the Availability Digest

Fire in the Computer Room, What Now?

June 2007

Heading to work, you hear on the radio that your office building has just burned down. Your cell phone rings. It's the CEO asking what your plans are to recover. How soon will it be before you have recovered the company records and are providing data processing services again? You *have* planned for this, haven't you?

The recovery from a disaster such as this requires extensive disaster recovery planning long before a disaster strikes. This book, *Fire in the Computer Room, What Now?*[1] walks us through the creation of a Disaster Recovery Plan to handle just this sort of situation.[2] Following such a disaster, if you have to ask, "What Now?," then it is already too late to create this plan.

What is a Disaster?

The authors define a *disaster* as "an extended service interruption of the data processing services of an organization which cannot be corrected within an acceptable predetermined time frame and which necessitates the use of an alternate site or alternate equipment for recovery."

Disasters include:

- Local site disasters (computer room) – fire, flood, catastrophic system failure, power failure, sabotage.

- Site disasters (building) – bombings, explosions, fire, flood, power outages.

- Area disasters (metropolitan) – bombings, earthquakes, environmental contamination, explosions, disease outbreaks, plane crashes, volcanic eruptions, wind or snow storms, terrorist attacks.

A study performed by the University of Minnesota showed that more than 60% of companies that faced a disaster and did not have a Disaster Recovery Plan were out of business in three years.

Another study by Contingency Planning Research Inc., showed the frequency of disasters to be:

• Fire	17.5%
• Terrorism	17.5%
• Hurricane/Tornado	14.0%
• Earthquake	10.5%

[1] G. Neaga, B. Winters, P. Laufman, *Fire in the Computer Room, What Now?*, Prentice-Hall; 1997.

[2] This book is a companion book to the *Continuous Availability Systems Design Guide* reviewed in the January, 2007, issue of the *Availability Digest*. That book focused on the planning required to move to a continuously available system. This book focuses on preparing a Disaster Recovery Plan to ensure continuous availability.

Power Outage	9.5%
Software Error	8.8%
Flood	7.0%
Hardware Fault	5.3%
Burst Pipe	3.5%
Network Outage	3.5%
Other	2.9%

The Disaster Recovery Plan

Recovery from a disaster that destroys a data center or renders it inoperable for an extended period of time can only be done by migrating to a backup data center. The Disaster Recovery Plan defines the needs of the backup data center and the procedures for switching over to it.

Creating a Disaster Recovery Plan is not a simple exercise, and its execution is not inexpensive. The plan:

- must be designed to match the business requirements.
- may involve the building or modification of a backup data processing site.
- will require the development and testing of many new procedures.
- will involve a tradeoff between cost, recovery time, the completeness of the recovery, and the scope of the disasters covered.

The authors describe a six-step iterative process for creating a Disaster Recovery Plan:

- Determine what the business requires.

- Determine the data processing requirement.

- Design the backup/recovery solution.

- Select the products to match the design.

- Implement the backup/recovery solution.

- Keep the solution up-to-date.

Determine What the Business Requirements Are

This first step involves the preparation of a *Risk Analysis* and a *Business Impact Analysis*. These analyses identify what the business process priorities are and what recovery time is required for each business process.

Risk Analysis

The Risk Analysis identifies the risks that the corporation faces. These risks include but are not limited to:

- the physical security of the facility itself.
- the protection of the business's data.
- disgruntled employees.
- backup and recovery systems.
- the vulnerability of the infrastructure, such as power, communications, and water.

- the location of the data center (that is, to what disasters is it susceptible).
- key skills upon which the business is dependent.

Business Impact Analysis

A *business process* is a group of related activities that support the successful operation of the business. A business process may include data processing applications. It is the recovery of critical data processing applications that is the subject of a Disaster Recovery Plan.

The Business Impact Analysis covers:

- which business processes should be included in the scope of the recovery.
- the cost in terms of lost revenue, lost customers, or other metrics of an outage to each business process.
- the maximum allowable time of an outage for each business process.
- the acceptable amount of unrecoverable transactions for each business process.
- the recovery priority for each business process.
- the dependencies between business processes.

Determine the Data Processing Requirements

Once the needs of the business are understood, these can be translated into the procedures and resources needed to support recovery and ongoing processing at the recovery site.

The data processing requirements at the backup site are determined by creating an inventory of all of the applications used by the business. Each application should be assigned recovery attributes that are necessary to meet the critical needs of the business. For each application, these attributes include:

- the maximum acceptable downtime.
- the maximum tolerable data loss.
- the data currency required upon resumption of processing at the backup site.
- hardware capacity requirements (processors, disks, tapes, printers).
- network requirements.
- service levels to be maintained during and after recovery.

Design the Backup/Recovery Solution

The disaster recovery design defines the scope of the recovery – that is, what is being recovered in what time frame. The design includes:

- what types of disasters are included in and excluded from the recovery.
- the sequence in which applications will be recovered.
- the maximum recovery time for each application.
- the data that will be recovered.
- the currency of the data once it is recovered.

The strategy for testing the backup site must be determined. Is the backup site company-owned or provided under contract with a third party? Can testing be carried on at the backup site without impacting applications which the backup site is normally running? How frequently should testing be done?

The backup and recovery processes are defined in detail during this step. These processes are the heart of the Disaster Recovery Plan. The most important process within this group is the backup and recovery of data. The corporate data is one of the company's most important assets. The plan should address how data will be backed up, how it will be stored safely off-site (physically or electronically), how it will be retrieved following a disaster, and how it will be recovered.

The data backup requirements will be different for different classes of data. Data classes include application data, metadata, and system data.

Some data will be lost following a disaster. This could range from data that had not yet been backed up to data stored on unreadable or lost magnetic tapes. Procedures should be determined for reconstructing lost data that is so critical that its loss cannot be tolerated. This includes correcting inconsistencies between related files or tables. Furthermore, data must be resynchronized between its electronic copy and other media, such as paper or microfiche.

The readiness of the backup site must be specified. Is it a cold site, a warm site, or a hot site with an up-to-date database and all applications running? Is the site unmanned, manned, or operable remotely? If the site is unmanned, there should be remote monitoring of environmental parameters such as power, temperature, and humidity.

Finally, the configuration of the backup site must be specified. This includes not only the obvious hardware requirements, but also the networking interconnections between the primary site and the backup site. In addition, there may be requirements imposed on the applications and databases. Applications should be designed to be portable, avoiding complex linkages with other applications. Databases should be designed so that they are self-contained, and they should consist of small, transferable partitions.

Select Products to Match the Design

The Disaster Recovery Plan designed up to this point is a logical design. It describes how backup and recovery will function, but it does not say how it will function.

The selection of hardware and software products will be dependent upon the processing platform at the primary site since the data and procedures used at that site will be transferred to the backup site following a recovery.

The first step is to choose a backup site if one does not already exist. Thereafter, product selection will be based on need, compatibility, functionality, product quality, and cost.

Implement the Backup/Recovery Solution

This is the implementation phase. The backup site is acquired and provisioned if necessary with the appropriate hardware and software. New data backup, storage, and recovery procedures may have to be put in place. The backup site configuration and procedures must be integrated into the company's change management processes.

Revised application design rules should be put in place to ensure that new applications and databases are backup/recovery-ready. Human resource procedures should be reviewed to cover issues such as employee notification following a disaster, who works and who doesn't, compensation, counseling, and ongoing status notification to employees.

At this point, the Disaster Recovery Plan should be formalized, approved, and documented. It should include:

- the recovery scope and assumptions.
- the process for recognizing a disaster and invoking the plan.

- the identification of the recovery teams and their members.
- the major tasks and responsibilities of the recovery teams.
- the owner of the plan.
- how the plan will be maintained.
- how the plan will be tested.

Keep the Solution Up-to-Date

Changes to the data processing environment in a data center are constant, and any change has the potential of rendering the Disaster Recovery Plan useless. Therefore, procedures must be put in place to keep it viable. These procedures include maintenance, auditing, and testing.

<u>Maintenance</u>

Changes may come from the development of new applications, changes to hardware configurations, changes to the network, organizational changes, system changes, and changes to the backup site.

The Disaster Recovery Plan must be part of the change management process to ensure that it is updated to reflect any change that impacts it.

<u>Auditing</u>

The plan should be independently audited periodically. Typical audit periods are six months or a year.

<u>Testing</u>

Perhaps the most essential part of a viable Disaster Recovery Plan is periodic testing of the recovery procedures. This provides several important functions:

- It is a rehearsal for the operations staff to keep them trained in recovery procedures.

- It verifies that all changes have been made to the plan to keep it up-to-date.

- It uncovers mistakes that have been made in the plan.

- It creates a comfort that the recovery plan really works.

Testing can either be planned or unannounced. A mix of these is optimal.

Testing the recovery plan can itself cause a disaster, especially in the early tests. During the design of the Disaster Recovery Plan, the potential negative impact of a test gone wrong should be considered, and procedures should be in place to quickly recover from such problems.

Summary

The process for creating a Disaster Recovery Plan detailed by the authors in *Fire in the Computer Room, What Now?* is not really a strict sequence of actions. Often, information from a future step is useful in an earlier step, such as having an idea of available backup and recovery hardware and software options when designing the plan. Furthermore, the development of one step may impact previous steps, which must be reworked.

Therefore, the creation of a Disaster Recovery Plan is highly iterative; and its development should be flexible.