Data Protection and Recovery Strategies

September 2004

Introduction

Data protection, disaster recovery and security are possibly the three most critical issues facing the IT industry today. Clearly, the impact of terrorist activity, natural disasters, a much higher fiscal responsibility and the globalization of our economy have repeatedly correlated the value of data to the survival of most businesses. Though determining the monetary value of data remains difficult and varies significantly based upon the business, knowing the relative value of data for a given business is becoming a more common practice. This enables a business to select the most appropriate high-availability strategy for its storage infrastructure. Not knowing the value of data to a business is becoming increasingly unacceptable as compliance requirements, governance boards and auditors now require a higher level of responsibility. Business continuance refers to the ability to maintain business operations in the face of any of a variety of problems, even if means that the service levels may be reduced. There are many data protection choices, and they all have tradeoffs that must be considered to achieve optimal availability based on the value of data. For reference purposes, four distinct levels of classifying data are most commonly used. These levels indicate which backup and recovery technology may be optimally suited and most cost-effective for each level.

Data classification	Description
Mission critical	Up to 15 percent of online data. Fundamental data required for business survival in the event of a disaster. Normally mirrored to disk and also backed up to tape in a different geographic location.
Vital	About 20 percent of online data. Data used in normal business processes but may not be needed immediately for a disaster recovery. Normally backed up to tape and/or replicated to lower-cost disk storage.
Sensitive	About 25 percent of online data. Data used in normal business processes that has an alternative source or can be reconstructed and may not be needed for hours or days after a disaster. Normally backed up to tape.
Non-critical	Typically 40 percent of online data. Data that is not needed for disaster recovery. Easily reconstructed or duplicated from prior backup copies.

Data protection considerations

Data replication covers a wide range of methods used to create additional copies of data either locally or in a remote location. Understanding the growing list of options available to implement a high-availability storage strategy is essential but can seem confusing, and existing solutions entail many tradeoffs. Nonetheless, a successful implementation of data-replication techniques can significantly improve the likelihood of surviving a business disaster.

Storage Pricing Guidelines

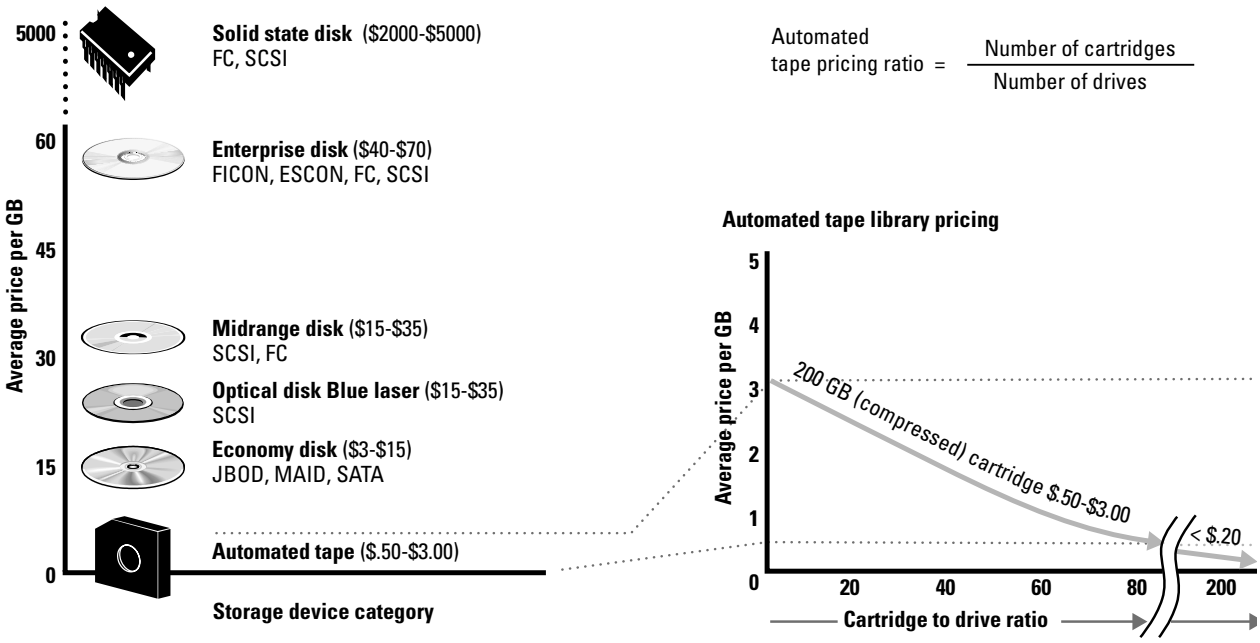

- All prices are ASP per gigabyte for working subsystems.
- The price per gigabyte decreases as the ratio of cartridges to drives increases, diverging from disk costs.
- Automated tape prices include drives, media and library and use a 2-to-1 compression factor.
- Tape cartridge capacity growing faster than disk drive capacity.
- Tape pricing per gigabyte now diverging from disk pricing.

Source: Horison Information Strategies

Replication options

Replication refers to the process of creating multiple copies of data at either local and/or remote locations. The copies may or may not be identical based on the type of replication used. There are several techniques available to implement replication.

1) Backup/restore is the most traditional disaster recovery method of writing data, usually a complete file or full volume, from primary disk to either disk or tape for backup and from tape or disk back to primary disk for recovery. These are sometimes referred to as D2D2T for disk-to-disk-to-tape, D2D for disk-to-disk or D2T for disk-to-tape. In most cases, traditional backup causes the application being backed up to be impacted or even stop. Data can be backed up locally or remotely. Some more advanced application-specific backup modules allow open files and databases to be backed up non-disruptively without stopping or interrupting the application, though write operations may be suspended. *Tradeoffs exist* when choosing an effective backup strategy.

Backing up full disk volumes or files can become very time consuming and may be difficult to schedule. In addition to full backups, incremental and **differential backups** represent further options. In a differential backup, the same data that was backed up on the previous differential backup is also backed up on the next differential backup. That's why differentials often grow in size each day between full backups. This means that daily backups gradually get larger, but the restore time is minimized compared to full or incremental backups. A full restore requires only the last full backup and the last differential copy.

For **incremental backups**, only the data that has changed since the last incremental backup is backed up. This minimizes the amount of data backed up and, therefore, reduces the time needed for the "backup window" making it different from differential backup. A full restore, however, takes longer as each incremental backup will have to be restored to get all files to their last known state. This is generally a more complex process. Often a full backup will be performed weekly while an incremental backup is performed daily. Incremental backup minimizes the backup time, a differential backup minimizes the restore time and the specific application may require one or the other. Often, storage administrators want to minimize both the backup and recovery process complicating matters.

2) Mirroring is implemented as a block-for-block replica of a file, a logical unit or a physical disk volume normally using disks for all copies. Once the mirrored data element is established by copying the original data element, the mirror is maintained by replicating all write operations in two (or more) places creating identical copies. The choices for mirroring increase as the storage administrator must choose to implement asynchronous or synchronous mirroring and *tradeoffs exist for each case*.

Synchronous mirroring is frequently used in z/OS (mainframe) environments given the critical nature of mainframe applications. In synchronous mirroring, both the source and the target devices must acknowledge that the write is completed before the next write can occur. This degrades application performance but keeps the mirrored elements synchronized as true mirror images of each other. For asynchronous mirroring, the source and target devices do not have to synchronize their writes, and the second and subsequent writes occur independently. Therefore, asynchronous mirroring is faster than synchronous mirroring, but the secondary copies are slightly out-of-synch with the primary copy. This is sometimes referred to as a *fuzzy copy*. Asynchronous mirroring is often used with an IP storage protocol to replicate data to locations hundreds of miles away. In reality, the secondary data element is usually no more than one minute behind or out-of-synch with the primary copy. This can be a significant exposure for write-intensive applications.

Mirroring is used for many mission-critical applications and it is the fastest way to recover data from a device or subsystem failure since restore operations can occur in no more than a few seconds by switching over to a mirrored copy. Mirroring **does not help** protect against a

data corruption problem (hacker, worm, virus, intrusion or software error) as it produces two or more copies of corrupted data. For best practices, mirroring should always be accompanied by point-in-time (PIT) copies of data. This will permit a restore to occur from clean data that existed before the corruption occurred. Mirroring is defined and also commonly referred to as RAID 1.

3) Snapshot copy is another high-availability feature for disk data. It has gained popularity as it provides a less expensive means for a business to experience some of the benefits of disk mirroring. **With snapshot copy, only one complete copy of the data exists at a time**. When using snapshot copy and write operations occur, the changed areas (writes) are saved in a separate area or partition of disk storage specifically reserved for snapshot activity. Here the old value of the affected area or block can be saved in case the new block(s) is corrupted or to permit a fuzzy data image that can be used for a non-disruptive backup. *Tradeoffs exist here also*. Every change to the primary copy of data generates additional write operations to the area on disk storage designated to contain the snapshots. This activity adds I/O overhead and increases disk storage consumption. The storage administrator also must manage the number and currency of snapshots. Snapshots provide data protection from intrusion and data corruption but not from a device failure.

4) Point-in-time (PIT) copy is the fourth type of replication technology and has been a popular replication method for years. PIT copy provides a view of data at a specific point-in-time. It enables a mode of providing data protection at specified times. Like a series of still images and slightly different from general backup/restore, PIT copies are complete data images and are taken at *specified* points in time. PIT copies generate more storage consumption and an additional I/O workload as frequency increases. PIT copies enable an administrator to go back in time to restore data from a stable state prior to when a corruption or other disruption occurred. Gaining popularity, this is the best method to protect from human errors, software problems, viruses, intrusions and data corruption, and should accompany any mirroring implementation.

Again, *tradeoffs exist*. The more frequently the PIT copy is taken, the more storage is required and the more time it takes to determine which copy is the correct one to restore from. Current pricing economics significantly favor using automated tape over disk for this data protection function, and the savings grow as the storage pool gets larger.

5) Journaling is another method to enable data recovery where every write and update operation is continuously written to another device that may or may not be the same as the primary device. Unlike mirroring, however, the secondary copy is a sequential history of write events. All write operations are queued to the secondary device, or the journal device, which may be disk or tape. Usually an asynchronous approach regarding write operations, journaling uses metadata to associate the write operation with the location in primary storage where the data belongs. *Tradeoffs exist in journaling*. In a typical IT environment, on average, 20 percent of all I/Os are writes. Journaling can reconstruct a full volume if the journal was initially created as a mirrored copy of the primary volume. If the journal copy was not created as a mirrored copy, journaling is not a substitute for a physical device failure as the primary copy is no longer available. Journals are typically kept as a continuous history for two to four days covering the period of maximum likelihood for a data-recovery action to occur. Journals also speed up the recovery process and reduce the backup window. Journals are especially good for protecting from intrusion and data corruption, enabling restores to go back in time to a point before the corruption occurred. Normally, journaling accompanies one of the above replication strategies to build a complete recovery strategy for a file, data set or database.

The Cost of Downtime

Leading causes of computer downtime for more than 12 hours

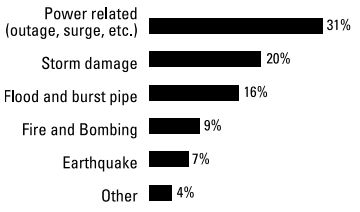

Source: Contingency Planning Research, a division of Eagle Rock Alliance

Cause of Data Loss

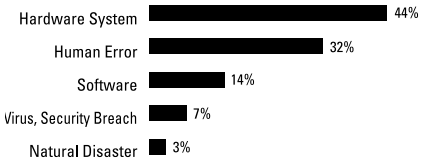

Source: Horison Information Strategies

Effect of Outage

IT outages take their toll

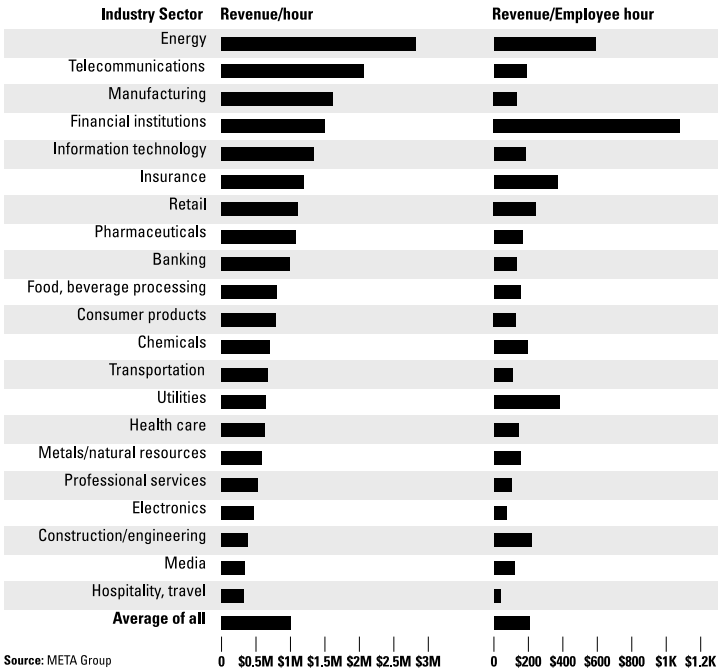

Source: META Group

Where to replicate?

Replication can be implemented at the server, the switch or network, or in the storage subsystem. *Again, there are choices and tradeoffs exist*. For server replication, the replication function runs on the server(s) and the servers almost always must run the same operating system. Data is usually transmitted between servers using an IP network. The storage devices do not need to be the same type on each server, however, as host computer resources are consumed. Switch- or network-based replication supports both heterogeneous servers and storage devices and doesn't consume host resources. This may be the most costly approach as proprietary hardware and software are normally required. Storage subsystem-based replication requires the devices to be the same or from the same vendor, but the switch and server can be from different vendors. This is the most proprietary implementation of replication and is just now beginning to appear with the new intelligent switches.

Conclusion

If you think choosing the optimal data protection strategy is confusing and that there are too many choices, you are definitely not alone. Why does a high-availability solution for disk require someone to select from full, incremental, differential, point-in-time, synchronous or asynchronous mirrors, snapshot copies, or journals running on either the host, the switch or in the storage subsystem? Why does just using tape for backup often seem so simple by comparison? Storage management for disk has become complex and businesses often choose the simplest approach after sustaining three years of downsizing and cutbacks. The simplest approach, however, may provide the highest availability. Today's IT environments are demanding a more comprehensive strategy for business continuance and high-availability than ever before. Replication options, including mirroring, along with a variety of point-in-time copies that match the applications' specific requirements offer the highest probability of success. The replication technologies are now in place to deliver high-availability data protection enabling a business to choose the one that best suits its needs. As data becomes more valuable every day, implement some type of data protection strategy. Doing nothing is a strategy, just not a very good one.

Horison Information Strategies
100 Arapahoe Ave., Suite 14
Boulder, Colorado 80302
303.417.9455 phone
303.939.9159 fax
fmoore@horison.com
www.horison.com

JT0007B