VIRTUALIZATION AND INFRASTRUCTURE OPTIMIZATION

REFERENCE GUIDE

How to make the most of your infrastructure with virtualized environments.

CDWG.com/virtualizationguide | **888.693.4239**

CDW·G®

The Right Technology. Right Away.®

VIRTUALIZATION AND INFRASTRUCTURE OPTIMIZATION REFERENCE GUIDE

TABLE OF CONTENTS

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
BUILDING FOR THE FUTURE

CHAPTER 1:

IT managers are constantly on the lookout for ways to make their networks more efficient, secure and flexible, as well as opportunities to cut down on complexity. This demand for improved performance is a strong motivator among IT manufacturers to fine-tune existing technologies and develop pertinent, new technologies.

This commitment to innovation really paid off with the spike in data center growth in the 1990s. This rapid expansion laid the groundwork for government and educational organizations' full embrace of and reliance on digital communications and operations today.

But this expansion was not without its drawbacks. Supporting a sprawling physical infrastructure entails high costs, not only in terms of hardware, power and cooling, but also in management and maintenance.

In addition, many organizations were underutilizing their servers — on average as low as 15 percent of their total capacity, according to the International Data Corporation (IDC). These servers often lacked sufficient disaster or failover protection, and they frequently lacked adequate security.

THE BIRTH OF VIRTUALIZATION

These numerous shortcomings became untenable for many organizations. Manufacturer innovation led to a solution for server sprawl: virtualization. Virtualization is a method of dividing a computer's resources into multiple environments to create a flexible, easy-to-manage, secure computing system.

IT managers began to see the potential for using server virtualization technology across other areas of the network, such as storage and client-level computing. While server virtualization remains the most common type of virtualization, this technology increasingly has applications across many other areas of IT.

Virtualization can improve many of the components that comprise today's computing environment. Two of the most important areas where virtualization can have a significant effect are servers and clients (such as desktops and notebooks).

VIRTUALIZATION VALUE

IT managers have embraced virtualization with enthusiasm because of its potential to reduce costs and data center complexity. Reduced costs have come in the form of lower utility expenses, hardware maintenance and support costs, as well as longer hardware refresh cycles.

With newer focal areas (such as client virtualization) now reaching maturity, IT managers are finding ways to employ virtualization technology throughout their organizations. For example, virtualization can improve performance and reduce costs in the disaster recovery process without incurring any additional costs.

Another great value of virtualization is that it allows for the separation of resources (or dependencies). This enables greater control and movement of resources to different situations 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.

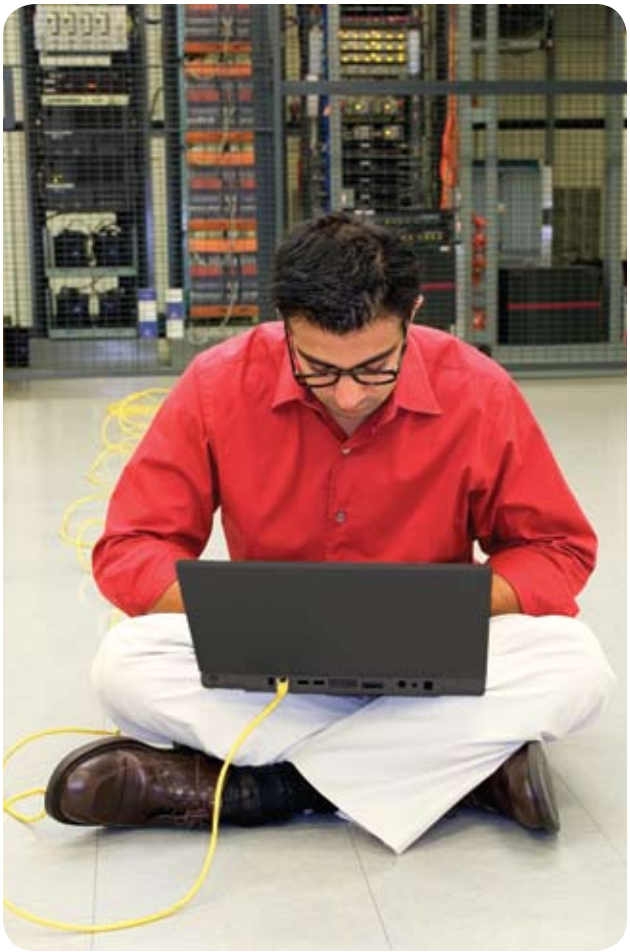

CH 01

OPTIMIZATION OFFERS EFFICIENCY

Optimization as a strategy to improve network performance is another IT practice that has improved the way that organizations operate. For many years, systems standardization had many 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 to between one and three models of PCs, servers and networking equipment became quite popular for organizations.

This approach certainly has its merits when running an IT department; it helps keep costs down. However, this approach also presents certain challenges. With PCs, every user benefits from having an individualized PC, but the numbers and types of applications used on these computers quickly proliferates.

Then 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.

The answer to this interconnected problem? The best approach is a strategy of optimization, implemented in conjunction with virtualization. A key tenant of optimizing the network is finding ways to do more with less. (Virtualization is one means of achieving that goal.) Using a much smaller set of hardware components (thanks to virtualization), optimization can be applied to PCs, servers, storage or networking traffic.

OPTIMIZATION GAINS

The benefits offered by optimization and virtualization far outweigh the value of keeping network systems separate. An optimized IT environment provides: a refined backup and disaster recovery (DR) methodology, keeping your organization safer; a much lower power consumption rate, helping your organization reduce power costs and go green; and a simpler suite of data center management tools, providing greater control over IT assets.

During the migration from mainframe computing to application servers and networked PCs, IT managers faced a revolutionary change in how their operations were organized and run. Today, IT teams confront 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.

Is Your IT Staff Prepared?

In a 2008 IDC Research Service survey of 464 IT managers, 44 percent said inadequate skill levels create the biggest hurdles for virtualization. Make sure your technicians are up to speed about what's required to go virtual.

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 (including virtualization) that have a tremendous return on investment. ◊

SERVER VIRTUALIZATION
INCREASING SERVER EFFICIENCY

CHAPTER 2:

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 CHALLENGES

The IT industry has evolved dramatically over the past decade, along with 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, each of which virtualization can potentially remedy. What follows is a listing of the numerous positive effects that virtualization can bring to an organization.

Increase TCO: Perhaps the most common requests heard by IT managers 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.

The primary reason so many organizations have adopted virtualization strategies over the past few years is because of the reduced TCO realized by implementing a virtual infrastructure. Some of the areas where savings occur are hardware and software costs, operational costs, and consulting services costs.

Reduce server sprawl: Data centers have come to rely on x86 servers to support organizational growth and today's more widely distributed operations. 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. Virtualization addresses this issue by reducing the number of physical servers in the data center.

Improve server utilization: Industry-standard x86 systems continue to be the best-selling server platform. Processor, memory, network and disk speeds continue to increase, while technologies such as multicore processor architecture continue to improve hardware performance. However, few operating systems and applications benefit from these performance increases.

This is why the average server utilization rate for most organizations falls between 5 and 15 percent, according to Microsoft. Low utilization on an endless group of servers has made many organizations rethink their IT strategies and implement a control structure to avoid server sprawl. Virtualization allows you to gain a higher degree of use out of each physical server.

Simplify 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. Even though newer physical-to-physical technologies exist today, they are often not cost effective and offer unpredictable results. Virtualization can greatly simplify this migration process for organizations.

Speed up 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. With virtualization, you can greatly reduce the time it takes to deploy applications to the organization.

Increase high availability for critical systems: High availability is a series of measures undertaken to implement minimal to near-real-time failover for a particular application. Because of 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 to 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). Virtualization technology can help increase the high availability of important systems.

Reduce disaster recovery complexity: Disaster recovery is a series of measures undertaken to implement minimal to near-real-time failover 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, building a disaster recovery site increases in complexity as the size of the data center grows, making it expensive to implement and maintain. Virtualization technology can simplify and greatly speed up the disaster recovery process.

Easier licensing management: Managing the number of server and application licenses to stay in compliance has become difficult for many organizations to track and maintain. As a result, you may not have a good handle on whether you are in compliance or not, and you may struggle to maintain different maintenance contracts that expire at different intervals. Server virtualization can contribute to a reduction in the amount of time you'll need to manage licenses.

Improve power and cooling efficiency: Data centers were originally built at a 1KW to 2KW per-rack average. Today, newer systems such as blade centers, multiprocessor systems and storage arrays 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 because of 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 prolong the lifecycle of the data center.

Going Green

Some utility companies in the United States are now issuing credits to organizations that implement virtualization, which allows them to reduce their power consumption. Similarly, there is a flurry of industry activity focused on finding 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.

SERVER VIRTUALIZATION BEST PRACTICES

Perhaps the most difficult part of rolling out a virtualized environment is crafting the right financial plan to convince stakeholders to do it. Once the environment is up and running, it's relatively easy to manage. Here are some best practices to help

your organization maintain a stable and functional virtualization environment.

Tap templates and customizations for rapid deployment. Virtualization gives you the capability of creating a virtual server or desktop once, saving it as a template, and using customization scripts to make a mass rollout of similar servers and desktops very easy.

Without using customization and scripts, you would have to clone the master image manually, then boot up and make changes to each cloned image. Depending on the image size, this process could take an hour or more. With customization and scripts, you will need only a few minutes to create your custom specifications and execute your script.

Build clusters of physical servers with shared storage. Take advantage of your virtualization software's native clustering and disaster avoidance capabilities by attaching your physical servers to a storage area network (SAN). If you're currently clustering database and application servers, you may be able to greatly reduce and simplify those clusters.

Hardware upgrades for a cluster are as simple as migrating the virtual machines (VMs) from the old hardware, shutting down the old server and powering up the new server. Using a SAN also allows faster and easier disaster recovery if your physical server fails.

Have your IT team brush up on their Unix skills. In most small IT shops, Microsoft Windows rules. But if you are (likely) virtualizing with VMware, IT team members will benefit from a Unix refresher course. Many of the tools require Unix code to carry out transactions or to complete operations.

Use pre-built virtual machines. Ask the manufacturer if it has any pre-built VM images that you can deploy directly to your virtual environment. There is a wide variety of pre-built VMs available for web servers, databases, systems monitoring and more.

You can save a significant amount of time and reduce the possibility of human error by using a pre-built image when migrating to a new platform or system.

Create a cost-effective test environment. You're going to need a dedicated test environment for your critical systems, even after they're virtualized. In fact, all critical systems should have a test and development environment, but even your non-mission-critical applications can benefit from snapshots.

A snapshot allows you to save a virtual machine at a point in time. You can then begin moving the previously tested system changes to production. If the new patch proves troublesome, fixing it is as easy as rolling back to an earlier snapshot.

Control virtual machine sprawl. One of the pitfalls of virtualization is the ease with which new virtual machines can be created. This can lead to significant inefficiencies in your virtual environment. It isn't necessary to create a separate virtual server for everything.

For example, you don't need to create a new virtual Apache web server for every web application; just use Apache's virtual host configuration. Every unneeded virtual machine uses system resources that can be used elsewhere and adds to the number of virtual machines that your staff will need to manage.

One approach to controlling virtual machine sprawl is to assess the performance requirements of a new service. If the load is expected to be low, and you have an existing virtual machine that can handle it, then use the existing virtual machine. You can use the built-in performance monitoring tools, which indicate when it's time to begin splitting services onto new virtual machines. Proper change and configuration management is key.

Migrate your existing physical servers to virtual servers without the hassle. Most virtualization software has some form of physical-to-virtual system migration tool, such as VMware Converter, which can copy a physical server bit for bit to your virtual environment.

There are typically a few different versions of the tool with varying capabilities. Often the migration tool will be free, if only in a limited version. Assess your needs to determine which tool can do the job well.

Consider doing a storage upgrade. For most organizations, the footprint reduction that virtualization allows pays dividends, such as reduced power and cooling requirements that add up to significant savings.

With virtualization, organizations should consider network-attached storage — not from a tech perspective, but because you miss many of the benefits if you don't. The chief benefit is the ease of migrating VMs on the fly, followed closely by the ability to automate VM movements to ensure high throughput and failover.

Although it might mean more of an initial investment, it's still a huge cost savings on the back end because of the reduced server requirements going forward and the immediate access to the full suite of virtualization capabilities.

To help justify the spending, be sure you factor in downtime. With virtualization, networked storage and high availability you avoid the cost of downtime. It's easy to overlook because it's not easily calculable, but it really needs to be factored in.

VIRTUALIZATION MANUFACTURERS

A server virtualization solution consists of three main components: software, servers and storage. VMware stands at the forefront

of virtualization software today. The manufacturer's 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 options to choose from including IBM, NetApp, EMC, and HP's LeftHand Networks.

TROUBLESHOOTING TOOLS FOR VMWARE ESX

Learning the finer points of managing a virtual environment can take some time. To help you along in this process, here are some go-to tools to assist in troubleshooting any problems that come up with a VMware rollout.

Virtual Infrastructure Client (VIC): VIC can give you full access to all of your virtual machines — to diagnose or power them down. VIC can be installed on any system and is available for download by connecting a web browser to an ESX host.

Once you have it installed, connect directly to an ESX host by pointing VIC at the IP address (or host name) of the ESX host server, and log in using the ESX host's root user login.

Virtual Center Server: The vCenter Server gives you the ability to treat your entire VMware infrastructure as one single unit or groups of units. One small tip: If you intend to use vCenter with Microsoft SQL Server, make sure you are running at least the Standard Edition of SQL Server. You will run into several problems if you have the Express Edition.

PuTTY: This third-party remote-session tool can be used to connect directly to the file structure of an ESX host server using the Secure Shell protocol to run console commands or modify configuration files in a Unix/Linux command format. For details about setting up remote shell access, consult VMware's online knowledge site at www.vmware.com/support.

Secure Copy: With Secure Copy applications such as Veam FastSCP and WinSCP, you can move files across the network to and from an ESX host, or to a storage area network that acts as shared storage for the infrastructure.

Remote Desktop Connection: Because VMs are nearly identical to bare-metal systems, why not treat them as such? Use RDC to connect to the IP address or host name of any VM to take advantage of native tools such as Event Viewer, Command Prompt and Control Panel.

Sysprep: Microsoft Sysprep is more about trouble prevention than troubleshooting. After using vCenter to create one VM, you can use that VM as a template to produce an endless army of clones. Sysprep lets you customize each one and change important system info, such as the host name, at the time of creation.

CONSIDERING HYPER-V

While VMware is the dominant virtualization manufacturer, Microsoft is gaining ground in this market. Each virtualization product has its benefits and drawbacks. Depending on the needs of your organization, you may find that one solution works better for you than others. So it's worth knowing a bit about what other solutions are out there beyond VMware.

Microsoft's release of Hyper-V in Windows Server 2008 illustrates the shifting conditions in the virtualization market. Virtual machines have become so prevalent today that Microsoft includes the technology as a standard tool within its server OS. And when Microsoft ships something standard, the technology has gone mainstream. Here are three reasons your organization may want to consider Hyper-V.

1. **The price is right.** When it comes to saving money, you can't beat free. The software is part of Windows Server 2008. You can also download Hyper-V and get started with it at no cost.

   In today's economic climate, money matters more than ever to IT professionals. The numerous inherent savings of going virtual, along with no up-front cost for starting with Hyper-V, make this product an option worth strong consideration.

2. **Better suited for your specific disaster recovery needs.**

   Hyper-V is not quite as automated as some of the competition. For example, a disaster recovery operation with Hyper-V might take you five minutes, whereas it might be nearly instantaneous with VMware. You may find that many of your operations don't require the advanced recovery capabilities that VMware offers. Your organization may be okay if users and apps are down for a few minutes.

3. **You know the drill.** That familiarity — knowing which button to press when working with Hyper-V — leads to the final reason why users opt for Microsoft's virtualization technology: If you're already running Windows, then your IT administrators already have competence using Microsoft's system management tools.

   VMware may offer a superior feature set, say users, but Hyper-V offers a more than adequate set of its own. You will have to determine what exactly your organization's virtualization needs are to decide if Hyper-V or another product is right for you. ◊

CLIENT AND APPLICATION VIRTUALIZATION

IMPROVING RESOURCE DELIVERY

CHAPTER 3:

Desktops, notebooks and other client devices are in use throughout the workspace. 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. For some organizations, client virtualization may be the answer, simplifying IT management and lowering administration costs.

THE BENEFITS OF CLIENT VIRTUALIZATION

Organizations that have a variety of client devices can look to gain the following benefits from a virtualization initiative:

- **Rapid deployment:** The timeframe to deliver virtualized 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.
- **Improved 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.
- **Reduced hardware costs:** Organizations can expect lower hardware costs as a result of moving to less-expensive thin clients, reduced support costs and extended systems lifecycles.

CLIENT VIRTUALIZATION COMPONENTS

The client virtualization market is quite diverse compared to the server virtualization market, with more technology options to choose from and a wider pool of manufacturers. Many organizations combine multiple virtualization components to meet the requirements that satisfy their operational needs.

DEVICES

When building a client virtualization solution, the most important consideration is the end-user experience. You need to determine the audience not only in terms of what devices end users will operate on (such as notebooks, desktops, thin clients, tablet PCs and mobile devices), 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, Windows Vista or Windows 7. However, there are numerous ways to present an operating system to the organization's 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, the 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.

Another possible solution is hosting the Windows images on a server and then streaming the operating system to existing desktops. This is 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.

APPLICATIONS

Citrix revolutionized the way applications are delivered with its WinFrame and MetaFrame software, and the manufacturer's 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.

These solutions host applications on a Windows server, then deliver them to users via sessions, enabling 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 the 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 Dynamic Link Library (DLL) conflict issues that are 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.

Another option to deliver applications is the traditional method of installing software locally, either manually or with a managed software delivery technology such as Symantec's (Altiris) 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

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

Printing: This operations 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, security is an essential consideration because you may be delivering access to critical applications and data via the Internet to devices that the organization may not have any control over. 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 more than a decade, many organizations have used imaging technologies to deliver a consistent image to desktops and notebooks. Symantec Ghost is probably the best-known tool, although Microsoft and a number of other manufacturers offer competitive 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. This solution 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 at the next access.

ADDITIONAL PLANNING CONSIDERATIONS

For most organizations, client virtualization can be a good fit. But it is a new approach to an important part of your organization's operations. As such, there are several important issues to consider regarding this approach as you architect your solution.

As you plan your *application architecture*, you will want to consider the following:

- **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, 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 applications (softphones or video conferencing, for example), 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 and biometric scanners. Test these devices properly with the virtualization solution of your choice.

As you plan your *infrastructure architecture*, you will want to consider the following:

- **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 is 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.

As you review your *remote access requirements*, you will want to consider the following:

- **Devices:** Organizations today have a variety of devices to choose from (such as desktops, notebooks, tablets and PDAs). 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.

As you plan out your *operating system and application licensing strategy*, you will want to consider the following:

- **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.

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.

APPLICATION VIRTUALIZATION

The enormous amount and diversity of applications used by organizations can make managing them quite a challenge. When you separate the application from its underlying components and view the work from a layered perspective, there's the application itself, the user workspace, storage, operating system and hardware.

Many managers today are convinced that virtualizing their desktop environment will also solve their application layer problems. But does it? A closer look at the problems affecting traditional software installations reveals that a different strategy, virtualizing at the application level, may be needed.

TYPICAL APPLICATION SHORTCOMINGS

Looking at the shortcomings of the traditional application layer model reveals why more IT organizations realize virtualization is an attractive option. Here are some examples.

- **Application incompatibilities:** The more applications installed on a machine, the higher the possibility of application incompatibilities because of varying prerequisites in some cases.
- **Patching or upgrading apps:** Unless the infrastructure is already in place to update or patch the installed apps, reimaging the computers to apply an upgrade is the optimal approach. Even with a solution that pushes out updates or patches, a tremendous amount of regression testing is needed before deploying the new image.
- **Failure to run different versions of the same application:** For many organizations there's a need to run multiple versions of the same application to ease the transition for users. Running multiple versions of the same application helps train users on the new version and reaching a certain comfort level before retiring the older version.

APPLICATION VIRTUALIZATION DEFINED

Application virtualization is the abstraction of the underlying layer (the operating system). This is accomplished by tricking the application into believing that it's residing and interacting directly with the OS, when in fact it's running in its own isolated virtual environment.

Virtualized applications can run in one of two environments: either on client machines using local resources or on the terminal server, where processing cycles are provided by the server and the end user establishes a remote session.

Most of the problems brought up earlier regarding the application layer can be resolved through application virtualization.

Because virtualized apps run in their own discrete environments, each application responds as if it's running alone in the OS — thus, application incompatibilities don't arise.

In addition, this "bubble" effect makes it possible for each application to be bundled with all its prerequisites, thus allowing the end user to run multiple versions of the same software on the same machine.

Because applications are installed locally on the computer, upgrading or patching an app becomes simply a matter of repackaging and distributing it to clients.

Furthermore, the modular/layered approach of decoupling apps from the OS offers IT managers the freedom to investigate solutions to better manage the other layers without the hassle of trying to find a nonexistent, all-encompassing solution.

Based on the vendor and application virtualization solution in place, some of the added benefits are the centralization of software distribution, application usage statistics, license management and the ability to manage application access and authorization.

Prior to application virtualization, one of the major obstacles to server-based computing was offline use. Most of the existing application virtualization solutions handle offline use by streaming and caching the application locally, or by creating a complete self-encompassed package that can run with no strings attached.

APPLICATION VIRTUALIZATION BEST PRACTICES

Any new system requires a learning curve and comes with some shortfall. Application virtualization is no different. Here are some best practices to help smooth the adoption process.

Decide carefully which applications to virtualize. Before jumping into application virtualization, understand that no matter which solution you choose, not all applications can be virtualized.

Similar to server virtualization, in which a physical server might be required for some deployments, some apps have to be installed and can't be virtualized. These are mainly apps that deal with system-level drivers or, in some cases, COM+ objects.

Review all your vendors and license agreements. The added benefits of license management, application usage stats and the centralization of application distribution will necessitate a revision of some license agreements, and in some cases may result in a reduction in the number of licenses required.

Many software manufacturers either aren't aware of application virtualization or haven't adjusted their licensing models to take it into consideration. Dealing with the manufacturer can be tricky when trying to determine the optimal number of licenses or renegotiating a license agreement.

Also, some manufacturers have different licenses and pricing to run their software on the terminal server.

Make provisions for training and maintenance. If the virtual app is running locally on the end user's machine, training for end users is quite minimal or perhaps even nonexistent. Regardless of whether an application is virtual or locally installed, all the average user needs to know is how to launch the app, usually by double-clicking on a shortcut.

Some training may be needed when a notebook is involved and offline access to an app is required. In that case, the user would need to cache the app locally before going offline. Training might also be required if the virtual app is hosted on a terminal server and connecting to a remote session is new to your environment.

Take a close look at workflow and software ordering. Application virtualization doesn't necessarily mean centralizing software distribution. But if that's the route you choose, then a re-evaluation of software ordering, distribution, assignment, updating and license negotiation (and possibly a new workflow) is in order.

APPLICATION VIRTUALIZATION SOLUTIONS

Among the virtualization markets, application may be the least familiar to many organizations, but there are many manufacturers that have robust product offerings. Here are some

of the major players in the application virtualization field.

Citrix XenApp: XenApp is an agent-based system that offers client and server-side application virtualization. With the new release of XenApp, if an application can't be virtualized it can be installed on a virtual machine (VM) and in that way provided to the end user.

Microsoft App-V: Microsoft's contribution to the application virtualization market is Microsoft Application Virtualization, or App-V for short. Part of the Microsoft Desktop Optimization Pack (MDOP) offered through Software Assurance, App-V is agent-based software that offers client-side and server-side application virtualization.

VMware ThinApp: Unlike XenApp and App-V, ThinApp is an agentless system that has an embedded virtual OS along with the packaged app.

ThinApp doesn't have an application streaming or distribution mechanism. The packaged application is self-sufficient and can be pushed out like any other file. ◊

GETTING
VIRTUALIZATION RIGHT

An effective virtualization strategy offers a number of potential benefits:

- **Simplified administration.** An effective server virtualization strategy allows an organization to make changes to the computing environment quickly and seamlessly. In many instances, installations and upgrades take place in seconds rather than hours or days.

- **Reduced server and power consumption costs.** Virtualization typically complements a server consolidation strategy and allows an organization to reduce the number of physical servers. Not only is the organization able to reduce the direct cash outlay for servers; it's able to slash its energy bill by running fewer servers and reducing the need to cool data centers.

- **Improved agility and flexibility.** Moving away from the one server-one application model opens the door to a far more robust computing environment — and one that runs different operating systems and applications on the same server. A virtualization strategy often boosts server utilization rates from below 20 percent to upward of 80 percent.

- **Reduced licensing costs.** A number of application providers have introduced virtual licensing agreements. These licenses allow an organization to run virtual instances of an operating system on multiple servers at no extra cost.

- **Improved availability and stability.** Virtualization boosts availability for key systems and applications. If a physical server crashes, other virtual servers can take over and applications running on them remain unaffected. Some virtualization software products automatically set up a new virtual server on a different machine — if a hardware malfunction takes place.

- **Improved disaster recovery.** Whether it's a natural disaster that impacts an entire region, or a transformer malfunction that affects a city block, a sudden power outage or disturbance can wreak havoc on your network. A virtualization strategy allows you to maintain an instant failover plan that provides operational continuity throughout disruptive events. With the right tools, you can enable automated backup, replication, and rapid movement of servers, desktops and applications.

- **Less data center space required.** Not surprisingly, fewer servers translate directly into a reduced need for data center space. In today's compute-intensive environment, gaining control of server sprawl often leads to a substantial savings in real estate and construction costs.

To learn more about server virtualization or to find out if server consolidation is the right choice for your organization, log onto CDW•G.com/virtualization today.

ORACLE

Oracle® Database 11*g* Standard Edition Solution

Increases flexibility and decreases costs

Named user license[1]

CDWG 440283

$350.00

Processor license

CDWG 440289

$17,500.00

- Open standards ensure easy integration with existing infrastructure, maximizing the value of prior investments
- Scalability and interface continuity minimize the cost of future expansion and training
- Centralization and automation reduces management and administration costs

[1]Minimum purchase of five named user licenses required; call your CDW•G account manager for details

redhat
AUTHORIZED VOLUME RESELLER

Red Hat® JBoss® Enterprise Linux® Advanced Platform

Server

Standard 1-year subscription[2]

CDWG 1182278

$1499.00

Standard 3-year subscription[2]

CDWG 1184110

$3971.99

Critical technologies in a single solution, this product is fully integrated and ready to go. No need to purchase additional products from multiple vendors.

- Unlimited virtualized guests — multiple operating systems, multiple versions, multiple performance and security settings, all running on the same server using Red Hat® virtualization technology
- Storage virtualization — multiple guests, running on the same or different servers, access and share logical storage volumes and file systems to reap the full benefits of server virtualization with shared storage
- High-availability clustering and failover — application and guest operating system failover
- Management — includes comprehensive management capabilities

[2]Standard subscriptions include multiple support services providing unlimited incidents and coverage up to 24x7 with one-hour response

Sun Partner Advantage

vmware

solaris

intel Xeon inside
Powerful. Intelligent.

Hard drives sold separately

Sun Fire™ X4270 Rack-mount Server

Optimized for virtualization and core operations applications

Quad-Core Intel® Xeon® Processor X5570 (2.93GHz)

CDWG 1860691

$4548.59

- Memory: 2GB std., 144GB max. (DDR3)
- Hard drives: none ship std.; up to 16 hot-swappable SAS/SATA disk drive bays available
- 8MB Level 3 Cache
- Four integrated 10/100/1000BASE-T Ethernet ports

EMC²
VELOCITY SIGNATURE PARTNER

EMC® Celerra® NX4

Unified storage platform

CDWG 1542955

CALL FOR PRICING

- Integrates EMC® CLARiiON® networked storage to deliver no-compromise availability and advanced data protection
- Easy to install and use, requiring as few as 15 minutes to go from power-up to production
- Included functionality like Celerra® Manager, Virtual Provisioning, Automated Volume Management and Celerra SnapSure™ simplify common administrative tasks

CDW•G Services

CONTACT YOUR CDW•G ACCOUNT MANAGER FOR DETAILS.

Product lifecycle support: We're committed to helping you optimize your IT investment throughout its lifecycle with a range of services designed to offer the support, assistance and protection you need for your technology solutions.

Solutions include:

- Virtualization assessments
- Server and storage virtualization and optimization
- Continuity of operations
- Backup/recovery services

ACHIEVE

MAXIMUM UPTIME

WITH A REDUNDANT NETWORK.

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.

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.

Redundancy, also known as resiliency, in a network provides a maximum level of network uptime. Redundant networks help to ensure the availability of critical applications to staff, partners and customers through an optimized network design. By providing multiple pathways to your data, a redundant network eliminates single points of failure in the event that network equipment fails or becomes unavailable.

A redundant network can help your organization achieve multiple benefits, including:

- Multiple pathways to ensure application availability
- Maximum bandwidth utilization
- Improvement of disaster recovery posture
- Increased network management and reporting

When considering the implementation of a redundant network, it's important to:

Conduct a network audit. Inventory all hardware and software components within your organization to identify missing product patches, multiple software versions, product end-of-life or end-of-support deadlines, and any factors that may indicate possible security risks or network performance issues.

Perform a network assessment. An assessment maps all network devices, links and protocols and provides a detailed view of the complete LAN and WAN design. It can identify critical network issues and provide a remediation to address them.

Develop a strategy plan. Consider all of your organization's operational needs to ensure your redundant network has the flexibility to support future growth. Evaluate timelines, availability expectations and prioritization of each project component.

To learn more about how network optimization solutions can improve the way your organization operates contact your CDW•G account manager or visit us online at CDWG.com/optimization

Cisco® 2811 Integrated Services Router

Secure network connectivity for data, voice and video

CDWG 699711

$1775.54

- Wirespeed performance for concurrent services such as security and voice, and advanced services to multiple T1/E1/xDSL WAN rates
- Antivirus defense support through Network Admission Control (NAC)
- Network Infrastructure Protection — protects the network infrastructure from attacks and vulnerabilities, especially at the network level
- Intrusion Prevention as well as stateful Cisco IOS® Firewall support and many more essential security features
- Analog and digital voice call support

Cisco Catalyst® 3560-24PS

24-port 10/100BASE-TX managed, rack-mountable switch with two SFP-based Gigabit Ethernet ports and Standard Multilayer Software Image (SMI) installed

CDWG 595015

$2690.25

- Ideal access layer switch for small enterprise wiring closets or branch office environments
- Fixed configuration, IEEE 802.3af and PoE with Fast Ethernet configurations provide availability, security and Quality of Service (QoS) to enhance network operations
- PoE for IP telephones, wireless access points, video surveillance, building management systems and remote video kiosks
- Advanced QoS, rate limiting, access control lists, multicast management and high-performance IP routing

Cisco Catalyst 3750E-48TD

48-port 10/100/1000BASE-T Gigabit Ethernet managed stackable rackmountable switch with two X2-based 10 Gigabit Ethernet ports

CDWG 1162497

CALL FOR PRICING

- Smartports enable fast and easy configuration of Cisco recommended best practice security and QoS features
- Dynamic Host Configuration Protocol (DHCP) autoconfiguration of multiple switches through a boot server eases switch deployment
- Automatic QoS (AutoQoS) simplifies QoS configuration in voice over IP (VoIP) networks by issuing interface and global switch commands to detect Cisco IP phones, classify traffic
- Dynamic Trunking Protocol (DTP) facilitates dynamic trunk configuration across all switch ports

Cisco ASA 5510 Adaptive Security Appliance

With 250 IPsec VPN peers, three 10/100BASE-TX interfaces and 3DES/AES

CDWG 792590

$2735.14

- Delivers advanced security and networking services for remote/branch sites in an easy-to-deploy, cost-effective form
- Allows you to manage and monitor these services through the integrated, web-based management application
- Maximize operational continuity — a license upgrade can maximize your continuity by expanding VPN capacity to support a greater number of concurrent VPN connections for mobile users and remote offices
- Encryption algorithms include DES, 3DES and AES

THE FACTS ABOUT
CLIENT VIRTUALIZATION

What is Client Virtualization?

Similar to the proliferation of servers in the data center, the increased use of desktops, notebooks and other devices has become the norm for most organizations. This creates a number of challenges including:

- Desktop/notebook device security
- Increased device management
- Increased support and maintenance costs due to sprawling client systems and images

The idea behind client virtualization is to decouple hardware and software components so that isolation can occur. Security can also be achieved since no data is actually stored on the end-user's device. All data can be securely stored in the data center, enabling both high availability and disaster recovery capabilities.

The Benefits of Client Virtualization

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.
- **More efficient control:** Thin clients centralize control of all desktop operating systems and applications, so virus outbreaks can be more easily managed and controlled, and issues can be resolved from a single console, enabling IT managers to focus on more strategic initiatives.
- **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.

CDW•G delivers much more than just products. Our comprehensive approach enables your organization to reap the benefits of the best technology, expertise and ongoing support services.

Call your CDW•G account manager today or visit us online at CDWG.com/optimization to find out if client virtualization is the right choice for your organization's needs.

CITRIX PARTNER Independent Software Vendor

NEW version

Citrix® XenDesktop™ 4

The virtual desktop — for everyone

CALL FOR PRICING AND LICENSING OPTIONS

- Simplify desktop management by centrally applying updates once for all users
- Improve security by centralizing desktops and data in the data center
- Deliver desktops instantly to users anywhere while providing a high-definition experience

hp invent

AMD Sempron™

HP SMART BUY

HP Compaq Thin Client t5730w

Provides convenient access to server-based computing solutions

CDWG 1781027

$519.99[1]

- AMD Sempron™ Processor 2100+ (1GHz)
- Memory: 1GB
- Windows® XP Embedded
- Three-year limited parts and labor warranty

[1]After HP Smart Buy instant savings; savings may vary based on channel and/or direct standard pricing; available as open market purchases only; call your CDW•G account manager for details

Microsoft

For display only

Microsoft® Windows Server® 2008 R2 Standard

Built-in server virtualization technology

Open License Government[2]

CDWG 1826220

$558.59

- Enables your organization to reduce costs, increase hardware utilization, optimize its infrastructure and improve server availability
- Security innovations provide high levels of protection for your data

[2]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

hp invent

intel Xeon inside™

Powerful. Intelligent.

Hard drives sold separately

HP StorageWorks® MSA2312i Dual Controller Array

CDWG 1779056

$8433.85

- 1Gb Ethernet (1GbE) iSCSI connected array
- Grow storage as demands increase up to 29.7TB SAS or 60TB SATA, supporting up to 32 hosts for iSCSI attach
- Supports the optional HP StorageWorks® 2000 Modular Smart Array Snapshot Software which offers controller-based snapshot and cloning functionality

PERFORMANCE MODEL

HP ProLiant DL380 G6 Rack-mount Server

Powerful management tools to optimize and control all aspects of your infrastructure

Two Quad-Core Intel® Xeon® Processors X5550 (2.66GHz)

CDWG 1723259

$5750.58

- Memory: 12GB std., 192GB max. (PC3-8500R DDR3)
- Hard drives: none ship std.; up to eight SFF hot-pluggable SAS/SATA drive bays available, 6TB max. storage
- 16MB Level 3 Cache
- Two Embedded NC382i Dual-Port Multifunction Gigabit Server Adapters

RAMP UP PERFORMANCE WITH APPLICATION VIRTUALIZATION.

What is application virtualization?

Application virtualization separates the application configuration layer from the OS in a desktop environment, reducing application conflicts, bringing patch and upgrade management to a central location and accelerating the deployment of new applications and updates.

How can application virtualization help my organization?

Managing numerous desktop images drives up IT cost and consumes significant IT resources, as does maintaining up-to-date applications and security patches. By separating the application layer from the operating system on the desktop image, an IT organization can remove much of the complexity and management challenges it faces every day.

Significant application-to-application regression testing is typically required before deploying a major application or OS update. In an environment where applications are virtualized and delivered on demand to the desktop, application-to-application conflicts are nearly eliminated, since the applications execute without impacting the desktop operating system — or other applications. In addition, application maintenance and security patching is streamlined, since it only needs to take place at the source, not hundreds or thousands of individual desktops. This not only reduces labor needed for desktop maintenance, but ensures that any security vulnerabilities are adopted immediately and comprehensively.

Why Citrix XenApp?

Citrix XenApp is an application virtualization solution that lets organizations reduce the cost of Windows application management by up to 50%.

Citrix XenApp enables organizations to improve application management by:

- Centralizing applications in the data center to reduce costs
- Controlling and encrypting access to data and applications to improve security
- Delivering applications instantly to users anywhere

Application virtualization is scalable and secure

To enable an enterprise-class infrastructure for any size organization and ensure that virtual application delivery and management is secure by design, Citrix XenApp integrates the tools and infrastructure needed to help you implement a scalable solution and to control, measure and monitor performance to ensure that service level agreements and security requirements are met.

Contact your CDW•G account manager to find out how application virtualization can benefit your organization.

CITRIX

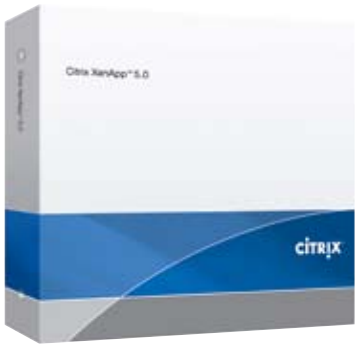

Citrix® XenApp™ 5.0

End-to-end applications-delivery system

Single-user connection license With Subscription Advantage

CDWG 1176409

$600.00

Citrix® XenApp™ 5.0 offers both application virtualization and streaming to enable access for any user, with any device, working over any network. Organizations can centralize applications and data in secure data centers, reducing costs of management and support, increasing data security and ensuring fast, reliable performance.

- Accelerate graphics performance with SpeedScreen™ Progressive Display
- Provide easy, secure application access from locked-down devices
- Record and report all Presentation Server configuration changes

lenovo

Lenovo ThinkPad® T400

Built for organizations — featuring Lenovo ThinkVantage® Technologies

CDWG 1894498

$898.99

- Intel® Centrino® 2 Processor Technology
- Intel® Core™2 Duo Processor P8600 (2.40GHz)
- Intel® Wireless WiFi Link 5100AGN
- Memory: 2GB
- 250GB hard drive
- DVD±RW drive
- 14.1" WXGA+ display
- Windows® 7 Professional

hp invent

Hard drives and blade chassis sold separately

HP ProLiant BL490c G6 Blade Server

Effectively removes the constraints that limit virtual machine hosts

Quad-Core Intel® Xeon® Processor X5570 (2.93GHz)

CDWG 1723369

$4003.96

- Memory: 6GB std., 192GB max. (PC3-8500R DDR3)
- Hard drives: none ship std.; up to two SFF non-hot-pluggable Solid State Drive bays available, 1.8TB max. storage
- 8MB Level 3 Cache
- Embedded NC532i Dual-Port Flex-10 10GbE Multifunction Server Adapter

Microsoft

For display only

Microsoft® Windows® 7 Professional

The most advanced Windows operating system for organizations

Open License Government Upgrade[1]

CDWG 1826212

$145.63

- Improve organizational agility
- Increase security
- Reduce cost and complexity
- Make users productive anywhere
- Streamline PC management

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

CDW•G Installation Services

IT departments are stretched thin. Technology costs are unpredictable. The design, implementation and protection of your new technology demands expertise and time that many in-house IT professionals simply don't have.

Our highly-certified specialists will help you get the most out of your technology by providing expert solutions that meet your unique needs with minimal disruption of operations. CDW•G partners with only the finest industry-leading service providers to bring solutions right to your door.

DATA CENTER OPTIMIZATION
FINE-TUNING YOUR DATA CENTER RESOURCES

CHAPTER 4:

Blade Servers

Load Balancing

Power and Cooling Management

Hosting and Managed Services

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 offers one approach to physical consolidation. What makes blade servers different than traditional rack-mount servers is that blade systems are built with a modular infrastructure, allowing them to 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 arrangement. Another key difference is that blades do not have traditional input connections typically 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.

Some setups can fit up to 16 blades into a 10U rack space. That's quite a space savings compared to traditional 1U "pizza box" servers of the past. More recently, 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 128GB 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 savings on power use over traditional servers.

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 for 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:

- the port count should match the organization's needs, especially if the blades will be used for virtualization;
- blade management is a key feature, so 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 to make the most effective use of all those resources, to optimize them.

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

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 are served faster.

Server Farms

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 websites and high-bandwidth FTP (file transfer protocol) sites.

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 back-end 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 20 to 30 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 will want to give careful consideration to failover features, availability and overall performance.

ADDITIONAL FEATURES

Load-balancing appliances today do more than simply balance loads. In addition to significantly improving performance and reducing server loads, these solutions are ideally positioned to examine traffic and find performance and security problems.

Many load-balance solutions also provide intrusion detection and prevention, enabling an organization to move web servers in from the dimilitarized zone and add a layer of security. Some load balancing solutions also perform higher-level application functions, including SSL encryption termination, TCP/IP offload and content-switching.

Load balancers can offer a variety of application-level features and functions. Here are some to consider.

SSL encryption termination: Rather than forcing the application server to deal with the overhead of terminating SSL encryption, most load balancers can offload such tasks, freeing up the application server to do what it does best: serve applications.

Compression: Load balancers save on bandwidth by ensuring that traffic loads are distributed evenly among back-end servers, but they can also perform significant bandwidth savings duties such as compression and traffic shaping.

Content switching: Another key security feature some products offer is content switching, where the load balancer filters for certain data strings (such as credit card numbers or Social Security numbers) and ensures they are blocked before being carried across the WAN. This helps prevent data leakage.

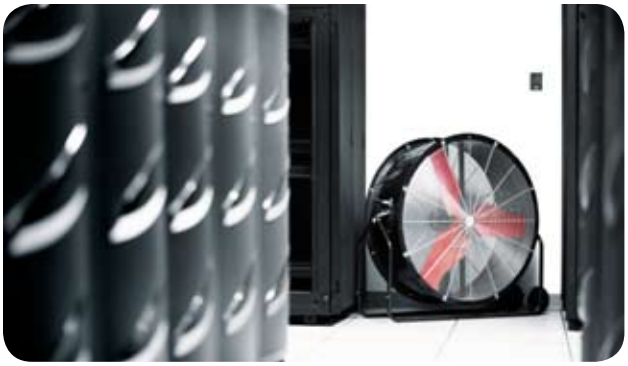

POWER AND COOLING MANAGEMENT

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 cooling solutions to choose from.

Room oriented: This is a traditional 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.

A fifth option is to simply raise the temperature in the data center. Many organizations keep their data center thermostats set in the mid-to-upper 60s, resulting in an increased workload for the air-conditioning units. Raising the temperature a few degrees to 70 to 72 degrees can reduce the cooling bill.

HOSTING 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 kinds of hosting 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

Hosting 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 hosting service works both ways: services increasing or decreasing in response to an organization's strategic initiatives. ◊

STORAGE OPTIMIZATION
MANAGING DATA DEMAND

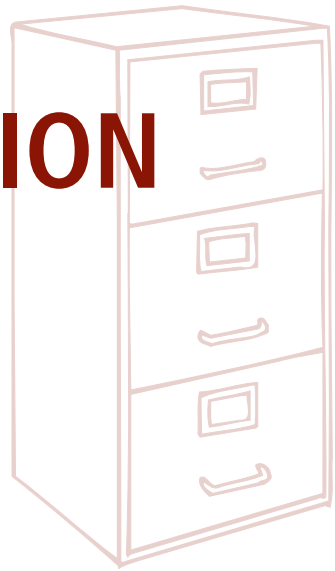

CHAPTER 5:

Storage Virtualization

Storage Area Networks

High Availability

Virtualization technology has revolutionized how IT departments can manage their networks. This is especially true in the storage market, where technological advances have steadily improved an organization's ability to keep up with the rapid pace of increasing storage demands. While virtualization is not the only viable approach to storage, it is worth taking a close look at to see if it's right for your organization.

STORAGE VIRTUALIZATION

The basic concept of storage virtualization is that it unites multiple storage devices into what appears to be a single storage pool. By abstracting the many physical devices into one logical layer, virtualization makes it easier to centrally manage and back up data.

When this approach is combined with server virtualization (referenced in Chapter 2), both the computing environment and its associated data can be seamlessly moved through the physical environment, creating little if any downtime and high availability in computer processing environments.

While storage virtualization is not a new concept, what's new and driving interest in it today are the numerous storage management features that are accessible through arrays and appliances that, when combined with storage virtualization, offer great benefits to organizations.

Technologies such as thin provisioning, data deduplication, continuous data protection, snapshot and replication technologies, and data tiering can all bring tremendous value to an organization's storage strategy in conjunction with storage virtualization.

Some of the benefits that organizations can derive from a storage virtualization strategy include the following.

- **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 turn over less frequently.

Organizations looking for faster data recovery often turn to storage virtualization. Virtualizing storage allows organizations to carry out a snapshot at the block level of data, facilitating faster file recovery and data backups that use less throughput on the network. You can also take virtualized storage, copy it over, and have a test environment within minutes.

Beyond the additional management functions, scalability is another popular feature driving the adoption of virtualization. Not wanting to either over- or under-allocate storage, organizations

need flexible storage options that give them what they need when they need it. Virtualized storage makes it easy for you to start minimal and increase easily as needs increase.

When additional storage is needed, it can be added or reallocated on the fly, without downtime or interruption. You simply add disks to the array and you have more available storage. You don't have to create new arrays or rely on software applications to group storage into logical units.

VIRTUALIZE AT THE FILE OR BLOCK LEVEL?

Perhaps the biggest decision that organizations will make in selecting a storage virtualization device is whether to go with file-level or block-level virtualization. Simply stated, file virtualization deals with network-attached storage (NAS) and file servers and is built around virtualization of files and file systems; block-level virtualization is focused on storage area networks (SANs) and virtualizes blocks of data as opposed to individual files.

Organizations tend to lean toward block-level virtualization, which lets administrators build any type of file system on top of it. The biggest advantage of this approach, though, is that there is a reduction in the actual overhead. File-based protocols are semantically rich, which means there's a lot more overhead. Block-level virtualization is particularly useful for database applications, which work best when they have unencumbered access to the raw disk.

But file-level virtualization and NAS may work better for other types of applications. Web work favored NAS deployments initially because they enabled you to coordinate access to files as you scaled out a web service by adding more servers. With block-based storage, you still need a file system in there that will allow that coordination. So file-level virtualization may be a better fit in that scenario.

DATA DEDUPLICATION

When having to do a series of incremental and full backups, you will inevitably capture the same data over and over again. With the high cost per gigabyte to manage and store this data, finding a way to eliminate that redundancy can save organizations lots of money. Enter data deduplication.

Data deduplication searches for duplicate segments of data and replaces them with a pointer to a single instance of that information. For example, if there are 50 copies of an e-mail attachment, the storage system will retain only one.

This dramatic reduction in stored data in turn drives big savings: Less data requires fewer servers to process the data, and fewer tape and disk drives to store it. Organizations also save on related expenses such as heating, cooling, space, personnel and bandwidth.

Data deduplication can take place either inline or post-processing. Inline data deduplication technology examines data as it enters the box before it's stored to disk, whereas the post-processing technique deduplicates data after it's stored.

The inline approach has proven valuable because everything is done in the CPU and memory and not on disk, so you only need a fraction of the disk space. The smaller data pool also speeds the time it takes to write data to tape, which makes disaster recovery more efficient.

STORAGE AREA NETWORKS

While storage virtualization's popularity has blossomed, networked storage continues to hold a firm presence in the storage market. With networked storage, the data already resides on a shared disk array, which speeds the storage process and allows organizations to take greater advantage of virtual machine technology.

There are two varieties of networked storage: basic NAS and higher-speed SANs. 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 its high speed. SAN solutions move applications across a network 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 growing more popular as components drop in cost and complexity. In addition, the proliferation of iSCSI (an IP-based connection standard) in storage devices makes a SAN even more desirable.

One disadvantage of SAN technology 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.

Also 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.

FAST FIBRE CHANNEL FOR SAN

Organizations adopting server virtualization may soon be looking to give their SANs a speed injection via 8-gigabit-per-second Fibre Channel. Manufacturers such as Brocade, Emulex, Hewlett-

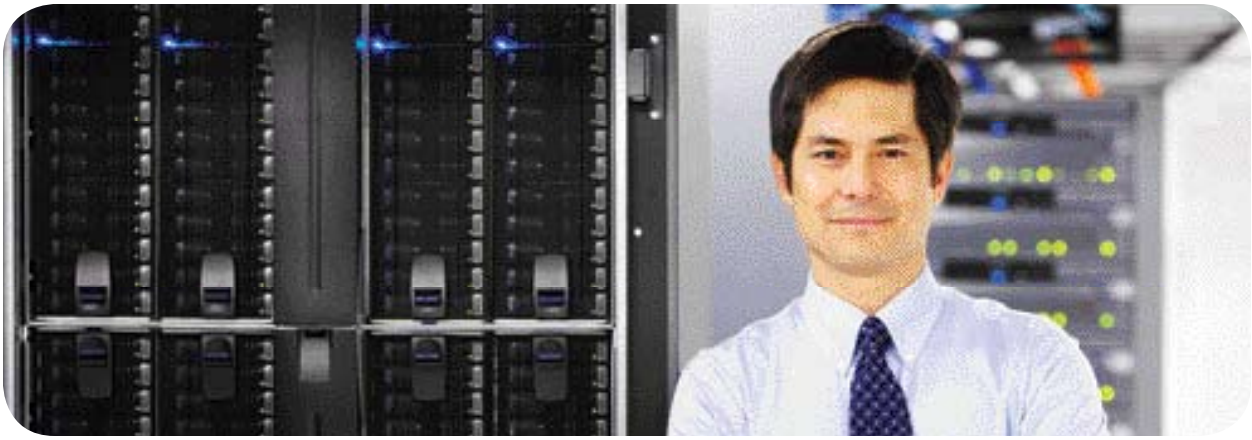

Packard and QLogic are now offering 8Gbps Fibre Channel gear. Organizations appear to be taking to it at a time when IT purse strings have tightened considerably.

Advanced traffic management features in some of the 8Gbps gear include the ability to identify duplicate files and prevent them from being backed up multiple times, as well as the ability to dynamically allocate shared resources as changes occur in the data flows between virtual servers and the storage pool. As IT managers pack more applications onto servers with virtualization, the last thing they want is for the new batch of virtual machines to run out of storage capacity.

An 8Gbps SAN infrastructure can be compared to an upgraded freeway system — it has twice the lanes, or pipes, to carry traffic compared with 4Gbps. Thus, an organization buying an 8Gbps SAN switch can boost bandwidth and fit more in the pipe to alleviate network congestion and speed response times.

The Fibre Channel Industry Association notes that exchanging two 4Gbps Fibre Channel adapters for one 8Gbps adapter lowers total cost of ownership. While these switches provide more functionality, they come at a price premium of up to 65 percent. Therefore, the ROI equation is more important than ever. But 8Gbps is an obvious choice for those who can take advantage of the added performance or the new management features.

Applications (such as entertainment, streaming video, audio, backup, archiving, surveillance, simulation, data movement and migration) that are large sequential I/O or are bandwidth intensive benefit from the increased throughput of 8Gbps Fibre Channel. Fast Fibre Channel also reduces latency of time-sensitive applications such as databases, e-mail, web and file servers.

HIGH AVAILABILITY

With more data being generated and stored for longer periods of time, organizations are facing increased pressure to ensure it is highly available and accessible. There are several options to consider when addressing this issue. The design of a high availability data storage infrastructure can be as varied as the environments and applications that it supports.

There are many techniques, technologies and best practices that can be aligned to diverse needs and budgets to counter threat risks and meet an organization's high availability requirements. These six essential steps will help enable high availability for storage environments:

- Develop strategies to address issues, threats and mission-objective requirements.
- Establish a plan that includes applicable technologies, techniques and ongoing activities.
- Implement a plan that includes technology deployment, configuration and day-to-day support.
- Document and integrate with change control and continuity-of-operations processes.
- Measure and rely on recurring testing to validate the plan and the technologies.
- Use problem determination, isolation, resolution and post-mortems to avoid future issues.

A basic tenet of high availability for storage (which applies as well to networks and servers) is fault isolation and fault containment. That is, eliminate single points of failure (SPOFs) and configure systems so that (if the SPOFs cannot be eliminated) any resulting fault or error condition will be contained to prevent a rolling disaster.

For example, you could configure a pair of networking or storage adapters to have separate paths to a shared storage system; in the event of a failure, you would have access to the storage on the surviving adapter.

Keep in mind that high availability is a balancing act between the availability needed to protect against the most likely scenarios (or scenarios that would have the most dire impact) and your budget. The perception is that components that have more "nines of availability" (99.999 percent availability) will enable high availability.

More nines of availability is good if you can afford it, but more important is how well the components work together. Overall availability is the sum of all of the pieces working together.

MEASURING AVAILABILITY

Availability is often discussed in terms of five nines, six nines or higher. It is important to understand that availability is the sum of all components and their configuration. The amount of downtime per year is calculated as a percentage: (100 – N)/100, in which N is the desired number of nines of availability.

Availability is the sum of all components combined with design for fault isolation and containment. How much availability you need and can afford will be a function of your environment, application and requirements, and objectives. Applications can be looked at from the standpoint of a specific layer or resource, or from end to end, which is what a user of IT services sees.

ANTICIPATING AND PREPARING FOR FAILURE

Availability is only as good as the weakest link. In the case of a data center, that weakest link could be the applications, software, servers, storage, network, facilities, processes, or best practices.

Virtual data centers rely on physical resources to function; a good design can help eliminate unplanned outages to compensate for individual component failure. A good design removes complexity while providing scalability, stability, ease of management and maintenance, as well as fault containment and isolation.

As part of the configuration, costs can be saved by using a single switch. But even with five or six nines of availability, that switch and its firmware or software still present a single point of failure. You should therefore configure a pair of switches, each on its own network, to avoid device failure, software or configuration errors, or network disruptions.

There is a tendency to try to reduce costs by replacing multiple smaller devices with a single, larger higher-availability device — for instance, using a large switch in place of two separate switches. In that scenario, even with a manufacturer that boasts support for more nines than the competition, the physical frame itself might have a common SPOF — for example, a backplane that creates the potential for multiple component failures.

The bottom line is that if something can fail, it will; it's just a matter of time. High availability is about mitigating risk while balancing the PACE (performance, availability, capacity and economics) of operations requirements.

Any technology (hardware, software, network or service) can fail at some point because of the technology itself, its configuration or from acts of nature or man. Most manufacturers will claim that their products have no single points of failure, and thus will not fail. But they also typically describe how to implement fault isolation and other capabilities so that if and when their products fail, they do so gracefully and predictably.

Look for a storage system that is resilient, yet scales with stability and flexibility. This means that as performance increases, availability does not suffer, or as availability increases, performance and capacity do not suffer.

Likewise, combine individual component availability with sound configuration best practices, keeping in mind that even highly available components can break down because of technical or human error. In the end, it's how you configure components that reduces the impact of a failure and maintains high availability. ◊

High Availability Considerations

A sound high availability strategy will include the some or all of the following:

- Fault containment and fault isolation designs
- Availability tools and technologies, including RAID, failover and clustering
- Reliable technologies, including individual components and entire solutions
- Point-in-time snapshots, copies and disk-to-disk backups
- Monitoring and diagnostics for both reactive and proactive analysis
- Server adapters, cabling, switches, routers and input/output data path network items
- Continuity of operations and disaster recovery techniques and processes
- Performance, availability and capacity-planning management
- Distance for survivability on a local or long-distance basis
- Divergent network paths that do not share common infrastructure items
- Configuration management databases and technology tracking
- Cross-domain infrastructure resource management tools and technologies
- Testing, change control and configuration document management

NETWORK OPTIMIZATION
IMPROVING NETWORK RESOURCES

CHAPTER 6:

Network Design

WAN Optimization

Application Networks

Storage Network Protocols

As more organizations look to lower operating costs and consolidate servers by utilizing virtualization, the network plays a bigger role in ensuring communication flows between fewer physical servers. As a result, 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 redundant network, the first step 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, redundant infrastructure.

This assessment should take into consideration the following points.

- **Applications and data on the network:** Be sure to include all of your network traffic, such as Voice over IP (VoIP), Structured Query Language (SQL), Common Internet File System (CIFS), Internet, e-mail and video-on-demand.
- **Virtualization and consolidation of server infrastructure:** You will want to include all efforts your organization has made to bring efficiency to the data center.
- **Network topology:** This should include 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:** Compile and review the data you have on your utilization rates.

Many tools exist to facilitate network assessment. These tools range from basic device information output tools that display the network device utilization to other 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 other application flows. Third-party tools that monitor networks, sniffers and Simple Network Management Protocol (SNMP) tools can also be used.

DESIGNING THE NETWORK

The third step in building a redundant 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 (SPOF).

Based on the 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.

REDUNDANCY AT LAYER 2 AND LAYER 3

In designing redundant networks, it is crucial to eliminate SPOFs. 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 seconds after a primary link failure.

A common best practice, incorporating aggregate redundant links at Layers 2 and 3, increases redundancy. 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.

WAN OPTIMIZATION

As organizations expand and grow, low-speed 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 install 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, backup, 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 dilemma.

WAN optimization helps overcome this challenge. Through WAN optimization, IT departments can provide service nearly as effective as 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.

Optimizing the WAN

When implementing a WAN optimization solution, consider the following:

- WAN connection speeds
- WAN latency
- WAN router lengths of service
- WAN router operating systems
- WAN utilization and TCP flows
- Network infrastructure
- Proprietary/uncommon applications traversing the WAN
- Multimedia or streaming content traversing the WAN

While many options exist for implementing WAN optimization, you should deploy WAN optimization as close to the WAN edge as possible. This ensures that only traffic destined to the WAN gets optimized.

Each WAN manufacturer has different recommendations for installing WAN optimization devices. Most will recommend one of the following: 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, you might utilize an out-of-path mode via WCCP for data centers and headquarters, while using inline mode or installing the WAN optimization appliance in the WAN routers.

WAN OPTIMIZATION BEST PRACTICES

WAN bottlenecks are a big challenge for the IT department. Web traffic is growing exponentially, making once enormous T1 bandwidth feel like dial-up service. Plus, as more staff telework, they put additional strain on already-overtaxed network resources. Here are suggestions to help improve your network's performance.

Identify big traffic sources: Every network is different, so this step is essential. Most networks handle multiple sources of traffic. Wireless access points are notoriously poor resource managers, and large patch or application downloads can create network bottlenecks at inopportune times.

Create traffic policies: User policies should be established so staff members understand the limits. Start with rules for what kinds of e-mail attachments should not be sent internally. Instead of sending e-mail with large attachments to large lists, host attachments on a website or intranet, then provide a link.

Also set rules that dictate which websites should be avoided. Using firewall rules and content-filtering packages can control unwanted data streaming.

Eliminate the unnecessary: By default, most network printers come with support for several protocols. Turn off unnecessary protocols and use the firewall to block ports you do not need. Check PCs for chatty applications.

Use spyware-removal tools to eliminate malware or spyware that may install with other apps. Run Msconfig to see which ones start up with Microsoft Windows, and eliminate everything you can. Turn off auto-updates from Windows, Adobe and other apps.

Prioritize traffic: Use your router as a traffic cop. Give high priority to the most important data packets, such as Voice over IP packets. Advanced switching can improve network speed and quality of service.

Consider outsourcing heavy-use applications: Some apps may be better off hosted by a third party rather than at a central location. For instance, consider a database or SharePoint application that many users will need to access every day. Placing applications onsite or offsite depending on workload can create more bandwidth for other apps and services.

For security and COOP, keep WAN optimization in mind: Security systems often throttle traffic. Some scan packets for content filtering, which can create a slowdown that may cause retransmissions and thereby increase latency.

Backups to offsite facilities should be done during off-peak hours. High-availability replication software can also reduce bandwidth if there are significant document changes on local servers. Always look for replication apps that transfer the delta changes only.

Consider cache servers and thin computing: A document-management cache server cuts traffic by replicating files so there is a copy of all documents on both the main office and remote servers. This also keeps the remote site working if the connection fails.

A thin server that delivers access to your central server cuts bandwidth requirements because it transmits only mouse clicks, keystrokes and an occasional screen refresh.

Separate network resources: Wireless systems are chatty devices, and video conferences can be bandwidth intensive. By giving such processes their own Internet connection, you can increase security and bandwidth simultaneously. Some firewalls let you redirect certain types of traffic to an additional Internet pipe, leaving more room on the main pipe for essential traffic.

Consider WAN acceleration: WAN acceleration uses several methods to reduce traffic. Two accelerators can be used to compress point-to-point traffic so it travels more efficiently. Protocol spoofing looks for similar traffic that can be bundled. Caching web pages reduces Internet traffic because users access a stored cache on the local network for repeat requests.

Add more bandwidth: Bandwidth costs have dropped. Aggressive DSL prices have forced down the cost of copper connections. Fiber and wireless Internet options now let you add fat pipes and redundant bandwidth for a fraction of the cost.

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 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 benefits:

- Disparate applications communicating by appropriately routing application messages in the format expected by that destination;
- Enforcement of consistent security policies for application access and information exchange;
- A high level of information flow visibility, including the monitoring and filtering of messages for both operations and infrastructure purposes;
- Enhancements to application optimization by providing application-level load balancing, offloading security and Extensible Markup Language (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 stand-alone 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 stand-alone device or integrate into a router module.

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 acts as an XML trust enforcement point to provide consistent authentication, authorization and accounting enforcement across all back-end 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 NETWORK PROTOCOLS

Centralized storage is key 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 (FC), 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 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 less expensive 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. ◊

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 type 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.

Application virtualization

Application virtualization as a process is an abstraction of the underlying layer (the operating system) wherein the application is "tricked" into operating as though it is residing and interacting directly with the OS, when in fact it is running in its own isolated virtual environment.

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.

Block-level virtualization

This term refers to the technique of virtualizing data at the block level rather than the file level. Block level virtualization allows administrators to build any kind of file system on top of it, reducing the actual overhead of data significantly.

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.

Continuity of operations plan (COOP)

A COOP is a necessity for all organizations in the event of a disruption of their operations. It involves having the technology and strategy in place that will allow for an immediate or near-instantaneous recovery of networks and systems following a disruption of services.

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.

Dynamic Link Library (DLL) conflict

DLL conflicts often arise on computers that are running multiple versions of the same application. When the computer calls upon a DLL for a particular application of which another version is already running, an error message often comes up because of confusion over the application's distribution.

Fibre Channel over Ethernet (FCoE)

FCoE enables SAN traffic to travel over an Ethernet network. This traffic moves across the link layer, and uses Ethernet to transmit the FC protocol.

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 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 failover.

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.

Multi-win capable

This term refers to applications that are designed to run under a single operating system with multiple sessions.

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.

Optimization

Broadly speaking, this term refers to any software or hardware upgrade that improves the efficiency, speed and/or capabilities of an IT system so that more can be done with fewer resources.

PuTTY

A helpful tool for use in server virtualization, PuTTY is an open source remote session application that can be used to connect directly to the file structure of an ESX host server using the Secure Shell protocol to run console commands or modify configurations files in Unix/Linux command format.

Serial-attached SCSI (SAS)

SAS is an inexpensive, disk-based approach to data storage that emphasizes higher density storage and a high transfer rate, but with decreased performance.

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 virtualization

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

Spanning tree protocol (STP)

STP is a Layer 2 protocol that provides path redundancy (while avoiding loops) for any bridged LAN when the initial link fails. Redundant links are provided via a tree that connects all of the network's switches.

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.

Thin provisioning

Thin provisioning is an approach for optimizing how the available space is used in a storage area network. This approach allocates disk storage space among multiple users, based on the minimum space needed by each user at any given time.

WAN optimization

WAN optimization refers to an approach to improve network services to branch users, working to enhance both the WAN itself and the travel of applications across it. This is done through an optimized version of Transmission Control Protocol and other common protocols.

INDEX

Disclaimer

CDW·G

CDWG.COM/VIRTUALIZATIONGUIDE
888.693.4239

ABOUT THE CONTRIBUTORS

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 about trends and directions in the server, client and storage virtualization spaces. »

SAR HAIDAR is an IT manager at the University of Memphis. He has been involved with application virtualization for more than four years. He also implemented and manages the university's program for remote access to licensed software. His personal blog is at: zephyrski.typepad.com. »

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), Check Point Certified Security Engineer (CCSE) and Cisco Certified Network Professional (CCNP). Eric has written three books for CiscoPress. »

« 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 pre-sales engineer. Todd's current focus is IBM System p and storage hardware, along with IBM's Tivoli Suite of storage products.

« HOWARD WEISS manages the CDW 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 pre-sales team, the Field Solutions team, the HP FieldSE team and now the Network Solutions team.

VIRTUALIZATION AND INFRASTRUCTURE OPTIMIZATION
REFERENCE GUIDE

100224 • Flyer 75745

LOOK INSIDE for more information on:

- Maximizing server virtualization gains
- Developing an application virtualization strategy
- Increasing storage efficiencies
- Designing a redundant network

ISO 9001 SRI CERTIFIED ISO 14001 SRI CERTIFIED

GSA Contract Holder

SmartPay™ GSA