Developing Backup Strategies That Work

Ted Saul
OpenVMS Product Competency Center

Overview

Today's computing environment today demands that backup strategies be finely tuned and designed to meet the specific needs of each company. It is no longer simply a matter of running a backup each day with the intent of being able to restore system files should a disk failure take place. A data protection strategy must be in place to enable your company to continue to do business in the event of any of the following:

- A lost file caused by accidental deletion or data corruption
- A lost disk perhaps caused by a hardware failure
- A total system failure where a server becomes inoperable and the total environment will need to be recreated
- A catastrophic site loss caused by nature or hostile activity

Also affecting the criticality of today's backups is the ever-growing list of governmental regulations that require that certain data be available and accessible for a specific period of time. Consideration must be given to the new types of media that are available and their life expectancies. Thought needs to be put into the redundancy of backups that must be retained and where they will be stored. Safeguards should be put in place to ensure that data that needs to be restored will be available and restorable with the least downtime possible. As well, there are steps to be taken to make it possible to restore a system as quickly as possible.

How can you be sure your backup strategy works? Will you be able to recover from a catastrophic event ranging from a disk failure to a complete loss of site computing due to natural disaster or terrorist activity? Developing a disaster-tolerant backup strategy is one of the most important tasks that can be carried out in the IT environment. The design must consider the nature of your business, the size of your budget and the speed with which the system must be recovered. You can take specific steps to ensure the integrity of your backups and that all required data can quickly be restored in the event of a catastrophic outage. This article will help system managers review current strategies to ensure they meet the needs of their company and answer the many questions that arise in building such a strategy. It is also a starting point for new computing environments to help guide the design of their backup strategy.

Understanding your Environment and Defining Data Zones

The first step in a solid backup strategy is to have a good understanding of your data environment. This knowledge goes beyond simply executing a SHOW DEVICE command to see what disks are on your system; it extends into logically grouping your data for backup purposes. You can design *data zones* to group certain types of data together for backup protection. For example, static data that experiences very little change may fall into one zone while data changing frequently may belong to another. These data zones may be physically set up by disks or arrays of disks, which makes for easy viewing and understanding. As an alternative, they may be "virtual" groups established by directories with like data. This strategy can be documented on paper or by software but must be easily interpretable.

You will need to decide what data zone scheme works best for your company. A simple example of how three data zones might work follows.

- **Data Zone 1:** Consists of the data that is critical to the business. This zone is targeted for more frequent backups and placed at a high priority on the schedule. The files in this zone may change quite often and may be used for transaction processing. These crucial files may also not be easy to recreate from scratch, such as a system disk.
- **Data Zone 2:** Consists of files such as your system disk. Although they are clearly important to your system, these files may not change often. System files may also be backed up as a part of a standalone backup scheme. Backups of this zone will need to be scheduled after an upgrade of the operating system or applications as well as when ECOs have been applied. Consideration should be given to administrative changes that may be made, such as changes in Authorize, the DECNET configuration files, or other such files that may be modified as a part of system management. Note that you need to decide whether retention of files such as the Operator log and error log normally located in system areas are considered critical to your business.
- **Data Zone 3:** This group of files may include those that can be easily recreated in case of loss, files that are not required to recover an application, and user temporary files that have no effect on the running system. Some of these may be candidates for backups managed by users themselves, assuming the hardware resources are available.

Identifying data zones will also help set up the schedule for backups. It will be evident what files need to be backed up first each night to ensure their completion. A non-critical zone may be skipped should there be a hardware outage or an unplanned error that reduces the backup window.

Data Zone Backup and Restore Policies

Once you have organized the data into workable groupings, you can develop your backup and restore polices. Assessing your environment and answering a number of questions can accomplish this.

How much data is there to be backed up in each zone?

To identify required backup resources, you will want to calculate how much data you have in each of the defined zones. To see individual file sizes use the command:

```
$ dir/size

Directory SYS$SYSROOT:[SYSMGR]

ACCOUNTNG.DAT;6         1915
B.DAT;1                    1
COPY.DAT;1                31
CRASH.DAT;1               60
FILE_TO_EMAIL.DAT;2        2
IOGEN$PREFIX.DAT;1         9
K.DAT;2                   41
KERRY.DAT;1                1
MIKE.DAT;1                 4
MONITOR.DAT;3              5
.
.
.
.
Total of 15 files, 275 blocks.

Grand total of 2 directories, 43 files, 5414 blocks.
```

To capture the entire directory and a total use the following command:

```
$dir/size/total


Directory SYS$SYSROOT:[SYSMGR]
Total of 784 files, 1846347 blocks.
Directory SYS$COMMON:[SYSMGR]

Total of 303 files, 54898 blocks.
```

Divide your number of blocks by 2000 to obtain an approximate number of megabytes. With OpenVMS Alpha Version 7.3-2, you can use the following command to obtain the number of kilobytes for each file and the total for all subdirectories:

```
$ dir/size=units=bytes  [.*] *.dat
Directory SYS$COMMON:[SYSMGR]

ACMEXCPAR_GENVAR.DAT;1   4KB
AMDS$DRIVER_ACCESS.DAT;1
                         9KB
APACHE$CONFIG_DEFAULT.DAT;1
                         4KB
DECW$AUTOCONFIG.DAT;1   22KB
DECW$FS_CONFIG.DAT;1     4KB
DECW$RGB.DAT;1          18KB
VMS$AUDIT_SERVER.DAT;1
                         9KB
VMS$IMAGES_MASTER.DAT;1 22KB
VMSIMAGES.DAT;1          4KB

Total of 9 files, 99KB
Grand total of 2 directories, 15 files, 3.33MB
```

```
$ dir/width=file=30/size=units=bytes *.dat
Directory SYS$SYSROOT:[SYSMGR]

ACCOUNTNG.DAT;1                  3.94MB
APACHE$CONFIG.DAT;5                 4KB
IOGEN$PREFIX.DAT;1                  4KB
VMSIMAGES.DAT;1                    18KB

Total of 4 files, 3.97MB

Directory SYS$COMMON:[SYSMGR]

ACMEXCPAR_GENVAR.DAT;1              4KB
AMDS$DRIVER_ACCESS.DAT;1            9KB
APACHE$CONFIG_DEFAULT.DAT;1         4KB
DECW$AUTOCONFIG.DAT;1              22KB
DECW$FS_CONFIG.DAT;1                4KB
DECW$RGB.DAT;1                     18KB
VMS$AUDIT_SERVER.DAT;1              9KB
VMS$IMAGES_MASTER.DAT;1            22KB
VMSIMAGES.DAT;1                     4KB
```

```
Total of 9 files, 99KB

Grand total of 2 directories, 13 files, 4.06MB
```

With this information and the expected throughput of your drive and controller, you can get an idea of how large a backup window will be needed to complete backup of each data zone.

Contingency plans can be put into place should a problem arise during one phase of the backup. This might include canceling the backups of a low priority zone or shifting to another set of tape drives. Tracking how much data is to be backed up will ensure that you have enough tape volumes available for your backups. This will help prevent the surprise of suddenly needing a new tape volume and finding your backup waiting on a tape when you arrive in the morning. It will also help those responsible to make sure enough tapes are loaded into a tape library before the backup run begins. As an added benefit, sudden jumps in the amount of data being backed up may uncover wasted disk space caused by an event such as the undetected growth in an application log or transaction file.

How should each data zone be protected?

OpenVMS is known for its non-stop processing and continuous uptime. This can present some challenges when it comes to making sure all these data are backed up. Tape backup is only one option of many available to companies today when it comes data protection. Other options to setting up a disaster-tolerant environment include:

- Hierarchical Storage Management (HSM) Software to allow for the shelving and un-shelving of files to and from near-line storage. The near-line storage may be backed-up as needed, preventing the need to back up the online data.
- Volume Shadowing for OpenVMS that provides an implementation of RAID technology to ensure that data remains accessible in the event of media deterioration or of controller, device, or interconnect failure. Shadowing plays a vital role in disaster-tolerant cluster configurations by allowing for duplicated data and configurations over long distances. Members of shadow sets may be taken off-line to be backed up while the others remain online and available for processing.
- RMS Journaling to allow for the ability to provide a transaction oriented file-level journaling operations for RMS files. When used in conjunction with redundant sites, a disaster tolerant environment can be created. The standby site can then be used as the backup source while leaving the on-line site running at optimal levels.
- Application-based transaction protection provided with products such as RTR and ACMS.
- Hot Backups provided by the vendors of databases. These backups write data while the database is in an open mode, allowing processing to continue.

The key to a successful backup policy is to understand how quickly a file may need to be restored and ensuring that the data is reachable in that amount of time. This means that if a disk goes bad and a hospital patient file is suddenly corrupted, most likely that file will be needed almost immediately. In this case, volume shadowing may have been a good strategy to implement so that a secondary member can be brought into place quickly. However, if shadowing is unavailable, a quick find of the last backup of the patient file will need to be located to start the restore. A user file that contains information that is needed on a project next week, however, perhaps can wait until the tape is brought back from off-site storage to have the file restored by an operator at a scheduled time.

At some sites, it is nearly impossible to shut down processing for any amount of time to perform a standalone or OpenVMS backup. If this is the case, one of the above solutions may be best for you. However, a great number of OpenVMS sites can carve out a backup window and are able to send their data directly to tape. These users should move on to define the schedule for backing up their data zones.

How often should each data zone be backed up?

Depending on the number of transactions taking place in your application, this calculation may vary. Some businesses may be taking in thousands of records per hour and the effect of an outage and perhaps three to four hours of lost processing might be devastating. This decision will be based on the trade-off between the restore time and any downtime that may be required to execute a backup. For example, if your processing takes place between 8:00 and 5:00 each day, you may lose more data as the day goes on. Will it cost you more to delay processing in the middle of the day to perform some type of backup or will there be more loss by requiring a restore of data and re-entry of data from more hours lost the day of the outage and restore time?

Also affecting this trade-off will be:

- The speed of the equipment used for backups and restores. New efficient tape devices allow for the random selection of files and directories from tapes that may be already loaded in an automated tape library. When used in conjunction with a backup application solution, the files can be easy to locate and the process can be done almost automatically and unattended.
- The location of the volumes used to store the backed-up data. If the latest tapes have been sent to an offsite location or service, the length of time it will take to retrieve these tapes must be calculated. This should be factored into the cost of downtime to recover a system.

A common schedule template is to do a Weekly Image with Daily Incremental backups. In its simplest form, it allows for a complete backup of the system one day a week followed by a backup of files that have been changed each day. To conserve tape space, the incremental backups can be stored on a defined set of tapes that can be reused. The image backup volume set may be sent off-site for protection.

What devices do you have to back up each data zone?

When it comes to tape devices, there are two issues to consider when looking at where to store your data:

- The speed of the device.
- The reliability of the device.

You may want to send your most critical data to your fastest device but consider its trustworthiness. If the device consistently shows numerous errors caused by verification of reads and writes, not only will performance suffer but also the level of data integrity may be questionable. The availability and age of the associated media needs to be considered as well. If your tape device has been consistently using its Compact III tapes for any extended period of time, data written to the tapes may be compromised.

When should the backup of each data zone take place?

By gathering the information of how much needs to be backed up (Question 1), how often it needs to be done (Question 3) and what devices are available for the backup (Question 4), you can start to define your backup window. If you find this window overlapping with production time, you will have to take steps to reduce its size. This can be accomplished by using techniques such as the following:

- Using an Image / Incremental pattern of backups. The Image may be written on the weekend when system availability is greater.
- Off-loading low priority data zones to weekends.

- Reducing the iterations of backups taken against the system disk. This may require the "lock-down" of the system disk and scheduling times of system management work to correspond with the next backup.
- Increasing the number tape devices available to handle the backup.
- Replacing slower tape devices with newer, faster devices and interconnects.
- If backups are dependent on an operator schedule, implement an unattended backup scheme using automated tape libraries and software to manage the backups.
- Use OpenVMS Backup to a disk and copy the savesets in a separate operation.
- Implement Volume Shadowing to allow for the temporary breaking of shadow sets.
- Obtain redundant hardware for appropriate disaster tolerance.

What is best for your site will depend on your budget, backup resources available and the ability of your business to absorb computer outages.

How long do backups need to be retained?

Decide how long your data must be maintained. This length will depend on the industry such as Healthcare, Securities Trading and Department of Defense. Be sure to investigate the legal requirements of keeping your data. Data zones may also be formed by retention lengths required on groups of files.

Keep in mind the type of media that you are using for backup. Choose a medium that will have a life expectancy consistent with the data you are saving. Keeping multiple copies of time-critical data may also be something to consider.

What Will Be Your Data Restore Policy?

First, who will be responsible for the restore? Be sure that person understands the directory structure and the privileges required to write to those directories. This person must be able to find documented processes for restores quickly and be able to execute them confidently.

Before any restore takes place, you should determine why the failure took place so as to prevent the incident from reoccurring. If the failure was caused by a hardware failure, be sure that alternative hardware is in place. If an application or software has failed, be sure the problem has been fixed or an appropriate workaround is in place. A user that has inadvertently deleted a file should be coached to ensure that the behavior is not repeated. Data restoration can prove to be very expensive making it imperative that it is not done unnecessarily.

Timing of restores is important as well. Those performing the restores need to be aware of conflicts that might arise, including overlap with scheduled backups, or attempts to restore at a time of heavy system use. They must also be able to find a tape in a vault. A review of all types of file restores should take place regularly to ensure that operations personnel are familiar with the procedures. One does not want to generate another interruption of the system by attempting to restore at the wrong moment, for example.

Your restore policy needs to define how to find files in your collections of backup tapes. A backup application can be quite helpful in this role as it can provide catalogs that record which tape volume a file or disk is located on as well as the current location of the tape. Depending on the complexity of your backup environment, a backup listing may be used to provide similar data. It is important to have some type of organization to your multitude of tapes in order to locate the needed object quickly.

A last area to consider is whether or not there are predictable instances where files or directories will need to be restored. For example, testing data gathered through the month may need to be restored at one point in order to process and generate summary reports. Such restores can be scheduled as a part of IT processing and, if necessary, plans can be made to bring appropriate volume sets back onsite.

How will backup and restore failures be addressed?

Develop a plan on how to troubleshoot backup and restore failures. An operator should know exactly where to look first, whether it is looking for error lights on a tape drive or for the existence of log files. A backup should not simply be restarted unless some explanation for the problem can be presented. Otherwise valuable time may be wasted as well as allocation of more volumes than necessary. Operators should also be aware of how to handle unusual Opcom errors that are generated. Replying incorrectly may cause valid data to be overwritten or lost. Phone numbers and pagers should be readily available for operators to use in order to escalate their problem. Support organization numbers should be available along with a list of the pertinent log files that may be required.

Vaulting

Vaulting is the process of moving retained data to a location away from your systems. In case of catastrophic outage such as fire, earthquake or other, these volumes can be brought back in to rebuild a system. Often vaulting is handled by a third-party service that provides the pickup and delivery of tapes on a pre-determined schedule.

For vaulting to work successfully, a schedule of when tapes are to move off site and back on site must be developed. Begin by thinking of a worst-case scenario where your environment is severely compromised shortly after a backup has been done. Where do you want those most recent backup tapes to be located? If they are located on site with your system, they may be lost as well. However, if they have been shipped offsite, they will be safe.

Consider the following when designing a Vault scheme:

Is there an offsite facility available to be used as the vault?

The vault needs to be a secure location away from the main computing site. It is also good to put the tapes into a fireproof vault to guard against fire, explosion and water damage. Be sure that where you send your tapes will be secure not only from harm but also from theft and unauthorized access. The data on critical zone tapes must be considered crucial to the existence of the company.

Is there an offsite service to be used?

Reputable vaulting companies will pickup your tapes on a scheduled day while returning expired tapes as well. They are experts in ensuring that stored data is safe from the natural elements of the regions and can provide facilities against known threats in the area. The environment of their vaults should also encourage the lifetime of the media by providing appropriate cooling and humidity levels.

How fast can volumes be retrieved?

It is equally important to know how quickly your volumes can be retrieved from their offsite location whether they are managed internally or protected by a third-party service. The success of your restore policy will depend on having the needed tape volumes in hand when its time to start the return of the data.

Backup Verification and Testing

The moment of an emergency restore is not the time to find out if your backups have worked successfully. It is better to perform disaster recovery drills regularly. This will require some free space to be done without affecting current production machines. A test or backup machine makes a good candidate for this type of work. A simple test of a backup is to mount a tape in a drive an issue a BACKUP/LIST against the tape. Files should be displayed. The command will need to be reissued for each saveset on the tape.

Practice runs on getting tape volumes back from offsite services should also be executed at appropriate intervals. These may be when your staff has seen a significant turnaround or there have been changes at your offsite service.

Special Considerations

There are a few areas that should not be forgotten. Not all these areas apply to all business sites, but they are all too easily overlooked.

Databases – Do you need to use a vendor's backup product to back up their database? This may allow you to execute hot backups in order to get saved data while the processing continues and the database is still open. However, the format of the data on tape may not be the same as written by OpenVMS Backup. It may require that the application is running during a restore.

One strategy for database backups includes using a disk-to-disk-tape process (D2D2T). In this case:

- The database is shut down.
- The Vendor Software is used to write their saveset to disk using their own format.
- The database is restarted and processing continues.
- OpenVMS backup is used to write the database savesets to tape.

A disk-to-disk backup will go more quickly than a tape backup, thus allowing for a faster return of the database to an available state.

Standalone Backups – Booting a system into standalone backup and running an image backup is the only method to ensure that a system disk is backed up safely. This procedure should take place regularly while taking into consideration changes that have been made to system files. These changes may include updates to user accounts, DECNET files and TCP/IP host files.

Note: The use of /ignore=interlock during an online backup does not ensure that all files will be available from a restore. Standalone backup is the only method by which OpenVMS engineering ensures a bootable disk can be created from a backup.

Remote nodes – Nodes that may be located outside of the cluster or normal processing area may have an impact on backup schedules should they be using cluster backup resources such as libraries and tape drives. This type of backup needs to be looked at to identify how much it may require the use of the cluster resources and be scheduled in such a way that it does not interfere with other backup streams. Products that allow tape drives to look local to a network node can introduce challenges of their own and should not simply be assumed to work in conjunction with local backups without conflict.

Backup Consolidation – Products such as Hewlett Packard's SaveSet Manager (SSM) have the ability to combine multiple backups into one. This product can take a series of incremental backups and merge them with an image restore essentially taking a week's worth of tapes (5 or 6) into one. This strategy may be used to reduce the number of tapes sent to long-term off-site storage by combining backups and reducing the costs of tapes necessary to provide site data protection.

Backup Applications

Backup applications such as the Archive Backup System and Storage Library System can prove helpful in managing your backups. Common features found in these applications include:

- Ability to set up backup policies to assist with load balancing.

- Scheduling of backups to execute at required times.
- Cataloging of data on backup volumes for easy location.
- Interaction with robotic libraries for unattended operations.
- Easy restore operations including scheduling during off-hours.
- Notifications of backup failures.
- Vaulting services.
- Writing of multiple savesets on single tapes.

Many of these features address issues that have been mentioned previously. Though the initial configuration of this type of application can appear to be overwhelming, once set into place the operations are typically hands-off. Documenting how to handle problems within the application and where to look for logging information becomes more of a priority. It is also important to have posted where support will come from in case of problems.

Conclusion

There is obviously no generic "how-to" manual for setting up OpenVMS backups. There are however many basic issues to watch out for. Taking the time to establish appropriate backup strategies that meet corporate standards can prevent serious loss of time and money should the worst scenario come about. Backups are the insurance policy for computing departments. A backup strategy should be developed carefully and reviewed to ensure that it provides adequate coverage at all times.