3.3 Understanding Disk Fault Tolerance

Windows May 15th, 2007

Fault tolerance refers to the capability of a computer or network to continue to function when some component fails. Disk fault tolerance refers to methods of storing data on the disk in such as way as to create redundancy of the data, so that it can be retrieved or recreated if a disk fails. Fault tolerance is not a substitute for backing up your data, but should be used in conjunction with a regular backup schedule that includes offsite storage. Generally, a fault tolerance solution will enable you to get up and running again more quickly than if you have to restore from backup, but backups are another line of defense in case the entire computer fails or is destroyed by fire, flood, etc.

There are several different ways to achieve disk fault tolerance. The most common implementation is known as RAID, or Redundant Array of Independent (or Inexpensive) Disks. Multiple disks can be configured in a number of different ways to create a fault-tolerant array. Data can simply be mirrored from one disk to another, or parity information can be stored that will enable the regeneration of lost data. RAID can be implemented either as a hardware or software solution. There are many different "levels" of RAID: 0, 1, 2, 3, 4, 5, 6, 7, 10, 0+1, and 53 are the most common. Some of these can be implemented only via the hardware. For more information about the different levels of RAID, see the RAID.edu Web site at www.acnc.com/04_01_00.html.

Windows Server 2003 has built-in support for three levels of software-implemented RAID:

- level 0 (disk striping, no parity)
- level 1 (disk mirroring)
- level 5 (striping with parity)

The biggest advantage of hardware RAID is performance; disk access is faster because you don't have the operating system overhead (the RAID disks appear as one to the operating system). The big advantage of software RAID is cost; you don't have to buy extra expensive RAID controllers or other additional hardware to use it.

Reference:

http://www.linglom.com/2007/05/15/understanding-disk-fault-tolerance/

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

RAID is commonly implemented for both performance and fault tolerance. With RAID, you can choose to assemble disks to provide fault tolerance, performance, or both, depending on the RAID level that you configure. Table 1.4 summarizes commonly available RAID levels.

Table 1.4 RAID Comparison

RAID Level	Description	Minimum Disks Required	Effective Capacity
0	Disk Striping. Two or more disks appear to the operating system as a single disk. Data is striped across each disk during read/write operations. Potentially increases disk access speeds 2X or better. Not fault tolerant.	2	S*N *N* = # of disks in array *S* = Size of smallest disk in array
1	Disk mirroring. Data is mirrored on two or more disks. Provides fault tolerance, but at a higher cost (space required is double the amount of data). Read performance is increased as well.	2	S *S* = Size of smallest disk in array
0+1	Combines RAID 0 and RAID 1; offers the performance of RAID 0 and the protection of RAID 1.	4	S*N/M *N* = # of disks in array *S* = Size of smallest disk in array *M* = Number of mirror sets
5	Disk striping with parity. Provides slower performance than RAID 0, but provides fault tolerance. A single disk can be lost without any data loss. Parity bits are distributed across all disks in the array.	3	S*(N-1) *N* = # of disks in array *S* = Size of smallest disk in array

From a design perspective, your choice of a RAID solution should be dictated by the type of data being stored. Although RAID 0 offers the fastest read and write performance, it does not offer any fault tolerance, so that if a single disk in a RAID 0 array is lost, all data is lost and will need to be recovered from backup. This might be a good choice for high performance workstations, but might not be suited to mission-critical servers.

RAID 1 allows you to configure two or more disks to mirror each other. This configuration produces slow writes, but relatively quick reads, and provides a means to maintain high data availability on servers, because a single disk can be lost without any loss of data. When more than two disks make up the mirror, the RAID 1 array can lose multiple disks so long as a complete mirrored pair is not lost. When planning a RAID 1 solution, remember that the amount of physical disk space required is twice the space required to store the data.

RAID 0+1 combines the performance benefit of striping with the fault tolerance of mirroring. Compared to RAID 0, writes are slower, but reads are equally fast. Compared to RAID 1, RAID

0+1 offers faster writes and reads but also requires additional storage to create the mirrored stripe sets. This configuration is often ideal for mission-critical database storage, because it offers both fast read access and fault tolerance.

RAID 5 provides fault tolerance: you can lose a single disk in an array with no loss of data. However, RAID 5 operates much more slowly than RAID 0 because a parity bit must be calculated for all write operations. RAID-5 volumes are well suited for reads and also work well in the following situations:

- In large query or database mining applications where reads occur much more often than writes. Performance degrades as the percentage of write operations increases. Database applications that read randomly work well with the built-in load balancing of a RAID 5 volume.

- Where a high degree of fault tolerance is required without the cost of the additional disk space needed for a RAID 1 volume. A RAID 5 volume is significantly more efficient than a mirrored volume when larger numbers of disks are used. The space required for storing the parity information is equivalent to 1/Number of disks, so a 10-disk array uses 1/10 of its capacity for parity information. The disk space that is used for parity decreases as the number of disks in the array increases.

3.4 Developing a Disaster Recovery Plan

- One way a company can prepare itself to face natural, infrastructure or operational disasters is to develop a well throughout and comprehensive disaster- recovery plan (DRP) that specifies how a company will maintain information system and services when a disaster strikes.
- The plan should clearly list and specify the specific situations that warrant the declaration of a disaster.
- Also should identify the specific course of action that employees must take when a disaster strikes.
- All critical system elements must be included in the DRP.

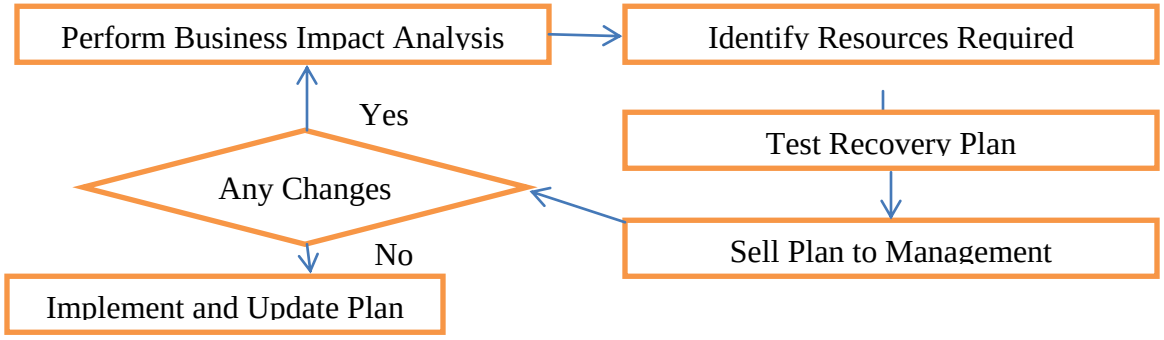

- The first is to conduct a thorough analysis of the impact of a disaster on the business called a **business impact analysis** or **business resumption plan**.
- The more detailed the business impact analysis is the smaller is the change of overlooking critical elements during a disaster.
- It should clearly identify the situations that qualify as disasters.
- It should specifically name individuals who have the right and the responsibility to declare a disaster, identify the specific process by which a disaster will be declared.
- Identify and prioritize the critical processes functions and resources of the organization which will be included in the DRP.
- It is driven from key question “What would we do if that key piece of technology were not there?”
- Three steps to develop a business resumption plan
    - Take an inventory of all corporate assets functions and resources that are essential to operating the business and prioritize them.
    - Identify the situations that qualify as disasters and the individuals responsible for declaring a disaster.
    - Develop the course of action that each employee in the organization must take to make the company operational after a disaster strikes
- The next step is to **identify the resources** required to recover from the disaster.
    - It includes money, time, personnel and facilities that should be specifically dedicated to disaster recovery.
- The third step is to test the recovery plan to make sure that it works and that no critical elements have been forgotten.
    - During a disaster recovery test, every step in the plan must be practiced several times.
    - The plan itself must be tested several times each year so that all employees know what to do if a disaster strikes.

  - The plan should also be updated frequently to include new assets and new risks that increase the vulnerability of the firm.
  - A DRP is good only if it is well tested well understood and well-remembered.
- The next step is to **identify the alternative for implementing the plan.**
  - For example, a company can use the service of a disaster service bureau that provides temporary disaster recovery service such as immediate access to time sharing environment access to utilities and temporary personnel.
  - Another alternative is to have reciprocal arrangements with other companies that use similar equipment and applications so that in case of an emergency one company can use the facility of the other.
  - Another alternative is outsourcing
  - Similarly another alternative is backup site.

Planning a Server Disaster Recovery

Types of disaster - do you have them all covered?

What would you do in the event of each of the following? Which order would you put them in? Most likely =1 least likely = 6.

1. Power failure
2. Hackers and security breaches
3. Stolen Kit, crime and vandalism.
4. Fire, storm, flooding, earthquake, which is most likely in your area?
5. Terrorist attacks, chemical attack.
6. Beware staff leaving - I was called in to help one company because no-one new how the system worked or even where the servers were! A disaster caused by an outsourcing deal that went bad.

Identify the most likely cause in for your situation. Eliminate two areas as extremely unlikely. Are you resources deployed according to your priorities?

Disaster Recovery - Planning

Plan to identify, then eliminate, single points of failures. Make sure you have duplicate systems for both hardware and software. Will you need replica servers (windows and Linux clustors) or even a whole mirror site in another location? At the very least store copies of your backup tapes offsite.

Define a strategy for each system. Windows Server 2003 has its own recovery tools, for example, system state. Exchange and SQL have their own specialist database recovery utilities. Failover clustering is great preventative measure. Similarly Linux also have recovery tools.

When it comes to a restore list the service dependencies and then sequence your recovery process. For example, operating system first, SQL program, finally database store.

I find targets are both measurable and motivating. Set targets for availability 99.9 or 99.99. Set timings for recovery. 2hrs for a server, 24 hrs for a site.

Consider the effect on your users and the effect on your customers. If your database goes down customers cannot order, but internal users can still use their workstations. If a virus cripples the email server users may grind to a halt but customers can still keep ordering.

Get executive enthusiasm. Lobby for a champion particularly when it comes to financing your disaster recovery plans.

Reference:

http://www.computerperformance.co.uk/w2k3/disaster_recovery_home.htm

3.5 Develop a backup plan

Hardware failure, hardware theft, virus attacks, or even operator error can result in lost data. Depending on the amount and type of data, your company's bottom line could even be affected, which makes it imperative that you have an effective backup strategy that you can stick to.

Different backup tools are available in market to backup data in servers. We can use dupm/restore, rsync, Amanda backup and many more open source tools in Linux and other unix based system to backup data and restore. Similarly Windows provides NT Backup utilities and Windows server Backup (in latest versions only). Finally there is specialized Backup server available in market which gives best performance and security like LaCie 5big Backup Server.

Factors that should be considered while backup are:

- What filesare needed to be backuped up?
- Where are these files?
- Who will backup the files
- Where, when and under what conditions should backup be performed?
- How often do these files changes?
- How quickly does an important missing or damaged files need to restored?
- Where will the data be restored?

Backup Strategies

- The simplest and most thorough backup scheme is to **copy all the files** on a system to backup disk or tape.
- A **full backup** does just that including every file within a designated set of files often defined as those on a single computer system or partition.
    - Full backup are time consuming and space consuming
    - Restoring a single file from a large backup spanning multiple tapes or disk is often inconvenient and time consuming.
    - When files are not changing very often, the time taken to complete a full backup may not be justified.
- **Incremental backup** usually done more frequently
    - In it, the system copies only those files that have been changed since some previous backup.
    - It is used when full backups are large and only minorities of the data changes within the course of say a day.
    - In such case backing up only the changed files saves a noticeable amount of space and time over performing a full backup.
- **Difference Backup**
    - In it the system copies only those files that have been changed since 1st backup instance and latest. Between 1st and latest backup instances there may exist many incremental instances.
- Unattended Backup is often used while backup is down in case of servers. Backup script is run with job scheduler (like cron in case of *nix system)

- In real world scenario, most of the business houses use both of Full backup and Incremental backup. For exam they can used incremental backup daily and 1 full backup per week to be safe in case incremental backup fails some how.

Backup Storage Considerations

- Properly storing the backup tapes, diskettes or other media once you have written the in an important part of any backup plan.
- Here are some things to keep in mind while deciding where to store the one for you system.
    - Know where things are
    - Make routing restorations easy
    - Write protect backup media
    - Environmental considerations
    - Handle media properly
    - Take security into account.

Creating a Data Backup Plan

- Determine what data to backup (see factors to consider)
- How often to backup (use of backup strategies)
- What medium to use (Storage conditions)
- Where to store data for safe keeping and how to label disk so that it case be easily found when and where needed.
- The data backup test should be done