

Don't Flirt With Disaster

Best practices for disaster preparedness.

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

As organizations become increasingly dependent upon rapid access to information — and subsequently less tolerant of failure — an increased focus must be given to disaster preparedness. Organizations that have a thought-out, detailed strategy for disaster preparedness are in a position to quickly go about righting their ship when problems occur.

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.

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. It is possible to virtualize both servers and desktops, which enables centralized storage of data and allows all instances to be replicated into the recovery facility and brought online in minutes, depending upon the design and architecture. But as an alternative to virtual desktops, terminal services (which are used predominantly today) can be used to serve up the applications or desktops to end users in either site.

When an organization moves to virtualization, 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 run 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 vendor during a refresh cycle. Furthermore, the end-user experience is almost identical when accessing applications at the recovery site.

The Impact of a Disaster

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



CDW-G Poll:

What do you feel is the biggest vulnerability in your organization's disaster recovery plan?

- 32% We rely too heavily on tape backups.
- 19% Insufficient funding to make it effective.
- 16% Our plan is not updated frequently.
- 14% We don't have a disaster recovery plan.
- 4% Our data is not centralized to facilitate better backups.
- 15% Other

CDW-G Poll:

What would you say is more of a driver behind your organization's disaster recovery and continuity of operations plans: hardware failures, natural disasters or both?

- 49% Hardware failures
- 33% Hardware failures and natural disasters
- 18% Natural disasters

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

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

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

A plan for disaster preparedness that helps operations recover quickly from a disaster with a minimal loss of data, information and productivity is a necessity.

Disaster Preparedness Lifecycle

Disaster preparedness should be considered more of a lifecycle than a process that can be followed end-to-end to completion. Although the solutions are often technology-centric, they must always be derived from a deep operations acumen, or a detailed and intimate knowledge of the operational needs. There are five key phases to the disaster preparedness lifecycle:

- **Analysis:** This is arguably the most critical component of developing a disaster preparedness plan. It is during the analysis phase that several examinations will be conducted to determine potential impacts,

identify likely threats and develop impact scenarios.

- **Solution design:** During this phase, the goal is to identify the most cost-effective and technically viable disaster preparedness solution.
- **Implementation:** This phase consists solely of the execution of the design elements identified in the solution design phase.
- **Testing and acceptance:** To be certain that disaster preparedness plans meet the needs of the organization, testing is required to assure process and acceptance.
- **Maintenance:** Once a disaster preparedness plan has been established, regular maintenance of the plan helps to ensure viability. The maintenance phase is the ongoing effort to address technical solution needs, recovery solution needs and organizational changes as they impact operational preparedness and a host of other factors.

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

Best Practice Tip #1: Establish Recovery Objectives

Understanding recovery time objectives (RTO) and recovery point objectives (RPO) becomes the foundation on which an effective business continuity solution is built. A few years ago, an organization might have established a two- or three-day window for getting systems back online. Today, the timeframe is usually a few hours. When an organization establishes clear RTO and RPO, it's possible to match hardware and software with business processes and data recovery needs. In some cases, it's important to have systems back online within minutes; in other cases, hours or days will suffice.

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

- **Restoration from tape media to similar hardware:** Although this is the best-case scenario for physical servers, servers must be purchased at the same time to guarantee that the system being recovered is identical.

Cold-Site Recovery:

If an organization has a wider recovery window, cold-site recovery is an option. With cold-site recovery, an organization rents rack space with a service provider and installs its network. In case of a disaster, the organization populates the rack space with the necessary hardware and software for recovery. Forty-eight to 72 hours is the best-case scenario for a complete recovery, and this estimate is based on how long it will take for the recovery of data.

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

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.

different types of backup can ensure that your backup strategy will accomplish what you need. If your data isn't centrally located, returning to your current state after a hardware failure or natural disaster will take much longer. Think about the impact on your organization if you lost critical data. A great deal of valuable data often resides on PCs and not in centralized locations. If you haven't already done so, the first step in disaster planning is to centralize vital information.

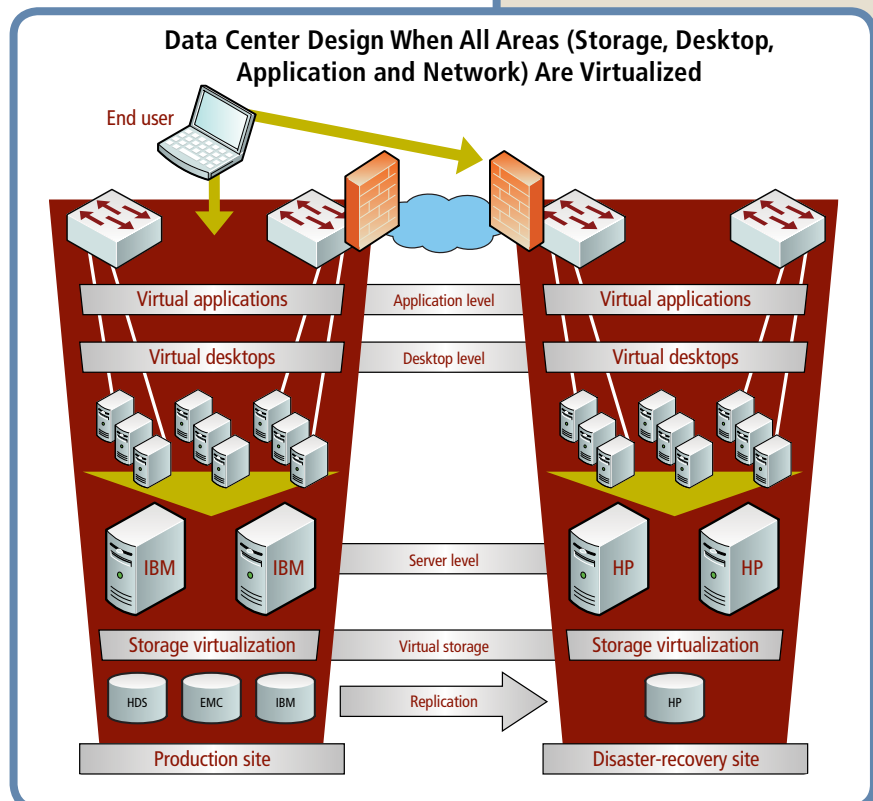
Additionally, an organization's ability to access mission-critical data after a disaster is key to getting it back up and running, and the more centralized the data, the easier the task. Because data storage spans a variety of technologies, it is important to start evaluating the different methods in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

Direct attached storage. DAS, where the storage device is attached directly to the server, offers two varieties of storage options — mirrored disk and the various redundant array of independent disks (RAID) solutions. Both of these options protect from the most common form of disaster: 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

Best Practice Tip #2: Rethink Your Data Storage Options

Your backup strategy is critical to contingency planning — continuity of operations in the event of a disaster (including server failure) and ease of data restoration. To that end, techniques such as centralizing data, establishing a recovery flow chart, determining restore parameters and using



Create a Disaster Flow Chart:

With centralized data stores in place, think about how to store and back up your data and consider an approach that will best suit your organizational needs. You also need to think about where you are geographically, what common threats are experienced in your locale and how your data would survive those threats.

Suppose there was a fire and the sprinklers went off and soaked the servers, or fire overtook the building. Imagine a flood or a gas explosion. How would your organization continue, and what's the probability of your data surviving? These are hard questions and require a lot of "what if" planning.

A basic flow-chart plan is a good way to start your planning. Think of the scenarios you might encounter and how each would affect your IT environment.

external cabinets that can expand to hold many disks and usually have dedicated processors and caches to control the corresponding RAID solutions. The advantages of storage subsystems over traditional primary storage solutions are higher performance and greater expandability. These newer disk subsystems also allow for advanced features such as remote replication, which involves all of the data on one subsystem being copied to another subsystem. The copying can be done over a wide area network (WAN). This subsystem can either be located locally to protect against a subsystem failure, or remotely.

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

Continuous Data Protection. CDP monitors an organization's files, and as a file is changed or "auto saved," a copy of the changed bytes/blocks is replicated to either a local directory or remote location. With this automatic reproduction happening constantly, an organization can have granularity of recovery up to literally the last second. When compared to the traditional previous night backup, this is a boon for organizations needing recovery within a tighter window of time.

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

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

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

Tape is portable (unlike disks, tape can be removed from a drive and taken to another location for recovery or storage), "green" (tape requires no power other than the read/write drives), dense and very fast. Although not widely noted, tape write rates can exceed 130MBps and can be faster than disk for certain types of data recovery. Tape also has a life expectancy that extends greater than 30 years, making it ideal for long-term archives that may or may not be accessed, yet need to be protected.

Today's most widely used tape standard is Linear Tape-Open. LTO is an open-format technology, making it compatible with a variety of products and media. It was developed to combine the advantages of a linear multichannel, bi-directional format with continued enhancements in server technology, data compression, track layout and error correction code to maximize capacity, performance and reliability. Currently, in its fourth generation, the LTO tape consortium releases a new version approximately every 18 to 24 months. The generations are readable two releases back, but writeable only one generation back. For example, an organization that upgrades to LTO-4 will only be able to read media that is LTO-4, LTO-3 and LTO-2, and write to LTO-4 and LTO-3 media.

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

Protecting data can be accomplished in many different ways, including a combination of disk and tape. One approach involves data de-duplication. Data de-duplication, often referred to as "intelligent compression" or "single-instance storage," is a method of reducing storage needs by eliminating redundant data. Only one unique instance of the data is retained on the storage media.

A disk-based solution featuring data de-duplication can allow for virtual compressed quantities of data with smaller amounts of real disk. Depending upon the data, de-duplication can provide compression in the 10x or greater ranges. Compare this to tape, which has greater density (1TB or more per tape), and can provide even greater long-term, cost-effective and efficient storage.

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

Another approach for long-term retention of archival data is a hierarchical storage management (HSM) solution, which involves migrating data from its production location to a lower cost/tier of storage while leaving a “stub” file behind. The stub file allows applications or file searches to see the file in its normal location, but when accessed, recall the file from its lower-cost location. This lower-cost location can be either a slower disk, such as SATA, or even a backup solution, such as tape.

Best Practice Tip #3: Evaluate Virtual Backup Options

Organizations that have begun the migration to a virtual infrastructure have already reaped the benefits of this technology.

Virtualization software allows you to run multiple virtual machines on the same physical host. Each virtual guest runs in a separate environment so errors with the operating system or application will not take down other guests running on the same physical host. Due to 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. Additionally, server virtualization can help heighten security, make restores quicker, improve fault tolerance and reduce server maintenance costs.

Multiple hosts can improve fault tolerance. By combining the power of Windows Server 2003 Distributed File System Replication (DFSR) and virtualization, you can easily move virtual guests to a different host in the event of a server crash. Assuming you perform an image backup of all of the virtual server guests at off-peak times and run differential backups regularly, you can move the virtual server guests to a different host and recover the servers faster than if they were running on dedicated machines.

With virtualization, it is possible to take the server guest files running on a host and move them to a folder that is replicated with DFSR. DFSR will then replicate these guest files to a different host server. DFSR has two features that

make it ideal to replicate virtual server guest files: remote differential compression (RDC) and cross-file replication (CFR). RDC will examine a file and only replicate the changes made to the file to a remote server.

Cross-file replication examines a file and searches for files that are available locally to create the desired file. This lends itself very well to replicating virtual server guest files because most of the information in the file remains static and they have a significant amount of the same information stored in them. After the initial replication, the actual time to replicate the virtual guest files should be short. You can create the virtual server guests on the other hosts, but leave them down and only bring them up if you have a problem with one of the host servers.

For example, in a three-host-server scenario, let's assume you lose host server A, which has three virtual servers running on it. Assuming that you already have the virtual guest files replicated to the other two host servers, you could bring up two guests on server B and one guest on server C. Then you would need to restore the latest differential backup to get the servers as current as possible.

If you design your server host farm with this in mind, make sure to have enough processor, memory and disk capacity to handle the failure of one host server. If you do lose a host server and use this strategy, you should be able to bring up all three guests within a few hours or possibly even faster, depending on how long it takes to restore the last differential backup.

The major drawback to this strategy is the additional storage required to keep “warm” virtual server guest files on the other host servers. This virtual-server configuration also lends itself very well to having a remote disaster recovery warm site or other remote location. Instead of copying the files to a local server, you could use DFSR to replicate the virtual server guest files to a remote server.

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

- **Restoration from tape media to any hardware:** Since 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

Three Tape Tips:

Numerous IT managers report that they distrust tape for backups but continue to turn to tape because of the relatively lower costs, compared with other storage media. If tape is an integral part of your backup plan, consider these three tape tips:

1) Clean Your Heads

Clean the recording heads on tape drives based on manufacturer guidelines, because the oxide-coated plastic wears over time, which causes read and/or write errors.

2) Check Tape Frequently

Never assume that your data is being backed up. On a regular basis, restore a file to make sure that it is actually there, and always keep the latest copy offsite.

3) Store Tape Offsite

Whether on tape or disk, those backup files won't help if disaster strikes and they are kept onsite.

Virtual Desktop Advantages:

- Offers ease of access from any device over any network (ICA protocol used by Citrix is very efficient over slow networks).
- Secure desktop sessions are unique and not shared with other users.
- Virtual desktops can be dynamically moved between different servers in the farm to increase performance based on predetermined thresholds.
- If a VMware Infrastructure (ESX) server becomes unstable, virtual desktops are rebooted on other servers in the farm.
- Licensing is fairly straightforward with the Microsoft Vista Enterprise Centralized Desktop (VECD) model. Note that each desktop also needs Windows Server CALs to access resources on Windows Servers, and VMware Infrastructure licenses for the server virtualization infrastructure.

hypervisor-based, and options are now available to replicate virtual servers directly from the hypervisor. This allows for the continuous or periodic replication from the hypervisor, with no impact on the virtual server.

- **Guest-based replication:** Rather than run replication products at the hypervisor level, replication can also occur from within each guest or virtual server operating system. In this instance, the virtual server can use host-based technologies defined for physical servers, and thus gives organizations more options to choose from.

Virtual Applications

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

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

Virtual Desktops. Server virtualization is now approaching maturity and widespread adoption. As a result, using the same technology to host desktops is quickly becoming the norm.

Virtual desktops involve an end user accessing a desktop operating system (with or without local applications) with a thin or thick client device using remote desktop protocol (RDP). The desktop operating system runs inside a virtual machine sitting on a virtualized architecture, making use of powerful server processing and SAN-based storage.

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

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

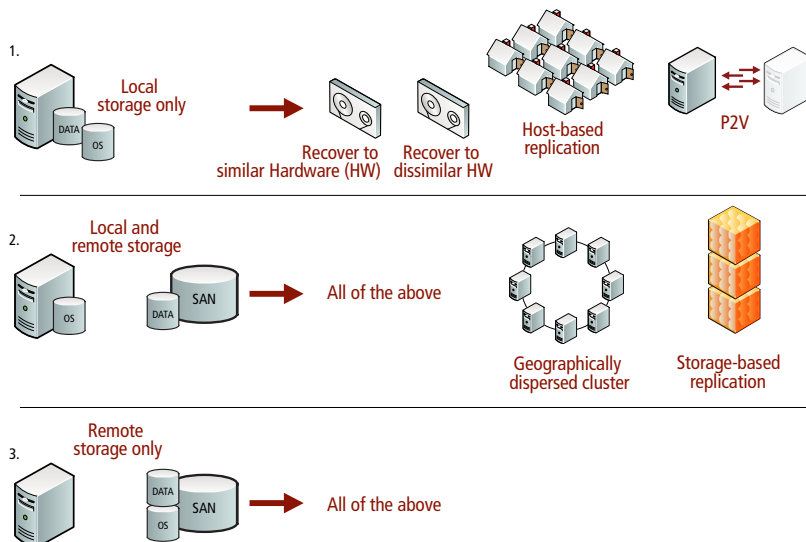
Although a number of solutions for client access are available, it's generally recommended that organizations use published applications or virtual desktops to access applications hosted at the recovery site. This will provide the least amount of configuration required to get end users access to the tools they need to continue doing work.

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

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 prior to the disaster. A single click opens a seamless window to the virtual desktop, and the end-user can begin working immediately.

Recovery Solutions Available for a Physical Server Environment



Best Practice Tip #4: Planned Redundancy

Organizations should start with internally redundant systems with dual processors, redundant disk drives and power supplies. Redundant servers/appliances are the next step to ensuring uptime in case of failure. Finally, in addition to redundant call processors, it is important to geographically separate those systems to protect from a large-scale disaster.

In addition to redundancy, it is important for organizations to have a stringent backup policy in place for configuration and data recovery in case any one system fails. Voice gateway configurations should be kept on file and backed up with revision history policy in mind. Server operating systems and applications should also be backed up for use, if necessary, in a system rebuild.

As with all disaster preparedness plans, each increment of redundancy requires an exponential increase in investment. When designing a plan for high availability and disaster recovery, it is important to understand what disasters your organization wants to be prepared for. Not a single one of these design best practices can provide bulletproof systems.

However, when combined, 99.999 percent reliability can be built into the communications system. When an organization heavily depends on the ability to communicate internally and externally, the price of a solid design is well worth the investment.

Organizations should start with internally redundant systems with dual processors, redundant disk drives and power supplies. Redundant servers/appliances are the next step to ensuring uptime in case of failure. Finally, in addition to redundant call processors, it is important to geographically separate those systems to protect from a large-scale disaster.

In addition to redundancy, it is important for organizations to have a stringent backup policy in place for configuration and data recovery in case any one system fails. Voice gateway configurations should be kept on file and backed up with revision history policy in mind. Server operating systems and applications should also be backed up for use, if necessary, in a system rebuild.

Best Practice Tip #5: Maintaining Power

One of the most basic but overlooked aspects of continuity of operations is maintaining electrical power during a blackout or disaster. Many data centers, computer rooms or wiring closets are prepared for power-related incidents that last a fraction of a second, or a few seconds, and some a few minutes or a few hours with some combination of power protection/conditioning and uninterruptible power supply (UPS), and possibly a generator.

A power supply can be interrupted in little and big ways. Voltage spikes and sags can briefly interrupt power and do serious damage to electronic equipment. Natural disasters can wipe out an organization's power supply completely. Power lines are routinely knocked out in inclement weather.

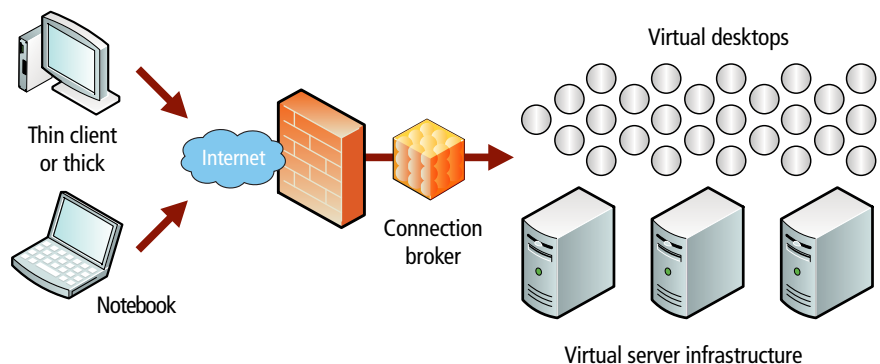
Such power disruptions no longer have to bring operations to a halt. Determining how to keep a steady power supply up and running is an essential element of a disaster preparedness plan. The key to keeping the power running is an uninterruptible power supply.

UPS systems fall into three main categories in roughly ascending order of price and performance.

1. Standby or Offline UPS: UPS powers IT equipment directly from the AC outlet. If a power disturbance occurs, whether it's a blackout, surge or sag, a standby UPS will switch to battery power to protect the technology. When it detects a drop (or a spike) in the electric power coming from the wall outlet, it switches over to its internal battery to power the connected equipment. There's a slight lag in the switchover, but typically not

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.

How end users use virtual desktops following a disaster



The role of a UPS is to protect equipment from brief outages or spikes (most include surge suppressors) and, in the case of an extended power failure, provide sufficient time to complete tasks in progress and properly shut down systems to prevent data loss or corruption.

Not every piece of equipment needs its own UPS, and it's often economical to consolidate protection using one or two larger models with a lot of high-end features, rather than individual UPSs for each device. Some organizations utilize individual UPS devices with dual-conversion online UPS systems to protect all the servers, equipment, data centers and computer rooms that power e-mail and network servers.

long enough to cause a shutdown or data loss. A standby UPS will usually suffice for most desktop computers, printers and other noncritical peripherals, while environments that suffer frequent power dips and surges, or are subject to electrical interference or line "noise" from other equipment will benefit from a line-interactive model.

2. Line Interactive or Automatic Voltage Regulation (AVR) UPS: When an overvoltage or undervoltage occurs, a line interactive UPS corrects the strength of the voltage without the device switching over to battery power. It's a step up from a standby UPS, which automatically switches to battery power for voltage problems. This UPS increases battery life as a result. A line-interactive UPS adds a transformer that can boost or moderate the incoming voltage to the right level without having to switch over to battery power.

3. Double Conversion Online UPS: This UPS is designed so the incoming power flows through the battery, which then powers the IT equipment. If there's an outage, the battery continues to power the equipment until it is drained. With other UPS systems, there is a short interruption in service as the

device switches from incoming power to the battery. This system takes the incoming power, converts the AC power to DC, then reconverts it to AC, so it filters out problems, such as electrical line noise, and provides clean, perfect power to IT equipment.

An online UPS powers the connected equipment from the internal battery all the time, and uses the incoming AC voltage to keep the battery charged. In the event of a power outage, there's no drop in voltage or switchover to the battery. The only change in state is that the battery starts to drain once recharging stops. An online UPS is a sensible investment to protect a telephone system and other gear that an organization relies on for mission-critical operations, such as transaction processing, Internet access or e-mail.

Most online UPSs are "dual conversion" — the AC from the wall is converted to DC for charging the battery, and then back again to AC to power the equipment. Some newer models may be "triple conversion," converting the DC power that goes to the battery to a different DC voltage, which can extend the life of the battery perhaps another year. Because of the constant use, the battery in an online UPS typically needs to be replaced every three to five years.

Though UPS systems can provide up to several hours of battery backup, they are not designed to keep an organization fully operational during a prolonged power failure. That usually requires a backup generator. ♦

Recovery Solutions Available for a Virtual Server Environment

