

Disaster Planning and Recovery Pack

Risk Identification



Prevent Tomorrow's Disaster Today!


While IT managers can't plan for every service interruption, they can take proactive measures to ensure that such a disturbance won't result in a full-blown IT disaster. Of course, proper planning and preparation are key to achieving this goal, and TechRepublic's *Disaster Planning and Recovery Pack* provides the resources you need to do just that.

This valuable set, consisting of the *Administrator's Guide to Disaster Planning and Recovery*, Third Edition book, CD-ROM tool kit, and laminated Disaster Planning & Recovery Checklist and Contact Sheet, will help you:

- Develop a comprehensive disaster contingency strategy
- Assess your organization's vulnerabilities
- Avoid facilities management mistakes

- Implement a formal crisis communications policy
- Test your disaster recovery plan
- Select the right vendors for server colocation and other backup needs

As an IT professional, you need to be 100-percent confident that your IT organization's systems and data are fully protected in the event of a crisis. If you have even a shred of doubt that your disaster recovery plan is as complete as it should be, you need to order TechRepublic's Disaster Planning and Recovery Pack today.



Avoid these facilities management mistakes to keep your computer center online

Jan. 8, 2003

By Change Tech Solutions Inc.

Facilities management in an IT infrastructure is analogous to the props, lighting, scenery, and sound in a major theatrical production. All the elements must work in harmony for users (the audience) to get the most out of their experience. Avoid these five facilities management mistakes and you can be sure that the IT show will go on.

Mistake # 1: Presuming major components of facilities management are all addressed

If you were to ask typical infrastructure managers to name the major components of facilities management, they would likely mention common items such as air conditioning, electrical power, and perhaps fire suppression. Some may also mention smoke detection, uninterruptible power supplies (UPS), and controlled physical access. Few of them would likely include less common entities, such as electrical grounding, vault protection, and static electricity.

Below is a more comprehensive list of the major components of facilities management:

- ▶ Air conditioning
- ▶ Humidity
- ▶ Electrical power
- ▶ Static electricity
- ▶ Electrical grounding
- ▶ Uninterruptible power supply (UPS)
- ▶ Backup UPS batteries
- ▶ Backup generator
- ▶ Water detection
- ▶ Smoke detection
- ▶ Fire suppression
- ▶ Facility monitoring with alarms
- ▶ Earthquake safeguards
- ▶ Safety training
- ▶ Supplier management

- ▶ Controlled physical access
- ▶ Protected vaults
- ▶ Physical location
- ▶ Classified environment

Temperature and humidity levels should be monitored constantly, either electronically or with recording charts, and reviewed once each shift to detect any unusual trends. Electrical power includes continuous supply at the proper voltage, current, and phasing as well as the conditioning of the power. Conditioning purifies the quality of the electricity for greater reliability. It involves filtering out stray magnetic fields that can induce unwanted inductance, doing the same to stray electrical fields that can generate unwanted capacitance, and providing surge suppression to prevent voltage spikes. Static electricity, which affects the operation of sensitive equipment, can build up in conductive materials, such as carpeting, clothing, draperies, and other noninsulating fibers. Antistatic devices can be installed to minimize this condition. Proper grounding is required to eliminate outages and potential human injury due to short circuits. Another element sometimes overlooked is whether UPS batteries are kept fully charged.

Water and smoke detection are common environmental guards in today's data centers, as are fire suppression mechanisms. Facility monitoring systems and their alarms should be visible and audible enough to be seen and heard from almost any area in the computer room, even when noisy pieces of equipment, such as printers, are running at their loudest. Equipment should be anchored and secured to withstand moderate earthquakes. Large mainframes decades ago used to be safely anchored, in part, by the massive plumbing for water-cooled processors and by the huge bus and tag cables that interconnected the various units. In today's era of fiber-optic cables, air-cooled processors, and smaller boxes designed for

nonraised flooring, this built-in anchoring of equipment is no longer as prevalent.

You should include emergency preparedness for earthquakes and other natural or man-made disasters as a basic part of general safety training for all personnel working inside a data center. They should be knowledgeable about emergency powering off, evacuation procedures, first-aid assistance, and emergency telephone numbers. Training data-center suppliers in these matters is also recommended.

Most data centers have acceptable methods of controlling physical access to their machine rooms, but this is not always the case for vaults or rooms that store sensitive documents, check stock, or tapes. The physical location of a data center can also be problematic. A basement level may be safe and secure from the outside, but it might also be exposed to water leaks and evacuation obstacles, particularly in older buildings. Locating a data center along outside walls of a building can sometimes contribute to sabotage from the outside. Classified environments almost always require data centers to be located as far away from outside walls as possible to safeguard them from outside physical forces, such as bombs or projectiles, as well as from electronic sensing devices.

In fairness to infrastructure managers and operations personnel, several of these components may be under the management of the facilities department for which no one in IT would have direct responsibility. But even in this case, infrastructure personnel and operations managers would normally want and need to know who to go to in the facilities department for specific types of environmental issues.

Mistake # 2: Believing that the roles and responsibilities of key individuals are clearly defined and understood

It's important to identify the key individuals who participate in facilities management, define their roles and responsibilities, and effectively communicate that information. Clearly defining the areas of responsibility and, more important, the degree of authority between these IT and facilities usually is the difference between resolving a facilities prob-

lem in a data center quickly and efficiently versus dragging out the resolution amid chaos, miscommunication, and strained relationships.

For example, suppose a power distribution unit feeding a critical server fails. A computer operations supervisor would likely call in electricians from the facilities department to investigate the problem. Their analysis may find that the unit needs to be replaced and that a new unit will take days to procure, install, and make operational. Facilities and IT need to brainstorm alternative solutions to determine each option's costs, time, resources, practicality, and long-term impact, and all this activity needs to occur in a short amount of time—usually less than an hour. This is no time to debate who has responsibility and authority for the final decisions. That needs to have been determined well in advance. Working with clearly defined roles and responsibilities shortens the time of the outage to the clients, lessens the chaos, and reduces the effort toward a satisfactory resolution.

The lines of authority between an IT infrastructure and its facilities department will vary from shop to shop depending on size, platforms, degree of outsourcing, and other factors. The key point here is to ensure that the two departments clearly agree upon, communicate to their staffs, and ensure compliance with these boundaries.

The mistake arises when one or more of these three parts—identification, definition, and communication—is believed to have occurred but, in actuality, has not. The mistake becomes fatal when the data center and the facilities department each believe the other is following up on an incident, resulting in major extended outages because neither group took action.

Mistake # 3: Thinking that the owner of the IT facilities management process is adequately qualified and trained

One person should be assigned the role and responsibility of facilities management process owner. Our next mistake occurs when infrastructure managers presume that this person is instinctively qualified and trained for the

assignment. The mistake becomes fatal if a major incident, such as physical disaster, pushes the person beyond the skill level he or she is able to handle.

The owner of the facilities management process almost always resides in the computer operations department. There are rare exceptions—small shops or those with unique outsourcing arrangements—in which the facilities management process owner is part of the facilities department and matrixed back to IT or is part of the IT executive staff. In any event, the selection of the person assigned the responsibility for a stable physical operating environment is an important decision.

Mistake # 4: Relying solely on environmental monitoring to eliminate supplemental analysis


IT facility managers sometimes believe that the more they automate the monitoring of their data centers and server room, the less effort they will need to expend to ensure stability. This is a natural conclusion but a flawed idea. A patient in intensive care hooked up to dozens of monitoring devices still requires doctors and nurses to periodically check the patient's vital signs. This serves to record the current condition of the patient and to verify the proper operation of the equipment. Similarly, IT facilities managers need to evaluate, rather than simply monitor, the current state of their data center's physical environment. The mistake of relying solely on monitoring systems can become fatal if sensors, alarms, or other types of annunciator systems fail during a major disaster.

There are a number of sources of information that can assist data center managers in evaluating the current state of their physical environment. Outage logs normally associated with availability reports should point to the frequency and duration of service interruptions caused by facilities. If the problem management system in use includes a robust database, it should be easy to analyze trouble tickets caused by facilities issues and to highlight trends, repeat incidents, and root causes.

Mistake # 5: Ignoring the nurturing of human relationships

Data center managers sometimes ignore the value of developing strong, personal relationships with key individuals outside of their own departments. These external individuals will vary from shop to shop, but usually include the managers and foremen of the company's facilities department and representatives of government inspecting agencies. IT managers responsible for data center facilities do not always view relationship building as an integral part of their jobs, focusing instead on the more technical, nonhuman aspects of their work.

This mistake can become fatal if data center facility managers alienate these key external individuals to such an extent that they delay critical physical expansions or upgrades required to sustain a stable operating environment.

Understanding and, more importantly, avoiding these five mistakes can help you sustain the continuous online services of a computer center. This, in turn, can prevent your online performance from coming to a screeching—and unnecessary—halt. 



Prevent Tomorrow's Disaster Today!

As an IT professional, you need to be 100-percent confident that your IT organization's systems and data are fully protected in the event of a crisis. If you have even a shred of doubt that your disaster recovery plan is as complete as it should be, you need to order TechRepublic's Disaster Planning and Recovery Pack today.

Be aware of utilities' impact on disaster recovery plans

March 12, 2004

By Mike Talon

Most disaster recovery plans include something to do with replication systems. These systems can offer the ability to immediately resume operations at another site or, at the very least, create data replicas for eventual restoration.

However, while replication systems are an almost mandatory portion of large DR plans, they can pose some interesting issues when you run utilities on your production systems.

One of the most common culprits is defragmentation software. Windows systems especially benefit from these tools, and they're also becoming a common utility in production data centers.

Unfortunately, replication tools don't always play nicely with these utilities. When the defrag kicks off and starts reading, writing, and moving data all over the disk, replication tools might see this data change and replicate it to the DR systems, causing massive traffic spikes.

The good news is that this is typically restricted to the world of hardware-level replication, so those using host-based products generally don't see the massive push of data across the WAN. That's due to the fact that most defrag systems work at the disk level—not the file-system level—and nonhardware-based replication tools therefore ignore them.

If you use a hardware-based replication tool, you'll have to take this traffic pattern into account when using defrag utilities. Unfortunately, you can't simply shut off replication during the defrag, which would leave you with horribly corrupted files at the DR site.

However, these hardware-based systems typically require a great deal of bandwidth to work properly, so you should have enough pipe to handle the load of a defrag, but you'll experience a pretty bad performance hit when the utility begins. Since defrags generally take place in the middle of the night, proper timing can keep end users from experiencing the hit, but you should still be aware of its occurrence.

Another usual suspect is antivirus software. Again, this is a larger issue in the Windows world than elsewhere, but it still affects all platforms in some way.


Production antivirus systems typically don't cause major issues, and nearly all replication tools can work with antivirus tools without incident. But issues arise on the DR systems, where constant scanning and rescanning of files (replicated either block by block or byte by byte) creates an enormous load on the antivirus systems.

Because of this load, replication streams can slow to a crawl and could force remirroring or other uncomfortable consequences. The fix for this issue is relatively easy: Shut off antivirus scanning in whole or in part on your DR systems—and only the DR systems.

If you're scanning everything at the production site, you don't need to scan that data again at the DR site. You can easily flip on the antivirus services during a fail-over procedure to allow for scanning to occur when users are directly accessing the files.

If corporate regulations don't allow you to completely shut off antivirus systems (or the thought of doing so gives you a case of the screaming heebie-jeebies), most software allows you to exclude directories and files. Simply ignore the directories and files you're replicating, and scan everything else. It may sound risky to shut off antivirus tools, but remember that excluding this data, which you're scanning at the main facility, while scanning everything else shouldn't leave your organization vulnerable to virus attacks.

Of course, these are only two examples of utilities that can interfere with DR planning and maintenance. There are a myriad of tools and utilities that can impact how your DR plan functions, especially when replication is part of the picture.

None of these headaches warrants eliminating DR tools from your toolkit. But each utility requires an awareness of its functionality and how it can impact your overall plan. 

A holistic approach to vulnerability assessment

April 21, 2004

By Ruby Bayan

There was a time when fending off hackers and viruses was regarded as mere exception processing. Today, keeping the network safe from rapidly evolving malicious intent is considered critical procedure with top-level priority. Firewalls, intrusion detection devices, antivirus applications, and vulnerability assessment tools are now vital munitions in every CIO's security arsenal.

Unfortunately, in spite of diligent efforts to thwart the growing number of attack tools and techniques, companies continue to lose millions of dollars to security incidents caused by product and system vulnerabilities.

Given the vulnerability assessment options available—manually implemented toolsets, consultant penetration testing services, and automated Web-based assessment—how should the CIO map out the best solution for his enterprise? The experts we interviewed suggested a holistic approach to vulnerability assessment.

Start with risk assessment

According to Stan Quintana, Vice President of Managed Security Services of AT&T, any kind of corporate assessment should be one that is risk-based and quantifiable. "The intent of a risk-based security assessment is to isolate corporate assets that generate the highest value to a corporation and, at the same time, present the highest potential threats and vulnerabilities associated with the assets," he said.

Quintana added that by understanding the company's risk profile—"value X threats X vulnerabilities"—it could more readily identify areas in which to invest its precious funds.

"And a CIO will likely have a request from the company's auditors for some form of risk model," Quintana explained. "An assessment will help to pinpoint the areas where security investigation is needed and where it's likely to be fraught with consequences. Such an assessment will also point out whether a business continuity/disaster recovery plan is needed. This, then, forms a holistic security architecture approach."

Understand the compliance obligations of your organization

"When purchasing any security solution, it is important to understand the regulatory obligations of your organization, as this will dictate your specific requirements," said Andrew Maguire, director of product marketing at nCircle, a provider of appliance-based vulnerability management solutions.

Regulatory compliance used to be a non-critical issue; it was enough to meet the minimum requirements, if the rules were pertinent at all. But as technology matured, so did the regulations, along with strict enforcement.

Maguire's example was the Sarbanes-Oxley Act, which holds company officers accountable for the enforcement of "best practice" in audit and compliance. "This has been done in an effort to increase the overall security of sensitive data. Compromise of that protected data and failure to prove best practices can lead to severe fines for the company and potential jail time for the company officers responsible."

George Lekatis, general manager of George Lekatis Inc., a firm that specializes in network security, computer forensics, and litigation, stressed the same point: "The CIO must customize the network vulnerability assessment according to his company's technical and legal needs."

He mentioned the need for compliance with regulations such as the Gramm-Leach-Bliley Act, the Health Insurance Portability and Accountability Act, and the ISO-17799. He also advised understanding the differences in the data protection directives in the USA and Europe.

Review security testing methodologies

Lekatis also said that in evaluating vulnerability assessment solutions, CIOs should look into tried and tested methods of security testing, such as the Open Source Security Testing Methodology Manual (OSSTMM), an open

standard that claims to be the most widely used, peer-reviewed, comprehensive security testing methodology.

He suggested checking out the National Institute of Standards and Technology, which provides access to other methodologies like the Computer Security Resource Center security testing systems and standards, and the Common Criteria for IT Security Evaluation (ISO International Standard 15408).

Study the pros and cons of vulnerability assessment tools

The next major step in determining the best vulnerability assessment solution is to scrutinize what's available. Quintana recommended finding answers to important questions.

When considering manual tools—whether commercial or open source—ask if the staff using them have the skills and time to wield them effectively and safely, Quintana advised. “What consequences are likely, and what mitigation plan is in place, if the tools are accidentally or intentionally misused?”

As to using consultants, Quintana said an important question would be whether they are skilled in *your* enterprise's chosen infrastructure. Also, “Are they ethical and trustworthy? Do they offer guarantees? Are they too expensive or too cheap?”

Quintana added, “For automated, Web-based assessments, the CIO should be aware that these are always less thorough than either of the other two options. They are, however, useful when a company has a significant outward-facing infrastructure and needs to ‘keep an eye’ on it. I would use the same criteria on automated assessment providers as on professional services firms: reliability, ethics, and skill.”

But would a specific tool be the one best solution for a particular enterprise? Not necessarily.

Consider an ensemble of tools

“Every solution is different for every customer,” said Quintana. “For a hypothetical *average customer*, I would recommend a program of security training for the existing staff, a detailed penetration test from an outside consulting firm, with a follow-up from a different firm six months later, and perhaps regular

scanning of the company's Web and server infrastructure that faces the Internet.”

Quintana said this might be the best and most effective use of company dollars. “Training will educate the staff to be more proactive. The external penetration test's first readout will baseline the current status and give a roadmap for improvement. The second iteration will validate the work done to fix problems found the first time around.” Regular scanning will keep a “weather eye on the Web-and-server farm,” he said.

“For my money, I would make external scanning more important if the customer's Web infrastructure is not running highly secure software packages,” Quintana added.


Explore vulnerability management

Maguire said that although “vulnerability assessment bolsters the last line of defense by helping make the target immune to attack,” it is important that “enterprises take into account that they need to do more than just identify vulnerabilities.”

“Vulnerability management technology extends the capabilities of vulnerability assessment by providing a framework for addressing vulnerabilities,” said Maguire.

“To take a truly proactive approach to network security, the IT team needs to take action to eliminate threats. In contrast with vulnerability assessment, vulnerability management provides for a structured security program where budget and resource planning can be executed based on measuring the company's level of exposure.”

Maguire advised that if investing in vulnerability management, CIOs should look for solutions that are scalable and will meet the complex requirements of their organization. That is, “offer seamless deployment across multiple locations.”

He also said a good vulnerability management solution has specific management features that include remote device configuration, role-based access, auditing, reporting and remediation management. More important, it provides for ease of management “so IT staff can manage a security program instead of it managing them.” 

Rank the importance of applications and systems

Dec. 8, 2003

By Mike Talon

Before you develop a disaster recovery plan, and before you begin talking to your company's business continuity and human resources departments and vendors, you must take one very important step. You must first determine just what it is you're trying to protect—and how you'll define protection.

Aside from determining the differing levels of DR (such as data vaulting, high availability, etc.), you need to find out basic facts, such as service-level agreements, before you can begin proper DR planning. To begin this phase of your planning process, you need to talk to the owners of the individual data systems to determine two important factors for each system.

The first factor is the recovery point objective (RPO), which is the measurement of how much data the company can conceivably lose to a disaster of any level that won't significantly impact business. RPO is a tricky factor to judge; nearly everyone in charge of a data system will first insist the RPO is zero bytes of data loss.

However, once people find out the prospective cost of the type of protection systems required for zero bytes of data loss, most managers quickly come up with more realistic RPO estimates. For example, since any Windows-based application must be able to withstand the loss of everything in the page writer cache, the RPO of a Windows-based system could be surprisingly large.


This isn't to say the company won't have applications that do require an RPO of zero bytes lost. For example, Linux and UNIX-based systems—especially financial applications—may require this level of protection.

Hardware-based solutions can offer this level of recovery, but they're exceptionally expensive. Therefore, for budgetary reasons, you must thoroughly justify the cost before deciding to implement these solutions.

The second factor you must define is the recovery time objective (RTO), which is the amount of time any system can remain offline without significantly impacting business. Once again, budgetary concerns will primarily determine your decisions. While you can have instantaneous failover for many systems, managers must make a very good case to cover the cost.

Most applications can easily withstand larger RTOs than you might imagine, mainly due to the complexities of end users. Except for lower-level emergencies, your end users will most likely be unable to access the data systems from their original terminals. This means you'll have quite a bit of time to fail over the systems while end users restore connectivity or move to a new location.

After determining the RTO and RPO of individual data systems, you can begin the process of determining which tools you'll need to implement to meet these needs. You'll probably need to use more than one type of DR solution, both hardware-based and software-based in addition to traditional tape backup resources.

Performing an analysis beforehand allows you to determine which systems require which solutions so you don't waste budget funds on systems that don't require more expensive solutions. 

Are you ready for a small catastrophe?

Jan. 31, 2002

By Bob Artner

For obvious reasons, IT organizations around the world have spent the last few months creating or reviewing their disaster recovery plans. They have reviewed their physical and virtual security, contacted backup hosting and data centers, and worked on contingencies for immediate office space, in case the entire organization had to move to a new location.

That's the kind of world we live in today, and an IT manager would be foolhardy not to plan for such events.

However, what if you were confronted by a small disaster, something that's more than an inconvenience but less than a total failure? Is there a gap in your planning? Do your procedures allow for only two possibilities: "business as usual" and "start from scratch"? In this article, I'm going to discuss a small disaster that happened recently at TechRepublic and the questions it raises about the planning and flexibility needed by technical managers.

It all started with some idiot on a tractor...

Last Friday morning, we lost our electricity at TechRepublic's main offices for almost an hour. Apparently, somebody with a backhoe cut a power cable and put a couple of city blocks in the dark. In the engineering building, the UPS in the server room kicked on, and the operations guys started shutting down the various servers. Unfortunately, before they could completely shut down the Exchange server, the UPS died. When power was restored and the techs brought the Exchange box back up, they discovered database corruption and attempted to repair the errors.

When the lights go out, know who has the flashlight

Instead of focusing on the specifics of what went wrong, and how our operations folks fixed it, I want to concentrate on what this

kind of minor crisis teaches us about the need for planning and flexibility. After all, the power was off for just an hour. None of our offices was actually damaged. In the minds of most of our employees, it just meant a wasted morning—until e-mail stayed offline for the rest of the day. As most of you know, e-mail is the mission-critical application for many organizations, and TechRepublic is no exception.

It ended up being a real problem, and yet it wasn't serious enough to trigger the company's formal disaster recovery plans. What do we call it: a small disaster? A minor train wreck? Whatever the term, IT managers have to be able to respond.

How to prepare for a "minor disaster"

Looking at what our operations team had to go through, here are some factors for you to consider when fighting through your own "small emergencies":

Contact information

While e-mail is vital to any IT organization, the telephone also comes in mighty handy. Take a minute and think of everyone our operations people had to contact:

- ▶ The landlord
- ▶ The electric utility
- ▶ The hardware manufacturer
- ▶ The tape backup support department
- ▶ Microsoft technical support
- ▶ Plus anyone else they could think of who'd encountered a similar problem in the past

Faced with the same kind of crisis, could you easily come up with all the contact information necessary? Consider a lightning strike that overpowers your UPS and fries some of your boxes. If your organization is like many others, you not only have to support a variety of hardware manufacturers, but you also have to

5.5.2 Law of intermediate crisis management

The only small problem is one that happens to the other guy.

consider a mixture of equipment that you own outright and other types of equipment that you lease. Obviously you have to treat the latter differently.

Internal client needs

Here is a tough test for you. Go into your server room and point to a box at random. Now assume that the box will be offline for eight hours. Do you know who in your organization would be affected by such an outage? Probably not. Do any of your people? You'd better hope so. Even in midsize companies, most internal networks are too complicated for any one person to understand the dependencies of each piece of hardware or every software application.


Keeping everything else running

While it would make life much simpler to be able to focus all your efforts on the current problem and let everything else fall by the wayside, that's not practical. You've got to be able to juggle a number of balls.

Cross-training

When the Exchange server goes down, you'd like to have more than one Exchange expert on the premises. All IT departments are stretched thin these days, but it's worth the effort to cross-train staff, particularly on mission-critical equipment and applications.

Keep calm

If a flood put your data center under water, chances are that everyone would understand that it could take some time to get an alternate data center up and running. People would cut you some slack. When the Exchange server is down, on the other hand, and everything else is working normally, people are more inclined to pick up the phone and holler, "I need Outlook—now!" That's the paradox: Small crises breed impatience. While you might not be able to soothe your internal clients, try to stay calm yourself. 

Prevent Tomorrow's Disaster Today!


While IT managers can't plan for every service interruption, they can take proactive measures to ensure that such a disturbance won't result in a full-blown IT disaster. Of course, proper planning and preparation are key to achieving this goal, and TechRepublic's *Disaster Planning and Recovery Pack* provides the resources you need to do just that.

This valuable set, consisting of the *Administrator's Guide to Disaster Planning and Recovery*, Third Edition book, CD-ROM tool kit, and laminated Disaster Planning & Recovery Checklist and Contact Sheet, will help you:

- Develop a comprehensive disaster contingency strategy
- Assess your organization's vulnerabilities
- Avoid facilities management mistakes

- Implement a formal crisis communications policy
- Test your disaster recovery plan
- Select the right vendors for server colocation and other backup needs

As an IT professional, you need to be 100-percent confident that your IT organization's systems and data are fully protected in the event of a crisis. If you have even a shred of doubt that your disaster recovery plan is as complete as it should be, you need to order TechRepublic's Disaster Planning and Recovery Pack today.



Guidelines for a disaster plan

Dec. 4, 2001

By Laura Taylor

Disaster recovery is becoming increasingly important for businesses aware of the threat of both man-made and natural disasters. Having a disaster recovery plan will not only protect your organization's essential data from destruction, it will help you refine your business processes and enable your business to recover its operations in the event of a disaster. Though each organization has unique knowledge and assets to maintain, general principles can be applied to disaster recovery. This set of planning guidelines can assist your organization in moving forward with an IT disaster recovery project.

Accountability and endorsement

A key factor in the success of your disaster recovery plan will be holding someone on the executive management team accountable. It could be the CIO, CTO, COO, or, if your company is small, the IT director. Whether this person manages the disaster recovery preparations or delegates them to someone else, it will be necessary for the entire organization to know that the disaster recovery preparations are deemed essential by executive management in order to obtain the cooperation you'll need from all the staff involved. Without endorsement from higher management, collecting all the information you'll need to make the disaster recovery project a success will be much more difficult. Even if the disaster recovery project is managed by someone who has had the task delegated, upper management needs to convey to the entire organization the importance and essentiality of the project.

Identify and organize your data

One of the first steps in putting together a disaster recovery plan is to identify your mission-critical data. Doing so will help you understand what data you need to back up for off-site storage. It will also prompt you to document why you need this data and plan how to put it back where it belongs in the event of a recovery operation.

Next, instruct your users to assist you in organizing the data in intuitive directories on a central server. Even if you plan to just back up everything, knowing which files are where is a key part of the recovery process. If, for example, when disaster strikes you have customer data spread all over your network on different users' hard drives, finding the data will not be easy. Restoring all the data from backup media is only half the battle. Once the data is restored, you don't want to be walking around the office saying, "Does anyone know where we keep the XYZ contract?" The data must be organized before you back it up.

Some data types that you should take into consideration for organization on a central repository are as follows:

- ▶ **Key customer files:** contracts, agreements, contact information, proposals
- ▶ **User login data:** profiles, UNIX .dot files, Config.sys files, Autoexec.bat files
- ▶ **Network infrastructure files:** DNS, WINS, DHCP, router tables
- ▶ **User directories**
- ▶ **Application data:** databases, Web site files
- ▶ **Security configuration files:** ACLs, firewall rules, IDS configuration files, UNIX password/shadow files, Microsoft Windows SAM database, VPN configuration files, RADIUS configuration files
- ▶ **Messaging files:** key configuration files, user mailboxes, system accounts
- ▶ **Engineer files:** source code, release engineering code
- ▶ **Financial and company files:** general ledger, insurance policies, accounts payable and accounts receivable files, incorporation registration, employee resource planning (ERP) data
- ▶ **License files for applications**

Asset inventory

Aside from the data itself, your company needs to have an up-to-date hardware and software asset inventory list on hand at all times. The hardware list should include the equipment make, model, and serial number, and a description of what each particular piece of equipment is being used for. The software inventory should be similar, with the vendor name, version number, patch number, license information, and what the software is being used for. The information for each piece of equipment and software on the list should be mapped to the corresponding devices on the company network map. Be sure to include all cables and connectors, as well as peripheral devices such as printers, fax machines, and scanners.

You might want to submit the asset inventory list to your insurance company once a year.

Restoration and recovery procedures

Imagine that a disaster has occurred. You have the data, now what should you do with it? If you don't have any restoration and recovery procedures, your data won't be nearly as useful to you. With the data in hand, you need to be able to re-create your entire business from brand-new systems. You're going to need procedures for rebuilding systems and networks. System recovery and restoration procedures are typically best written by the people that currently administer and maintain the systems. Each system should have recovery procedures that indicate which versions of software and patches should be installed on which types of hardware platforms. It's also important to indicate which configuration files should be restored into which directories. A good procedure will include low-level file execution instructions, such as what commands to type and in what order to type them.

Document decision-making processes

Recovering your data, systems, and networks is one thing, but when you lose staff, recovering the knowledge they held is quite different. You will never be able to recover that knowledge completely. However, you can mitigate this

loss by documenting decision-making processes in flowcharts. To do this, have each of your staff identify decisions that they make and then create flowcharts for their thought processes. Sample decisions could be:

- ▶ How much do you charge for a new service?
- ▶ How do you know if taking on a particular new project is worth the return?
- ▶ How do you evaluate new business?
- ▶ How do you decide whom you should partner with?
- ▶ How do you decide who your sales prospects are?
- ▶ How do you decide who your suppliers are?
- ▶ When a call comes in to your help desk, how does it get routed?
- ▶ What are your QA procedures for your product?

It's impossible to document every decision your staff is capable of making. To get started, don't ask your staff to document every possible decision-making scenario. Ask them to document the three most important decision-making processes that they use on a consistent basis. You can add new processes to your disaster recovery plan in the future, and you may want to have employees write three new decision-making flowcharts each year at the time of their annual reviews.

Backups are key

As an IT or network administrator, you need to bring all your key data, processes, and procedures together through a backup system that is reliable and easy to replicate. Your IT director's most important job is to ensure that all systems are being backed up on a reliable schedule. This process, though it seems obvious, is often not realized. Assigning backup responsibilities to an administrator is not enough. The IT department needs to have a written schedule that describes which systems get backed up when and whether the backups are full or incremental. You also need to have the backup process fully documented. Finally, test your backup process to make sure it works. Can you restore lost databases? Can you restore lost source code? Can you restore key system files?

Finally, you need to store your backup media off-site, preferably in a location at least 50 miles from your present office. Numerous off-site storage vendors offer safe media storage. Iron Mountain is one example. Even if you're using an off-site storage vendor, it doesn't hurt to send your weekly backup media to another one of your field offices, if you have one.

Disaster strikes

Let's say for a moment that the worst occurs and your business is devastated by a disaster, to the point where you need to rebuild your business from scratch. Here are some of the key steps you should take to recover your operations:

1. Notify your insurance company immediately.
2. Identify a recovery site where you will bring your business back up.
3. Obtain your asset inventory list and reorder all lost items.
4. Distribute a network map and asset inventory list to your recovery team.
5. As the new hardware comes in, have your recovery team connect the pieces.
6. Restore your network infrastructure servers first (DNS, routers, etc.).
7. Restore your application servers second.
8. Restore your user data third.
9. Perform any necessary configuration tweaks according to your guidelines.
10. Test all applications for functionality.
11. Test all user logins.
12. Put a notice on your Web site stating that your business was affected by a disaster.


Summary recommendations

It's likely that in the event of a real disaster, not everything will be recoverable. Your goal should be to recover enough data, processes, and procedures so that your business can be up and running as quickly as possible, once you're in a new office.

Testing your plan is key to ensuring its success. A good way to test your plan is in a lab setting. With uninstalled systems that aren't connected to the network, see how fast you can install your systems, configure them, and restore essential data. The best test is to use a recovery staff other than the everyday staff that uses and administers the systems. By using staff that aren't familiar with everyday usage of your systems and applications, you'll uncover deficiencies in the processes and procedures you've documented. Time your recovery scenario and see if you can improve the time it takes for recovery each time you hold a practice drill.

Conclusion

A disaster recovery plan is essential to your company's long-term success. Even if you never have to use the plan, the process of putting it together will, by its very nature, increase the security of your assets and improve your overall business efficiency. The preparation of a disaster recovery plan will teach you what data is important and will necessitate that you understand how your business works from a decision-making standpoint. Disaster recovery can be more easily achieved if you follow this simple outline:

- ▶ Hold someone accountable for disaster recovery.
- ▶ Identify mission-critical data.
- ▶ Organize data on a central repository.
- ▶ Create procedures for recovering mission-critical servers.
- ▶ Create knowledge-based, decision-making flowcharts.
- ▶ Back up your data on a regular schedule.
- ▶ Store your data off-site.
- ▶ Test your recovery plan. 

Include messaging systems in your DR efforts

Dec. 8, 2003

By Mike Talon

Once considered a peripheral function of the enterprise environment, messaging systems are quickly becoming a vital part of most modern organizations. Due to the increasing reliance on messaging systems, the systems themselves have become integral components of a disaster recovery plan.

In addition, many other programs are now relying on messaging systems in order to properly function. Common examples include voice mail systems, customer relationship management systems, and fax systems.

No matter which messaging platform your enterprise currently uses, or what other applications rely on it, messaging has become an essential part of the enterprise. That means it's also an application that you can't afford not to focus on when developing your organization's DR plan.

Messaging systems come in many versions and flavors. The most popular are currently Microsoft Exchange and Lotus Notes. Both of these systems run on the Windows platform, but there are many messaging systems that run on UNIX and Linux systems as well.

Regardless of what platform you choose to run your messaging system on, its failure could cause the failure of other applications that rely on it. Therefore, it becomes even more necessary to ensure that the messaging system does not fail. How do you secure messaging systems that require dedicated connections to the Internet and generally have end users connected at all times?

Of course, standard DR methodologies apply to messaging systems. Redundancy of hardware will allow you to survive single server failure, and redundancy of data via replication or another methodology of moving data from one place to another can help you survive hardware-related tragedies.


But many of these methodologies don't account for the fact that end users remain con-

stantly connected to the systems. In this somewhat unique situation, end-user client software expects to find a steady state connection, which presents a rather difficult problem. How do you successfully fail over these operations from one hardware platform to another—or worse yet, from one data center to another?

Technologies exist to assist you in this plight. Whether you're using software based on Windows, UNIX, or Linux, some form of clustering should be available to you. Clustering solutions can allow you to fail over between physical hardware devices without losing end-user connectivity for more than a few seconds. Although you do lose the connection state for a brief period of time, most applications designed to work with these clustered messaging systems are more than capable of dealing with this particular kind of outage.

While clustering gives you the opportunity to endure a singular hardware failure at a singular data center, it usually doesn't take care of the situation when an entire data center goes offline. In these instances, it's necessary to use a WAN-based DR solution, which may or may not offer high availability.

This is not to say that there are no vendors that provide HA solutions across all WANs. But make sure you ask your DR vendors about the full capabilities of their products.

Procuring funding for DR and HA specifically for messaging systems may initially be an uphill battle. Many nontechnically trained personnel may not recognize the vital nature of the messaging system. They may simply see it as e-mail, in which case you might have to drill home the fact that many other applications within the enterprise rely on the messaging system. Once that distinction becomes clear in their mind, you'll probably find that procuring necessary budget funds becomes a great deal easier. 

Plan for international recovery

Dec. 8, 2003

By Mike Talon

Recently, while I was on an overseas mission, I got the opportunity to speak about disaster recovery to an audience of European technical professionals. Some DR specialists have clients in overseas locations, so it's important to examine how they handle or mishandle their DR and high availability (HA) situations.

There are major issues that exist when you plan to provide DR and HA for satellite offices in other countries. Most noticeably, you almost always run into high bandwidth costs between continents. In addition, it's difficult getting technical staff to the various sites to perform setup and installation tasks. But you can overcome these issues with careful planning.

Bandwidth between countries will remain expensive in the foreseeable future. Try to get the best deal that you can, and minimize the bandwidth usage of the DR and/or HA solution. If possible, use different offices within the same country to back up each other. However, you should be prepared to shell out some cash to bring data back to the main office if your business continuity plan (BCP) requires you to do so.


Another major issue to overcome is technical staff located in the wrong country for a particular install. It's extremely expensive to fly staffers from one country to another, so you must either budget the proper funds into your DR planning sessions or make alternate plans.

One example of an alternate plan is to leverage vendor professional services to perform installations. Or, you can rent technical staff from independent contractors in the country in question. While both alternatives cost money, some analysts contend that this costs less than sending your own people to the remote location.

Don't forget to take foreign languages and customs into account. There was a client who set up a perfect DR plan, only to find out that no one in the remote office could read the DR documentation. Between time differences and the language barrier, it took the company nearly 48 hours to enact a DR plan that called for two hours of downtime.

As for local customs, let's just say that asking certain employees to be available 24/7 may run you afoul of religious regulations and even civil law in some countries. This is why it's important to make contingency plans.

Planning for overseas DR and HA is by no means impossible, but it takes a great deal more planning than domestic operations. Do your homework, find the best deals, and create contingency plans that all of your employees can understand. These preparations will help if disaster strikes. ☹



Prevent Tomorrow's Disaster Today!

As an IT professional, you need to be 100-percent confident that your IT organization's systems and data are fully protected in the event of a crisis. If you have even a shred of doubt that your disaster recovery plan is as complete as it should be, you need to order TechRepublic's Disaster Planning and Recovery Pack today.

Evaluate data centers as part of a disaster planning and recovery strategy

Jan. 28, 2004

By David Southgate

Companies with mission-critical data systems and near-zero tolerance for downtime must have disaster recovery measures in place, often including one or multiple external data centers. But selecting a cost-effective data center with five nines of availability (or 99.999 percent uptime) is wrought with challenges. With so many possible points of failure at a data center, CIOs need a guide to help them evaluate collocation providers.

The following report highlights key questions CIOs should ask when evaluating data centers, with advice gathered from three top data center experts: David F. Locke, contingency planner with San Mateo, CA-based financial services company Franklin Templeton; Lou Kirchner, president and CEO of Sacramento, CA-based data center Herakles LLC; and Ron Hughes, president of California Data Center Design Group, who has been involved in the design, construction, and operation of data centers for over 20 years.

Outsourcing: A low-cost alternative

With the cost of building your own data center running upwards of \$1,000 a square foot—well out of reach of many enterprises even if the construction expense is amortized over a 20-year period—some companies opt for external data center providers. These providers once offered facility leases for as much as \$150 a square foot per month, according to Hughes. Because many data center providers built facilities under faulty financial and market assumptions, several have been forced into bankruptcy and sold for pennies on the dollar. The new owners can now provide rock-bottom data center leases priced at \$8 to \$12 a square foot per month.

Low prices, however, can't be the only measure when selecting a collocation center.

After evaluating their needs from a data center, CIOs should consider the following:

Geography

When selecting a data center provider, one of the first things to consider is the geographic location of the data center itself.

- ▶ How close does the staff need to be to the data center in the event of an emergency?
- ▶ How quickly can staff get to an external location via alternative forms of transportation in the event that airplanes are grounded?
- ▶ Is the data center located in a geographic threat zone, a place prone to seismic disturbances or severe weather, such as hurricanes, or is it located on a 100-year flood plain? (All of these pose a real threat to a data center's uptime.)

Facility design

In an evaluation of a facility's design, it's crucial to look for any single point of failure. Examine building construction plans with an eye to determine if the site can be operated and maintained without an interruption in service.

- ▶ Is the facility designed for five nines of availability?
- ▶ Is there adequate redundancy and reliability designed into the systems?
- ▶ What types of mechanical systems are installed?
- ▶ In terms of air-conditioning units, does the facility use chilled water or air cooled units? (Air cooled units are less efficient and drive up facility costs and your end price.)
- ▶ Are there redundancies built in to the facility's electrical switches that power the floor? (If the electrical system for the building shorts out, no form of electric power generation will fix the problem.)

- ▶ Are there at least two fiber carriers entering the building?
- ▶ Are these carriers entering the building in different fiber vaults?
- ▶ What is the present electrical load in the facility?
- ▶ What capability does the facility have to support itself if power from the main grid goes down?
- ▶ How many backup generators exist? (There must be at least two generators, each one with the capacity to generate enough electricity at full load to power the facility.)
- ▶ Has the company used best-of-breed data center practices as defined by associations versed in business recovery and emergency planning, such as BRMA or AFCOM?
- ▶ Has the infrastructure been tested in a real emergency or simulated emergency?

Site security

Security of the data center itself is an important consideration. For that reason, before examining the inside of a data center, start on the outside by asking the following types of questions:

- ▶ Does the facility stand alone or does it share a building with other tenants?
- ▶ Do other tenants pose any kind of threat to security (direct or indirect competitors, many people entering and exiting, frequent shipments of products)?
- ▶ Is the site surrounded by fences, berms, or shrubs that pose easy access points to and from the perimeter?
- ▶ Is the parking lot gated with guards at the checkpoint?
- ▶ Are the security guards professionals employed by a security firm or are they merely employees of the data center? (Data center employees may not be properly trained.)
- ▶ Are the security guards present 24 hours, seven days a week?
- ▶ Are there adequate security cameras covering the perimeter and entrances?

- ▶ Is activity on these cameras being recorded?
- ▶ How long are video surveillance tapes archived?
- ▶ Does the building have controlled access with biometric security?

Operations

In terms of operations, look for any single point of failure and don't overlook things that seem small or insignificant. An overlooked detail, such as a poorly designed fluorescent lighting circuit can cause a power panel to trip, leaving an entire facility without electricity. This particular scenario may sound far-fetched, but one of Franklin Templeton's data centers experienced just that, according to Locke. Locke's team traced the problem back to a failed fluorescent lighting ballast, a kind of transformer.

Insufficient maintenance of the facility's power generators or air-conditioning units might cause downtime, too. The risk of failure rises as data centers attempt to save money by trimming maintenance schedules. To begin an evaluation of the operations, examine the following:

- ▶ Who has actual accountability for operating the facility?
- ▶ How often is maintenance of air conditioning, heating units, power generators, and other systems performed? (Ask to see the records.)
- ▶ In the event of a failure in the air conditioning, power generator, or other system, are there sufficient spare parts on site to fix a variety of common or likely problems?
- ▶ How often are backup generators tested?
- ▶ Are the backup generators tested under full load to simulate real power demands for several hours?
- ▶ What expertise does the support staff have, and is the same level of support available every hour of every day of the year?

Finances

Many collocation centers built in the last decade were developed under business assumptions that proved to be unrealistic. When the Internet went bust, the predicted

demand for collocation centers failed to materialize—and so did the revenue. Many facilities were put on the auction block while others cut back on maintenance and staff in an effort to remain solvent. Under these circumstances, companies interested in leasing space at a data center would be wise to look at the provider's financials as well.

Ask these questions:

- ▶ Was the facility sold in a fire sale, making it easier for new owners to actually make a profit?
- ▶ Is the company profitable?
- ▶ What are the operating expenses?
- ▶ What are utility costs and maintenance costs?
- ▶ Do the operational costs, including full-time staffing of engineers, security, maintenance, and utilities, seem realistic, or does it appear that the company is trimming corners somewhere?

A final precaution

Finally, even after a thorough evaluation of the preceding points, if the facility passes muster, it doesn't hurt to bring in a consultant or interview the original designers of the facility. Performing the preliminary investigation with your own staff helps to familiarize them with a facility and can save on consulting costs. But external expertise can assure top management that the facility really has no single points of failure, says Locke. After all, when the servers are installed and the data is mirrored, a failure in the data center can be an embarrassing "gotcha" that no one wants to experience. ☹

Prevent Tomorrow's Disaster Today!

While IT managers can't plan for every service interruption, they can take proactive measures to ensure that such a disturbance won't result in a full-blown IT disaster. Of course, proper planning and preparation are key to achieving this goal, and TechRepublic's Disaster Planning and Recovery Pack provides the resources you need to do just that.

This valuable set, consisting of the *Administrator's Guide to Disaster Planning and Recovery, Third Edition* book, CD-ROM tool kit, and laminated Disaster Planning & Recovery Checklist and Contact Sheet, will help you:

- Develop a comprehensive disaster contingency strategy
- Assess your organization's vulnerabilities
- Avoid facilities management mistakes

- Implement a formal crisis communications policy
- Test your disaster recovery plan
- Select the right vendors for server colocation and other backup needs

As an IT professional, you need to be 100-percent confident that your IT organization's systems and data are fully protected in the event of a crisis. If you have even a shred of doubt that your disaster recovery plan is as complete as it should be, you need to order TechRepublic's Disaster Planning and Recovery Pack today.



Members share why planning, and more planning, is key to disaster recovery

Oct. 11, 2002

By David Southgate

CIOs typically plan for systems recovery in the event of some kind of network outage. But they often don't take the next critical step—ensuring business continuity and operational functions. In stopping short, they put the company at risk of losing customers, reputation, and revenue.

To develop a sturdy disaster recovery and business continuity plan, experts and TechRepublic members advise CIOs to work alongside business partners to identify mission-critical business processes. They also must consider operational and logistical scenarios. And lastly, enterprises must test and maintain disaster response plans regularly.

Analyzing business processes and IT systems

To overcome the traditional IT-centric approach to disaster planning, **Scott J. Darling**, a TechRepublic member who manages product development at an accounting and tax consultancy, recommends a two-pronged, team-based approach:

1. The first prong features teams from each operational division. Each team is responsible for outlining the essential business processes for its respective division.
2. The second prong focuses on core corporate processes—sales, regulatory, or cash management. CIOs might consider structuring teams around each of these core corporate processes if the staff number and complexity warrants multiple teams.

As the various teams analyze the business processes, it becomes clear how the underlying technologies relate to each other, according to Darling. This, he added, also demonstrates the interconnectedness of IT processes and businesses strategies.

After leaders have completed the analysis, the business unit teams and the core process

teams should then develop recovery plans around three scenarios:

1. Manual recovery, or a means of staying in business without any of the IT systems in place
2. Interim recovery, in which only a few workers have limited access to systems and desktops
3. Full recovery, in which all systems are restored, and seats and access have been established for everyone

This methodical approach is certain to identify IT's role in the total recovery process, explained Darling.

Yet as another TechRepublic member explains, the job doesn't end there.

Include operational and logistical considerations

A complete disaster and business availability plan must prepare for upheavals within physical operations as well as anticipate the logistical needs in overcoming them, as this anecdote from TechRepublic member **Kyle Kuda** illustrates.

Late one Friday evening, the IT specialist received a call at home from his company's security monitoring company. A power spike had tripped the surge protector, causing the UPS backup to kick in. "When the batteries were drained, the system went into a bypass mode, allowing nonconditioned power into our UPS," Kuda explained. A second power surge started a fire in the UPS.

Kuda's business continuity response team quickly declared the incident a disaster, triggering the planned response: phone calls to hardware suppliers, employees, a cleaning company, and phone system reps, and a plane trip to the cold storage site in another city to restore lost data—an activity the company had tested the year prior.

By Monday morning, the team had restored the server room and most functionality. To customers, it was as if nothing at all had happened, added Kuda.

The need to test and update plans

One important part of Kuda's plan included a test of the IT and logistical response. Such drills should occur at least once a year, according to **Eric Beser**, TechRepublic member and CEO of Owings Mills, Maryland-based Ennovate, Inc. Beser's firm, which provides business continuity services, recently published a free e-book, *A Guide to the Perplexed Business Owner: Helping Your Business Survive the Unexpected Shutdown* (<http://www.expertpractices.net/article.aspx?cid=1&y=2002&m=7&d=7&ebook=1>). The book provides tips on how to do business continuity and emergency planning.

Beser said plans should be tested when a company is moving operations or relocating a data center, or else annually. He also recommends having the response team leader randomly select people from the response team to be "no longer available." This prompts team members to crosstrain in each other's domain areas.

"During these tests, I'll ask a CFO of a company if he can recover financial data," explained Beser. Usually the answer is no, which demonstrates a possible weakness in the logistical response to the plan, he added.

Following mock drills and test, companies often learn important lessons—for instance,


that they have not provisioned for an emergency meeting place if the office suddenly becomes incapacitated. Response tests might also reveal the need to maintain and update documentation, including service phone numbers, service plans, account numbers for all systems, and vendor relationships including financial services.

These disaster recovery plans should be updated any time a change is made in company policies, procedures, equipment, or key personnel, and then shared with all employees and involved vendors.

Laying the groundwork can save millions

The bottom line of proper disaster recovery and business continuity planning is the difference of a few thousand dollars to restore lost equipment or the possibility of losing millions in equipment and revenue if the disaster affects client business.

"If we had not been prepared and been able to get back up and running in less than 48 hours, the cost...would have been much more," wrote Kuda about his server room disaster. "The loss of confidence in our business continuity would have caused clients to pull their business. A few thousand dollars in insurance vs. a few million in lost revenue....It's no contest." ☞



Prevent Tomorrow's Disaster Today!

As an IT professional, you need to be 100-percent confident that your IT organization's systems and data are fully protected in the event of a crisis. If you have even a shred of doubt that your disaster recovery plan is as complete as it should be, you need to order TechRepublic's Disaster Planning and Recovery Pack today.

Notes

Prevent Tomorrow's Disaster Today!

While IT managers can't plan for every service interruption, they can take proactive measures to ensure that such a disturbance won't result in a full-blown IT disaster. Of course, proper planning and preparation are key to achieving this goal, and TechRepublic's Disaster Planning and Recovery Pack provides the resources you need to do just that.

This valuable set, consisting of the *Administrator's Guide to Disaster Planning and Recovery, Third Edition* book, CD-ROM tool kit, and laminated Disaster Planning & Recovery Checklist and Contact Sheet, will help you:

- Develop a comprehensive disaster contingency strategy
- Assess your organization's vulnerabilities
- Avoid facilities management mistakes

- Implement a formal crisis communications policy
- Test your disaster recovery plan
- Select the right vendors for server colocation and other backup needs

As an IT professional, you need to be 100-percent confident that your IT organization's systems and data are fully protected in the event of a crisis. If you have even a shred of doubt that your disaster recovery plan is as complete as it should be, you need to order TechRepublic's Disaster Planning and Recovery Pack today.

