



Data Storage and Backup

Sanjay Goel

School of Business

University at Albany, SUNY

Data Backup

Data Backup

Why?

- Files can be accidentally deleted
- Mission-critical data can become corrupt.
- Natural disasters can leave your office in ruin.
- Backup is the best insurance against disasters
 - A backup is the most cost effective technique for managing disasters
 - You need to figure out the backup strategy that suits the organization

Data Backup

Planning

- Planning involves
 - Figuring out what data needs to be backed up
 - How often the data should be backed up
- Boils down to risk analysis
- There are several pertinent questions that need to be answered
 - How important is the data on your systems?
 - What type of information does the data contain?
 - How often does the data change?
 - How quickly do you need to recover the data?
 - Do you have the right equipment for backups?
 - Who will be responsible for backup and recovery plan?
 - What is the best time to schedule a backup?
 - Do you need to store information off-site?

Data Backup

Options

Type	Description	Pros	Cons
Full Backup	A complete set of all files backed up.	Provides a complete copy of all your data; makes it easy to locate files for restoring.	Takes a long time and the most space on backup media; redundant backups created, as most files remain static.
Incremental Backup	A backup of files changed since the last backup of any type.	Uses the least time & space as only files changed since the last backup are copied; lets you back up multiple versions of the same file.	Makes the job of restoring files difficult since the last full backup and subsequent incremental backups have to be reinstalled in correct order. Also makes it hard to locate a specific file in backup
Differential Backup	A backup of files changed since the last full backup.	Takes up less time and space than a full backup; provides for more efficient restoration than incremental backups.	Redundant information stored, as each backup stores much of the same information plus information added since the last full backup. Subsequent differential backups take increasingly longer as more files are changed.

Sanjay Goel, Information Technology Management Department,

School of Business, University at Albany, SUNY

Data Backup

Types

- With differential backups all the files that have changed since the last full backup are backed up (which means that the size of the differential backup grows over time).
- With incremental backups, only files that have changed since the most recent full or incremental backup are backed up (which means the size of the incremental backup is usually much smaller than a full backup).

Data Backup

Selecting Media and Devices

- Capacity: Amount of data that needs to be back up routinely
 - Can the backup hardware support the required load
- Reliability: Failure rate of hardware and media
 - Reliability needs to be balanced with cost and time.
- Extensibility: The extensibility of the backup solution.
 - The solution needs to be scalable to growing needs of organization
- Speed: The speed at which data can be backed up & recovered.
 - Need to balance cost of down operations versus cost of equipment
- Cost: The cost of the backup solution.
 - Does it fit into your budget?
 - Does it commensurate with the loss of data and services?

Data Backup

Selecting Media and Devices

Device	Media Capacity	Speed	Comments
3.5" Floppy Drive	1.44 MB Removable Media	Slow	Nice for small amounts of data. Cheap and portable media.
CD-R/W	Up to 700 MB Removable Media	Moderate	Great backup device and wonderful for making your own music CDs too. Large backups will require multiple CDs.
DVD-R/W	Up to 4.7 GB Removable Media	Moderate to Fast	Similar to CD-R/W with greater storage space.
Hard Drive (Primary)	Up to 160 GB and growing.	Fast	Backup on current hard drive. Good for recovering files but insufficient against system failures.
Hard Drive (Alternate)	up to 160 GB and growing.	Fast	New hard drives are cheap and somewhat easy to install. External USB drives very efficient
ZIP® Drive	100 MB or 250 MB	Slow	This is a floppy on Steroids. The most popular high-capacity floppy-disk type device.
Tape Drive	4GB to 110 GB	Fast	A great high-capacity removable media. Generally used by more sophisticated users.
Internet Backup	Unlimited storage	Moderate	Depends mostly on internet connection speed. No devices to mess with. Data is off-site.
Printer	Unlimited pages	Very Slow	A paper backup is sometimes very effective

Data Backup

On-site vs. Off-site

ON-SITE

- Advantages:
 - Its under your control
 - You know what is happening with the media
 - You know it is working
 - Good for large data files
- Disadvantages
 - Fire means data loss
 - If you have an accident while backing up data is lost
 - Restoring backups requires effort

OFF-SITE

- Advantages
 - Saves organization time
 - Service Level agreement ensures guaranteed backup
- Disadvantages
 - No control
 - You don't know what they do with your data
 - Large data stores can't backed-up
 - Network can be expensive

Data Backup

RAID

- RAID: Redundant Arrays of Inexpensive Disks
- Uses multiple hard drives to enhance I/O performance, reliability, and capacity
- Metric of performance: Mean Time Between Failure (MTBF)
- MTBF of the raid is the MTBF of a single drive divided by the total number of drives used in the RAID
 - Fault-tolerance may be increased by adding redundant disk arrays

Data Backup

RAID

- RAID encompasses any basic concepts that attempt to combine physical disk space for either reliability, capacity, or performance
- There is both hardware and software implementation of RAID
 - Hardware RAID requires a RAID Controller
 - Software RAID requires CPU power to run

Data Backup

RAID Cont'd.

- A basic disk may be transformed into a dynamic disk using by creating volumes the Windows disk management options
- A volume is a storage unit made from free space on a disk
- Volumes can be formatted with a file system and assigned a drive letter
- There are five available types of volumes for dynamic disks: simple, spanned, mirror, striped, and RAID-5

Data Backup

RAID Cont'd.

- Simple
 - Uses free space on a single drive
 - Not fault tolerant
- Spanned
 - Uses free space on multiple drives (32 max)
 - Not fault tolerant
- Striped (RAID-0)
 - Data is allocated alternately and evenly across multiple physical disks
 - Not fault tolerant, cannot be mirrored

Data Backup

RAID Cont'd.

- Mirrored (RAID-1)
 - All of the data is redundantly duplicated on two disk arrays
 - Fault tolerant
- RAID-5
 - Data is striped across an array of three or more disks along with a value that can be used to reconstruct the array called parity
 - Failing disks can be reconstructed from parity and remaining data
 - Fault tolerant

Data Backup

Storage Area Network

- Storage Area Network (SAN): A SAN is a high level network connecting servers and storage devices for block level I/O.
 - Supports disk mirroring, backup and restore, archival and retrieval of archived data, data migration from one storage device to another, and the sharing of data among different servers in a network.
 - SANs can incorporate sub networks with network-attached storage (NAS) systems.
- NAS: Network Attached Storage is a hard disk set up with its own network address rather than being attached to a specific computer.

Data Backup

Hard Drive Protection

- Drive Fitness Test (DFT) uses a PC-based program that accesses special hard drive microcode, enabling users to monitor hard drive operation.
- Self Monitoring Analysis and Report Technology (SMART) is a monitoring system for computer hard disks to detect and report on various indicators of reliability, in the hope of anticipating failures
 - Useful for predictable failures when some failure modes, especially mechanical wear and aging, happen gradually over time.
 - Not very useful for unpredictable failures, such as an electronic component burning out.

Data Backup

Hard Disk Specifications

- Capacity
 - The whole capacity or capacity of one disk
- Rotate Speed
 - The speed a disk rotates in RPMs
- Average Seek Time
 - Measure of drive speed in multi-user environments where read and write request are uncorrelated
 - 10ms is common for hard drives
- Average Latency
 - The time it takes for the head of a hard drive to meet the correct sector of the drive once on the correct cylinder
 - Faster rotation speeds equals lower latency

Data Backup

Hard Disk Specification Cont'd.

- Average Access Time
 - Command Overhead Time + Seek Time + Latency
- Buffer Size (Cache)
 - A small, fast memory holding recently accessed data for quick access
 - Cache processors are much faster than main processors and stores information based on temporal and spatial locality
 - Temporal and spatial locality refer to when data is accessed and what data is in close proximity to that data

Data Backup

Hard Disk Specification Cont'd.

- Noise and Temperature
 - Usually derived from the motor
 - Lower temperatures mean happy and healthy hard drives

Disk Structure

Disk Structure

Physical Structure of Hard Disk

- Hard disks are made up of multiple platters spinning on a spindle that are read and written to by electromagnetic pins called heads
- Data on the platters are stored in circular bands and heads can read and write to a single band called a band track
- Sections within a track are called sectors

Disk Structure

Physical Structure of Hard Disk Cont'd.

- Platters rotate at a constant speed
- Tracks on the inside of the platter are moving faster than those at the edges
- To compensate for this data is dense on the inward tracks and sparse on the edges to create smooth read times
- One side of one platter will always be reserved for hardware track positioning information

Disk Structure

Logical Organization of Hard Disk

- A sector is typically 512 bytes
- A cluster is a space reserved for data and are typically the same size as a sector
- Since many files are larger than 512 bytes, not all data fits into a single sector
- When a file is written onto consecutive clusters, the clusters are called contiguous
- Contiguous clusters are read faster than fragmented clusters

Disk Structure

Logical Organization of Hard Disk

- Larger cluster sizes allow for less fragmentation, but increases the potential for more unused space within the clusters
- Reducing fragmentation decreases the amount of memory used to store the location of used and unused portions of the hard disk

Disk Structure

Hard Disk Interfaces: IDE/ATA

- IDE refers to Integrated Device Electronics, but ATA, Advance Technology Attachment, is the real industry standard name
- Most PCs have two IDE controllers that may support two devices for a maximum of four hard drives
- PATA or parallel ATA uses ribbon cables to connect hard drives to the motherboard
- When two hard drives are connected to the same IDE controller, one must be designated the master and one the slave through jumper cables

Disk Structure

Hard Disk Interfaces: SATA

- Is replacing the more common PATA
- Replaces ribbon cables and master/slave designation with more airflow friendly cables
- Differences in transfer rates are expected to hit 600MB/s in 2007 with SATA II
- Speed increases are held back by hard drive mechanic speed and the use of PATA controllers

Disk Structure

Hard Disk Interfaces: DMA

- Direct Memory Access is a function of the memory bus that allows for direct transfer between hard drives and memory through the IDE controller ignoring passage through the CPU
- Bus Mastering DMA ignores the IDE controller all together and makes direct transfers between hard disks and memory

Disk Structure

Hard Disk Interfaces: USB

- Universal Serial Bus is a hardware bus that is supported on many motherboards
- Allows many devices to connect to the bus
- USB transfer rates are about 12 Mbits/s while USB 2.0 rates are 480 Mbits/s
- Firewire is a lesser known, but common alternative to USB that is better than standard USB
- USB 2.0 leveled the disparities specifically by increasing bandwidth amongst other things

Disk Structure

Hard Disk Interfaces: SCSI

- Small Computer System Interface that is less common than ATA on PCs but more prominent on servers
- SCSI tends to be faster, more reliable, and more expensive than ATA
- Another advantage is that it can handle seven devices to ATA's two

Disk Structure

Primary Formatting of Hard Disk

- Before restoring data, hard disk logical structures must be set up through low level formatting, partitioning, then high level formatting
- The disk is divided into MBR, DBR, DIR, FAT, and DATA

Disk Structure

Low Level Format (LLF)

- Functions
 - Test hard disk media
 - Partition tracks for hard disk
 - Arrange sectors on track by interleave
 - Assign sector IDs and finish setting sectors
 - Test the hard disk surface for damaged sectors and mark them “bad”
 - Write certain ASCII to each sector
- Techniques
 - CMOS, disk tools, and debug programs on older systems
 - Hard disk manufacturers now provide tools to handle the tasks

Disk Structure

High Level Format (HLF)

- After the low level format, logical drives are created
- Drives are commonly named after alphabet letters such as C: or D:
- Trying to access the drives will currently result in a DISK MEDIA ERROR because they are empty
- To use them a file system must be created
- A high level format of DOS logic disk can be initiated by the “format” command

Disk Structure

High Level Format Cont'd.

- Functions
 - Assign local serial numbers to sectors from cylinder that assigned by each logical drive
 - Establish DBR in basic partition and load 3 system files of DOS if there is “/S” parameter in command
 - Establish file allocation table (FAT) in each logical disk
 - Establish File Directory Table in each logical disk

Disk Structure

High Level Format cont'd.

- When using the “format” command:
 - Activate the DOS partition by “Format C:/s”
 - Format other logical disks by “Format [d:]”
 - Note that formatting a disk will lose all information stored on the disk
 - For the using disk without adjusting the partition, also may carry on the fast format command “Format C:/Q”

Disk Structure

High Level Format: Windows

- Explore windows will show you different partitions in different colors
- Click right key in partition you wish to format
- Select the type of format you wish to execute
- Types range from format, fast format, complete format, etc.

Disk Structure

High Level Format: Partition Magic

- The program will show you different partitions in different colors
- Click right key in partition you wish to format
- Choose “format”
- Confirm and acknowledge that formatting will erase any existing data on your hard drive

Disk Structure

Data Storage Region of Hard Disk

- Hard disks are divided into divided into MBR, DBR, DIR, FAT, and DATA
- MBR is created by the partition software
- DBR, DIR, FAT, and DATA are created by the high level format
- The file system writes in data by rewriting FAT, DIR, and DATA areas

Disk Structure

MBR

- The first physical sector of the first hard drive (cylinder 0, head 0, sector 1)
- Each hard drive has an MBR but not every BIOS can start the running OS from every drive
- MBR is then loaded to a fixed point in memory where it loads the OS

Disk Structure

DBR

- DOS Boot Record (cylinder 0, column 1, sector 1)
- First sector an OS visits
- Contains a boot program and a BIOS Parameter Block (BPB)
- Boot program determines if the first two files in the root directory of this partition are the root files for the OS

Disk Structure

FAT

- File Allocation Table
- File system for MS-DOS
- Nearly universal OS support
- Reading and writing slow relating to fragmentation on creation and deletion
- The numbers after FAT indicate the number of cluster bits (FAT12, FAT16, FAT32, etc.)

Disk Structure

FAT Cont'd.

- Directory or File Directory Table (FDT)
- The root sector after a backup FAT
- Records each start cell, files
- OS can locate files on the outset of FAT and FAT

Disk Structure

Cont'd.

- Where files are stored
- Largest portion of hard disk space

Disk Structure

MBR

- Given no hardware damage, MBR recovery is the first step in partition recovery
- MBR may be recovered using the “Fdisk” command, Fixmbr from Microsoft, and other similar programs

Disk Structure

Partition Recovery

- In minor cases, the partition can be restored automatically
- In other cases, it must be built up manually using tools such as Norton Utilities 8.0, DiskMan, PartitionMagic, or Partition Table Doctor

Disk Structure

DBR Recovery

- Partition OS cannot be booted if the DBR is damaged
- Functions of DBR are different for FAT and NTFS partitions
- Formatting will restore DBR, but not data
- Partition Table Doctor and WinHex are examples of programs that can help recover DBR

Disk Structure

FAT Recovery

- If FAT1 is damaged and FAT2 is not, FAT2 may be used to cover FAT1
- This includes finding the start sector of FAT2 and finding the total length of the FAT table using programs like DiskEdit and WinHex

Summary

Acknowledgements

- I would like to thank Daren Pon for helping prepare these slides for this presentation
- Materials for the lecture have been taken from several sources including
 - Helix and Knoppix web sites.
 - Book on Data Recovery