hp strategy brief

developing a business continuity plan

contents

introduction

In the business world, computer disaster recovery planning is evolving toward business continuity planning. In recognition of this trend, in 1995, Disaster Recovery Institute (DRI) International, an organization founded in 1988 to provide a base of common knowledge in continuity planning, replaced the designation *Certified Disaster Recovery Planner (CDRP)* with the designation *Certified Business Continuity Planner (CBCP).*

What is the difference between disaster recovery and continuity planning? A disaster recovery plan is reactive and usually focuses on recovering the computing environment. Although work may be done to harden the computing infrastructure to prevent a disaster, the plan's main purpose is to recover from damage to the infrastructure. In contrast, a business continuity or contingency plan is not only proactive, but it is also targeted at keeping the business running during an event, not just recovering the computers after the fact.

As part of its continuity planning process, a company needs to review the continuity or recovery of manufacturing, packaging, warehousing, shipping, customer support, and any other facilities or operations that are critical to the company's survival.

Many companies today do not have a working continuity plan. Of those companies that do develop a plan, many proceed without sufficient knowledge or input from end users. Because end users are not involved in developing the continuity plan, their manual procedures, physical facilities, hard-copy records, and other special needs are often overlooked. Thus hardware and applications might be recovered, but not the business processes that use them.

One way to approach a business continuity planning program is called the *Delphi method.* Experts in each business function identify their critical business processes and develop separate (but coordinated) continuity plans for each process. The benefits of this distributed approach are many:

- Business processes that are not critical do not hinder those that are, so limited resources can be used effectively.
- An infrastructure that supports noncritical business processes does not get recovered.
- Multiple critical business processes or applications can be recovered in parallel.
- Applications that normally run on different systems can be recovered on the same system, if necessary.

The survival of your business after a disaster depends on having a continuity plan in place. This brief is intended to help you develop and deploy one. Directed at a hypothetical employee, a corporate continuity planner, it details the procedures for launching a continuity planning program.

beginning the continuity planning process

Developing a corporate continuity plan involves a number of steps as well as some basic concepts.

where do I start?

Suppose you are beginning a corporate continuity planning process for the Absolutely Best Company (ABC), a manufacturer of top-quality widgets (an imaginary company with an equally imaginary product). There are several steps involved, as detailed in the following paragraphs.

obtain management commitment

For the continuity plan to be successful, management must be committed at the highest level. The plan must be part of the strategic business plan, and the company must budget appropriately and separately for the continuity planning program. A top-level policy statement should be issued that

- Affirms the value of business continuity
- Acknowledges and accepts the associated costs
- Documents management responsibilities
- Includes the goals and expectations of the plan, as well as any organizational assumptions or parameters

identify critical business functions

Assuming that you have management support, the next step is to identify how the company obtains its revenue in terms of business functions. For example, ABC first takes orders for widgets and then builds the widgets to meet those orders. Next, the widgets are installed, and the customers are billed. Finally, the employees are paid so that the process will continue. Other revenue comes from service and support, but because those are distributed functions, an event that has an impact on the corporate site should not directly affect them.

There also are numerous ancillary business functions within ABC, such as ordering parts from suppliers and paying them. Because parts are usually ordered months in advance and payables probably could be postponed until normal processing resumes, these business functions are usually not a critical concern. Note that a just-in-time supply chain would require substantially more complicated planning, which is beyond the scope of this paper. HP Services consultants can assist you with this planning.

Once you have defined the gross critical business functions (not the infrastructure, such as networks or computer applications, that supports them), a risk assessment and a business impact analysis should be performed for each of the business functions and then, if appropriate, for the infrastructure supporting them. Remember to analyze dependencies, because a business function that appears noncritical could be supporting one that is critical.

This information should be communicated to the responsible senior managers so that they can confirm and then rank the criticality of each business function. The executive staff or office of the CEO should have final approval of the list, because those individuals are in a better position to understand the company in its entirety.

After you have determined what the critical business functions are, it is time to contact the people who will write the plan to ensure the function continues or is recovered. Following a distributed planning methodology, you will not actually write any continuity plans—you will be acting as a consultant and a facilitator. The daily users know best the tasks involved in the business functions.

build business process core teams

You start by building business process core teams consisting of information technology (IT) operations management, end-user management, applications support staff for each critical business function, and the records management department. This team technique is called the *Delphi method*.

Through the Delphi teams, you develop a clearer view of the infrastructure (for example, processes, records, IT applications) the teams believe is critical to performing their business functions.

build a corporate team

You should also build a corporate team, consisting of members from core infrastructure support functions such as accounting, auditing, information technology, facilities, human resources, legal, public relations, investor relations, purchasing, postal services, records management, risk management, safety, security, shipping/receiving, and telecommunications. After a disaster, not only will these departments be required to continue their support roles, but also they may have to implement major infrastructure changes to support the affected areas. The legal, public relations, and investor relations departments need to keep the public and stockholders informed of the company's operational status after an event has occurred.

risk assessment and business impact analysis

The purpose of a risk assessment and a business impact analysis is to answer the following questions:

- What am I trying to protect? (system inventory and definition)
- What am I trying to protect them from? (vulnerability and threat assessment)
- What controls are currently in place or needed to prevent or minimize the effects of potential loss? (evaluation of controls)
- How much am I willing to spend on those controls? (decision)
- Is the money I am spending effective? (communication and monitoring)

The continuity planner (or planning coordinator) helps each Delphi team answer these questions as depicted in figure 1.

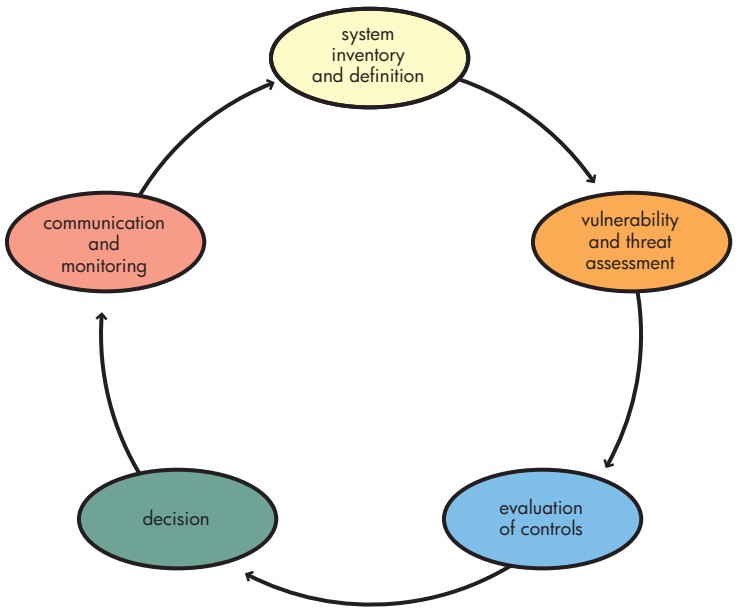

Figure 1. Disaster risk assessment.

The risk assessment involves identifying threats, vulnerabilities, risks, and the business impact of a disruption for each entity. *Threats* are events that could disrupt an entity. Some examples are natural disasters (wildfire, flood, earthquake), disasters created by people (burst pipe, fire from an electrical short), acts perpetrated by disgruntled employees, and mistakes. *Vulnerability* is the susceptibility to the threat (chances that an entity can be affected by a threat). For example, the closer a building is to an earthquake fault, the more vulnerable it is to an earthquake. The *business impact* from these risks can be loss of revenue, customers defecting to the competition, damaged reputation, or disgruntled employees if the company cannot pay them.

A risk assessment can show that an entity thought to be critical is not or vice versa. It is pointless to perform a formal risk analysis if the informal analysis shows that there is little vulnerability to the risk (for example, a snowstorm in New Delhi or a drought in London).

Before you begin the ranking process, determine what criteria to use. Generally, criteria are split between quantitative and qualitative. Quantitative losses can be expressed as a number, such as an annualized loss exposure (ALE). The simplest ALE equation is shown in the sidebar.

More extensive ALE calculations are outside the scope of this paper but are a standard part of any insurance company's activities. HP Services can assist with ranking risks, or you can use any of several well-defined methodologies (see "Conclusion").

At this time, ABC is interested in

- *Monetary exposure:* Would it cost ABC money if the function could not be performed?
- *Customer exposure:* Could ABC lose market share if the function were suspended?
- *Legal and regulatory exposure:* Is ABC required to perform the function?
- *Intracompany dependencies:* Would suspension of the function affect the critical activities of another area?

To start the risk assessment, rank all of the entities whose loss could negatively affect ABC's business, gain consensus from each Delphi team, and then merge the results for presentation to and concurrence by upper management. Note that some threats have a time component as well. For example, a power failure that lasts a few minutes may not be a disaster, but one that lasts hours could well be.

a simple ALE equation

(R)isk = f x E

(r)isk = f x e

B = R – r – c

B = f x (E – e) – c

legend

f = frequency

E = exposure to threat without control

e = exposure to threat with control

B = benefit

c = cost of maintaining the control (such as an uninterruptible power supply to guard against power failures)

recovery time and recovery point objectives

As part of the risk assessment, the Delphi teams estimate how long an entity can be unavailable, how old the information supplied by the entity can be, and how much of it can reasonably be lost when it is made available again. That is, they determine the recovery time objective (RTO) and recovery point objective (RPO).

Recovery time objective refers to the time from when the event occurs until the business process (for example, the accounting department, the accounting application, or manual procedures used by accounting) must become active again (recovered). This can also be called the *recovery window.*

Recovery point objective describes the point in time to which the data must be recovered—stale (old or obsolete) information no longer reflects the state of the company. This can also be thought of as the *freshness window.* The teams consult with management to confirm their decisions.

Believe it or not, these two goals are not tightly coupled, nor are they completely decoupled. The RTO and RPO should be determined independently, although it may be determined at some later point that they are interrelated due to infrastructure or technology issues.

If a plastics manufacturing plant cannot make any products for a day, that may not in itself be a problem; however, if liquid plastic cools in the pipes and the machinery while systems are down, the recovery costs could be enormous. This might indicate a long RTO but a very short RPO. In this case, uninterruptible power might be required to keep the plastic warm, though the machinery may not need to be running.

One example often used to demonstrate an application that can never fail is a system running a stock exchange. This implies an RTO of zero or very near to zero. In addition to minimal RTO, it is an absolute requirement that the RPO be zero. Transactions can be valued in the millions of dollars, and the last transaction against a specific security sets the price for the next transaction.

Because they build upon each other, transactions must be serialized. For example, a client might sell one stock so that she can buy another, or might buy a stock and sell it within minutes. After cutover to a backup system, one missing transaction could mean that every transaction following it is wrong and the database can diverge further and further from reality. No matter what, transactional integrity must be maintained. That is to say, if you buy a security from me, it must be removed from my account and added to your account. If either part of this transaction fails (for example, because of a locked record), the entire transaction must fail. (Figure 2 illustrates this situation using a transfer from a checking account to a savings account.) Because trading floors are disappearing from the securities markets and most large firms make their money through "programmed trading" run by computers, there is no manual backup. If the exchange's computers are down, the exchange itself is down.

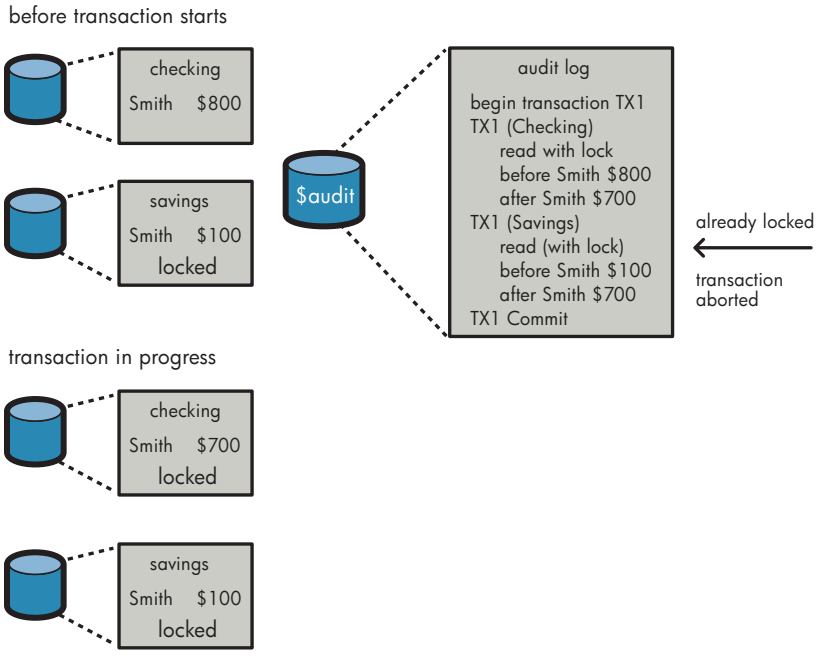

Figure 2. Transaction failure caused by a locked record.

Whereas the RTO for an automated teller machine (ATM) application might also need to be as close to zero as possible, the RPO is not nearly as critical. Transactions made at ATMs tend to be well under US$1,000, and they are usually confined to one account holder. Again, you want transactional consistency so that if money is transferred between accounts or is withdrawn, there is an accurate record. But the RPO does not need to be zero because the database will not wildly diverge if one or more transactions are missing, and the ATM keeps a log that can be used for reconciliation later on.

We have seen two examples in which the RTO must be close to zero, but what about systems where the RTO can be longer? One example might be an electronic funds transfer (EFT) system. If the system is down, there are manual backup procedures, such as using a telephone or fax machine. However, the RPO is zero because the loss of even a single multibillion yen transaction can ruin your day (not to mention your career!).

Whereas the RTO of a Web-based sales system must be close to zero, the RTO of the back-end systems (shipping and billing) can be much longer. If the back-end system is down, orders can be printed out and mailed or faxed to the warehouse and credit cards can be cleared manually. But if a Web-based system is down, customers will go elsewhere to do their business. Similarly, if the order entry system for a mail order business is down, orders can be processed manually.

The longer the RTO for a particular entity, the less the cost will probably be to recover it (figure 3). Unfortunately, at the same time, the losses from the entity that is unavailable are escalating (figure 4).

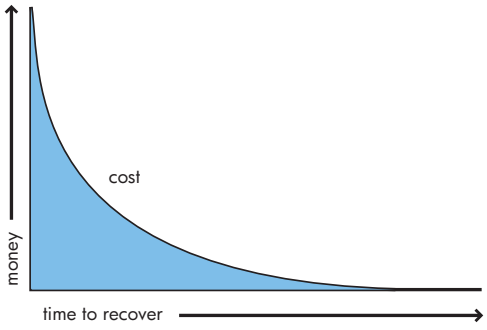

Figure 3. Cost of recovery.

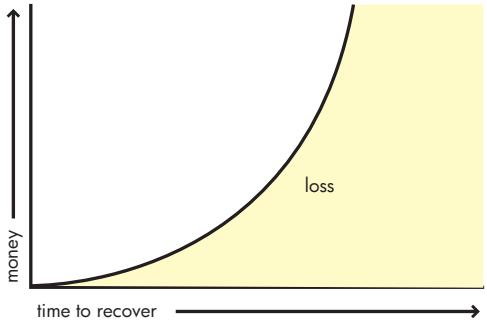

Figure 4. Loss due to the time an entity is unavailable.

By overlaying the curves in figures 3 and 4, you can see that there is a point at which the potential loss equals the cost of recovery (see figure 5). This represents the starting point of the maximum cost of the business continuity plan for each entity. With this viewpoint, you can determine how long the function will be allowed to be unavailable and how much to spend on your plan. Note that this is only a starting point—remember to include qualitative losses as well.

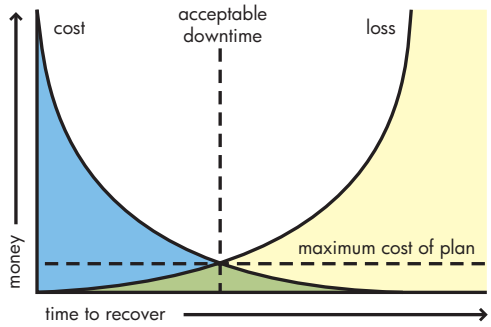

Figure 5. Maximum justifiable cost of plan.

There's one final point. Remember that the RTO must include the time needed to make a decision to deploy the plan. If the RTO for a particular business process is 10 minutes and you spend 8 minutes deciding whether to deploy the plan, then the steps needed to restore the business process can only take 2 minutes. Unfortunately, in a crisis decisions must be made quickly while the information required to make them trickles in (see figure 6).

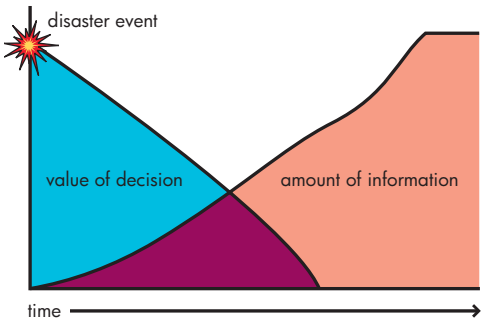

Figure 6. Although the information needed to make a decision about whether to deploy your plan accumulates over time, the value of that decision decreases the longer you wait.

the continuity and recovery continuum

RTO and RPO exist along a continuum, and part of your plan should be to determine where each critical business function should be placed (see figure 7).

The most traditional disaster recovery medium is magnetic tape. At some interval, you copy your computing environment from the disk to a tape and keep the tapes somewhere safe. The shorter the RPO, the shorter the backup interval. If at some point the backups are running continuously, a tape system may no longer be the right answer.

If the RTO is long, you simply order a new computer, load the tapes, and start your application. If the RTO is short, you may pay to have a computer on standby at a hot site or computer rental company. There are many factors to consider when evaluating the use of tapes.

- Where is your tape backup hardware?
- Where are tapes stored until they go off-site?
- How quickly do your tapes go off-site?
- Are duplicate tape copies sent via different routes?
- Do you perform tape retrieval and restore tests?
- Do you perform backups logically and ship recovery tapes in "waves" (just in time) or all at once?
- Is your application in a quiesced state when you make backups?

The objective of making tape backups is to create a safe copy of the information resident on your computer in case of a disaster. A disaster can be anything from someone accidentally purging or changing a file to total site destruction. The tape hardware should be located physically away from the computer that it is protecting so that an incident affecting the computer doesn't affect the tape drive as well. Tapes should be moved away from the backup device into a fireproof safe somewhere else until they can be sent off-site. This protects the tape from a disaster affecting the backup hardware itself.

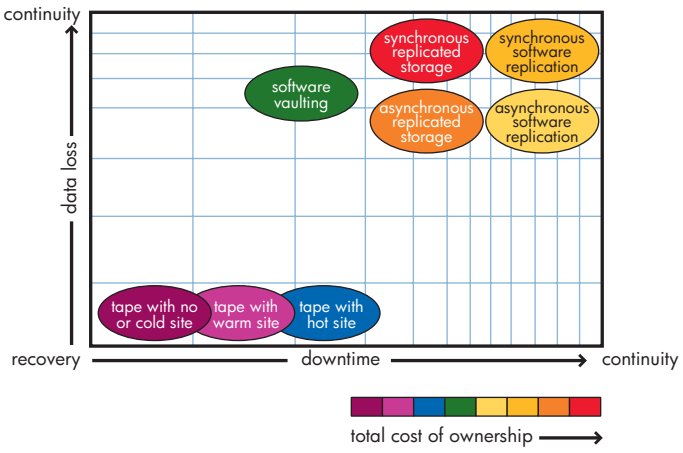

Figure 7. Different technologies can be used to meet differing downtime and data loss requirements.

disaster tolerance: closing the freshness window

If you find that tape backups are running continuously, then real-time technologies, such as online data *replication* or *vaulting,* may be the answer. These technologies can practically close the freshness window. Data replication is more expensive than routine backups but could be deemed necessary in some situations, such as Web transactions, wire transfers, and supply chain applications. Replication or vaulting duplicates data onto an off-site location as it is manipulated on the primary system. The semantic difference between these two terms usually is that vaulted data is batched and/or stored offline and needs to be moved to the backup hardware, whereas replicated data is sent in near real time, possibly directly to the backup computer system's storage, and is ready to run.

If your company is already geographically dispersed and application uptime is imperative, you can create *application domains* at more than one site and distribute the load. With load balancing routers, redundant communication lines, or other methods, transactions can be split between multiple servers running in multiple sites.

When the load is being shared in this manner, you do not actually have primary and secondary systems or sites. What you do have are the beginnings of indestructible, scalable computing. New servers can be added at any time, and applications and database files can be migrated between the servers and sites as needed, so that any server or any site can be taken offline for testing, maintenance, or upgrades. In a properly designed application domain, a failure is undetectable, except perhaps for application slowdown.

HP and its partners provide several disaster-tolerant and data replication products. Contact the HP sales team for more information.

alternative plans and controls

Once risks are assessed and recovery windows are determined, the planner can ask the Delphi teams to begin outlining possible continuity plans for their business functions, starting with the most critical. The baseline is the before risk or ALE (no plan or controls in place). For each alternative plan or control, the Delphi teams need to calculate the *after* risk or ALE (plan and controls in place) along with the cost of the plan and controls.

The plan execution time is important. If the entity can be unavailable for only a few hours, but the execution time is two days, reevaluate either the acceptable recovery time objective or the plan itself.

Because computer resources can be at a premium in a disaster, replacing computer-based processes with manual processes is an option that should be explored. Besides computing costs, the planner should also address such matters as alternative sites, temporary personnel, hotel and meal costs, off-site records and forms storage, and installation of new phone lines. The purpose of this step is to calculate the cost/benefit ratio of plans and controls for different recovery objectives.

At some point, the potential loss reduction (savings) will be less than the expenditures required to develop and implement the continuity plan. (See figure 8 for a comparison of four hypothetical plans.) At this point, the executive staff needs to determine which controls and recovery window each plan should address based on cost versus savings and time to deploy. Remember that some risks, such as loss of customers or reputation, can cause unforeseeable losses in the future. Possibly a continuity plan that costs more than the expected loss should still be implemented, based on potential loss of future revenue and corporate image.

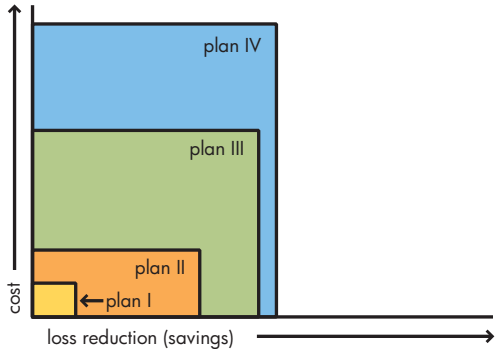

Figure 8. Plans I and III would break even (at different levels). Plan II would generate savings in excess of its implementation cost, and Plan IV would cost more than it would save.

When an actual crisis occurs, a crisis response team—composed of executive staff members and guided by one or more emergency response teams consisting of Delphi team members—determines if a disaster should be declared. The team considers the estimated time the affected entity will be unavailable, the recovery window, and the time it will take to execute the plan. If, for example, the recovery window is two days and the entity will be unavailable for only one day, it may not make sense to declare a disaster.

Declaration of a disaster needs to be reevaluated if the entity is down longer than initially expected, based on the time remaining before the entity can be made available again and the time required to execute the continuity plan.

documentation and standards

As a continuity planner, you assess current documentation.

required documentation

Before developing a continuity plan identifying activities to be performed during a disaster scenario, you need to understand how those same activities work every day. This means that you need access to some basic documentation for each business process. Determine if the following exist:

- Change control process
- Standard operating procedures (SOPs) for end users
- Run books
- Special forms requirements and special peripheral requirements for the operations staff
- Data flow diagrams and problem isolation procedures (for the business functions as well as the computer applications)
- Tape backup/rotation schedule used for records management

You also need to gather more specific information about your company's business functions. For example, find out if sufficient application downtime is scheduled to back up the databases used by the business functions, or, if the software allows it, if the databases are being backed up online. (If your company runs 24 x 7, you don't have time to take your system down for backups.) Determine if there is an archival process to remove inactive records from hard-copy files and databases so that they are kept at a manageable size. Also, identify where critical records are being stored: on-site, off-site, or out of the region.

If any of this information is missing, a continuity plan can be started, but it may not be effective. Without knowing how the business functions operate and what is required for normal processing, no one can guess the requirements for exception processing. Actually, the bulk of the work of developing a continuity plan lies in ensuring that standard business practices are documented and followed.

For a plan to be effective, you also need a commitment to gather or create the necessary documents and to begin a consistent backup program, if one does not exist. If you encounter opposition by pointing out that commitments are lacking, you can call in a consultant to identify the deficiencies so that you retain cooperation for developing a plan. You can also point out that by having these documents and processes in place the company will be more competitive in the long run due to

- Faster isolation of application bugs
- Fewer operational mistakes
- Reduced support requirements
- Faster training of new staff
- Easier maintenance and enhancement
- Fulfillment of industry and legal regulations
- Easier auditing

The Delphi teams must be especially aware of any non-IT-related infrastructure required to keep each of their functions running—for example, if they need special paper, printers, or inks. Some questions to be answered include those relating to the storage and accessibility of important phone numbers, PC hard drives, and removable media.

developing standards: a "cookbook"

The continuity planner either develops or purchases a standard set of forms and procedures to be used by each function: a continuity plan "cookbook." Without standards, each business function's plan will look different, making coordination difficult.

Different practitioners divide and name the phases of the continuity planning program differently. The DRI International defines seven phases of a continuity planning program as follows:

Phase 1: Project initiation

Phase 2: Functional requirements

Phase 3: Design and development

Phase 4: Implementation

Phase 5: Testing and exercise

Phase 6: Maintenance and update

Phase 7: Execution

Your cookbook should be developed or software should be selected during phase 3 to be used by the business function teams during phase 4. At a minimum, it should contain the following:

- Step-by-step approach for each group to follow in writing the continuity plan
- Corporate team description, stating which corporate resources will be available to assist each business function in developing its plan
- Notification process
- Plan considerations
- Responsibility list

As discussed previously, the corporate team performs the centralized tasks common to all business functions. By building the corporate team, you eliminate duplication of effort among the business functions while the plans are being written, as well as ease resource contention during an actual disaster. The corporate team should consist of a primary and secondary contact from each department involved.

The corporate team maintains an overall list of resources required for each business function's plan and resolves conflicts. Resource requirements are submitted via *corporate support forms.* Corporate support forms provide a standard way to list the resources that will be required of the corporate team members should a disaster occur. Each sheet is forwarded to the appropriate corporate team member, who then coordinates the requirements necessary for all business functions. If you have chosen a PC-based disaster recovery package, the forms can and should be modified to fit the format of the software. It is important that you find software that fits your planning program, rather than attempt to fit your planning program to some arbitrary product.

Although the planner maintains the first-level response part of the plan, including lists of important phone numbers, at the corporate level, the security department probably should start the actual notification process in the event of a disaster.

The plan considerations section identifies the issues that must be addressed by business functions as they write their plans. This section consists of details that might be overlooked by people who have not previously written a continuity plan.

A responsibility list is a script or checklist, by job function (not by name, because you don't know who will be available in a disaster situation), of what each person will be required to do during the seven phases of plan execution:

- *Evaluation and declaration:* Using information about the situation as well as the criteria for declaring a disaster, determine if a disaster should be declared and what parts of the plan to deploy.
- *Notification:* Deploy the notification process detailed in the plan.
- *Emergency response:* Deploy the parts of the plan that activate your hot site, relocate people and equipment, pull supplies from storage, hire temporary personnel, notify the media, stop trading in your stock, and so forth.
- *Interim processing:* Continue running your business as effectively as possible.
- *Salvage:* The damage assessment team works to salvage as much as possible from your primary site after pictures are taken and insurance issues are handled. Remember that equipment, microfilm, paper, and magnetic media degrade rapidly if not properly removed, stored, and recovered. This team also estimates the cost and time required to restore the primary site. It may be decided that restoring the site is not feasible, so the company should relocate to another primary facility.
- *Relocation/reentry:* Your company needs to move out of the emergency site to its previous facility or to new facilities. Enter any manually generated information into your automated systems.
- *Resumption of normal processing:* At the end of plan execution, your company's business functions return to normal. Remember to debrief everyone involved in plan execution and update and test the plan as necessary.

It is much easier to divide and assign the tasks necessary for recovery while the plan is being written, rather than in the middle of a disaster. The cookbook should contain sample responsibility lists and layout hints. Note that tasks should be assigned by teams or by job titles and not to individuals.

writing the continuity plan

Finally, you are ready to facilitate the development of a written continuity plan. If employees within each business function write the plan for that function, you achieve multiple goals:

- Employees know what the day-to-day activities are.
- Different functions can be generating plans at the same time.
- You gain buy-in of the business function employees.

You or your staff should be available to answer questions as each group writes its continuity plan. Not only does each business function need to generate a plan, but also each department represented on the corporate team needs to have a plan, in case those departments are also affected by the disaster.

Although the plan can be written online, it must also be stored on paper to help ensure accessibility. Some continuity planning software vendors sell Web-based services that make your plan available from any PC. The drawback to these services, however, is that if you do not have access to a PC or if the disaster also makes the vendor's site inaccessible, you may not be able to access the plan.

corporate continuity planner responsibilities

As the corporate continuity planner, you supply each business function's continuity planner and each corporate team member with a continuity plan binder containing plan dividers and a basic phone list. The phone list contains the direct phone numbers for police, fire, ambulance, hospital, hazardous materials team, government authorities, and utilities. The binders given to the corporate team members also contain the office, home, mobile, and pager numbers for each corporate team member, backup, and manager. These binders are used for storing each business function's continuity plans after you and the corporate auditors have approved the plans.

Each executive staff member is given a small packet for his or her briefcase, containing an emergency phone list, an executive staff phone list, and a plan execution "quick start" document.

Corporate team members keep a copy of the binder, with their own continuity plan and the requirements sheets from each business function appended, at home and in their offices. The various business functions select a primary contact and a backup contact to care for their own plans, which should also be stored at home, locally, and, if applicable, at the backup site.

The corporate continuity planner has copies of each corporate team member's and business function's continuity plan. As keeper of the entire corporation's continuity plan, you ensure that these copies can be retrieved in case all other copies of a particular continuity plan are destroyed. Copies should be maintained on-site in data safes and at the company's off-site data storage vaults.

exercising the plan

Continuity plans should not be tested—they should be exercised. Tests are passed or failed, whereas exercises are conducted for practice. Exercise the plans thoroughly to ensure that they work. During the exercise, note any problems that occur and encourage feedback from participants. The purpose of the exercises is to reveal any defective or missing components in the plans. It is counterproductive to reprimand someone for pointing out errors or omissions. Some companies spend hundreds of hours testing and refining specific parts of their plans until they are satisfied with the time or accuracy of execution.

It is best to exercise and update the plans at least annually, or when major changes occur. Update call lists quarterly. It is better to have no plan at all than to have an out-of-date plan—such a plan lends a false sense of security and wastes time during an actual disaster.

pitfalls of a continuity planning program

The primary pitfall of a continuity planning program is that often its expenditures are charged directly against profits and appear to provide no immediate benefit. Other pitfalls include

- The program is viewed as negative or depressing.
- Continuity planners usually work with employees who lack awareness of the problem.
- Employees often are expected to work on the plan in addition to their other tasks, possibly overburdening them.
- Many companies do not have adequate documentation of their daily activities or information flows to begin plan development, so those must be developed first.

tips for successful continuity planning

Here are some practical tips.

contact an amateur (ham) radio organization

If your company has an amateur radio club, that would be a great start to help with your recovery efforts. Contact the American Radio Relay League (www.arrl.org) to locate a local group if you are in the United States, or your national ham radio association if you are in another country. While you use your security radios to recover the business, amateur radio volunteers can be called on to assist with human safety issues. Their license restricts the use of ham radio to life safety and non-business-related communications. Volunteers cannot have forms delivered but can handle almost anything relating to people, including location, medical needs, and safety. Regional authorities can curtail the use of cellular phones and some or all business radio bands as well as commandeer phone and radio equipment during an emergency.

assess application availability requirements

Five nines, or *99.999 percent uptime,* is usually calculated based on unplanned downtime. This number isn't always relevant, especially because scheduled maintenance and installing patches is usually excluded from the measurements. According to *Interactive Week,* 30 percent of all application downtime is attributable to scheduled maintenance. Some vendors advertise 99.999 percent hardware uptime, but this does not necessarily mean *application* uptime. If your application is down, so is your company.

If the risk analysis shows that your computer application cannot be down even for routine maintenance, you need to select systems that allow for online upgrade and maintenance—for example, to add processors or disks, introduce new application code, and reorganize database files. HP NonStop servers provide that continuous availability.

refine your backup-and-restore process

As noted previously, many companies spend hundreds of hours refining their backup and restoration procedures. Some methods of decreasing your recovery time are as follows:

- Separate and back up files in the order that they need to be restored. One such partitioning could be operating system, application, database files, and transaction logs.
- Pull and ship files from your off-site storage in "waves" so that the most critical files arrive sooner. While the next set of tapes is being pulled, the first set can be on its way to the recovery site. If you have duplicate tapes, ship them separately so that a transportation mishap doesn't derail your recovery.
- Determine how many tape drives are required at the backup site to meet your RTO. One company uses 64 drives in parallel.
- Tape is inexpensive. Filling tapes to save money could prevent you from recovering in the allotted time because multiple restore threads might need different files on the same tape. One company cut its file restoration time by 50 percent by optimizing tape usage.
- If at all possible, automate your recovery tasks so that mistakes are minimized. One company has 800 batch jobs that restore the application and database from tape, roll the database forward with the logs, and check consistency.

enhance your plan with everyday procedures

Ensure that your change control process has links to your continuity plan. If processes, procedures, or infrastructure change, the plan may need to be altered.

Instead of storing special forms or other supplies in a warehouse or ordering them "just in time," rotate them from the vendor to your backup site or off-site storage and then to production. In this way, your supplies are always fresh, and if you declare a disaster, they are on hand. When on-site gear becomes obsolete, rotate it to the backup site or off-site storage for use during recovery.

work with local authorities

It is helpful to contact other companies and regional civil disaster management authorities during development of your plan. Schedule facility walkthroughs with local police, fire, hazardous materials, and utilities (such as power and water) personnel so that they become familiar with your facility before you need to call on them in a disaster situation. They need to know where your computer rooms, data closets, flammable or toxic chemicals, and power supplies are located.

By working with local authorities before a disaster, not only can you gain insight and expertise, you can also find out if elements of your plan conflict with those of other companies or with rules in effect during a regional disaster declaration.

When it is time to exercise the continuity plan, consider inviting the authorities to participate. The exercise will be more realistic, and you can find out if there are any coordination issues while there is still time to make corrections.

conclusion

A continuity plan should not and cannot be written by the IT department alone, nor should it be written solely for a given computer or data center. Developing such a plan is a long-term process that requires substantial human and monetary resources throughout the company. Without a long-term commitment to continuity planning from the highest executive levels, efforts to develop such a plan are bound to fail.

The planning process cannot even begin without documentation for everyday processing already in place, including change control procedures, standard operating procedures, run books, data flow diagrams, problem isolation procedures, and a tape backup or rotation schedule. As a byproduct, continuity planning forces more formalized standard documentation across the entire company, resulting in faster isolation of application bugs, fewer operational mistakes, reduced support requirements, faster training of new personnel, and easier maintenance and enhancement of current applications.

Not only is a continuity plan required by many regulatory agencies, it could also mean the survival of your company.

for further information

HP offers a wealth of products, solutions, and services for all of your business continuity and disaster recovery needs, including

- Fault-tolerant and high-availability computer systems and storage modules
- Data replication software and hardware
- Outsourced (or your own company's) hot-, warm-, and cold-site planning, design, implementation, and operation
- Professional services for all phases of your planning program

For more information on any of these offerings, contact the HP sales team.

DRI International was founded in 1988 to provide a base of common knowledge in continuity planning, a rapidly growing industry. It also administers the industry's only global certification program for qualified business continuity or disaster recovery planners. The organization's Professional Practices for Business Continuity Planners, a common base of knowledge, serves as the industry's best practices standard. Refer to the website at www.dr.org.

The *Disaster Recovery Journal* was the first publication dedicated to the field of disaster recovery and business continuity and now has more than 50,000 subscribers. The *Disaster Recovery Journal* also sponsors two annual conferences, one on each coast of the United States. More than 2,500 continuity planning professionals from all over the world attend these two conferences combined, making them the largest in the entire industry. The journal's website is www.drj.com.

Some helpful books on continuity planning are

- *Designing Controls into Computerized Systems.* FitzGerald, Jerry, and FitzGerald, Ardra F. 1990, Jerry FitzGerald and Associates, ISBN 0-932410-40-5
- *Disaster Recovery Planning.* Toigo, Jon William. 1999, Prentice-Hall, ISBN 0-13-084506-X
- *Data Replication: Tools and Techniques for Managing Distributed Information.* Buretta, Marie. 1997, John Wiley & Sons, ISBN 0-471-15754-6

For more information, go to www.hp.com/go/nonstopcontinuity.

5981-3912EN