

CONTINUITY OF OPERATIONS REFERENCE GUIDE



Keeping your resources highly available

CDWG.com/coopguide | 888.559.4239



The Right Technology. Right Away.®

CONTINUITY OF OPERATIONS REFERENCE GUIDE

TABLE OF CONTENTS CHAPTER

WHAT IS A CDW•G REFERENCE GUIDE?

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

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

01	Preparing for Future Threats	3
	• Confronting the Online Threat	
	• The Disaster Recovery Lifecycle	
	• The IT Department's Role in COOP	
02	Data Access and Storage	5
	• Data Storage Solutions	
	• Data Efficiency, Protection and Archiving	
03	High Availability for Servers	8
	• Consolidation for Easy Recovery	
	• Solutions for Physical Servers	
	• Solutions for Virtual Servers	
	• Making the Financial Case	
	• Migrating to Virtualization	
04	Networking and Disaster Recovery	21
	• Remote Network Management	
	• Network Load Balancing	
	• WAN Acceleration	
	• Carrier Services	
05	Client Access	25
	• Understanding Client Access	
	• Thin Clients	
	• Crafting the Right Strategy	
06	Other Disaster Recovery Considerations	29
	• Unified Communications	
	• Power and COOP	
	• Environmental Concerns	
	GLOSSARY	33
	INDEX	35



PREPARING FOR FUTURE THREATS

HAVING A PLAN IN PLACE TO RECOVER OPERATIONS IS ESSENTIAL

CHAPTER 1:

.....
Confronting the Online Threat

.....
The Disaster Recovery Lifecycle

.....
The IT Department's Role in COOP
.....

There's a great deal more to an IT infrastructure than hardware and software. This infrastructure also encompasses data, systems, processes, strategies and even people. It reaches into just about every organizational department and operation. Often, the full value of the IT infrastructure is not noticed until it is under threat by some form of disaster.

So critical have electronic applications become to the minute-by-minute operations of government agencies and educational institutions that even a few moments of downtime is no longer acceptable.

Users have come to expect that information and access to IT benefits and services will be available online, 24x7x365. Such high expectations for network availability have increased the importance of having a continuity of operations plan (COOP).

This reference guide discusses the tools and technologies that will help keep your operations running in the face of an emergency. It will focus particularly on servers, networking, data storage, power protection and client access. Prepared with the knowledge and background that you'll need to craft a viable COOP — and a source for the technology — you will be ready when the worst-case scenario becomes front-page news.

CONFRONTING THE ONLINE THREAT

Internet activity has infiltrated nearly every aspect of our lives: college admissions applications, state and local government vital records, financial benefits for government and educational organizations, business license applications, and regulatory information. Online government and educational operations have become a big part of our everyday life.

For staff and administrators with long careers in the government and education fields, a common observation by many is that each year seems to bring more work online and correspondingly less paperwork. As all of these operations migrate online, they become vulnerable to a variety of complex threats. Here are a few of the concerns that have emerged:

- Cybersecurity threats are morphing from “gotcha” interruptions or web announcements into more stealthy — and more costly — financially motivated attacks.
- Greater staff mobility, coupled with the increasing use of smartphones as network application endpoints, means more data is in motion to devices that are more of a challenge for IT to secure.
- The looming threat of disruptions to staff getting to work (such as pandemic and bioterror situations) means that remote access systems must always be kept in a state of readiness.
- All government and educational institutions should be concerned about threats to the power grid itself. This may be the ultimate challenge to any organization's COOP.

Your COOP cannot leave out planning for natural disasters: fires, floods, earthquakes, hurricanes and tornados. The far reach of these perennial threats can put both local and remote operations at risk.

With so many threats to consider, it quickly becomes clear that an organization's COOP must incorporate elements of disaster recovery. But a COOP's vision must be even broader, encompassing threats to software systems and internal staff actions that compromise ongoing operations.

THE DISASTER RECOVERY LIFECYCLE

Creating a COOP should not result in a document that your organization puts on a shelf and forgets about. Like the systems and processes it protects, a COOP has a lifecycle. The plan needs to be re-evaluated periodically, preferably after each test. And because the COOP is derived from operations acumen, it really needs to be tested on a schedule.

There are five unique phases in the disaster recovery lifecycle.

1. **Analysis:** The foundation of the COOP, the analysis phase is when the IT team determines the most potent and the most likely threats to occur, what the impact of each threat could be, and what the organization should do if the threat occurs. This is also known as the risk calculation phase.
2. **Solution design:** The goal here is to identify the most cost-effective solutions to identified threats that are technically viable.
3. **Implementation:** This phase consists solely of the execution of the design elements identified in the solution design phase.
4. **Testing and acceptance:** To be certain that the disaster recovery plans and/or the COOP meet the needs of the organization, testing is required to assure processes and acceptance. Testing is not merely an IT exercise, but should involve all stakeholders in the system.
5. **Maintenance:** Once a COOP and/or a disaster recovery plan are established, regular maintenance of the plans helps ensure continued viability. The maintenance phase is the ongoing effort to address technical solution needs, recovery solution needs and organizational changes as they affect operational preparedness.

From the very beginning and throughout the disaster recovery lifecycle, the focus must always remain on managing risk to the operations environment and maintaining continuous availability. Threats such as data loss, service failure, power outages, software incompatibility and security concerns put the operations environment in a continuous state of risk.

THE I.T. DEPARTMENT'S ROLE IN COOP

The design and execution of a COOP is the responsibility of the IT department under the CIO, but IT cannot create the plan in a vacuum. Organizational leaders must be active participants from the outset.

The organization's top management gets the assignment of brokering resource decisions between IT and the operations units. With detailed knowledge of their users' access needs, the legal and regulatory environment in which they are operating, and the specific costs of downtime, management is in the best position to establish how long systems can reasonably be down.

This information can help the IT department develop the recovery point objective (RPO). Calculated as the time between scheduled data backups and the data created between backups, RPO is an organization's tolerance for data loss. A system failure occurring a moment before the next backup would result in a loss equal to the RPO.

A logical response to the threat of system failure would be to do more frequent backups, and thus reduce the RPO. But this approach is financially shortsighted, as more frequent backups will drive up costs. So organizations eventually strike a balance between reducing the RPO and keeping the costs of backing up data within budget.

Keep in mind that different organizational systems can have different RPOs. The more mission-critical systems will have a more aggressive RPO, and the less mission-critical systems will allow for a longer RPO. This lets an organization manage costs while at the same time staying highly available where needed.

Whether the data center is outsourced or operated by staff, another critical piece of the COOP is the recovery time objective (RTO) specified in the plan and, when applicable, in the service-level agreement with an organization's IT service provider. The RTO presumes that manual workarounds for the nonfunctioning electronic transactions are still generating data.

Once the organization has calculated and agreed on the parameters of the COOP, the goal for your plan is to ensure that total recovery falls within the limits you've set. ♦

Toward Total Availability

The different categories of technology listed below can all contribute to the goal of providing 100 percent availability and ensuring that recovery from any disaster is quick and complete. To choose specific manufacturers and products, it's best to rely on research, information from colleagues and field trials whenever you can get them.

Total availability is supported by the following technologies:

- Uninterruptable power supply (UPS) units
- Backup hardware
- Backup software
- Storage area networks (SANs)
- Replication and clustering technologies
- Virtualization



DATA ACCESS AND STORAGE

PLANNING FOR AND PROTECTING YOUR MOST VALUABLE RESOURCE: DATA

CHAPTER 2:

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 options available in order to put a data storage solution in place that fits your organization's unique needs.

DATA STORAGE SOLUTIONS

Disaster recovery plans and COOPs tend to focus heavily on technological elements, while neglecting the enormous investment most 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 successful disaster recovery plan or COOP.

After all, electronic documents are more easily backed up, stored offsite and recovered in the event that 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, which allows 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 available 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 most organizations' storage needs continue to 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 your data 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 include 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 logical unit numbers (LUNs) 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 every time the file is changed. This allows for file recovery from the most recent version of a file after a failure. 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 rates are 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 failover in a disaster situation.

Continuous Data Protection

Continuous data protection (CDP) is a storage system that 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. For mission-critical systems that need an aggressive RPO, this approach is invaluable.

Higher functioning CDP recovery is also capable at the disk subsystem level and even the server level. Many storage vendors offer applications that continuously copy 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. This allows an organization to recover back to the time right before a corruption occurs, and even to an earlier version of the backup write.

Tape Storage

Tape is a data storage device that reads and writes data onto magnetic tape. It is typically used for low-cost, long-term storage and facilitates 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.

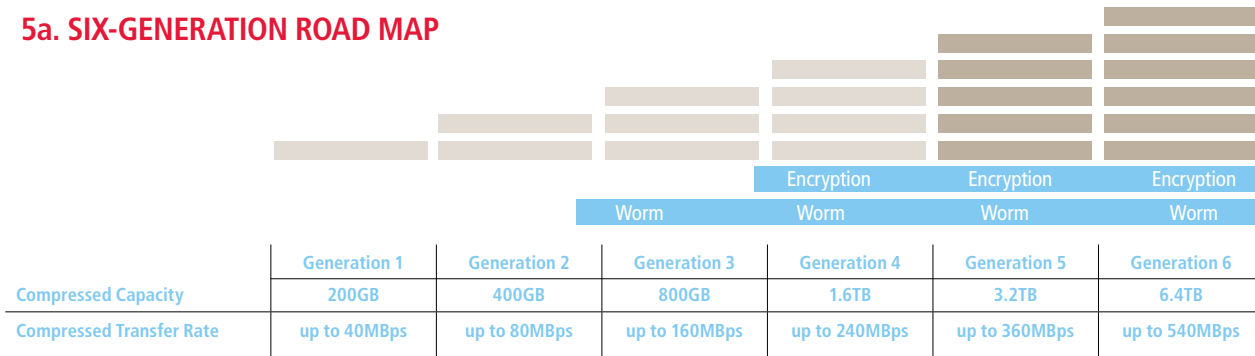
Tape is portable (unlike disks, tape can be removed from a drive and taken to another location for recovery or storage), power efficient (tape requires no power except during read/write sessions), dense and very fast.

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 exceeds 30 years, making it ideal for long-term storage that may or may not be accessed, yet needs 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 technology was introduced in 2000. Currently on its fourth generation, a fifth generation of LTO capable of 3.2TB of compressed storage and data transfer rates of up to 360MBps is scheduled to be standardized sometime in 2010. The generations are readable two releases back, but writeable only one generation back. For example, an organization that upgrades to LTO-5 will only be able to read media that is LTO-5, LTO-4 and LTO-3, and write to LTO-5 and LTO-4 media.

5a. SIX-GENERATION ROAD MAP



Source: The LTO Program

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

Figure 5a shows the current development schedule for LTO. If an organization is currently running generation 2 and wants to migrate to generation 5, it will not be able to use the new drives to do the migration. 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 is an ideal technology for meeting 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.

COOP: Data Deduplication vs. Archiving

Data deduplication is a widely used approach to data protection, and with good reason. It greatly reduces the amount of data that needs to be stored. Often referred to as "intelligent compression" or "single-instance storage," data deduplication is a method for eliminating redundant data. Only one unique instance of the data

is retained on the storage media.

Data deduplication works great for compressing long-term, limited-access data, such as archives or for VMware VMDK files, because typically all of the data is similar. But data deduplication might not be a good fit for your COOP. It's not generally recommended for high-access, high IOP-type data — the kind of day-to-day operations data that you need to get your organization up and running following a disaster situation.

And from a strategic perspective, data deduplication focuses on reducing the amount of data you have, while with your COOP, you are focusing on replicating copies of data to multiple sites.

An alternative, COOP-friendly approach to reducing data would be archiving. Archiving includes single-instance storage (similar to data deduplication), but because the process is an integral part of your backup plan, the risk of losing that single data copy is greatly diminished. Archiving has the additional benefits of supporting the lifecycle management of data, meeting compliance regulations and helping create a strong foundation for e-discovery.

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, it can 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. ♦



HIGH AVAILABILITY FOR SERVERS

IMPROVING SERVER AVAILABILITY FOR OPERATIONS CONTINUITY AND DAY-TO-DAY USE

CHAPTER 3:

.....
Consolidation for Easy Recovery

.....
Solutions for Physical Servers

.....
Solutions for Virtual Servers

.....
Making the Financial Case

.....
Migrating to Virtualization
.....

Developing a comprehensive and cost-effective recovery strategy for servers is a challenging undertaking for organizations. In fact, some organizations are surprised by the high cost of implementing a recovery site and instead choose to do nothing, hoping 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 server loads onto fewer servers through virtualization has been the consolidation method of choice for the past few years. This technology has allowed data centers to better utilize computer resources while reducing risks associated with planned and unplanned downtime in the infrastructure. Virtualization also reduces power and cooling consumption and frees up space in the data center.

Recent resource usage studies verify very low utilization rates among servers in use today, suggesting there is room for many organizations to achieve higher usage numbers through virtualization. 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-to-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 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 after a disaster.

But having a simplified server environment is not the only COOP benefit. With virtualization, you also significantly improve server availability. Virtualized server environments from one compromised physical server can quickly be brought back up at another safe location following a disaster — a boon for your COOP.

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).

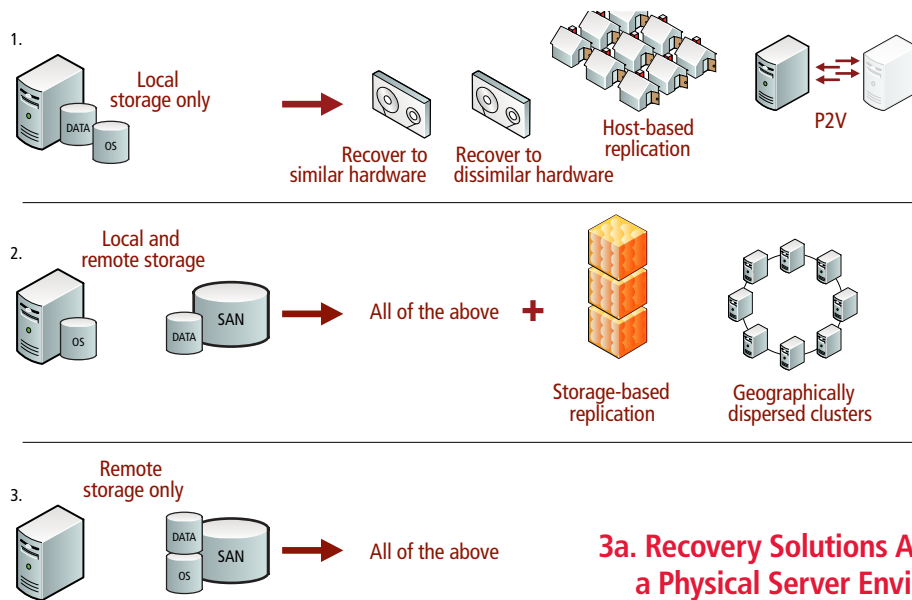
SOLUTIONS FOR PHYSICAL SERVERS

Although the future of the data center is clearly virtual, the number of physical servers deployed still outnumbers that of virtualized ones. Also, 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 failover of the physical servers.

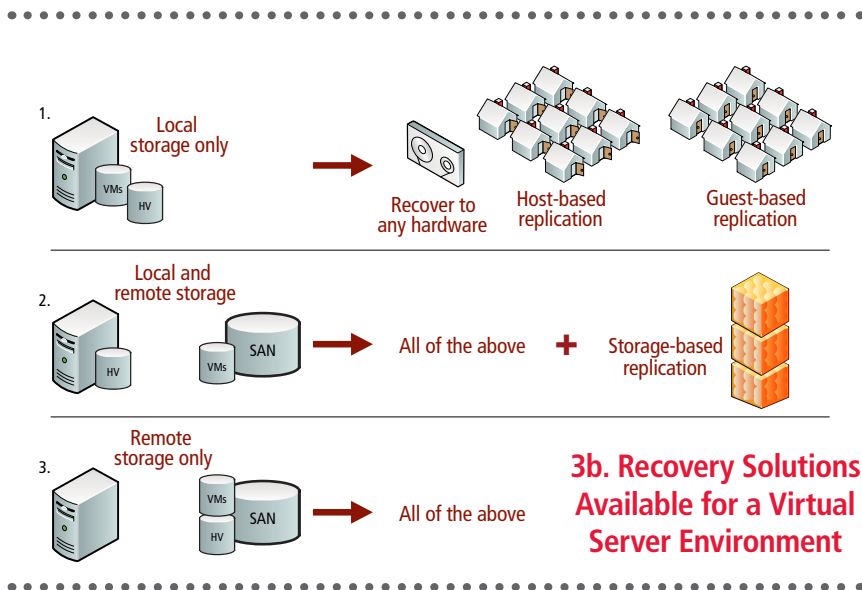
Figure 3a shows recovery solutions available for a physical server environment and is divided into three options that are explained as follows:

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

- **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 perfect 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



3a. Recovery Solutions Available for a Physical Server Environment



SOLUTIONS FOR VIRTUAL SERVERS

Organizations that have migrated 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:

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 from which to choose.

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.

same time) are slim to none.

- **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.

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.

MAKING THE FINANCIAL CASE

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. Here are some tips and suggested strategies to help gain support for budgetary outlays for upgrading your data center.

- **Operations recovery assessment:** Finding out the true cost of downtime for an organization can be a very difficult process, which is why most organizations outsource their operations recovery assessments 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 shrinks with regard to power, cooling and backup windows.
- **Widespread virtualization:** 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 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 then 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.

MIGRATING TO VIRTUALIZATION

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 utilized 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. ♦



NETWORKING AND DISASTER RECOVERY

FINE-TUNE YOUR NETWORK TO WITHSTAND A DISASTER

CHAPTER 4:

Remote Network Management

Network Load Balancing

WAN Acceleration

Carrier Services

Much of the thinking around your COOP involves planning for (hopefully) rare, worst-case scenarios that may affect your operations and network. With all this focus on future events, it's important not to lose sight of your network's day-to-day operations and look for opportunities where your COOP strategies might overlap with practices for running an efficient network.

When considering your network as you formulate your COOP, you are well advised to keep two goals in mind.

Disaster recovery should be your primary goal if flood, fire, theft, vandalism, component failure or some other calamity strikes a data center, an outside part of the network infrastructure or a remote site. Given the distribution of modern networks, chances are good that critical pieces of your fine-tuned network are located outside of one of the main data centers.

A second goal should be maintaining workable service when the network is overloaded or when failover mechanisms increase the load on the remaining parts of the infrastructure. That is also part of continuity of operations planning. In this chapter, we'll explore the main aspects of COOP thinking as they apply to the network.

REMOTE NETWORK MANAGEMENT

Although enterprise networks are often likened to the human nervous system, if it really were similar to a web of nerves, the network would communicate with the system administrator automatically when something goes wrong. When you stub your toe, it doesn't take long for the error message to register in your gray-matter console. For a network to similarly convey a problem,

it must be monitored.

A network requires monitoring to manage performance and detect threats, which can present the ultimate performance problem. Few enterprise networks can be monitored directly by the systems administrator, so network management is primarily concerned with remote monitoring (RMON).

An efficient RMON system takes into account the various vulnerabilities of devices on the network and their potential for causing damage or halting operations. For example, a compromised firewall is a far more serious matter than a jammed printer. Your RMON system should balance security considerations with the need for speed that the organization is paying to maintain.

As an organization strives to maintain this balance, it is important to integrate cybersecurity, physical security and the performance parameters subject to RMON. Integration also extends to the monitoring of application performance, not just network traffic.

Cybersecurity threats are increasingly focused on hijacking applications, rather than the hijacking of an entire network. One example of this phenomenon is a man-in-the-browser attack, in which a Trojan horse attaches itself to applications through which credential and identity management information flow and gathers this confidential data. The profiles and presence of such attacks can only be detected remotely.

Effective RMON requires more than simply installing an off-the-shelf package on the management console and its agents on remote devices. It requires planning to establish acceptable

performance parameters of the devices to be monitored, the deviations sufficient to set off alarms and the assignment of priority designations so that an out-of-toner printer doesn't generate the same level of alarm as a distributed denial-of-service attempt on a crucial router.

RMON also requires staff trained to understand what the monitor console is telling them. Your staff is often the most expensive component in the RMON chain and, in some ways, the most vulnerable. At some point, there is simply too much to watch, so a well-designed RMON scheme should spare the operators the need to pay attention to all but the most critical deviations from normal operations.

For many organizations, the answer to the staffing question is to outsource this task to contractors who specialize in remote monitoring. Outsourcing remote network monitoring converts the staff and overhead calculus into a managed service for which the agency sets the service-level agreement (SLA). The SLA should support your monitoring parameters as well as the maximum tolerable downtime established in your COOP.

An outsourcing option can allow for easier scaling (up and down) as branch offices, bureaus and remote users are added or subtracted. But it's important to do the requisite research for the contractors who bid on your organization's proposal. Pay particular attention to the certifications of the staff specifically assigned to your network and whether they have experience in your application environment.

NETWORK LOAD BALANCING

In many ways, network load balancing is at the technical center of the real objective of your COOP, namely organizational resiliency and the ability to maintain critical processes during a crisis, whatever kind it might be.

Load balancing is not a difficult concept to understand. Basically, it is the routing of network traffic among physical resources in such a way as to maintain optimal performance. But load balancing also allows for failover from one IT asset to another in the event of a crash or other interruption.

It also produces redundancy of resources, again both for Quality of Service considerations and for continuity. Load balancing can also enhance cybersecurity by presenting a logical picture to users (including those with bad intentions) that hides the specifics of network resources, the visibility of which can reveal your network's vulnerabilities.

One way to look at load balancing is to divide it into two basic elements. Network load balancing makes efficient work of traffic without regard to the application. It sees packets without context, operating at the lower levels of the network protocol

stack. Application load balancing operates above Level 4, with awareness of application-layer packets.

As more applications are developed as front ends with browsers and multiple servers, the requirements of session persistence and reliability mean application balancing must be part of your overall load balancing setup.

With public sector organizations embracing Web 2.0 applications, both internally and when interacting with the public, the nature and volume of network traffic increases the need for load balancing. Web 2.0 (also known as Rich Internet Application) may offer a terrific end-user experience and enhance collaboration, but it makes for more frequent and more granular back-and-forth network traffic.

Load balancing is more complex in actual implementation because there are several techniques for achieving it. These techniques may make use of network address translation, load allocation among multiple servers, failover automatic partitioning and global load balancing across a geographically scattered WAN.

Adding to its complexity, load balancing as a function works in various scenarios: within a single data center, across multiple data centers, across multiple IP links (often to the same website) and among storage subsystems.

Finally, load balancing applications themselves can operate in virtualized situations alongside other applications. But most administrators prefer dedicated hardware because the load balancer itself may have to play traffic cop in the event of a server crash, and you don't want it mounted on the same hardware that is causing trouble in the first place.

With a virtual load balancer trying to keep up with virtual servers, bottlenecks can develop. This is especially true in situations where different departments are adding virtual servers without the assistance of IT for noncritical reasons, which happens often enough in virtualized environments.

Successful deployment of load balancing technology itself requires striking a balance that gives equal weight to the cost savings from efficient use of network resources, the Quality of Service objectives for your applications and classes of users, and the need to maintain your COOP objectives of resiliency, redundancy and reliability.

WAN ACCELERATION

Over the years, wide area networks have grown exponentially faster and more critical to public sector organizations that are deploying applications to the public, have a remote workforce or are operating in more than one location. These conditions coexist within many government agencies and educational institutions,



and the WAN is the common communications denominator.

Also adding criticality to the WAN is the operations continuity concern of locating computing resources in geographically dispersed locations. Physical disaster remains a constant threat.

Optimizing the performance of the WAN also benefits operations continuity by keeping bottlenecks and performance snafus from occurring at critical times, such as when large numbers of staff need to be notified at once of some condition or event, or when large numbers of users might be expected to access your network resources simultaneously.

The term “WAN acceleration” is really an amalgam of two different but related technical approaches to keeping the WAN operating smoothly and fully utilizing the bandwidth that the organization is paying for. Put another way, WAN acceleration allows you to give users the experience you intend for them, while running your infrastructure at optimum efficiency.

The two approaches that make up WAN acceleration are application acceleration and WAN optimization. Without apps, a WAN merely has potential, with no need for optimization. And networked applications don't deliver value until they can be accelerated to work smoothly over the network.

Application acceleration software filters applications so that they can travel across the WAN more efficiently. This is necessary because many applications with a web focus are built on traffic-intensive serial protocols such as Common Internet File System (CIFS) or Messaging Application Programming Interface (MAPI).

On today's networks, these protocols are found in a variety of applications that are being used in mission-critical ways, such as Voice over Internet Protocol, video and other rich data; thin client and mobile device deployments; and software-as-a-service or cloud applications.

Many organizations with stringent privacy and security concerns are employing data encryption, which can add to network traffic (although the processing overhead of the encrypt/decrypt

application typically occurs at the endpoints of the network).

Application accelerators often use a technique called protocol optimization. TCP/IP and other protocols can be “chatty,” generating a lot of back and forth traffic. By analyzing the traffic and anticipating user requests, optimizers can smooth out the traffic by reducing latency.

Object caching is an application- and protocol-specific technique that reduces excess traffic by analyzing parts of packets that can be stored at the client or server. This reduces the need to move the packets back and forth across the WAN. It also reduces the need to re-create or rebuild them, cutting down on processor overhead as well as bandwidth demand.

By reducing the overhead otherwise associated with traffic-intensive applications and protocols, application accelerators have the added benefit of enabling a higher degree of virtualization.

WAN optimization is used for more traditional file transfer applications, such as the movement of spreadsheets and other documents, or for the transfer of batch processing results.

One technique to smooth out traffic is object caching. Unlike byte caching, this technique works without regard to the application or protocol. It also has the effect of reducing network traffic, because needed pieces of calls and responses are stored for reference at appliances at each end of the network.

Another way to streamline traffic is by using file compression. There are many techniques available with varying ratios of compression depending on where in the chain of file movement they occur, and whether they include the cached information.

Network managers also have a few brute force techniques that prevent WAN clogging. For example, simple communication limits can be placed on “bad” applications that you don't want moving across the WAN at all, let alone optimized or accelerated. They include, among others, web advertising, Skype yakking, spyware and peer-to-peer file sharing.

Network managers can also impose transfer rate limitations on

certain applications, users or types of traffic, or traffic-shaping measures suited to point-to-point applications. These measures, while freeing up bandwidth, are likely to be noticeable to users.

Acceleration of applications and optimizing the WAN go hand in hand. The primary day-to-day goal of using these technologies is maximum efficiency of bandwidth and minimum hassle to application users. But they play an important role in your COOP.

CARRIER SERVICES

Just as networks of semi-trailer trucks are dependent on oil refineries for fuel, computer network operators depend on telecommunications carriers for bandwidth, the fuel of networks. The analogy goes further: Neither diesel fuel nor telecom bandwidth is inexpensive or unlimited, and both are subject to spot shortages or interruptions.

So no matter how well thought out and tested your COOP might be, to some degree, you are at the mercy of your primary telecommunications carrier. The same is true for your electrical utility, which is discussed in Chapter 6. That's why carrier performance and backup are integral to COOP design.

Luckily, enterprise network operators have some say in the service-level agreements they exact from carriers (a great deal more, in fact, compared to the contracts that the average person agrees to for telecom service).

Carrier diversity — more precisely, the potential of diversity — is central to the continuity plans of many public sector organizations. This typically includes collocation, or the occupying of a cage or a cabinet in a multiple-carrier facility. Should a failure of a particular carrier occur, backup carriers are available on the premises.

Keep in mind that switching carriers can be a time-consuming activity. It may mean you have to remap all your network address translation rules as well as require a physical switchout of network interface devices. This means that a periodic test of procedures should be part of the COOP.

Another question to consider: If you temporarily switch carriers, is it necessary to have the same amount of bandwidth on tap as you have with the primary carrier? For single-campus situations, it may not be necessary if the crucial operations processes can carry on.

For example, small to midsize colleges can continue most classroom activities, as well as many K-12 schools. State and local government-run hospitals relying on electronic patient records and clinical systems may not fare well during a carrier interruption. And government agencies with, for example, call centers for tax, unemployment or healthcare beneficiaries, would likely not be able to function without full bandwidth replacement.

This means that when analyzing risks, the planning team must

make careful tradeoffs between costs and applications.

In some instances, it might be economical for a second carrier to simply feed a secondary circuit right into the building when dealing with a remote office or standalone facility. At the other end of the expense and capability spectrum is a geographically distant backup data center that might share some of the daily processing load but can come online by itself in the event of disaster at the main center.

Good network management, like effective COOP design, requires a clear picture of how the applications, network components and bandwidth services work together to keep data flowing, regardless of the circumstances. ♦

Networking Checklist

Here are five tips to help strengthen your network while also supporting your COOP.

- **Pay close attention to network latency and bottlenecks.** Simply adding more capacity can be cost prohibitive and inefficient, because you are building out capacity that might be idle much of the time. A better approach is to apply acceleration techniques that tame the behavior of modern, traffic-intensive applications. These techniques increase the odds of application and data availability in the event your COOP is invoked.
- **Don't overlook the security implications of load balancers.** Remember that load-balancing network traffic among physically disparate resources boosts your cybersecurity by making the network topology less visible to would-be intruders. It also reduces the risk of physical damage to one location or set of resources bringing your WAN to a halt.
- **Integrate remote monitoring of device performance with physical surveillance.** In a COOP situation, the two can be related. For example, if a physical occurrence such as a fire or flood causes a device to stop operating, there are two possible means by which you can be notified.
- **Use two carriers, or at least have a backup carrier.** For an enterprise WAN, a single carrier means a single point of failure that renders all other COOP measures moot. So it's invaluable to have a second carrier to fall back on should the need arise.
- **Monitor the wireless devices that are issued by your organization.** To the extent that the traffic they generate is linked downstream to the wired WAN, wireless devices have a hidden potential to cause unexpected loads that may affect your network's performance.

CLIENT ACCESS

ARRANGING FOR END-USER ACCESS TO THE NETWORK



CHAPTER 5:

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 typical 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 staff members 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 in place to replace them as quickly as possible.

UNDERSTANDING CLIENT ACCESS

Here are some common devices, technologies and strategies that will ensure safe, on-demand remote access for clients:

- Smartphones
- Smart cards
- Key fobs
- Encryption software
- Client virtualization
- notebooks

The most common implementation of all the deployment models is assigning physical desktops and notebooks 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 (now known as XenApp) 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 simple operating system and connects to a remote operating system and applications, thereby only sending screen pixel changes, mouse clicks and keyboard key strokes across the network. 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.

Thin Clients: Greater Productivity, Less Downtime

According to an IDC survey, organizations that utilize thin clients see a 78 percent increase in IT staff productivity and an 88 percent decrease in end-user downtime (on average).

Terminal Services

Both Citrix XenApp and Microsoft Terminal Services produce similar products that can host an end user's desktop sessions 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 is to understand the user's needs and develop a flexible solution to provide application access to many types of users. Fortunately, there are a number of methods for actually accessing the published applications.

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 presented with a desktop 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 has a consistent experience, no matter what endpoint 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, VMware and Symantec that isolate applications from the operating system. This allows applications to run completely isolated from each other, with minimal effect 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 has now reached maturity and widespread adoption. As a result, using the same technology to host desktops

is quickly becoming an approach to address issues with managing PCs. Virtual desktops work by allowing an end user to access a desktop operating system (with or without local applications) with a thin or thick client device using a remote connection protocol.

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 endpoint device while the data center is recovered. This allows the users to remain productive.

CRAFTING THE RIGHT STRATEGY

The development of a solid 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 infrastructures 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 configuration required to provide end users access to the tools they need to continue doing work.

Published Applications

Citrix XenApp 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 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 XenApp has the following advantages:

- It's a mature and stable technology.
- It allows for ease of access from any device over any network. (Independent Computing Architecture [ICA] protocol is very efficient over slow networks.)
- It offers reliable security.
- It publishes applications directly to end users.


Citrix XenApp has the following drawbacks:

- Bottlenecks can occur, as all users share the same server operating system and applications.
- If a server becomes unavailable or unstable, the session needs 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.

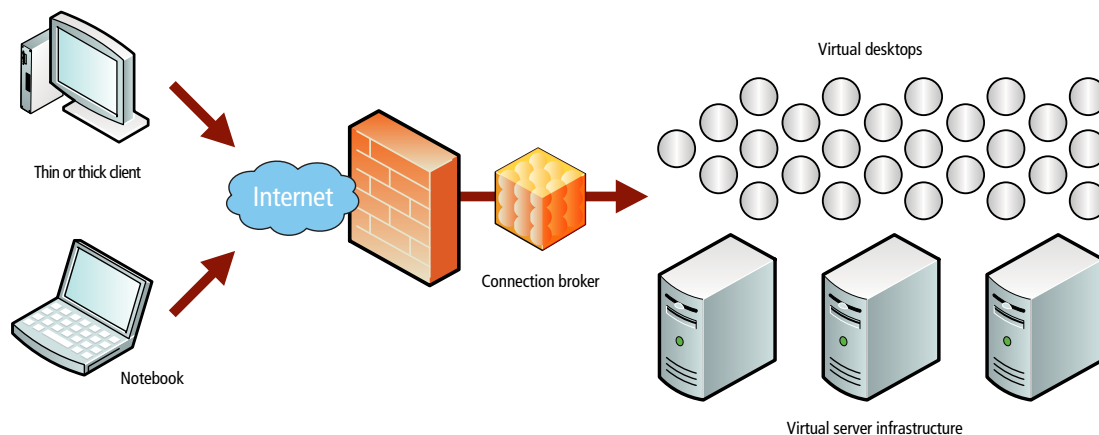
Using Virtual Desktops

Although still a relatively new concept, the virtualization of desktops makes sense for certain environments. Unlike Citrix,

2a. Citrix Application Delivery



2b. Virtual Desktop Access



which shares a slice of a Windows Server operating system and applications to each end user, a virtual desktop is its own unique environment, usually loaded with a Windows Vista/Windows 7/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.

Figure 2b 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.

A virtual desktop solution has the following advantages:

- It offers ease of access from any device over any network.
- It possesses reliable security.
- Desktop sessions are unique and not shared with other users.
- Virtual desktops can be dynamically shuffled 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 Windows 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.

A virtual desktop solution has the following drawbacks:

- It is not a technically mature solution.
- Even though virtual desktops are automatically rebooted after a server failure via another server, this failure 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.
- Only Windows, Linux and Mac web-client links to the connection broker are supported.
- The virtual desktop user experience over the WAN is challenging.
- Virtual desktops require additional server and storage resources to be used or added to the infrastructure.
- 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. ♦

OTHER DISASTER RECOVERY CONSIDERATIONS

PLANNING FOR COMMUNICATIONS, POWER MANAGEMENT AND ENVIRONMENTAL FACTORS

CHAPTER 6:

Unified Communications

Power and COOP

Environmental Concerns

In the midst of a disaster situation, how do you let your staff, contractors and public constituencies know what is going on? How can you ensure that your organization will have the power it needs to get your systems up and running in a timely manner? Finally, what is unique about the local terrain or temperature conditions in your area that may affect your COOP? All are important questions to consider as you prepare your plan.

Given the many communications channels that exist today, people expect to be notified by the means that they have at hand. A complete disaster recovery plan must therefore include the means to ensure universal and near instantaneous communication as to the nature and scope of the threat. And this information must be pushed out via any and all communications modes, including people's eyes and ears.

Your COOP also needs to take into consideration how your organization will restore power. After all, recovery cannot take place without sufficient power. Environmental factors such as heat and humidity are also important. In this chapter, we'll take a closer look at how implementing unified communications, crafting a power management strategy and making provisions for environmental factors can strengthen your COOP.

UNIFIED COMMUNICATIONS

It's quite remarkable how long switched-circuit voice communication has persisted into the Internet age. Voice over Internet Protocol (VoIP) options are plentiful, and they've been around for a while. But they have by no means supplanted

traditional telephone technology.

One reason may be that switched-circuit service has experienced a steep decline in price. Another reason is simply that switched service works; it's a reliable and well-understood technology, backed by a mature support industry.

But maintaining switched-circuit service might not be the best investment when other communications options are gaining maturity and dropping in price.

Operating a system with a PBX and related wiring keeps voice telecommunications (people talking to one another) on a separate infrastructure from the rest of an organization's systems. The phone apparatus is, in many cases, still maintained and operated by a staff separate from the IT department, running up additional costs for an organization.

At first glance, keeping voice as a separate facility makes sense as a disaster recovery strategy. After all, if the data system is knocked out, you'll still have voice communications, right?

But think this through: If the COOP and disaster recovery plans are sufficiently robust to ensure that data systems remain running, then the voice service sharing bandwidth on the data network would be preserved under the same COOP umbrella. When you couple this fact with the cost savings of retiring switched-voice service and the increased functionality of VoIP, it's easy for an organization to make a case for consolidating communications media.

Having fewer media channels makes even more sense when you consider the total communications needs of your organization.

Disaster scenarios mean all stakeholders must be notified, by whatever means are available. Regardless of the diversity of media within the organization, it is not strictly necessary to consolidate in order to reach everyone you need to reach during an emergency.

This capability, known as unified communications, might be enhanced or achieved more efficiently with fewer media, but the use of the right UC products should reduce complexity for both originators and recipients of emergency messages.

UC allows for the fast creation of a single message that can be received regardless of the recipient's device. Organizations with a PBX can, at the least, use off-the-shelf circuit packs to add IP interoperability to their communications system.

Emergency Notification

Some of the most difficult, high-profile events with the potential for disrupting operations take place at public sector facilities.

During any event that would invoke your COOP or disaster recovery plan, notification of people is an absolute requirement. Today, there are multiple ways to reach nearly everyone: public address systems, cell phones, landlines, closed-circuit or broadcast television and computers.

Consider the logistics required to send messages via each of these communications forms if they each entail different media. For example, landlines are hooked to the switched-circuit phone PBX; e-mail travels on the organization's data network; a public-address system might have its own 12-volt wiring; and CCTV has its own coaxial cabling.

Assuming you have the medium-to-medium translation technologies available for telephony, if you need to reach 1,000 or 10,000 mobile devices, it would take a very long time for your PBX to churn through that call list.

In comparison, the case for UC to support your emergency notifications is very strong. Unified communications can provide easy-to-access notifications to all devices. Any number of products can ride on your IP network for purposes of notification, including e-mail, voice calls and voicemail, as well as text messages to PBX-connected phones, cellular phones and VoIP handsets.

Some organizations (if they have a large gathering place such as an auditorium, cafeteria or call center) may want to deploy other output devices that can be accessed by the IP network, such as digital televisions and display panels, and IP loudspeakers that receive their amplification power over Cat-5 cable. (You can even power on other COOP-related devices, such as a 130-decibel mechanical siren, using an IP-addressable AC power switch.)

Thanks to advances in the power and cybersecurity capabilities of 802.11 wireless extensions to LANs, emergency notification is not

limited to areas connected by Cat-5 cable. A benefit of this radio technology is that it can enhance physical and financial security at a time when the organization is stressed by the emergency itself.

For example, valuables such as computer hardware or laboratory gear can be equipped with radio-frequency ID (RFID) tags, either passive or the more powerful active. The tags can let the system know when the gear is on the move. Keep in mind that active RFID tags require batteries and maintenance and plan accordingly.

Don't overlook UC as a way to foster collaborative planning among disbursed managers during activation of the COOP. Whether you video conference or simply exchange text messages, a well-thought-out UC strategy can ensure that everyone has the same information when it matters most.

Ultimately your UC platform, while technically braiding the strands of your communications system together, improves the speed and agility of your organization, especially during (but not limited to) times of emergency.

Notification Best Practices

Here are five tips for developing an emergency notification plan.

- **Maintain a complete inventory of the resources you have available.** Phones and e-mail are obvious, but unified communications systems today can also invoke what were once less-obvious avenues to reach large numbers of people, such as text messaging and even Twitter feeds, if you provide a mechanism for signing people up.
- **Don't just write up a COOP, put it in a binder and place it on a shelf.** A bound, written plan is nothing if you can't flawlessly execute it in a zero-day situation. In other words, test your plan.
- **Plan for the worst case.** When planning and running exercises, take into account the fact that in a real situation you might have to make decisions in a context of chaos, where you can't wait (or hope) for complete information.
- **Include in your disaster recovery plan a clear assignment of duties and chain of authority.** You must also empower line staff, as well as local (and remote) managers, to make decisions based on what they can discern. Headquarters may not always have the clearest picture.
- **Also include the capability of customizing outgoing notifications to the medium and device on which they will land.** Text messages, for example, contain only a few bytes, while video instruction might be several megabytes. But don't scrimp on detail; overly simple warnings to evacuate or take shelter could make a touchy situation worse.

Non-IP Network Devices

A PBX may not be the only non-IP notification device that must be connected for emergency notification under your UC setup. Analog televisions and audio systems are often in place. Numerous interface products are available to connect such output devices to the IP network.

You might have staff who are logged onto LANs, but not using typical operations applications such as e-mail or web browsers. They may be using medical, transactional or academic applications. In cases such as these, your UC platform could enhance communications using a desktop agent that, when receiving a message over the LAN, produces a notification on the user's PC.

A continuity plan should not necessarily require the replacement of all legacy communications devices and systems, or assume they can't be integrated into the data system. A carefully thought-out unified communications framework can extend notification capabilities while preserving existing legacy pieces of the building or campus infrastructure.

POWER AND COOP

When planning for operations continuity, a basic consideration is power. Every disaster scenario includes at least the possibility of an interruption in electricity. Whether it's extreme weather, fire, explosion, flood or earthquake, all are capable of causing a blackout. In the worst-case scenario, failure of the utility grid alone would trigger the activation of your COOP.

Operations recovery is impossible without data; therefore, the data center should be the central point of planning for power. It is also the hub of emergency notification capabilities at most organizations, given that the main channel is the IP network.

Highly critical facilities should not only receive adequate power but also receive it from two separate physical sources (where such service exists). Often, data centers and secure buildings are configured so that two utilities actually enter the premises from different directions. Unfortunately, while many areas have competitive suppliers of electricity, competing vendors often use the same, local electrical grid.

In Chapter 1 of this reference guide, we cover the importance of having an RTO specified in your COOP for your data infrastructure. Better still is to avoid downtime altogether. The only way to ensure this is with an approach to power that ensures that it flows uninterrupted, of course, but also that it provides high-quality electricity in terms of steady voltage and a clean AC waveform that is kind to sensitive electronics.

Types of UPS Systems

Uninterruptible power supply (UPS) systems range from units roughly the size of a loaf of bread serving a single PC to industrial units that can power an entire data center.

Here's what UPS units typically include as basic features:

- The equipment is isolated from the utility source and instead runs off the battery and associated inverter circuit of the UPS. The battery is constantly charged by the utility source.
- At the individual machine or LAN level, UPS systems include software that, within a specified time depending on the battery capability, proceeds with an orderly closing of applications and powering off the PC (or other equipment) when the power goes off.
- A UPS can communicate with the management console via the LAN, so power problems are known beyond the particular location where they occur.

UPS systems aren't all architected the same way. The least expensive are usually standby units, priced at consumer equipment levels, and may have only minimal power conditioning.

In some ways, they are little more than surge protectors with a battery. But they will adequately protect PCs, point-of-sale equipment and similar gear. They don't perform conditioning on the incoming power and are therefore prone to being activated easily during momentary interruptions.

Line-interactive units use specialized transformers on the incoming line to even out voltage fluctuations, so they are more tolerant to brownouts and less trigger-happy to switch over to pure battery.

At the high end are double-conversion (or online) UPS units, which include circuitry on the equipment side that conditions the power flowing to them in such a way as to avoid the switchover to battery, which occurs with other lower-end UPS systems.

They are more expensive than standby UPS units, but the protection they offer sensitive equipment such as VoIP telephones can make them worth the investment. Nothing is more detrimental to a disaster recovery plan than finding out that the backup equipment itself is causing damage to crucial communications.

Data center-class UPS systems are of the online variety, but online technology is also available in smaller devices and might be appropriate for remote locations that have their own servers (and are locally sensitive to a power failure).

Campus or data center UPS units are usually backed up to a generator (typically diesel powered) that can supply AC for hours or even days. Some organizations go as far as to maintain a diesel generator that can power a data center for two weeks.

Testing is an important part of operating a UPS network.

You don't want to risk critical data when doing so, but regularly schedule tests where, in effect, you cut the cable from the utility to see if failover sequences occur as you've planned.

ENVIRONMENTAL CONCERNS

Besides supplying electricity to keep the IT infrastructure running, a COOP must also take into account environmental considerations. Excessive heat or humidity caused by the loss of your heating, ventilation and air-conditioning system (because of a power outage) can quickly translate into a secondary failure if power is supplied only to the computers.

A failure in any part of the HVAC chain of components not caused by a general power failure can also cause trouble. So when planning the COOP in the data center, IT should work with the building facilities staff and its contractors to ensure that the necessary sensors for environmental conditions are installed and architected to notify the IT management console in the event of an out-of-spec situation.

The threat of water supply or wastewater leaks should also be taken into consideration.

Access Controls

Traditionally, most data centers are locked facilities with limited access. But many organizations, with their approach to security

having been shaped by the minicomputer and LAN era of IT, are somewhat more lax in overseeing the location and accessibility of their IT facilities.

When planning for power and other physical attributes, you really should consider access control, plus RFID and motion sensors, to guard against unauthorized access to the data center. Motion detection can also indicate if a window breaks and if wind-blown debris enters critical unsupervised areas.

COOP Conclusion

Ultimately, your COOP must be tailored to the characteristics and needs of your organization. There is no such thing as an off-the-shelf continuity plan.

COOP planners should be careful to deploy finite resources in a way that ensures your organization the most security for your specific needs. It is the mission that drives the shape of the COOP.

Those responsible for operations continuity must clearly understand the central functions and map them to the details such as recovery time objectives, the choice of a unified communications platform, or the server and storage redundancy setup. The technical goal is operations continuity, but the greater purpose is maintaining your constituents' confidence in your organization because of its resiliency and reliability. ♦



GLOSSARY



This glossary should serve as a quick guide to the most important terms. Note that like in many technology-related fields, acronyms are common.

.....

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.

Carrier diversity

Carrier diversity, access to the resources of more than a single service carrier, is essential to a robust COOP. This diversity can be maintained either through a collocation strategy or by utilizing the services of a multiple-carrier facility.

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.

Continuity of operations plan (COOP)

A COOP is a unique plan tailored to the resources and recovery needs of an organization. A COOP should lay out how systems will either be recovered or kept operational should a disaster scenario or some other threat arise.

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 to-the-second recovery.

Data 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.

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.

Document capture and management

Document capture and management is a data process that includes transferring paper documents into digital files, which then go through metadata tagging, allowing for easy search and retrieval in a central data repository.

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.

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 recovery assessment

An operations recovery assessment is a review carried out by a third party that determines the cost of downtime for an organization, important information when crafting a COOP.

Quality of Service (QoS)

QoS refers to network mechanisms that assign different priorities to different applications, users or data flows, or that guarantee a certain level of throughput to the data flow.

Recovery point objective (RPO)

An RPO calculates an organization's tolerance for data loss, and is based on the time period between a scheduled data backup and the data created between backups. A low RPO is desirable.

Recovery time objective (RTO)

An RTO is an organization's desired timeframe for getting a system back up and running following a failure. An RTO should be specified in an organization's COOP and in any SLA.

Redundant array of disks (RAID)

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

Remote monitoring (RMON)

Effective RMON balances network performance with threat detection, integrating cybersecurity, physical security and an organization's network performance parameters into its management strategy.

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.

Service-level agreement (SLA)

An SLA is a contract onsite that specifies the level of service that the service provider will provide, as well as support options, enforcement and penalty provisions for services not provided, a specified level of customer support and what software or hardware will be provided and for what fee.

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.

Storage resource management (SRM)

Prior to rolling out a data deduplication strategy, organizations will want to do an SRM assessment, which determines the files that are being accessed and edited the most. This knowledge will allow for more efficient data deduplication.

Thin clients

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.

WAN acceleration

WAN acceleration is a combination of two technical approaches, application acceleration and WAN optimization, which keep the WAN operating smoothly and efficiently using its bandwidth, a key component to robust network performance.

INDEX



Application virtualization	26	Physical server recovery	9-10
Archiving.....	7	Power.....	4, 29, 31, 32
Carrier services.....	23-24	Recovery point objective (RPO).....	4, 6
Client access	25-28	Recovery time objective (RTO).....	4
Consolidation.....	8-9, 11	Redundant array of independent disks (RAID)	6
Continuous data protection (CDP)	6	Remote monitoring (RMON)	21-22, 24
Cybersecurity	3, 21, 22, 24, 30	Serial advanced technology attachment (SATA).....	6, 7
Data deduplication.....	7	Serial-attached SCSI (SAS)	6
Direct-attached storage (DAS)	5-6	Service-level agreement (SLA).....	22
Disaster recovery.....	3, 5, 6, 11, 21, 25, 27, 29, 30, 31	Snap shooting.....	6
Disaster recovery lifecycle	4	Split-brain data center.....	11
Document capture	5	Storage area network (SAN)	4, 10, 11, 26
Document management	5	Thin clients.....	23, 25-26
Emergency notification	30-31	Unified communications	29-31, 32
Fibre Channel.....	10, 11	Uninterruptible power supply (UPS) systems.....	4, 31
Hierarchical storage management (HSM)	7	Virtual desktops.....	26-27, 28
Linear Tape-Open (LTO)	6-7	Virtual server recovery	10
Load balancing.....	22, 24	Virtualization.....	4, 8-11, 23, 26, 28
Object caching	23	Voice over Internet Protocol (VoIP).....	23, 29, 30, 31
Operations recovery assessment.....	11	WAN acceleration	22-23

Disclaimer

The terms and conditions of product sales are limited to those contained on CDW's website at CDW.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 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 LLC. All other trademarks and registered trademarks are the sole property of their respective owners. CDW and the Circle of Service logo are registered trademarks of CDW LLC. Intel Trademark Acknowledgement: Celeron, Celeron Inside, Centrino, Centrino Inside, Core Inside, Intel, Intel Logo, Intel Atom, Intel Atom Inside, Intel Core, Intel Inside, Intel Inside Logo, Intel 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. Intel's processor ratings are not a measure of system performance. For more information please see www.intel.com/go/rating. AMD Trademark Acknowledgement: AMD, the AMD Arrow, AMD Opteron, AMD Phenom, AMD Athlon, AMD Turion, AMD Sempron, AMD Geode, Cool 'n' Quiet and PowerNow! and combinations thereof are trademarks of Advanced Micro Devices, Inc. This document may not be reproduced or distributed for any reason. Federal law provides for severe and criminal penalties for the unauthorized reproduction and distribution of copyrighted materials. Criminal copyright infringement is investigated by the Federal Bureau of Investigation (FBI) and may constitute a felony with a maximum penalty of up to five (5) years in prison and/or a \$250,000 fine. Title 17 U.S.C. Sections 501 and 506. This reference guide is designed to provide readers with information regarding continuity of operations 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 continuity of operations 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. ©2010 CDW Government LLC. All rights reserved.



CDWG.COM/COOPGUIDE
888.559.4239

ABOUT THE CONTRIBUTORS



« SPENCER CAGLE is a Solutions Architect for CDW, with a focus on networking and load balancing.



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



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



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



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



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



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



« TOM TEMIN is a writer, editor, broadcaster and consultant, with 30 years of B2B media experience related to media and information technology products and services.

CONTINUITY OF OPERATIONS REFERENCE GUIDE

100421 • Flyer 75752

LOOK INSIDE for more information on:

- Maintaining client access
- Keeping servers highly available
- WAN acceleration techniques
- Gaining data storage efficiencies



GSA Contract Holder

