Paper Number: HPW97CZD
Disaster Recovery—"The Big Picture"

Michael Patrimonio
Melillo Consulting, Inc.
285 Davidson Avenue, Suite 202
Somerset, NJ 08873
908-563-8400 Ext. 8446

Many MIS professionals define a computing disaster as a catastrophic loss of data and believe that a good backup plan will protect them against any disaster. When, in reality, an MIS disaster should be defined relative to the loss of a business's "mission critical" service(s). Service interruptions can occur during events which range from a power outage to a disk crash to the destruction of an entire data center. The business needs of your organization will dictate the criticality of data processing services and their tolerance to interruptions. There are some environments, where the loss of some or even all computing services for even a few hours would be acceptable and not termed as disastrous. While in other cases, the loss of a critical server for any period of time would be unacceptable and considered a problem. It's when you are unprepared for these problems that they become disasters.

As is the case with most IT endeavors, a key to success will be gaining the proper insight into your organization's business needs. This can only be done by going out to the people who have those needs, throughout the organization, and soliciting them for this information. The assembly of a Disaster Recovery Planning team would be one way to facilitate this process. Since this will be the group that is going to define the business needs of your company, in the event of a data processing service interruption, inter-departmental cooperation and communication should be of prime concern. Once the group has been assembled, in order to develop a more thorough understanding of the organization's business needs required for a comprehensive disaster recovery plan, here are some of the questions that will need to be addressed:

- If the system(s) were to be turned off right now, who would be effected and why? For example, lost access to an application or data, loss of computer driven services (i.e., robotics, conveyor belts), etc.
- Of those people that would be immediately effected by an outage, how severe would it be? Could users continue working? Will the sale or manufacture of the company's product slow down or stop? Would there be a loss of revenue?
- Who would be effected by the disruption in a service provided by one of the directly impacted groups? For example, the Shipping department can't work because the Order Processing group has been taken off-line.
- What would each group effected by the disruption in service, both directly and indirectly, need to

continue operating at full capacity? At a minimum acceptable capacity? For example, connectivity for all their users, connectivity for a couple of the key users, a special form that can be used to perform their jobs in writing and have the information entered into the system when it becomes available.

- During the recovery of a system, if there was going to be a loss of data, how much of a data loss would be acceptable? Note: This will determine the level of fault tolerance and/or redundancy you may need built into your system.
- Based on the level of fault tolerance the group wants to achieve, what are they willing to give up in terms of scheduled downtime for system maintenance? If anything at all.

There are those individuals that believe any type of disaster recovery planning or plan should lie entirely within the IT department's bailiwick. They do not see the need for a group of people, representative of the areas of the company that would be adversely effected by a system outage, to have a say in the planning process. I strongly disagree with this philosophy and think that they are actually doing themselves a great disservice by excluding such valuable feedback. While the onus of recovering and restoring the business's data processing services is on you and the Technical Services group, the burden of recovering and restoring the business's operation, as a whole, has to be shared by the entire organization.

This next phase consists of three parts, the collection of information about the applications and other resources the Technical Services group is currently supporting, the merging of this data with the information collected in the previous step and the prioritizing of the entire list. The prioritizing is important because this is where expectations will be set, if there is ever a disruption in service, everyone will know what needs to be done and is going to happen when—from the first step through to the last. Once again, this should be a phase in the planning schedule where the whole group should have some say, especially since there may be at least one group that will receive less than what they ask for and if issues like this are not handled in an open and diplomatic manner, the problems that this has the potential of causing could condemn your plan to failure. Senior management should be involved as chief arbiters as well as the group responsible for prioritizing any items related to the loss of revenue and major expenditures. Some of key items that the Technical Services group will need to investigate include:

- What applications do we have running and how critical are they to the organization?
- What resources do we have in-house that can provide the organization with a knowledge base that's at the same level as the technology currently in service?
- Are the company's technical support agreements with the specific vendors at an appropriate level?

- Is the IT department's infrastructure set up so that any single points of failure have been eliminated or secured?
- Does the organization have multiple production and/or development systems on site or within the organization and are they comparably sized? Could any one or more of these servers be used to compensate for the loss of any other server?

Once the business needs have been defined and prioritized, a comprehensive set of requirements based on these needs must be outlined—these are going to become the building blocks for the Disaster Recovery Plan. The list of requirements will be comprised of policy, procedure and purchasable solutions. Some of the key items to be considered:

- Problem Escalation—Who should be called, when? When should the disaster recovery plan be activated and who can activate it? Who will have the authority to suspend activation of the plan?
- Downtime Tolerance—How tolerant will your business be to downtime? At what point will your business begin losing revenue? Up to what point in time should the system be recovered—within one hour of the failure? As of the last full or partial backup?
- Hot Site Services—In the event of a catastrophic loss (i.e. the entire data center), it will allow you to resume data processing services while the "data center" is being rebuilt. The costs include a monthly retention fee, guaranteeing the availability of the site in the event of a problem as well as a substantial charge for activating the site, which could make this an expensive proposition.
- Backup Strategy—Full, Partial or Incremental? Online or Offline? Validation of tapes? Off site tape storage? Note: Full backups are complete system backups, Partials will store everything that has been modified since the last FULL backup and Incremental backups will store everything that has been modified since the last backup. When evaluating the different backup methods always consider the recovery process for each method. Perform benchmarks of both the backup and the restore.
- Backup Software Solutions—On-line or hot backup software solutions require a thorough evaluation (backup and restore). There are a number of products available with specific features (good and bad), so choose wisely.
- Backup Hardware Solutions—DDS2, DDS3, DLT, Optical Disk, Reel Tape. What are the store and restore throughput rates? Evaluate your needs and choose wisely.
- System Redundancy—A spare system is going to be expensive, unless the cost can be justified by using it for development or load balancing. Also, a spare system on-site with the production system is not going to do any good if the site is destroyed.

- High-Availability Solutions—Software/Hardware combinations are expensive, have complex implementations and require a high level of maintenance.
- Disk Mirroring—Requires double the amount of disk and I/O hardware but it's excellent protection against a disk, cable or I/O channel failure. Data corruption and user errors, however, will be mirrored.
- Raid—Expensive but good protection against hardware failures. Still vulnerable to data corruption and user errors.

The information has been collected, priced and prioritized. You are now ready to move into the planning phase, which is made up of two parts. The first step will involve the planning of any modifications or additions to the current standard operating policies and procedures that me be required, as a prerequisite for your Disaster Recovery plan. The design of the plan itself, which will be the next step, consists of a series of documents that will become the core of your system's Disaster Recovery Plan—Policies and Procedures manual. Since plans will vary dramatically from site to site, or even from system to system, I have developed a template that can be used in the construction of this manual.

I. Disaster Recovery Plan—Policies and Procedures Manual
   A. System Information—If a single manual is created for all of the systems, the information specific to each system should be in its own section.
      1. Identification "Page One" Information
         a) Host or Node Name—qualified with domain.organization, if applicable.
         b) System Model Name & Number
            (1) Example: HP3000 Series 995, HP9000 T-500
         c) Network Connection Information
            (1) Connection Types—Ether, X.25, FDDI, etc.
            (2) Address Information—IP Address, X.25 Address, etc.
         d) Operating System Names and Version Numbers
            (1) Example: MPEiX 5.5 (HP3000), HP-UX 10.20 (HP9000)
         e) System Serial Numbers
            (1) Hardware Serial Number—On the CPU or main system cabinet.
            (2) Operating System Software Serial or ID Number
               (a) HPSUSAN Variable—HP3000
               (b) "uname -i" output—HP9000
         f) Any other information you may feel is relevant and would like to have available for quick reference.

(1) Model Names & Numbers of other critical components—Bus converters, Disk arrays, etc.

(a) Serial Numbers and any relevant, but brief, configuration information. Example: Raid level of disk array, etc.

2. Configuration

a) Total Number of CPUs (If applicable)

b) Total Physical Memory

c) Total Physical Disk Space

d) Configured Swap—Size & Type (HP-UX Only)

e) System Boot Paths—Primary, Secondary and Others

f) LAN Type(s)—FDDI, Ether, X.25, etc.

(1) Appropriate addressing information—IP address, X.25 address, etc.

g) I/O Configuration Summary

(1) Example: 4 Bus Converters, 20 2GB SCSI Disk Drives, etc.

h) Any other information you may feel is relevant and should be available for quick reference.

(1) Model Names & Numbers of other critical components—Bus converters, Disk arrays, etc.

(a) Serial Numbers and any relevant, but brief, configuration information. Example: Raid level of disk array, etc.

3. Prerequisite System Reports—These should be run, printed and placed in the manual on a regular basis. How often your system configuration changes will determine the frequency of the run times.

a) MPE

(1) SYSGEN

(a) I/O Configurator

(i) LPATH—Lists, in path order, all of the systems I/O devices

(ii) LDEV—Lists, in device number order, all of the systems I/O devices

(iii) LCLASS—Lists all of the system's configured device classes, their attributes and the logical devices in the class

(2) VOLUTIL

(a) SHOWSET <set name> DSTATUS—lists, for <set name>, the logical device numbers and their paths for each disk in the volume set

(3) DSTAT ALL—lists, by ldev number, all of the disks, with their status, volume name and volume set affiliation

(4) BULDACCT.PUB.SYS—When run for @ or /, it generates two job streams that can serve two purposes.

(a) BULDJOB1 can be used as a reference of all users, groups and accounts on the system. With capabilities, volume set affiliations, access privileges, etc. BULDJOB2 can be used as a reference for all of the user defined catalogs (UDCs) currently set at every level on the system.

(b) As job streams, both can be used to rebuild your system after an install of

the operating system. BULDJOB1 will recreate the accounting structure and BULDJOB2 will reset the UDCs. This should be run just prior to the full backup so that it can be kept on tape and restored in the event of a failure.

(5) REPORT @.@—Reports on the groups and accounts configured on the system.

(6) HPSWINFO.PUB.SYS—If all HP products and patches have been installed correctly, this ASCII files will contain a list of all the HP products and patches installed on your system.

(7) Any other system information that may be appropriate.

b) HP-UX

(1) IOSCAN -FN—Scan the system and reports on all of the I/O devices found, with I/O paths, device file names, etc.

(2) VGDISPLAY -V—Reports, verbosely, on each of the systems volume groups. The information includes a list of all the logical and physical volumes in the group.

(3) LVDISPLAY—Lists more specific information for each individual logical volume. Caution: For volumes that are mirrored the "-v" will list the two physical volumes in the pair. Unfortunately, it will also list the status for each physical extent regardless of mirroring, making the report extremely voluminous.

(4) SWAPINFO—Reports on the amounts, types and locations of all configured swap space.

(5) Copies of /etc/fstab (HP-UX 10.x) or /etc/checklist (HP-UX 9.x)—Associates the various logical volumes with their mount points

(6) Copies of /etc/passwd and /etc/group—Provides a list of all configured users and groups on the system

(7) A copy of /etc/hosts—The local system's host name database

(8) NETSTAT -RN—Displays information on the system's routing tables with the IP addresses

(9) LVLNBOOT -V—Scans the systems disk drives for bootable disks

(10) SYSDEF—Produces a list of tunable kernel parameters with their current values, the minimum and maximum values, unit of measure, etc.

(11) SWLIST -L PRODUCT (HP-UX 10.x) or LS -L /system (HP-UX 9.x)—Produce a list of HP products and patches currently installed on the system

(12) Any other information that may be appropriate. Specifically, if a special service is being used, the configuration file in the /etc directory should be printed. Example: Service: NFS, Config File: /etc/exports

4. Third Party Software and Hardware

a) Account or ID numbers, product names and version numbers, Technical support telephone numbers, etc.

B. Escalation and Contact Information

1. In the event of a system failure, who should be called when . . .

a) Example: System crashes, operator will initially notify immediate supervisor and users and place a call into a vendor's technical support group if necessary. System is down for X, senior department management will be notified and users will be

updated. System is down for Y, company management is notified, the Disaster Recovery team is put on standby and users are updated. System is down for Z, the Disaster Recovery plan is activated.

2. Company Contact Information
    a) Support Personnel
    b) Management
3. Vendor Contact Information
    a) Company name, product name, version numbers, account numbers and technical support telephone numbers

C. Miscellaneous Policy and Procedure Documents—This item will vary greatly from site to site because it will depend not only on the defined business needs but also on the operations and systems staff and their level of expertise. Here are some examples:

1. How to report a problem. The information to record in the event of a system failure. Error numbers and messages, strange sounds or any other abnormalities
2. How to take and load a memory dump
3. If necessary, how to boot from an alternate boot path or the Support media
4. How to place a call to the HP Response Center or any other vendors that may need to be called
5. General system recovery procedures. Example: On the HP3000, repair of KSAM files
6. Application recovery procedures. Example: Recovering an application's database
7. Creating a Custom System Load Tape (HP3000) or using COPYUTIL on the system disk(s) (HP9000)

II. System Configuration Backups

A. MPE

1. CSLT (Custom System Load Tape)—The TAPE command in the SYSGEN subsystem will create a tape that can be used to boot the system. It can also be used to update or re-install a corrupt operating system, with the current configuration. If this is used in conjunction with private volumes, a re-install of the operating system does not mean a reload of the entire system. Tip: Before creating your CSLT, run the BULDACCT command to create the two BULDJOBs. Then use the STORE option of the TAPE command to backup these two files or the entire SYS and TELESUP accounts for added security. This could reduce your system recovery to a series of less than ten steps.

B. HP-UX

1. COPYUTIL—Creates an image of a disk on tape. In the event of a disk failure, the image can be laid back onto the disk. Warning: The source and destination disks must be the same size. The system must be brought down and booted off of the support media, and it will remain down for the duration of the copy. The same will apply to the restore.
2. VGCFGBACKUP/VGCFGRESTORE—This command will backup, to disk, a volume group's configuration information. If the LVM information for the system is ever lost or corrupted, this will provide a way of rebuilding it while leaving the user data intact.

Once the plans have been completed and approved, by management and the Disaster Recovery Planning team, the implementation and testing phases can begin. Any new or modified procedures that are being incorporated into the Standard Operating Procedure should have already been tested, but their

installation into the production environment should occur one at a time, to ensure that they perform as tested. The same will apply to any new technology. So if a problem does come up it will be easier to debug.

Finally, while this may be the last step, this is by far the most important—a regular review of the Disaster Recovery plan is essential to ensure that the ever changing IT environment is properly protected. In fact, if the plan ever failed, it would probably be due to a problem caused by some subtle change in the environment that would have been detected with a regular review of the procedures.

What I have tried to provide you with in this paper could have been done in a more concise manner had the subject matter been broken down into its more basic components, but I wanted to show the importance of "big picture" thinking and how necessary it is for groups within an organization to cooperate and communicate.