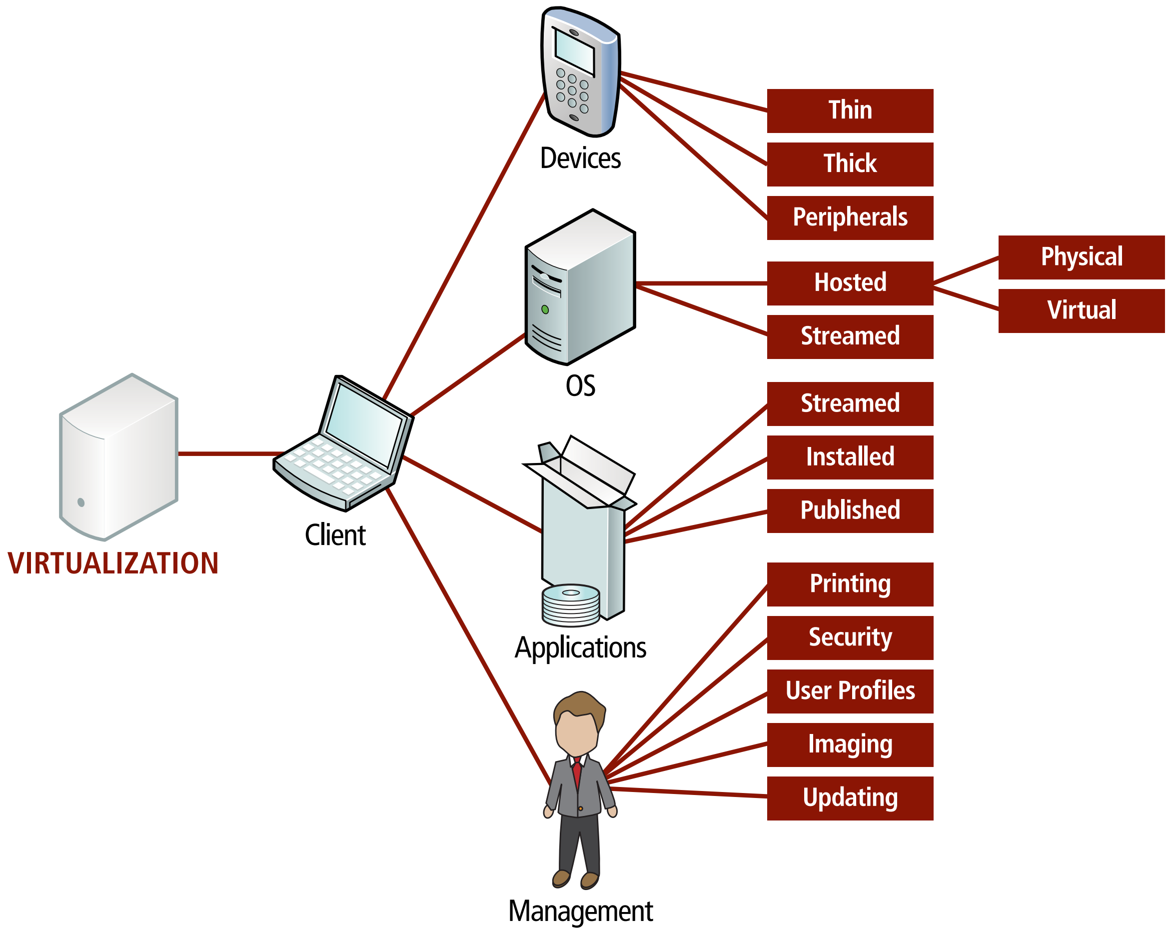

CH 03

APPLICATIONS

Citrix revolutionized the way applications are delivered with WinFrame and MetaFrame, and its solutions remain the standard in many organizations (with its XenApp product). Two solutions, Microsoft Terminal Services and XenApp technology, provide what is known as server-based computing.

They host applications on a Windows server, then deliver them to users via sessions. These application-delivery solutions enable all of the processing to take place at the server itself. This is referred to as publishing applications.

Application streaming, a newer technology available from a number of manufacturers, focuses on isolation and streaming. Isolation refers to technology that installs the application locally on a desktop but isolates it from other applications and processes so that it is completely contained.

This approach avoids the dynamic link library (DLL) conflict issues normally associated with running multiple versions of the same application on the same desktop.

Streaming technology then delivers these applications on demand to the desktop and facilitates removal, upgrades and metering. Microsoft App-V, a part of the Microsoft Desktop Optimization Pack (MDOP), Citrix XenApp (client side) and VMware ThinApp are all examples of application isolation technologies. VMware's ThinApp, although it can be centrally located, remains the only one of the three that doesn't currently offer streaming.

Another option to deliver applications is the traditional method of installing software locally, either manually or with a managed software delivery technology such as Altiris's (Symantec) Software Delivery Solution or Microsoft's System Center Configuration Manager. Some applications, such as antivirus technology, may require local installation because of the way they interact with the desktop operating system.

MANAGEMENT

In determining what devices, operating systems and applications make up your client virtualization solution, a number of important components related to management demand consideration.

Printing: This component has always represented a challenge for both physical and virtual worlds. Depending on the type of client virtualization solution used, the design elements have to be considered carefully.

While printing in a server-based computing environment may prove tricky, hosted operating systems make it a little easier because the printing model doesn't change much. The only possible difficulty that may occur concerns local printing, although most of the virtual desktop solutions have intelligent printer mapping technology that enables printing to a local attached printer.

Security: When designing a client virtualization solution, the consideration of security becomes essential for most organizations because you may be delivering access to critical applications and data via the Internet to devices over which the organization may not have any control. So security should be scrutinized end-to-end to ensure the solution meets the organization's operations objectives.

User Profiles: If delivering a solution via a streamed or hosted desktop, the centralization of user profiles may become important in order to lock down desktops and deliver a consistent user experience across any device. Microsoft delivers local, mandatory and roaming user profiles in its Windows operating systems, but many other technologies now exist to make user profile management easier.

Imaging: For over a decade, many organizations have used imaging technologies to deliver a consistent image to desktops and notebooks. Symantec Ghost probably remains the best-known tool, although Altiris (Symantec), Microsoft and a number of other manufacturers offer solutions. However, when considering client virtualization, most solutions already include this technology.

For example, Citrix Provisioning Server not only streams an operating system (via PXE Boot) to physical and virtual desktops, it also enables the streaming of a single image to those devices. It thereby saves a tremendous amount of disk space.

VMware View Composer, part of the View Premier bundle, has similar technology for virtual desktops, allowing a single image to power on multiple desktops. Therefore, imaging will become an important consideration as the organization develops its client virtualization architecture.

Updating: All desktops, notebooks, thin clients, applications, operating systems and their management applications will need patches and updates at some point during the lifecycle. Developing a strategy for deployment and rollback proves essential in a client virtualization environment because of all the intricate dependencies.

Although device updates will still require OEM tools, operating system and application patching will grow much easier. For example, if virtual desktops are deployed, a single image can be patched, and then every linked virtual machine can be instantly patched as well. With applications, the patch can be applied to a single application image and then the update can be forced centrally for delivery to all users on next access.

CLIENT VIRTUALIZATION BENEFITS

For most organizations that have a variety of client devices, the following benefits of employing a client virtualization solution apply:

Rapid Deployment: The timeframe to deliver new desktops and applications can range from days to minutes.

Efficient Patching: Both operating system and application patching can be done quickly and reliably in a short period of time, as can rollback.

Easier Imaging: With client virtualization, the need to have multiple images for different hardware becomes a thing of the past. Maintaining fewer images, and the ability to deploy them on demand, makes imaging a painless task.

Security: With the proper design, client virtualization can enable stringent security policies (once difficult to maintain) with a variety of client devices. This approach results in a consistent security policy from end to end.

Reduced Energy Costs: If the solution employs thin clients, the power/cooling footprint will reduce dramatically because most thin clients have no moving parts and require very little power.

CLIENT VIRTUALIZATION CONSIDERATIONS

While most organizations can take advantage of client virtualization, there are some important issues to consider as organizations plan to virtualize.

APPLICATION ARCHITECTURE

Multi-win Capabilities: Some applications are designed to run under a single operating system with multiple sessions, also known as multi-win. For those that cannot, you may need to consider isolation or another form of virtualization.

User Profile Dependencies: Some applications may use several locations on an operating system to run properly, which may become an issue if profiles are locked down or centrally located.

Graphics and Rendering Requirements: Most client virtualization technologies cannot handle high fidelity video or advanced rendering delivery. This circumstance remains mostly limited by the remote protocol used. But HP Blade workstations, for example, use their own protocol called Remote Graphics Software (RGS), which allows this capability.

Multimedia Requirements: The delivery of audio and video can present challenges in a virtualized environment. In some cases such as unified communications (soft phones, video conferencing, etc.), the solution may not work and the manufacturer may not support it. So take care when developing solutions that require specific multimedia capabilities.

Special Devices: Give consideration to the devices on which end users rely, such as USB flash drives, headsets, PDAs, biometric scanners, etc. Test these devices properly with the virtualization solution of your choice.

INFRASTRUCTURE ARCHITECTURE

Shared Storage: Achieving an optimal, highly available environment may require shared storage or a storage area network (SAN). This need, however, only becomes prominent when deploying hosted virtual desktops. Blade workstations and server-based computing solutions remain exempt from this requirement.

Workload Sizing: Hardware architectures should be sized to maintain 80 percent or lower CPU and memory utilization. The reason for this strategy has to do with the fact that most of the solutions could see a decreased performance for some or all of the end users hosted on the solution. This may be true for almost any application today, and it's usually operating-system related.

Backups: Depending on how you choose to present user profiles and shared data to desktops and applications, organizations may not need to back up individual desktops. However, each solution is different and should be analyzed independently.

REMOTE ACCESS REQUIREMENTS

Devices: Organizations today have a variety of devices to choose from: desktops, notebooks, tablets, PDAs, etc. Remote access technologies may vary by device, so be sure to determine the device audience and lock it down early in the design process.

Security: Encryption, two-factor authentication devices and biometric scanners are just some of the security technologies used to control remote access. Some technologies may not work with a particular client virtualization solution, so this is also an important consideration.

Bandwidth Limitations: Some remote desktop protocols work well on high latency networks, while some do not. Identifying the minimum bandwidth required for a chosen solution to work is important.

Access Methodology: A client virtualization solution may be delivered by a browser, by a client or by both. Identifying the optimal access method for your organization will help eliminate potential solutions.

TESTING/PROOF OF CONCEPT

Testing each solution for a specific application all the way down to a specific client via remote access technology is critical to a successful implementation. It is highly recommended to get input from end users with real workloads in order to achieve end-user acceptance early on in the testing cycle.

OPERATING SYSTEM AND APPLICATION LICENSING

Desktop Operating Systems: Microsoft now offers a subscription-based model to deploy desktop licenses hosted on a virtualization platform. This model, Vista Enterprise Centralized Desktop (VECD), is based on an annual subscription tied to a license agreement.

Desktop Applications: Licensing for most applications, including Microsoft applications, is based on a per-device model. Virtualizing applications implies no required changes to this model and therefore stays the same.

Server-based Computing: If you plan to host applications via Microsoft Terminal Services or via Citrix XenApp (server side), each device will require a Microsoft Terminal Services client access license (CAL) in addition to the Windows Server CAL. ◊

DATA CENTER
OPTIMIZATION

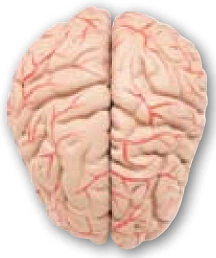

One way to address the problems associated with server sprawl is through physical consolidation. Through physical consolidation, an organization reduces the total number of servers in use by merging the workload onto fewer servers.

Organizations are able to make more efficient use of their computing resources, not to mention free up data center space, reduce power and cooling costs, and reduce complexity to make management easier.

BLADE SERVERS

The use of blade servers in the data center is one approach to physical consolidation. Blade servers are different than traditional rack-mount servers. Blade systems are built with a modular infrastructure that can share many resources with fewer components. Unlike rack-mount servers, blades do not have their own power supplies or fans.

Rather, they share these components with other blades in a chassis that they all reside in. Another key difference is that blades do not have traditional input connections as can be found on the back of a rack-mount server. Communication is done through the chassis midplane, which interconnects all of the blades to the chassis.

In their early development, blades were criticized for limitations with regard to processors, memory and I/O capability (such as the number of network interface cards [NICs] allowed per blade). Advancements in blade technology have overcome these early challenges and they continue to innovate.

In 2008, some manufacturers began to offer a blade with two motherboards in it, allowing for 32 server nodes in a 10U space. New half-height virtualization blades can take up to 128GBs of RAM and possess up to 24 10GB onboard NICs. And some blades now possess serial attached SCSI (SAS) switching, which allows blades to access external direct-attached storage.

BLADE BENEFITS

Blade servers offer many of the highlighted benefits of physical consolidation within the data center. Their consolidated design is more affordable to purchase and maintain than traditional rack-mount servers. Less rack space is used. Blades offer an almost 20 percent reduction in server airflow and around a 30 percent power savings over traditional severs.

With some manufacturers, up to 16 blades can fit into a 10U rack space. That's quite a space savings compared to the traditional 1U "pizza box" servers of the past. Blade deployment is much easier too. A blade server can be set up in minutes. And best of all, the IT team can add resources to a blade setup without having to rewire the entire rack. That's a tremendous time saver.

Blade servers may not be a good fit in every situation. If an organization is only looking to replace one to three servers, blades are probably not applicable. Or if an organization's storage needs are great and there's no budget for a centralized storage system, then rack-mount servers may be a better fit given their greater internal storage space.

When an organization does decide on a blade solution, it should consider the following:

- Be sure that the port count matches your organization's needs, especially if the blades will be used for virtualization.
- Blade management is a key feature. Be sure to pick a management solution that is well designed and sets up with no hassle.

LOAD-BALANCING

When an organization has a fully stocked data center, it has a great deal of computing resources to call upon. The next step is how to make the most effective use of all those resources — how to optimize them.

Load-balancing has become an increasingly popular option for optimizing in the data center, maximizing resource utilization and throughput and minimizing response time, while providing high availability with fail-over.

Load-balancing divides the amount of work that a computer has to do between two or more processors or computers so that the work gets done in the same amount of time and all network users get served faster.

One of the most common applications of load-balancing is providing a single Internet service from multiple servers. This is known as a server farm. Common load-balanced systems include popular website and high-bandwidth FTP (file transfer protocol) sites.

CH 04

Load-balancing is especially helpful in situations where it's difficult to predict the number of requests that will be issued to the network's server. Popular websites typically employ two or more web servers in a load-balancing setup. If one server gets swamped, the requests are forwarded onto another server with more capacity.

LOAD-BALANCING OPTIONS

Organizations can choose between hardware and software load-balancing solutions. A hardware load-balancing device directs computers to individual servers in a network. The device's directions are guided by considerations such as server-processor utilization, the number of connections to a server and overall server performance.

A software solution is often used for Internet services. The software monitors the port where external clients connect to access network services. It forwards requests to one of the backend servers, which then responds to the load-balancer. This allows the load-balancing software to reply to the client without revealing the internal separation of functions.

A security benefit is gained by having the load-balancer serve as mediator. The internal structure of the network remains hidden from the outside client, preventing an opening for attacks on the kernel's network stack or unrelated services running on the other ports.

Load-balancing creates the opportunity for server clustering. Combine load-balancing with server virtualization, where you can run multiple servers on one physical machine, and organizations can cluster a couple of dozen virtual machines onto a handful of physical machines.

If there is a failure of one of the physical machines, the workload is distributed to the remaining physical machines without a loss of service to the end user.

When researching load-balancing solutions, organizations should consider the following features:

- Fail-over features
- Availability
- Overall performance

POWER AND COOLING MANAGEMENT

Much of the efficiency of a data center can be controlled and leveraged through smart management practices. In particular, keeping a close eye on power distribution and hot spots in the data center can pay dividends down the road. Power and cooling management can be defined as the process of designing a modular, energy-efficient system to reduce energy costs and limit server downtime.

A fully integrated power and cooling strategy can help an organization address the following concerns:

- **Growing power demands:** Consolidation strategies create higher density data centers and increased individual power consumption.
- **Increasing power costs:** Every dollar spent on new hardware in the data center requires an additional 50 cents spent on power and cooling.
- **Excessive heat:** Higher density data centers create more hot spots.

In a typically dense data center, both rack-mount and blade servers often have trouble staying cool. These hot spots need direct cooling solutions to ensure that the servers continue to run properly. Putting an effective power and cooling strategy in place gives an organization's data center greater flexibility and maximizes its hardware lifecycle.

Organizations have four main power and cooling solutions to choose from:

- **Room oriented:** This is a standard cooling option that many organizations use. One or more air-conditioning units distribute air through a duct system to lower the temperature of the entire room. This design is significantly affected by equipment configuration, thermostat placement and airflow.
- **Row oriented:** This approach to cooling mounts air-conditioning units either directly above or below equipment racks, and can be adjusted to the unique needs of each row. Airflow is greatly improved compared to a room-oriented solution.
- **Rack oriented:** These cooling systems are dedicated to specific racks and are mounted within each rack. This allows for the best direct airflow to remove the server's hot air exhaust. This is the best option for high-density blade servers for removing hot spots.
- **Mixed cooling design:** The most effective solution combines all of the above options. These mixed environments allow organizations the flexibility to add additional equipment without greatly impacting the data center's cooling system.

HOSTED AND MANAGED SERVICES

Some organizations may opt to completely forgo having a data center and turn over the operation of their network to a commercial hosting center. These kind of hosted services operate much like a utility, offering enterprise-class service in an on-demand and usage-based model.

Here are some of the common services available through hosted services:

- **Internet:** Connectivity through a 100Mbps or 1Gbps connection to the network
- **WAN:** Configure, troubleshoot and optimize WAN network connections
- **Firewall:** Virtual firewalls and multiple firewall interfaces
- **Data storage:** On-demand storage capacity and management services
- **Data backup:** 24x7 availability of mission-critical data and applications
- **Disaster recovery:** Hot-site recovery, virtualized hot server and cold-site recovery
- **Infrastructure:** High-performance, fault-tolerant network and systems infrastructure

Hosted services offer increased performance and additional layers of redundancy that might be difficult for an organization to afford and manage on its own. Additionally, organizations benefit from a single point of accountability, including compliance audits for standards and best practices that are required of these service providers.

And the scalability offered by a hosted service works both ways: services increasing or decreasing in response to an organization's strategic initiatives. ◊

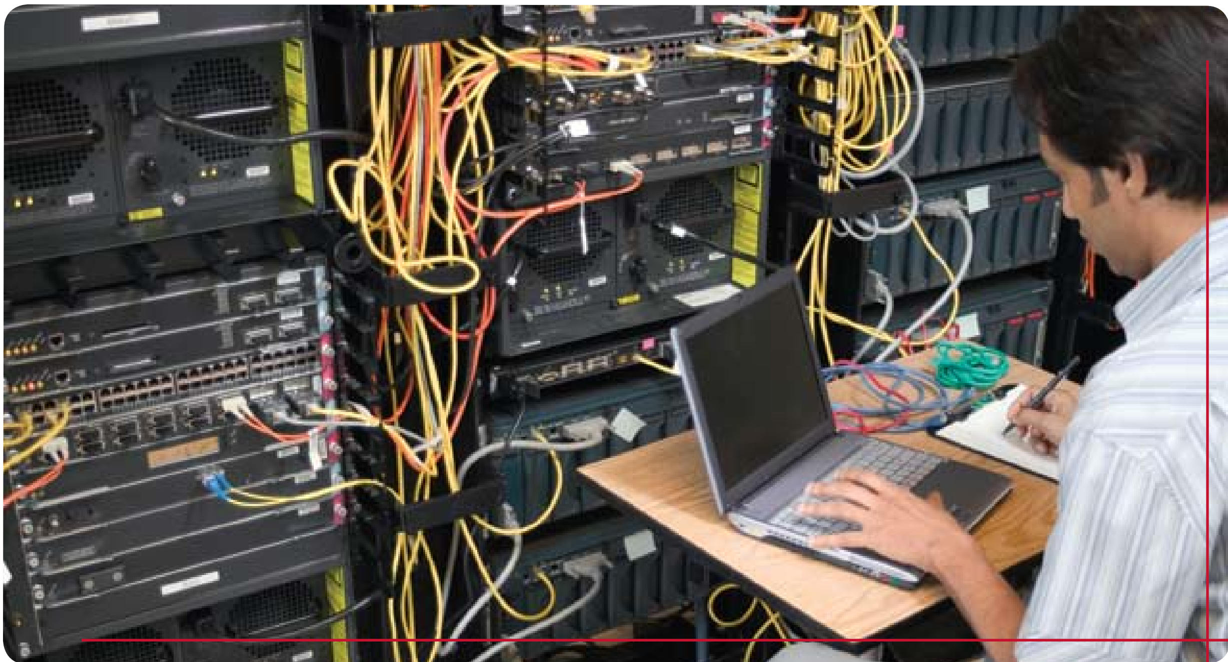

THE

BENEFITS

OF SERVER VIRTUALIZATION

There are many benefits to virtualizing your data center. There's less hardware and software to manage. You'll gain real estate in what's probably an already crowded server room. Labor costs will go down. But what may be one of the least appreciated aspects of virtualization may also be one of the most important: power savings.

By simply replacing physical servers with virtual servers, and thereby reducing the amount of equipment on the floor, you can achieve net energy reductions of 50 percent or more for comparable IT workloads. So what's the big deal with cutting down on power? As the amount of data you manage continues to inevitably increase, cognizance of your power usage becomes increasingly important. Data center expansions and blade server implementations can create power and cooling demands comparable to the amount of power used by small residential neighborhoods or strip malls. So even if you're not paying attention to the rising power costs, you can bet someone in management surely is.

Adding to the issue, servers tend to be underutilized. Organizations tend to utilize their servers somewhere in the 5 to 10 percent range. Even if you're way ahead of the curve and running at 35 or 40 percent capacity, the vast majority of your server space is still underutilized. Keep in mind these underutilized servers are still incurring costs for IT labor, software license updates, warranties, service contracts, and of course, power and cooling costs.

Server virtualization can stop the sprawl and the unnecessary power usage. By virtualizing, you're dividing multiple physical servers onto one — with multiple isolated virtual environments. You can then power down the servers you're no longer using, which clearly cuts your power consumption.

Your CDW•G account manager can help you stop server sprawl in its tracks. Call your CDW•G account manager today to learn more about what server virtualization can do for your organization.

Microsoft

Microsoft® Windows Server® 2008 Standard

With Hyper-V Virtualization Capabilities
Open License Government[1]

$558.59
CDWG 1418186

Microsoft® Windows Server® Hyper-V™, the next-generation, hypervisor-based server virtualization technology, allows you to make the best use of your server hardware investments by consolidating multiple server roles as separate virtual machines running on a single physical machine. You can also efficiently run multiple operating systems — Windows®, Linux® and others — in parallel on a single server.

- Powerful tools give you greater control over your servers and streamline configuration and management tasks
- Enhanced security features protect your data and network and provide a solid, highly dependable foundation for your organization
- Enables your organization to reduce costs, increase hardware utilization, optimize its infrastructure and improve server availability

[1]Purchase five licenses to qualify for the Microsoft Open License Government program; media must be purchased separately; call your CDW•G account manager for details

C7000 Enclosure shown with 495c blade servers; all sold separately

Call your CDW•G account manager for complete blade configuration

HP ProLiant BL495c G5 Blade Server

Quad-Core AMD Opteron™ Processor Model 2384 (2.70GHz)

$2805.96
CDWG 1627949

The HP ProLiant BL495c G5 Blade Server is a virtualization blade designed to eliminate the performance bottlenecks of a virtual machine host.

- Memory: 4GB std; 128GB max
- Hard drives: None ship std.; 2 non-hot-pluggable SSD internal drive bays available
- 6MB Level 3 Cache
- Embedded NC532i Dual Port Flex-10 10GbE Multifunction Server Adapter

HP 1/10GB Virtual Connect Ethernet Module

$5661.80
CDWG 1077604

Simplify and make your data center change-ready. The Virtual Connect Ethernet Module is the simplest, most flexible connection to your networks. The Virtual Connect Ethernet Module is a new class of blade interconnect that simplifies server connections by cleanly separating the server enclosure from LAN, simplifies networks by reducing cables without adding switches to manage, and allows you to change servers in just minutes.

HP Care Pack Services

Decrease hardware downtime with committed six-hour call-to-repair for hardware problem resolution.

For HP ProLiant BL495c G5 Blade Server
Five-year, 24x7, onsite, six-hour Call-To-Repair Hardware Support

$1601.00
CDWG 1595434

To fully capitalize on the capabilities of your HP ProLiant servers, you need a service partner who thoroughly understands server technology and how it behaves in a multi-vendor environment.

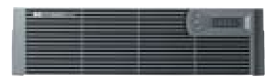

HP R5500 UPS

$3534.13
CDWG 1000449

- 5500VA, 4500W
- Attached line cords for easier and faster connection to your systems
- Longer runtime with Digital Signal Processing (DSP) technology
- Rack-mount kit included
- Support for two load segments
- Enhanced front panel display
- Enhanced Battery Management (EBM)

hp invent

OPTIMIZING THE NETWORK THROUGH VIRTUALIZATION

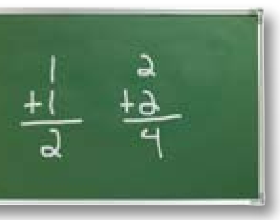

Server virtualization yields tremendous savings when it comes to simplifying the deployment and ongoing maintenance of equipment, freeing up staffing resources and allowing IT personnel to focus their energies elsewhere.

By offering a new way of managing IT infrastructure, virtualization can help administrators spend less time on repetitive tasks. These include such things as provisioning, configuration, monitoring and maintenance.

Considering the various network connectivity requirements, a traditional server will typically take one to two days to fully provision — with a virtual device that time frame drops to just one to two hours. Virtualization reduces the amount of time it takes to provision new servers by up to 70 percent. Which means your organization spends less time installing and configuring, and less money on networking, storage infrastructure and connectivity.

Do The Math

See just how much you can bolster the bottom line with server virtualization.

Server virtualization can make an organization really see green — and we're not talking about the technology's impressive energy savings. Significant cost savings can be achieved by journeying into the virtual world. HP compared the cost of maintaining an infrastructure of 20 physical servers with the cost required to consolidate the same environment with three ProLiant iVirtualization servers running 20 virtual machines.

HP found that over a three-year period, a customer deploying a completely physical environment would have to spend an average of $57,640 on new servers, plus $8,000 in provisioning costs, and around $48,000 in power and cooling charges.

Conversely, an organization that consolidated through virtualization would spend $25,566 on the new virtualized servers, $1,500 on virtualization software, $800 on provisioning, and $20,000 on power and cooling over the same three-year cycle. The end result? The customer that chose server virtualization would pocket $65,474 over three years.

CDW•G's Server Specialists can show you what virtualization can do to simplify management of your network infrastructure.

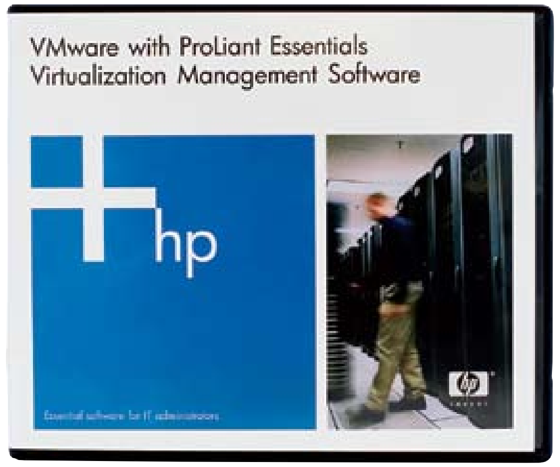

CALL FOR PRICING

HP VMware® Infrastructure 3 Enterprise Edition

CDWG 1005579

- Increase hardware utilization by 50 percent to 70 percent
- Decrease hardware and software capital costs by 40 percent
- Improve server-to-server administrator ratio from 10:1 to 30:1
- License plus one-year 9x5 support for two processors

HP Integrated Lights-Out (iLO 2), featuring iLO Advanced

$359.09

CDWG 1644616

HP Integrated Lights-Out (iLO 2), the fourth generation of Lights-Out remote management, delivers industry-leading remote control for ProLiant servers and ProLiant BladeSystem server blades. HP iLO 2 is an intelligent management processor and firmware integrated on most ProLiant servers.

- Dynamic Power Capping now provides users the ability to cap server peak power, in addition to existing power measurement and regulation capabilities
- iLO Text Console now offers customers a choice on accessing remote servers — through the graphical remote console or the new text based iLO Text Console, perfect for Linux® users

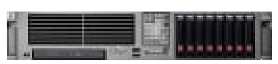

Hard drives sold separately

HP ProLiant DL380 G5 Rack-mount Server

Quad-Core Intel® Xeon® Processors E5450 (3.0GHz)

$5328.64

CDWG 1525488

Introducing a Quad-Core Intel® Xeon® version of the HP ProLiant DL380 G5 Rack-mount Server, with improved performance and enhanced multitasking capabilities for your most demanding applications and virtualization projects.

- Memory: 4GB std., 64GB (PC2-5300)
- Hard drives: none ship std.; up to eight hot-pluggable SAS/SATA hard drives
- 24MB Level 2 cache
- Two Embedded Dual NC373i Multifunction Gigabit network adapters

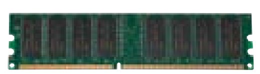

HP SMART BUY

HP Memory Upgrades

For HP ProLiant DL380 G5 Rack-mount Series Servers

2GB (2x1GB)
$148.63[1] CDWG 1133904

4GB (2x2GB)
$208.48[1] CDWG 1227721

8GB (2x4GB)
$427.93[1] CDWG 1528781

16GB (2x8GB)
$2383.03[1] CDWG 1528785

[1]See page 41 for HP Smart Buy details

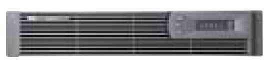

HP R3000 UPS

$1448.36

CDWG 1000446

- 3000VA, 2700W
- Enterprise-wide intelligent manageability with bundled HP Power Manager management software
- Transformer-less technology
- Rack-mount kit included
- Support for Remote Emergency Power Off (REPO) circuitry
- Enhanced front panel display
- Enhanced Battery Management (EBM)

STORAGE
VIRTUALIZATION

Historically, IT departments have deployed storage area networks (SANs) to create a common pool of shared storage resources — as opposed to utilizing direct attached storage (DAS), which is dedicated to a single server. When using SANs, it's necessary to reserve large blocks of storage for individual applications — which often requires additional hardware.

Without vigilant monitoring, good forecasting or the ability to quickly shift capacity to meet requests, your utilization is much lower than what the technology promises. When that allocated space is not being used or if it's never needed, those large blocks of storage are wasted — which means you probably have equipment on the floor that's underutilized and sucking up power.

Consolidated, virtualized storage comes with benefits similar to those you gain from server virtualization: a simplified system, easy creation and backup capability, and enhanced application performance and disaster recovery. Once an organization puts storage networking technology in place, utilization rates generally go up to the 50 percent to 70 percent range.

Benefits of Storage Virtualization

Storage virtualization helps the IT manager or storage administrator perform the tasks of backup, archiving and recovery more easily.

Other value includes:

- Single point of administration
- Non-disruptive data migration
- Faster, more agile computing
- Less physical area to cool
- Power supply efficiencies
- Information lifecycle management
- Heterogeneous replication

CDW•G can help you determine if storage virtualization is the right choice for your organization. Call your dedicated CDW•G account manager today or visit the CDW•G solutions center at CDWG.com today.


Hard drives and chassis sold separately

IBM® BladeCenter® HS21 Server
$2423.92
CDW 1392212

Organizations need reliable servers that offer consistent performance with minimal maintenance. IBM® BladeCenter® HS21 is simple to install and maintain, integrates into all BladeCenter chassis at power-efficient envelopes and offers key features for fewer maintenance hassles.

Quad-Core Intel® Xeon® Processor E5430 (2.66GHz)
- Memory: 2GB std., 32GB max. (PC2-5300)
- Hard drives: none ship std.; up to two SAS drive bays available, 292GB max. storage
- 12MB Level 2 Cache
- Integrated Dual Broadcom 5708S Gb Ethernet controllers

Eaton Evolution UPS 1550VA 1U Rack Mount
$518.69
CDWG 1426328

- 1440VA, 1100W, 4 output connector
- Hot-swappable batteries for high availability
- Automatic power off for non-critical devices on battery operation to conserve runtime for the most demanding systems
- High density, 1U form factor

Citrix® XenAPP™ 5.0
Single-user connection license with Subscription Advantage
$600.00
CDWG 1176409

Citrix® XenApp™ 5.0 is an end-to-end applications-delivery system that offers both application virtualization and streaming to enable the best access experience for any user, with any device, working over any network.
- Simplifies Windows® application delivery
- Optimizes security and compliance for sensitive information
- Accelerates application and network performance

Acronis® True Image Echo Virtual Edition™
One Server license with Standard Support[1]
$2683.27
CDWG 1697746

One Server license with Premier Support[2]
$2773.04
CDWG 1697748

- Complete system restoration within minutes after any unforeseen event
- Supports P2V and V2P migration
- Recover to dissimilar hardware

[1]Support available via e-mail, phone and chat during normal business hours 8 a.m. - 8 p.m. EST, M-F [2]Support available via e-mail, phone and chat 24-hours a day, 7 days a week

CALL FOR PRICING

IBM System Storage™ DS5000 Series
CDWG 1558213

- Designed to meet today's and tomorrow's demanding open-standards for lifecycle longevity
- Adept at supporting transactional applications such as databases and OLTP, throughput-intensive applications such as HPC and rich media, and concurrent workloads for consolidation and virtualization
- Pay-as-you-grow scalability, initially 256 drives, for demanding capacity

VIRTUALIZATION

BRINGS FLEXIBILITY

AND PROTECTION

Nearly every aspect of today's government is expected to be available continuously without interruption, regardless of circumstances. When disaster strikes — whether a natural disaster or technological failure — operational services and technologies are expected to be available. Most organizations need to place a high value on being prepared for disasters of any kind, because the practical ramifications of failing to do so can be devastating.

When a government organization experiences an interruption in services or suffers a loss of data, the public can lose confidence in that organization's viability in a crisis and its ability to protect their personal information. Even the loss of a single mission-critical service, such as e-mail or web connectivity, can cost some organizations millions of dollars in direct costs. Avoiding this downtime through careful disaster planning is a clear benefit.

Today's rising electricity costs, shrinking power supplies, and mounting social and economic pressure to go "green" are forcing organizations of all shapes and sizes to rethink their IT strategies.

For a growing number, server virtualization holds the keys to effectively meet these challenges.

The flexibility afforded by virtualization will pay dividends in facilitating business continuity and disaster recovery initiatives. Having modular building blocks that can be easily isolated from other pieces of the IT infrastructure or data center means those building blocks can be easily moved, turned on or turned off in the event of a disaster.

You can protect your network and your reputation with powerful solutions from CDW•G.

APC® NetShelter® SX Enclosures

Next generation rack enclosure includes advanced cooling, power distribution, cable management and security features

AR3100 42Ux24"x42"

$1228.91

CDWG 920138

- 42Ux24"x42", 42U base enclosure
- Removable doors and side panels
- Ships fully assembled
- Perforated front and rear doors

APC Smart-UPS® 1000VA USB and Serial RM 1U 120V

Performance power protection for rack-mounted servers and networks

$624.43

CDWG 434905

- 1,000VA, 640W, 4 outlets
- Automatic Voltage Regulation (AVR)
- Intelligent Battery Management
- Overload Indicator

CALL FOR PRICING

Sun Storage 7110 Unified Storage System

Easy-to-use enterprise storage features at entry-level prices

CDWG 1631059

The Sun Storage 7110 Unified Storage System makes it easier than ever to simplify your storage for less. The easy-to-use appliance is ideal for enterprise workgroups.

- Enterprise data services with no additional fees
- Install and configure in minutes
- Rapidly troubleshoot, locate and isolate problems
- Provides 2TB in a 2U form factor
- 8GB RAM
- 10K rpm SAS drives ensure higher performance

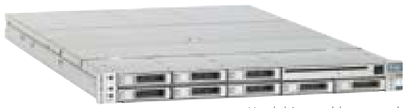

Hard drives sold separately

intel Xeon inside™ Powerful. Efficient.

Sun Fire™ X4150 Rack-mount Server

Energy-efficient system design saves on power and cooling

$3231.89

CDWG 1477630

The outstanding performance of the Sun Fire™ X4150 reduces development time, speeds time to market and runs memory- and disk-intensive applications more efficiently, making it especially suited for web serving, database and storage-intensive applications.

- Two Quad-Core Intel® Xeon® Processors L5420 (2.50GHz)
- Memory: 4GB std., 64GB max. (PC2-5300)
- Hard drives: none ship std.; contains six hot-swappable SATA disk drive bays available for expansion
- 12MB Level 2 Cache
- Four integrated 10/100/1000BASE-T Ethernet ports

STORAGE
CONSOLIDATION

One way to gain an understanding of the advantages offered by today's storage solutions is to look back at how these gains came about. Since the beginning of computing, there has always been a need to access more data than could be stored in servers. This became apparent in the 1990s as organizations began to increasingly rely on technology and easy access to information.

All that information needed to be stored somewhere. So organizations acquired more servers. But this approach proved to be inefficient. If a server needed more storage space, another disk drive or disk subsystem had to be procured, installed and then configured for accessibility. This process could take days at best — and weeks at worst.

ADVANCES IN STORAGE

With the introduction of consolidated storage, a larger quantity of storage could be acquired cost effectively and then shared between several systems. Each server still had its own storage, but if it required additional space, another LUN (logical unit number; referring to a virtual disk drive) could be carved out of the total pool of storage space and assigned to the individual server.

Storage continued to grow, and individual storage subsystems started to proliferate with 10s and 100s of servers being attached to each subsystem. With this development came the introduction of storage networking. Storage networking allowed servers to acquire storage from multiple storage subsystems if needed — again leading to greater storage efficiency.

While there are greater efficiencies, storage systems are often underutilized. Traditional mainframes had storage utilization rates in the 70 to 80 percent range. And the typical server averages only 25 to 35 percent storage utilization.

Another concern with storage networking is that if a storage subsystem needs to be replaced, a lengthy outage, typically a weekend, is required to move the data from one subsystem to another and then reconnect the associated servers.

The next major advancement in storage was storage virtualization. Storage virtualization abstracts the physical storage subsystem and allows the creation of a new virtual type LUN that is disconnected from the physical storage it is attached to. With this advancement, data could be moved from one subsystem to another without incurring any outage or the need to physically move the data.

Combined with server virtualization (referenced in Chapter 2), both the computing environment and its associated data could be seamlessly moved through the physical environment, creating virtually no downtime and a high availability in computer processing environments.

STORAGE AREA NETWORKS

As storage virtualization becomes more popular, networked storage continues to hold a presence in the storage realm. With networked storage, the data already resides on a shared disk array, which speeds the process and allows organizations to take greater advantage of virtual machine technology.

There are two flavors of networked storage: basic network attached storage (NAS) and higher-speed storage area network

(SAN). A SAN is a type of network attached storage, connecting separate disk arrays scattered over a network into one monolithic virtual storage device that can be managed uniformly.

What differentiates a SAN from a standard NAS system is high speed. SAN solutions move applications across an enterprise instantaneously and transparently to end users. They are ideal for high-performance applications. And they provide flexibility, higher utilization rates and lower costs, especially when combined with a server blade system.

SAN technology is becoming more widely used as components drop in cost and complexity. In addition, the proliferation of Internet Small Computer System Interface (iSCSI) — an IP-based connection standard — in storage devices makes a SAN even more desirable.

One disadvantage of SAN has been the complicated Fibre Channel I/O fabric necessary to pass data at high speeds between servers and storage devices. iSCSI is slower than Fibre Channel but fast enough for 80 to 90 percent of Windows environment applications.

It's a familiar interface for IT workers accustomed to Ethernet and IP. And it allows secure SAN extension over IP networks, making it possible to tie together storage devices in far-flung locations.

DATA STORAGE

Up to this point we have only been talking about the raw storage (LUNs, disk drives, storage subsystems) and not the data residing on top of the storage infrastructure. The data itself still needs to be backed up to additional devices in case of drive or subsystem failure. In the past this backup was typically done to tape.

As the volume of data organizations maintained continued to grow, backup systems became more sophisticated. Early on, a full backup of all data was usually taken daily — or weekly at worst. These backups occurred whether the data changed or not.

Backup software then became more intelligent and advanced so that it would only backup files that had changed. This advancement drastically reduced the amount of data that needed to be backed up on a full backup, and made it much more feasible to do a full backup nightly.

As the volume of data continued to grow, IT administrators realized that significant portions of data, though required to be saved and accessible, weren't being accessed with much frequency. Backup software manufacturers improved their technologies and introduced archiving software. Utilizing archiving software, data that is intermittently accessed could be moved to lower cost alternatives such as tape or slower disk.

One of the only noticeable drawbacks to this approach is a brief delay while the data is being brought back from the slower device. This archiving software drove down the cost of storing long-term data such as tax returns, historical records, etc.

So far, this review of the growth of data storage has only covered individual, unconstructed data without touching on e-mail. The vast volume of e-mail compounds all of the difficulties in providing adequate storage to organizations. The need for e-mail archiving has grown as organizations comply with legislation such as the Sarbanes-Oxley Act (SOX) the Health Insurance Portability and Accountability Act (HIPAA).

DATA DEDUPLICATION

The latest technology advancement is the ability to eliminate duplicate data from even the subfile level. With earlier storage options, if a small amount of data changed in a larger file, the entire file was rewritten and viewed as new. Imagine 20 people all have the same file saved, a small amount of data is changed and sent out to those same 20 people, you now have 20 new files.

Data deduplication not only recognizes these files as the same, it saves only one copy of the file and references the others through a sophisticated set of algorithms. And if one person changes a portion of one file, data deduplication allows organizations to only save the small portion of changes and not the entire file.

With this example, a nearly 20:1 savings in space can be gained. Data deduplication continues to drive storage costs down for organizations.

INFORMATION LIFECYCLE MANAGEMENT

Information lifecycle management (ILM) is a strategy for classifying different types of data and migrating those data categories to the right storage medium at the right time. Each organization's strategy is unique, but basic ILM principles apply across the board.

The first step in creating an ILM strategy is for the IT department to gain a firm understanding of how the organization uses its applications to run its operations. This understanding will allow the IT team to classify data between mission-critical information and nonmission-critical information.

A SAN or NAS system utilizing online disk arrays belongs in the highest tier and should reside in an online storage system with highest performance and highest availability. Background applications, such as payroll and administration tools, may be used for second-level data.

Disk-based backup systems would be appropriate storage media for this data class. E-mail, particularly older messages, would likely belong at the bottom of the totem pole and be stored in low-cost tape drives or write-once media.

But ILM isn't just about classifying types of data. The value of information often changes over time, so depending on the organization, e-mails less than 30 days old may qualify as first-level data, then move down the line over specified periods.

Establishing a hierarchical approach to data storage helps reduce total cost of ownership (TCO) of a storage solution. Low-priority data is moved to more economical storage media, and fewer pieces of information on each level of storage means faster backup and recovery time for each level.

That means organizations consume fewer IT resources. With a hierarchical approach to storage in place, organizations can reduce backup time by 30 to 40 percent and reduce recovery time by 70 to 80 percent.

DATA RELIANCE

Aside from compliance and cost issues, ILM is a sound strategy simply because data is what drives organizations today. The more control an organization has over that data, the better the organization will be run.

Industry analysts believe that 50 to 60 percent of an organization's knowledge can be found within its e-mail system. This statistic drives home the point that maintaining this information is very important, especially for continuity of operations plan (COOP) logistics.

ILM can be a daunting task for an organization that has long stored every bit of data on its networked disk-based storage backup system. Software is as important as hardware in developing an ILM strategy, both from a planning perspective and logistical standpoint.

Storage resource management software allows organizations to take a snapshot of the existing storage environment. The software records storage resources, files and database systems; it provides information regarding the size and utilization levels of files; and gives usage attributes of that information. With this data, organizations can set policies to move data to the right place at the right time.

STORAGE BEST PRACTICES

Technological advancements continue to drive rapid changes in the storage field. Similar to what is being seen in the realm of server virtualization, improvements in storage virtualization, as well as how data is stored and accessed, are developing rapidly. What follows here is a rundown of some of the gains offered by upgrading an organization's storage environment.

STORAGE VIRTUALIZATION

With hundreds of server instances (both physical and virtual) requiring 24x7 access to data, a consolidated storage subsystem is a must. By virtualizing at the storage subsystem level, an organization gains many benefits:

- **Manufacturer independence:** Storage virtualization allows seamless migration between different subsystems regardless of the manufacturer. At any point an organization can switch manufacturers/subsystems to better meet current needs.
- **Storage tiering:** Data can move from high-priced/high-performance storage to low-priced/lower-performing storage as the data's value decreases over time.
- **Storage efficiency:** Because data is easy to assign and move around, organizations can run subsystems at higher efficiencies without concern regarding procurement cycles.
- **Storage management/personnel:** Administrative tasks that traditionally were done during weekend outages can now be performed during the production day. Storage administrators are more satisfied with their workloads and tend to turnover less frequently.

FLEXIBILITY, QUICK RECOVERY

Today, rapid expansion in constructed, unconstructed and e-mail data is forcing storage architectures to be extremely flexible, yet recoverable nearly immediately.

The backup software that organizations use needs to be incremental. Most organizations do not have time for traditional full backups. These backups should be of the disk-to-disk-to-tape (D2D2T) architecture. D2D2T provides for extremely fast backup performance while still allowing high restore performance.

Archiving is a subset of backups that peels off the data that has a low probability for being restored and archives it on an even lower cost solution.

In conjunction with the software side of backup/archive management is the physical side. A best practice approach will include snapshots of the organization's primary data so that a simple file deletion can be restored without even needing to access backup software. All of this data can be brought back directly from the organization's file tree.

Once backed up, the secondary disk (of disk-to-disk) should be capable of deduplication. Deduplication virtually compresses data by eliminating duplicate blocks of data and referencing them instead. Compression ratios of 20:1 are not uncommon for data overall.

Tape is not dead. While designing storage for performance and recoverability is important, the need to be power conscious is still necessary. For data that has to be stored for years but has no urgency in the speed with which to access it, tape is less than one-tenth the cost of the least expensive storage, not even accounting for power savings. ◊

Gigantinormous.

It’s the only word we could come up with to describe our Configuration Center capabilities.

The CDW® Configuration Center can handle your configuration needs. Our A+ and vendor-certified technicians can custom configure, image and test hundreds of systems every day. And that’s just the beginning. We handle everything from basic installation of hardware and software to high-end enterprise configuration services at one of our three Configuration Centers. We offer a variety of services to keep your organization up and running at optimum efficiency.

Get configured. Call 800.808.4239
or visit CDWG.com

The Right Technology. Right Away.®

CDW·G®

NETWORK
OPTIMIZATION

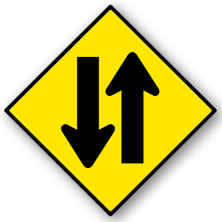

CHAPTER 6:

Network Design

WAN Optimization

Application Networks

Storage Networks

As more organizations look to lower operating costs and consolidate servers by utilizing virtualization, the network plays an increased role in ensuring communication flows between fewer physical servers. So organizations use network links more than ever before, and network redundancy becomes increasingly important.

NETWORK DESIGN

A network's primary purpose is to support the organization's functions. So when designing a resilient network, the first step for organizations is to determine the requirements to support operations functions and develop a network strategy accordingly.

Upper level management and the various departments of the organization can help determine operations requirements. In most organizations, an IT governance committee consisting of upper-level management helps establish the operations requirements of the network. Through these established requirements, network architects can determine the level of redundancy that needs to be incorporated into the network.

NETWORK ASSESSMENT

Once the requirements are set and a network strategy is developed, the second step, the planning phase, can begin. Planning includes both an accurate assessment of the current environment and a gap analysis to determine if the existing infrastructure, sites and production environment can scale to include a new, resilient infrastructure.

This assessment should take into consideration the following:

- Current applications and data on the network, such as Voice over IP (VoIP), e-mail, structured query language (SQL), common Internet file system (CIFS), Internet and video-on-demand
- Virtualization and consolidation of server infrastructure
- Current network topology, including but not limited to: network devices, physical and logical links, external connections, frame types, routed and routing protocols, application specific protocols and IP addressing schemes
- Traffic and network utilization analysis

Many tools exist to facilitate network assessment. These tools range from basic device information output tools that display the network device utilization to third-party tools.

For example, with Cisco devices, Network Based Application Recognition (NBAR) allows for the viewing of interface statistics, CPU and memory utilization, NetFlow and application flows. Third-party tools that monitor networks, sniffers and SNMP tools can also be used.

DESIGNING THE NETWORK

The third step in building a resilient network is laying out the actual design of the network. The design must incorporate all gathered information concerning operations and technical requirements. It must also include specifications for availability, reliability, security, scalability and performance.

Network engineers commonly recommend designing a resilient network in modules. Modules allow an organization to provide the highest degree of resiliency by segmenting traffic and preventing a single point of failure.

For example, in diagram 6a the campus network consists of access, distribution and core modules. The access module provides Layer 2 connectivity to workstations and end users. It connects through redundant links to the distribution module.

The distribution module provides routing between users in the access module and other advanced features such as quality of service (QoS), gateway redundancy and other security features. It connects to the core module via multiple high-speed Layer 3 links. In doing so, it takes advantage of the high convergence of advanced routing protocols. The core module provides high-speed switching to other modules in the infrastructure.

Based on an organization's operations requirements, network redundancy is designed for each network module. Since the core and distribution modules provide connection to other services, they typically incorporate redundancy in the form of multiple network devices.

RESILIENCY AT LAYER 2 AND LAYER 3

In designing resilient networks, it is crucial to eliminate single points of failure. This effort includes having redundant links to critical servers and network devices. However, redundant links can create problems.

For instance, in Layer 2 switched environments, redundant links can cause switches to flood packets throughout the network, effectively halting the switching of production traffic.

A Layer 2 protocol designed to prevent such flooding, Spanning Tree Protocol (STP) places one of the redundant links in a blocking state. Although STP prevents Layer 2 loops, it remains slow to converge. STP improvements (such as Rapid STP) help decrease the convergence time.

At Layer 3, advanced routing protocols enable the highest level of network resilience when utilizing redundant links. Not only can advanced protocols load-balance traffic over redundant links, they can converge in a matter of seconds in the event of a primary link failure.

A common best practice, aggregate redundant links at Layers 2 and 3, increases resiliency. Technologies, such as EtherChannel, that combine switched or routed links into one logical link effectively double the bandwidth on the link and minimize the convergence. Because the switch or router sees aggregated links as a single link, traffic continues to flow through the other links if one should fail.

6a. AN ACCESS, DISTRIBUTION AND CORE MODULES SETUP

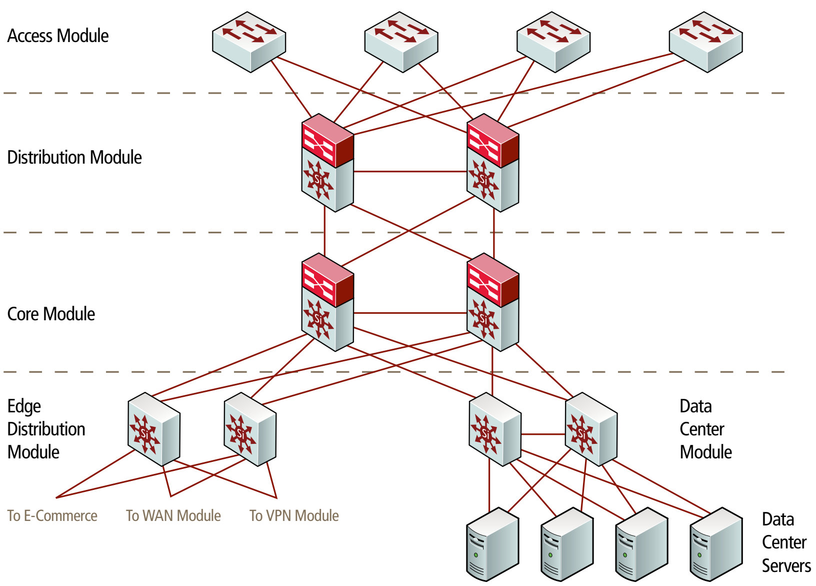

WAN OPTIMIZATION

The 21st century has brought tremendous change to how organizations operate. As organizations have expanded, low-speed wide area network (WAN) links often serve to connect branch offices to one another. In order to provide end users with a high level of service, IT departments typically installed local servers at each branch office.

At first, this proves to be an adequate solution. However, as organizations continue to grow, it becomes increasingly difficult to manage all of these remote servers. Major issues arise around patching, back up, repairs and redundant data.

To lower costs, IT departments often look to consolidate servers. Yet the dilemma of providing an acceptable level of service to branch users over WAN links remains a concern. Applications generally do not work well over WAN links, with WAN latency further contributing to this problem.

WAN optimization helps overcome this dilemma. Through WAN optimization, IT departments can provide service nearly as effective as local area networks (LANs) to branch offices over low-speed WAN links. WAN optimization accomplishes this feat in two ways: by optimizing the WAN and optimizing applications for WAN communications.

WAN optimization devices optimize WAN links in several ways. First, they implement a WAN optimized version of transmission control protocol (TCP). This version of TCP maintains large initial windows and enhancements to deal with WAN congestion. Additionally, advanced compression and disk caching contribute to minimizing WAN traffic.

WAN optimization devices also provide enhancements to optimize applications for WAN communications. Most commonly they optimize common protocols, such as Windows file and print, network file system (NFS), HTTP/S, file transfer protocol (FTP), SQL, messaging application programming interface (MAPI) and others for WAN communications.

The resulting optimization reduces application ping-pong over the WAN, while advanced disk caching helps prevent redundant data traffic.

When implementing WAN optimization, important things to consider include:

- Existing WAN connection speeds
- Current WAN latency
- The age of WAN routers
- The Internet operating system (IOS) of the WAN routers
- Current WAN utilization and transmission control protocol (TCP) flows
- Proprietary or uncommon applications that may traverse the WAN
- Network infrastructure
- Multimedia or streaming content that traverses the WAN

Many options exist for implementing WAN optimization in a network. However, deploying WAN optimization as close to the WAN edge as possible is recommended. This approach ensures that only traffic destined to the WAN gets optimized.

Each WAN manufacturer offers different recommendations on how to physically install WAN optimization devices. The main possibilities include installing WAN optimizers inline, out-of-path (using web cache communication protocol [WCCP] to intercept data to be optimized) or within a router. Some manufacturers offer all three options while others only offer one or two.

Which option to utilize depends on each network environment. Sometimes a network will use all three. For example, when designing and implementing a WAN optimization solution, an out-of-path mode utilizing WCCP for data centers and headquarters while using inline mode or installing the WAN optimization appliance in the WAN routers might be recommended. Diagram 6b shows such a deployment of WAN optimization appliances.

APPLICATION NETWORKS

Today's networks have grown more complex and intelligent than ever. Required to support data, voice, video and other critical applications, networks must provide security, quality of service, acceleration and high availability to these services.

Traditional networks perform network-related decisions based on the packet. On the other hand, application networks provide enhanced intelligence by examining packets at the message level to make network-related decisions based on the

6b. WAN OPTIMIZATION DESIGN

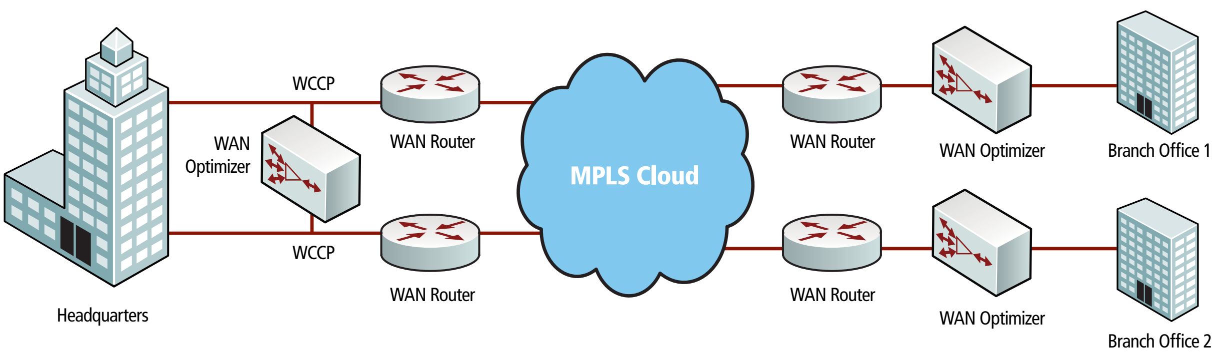

6c. APPLICATION NETWORK EXAMPLE

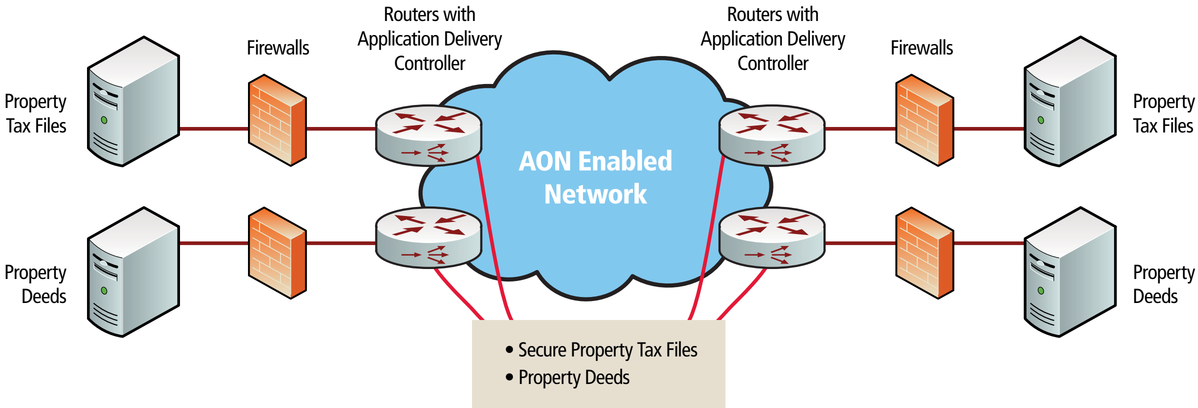

organization's policies.

Application networks have the ability to inspect the full message, including all content and headers, and can thus apply gradual policies to different types of applications within the same protocol. Through this deep inspection, application networks provide improved application availability, security and accelerations.

Through network intelligence, application networks provide the following:

- Enable disparate applications to communicate by appropriately routing application messages in the format expected by that destination
- Enforce consistent security policies for application access and information exchange
- Provide a high level of information flow visibility, including the monitoring and filtering of messages for both operations and infrastructure purposes
- Enhance application optimization by providing application-level load-balancing, offloading security and XML operations, achieving TCP optimization and offering application-level caching and compression services similar to WAN optimization

The application delivery controller serves as the core of the application network. Typically installed at the data center, this controller can function as a standalone device or as a module that installs in an enterprise class router. The application controller distributes traffic among a number of servers or brand offices based on application-specific policies.

A smaller version of the application delivery controller can be installed at branch offices. Like the data center application delivery controller, branch controllers can function as a standalone device or integrate into a router module. Diagram 6c shows an example of an application network.

APPLICATION NETWORK DEPLOYMENT SCENARIOS

At branch or remote offices, application network devices can be deployed at the edge. A single device can provide all the services required by the branch to effectively communicate with the central office.

At the enterprise edge, application network devices can act as an application-security or business-to-business (B2B) gateway. As an application-security gateway, the application network device has the ability to act as an XLM trust enforcement point to provide consistent authentication, authorization and accounting enforcement across all backend services and applications.

As a B2B gateway, the application network device provides a transparent interface with partners by providing trust, policy enrichment and enforcement, protocol bridging, message validation and transformation services.

At the enterprise core or data center, application network devices provide transparent interapplication communication and can intercept and analyze traffic in message formats. These devices can also provide a network-embedded communication bridge between protocols and applications. Finally, application network devices at the core can offload infrastructure functions, such as message-level load-balancing, to the network.

STORAGE NETWORKS

As has been outlined in earlier chapters, many organizations have turned to centralization, consolidation and virtualization to help reduce costs and better utilize scarce resources. Centralized storage stands as one of the keys to any consolidation project.

Storage networks attach remote storage devices to servers in such a way that storage appears locally to the host operating system. To achieve this result, they interconnect storage devices and servers through an array of different protocols, most commonly Fibre Channel, iSCSI or InfiniBand.

Although they function in different ways, each of these protocols has the same goal: to provide lossless, low-latency access to shared storage.

As the name implies, Fibre Channel (FC) operates over physical fiber. FC provides speeds of 1Gbps, 2Gbps, 4Gbps and 8Gbps of lossless low-latency access. In small environments, FC can connect a host to shared storage via a point-to-point connection.

However, in larger environments, FC typically interconnects hosts and shared storage through a switched fabric. Switched fabric has similarities to Ethernet switching in that all devices connect to Fibre Channel switches.

iSCSI allows clients to send SCSI commands to SCSI storage devices over IP networks. Unlike FC, iSCSI does not require special fiber cabling and can run over long distances using existing network infrastructure.

Traditionally cheaper to implement than FC, iSCSI interconnects hosts to storage devices. It does so in a manner similar to FC, either point-to-point or through a switch. However, instead of using an FC switch, iSCSI operates over the same mediums as IP. Thus, iSCSI typically employs Ethernet switches.

Nevertheless, storage behaves differently than Ethernet or IP. Thus, when implementing iSCSI and using Ethernet switches, best practice recommends enabling jumbo frames on the Ethernet switch. Doing so changes the maximum transmission unit (MTU) of Ethernet from 1500 to 9000, allowing the switching of larger packets without fragmentation.

InfiniBand, a highly scalable high-performance protocol primarily used in high-performance computing, not only provides connection between servers and high performance storage devices, but can also connect processors and RAM between nodes. Like FC, InfiniBand uses a switched fabric topology and requires a special cable for interconnections.

Although these are the primary protocols for connecting to storage networks, emerging technologies provide a unified fabric for IP traffic and FC traffic. For instance, Fibre Channel over Ethernet (FCoE) and data center Ethernet provide lossless high-speed access over 10Gbps links, thereby converging FC, Ethernet and IP traffic on the same medium.

The underlying principles of storage network design remain relatively straightforward:

- Plan a storage network topology that can handle the number of ports necessary now and into the future.
- Design with a given end-to-end performance and throughput level in mind, taking into account any physical requirements.
- Provide the necessary connectivity with remote data centers to handle the operations requirements of COOP and disaster recovery.

Adhering to these principles will ensure a storage network that scales well into the future. ◊

GLOSSARY

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

Application network

An application network is a form of network that provides enhanced intelligence by examining packets at the message level. An application network inspects the full message, including all content and headers, allowing it to be able to apply gradual policies to different types of applications within the same protocol.

Application streaming

Application streaming is a form of server-based computing with two basic processes: isolation and streaming. Isolation refers to the process of installing the application locally on the desktop, but isolated from other applications and processes. Streaming refers to the delivery of these applications on demand to the desktop and facilitates removal, upgrades and metering.

Blade server

Blade servers are a newer alternative to the traditional rack-mount server form. Blades are built with a modular infrastructure that can share resources and has fewer components. Blades do not have their own power supplies or fans. All communication with the blade is done through the chassis midplane that interconnects all of the blades sharing the chassis.

Client virtualization

Client virtualization refers to a form of virtualization that partitions end-user devices into multiple isolated virtual environments. Instead of the operating environment being located on the client device, the operating environment goes into a "virtual machine" that resides on shared hardware, like a server.

Data deduplication

Data deduplication is an approach to protecting data in which an organization's files are monitored and every time a file is changed or auto-saved, a copy of the changed bytes/blocks is replicated to either a local directory or a remote location, allowing for to-the-second recovery.

Disk-to-disk-to-tape (D2D2T)

D2D2T is an archiving methodology in which data is copied from one disk to a second disk and from the second disk to tape. This approach helps eliminate data loss problems experienced through tape drive or tape failure. It also allows for multiple copies to be made, kept both onsite and offsite.

Distributed computing

Distributed computing refers to the idea of spreading out server application loads across multiple tiers, platforms and locations. Many IT departments took up this approach to take advantage of the proliferation of underutilized x86 servers in data centers.

EtherChannel

EtherChannel is a Cisco-based switch technology for port trunking. EtherChannel brings together multiple physical Ethernet links to fashion one logical Ethernet link that provides greater bandwidth and minimizes convergence.

High availability

High availability is a series of measures undertaken to implement minimal to near-real-time fail-over for a specific application. Typically, systems that serve the network backbone, such as directory services, file and print sharing, e-mail and enterprise applications, are designated the most important applications and need to be highly available for an organization.

Hypervisor

A hypervisor is a virtualization software program that allows multiple operating systems to share a host computer at the same time. While it appears that each operating system has the host computer's full resources allocated to it, the hypervisor is actually controlling the resources, distributing what is needed to each operating system in turn.

Imaging

Imaging is the process of delivering an operating system or software program to a desktop or notebook computer. This process allows the exact same content to be delivered to multiple machines, saving time and resources that would be needed for manual delivery.

InfiniBand

InfiniBand is a switched fabric link for data flow between processors and I/O devices. It has throughput of up to 2.5GBps, is designed to be scalable, and supports quality of service and fail-over.

Information lifecycle management (ILM)

ILM is a hierarchical strategy for classifying different types of data and migrating those data categories to the right storage medium at the right time. Data is identified as either mission-critical or non-mission-critical. High-priority data is kept on a system with high performance and high availability, while low-priority data is moved down to more economical storage solutions.

Load-balancing

Load-balancing is a data center technique where processing work is split between two or more computers so that the work gets done in the same amount of time. All network users receive faster service as a result.

LUN

LUN refers to logical unit number. A LUN is a unique identifier used on a SCSI bus to distinguish between devices that are sharing the same bus.

Network attached storage (NAS)

NAS is a hard disk storage setup with its own network address rather than being attached to a server. It appears on a network as a shared drive, accessible by multiple computers for simultaneous file sharing.

Network interface card (NIC)

A NIC is a computer hardware interface that connects the computer to the network.

Serial attached SCSI (SAS) switching

SAS is an inexpensive, disk-based approach to data storage that emphasizes higher density storage and a high transfer rate, but with decreased performance. A SAS switch enables servers to connect to multiple SAS storage arrays.

Server-based computing

Server-based computing is a technology where applications are deployed, managed, supported and executed from a server. The only thing the desktop or notebook computer receives transmission of is the screen information.

Server cluster

A server cluster is a group of linked servers that operate as one server. Often connected via a LAN, server clusters improve performance, scalability or availability while giving organizations a greater return on their investment, and thus cost savings.

Server virtualization

Server virtualization is the virtualization of multiple application servers onto one physical server while remaining logically distinct with consistent hardware profiles.

Storage area network (SAN)

A SAN consists of a high-speed special purpose network (or subnetwork) that interconnects different kinds of data-storage devices with associated data servers on behalf of a large network of users. Although the storage devices are remote, they appear to be locally attached to the operating system.

Storage virtualization

Storage virtualization allows physical storage to be shared across multiple servers. The physical storage devices behind the virtualization layer are viewed and managed as if they were one large storage device. This technique hides the storage content from servers that are not authorized to access that storage, providing better security.

Switched fabric

Switched fabric is a network topology where network nodes connect to each other through one or more network switches. It includes the switching units in a node, the integrated circuits they contain and the programming that allows the switching paths to be controlled.

INDEX

ABOUT THE AUTHORS

NATHAN COUTINHO is a Solutions Manager for CDW, with a focus on virtualization. He has more than 10 years experience in IT, covering various roles in management, technical sales and consulting. His current responsibilities include evaluating and educating clients on trends and directions in the server, client and storage virtualization spaces. »

ERIC RIVARD is a Network Solutions Architect for CDW. Eric holds a B.S. in information technology and is finishing his M.B.A. He also holds numerous certifications, including: Microsoft Certified Systems Engineer (MCSE), CheckPoint Certified Security Engineer (CCSE), and Cisco Certified Network Professional (CCNP). Eric has written three books for CiscoPress. »

HOWARD WEISS manages the Network Solutions team in the western half of the United States. Throughout his 10-year tenure as a technologist at CDW, Howard has helped build multiple teams from the ground up, including the IBM presales team, Field Solutions team, the HP FieldSE team, and now the Network Solutions team.

TODD SCHNEEBERGER is a Product Manager for CDW. With the company (formerly Berbee) for more than 11 years, Todd has a long history in system and storage technology including roles as a circuit board/packaging engineer, a CAD/CAM technical SE, and a System p and TSM presales engineer. Todd's current focus is IBM System p and storage hardware, along with IBM's Tivoli Suite of storage products. He works out of Milwaukee. »

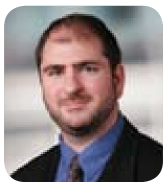

VIRTUALIZATION AND INFRASTRUCTURE OPTIMIZATION

REFERENCE GUIDE

090506 • Flyer 59330AB

LOOK INSIDE for more information on:

- Maximizing server virtualization gains
- Developing a client virtualization strategy
- Increasing power and cooling efficiencies
- Designing a resilient network