

DISASTER RECOVERY AND
CONTINUITY OF OPERATIONS PLAN
REFERENCE GUIDE



Prepare
for success before
problems arise.

CDWG.com/disasterrecoveryguide | 888.667.4239



The Right Technology. Right Away.®

RECOVERY

DISASTER RECOVERY AND CONTINUITY OF OPERATIONS PLAN

REFERENCE GUIDE

TABLE OF CONTENTS CHAPTER

01	Preparing for the Worst 3
	• Preparing for Disaster
	• Disaster Recovery Lifecycle
02	Client Access 5
	• Understanding Client Access
	• Thin Clients
	• Crafting the Right Strategy
03	Recovering Servers 9
	• Consolidation for Easy Recovery
	• Solutions for Physical Servers
	• Solutions for Virtual Servers
	• A Smart ROI
	• The Right Strategy
04	Networking and Disaster Recovery 13
	• Remote Network Management
	• Network Load Balancing
	• WAN Acceleration
	• Carrier Services
05	Data Storage 26
	• Data Storage Solutions
	• Data Efficiency, Protection and Archiving
06	Other Disaster Recovery Considerations 29
	• UC for Disaster Recovery
	• Power and Disaster Recovery
	• Security and Disaster Recovery
	GLOSSARY 33
	INDEX 35

WHAT IS A CDW•G REFERENCE GUIDE?

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

» We are pleased to also offer the **CDW•G IT Investment Guide**. It includes products and information to help you meet your disaster recovery and continuity of operations planning objectives.

PREPARING FOR THE WORST



CHAPTER 1:

Preparing for Disaster

Disaster Recovery Lifecycle

For any government organization or educational institution to fulfill its mission, it is vitally important for data to be continually secure and available. In the United States, numerous events over the past decade have led to a heightened awareness of the need to be prepared for the worst.

Chief information officers, infrastructure services managers and IT administrators are now faced with the pressing concerns of growing demand for and reliance on access to information, along with ensuring continuous operation during a disaster.

As organizations are increasingly dependent upon rapid access to data (and subsequently less tolerant of failure), an increased focus must be given to disaster recovery and having a comprehensive continuity of operations plan (CoOP) in place. Organizations that have a thought-out, detailed strategy for disaster recover and continuity of operations are in a position to keep their systems running when problems occur.

PREPARING FOR DISASTER

The demands placed upon government and education information systems are significant. Both internal and external parties expect operations, services and technologies to be available 24 hours a day, seven days a week, 365 days a year, with no exceptions and no excuses.

Nearly every aspect of today's government is expected to be available continuously without interruption, regardless of the circumstances. When disaster strikes — whether a natural disaster or technological failure — operational services and technologies are expected to still be available.

Most organizations need to place a high value on being prepared for disasters of any kind because the practical ramifications of failing to do so can be very high indeed:

- **Public confidence:** When a government organization or educational institution experiences an interruption in services or suffers a loss of data, the public can lose confidence in that organization's viability in a crisis and its ability to protect their personal information.
- **Public safety:** Organizations that manage public safety operations (or manage data that could be potentially vital to intelligence) and law enforcement organizations have a heightened responsibility to ensure continuous availability of any system that might directly or indirectly impact public safety.
- **Staff confidence and effectiveness:** As technology becomes an even greater part of government and educational operations, users have come to rely more and more on services and technologies to do their jobs. When those services or technologies become unavailable, even for short periods of time, users suffer major productivity losses.

In addition to the direct costs of lost productivity, long-term damage can result in low staff morale and confidence in the organization, extending the monetary damages well into the future, even after services have been restored.

- **Cost:** Even the loss of a single mission-critical service, such as e-mail or web connectivity, can cost some organizations millions of dollars in direct costs. Avoiding this downtime with a disaster recovery plan or CoOP in place is a clear benefit.

A plan for disaster recovery or continuity of operations that helps operations get up and running quickly or keep them up and running with a minimal loss of data, information and productivity is a necessity. Tools and technologies are available to help mitigate various forms of disaster and to help keep downtime to a minimum.

Even if organizations do have a disaster recovery plan or CoOP in place, they can still face problems. Often that plan doesn't reflect real-world costs.

Maybe this sounds familiar. An organization's management knows a disaster recovery plan is needed, even required, but because they believe the probability of a disaster actually occurring is very low, the budget for the disaster plan is woefully under what the costs are to implement a plan that would get them up and running.

There is a cost for effective disaster recovery and continuity of operations. And certainly this cost is worth it compared to the alternative: your organization being unable to function.

Laws and regulations on how organizations must put together disaster recovery plans and continuity of operations plans differ across federal, state and local levels. Regardless of the organization, it is critical to get buy-in from key departments so that disaster recovery and continuity of operations are not relegated secondary issues.

DISASTER RECOVERY LIFECYCLE


Maintaining a disaster recovery plan or CoOP is an ongoing process and should be considered more of a lifecycle than a once-a-year checklist that can be checked off to completion. Although the solutions are often technology-centric, they must always be derived from a deep operations acumen, or a detailed and intimate knowledge of the operational needs. There are five key phases to the disaster recovery lifecycle:

- **Analysis:** This is arguably the most critical component of developing a CoOP or disaster recovery plan. It is during the analysis phase that several examinations will be conducted to determine potential impacts, identify likely threats and develop impact scenarios.
- **Solution design:** During this phase, the goal is to identify the most cost-effective and technically viable solution.
- **Implementation:** This phase consists solely of the execution of the design elements identified in the solution design phase.
- **Testing and acceptance:** To be certain that disaster recovery plans and CoOP meet the needs of the organization, testing is required to assure process and acceptance.
- **Maintenance:** Once a disaster recovery plan or CoOP has been established, regular maintenance of the plan helps to

ensure viability. The maintenance phase is the ongoing effort to address technical solution needs, recovery solution needs and organizational changes as they impact operational preparedness and a host of other factors.

From the very beginning and throughout the lifecycle, the focus must always remain on managing risk to the operations environment and maintaining continuous availability. Losses related to data loss, service failure, power outages, software incompatibility and security concerns continually threaten the operations environment.

This reference guide discusses the tools and technologies that will enable your operations to continue when facing adversity. It will especially focus on servers, networking, data storage, power protection and client access. When the worst-case scenario becomes front-page news, your organization will not only have an accurate and complete disaster recovery plan or CoOP, it will also be prepared with the technology to make it work. ♦



CLIENT ACCESS



CHAPTER 2:

Understanding Client Access

Thin Clients

Crafting the Right Strategy

Many organizations put a tremendous amount of time and effort into developing a disaster recovery plan, but most overlook the development of a client access plan. A sound disaster recovery plan replicates all server data from the organization's data center directly to the disaster recovery site, so in an instant all systems and applications can be brought online.

However, with little or no strategy as to how end users actually connect to the systems in the disaster recovery site, a good plan can quickly be rendered ineffective.

When preparing your organization for disaster or continuity of operations, it is essential to plan for desktop, notebook and thin client replacement. Depending on the situation, many workstaff may lose their smartphones as well. Given our dependence on these mobile devices to conduct work on a day-to-day basis, a strategy needs to be put in place to replace them as quickly as possible.

UNDERSTANDING CLIENT ACCESS

Client access is the method used to present applications to end users. There are many solutions developed around application delivery. The most common are:

- PDAs
- Encryption
- Smart cards
- Key fobs
- Keeping desktop images in a safe location; deploying them on demand
- Thick clients

The most common implementation of all the deployment models is assigning physical desktops and notebook computers to end users. These systems are loaded with local and remote applications that can save data both locally and remotely.

For example, an end user may have a desktop computer with spreadsheet software loaded locally. The end user can have spreadsheets that are stored locally, on a file server somewhere on the network, or both. Although today's network technology can force end users to save all their documents on a centralized file server, seldom is this actually implemented properly.

Individual desktops and notebooks ultimately allow for personalization of the operating system and applications that they access, and although it may be a benefit to an individual end user to have a personal environment, it makes it extremely difficult to duplicate these environments in a disaster.

THIN CLIENTS

The advent of Citrix Presentation Server and Microsoft Terminal Services paved the way for thin clients to be produced in all shapes and sizes.

The concept was simple: a small, inexpensive device with no moving parts that runs a light operating system and connects to a remote operating system and applications, thereby only sending mouse clicks and keyboard information back and forth. End users would always be presented with a common desktop and common applications with little-to-no customization.

Today, thin clients are being used to access virtual desktops as well as blade PCs and workstations. Although the method of application access and delivery is different, the thin client is still an

ideal, low-cost access device that can be deployed fairly easily and replaced with no loss of end-user productivity.

TERMINAL SERVICES

Both Citrix Presentation Server and Microsoft Terminal Services produce similar products that can host an end user's desktop as well as their applications. Whether your organization decides to deliver a published desktop or a published application is completely dependent on the user's needs.

In general, the recommendation has always been to publish applications only. Fortunately, there are a number of methods for actually accessing the published application.

For example, an end user can access a published application using a browser on a home computer, which could be a Mac, a Linux or a Windows desktop. A client can also be loaded directly on the operating system so that application icons can be readily available on the actual desktop and only a single click away. The bottom line is that the end user always gets a consistent experience, no matter what device is used.

VIRTUAL APPLICATIONS

Application deployment isn't always an exact science. Although many tools are available to automate application delivery, each tool has its own set of advantages and disadvantages.

Application virtualization solutions are available from Microsoft, Citrix and Symantec that isolate applications on each operating system. This allows applications to run completely isolated from each other, with minimal impact on the operating system itself. This means you can effectively stream applications to each

desktop operating system and have applications available offline.

VIRTUAL DESKTOPS

Server virtualization is now approaching maturity and widespread adoption. As a result, using the same technology to host desktops is quickly becoming the norm. Virtual desktops involve an end user accessing a desktop operating system (with or without local applications) with a thin or thick client device using remote desktop protocol (RDP).

The desktop operating system runs inside a virtual machine sitting on a virtualized architecture, making use of powerful server processing and storage area network (SAN)-based storage.

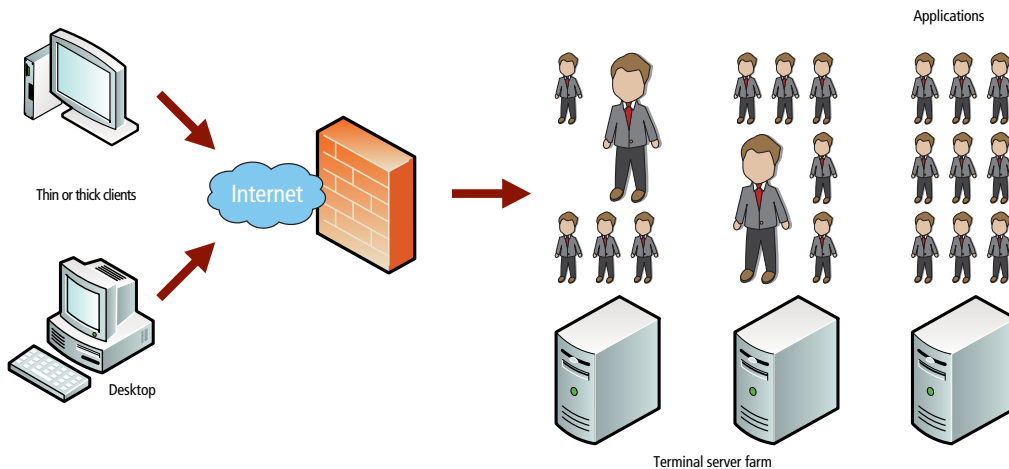
From a disaster recovery perspective, this makes a lot of sense. If the entire data center's servers and desktops are on the storage array, why not replicate it "as is" to the recovery site, and bring it all up together? End users can then access their virtual desktops from any system until the data center is recovered.

CRAFTING THE RIGHT STRATEGY

The development of a sound client access recovery plan should revolve around the applications needed to run the organization. Once these applications are identified and the server, storage and network infrastructure are replicated, the last piece that needs to be put into place is how an end user will access those key applications.

Although a number of solutions for client access are available, it's generally recommended that organizations use published applications or virtual desktops to access applications hosted at the recovery site. This will provide the least amount of

2a. Citrix Application Delivery



configuration required to get end users access to the tools they need to continue doing work.

PUBLISHED APPLICATIONS

Citrix Presentation Server is a leading product for centralized application delivery. Designing a client access solution around Citrix enables a secure, reliable application access method over any device and any network. This enables end users to access their applications immediately from any web browser. Many organizations implement client access using this technology for not only disaster recovery but also day-to-day use.

Figure 2a (on the previous page) depicts how end users can access applications following a disaster.

Users can open up a web browser on a Windows, Linux or Mac system and navigate to the organization's website, log in with their existing security credentials, and then Citrix Web Interface will display the available applications to that particular end user. A single click will open a seamless window to that application, and the end user can begin working immediately.

Citrix Presentation Server has the following advantages:

- It is a technically mature solution.
- It offers ease of access from any device over any network (independent computing architecture [ICA] protocol used by Citrix is very efficient over slow networks).
- It is a very secure solution.
- It publishes applications directly to end users.

Citrix Presentation Server has the following disadvantages:

- All users share the same server operating system and applications, so bottlenecks can occur.
- If a server becomes unavailable or unstable, the session will need to be restarted on another server in the farm. Although this can be automated, it can cause some disruption to the end user.
- User sessions cannot be dynamically moved or load balanced, they must disconnect and reconnect to generate a session to another Citrix Presentation Server in the farm.
- Each user or device needs to have a Windows Server client access license (CAL), a Terminal Services CAL and a Citrix Presentation Server CAL, in addition to each Windows Server license that needs to be installed to host the sessions. Also, applications loaded on each server need to adhere to licensing policies for each vendor.

USING VIRTUAL DESKTOPS

Although still a relatively new concept, the virtualization of desktops makes sense for certain environments. Unlike Citrix, which shares a slice of a Windows Server operating system and applications to each end user, a virtual desktop is its own unique sandbox usually loaded with a Windows Vista/XP operating system and applications that can be locked down or left open to end-user customization.

Ideally, every virtual desktop should be completely locked down so that end users cannot save any data on the virtual desktop, but rather to a folder on a shared file server. This ensures a consistent environment for every end user, even if they get a different virtual desktop every time they access the system.

2b. Virtual Desktop Access






Figure 2b (on the previous page) demonstrates how end users access virtual desktops following a disaster. Although the paths to access the system are very similar to methodology used for published applications, the end user receives a full desktop rather than just links to individual applications.

Users can access these virtual desktops on a Windows, Linux or Mac system through a similar redirection of the organization's website to the connection broker's website, using the same network credentials they had before the disaster. A single click opens a seamless window to the virtual desktop, and the end user can begin working immediately.

Virtual desktops have these advantages:

- They offer ease of access from any device over any network (ICA protocol used by Citrix is very efficient over slow networks).
- They are a very secure solution.
- Every desktop session is unique and not shared with other users.
- Virtual desktops can be dynamically moved between different servers in the farm to increase performance based on predetermined thresholds.
- If a VMware Infrastructure (ESX) server becomes unstable, virtual desktops are rebooted on other servers in the farm.
- Licensing is fairly straightforward with the Microsoft Vista

Enterprise Centralized Desktop (VECD) model. Note that each desktop also needs Windows Server CALs to access resources on Windows Servers, and VMware Infrastructure licenses for the server virtualization infrastructure.

Virtual desktops have these disadvantages:

- They are not a technically mature solution.
- Even though virtual desktops are automatically rebooted after a server failure on other servers, this still causes a temporary outage and the loss of the end user's session.
- Publishing a desktop to every user may appear more complicated than publishing applications to the end user.
- Only Windows, Linux and Mac web-client links to the connection broker are supported.
- Microsoft RDP is not as efficient as Citrix ICA.
- All users share the same server operating system and applications, so bottlenecks can occur.
- If a server becomes unavailable or unstable, the session will need to be restarted on another server in the farm. Although this process can be automated, it can cause some disruption to the end user. ♦

RECOVERING SERVERS



CHAPTER 3:

Consolidation for Easy Recovery

Solutions for Physical Servers

Solutions for Virtual Servers

A Smart ROI

The Right Strategy

Developing a comprehensive and cost-effective recovery strategy for servers is a challenging undertaking for organizations. In fact, many organizations are surprised by the high cost of implementing a recovery site and choose to do nothing, falsely believing that a disaster will never affect them directly.

Fortunately, stronger regulation and the impact of recent disasters has forced many organizations to develop recovery plans and procedures in order to minimize data loss and ensure operations continuity.

CONSOLIDATION FOR EASY RECOVERY

Server consolidation, with the goal of reducing data center cost and complexity, should be explored before developing a recovery plan. Not only is the goal of server consolidation to reduce the number of servers, but it also aims to make the environment simpler to administer and maintain. This in turn makes it easier to recover.

The four most common forms of server consolidation carried out today are physical server consolidation, virtual server consolidation, data center consolidation and application server consolidation.

PHYSICAL SERVER CONSOLIDATION

Physical consolidation was developed shortly after the surge in server deployments following the early success of the Internet.

Initially, most organizations began consolidating multiple file, print and database servers onto fewer, larger, clustered servers to reduce the physical footprint and lower administration costs.

For example, an organization might consolidate 15 departmental file servers onto a two-node file server cluster for redundancy and increased performance.

VIRTUAL SERVER CONSOLIDATION

The ability to migrate multiple physical servers onto fewer servers through virtualization has been the consolidation method of choice for the past few years. This technology has allowed data centers to reduce their power and cooling consumption, in addition to freeing up rack space for future growth.

For example, in certain situations virtualization can allow an organization to consolidate 100 servers onto 10 or fewer physical servers in a data center, yielding a 10:1 consolidation ratio.

DATA CENTER CONSOLIDATION

Organizations with multiple data centers and remote sites have begun consolidating servers and storage devices to a centralized data center. Although servers can be completely eliminated in certain scenarios, some organizations may need to keep servers at remote sites if loss of network connectivity is an issue.

As an example, an organization with 100 file servers in branch offices might consolidate to a central location, while using wide

area network (WAN) optimization devices to eliminate network bottlenecks and increase throughput to ensure a consistent end-user experience.

APPLICATION SERVER CONSOLIDATION

Today's 64-bit hardware and software can host more users with more processor cores and memory than ever before. This allows for the consolidation of multiple application servers onto fewer physical or virtual systems.

For example, an organization with 10 Microsoft Exchange 2000 servers could migrate to a clustered pair of Microsoft Exchange 2007 servers and utilize more memory, more cores and 64-bit technology, resulting in better performance and a higher user yield per server.

Regardless of the method of consolidation, the goal is still the same: Simplify the server environment in order to make it easier to recover following a disaster.

Note: Although content in this chapter may apply to UNIX, midrange and mainframe systems as well, the primary focus is x86 servers (Intel Xeon and AMD Opteron).

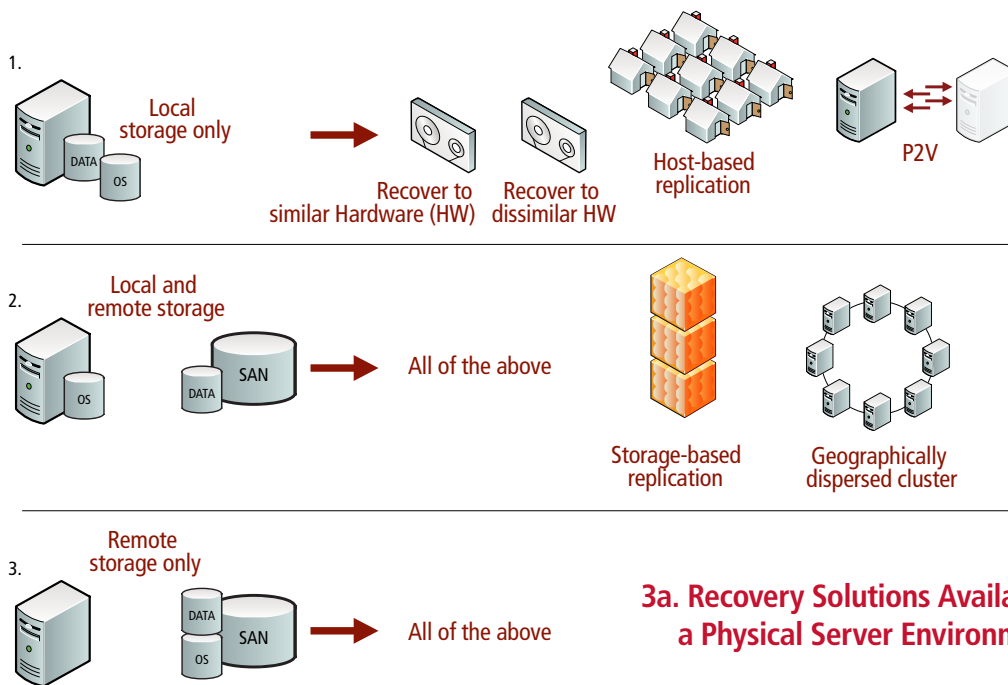
SOLUTIONS FOR PHYSICAL SERVERS

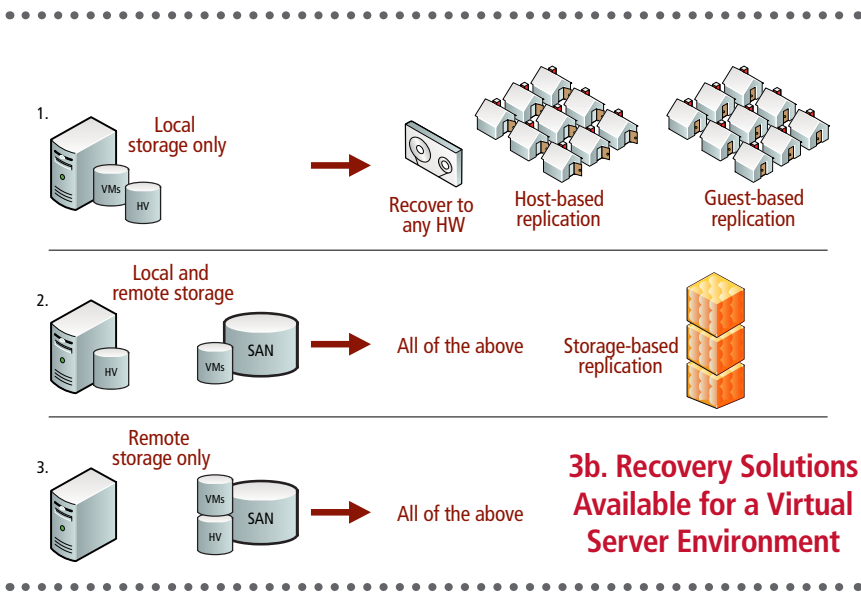
Although the future of the data center is clearly virtual, the number of physical servers deployed still outnumbers that of virtualized ones. Also, experience has shown that on average only 80 percent of a data center can be virtualized, which means that solutions still need to be developed for the replication and fail-over of the physical servers.

Figure 3a shows recovery solutions available for a physical server environment, and is divided into three options that are explained as follows:

1. For servers deployed with local storage or direct-attached disk, recovery options include:

- **Restoration from tape media to similar hardware:** Although this is the best-case scenario for physical servers, servers must be purchased at the same time to guarantee that the system being recovered is identical.
- **Restoration from tape media to dissimilar hardware:** This is not a full-proof technology, but many options from various manufacturers are now available to perform a tape restore to dissimilar hardware. In comparison to restoration to similar hardware, this option seems more feasible, given that server hardware is refreshed every three to six months, so the odds of finding identical hardware (if not purchased at the same time) are slim to none.





SOLUTIONS FOR VIRTUAL SERVERS

Organizations that have begun the migration to a virtual infrastructure have immediately reaped the benefits of this technology. Because of its isolation and encapsulation capabilities, virtualized servers can be moved and restored between different physical servers and storage hardware, with no need for any kind of migration.

Figure 3b shows recovery solutions available for a virtual server environment and is divided into three options, which are explained as follows:

- **Host-based replication:** Many products are now available to replicate data directly from the operating system or the application layer. This allows for the continuous or periodic replication of the server to a similar or completely different server in the recovery site.
 - **Physical to virtual (P2V):** Physical to virtual technology has advanced over the last five years, and products are available to continuously convert and replicate physical servers to other physical or virtual servers at the recovery site.
2. In addition to the previously mentioned recovery options, servers deployed with a combination of local and remote storage can also make use of the following options:
- **Storage-based replication:** The data residing on externally attached network file system (NFS), Fibre Channel or iSCSI storage can be replicated to another similar storage device in the recovery site. However, this only protects the data on the SAN; a plan is necessary to revive the operating system.
 - **Geographically dispersed clusters:** Clusters go hand in hand with storage-based replication and enable certain applications to be continuously available after a disaster. A cluster of this magnitude could potentially bring up the application within minutes of a major disaster, and therefore, no restoration of the operating system is necessary.
3. For servers deployed with no local storage, with both the operating system and data drives stored on remote storage, all of the previously mentioned options are available for recovery.

3b. Recovery Solutions Available for a Virtual Server Environment

1. For servers deployed with local storage or direct-attached disk, recovery options include:
 - **Restoration from tape media to any hardware:** Because virtual hardware is identical, virtual servers can be restored to a virtualization platform on any server or storage hardware.
 - **Host-based replication:** Enterprise virtualization technologies are hypervisor-based, and options are now available to replicate virtual servers directly from the hypervisor. This allows for the continuous or periodic replication from the hypervisor, with no impact on the virtual server.
 - **Guest-based replication:** Rather than run replication products at the hypervisor level, replication can also occur from within each guest or virtual server operating system. In this instance, the virtual server can use host-based technologies defined for physical servers and thus gives organizations more options to choose from.
2. In addition to the previously mentioned recovery options, servers deployed with a combination of local and remote storage can also make use of storage-based replication.

Similar to solutions defined for physical environments, data residing on externally attached NFS, Fibre Channel or iSCSI storage can be replicated to another similar storage device in the recovery site. The difference for virtualized environments, however, is that both the operating system and data volumes for every virtual server are replicated and can be brought online in minutes.
3. For servers deployed with no local storage, with both the operating system and data drives stored on remote storage, all of the previously mentioned options are available for recovery.

A SMART ROI

Duplicating a server infrastructure at an alternative facility can be expensive. However, if the operations requirements state that systems need to be recovered as soon as possible, then incurring the costs of a second data center cannot be avoided, as well as the costs that go along with the management and maintenance of that facility.

However, the cost of every solution needs to be justified, and a recovery solution is no exception. So here are some tips to help support the financial justification:

- **Risk assessment (RA):** Finding out the true cost of downtime for an organization can be a very difficult process, which is why most organizations outsource their RAs to outside consultants. This process can take quite a long time, but its results are invaluable. Once your organization determines the cost of downtime, it will be easier to make a case for the cost of the execution of a disaster recovery plan.
- **Split-brain data center model:** One way to ensure that the recovery site is properly utilized is to split the server farm between both the production and the recovery data center. This has some advantages. First of all, if a disaster occurs only half of the server farm is affected. Second, the recovery site is 50 percent utilized at all times. Third, the load on the production data center in terms of power, cooling and backup windows shrinks.
- **Virtualize everything:** Although not completely possible yet, virtualizing a production data center enables an organization to cut power and cooling costs, reduce rack space, and introduce new methods of management, maintenance, backup/restore and, of course, recovery.

The realized savings from virtualization can then be used to build out (or outsource) a recovery center. Some organizations use virtualization in recovery sites first, and then later come back to virtualize the production data center. Although this can be done very effectively, the most cost-effective method would be to virtualize as much as possible.

- **Testing and development:** Once an organization has virtualized its testing and development servers and is continuously replicating them to the recovery site, these servers can be cloned and used for all testing and development needs.

This will reduce the virtual server instances running in the production data center, as well as the need to clone virtual servers periodically to test updates, such as patches and upgrades. Because the virtual servers are being continually replicated to the recovery site, an up-to-date copy is always available and can be used immediately for testing and development.

THE RIGHT STRATEGY

Server virtualization is now the leading technology used for disaster recovery. Organizations have begun using this technology not only because of its immediate cost savings, but also because of its flexibility.

Smaller organizations that don't usually have a shared storage subsystem (such as iSCSI, Fibre Channel or a SAN) can use virtualization in both production and recovery sites using host-based replication software. This enables any organization with server virtualization technology to implement the right solution for its operations and still stay within budget.

The future of consolidation is clearly virtual, and although virtualization across the entire data center hasn't fully matured, we can easily predict where this technology is going. Data center design can become very simple once all areas (storage, server, desktop, application and network) are virtualized.

Starting in the production site, multiple shared storage subsystems can be used with storage virtualization in front of them. This allows servers to be moved on demand between different storage devices for performance reasons, or at the end of a lease cycle. All servers and desktops can be completely virtualized, which allows all instances to be replicated into the recovery facility and brought online in minutes.

The end-user experience is almost identical when accessing applications in either site. In many instances, users could be redirected to the recovery site automatically, but it depends on the design and architecture. As an alternative to virtual desktops, Microsoft Terminal Services (which is used predominantly today) can be used to serve up the applications or desktops to end users in either site.

The biggest value point for organizations once the data center is virtualized is that all of it can be replicated to the recovery site and brought online immediately.

Virtualization removes hardware dependencies. This enables completely different servers and storage subsystems to be used in the recovery site. Removing hardware dependencies enables organizations to reuse existing hardware for their recovery sites, as well as allow for a smooth transition to a different hardware manufacturer during a refresh cycle. ♦

DATA CENTER DESIGN

can become very simple once all areas (storage, server, desktop, application and network) are virtualized.

NETWORKING AND DISASTER RECOVERY



CHAPTER 4:

Remote Network Management

Network Load Balancing

WAN Acceleration

Carrier Services

A well-maintained highway infrastructure functions as a critical element for moving goods and services, which is essential to the functioning of a nation's economy. Similarly, a well-maintained network is a key resource that ensures an organization's systems availability and, in essence, directly supports operations objectives in today's connected enterprise.

A typical modern enterprise architecture comprises many devices that can span national and often international boundaries across multiple time zones. Complexity of these networks aside, the scale alone justifies a well-defined and well-managed network management infrastructure.

This infrastructure should not only provide visibility into all areas of the IT footprint down to the device configuration level, but it should also be remotely accessible in a secure, robust and resilient manner, both in- and out-of-band.

REMOTE NETWORK MANAGEMENT

Imagine a well designed environment where the necessary computer infrastructure is in place and functioning as desired. This network can be managed remotely, whether from across the room or from another room, building, city, state or even country. Also, organizations have the option of managing it themselves or outsourcing.

IN-HOUSE MANAGEMENT

Some organizations choose to do all of their network management in-house. To be effective, this approach requires a sizable

investment in people, processes and technology. It requires properly trained staff available around the clock to respond to alerts; a strong set of standard operating procedures defined to handle various events; and a fault-tolerant infrastructure to monitor and generate alerts based on pre-established parameters.

Depending on their requirements, organizations can find management solutions that offer varying levels of visibility, control and application. The closer you are to the management device, the more perceivable control you have over the network. However, this arrangement does not scale particularly well and does not offer a flexible way of handling network management.

At the console of the device, you can literally see what is transpiring by looking at the screen and/or panel and lights on the device itself. You can tell whether the environmental conditions of the space are conducive to the operating requirements of the network's assets. You can also personally observe any physical conditions that may be interfering with the proper operation of the equipment.

With this kind of insight available, where an organization has visibility into these conditions remotely in an intelligent and actionable way, it becomes impractical to devote human resources to these more mundane tasks. In fact it's this very capability that opens up the possibility of outsourcing network management to a managed service provider.

OUTSOURCED MANAGEMENT

Two key considerations for outsourcing your network management

are price and contractual obligations. If terms can be defined and agreed to, outsourcing can be far more effective and efficient (as well as less costly) compared to an in-house management strategy.

Outsourcing also offers a key benefit in relation to disaster recovery. It allows organizations to concentrate on recovering their operations and the communication around it rather than spending precious time worrying about systems recovery.

Many organizations have turned to hosted managed services (HMS), which are usually based on service level agreements, rather than continuing to devote human resources to these routine, yet critical, tasks.

Service level agreements (SLAs), in particular, allow the organization to determine the degree of outsourced management that best suits it. This model scales well and ultimately proves more cost-effective than self-maintaining the environment 24x7.

Just as you would qualify a doctor by checking references, it's worth the effort to investigate how an HMS provider runs its network operations center (NOC). Only in doing so can you responsibly select a provider to monitor and manage your network.

For example, find out what certifications the provider has and what its overall reputation is in the marketplace. How well is its NOC staffed? What kind of SLAs does the provider offer, and will any of them aptly fit your needs?

You may want to inquire as to what green energy measures the facility has implemented. And it's critical to know what the service will truly cost. Answers to such questions will help you select the proper vendor for your remote network management needs.

NETWORK LOAD BALANCING

Network load balancing is the process of spreading the work or balancing the workload between two or more servers, network links or other devices. In addition to efficiently allocating the load on the network, load balancing can increase resiliency within the network by potentially eliminating single points of failure and by providing seamless continuity should a network link, server or other device fail.

Network load balancing can be accomplished within the network using software (such as certain operating systems) or through purpose-built hardware. It is commonly deployed in a variety of scenarios. These scenarios may include link load balancing, local load balancing, global load balancing and storage load balancing.



In addition to the efficiencies gained through load balancing as previously mentioned, each of these types of load balancing, when properly deployed, can add critical elements to an organization's disaster recovery plan or CoOP.

Local load balancing: This type is utilized to load balance traffic, often HTTP traffic, within a single data center to ensure users are directed to a functional server. Local load balancing typically ensures high availability for your applications by monitoring the health and performance of individual servers in real time.

Link load balancing: This technology is used to manage multiple Internet links to a single data center, eliminating the need to depend on a single ISP. When properly configured, link load balancing provides the ability to manage multiple redundant links as well as the ability to choose the best performing or least-utilized link.

Global load balancing: This approach allows users to be directed to the closest or best-performing available data center. This decision can be based on factors such as proximity, performance and availability. Global load balancing can not only protect against

the loss of a data center but can also ensure users receive the best available application performance.

Storage load balancing: This technology provides the ability to consolidate storage resources into a single virtual file server. This process can eliminate bottlenecks and improve the overall efficiency of the data center.

Whether the situation requires balancing server loads, multiple links or traffic between data centers, network load balancing used as a portion of the overall data center strategy can effectively improve the performance and resiliency of network resources as well as the experience of the end user.

WAN ACCELERATION

The valuable data that IT systems carry is essential, not just for fulfilling an organization's mission but also for it to function at all. This data may span terabytes of storage and require dedicated disaster recovery space. An organization should ideally locate this space away from its central site in order to hedge against disasters, natural or otherwise.

However, a second potential challenge exists as organizations strive to properly protect their data. The normal growth of an organization frequently leads to multiple branches. These branches often end up having limited technical resources, leading to difficulties in managing and protecting the data that they harbor.

Consolidating this data to a data center as well as a disaster recovery site resolves this problem. However, such a consolidation can create subsequent problems when accessing vital information quickly. Wide area network (WAN) optimization resolves this quandary by accelerating protocols, compressing traffic and caching files. It thus speeds up applications such as Exchange, Citrix and web applications.

The WAN acceleration market has become fiercely competitive, with many manufacturers even offering a "try-and-buy" approach. This allows customers to install the hardware in their production environments and measure the results before deciding whether to purchase. Organizations can also accomplish WAN optimization via managed services if they prefer not to purchase and manage the necessary hardware.

Uncertain economic conditions pressure organizations to make do with less and to operate more efficiently. WAN optimization offers a strong return on investment when it is properly implemented.

Acceleration statistics vary by manufacturer, but backup and recovery of Windows data processes can be accelerated up to three times their normal speed, while bandwidth utilization can be decreased by more than 60 percent, according to some manufacturer claims.

The amount of time needed to configure a WAN accelerator varies by the manufacturer and the organization. They should be configured specifically per site to fully capitalize on their benefits. Building an application-specific baseline before the installation is an important consideration in order to be able to measure improvements.

CARRIER SERVICES

Disaster recovery and CoOP can be two of the most challenging and frustrating projects for IT departments. These projects test not only the team's skill, but most of all, its perseverance. Accountability and sharing the load are two very important attributes to a successful implementation of disaster recovery and CoOP.

A disaster by definition is an adverse, unfortunate and unforeseen event. There are three common types of disasters:

- Natural (fire, flood, wind, earthquake)
- Malicious intent (viruses, burglary, vandalism)
- Accidental (outages of power, telecom, hardware or software systems)

The following areas should be considered the core of a successful and ongoing disaster recovery solution:

- **Redundancy:** Whether this is at the carrier level, server level, application level or hardware level, redundancy should be a primary focus when designing and implementing your plan.
- **Consistent testing:** Organizations that do regular disaster recovery testing are quickly back up and running when real disasters happen. Depending on the organization's size, number of applications and number of changes that occur, testing should occur one to three times per year.
- **Change management:** A very strong change management process should be in place to ensure that disaster recovery procedures are kept consistent. If network equipment, applications and servers are thrown in at random, when it comes time to test, the results will not be positive. Developing a plan in the midst of a crisis creates more chaos and adds to existing pressure.
- **Risk assessment:** Organizations should assess their operations models and the infrastructure that runs those models by running an operations impact analysis (OIA). An OIA will assign a dollar amount to determine how much money could be lost per day or per week. Thus, this financial number should drive the importance of the recovery for specific applications/hardware.
- **Staff training:** It is important to engage department heads, making them a part of the disaster recovery planning and

empowering them to make decisions throughout the process. This ensures that staff are aware of the procedures and have a vested interest to make sure their department is “disaster free.”

- **Protect organizational assets:** Along with the risk assessment, organizations should identify those networks, applications and devices that they deem critical. Once identified, organizations should then do a vulnerability study on these critical systems. Approaches to alleviating those vulnerabilities can include supplying redundant hardware, implementing virtualization and/or creating a completely separate environment.

CARRIER DIVERSITY

In many cases, carrier dependency is an organization’s single greatest vulnerability. When it relies on a single Internet connection to supply Internet, e-mail and customer access to web-based applications, the organization is vulnerable. It isn’t that organizations don’t know of the danger; they’re often concerned with the costs associated with dual connections.

An OIA should be the main driver in this situation. Is the cost of load balancing the applications (adding another telecom circuit) more than the cost associated with the loss of operations for a certain amount of time? For most organizations, the cost of public or user dissatisfaction far outweighs the cost of the extra circuit.

SINGLE LOCATION

If an organization has only one location, it may find through an OIA that it is not cost productive to have another fail-over or disaster recovery location. In this case, it’s a good idea to have a

secondary circuit for each primary located on the opposite side of the building and using a carrier other than the primary carrier on it.

It is crucial that the primary and secondary circuits not use the same local “central office.” Depending on the OIA, the circuits may not need to be the same size and each circuit may not need to be duplicated, but you should cover the identified critical systems.

MULTIPLE LOCATIONS


If an organization has multiple sites, an OIA may reveal that it is cost effective to build out one of these existing sites (if possible) as a disaster recovery site. This may include replications, virtualization, DNS fail-over and adding circuits that can handle what the primary site deals with on a daily basis.

Generally, these sites should be located more than 20 to 30 miles from the primary site. In this example, organizations have to take into account physical distance, cost and improbability (how this provider fits into the existing network).

A COLLOCATION SITE

Some organizations that do an OIA may find that the cost of having a completely separate data center, which they can fail-over to in an instant, is worth the cost of duplicating most of the primary environment. This kind of setup is somewhat labor intensive and expensive at first, but if built correctly, the labor portion is easily managed or can be contracted out.

In this example, many organizations have found that using a collocation vendor that leases rack space, cages or rooms is the optimal way to go because power, cooling and circuit needs are fully redundant. ♦





Finding technology to meet your needs can be difficult. It's a good thing we enjoy a challenge.

When you have huge demands on technology, it's like music to our ears. CDW·G is there to take on your technology challenges. And to help us succeed, we have one of the largest selections of top-name products around, along with a personal account manager to guide you through all your options. If you have in-depth questions, we have technology specialists ready with answers. And when you need something a little more specific, we have a custom configuration center to get you new customized technology. So give CDW·G a call today, because we're up for the challenge.



The Right Technology. Right Away.®
CDWG.com • 800.808.4239

DATA STORAGE



CHAPTER 5:

Data Storage Solutions

Data Efficiency, Protection and Archiving

An organization's ability to access mission-critical data after a disaster is key to getting it back up and running. Because data storage spans a variety of technologies, it is important to start evaluating the different methods in order to put a data storage solution in place that fits your organization's needs.

DATA STORAGE SOLUTIONS

Disaster recovery plans and CoOP tend to focus heavily on technological elements, while neglecting the enormous investment organizations often have in paper-based resources. A document management system that transfers those paper-based resources into a space-saving, more-manageable electronic format is a critical component of any disaster recovery plan or CoOP.

After all, electronic documents are more easily backed up, stored offsite and recovered in the event the originals are lost or destroyed.

ELECTRONIC DOCUMENT CAPTURE AND MANAGEMENT

Document capture is the first step in the document management process. The capture involves making digital copies of the original paper documents. This is typically done with a scanner or multifunction device. Once the content is in an electronic format, it is indexed and stored in a central storage repository with the organization's other electronic documents.

Most document management systems utilize some form of metadata tagging (also known as indexing) of the content to allow for easy search and retrieval within the repository. Electronic document storage is often a tiered process, with content that can be archived being moved to a secondary storage device at an offsite location, while content that needs to be more readily

accessed is kept in a primary, high-performance storage device onsite.

Being able to quickly and easily locate important documents in the aftermath of a disaster makes a document capture and management system invaluable to a disaster recovery plan. Electronic document management also allows access to the same documents by staff from multiple sites.

DOCUMENT MANAGEMENT SOFTWARE AND HARDWARE

Document management software provides the framework for turning physical documents into electronic copies and for managing the electronic documents' security and availability. From such basic software as Adobe Acrobat to full-fledged document management systems such as EMC's Documentum and Extensis Portfolio, the choices for effectively managing documents and helping ensure their continuous availability are broad.

Choosing the right document management hardware can make all the difference. For example, a high-speed, sheet-fed scanner can quickly convert reams of paper documents into electronic ones. Scanners from the Xerox DocuMate family, the HP Scanjet family, Fujitsu, Canon and other manufacturers can help speed the process of digitizing paper documents.

DISK-BASED STORAGE

As storage needs grow at an ever-increasing pace, the means of providing storage have changed drastically, especially from a disaster recovery perspective. Direct-attached storage (DAS), where the storage device is attached directly to the server, is one option.

DAS offers two varieties of storage options: mirrored disk and the

various redundant array of independent disks (RAID) solutions. Both of these options protect from the most common form of system interruption, a drive failure. With DAS, there's no risk of network interruptions affecting an organization's storage process.

Storage technologies have advanced with the development of storage subsystems, which are external cabinets that can expand to hold many disks and usually have dedicated processors and caches to control the corresponding RAID solutions. The advantages of storage subsystems over traditional primary storage solutions are higher performance and greater expandability.

Primary storage units have advanced as well, especially the means of protecting them. Besides the various RAID solutions available, such as RAID 5 or RAID 1, organizations now have access to advanced functions, such as snap shooting and cloning.

These functions can automatically copy production LUNs (logical unit numbers) to another location on the system so that if the production copy failed or became corrupted, the cloned copy could be brought into action for nearly immediate recovery.

Snap shooting copies file changes to a disk from a particular point in time forward. So a file can be recovered from the point in time where it is deleted or changed. Cloning involves creating an exact replica of a disk. So unlike snap shooting, cloning can protect against complete data loss.

The introduction of less-expensive disk types such as serial advanced technology attachment (SATA) and serial-attached SCSI (SAS) has lowered the cost of these options by providing higher-density storage, but with decreased performance. The main advantage to SATA and SAS is that the transfer rate is much higher compared to conventional disk.

These newer disk subsystems also allow for advanced features such as remote replication, which involves all of the data on one subsystem being copied to another subsystem. The copying can be done over a wide area network (WAN). This subsystem can either be located locally to protect against a subsystem failure or remotely.

Remote storage subsystem replication can be designed to provide high availability up to and through true disaster recovery and operations continuance by replicating live data over great distances to a remote location designed for hot fail-over in a disaster situation.

CONTINUOUS DATA PROTECTION

Continuous data protection (CDP) monitors an organization's files and as a file is changed or "auto saved," a copy of the changed bytes/blocks is replicated to either a local directory or remote location.

With this automatic reproduction happening constantly, an

organization can have granularity of recovery up to literally the last second. When compared to the traditional previous night backup, this is a boon for organizations needing recovery within a tighter window of time.

Higher functioning CDP recovery is also capable at the disk subsystem level and even the server level. For instance, EchoStream is an AIX (Advanced Interactive eXecutive) application that continuously copies all "writes" to disk and replicates them to an alternative location.

At the alternative location these writes are logged and not only provide the capability of to-the-second recovery, but allow an organization to go back and forth through recovery time and retrieve an earlier version of a particular write. So an organization can recover back to the time right before a corruption occurs and even to points beyond, both near and far.

TAPE ARCHIVING

Tape is a data storage device that reads and writes data onto magnetic tape. It is typically used for archiving and allows for access to data sequentially rather than randomly. Although industry publications have predicted the demise of tape as a primary data storage solution for years, it is the most economical solution for long-term data storage.

Tape is portable (unlike disks, tape can be removed from a drive and taken to another location for recovery or storage), green (tape requires no power other than the read/write drives), dense and very fast.

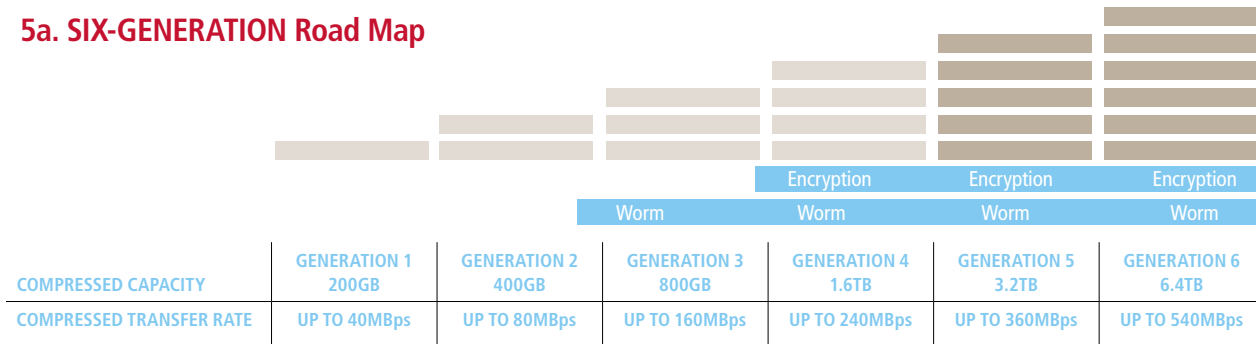
Although not widely noted, tape write rates can exceed 130MBps and can be faster than disk for certain types of data recovery. Tape also has a life expectancy that extends greater than 30 years, making it ideal for long-term archives that may or may not be accessed, yet need to be protected.

Today's most widely used tape standard is Linear Tape-Open. LTO is an open-format technology, making it compatible with a variety of products and media. It was developed to combine the advantages of a linear multichannel, bidirectional format with continued enhancements in servo technology, data compression, track layout and error correction code to maximize capacity, performance and reliability.

LTO was introduced in 2000. Currently, it is on its fourth generation. The generations are readable two releases back, but writeable only one generation back. For example, an organization that upgrades to LTO-4 will only be able to read media that is LTO-4, LTO-3 and LTO-2, and write to LTO-4 and LTO-3 media.

To be protected and current (keeping in mind that speed and density typically double with each version), it is recommended that an organization upgrade to the latest version approximately every five years.

5a. SIX-GENERATION Road Map



SOURCE: THE LTO PROGRAM

Figure 5a shows the current development schedule for LTO. If an organization is currently running generation 1 and wants to migrate to generation 4, it will not be able to use the new drives to do the migration. So organizations need to be aware of the importance of not having a media that's more than two generations older than their current drive level. By staying "current 2," you will always be able to use your new drives to migrate your old tapes.

Tape also helps meet the growing challenge of regulation. It is often the first choice for addressing regulation and compliance issues for how electronic data is stored. Aside from being inexpensive compared to other storage formats, tape also offers easy encryption and "read only" features.

DATA EFFICIENCY, PROTECTION AND ARCHIVING

Cost effective storage of archives can vary a great deal from organization to organization. Each has its own unique requirements that need to be explored and understood in order to develop a cohesive data storage strategy. Protecting data can be accomplished in many different ways, including a combination of disk and tape.

DATA DEDUPLICATION

One approach involves data deduplication. Data deduplication (often referred to as "intelligent compression" or "single-instance storage") is a method of reducing storage needs by eliminating redundant data. Only one unique instance of the data is retained on the storage media.

A disk-based solution featuring data deduplication can allow for virtual compressed quantities of data with smaller amounts

of real disk. Depending on the data, deduplication can provide compression in the 10x or greater ranges. Compare this to tape, which has greater density (1TB or more per tape) and can provide even greater long-term, cost-effective and efficient storage.

Either solution, when architected with enterprise caliber software, provides a very protective approach to storing an organization's data.

Before rolling out a data deduplication strategy, an organization must understand its data access characteristics. This can be done with a storage resource management (SRM) tool, which determines what files are being accessed and edited the most. Organizations can then design a multitiered solution to meet their particular needs.

Data deduplication works great for compressing long-term, limited-access data, such as archives, or for VMware VMDK files, because all of the data is typically similar. Data deduplication is not recommended for high access, high IOP-type data, situations where data is constantly being read and written.

HIERARCHICAL STORAGE MANAGEMENT

Another approach for long-term retention of archival data is a hierarchical storage management (HSM) solution, which involves migrating data from its production location to a lower cost/tier of storage while leaving a "stub" file behind.

The stub file allows applications or file searches to see the file in its normal location, but when accessed, recall the file from its lower-cost location. This lower-cost location can be either a slower disk such as SATA, or even a backup solution such as tape.

The goal of an HSM approach is to reduce the cost of the storage environment by placing data that is infrequently accessed down the performance curve, where slower access is less expensive. ♦

OTHER DISASTER RECOVERY CONSIDERATIONS



CHAPTER 6:

UC for Disaster Recovery

Power and Disaster Recovery

Security and Disaster Recovery

At one time or another, many organizations have implemented various aspects of a unified communications (UC) strategy within their organization: e-mail, voicemail, cell phones, public address (PA) systems. These are all things that are in common use today.

UC FOR DISASTER RECOVERY

The adoption of “one-wire” infrastructures has combined many of these systems onto one data and voice IP network, and the use of wireless technology has further extended the network beyond the four walls of a building.

Computers and IP phones now share a common network along with IP cameras, remote sensors, printers and a whole host of systems and other technologies. With this new infrastructure, the disaster recovery challenge is adapting the network for mass communication and emergency notification.

EMERGENCY NOTIFICATION

In recent years, emergency notification has taken on a new meaning: sending out messages to entire office buildings or an entire campus of buildings. IP phones, sitting on the desks of every staffer, have become more than just phones.

They can also be used to display a text message and play audio broadcasts from the handset speaker. In a matter of seconds, a public safety official can log onto an internal website and send out a prerecorded and pretyped emergency message to any building within the organization.

IP-addressable speakers can be deployed throughout a building or across a campus. These IP speakers are Power over Ethernet (PoE) devices that can be plugged into any CAT 5 jack on the network, turning a voice and data network into a public address system.

There is no need to run separate and redundant low-voltage cable systems or install AC power for speaker amplifiers. Plug in an IP speaker and it joins the network alongside the computers, servers, printers, phones and all the other devices already on the network.

For remote outdoor locations, IP speakers can also be deployed with wireless radios to extend the network. These popular devices can be used everywhere from stadium settings to pedestrian walkways to outdoor common areas.

In the same way that networks can be used for notification, they can also be used for detection. Radio frequency identification (RFID) tags can be placed on expensive technology in an effort to help manage inventory.

When a wireless network is deployed, RFID tags can also be used for security, activating IP cameras or a security notification network to communicate when something is being removed from an authorized area.

Contact closures, such as infrared detectors, pressure sensors or practically any device that registers an open or closed setting can be placed on the network and, when implemented with an InformaCast-enabled network, deliver a highly customized emergency notification communication, addressing the specific needs of an organization.

Activating systems outside the network can also be a requirement of an emergency notification system. Integration with third-party phone and e-mail notification systems enables messages to reach a network of cell phones via prerecorded voicemail or text messages. Notification outside the network can be an important part of emergency protocol, especially when a need exists to have people keep some distance from an emergency situation taking place.

NON-IP NETWORK OPTIONS

A desktop notification system (DNS) can be deployed throughout an organization where IP phones or IP speakers are not yet in use. This small desktop client sits in the system tray of workstations until an emergency notification is sent. Once activated, the DNS displays a text message about the alert while playing an audio broadcast via desktop computer speakers if available.

Integration between legacy paging systems and IP networks is another option. In many cases, these legacy systems can be brought on the network with the installation of a network zone controller paired with a ControlKom platform. This enables audio messages initiated on the IP network to be distributed on the legacy analog PA system, which is a cost-effective alternative for organizations to quickly expand coverage for emergency notification when a full network upgrade is not practical.

Implementing emergency notification takes more than just technology, it takes a coordinated effort between IT staff, administration and the public safety department. Organizational alignment around common goals and well-designed emergency protocols is paramount. Once these are in place, the technology exists today for communicating to staff throughout the organization.

POWER AND DISASTER RECOVERY

When planning for continuity of operations, organizations often devote large portions of their budget to storing and backing up their data. It's important to recognize that these backup devices are still only as reliable as the power that is being supplied to them. The best way to ensure that important IT equipment is receiving consistent, clean and reliable power is to back that equipment up with an uninterruptible power supply (UPS).

A UPS is a device that provides backup power via battery to electronic equipment. Not all UPS products are the same. There are different types of UPS that offer varying types of protection and various additional features. Each application will have its own specific requirements, and choosing the right UPS can mean the difference between seamless operation and outright system failure during a power emergency.

Standby UPS: This is the most basic level of protection. When the UPS detects a drop or spike in voltage or the complete loss of power from the source, it switches over to its internal battery. The standby UPS is designed to protect desktop and workstation systems, retail/POS equipment and any other equipment that is remote to the main data center and requires just basic protection for a period of time long enough to safely save open applications and shut down the equipment.

Line-interactive UPS: This setup offers mid-level protection. In addition to the battery backup supplied by a standby UPS, the line-interactive UPS also contains a transformer that can regulate the incoming voltage, bringing it up or down to an acceptable range rather than failing over to the battery at the first sign of trouble.

Additionally, many models in this category have the ability to run for extended periods of time using additional optional battery modules. This can allow the equipment to ride out most power outages without having to be shut down.

A line-interactive UPS is best suited for environments without a backup generator. They are well positioned to protect servers, storage, and non-VoIP LAN and WAN equipment. It's very common to find this type of UPS in smaller operations, computer rooms and wiring closets.

Online or double-conversion UPS: This is recommended for the highest level of protection. This technology, takes the battery backup and voltage regulation of the line-interactive UPS and adds power conditioning via rectifiers and inverters that convert the incoming power to its purest form before releasing it out to the load.

This technology is critical when protecting a load that is also supported by a backup generator as well as VoIP applications, central data centers and any other equipment that is considered mission critical. And in regions of the country that are known for their harsh power environments, an online UPS is a must.

UPS CONSIDERATIONS

When looking for a UPS solution, there are several factors to take into consideration. Most important, you want to determine the amount of power the organization's devices consume, determine how much future expansion the organization anticipates and select the UPS that will support the capacity needed.

A general rule of thumb is to leave at least 20 percent or more room for additional capacity. UPS capacity can be communicated in VA (volt amperes) or W (watts). Review the wattage ratings between UPS units for a better apples to apples comparison of true capacity.

Next, consider how much runtime (or battery uptime) you will

need. Runtime is the amount of time the UPS will continue to provide backup battery power to your devices.

On average, most organizations are looking for 15 to 30 minutes of runtime. This provides adequate time for applications, operating systems and devices to shut down. With locations that have a generator, a UPS would need to provide a relatively short amount of runtime during the transition of utility power failure to backup generator power, typically no more than five minutes.

Alternatively, some mission-critical applications are in environments without a generator and need an extended amount of runtime. In some cases this can be 60 minutes or more. Most UPS solutions allow you to extend the amount of runtime by adding additional external battery cabinets to them.

When selecting a UPS solution, you will want to figure out the voltage of the devices being supported as well as the voltage of the power available at the site. In the United States, 120V, 208V, 240V, 208V three phase and 480V three phase are the most common types you will run into.

Most often, IT devices will run off of 120V or 208V single phase; however, blade servers and larger equipment may require three phase power. In addition to reviewing the voltage, if you are looking to expand the amount of devices being supported on a UPS, you will want to make sure there is additional power available at the site.

Electrical panels, similar to what you will find in a house, provide for a certain amount of power to be drawn from them. If the organization's electrical panel is at or near capacity, an additional power source will be needed.

Here are other important things to take into consideration:

- **Redundancy:** Similar to a RAID set with hard drives, having a mirrored (or 2N) UPS solution will provide the organization with additional availability in the event the primary UPS fails.
- **Form factor and physical size/space availability:** Take a look at your environment and determine how much room you have for the UPS unit. There are both rack-mount and freestanding UPS solutions.
- **External bypass:** This feature will allow you to disengage power from the UPS while maintaining power to your equipment.
- **Power distribution:** UPS units can provide output power in two ways: outlets on the back of the UPS or a direct connection to an electrical panel via conduit.
- **Services:** Similar to a furnace or air-conditioning system in a house, regular maintenance on the UPS unit is important to prolong the life of the system and to prevent unexpected failures. UPS manufacturers provide startup services to inspect

the wiring, boot up the unit, run diagnostics and ensure the system is functioning properly.

SECURITY AND DISASTER RECOVERY

Information security aims to provide three things for an organization: confidentiality, integrity and availability. Organizations need to keep their secrets secret, ensure that their information is reliable and guarantee access to services when users need them. This last category is often downplayed, especially in recent times when such a large portion of the security effort has focused on compliance initiatives.

When organizations focus primarily on compliance, they emphasize different aspects of security than those engaged in disaster recovery or continuity of operations planning. Compliance efforts tend to focus on information confidentiality and integrity. Disaster recovery, however, deals with making sure that information assets remain reliable and available in the event of trouble.

If a disaster can be defined as any interruption in service that goes beyond an organization's tolerance threshold, then certain security incidents should be considered disasters, and therefore, need to be planned for. Security has two primary concerns with respect to disaster recovery planning: minimizing the likelihood of disasters and minimizing their severity.

DISASTER PREVENTION

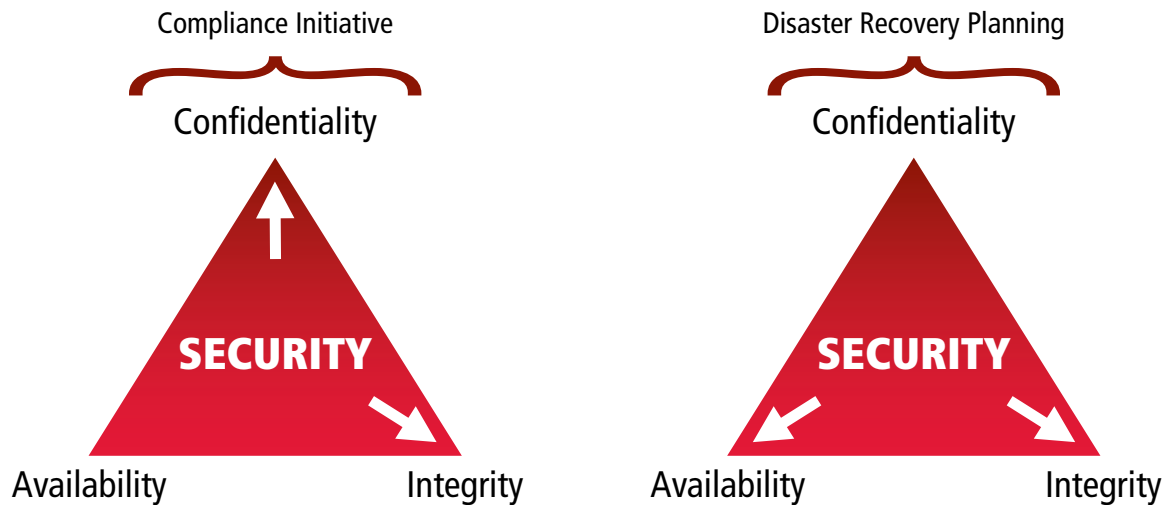
The technologies and processes that we use have flaws, and these flaws, if they are triggered, can precipitate losses. For example, a piece of web server software might have a weakness that causes it to crash if it encounters a specially crafted series of requests. For some organizations, a massive web server failure would be a disaster.

A healthy security program should reduce the probability that an attacker will crash the web server. The most common approaches to solving this particular problem include patching the web server or front-ending it with a proxy or application-layer firewall that prevents the harmful requests from being passed to the vulnerable server.

These two approaches have similar end results, but they call attention to the fact that (in this context) security's goal is to prevent disasters, not merely to chase down vulnerabilities.

Vulnerabilities alone don't make disasters. In fact, many dormant vulnerabilities exist in our systems, all of which remain harmless if never triggered. An actual incident consists of both a vulnerability and the threat that activates it. Threats might include a natural emergency, a hacker or a bot. Vulnerabilities, on the other hand, are exploitable conditions such as a missing patch, a weak password or unwary system administrators.

6a. Compliance vs. DR Planning: Different Focuses



So when an organization adopts disaster prevention as a security goal, there are three basic focuses:

- **Eliminate threats:** This option is not always feasible, but many organizations at least take steps to make themselves less likely targets. For example, they will move critical data centers to areas not plagued by earthquakes or tornadoes.
- **Eliminate vulnerabilities:** Again, while an organization can't eliminate all existing vulnerabilities, it can make an effort to stomp out critical ones.
- **Make vulnerabilities unexploitable:** This approach interposes some measure that prevents a threat from triggering a vulnerability.

Focusing on the various options for preventing disasters is an important part of disaster recovery. Thinking through the prevention strategy is a critical exercise because it provides insight into what combinations of threats and vulnerabilities currently remain uncountered. Information about the types of disasters from which an organization might have to recover obviously serves as a helpful starting point in the recovery plan.

DISASTER MITIGATION

Disaster recovery planning as a whole should aim to minimize the impact of disasters and allow the organization to resume normal operations with as little loss as possible. This objective is a familiar one in security, where the organization strives to manage risk.

Like disasters, risk has more than one component. The following

equation is a simple and common formulation of risk, expressed as an annualized loss expectancy:

$$\text{risk} = \text{frequency} \times \text{impact}$$

Risk is expressed as a figure in dollars per year. That figure represents the product of the number of incidents per year and the average loss expectancy associated with an incident. In the case of disaster-scale incidents, we hope that frequency remains low, but we understand that they have high impact.

Although most organizations devote the bulk of their security dollars to minimizing the frequency of incidents, other measures can lessen the impact of incidents when they do occur. User education is a prominent example. If users know how to respond to signs of trouble (such as a worm outbreak), a potentially serious incident could be nipped in the bud before it blossoms into a full-fledged disaster.

Likewise, compartmentalization of information and assets helps to contain incidents within manageable boundaries instead of allowing them to spread across the enterprise. Investments in fault tolerance, redundancy and capacity are equally investments in system availability.

Finally, a proper incident response plan includes the preservation of key information about the incident. Keeping detailed and accurate records of what went wrong doesn't lessen the severity of a disaster, but it can help greatly when an organization goes to recover costs either through insurance claims or litigation. ♦

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 VIRTUALIZATION

Application virtualization is a technology that improves application compatibility and manageability by encapsulating applications from each other and the underlying operating system. This allows for the streaming of applications to each desktop.

CLIENT ACCESS

This term refers to the method by which applications are delivered to users. Client access is usually built around the deployment of desktops and notebooks that are bundled with local and remote applications. Thin clients, PDAs and workstations also facilitate client access.

CLONING

Cloning is the automatic copying of production LUNs (logical unit numbers) to an alternate location on the network. With cloning, an exact replica of a storage disk is made, helping to protect against complete data loss.

COLLOCATION SITE

A collocation site is a secure, dedicated facility to set up hardware and equipment that usually has a secured cage or cabinet, regulated power, an Internet connection, security and support.

CONTINUOUS DATA PROTECTION (CDP)

CDP is a form of data protection in which 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.

DEDUPLICATION

Data deduplication is an approach to protecting data that eliminates redundant instances of data so that only one unique instance is retained in the storage media. This is a good approach for compressing long-term, limited-access data.

DESKTOP NOTIFICATION SYSTEM (DNS)

This term refers to a system that broadcasts an audio and text message to non-IP-connected workstations. Each workstation is outfitted with a small client that receives and displays the emergency notification.

DIRECT-ATTACHED STORAGE (DAS)

A DAS is a storage device that connects directly to the server, usually either a mirrored disk or a redundant array of independent disks (RAID) solution. This approach offers excellent protection against the most common type of disaster, drive failure.

DISASTER RECOVERY LIFECYCLE

This is the ongoing process of being prepared for disaster recovery that involves five key phases: analysis, solution design, implementation, testing and acceptance, and maintenance.

HIERARCHICAL STORAGE MANAGEMENT (HSM)

HSM is a storage approach for long-term archival data that involves migrating a file from its production location to a lower cost/tier of storage. A stub is left behind and the file appears to be at its original production location, but when accessed, the file is recalled from its tiered location.

HOSTED MANAGED SERVICES (HMS)

HMS is an option for organizations that opt to turn over the operation of their network to a commercial hosting center. These services can include Internet, WAN, firewall, data storage, data backup and disaster recovery. HMSs are typically offered in an on-demand, usage-based model.

HYPERVERSOR

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

LINEAR TAPE-OPEN (LTO)

LTO is an open-format technology that is usually upgraded every 18-24 months. It has the advantages of a linear, multichannel, bidirectional format combined with continued enhancements in servo technology, data compression, track layout and error-correction mode for maximum capacity, performance and reliability.

LOAD BALANCING

Load balancing is a data center technique where processing work is split between two or more servers so that the work gets done more efficiently. All network users receive faster service as a result, and the network is more resilient.

OPERATIONS IMPACT ANALYSIS (OIA)

An OIA is an essential part of a disaster recovery or continuity plan. It has an exploratory phase to reveal vulnerabilities and a planning phase to develop strategies for reducing risk. The OIA allows the organization to identify the monetary costs related to failures. It quantifies the importance of the various operations components and suggests appropriate budget allocation.

REDUNDANT ARRAY OF INDEPENDENT DISKS (RAID)

RAID is a category of disk drives for data storage. Two or more drives are used for increased performance and fault tolerance.

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. A SAS switch enables servers to connect to multiple SAS storage arrays.

SERVER CONSOLIDATION

Server consolidation is an efficient approach to utilizing server resources in order to decrease the total number of servers and/or server locations that an organization is using.

SNAP SHOOTING

A disk-based data storage method in which changes to a file are copied from a particular point in time forward, allowing the file to be recovered from the exact point in time where it was deleted or changed.

SPLIT-BRAIN DATA CENTER

A split-brain setup is an arrangement of the server farm where it is split between the production and the recovery data centers. This setup allows for greater utilization of servers, and in the event of a disaster only half of the servers are affected.

THIN CLIENT

A thin client is a low-cost display-only computer where the bulk of the data processing occurs on the server, including applications and data. Thin client devices have no moving parts.

UNIFIED COMMUNICATIONS (UC)

UC refers to a "one-wire" infrastructure where numerous systems such as e-mail, voicemail, cell phones, PAs, printers and Internet all reside on a single data and VoIP network. Emergency notification systems can make use of a UC setup as well.

VIRTUAL DESKTOP

A virtual desktop refers to a desktop that is hosted on a server and published to the actual end-user computer. It's the end result of a server-based computing model that is designed to give administrators the ability to manage desktop virtual machines in the data center while giving end users a full desktop experience.

WAN OPTIMIZATION

WAN optimization improves the performance of a WAN by accelerating network protocols, compressing traffic and caching files. This has the effect of speeding up core applications such as Exchange, Citrix and web-related applications.

INDEX



Application access/delivery	5-7	Hierarchical storage management (HSM).....	28
Carrier diversity	16	Hosted managed services (HMS).....	14
Carrier services.....	15	Load balancing.....	14-15
Client access	5-7	Remote network management	13-14
Continuous data protection (CDP)	27	Risk assessment	11, 15-16
Data deduplication.....	28	Risks (of not being prepared)	3-4
Data storage	26-28	Server consolidation.....	9-10
Desktop notification system (DNS).....	30	Server recovery.....	10-11
Direct-attached storage (DAS)	26-27	Server virtualization	6, 8, 12
Disaster mitigation	32	Tape archiving.....	27-28
Disaster prevention	32	Thin clients.....	5-6
Disaster recovery lifecycle	4	Unified communications	29
Disk-based storage	26-27	Uninterruptible power supply (UPS).....	30-31
Document management	26	UPS considerations	30
Emergency notification	29-30	Virtual desktops.....	6-8
Financial justifications (for recovery plan).....	11-12	WAN acceleration	15

Disclaimer

The terms and conditions of product sales are limited to those contained on CDW•G's website at CDWG.com. Notice of objection to and rejection of any additional or different terms in any form delivered by customer is hereby given. For all products, services and offers, CDW•G® reserves the right to make adjustments due to changing market conditions, product/service discontinuation, manufacturer price changes, errors in advertisements and other extenuating circumstances. CDW®, CDW•G® and The Right Technology. Right Away.® are registered trademarks of CDW Corporation. All other trademarks and registered trademarks are the sole property of their respective owners. CDW and the Circle of Service logo are registered trademarks of CDW Corporation. Intel Trademark Acknowledgement: Celeron, Celeron Inside, Centrino, Centrino Inside, Core Inside, Intel, Intel Logo, Intel Atom, Intel Atom Inside, Intel Core, Intel Inside, Intel Inside Logo, Intel Viviv, Intel vPro, Itanium, Itanium Inside, Pentium, Pentium Inside, Viiv Inside, vPro Inside, Xeon and Xeon Inside are trademarks of Intel Corporation in the U.S. and other countries. AMD Trademark Acknowledgement: AMD, the AMD Arrow, AMD Opteron, AMD Phenom, AMD Athlon, AMD Turion, AMD Sempron, AMD Geode, Cool 'n' Quiet and PowerNow! and combinations thereof are trademarks of Advanced Micro Devices, Inc. This document may not be reproduced or distributed for any reason. Federal law provides for severe and criminal penalties for the unauthorized reproduction and distribution of copyrighted materials. Criminal copyright infringement is investigated by the Federal Bureau of Investigation (FBI) and may constitute a felony with a maximum penalty of up to five (5) years in prison and/or a \$250,000 fine. Title 17 U.S.C. Sections 501 and 506. This reference guide is designed to provide readers with information regarding disaster recovery and continuity of operations planning technology. CDW•G makes no warranty as to the accuracy or completeness of the information contained in this reference guide nor specific application by readers in making decisions regarding disaster recovery and continuity of operations plan implementation. Furthermore, CDW•G assumes no liability for compensatory, consequential or other damages arising out of or related to the use of this publication. The content contained in this publication represents the views of the authors and not necessarily those of the publisher. ©2009 CDW Government, Inc. All rights reserved.



CDWG.COM/DISASTERRECOVERYGUIDE
888.667.4239

ABOUT THE AUTHORS

» LANCE CASEROTTI is a Network Solutions Architect for CDW, and has designed, implemented or managed numerous data center builds and disaster recovery sites.



« PEYTON ENGEL is a Technical Architect for CDW, and leads a team of security engineers.

» NATHAN COUTINHO is a Solutions Manager for CDW, with a focus on virtualization.



« GURPREET SACHDEVA is a Solutions Architect for CDW, with a focus on LAN/WAN solutions and network-based security.

» SPENCER CAGLE is a Solutions Architect for CDW, focusing on networking, application optimization and WAN optimization solutions.



« BARI QURESHI is a Network Solutions Architect for CDW, with more than 10 years experience designing and building internetwork architectures.

» RANDALL FOLTYNIEWICZ is a Power and Cooling Sales Manager for CDW, and leads a team of power and cooling specialists.



« BRAD PARKEL is Director of Marketing for Singlewire Software, with extensive experience in technical marketing and a deep knowledge of the InformaCast system.

DISASTER RECOVERY AND CONTINUITY OF OPERATIONS PLAN REFERENCE GUIDE

090701 • Flyer 59337AB

LOOK INSIDE for more information on:

- Maintaining client access
- Preparing for server recovery
- Utilizing network load balancing
- Managing data storage

