Gamze Zeytinci
CSIS-550 History of Computing
Spring-2001

Evolution of the Major Computer Storage Devices

From Early Mechanical Systems to Optical Storage Technology

CONTENTS

Gamze Zeytinci
CSIS-550 History of Computing
Spring-2001

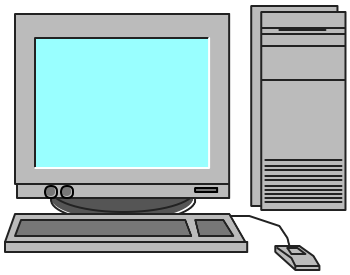

Evolution of the Major Computer Storage Devices

From Early Mechanical Systems to Optical Storage Technology

Density of data stored on a magnetic hard disk increased 1.2-million-fold in the four decades after IBM's introduction of the first commercial disk drive in 1957. Improvements in the miniaturization have been the primary catalyst for this spectacular growth. Because of this, computers are no longer simply relegated to the desktop. They are in our cars, our TVs, VCRs, Stereos and toasters. Increasingly, we are doing business and accomplishing everyday tasks over vast computer networks. Our world is changing from the physical to the digital. This transformation is no small task and the transition from the present world to the digital one is highly dependent on smart, inexpensive and abundant digital storage [.1.].

Introduction

Since the inception of the modern computers in the late 1940s, computing systems have continually grown in complexity, both in hardware and in the software associated systems. This complexity is due to two factors: first, the tradeoffs in cost/performance versus size of various memory–storage systems, and second the way users organize and use data.

Modern computing systems consist of four hardware components: the central processing unit (CPU), the main memory, the secondary storage (disks, tapes, CDs, etc.) and the input-output devices that interfaces with humans. Since the beginning of the modern computer era in the late 1940s and early 1950s, the need for larger storage capabilities has become increasingly apparent. The need for larger storage is due mainly to the vast amount of digital data including graphic, audio and video media.

All computations, either mental, mechanical, or electronic require a storage system of some kind, whether the numbers be written on paper, remembered in our brain, counted on the mechanical devices of a gear, punched as holes in paper, or translated into electronic circuitry.

This study reviews the evolution of the storage systems from early mechanical systems to optical storage technology. Because of the importance of the storage systems in computing, evolution of the storage systems is directly related the CSIS 550-History of Computing.

Floppy Disk Drives / Floppy Disks

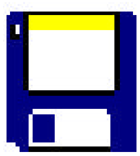

Floppy disk drives were originally introduced commercially as a read-only device in the early 1970s. These early floppy disk drives were used to hold microcode and diagnostics for large IBM mainframe computer systems. These disk drives were using 8-inch floppy diskettes recorded on only one side. By changing these diskettes inside the floppy drive, technicians could easily update the microcode to the latest revisions or load diagnostics easily. The storage capacity of these early read-only drives was less than 100 kilobytes.

In 1973 a new upgraded 8 inch drive with read/write capability and a capacity of about 250 kilobytes began shipping which IBM used in data entry systems. This drive incorporated many technical improvements and became a model for drives still in use today. As time went on, designers learned how to reliably record on both sides of the diskette as well as increase the density of the data recorded on the diskette.

In 1976 smaller 5.25 inch size floppy drives were introduced by Shugart Associates. In a cooperative effort, Dysan Corporation manufactured the matching 5.25 inch flexible floppy diskettes. Originally these drives were available in only a single-sided low density format, and like the first 8 inch models, stored less than 100 kilobytes. Later they received many of the same improvements made to the 8 inch models, and eventually 5.25 inch floppy drives settled at a double-sided, "double density" formatted capacity of about 1.2 megabytes. This drive was used in the IBM-AT personal computer.

Modern floppy drives and diskettes (3.5 inch) have evolved to a much smaller size with larger capacities as well. In 1980, the 3.5 inch floppy drive and diskette was introduced by Sony. During the early 1980's many competing formats were tried to compete with the 3.5 inch drives. Over time the industry settled on the 3.5 inch format which was standardized and manufactured by many companies. Today's standard 3.5 inch diskettes hold a formatted capacity of 1.44 megabytes while still using the same basic technology of the second generation 8 inch drives.

The primary factor that caused engineers to reduce the size and cost of floppies was the introduction and evolution of the personal computer. It was in the personal computer market that the low cost, mass produced floppy drive found its first real home. Very quickly, the floppy became the standard method of exchanging data between personal computers. It also became the popular method of storing moderate amounts of information outside of the computer's hard drive. Floppy diskettes are small, inexpensive, readily available, easy to store, and have a good shelf life if stored properly. It is a round, flat piece of Mylar coated with ferric oxide, a rustlike substance containing tiny particles capable of holding a magnetic field, and encased in a protective plastic cover, the disk jacket. Data is stored on a floppy disk by the disk drive's read/write head, which alters the magnetic orientation of the particles. Orientation in one direction represents binary 1; orientation in the other, binary 0.

Hard-Disk Drives (Hard Disks)

The hard drive, also called the hard disk or fixed disk, is the primary storage unit of the computer. It is always labeled the C drive. Additional drives are labeled after it as the D, E, F, etc. It has several read/write heads that read and record data magnetically on platters, a stack of rotating disks inside the hard drive. Hard drive is important for the following reasons:

- It stores programs that must be installed to the hard drive before they can be used.
- It stores data files that can be accessed later.

- It organizes files like a file cabinet so they can be accessed more easily.

The Hard Drive can store a large amount of computer data on it. Many advancements have made it possible to store a large amounts of data in a small space. The hard drive's speed is discussed in terms of access time. This is the speed at which the hard drive finds data. The average access time is measured in milliseconds. One millisecond equals 1/1000 of a second. The average drives had 9 to 14 ms access time. The lower the access time the faster the hard drive. The capacity, or amount of information that a hard drive can store, is measured in bytes. Today many computers come with 20-80 GB (Giga Byte=1,000 Mega Bytes) hard drives.

Types of Hard Drives

(1) Bournolli has a mechanism that acts as a fly wheel. When the HD spins the disk follows gravities course and is lifted up a few centimeters causing the Read/Write heads to touch. When it stops spinning, the HD will fall back down the few centimeters.

(2) Magnetic Optical Drive has the best storage capacity. It has the same principles as the Bournolli, but is mixed with the Compact Disc technology so that the Read/Write head puts data in order and the laser reads off of it.

(3) Standard Magnetic Drive The standard magnetic drive is less complex and less expensive than the others. It stores data with a read/write head which sends a pulse of electricity through causing the magnetic films electrons to line up in a certain way.

History of Hard-Disk Drives

The hard-disk drive is a wonder of modern technology, consisting of a stack of disk platters. each one an aluminum alloy or glass substrate coated with a magnetic material and protective layers. Read-write heads, typically located on both sides of each platter, record and retrieve data from circumferential tracks on the magnetic medium. Servomechanical actuator arms position the heads precisely above the tracks, and a hydrodynamic air bearing is used to "fly" the heads above the surface at heights measured in fractions of microinches. A spindle motor rotates the stack at speeds of between 3,600 and 10,000 revolutions per minute.

This basic design traces its origins to the first hard-disk drive-the Random Access Method of Accounting and Control (RAA/IAC)-which IBM introduced in 1956. The RAMAC drive stored data on 50 aluminum platters, each of which was 24 inches in diameter and coated on both sides with magnetic iron oxide. (The coating was derived from the primer used to paint San Francisco's Golden Gate Bridge.) Capable of storing up to five million characters, RAMAC weighed nearly a ton and occupied the same floor space as two modern refrigerators.

In the more than four decades since then, various innovations have led to dramatic increases in storage capacity and equally substantial decreases in the physical dimensions of the drives themselves. Storage capacity has jumped multiple orders of magnitude during that time, with the result that some of today's desktop PCs have disk drives containing more than 80 gigabytes. According to some analyists the industry has achieved these improvements largely through straightforward miniaturization. "Smaller heads, thinner disks, smaller fly heights (the distance between head and platter.)

Growth Rate of Hard-Disk Drives

Many corporations find that the volume of data generated by their computers doubles every year. Extremely large databases containing more than a terabyte --that is, one trillion bytes--are becoming the norm as companies begin to keep more and more of their data on-line, stored on hard-disk drives, where the information can be accessed readily. The benefits of doing so are numerous: with the right software tools to retrieve and analyze the data, companies can quickly identify market trends, provide better customer service, modify manufacturing processes, and so on. Meanwhile individual consumers are using modestly priced PCs to handle a data glut of their own, storing countless emails, household accounting spreadsheets, digitized photographs, and software games.

All this has been enabled by the availability of inexpensive, high-capacity magnetic hard-disk drives. Improvement in the hard disk technology has been quite impressive. The capacity of hard-disk drives grew about 25 to 30 percent each year through the 1980s and accelerated to an average of 60 % in the 1990s. By the end of 1999 the annual increase had reached 130 percent. Today disk capacities are doubling every nine months, fast outpacing advances in computer chips, which obey Moore's Law (doubling every 18 months).

Price of Magnetic Hard-Disk Drives

The sales of hard-disk drives have soared as costs per megabyte have decreased. Sales revenues are expected to grow to $50 billion in 2002. At the same time, the cost of hard-disk drives has plummeted. Disk/Trend, a Mountain View, California-based market research firm that tracks the industry, reports that the average price per megabyte for hard-disk drives plunged from $11.54 in 1988 to $0.04 in 1998, and the estimate for last year is $0.02. Some experts predict that by 2002 the price will have fallen to $0.003 per megabyte.

This remarkable combination of rising capacity and declining price has resulted in a thriving hard-disk market. The industry shipped 145 million hard-disk drives in 1998 and nearly 170 million last year. That number is expected to surge to about 250 million in 2002, representing revenues of $50 billion, according to projections.

SPE barrier: Limit of Technology in Hard Disks

But whether the industry can maintain these impressive economics is highly questionable. In the quest to deliver hard disks with ever increasing capacities, IBM, Seagate Technology, Quantum Corporation and other manufacturers have continually crammed smaller and smaller bits together, which has made the data more susceptible to *Super-Paramagnetic-Effect,* or simply SPE. In the coming years the hard-disk technology could reach a limit imposed by this effect. Simply described, SPE is a physical phenomenon that occurs in data storage when the energy that holds the magnetic spin in the atoms making up a bit (either a 0 or 1) becomes comparable to the ambient thermal energy. When that happens, bits become subject to random "flipping" between O's and 1's, corrupting the information they represent.

As one shrinks the size of grains or crystals of magnetic material to make smaller bits, the grains can lose the ability to hold a magnetic field at a given temperature. "It really comes down to the thermal stability of the media," One researcher explains. "You can make heads more sensitive, but you ultimately need to consider the properties of the media material, such as the coercivity, or magnetic stability, and how few grains you can use to obtain the desired resistance to thermal erasure."

Traditionally, a minimum of about 500 to 1,000 grains of magnetic material was required to store a bit. In March 2001, IBM scientists announced a process for self-assembling magnetic particles into bits that could provide areal densities as high as 150 gigabits per square inch.) Currently researchers are actively looking for improved materials that can hold a detectable magnetic charge and resist SPE with fewer grains. Also, the industry has been developing better manufacturing processes to decrease the impurities in the storage medium and thereby enable smaller bits.

In lieu of improvements of this type, the limit of bits per inch will remain in the range of between 500,000 and 650,000, according to a storage technologist for Seagate Technology's research division. But this parameter, which is for data stored in a particular track on a platter, is only one determinant of areal density, which is the number of bits per square inch.

Strategies for Avoiding the SPE Barrier

The capacity-versus-performance debate could become acute as the industry considers various ways to avoid the SPE barrier. Experts agree that moving beyond areal densities of 150 gigabits per square inch will require a significant departure from conventional magnetic hard disks. Some of the alternatives boast impressive storage capabilities but mediocre speeds, which would limit their use for certain applications. At present, the main strategies include:

- Change the orientation of the bits on the disk from longitudinal (circumferential) to perpendicular, or vertical, to cram more of them together and to prevent them from flipping.

- Use magnetic materials, such as alloys of iron/platinum or cobalt/samarium, that are more resistant to SPE. If the magnetic "hardness" of the material is a problem for recording data, heat the medium first to "soften" it magnetically before writing on it.

- Use a radically different storage material, such as holographic crystals, phase-change metals, or plastic.

- Imprint a pattern lithographically onto the storage medium to build microscopic barriers between individual bits.

Although several of these approaches have attracted large investments from the leading manufacturers, most remain in the very early stages of testing. Some of the concepts await research breakthroughs or key advances in supporting technologies before work can begin in earnest on prototypes.

Head Improvements in Hard-Disks

Many of the past improvements in disk-drive capacity have been a result of advances in the read-write head, which records data by altering the magnetic polarities of tiny areas, called domains (each domain representing one bit), in the storage medium. To retrieve that information, the head is positioned so that the magnetic states of the domains produce an electrical signal that can be interpreted as a string of 0's and 1's.

Early products used heads made of ferrite, but beginning in 1979 silicon chip-building technology enabled the precise fabrication of thin-film heads. This new type of head was able to read and write bits in smaller domains. In the early 1990s thin-film heads themselves were displaced with the introduction of a revolutionary technology from IBM. The innovation, based on the magnetoresistive effect (first observed by Lord Kelvin in 1857), led to a major breakthrough in storage density.

Rather than reading the varying magnetic field in a disk directly, a magnetoresistive head looks for minute changes in the electrical resistance of the overlying read element, which is influenced by that magnetic field. The greater sensitivity that results allows data-storing domains to be shrunk further. Although manufacturers continued to sell thin-film heads through 1996, magnetoresistive drives have come to dominate the market.

In 1997 IBM introduced another innovation--the giant-magnetoresistive (GMR) head--in which magnetic and nonmagnetic materials are layered in the read head, roughly doubling or tripling its sensitivity. Layering materials with different quantum-mechanical properties enables developers to engineer a specific head with desired GMR capabilities. Currie Munce, director of storage systems and technology at the IBM Almaden Research Center in San Jose, says developments with this technology will enable disk drives to store data at a density exceeding 100 gigabits per square inch of platter space.

Interestingly, as recently as 1998 some experts thought that the SPE limit was 30 gigabits per square inch. Today no one seems to know for sure what the exact barrier is, but IBM's achievement has made some assert that the "density demon" lives somewhere past 150 gigabits per square inch.

Types of Hard Drive Connections

Basically there are three types of Hard Drive Connections:

(1) Integrated Device Electronics (IDE)

(2) Small Computer Interface (SCSI)

(3) Enhanced Integrated Drive Electronics (EIDE)

The Issue of Speed in Hard-Disk Drives

Storage capacity is not the only issue when dealing with hard disks. The rate with which data can be accessed is also becoming an important factor that may also determine the useful life span of magnetic disk-drive technology. Although the capacity of hard-disk drives is surging by 130 percent annually, access rates are increasing by a comparatively tame 40 percent.

To improve on this, manufacturers have been working to increase the rotational speed of drives. But as a disk spins more quickly, air turbulence and vibration can cause misregistration of the tracks, a problem that could be corrected by the addition of a secondary actuator for every head. Other possible enhancements include the use of fluid bearings in the motor to replace steel and ceramic ball bearings, which wear and emit noticeably audible noise when platters spin at speeds greater than 10,000 rev. per minute.

Many industry onlookers foresee a possible bifurcation in the marketplace, with some disk drives optimized for capacity and others for speed. The former might be used for mass storage, such as for backing up a company's historical files. The latter would be necessary for applications such as customer service, in which the fast retrieval of data is crucial.

In the past, customers typically preferred a bigger drive at the lowest possible cost, even if the product had slower performance.

But new applications are demanding faster drives. With electronic commerce over the World Wide Web, for example, companies need to store and retrieve customer data on the fly. In addition, businesses are deploying an increasing number of dedicated file servers for information that needs to be shared and accessed quickly by a number of employees.

Tracks in Hard Disks

Storage capacity also depends on the narrowness of the tracks, and so far manufacturers have been able to cram up to 20,000 tracks per inch. This number is limited by various factors, such as the ability of the recording head to resolve the different tracks and the accuracy of its position-sensing system. Squeezing in additional tracks will require significant improvements in several areas, including the design of the head and the actuator that controls that head. To achieve an overall density of I 00 gigabits per square inch, the industry must somehow figure out a way to fit about 150,000 tracks or more per inch.

With the existing technology, tracks must be separated by gaps of 90 to 100 manometers, according to analyists. Most write heads look like a horseshoe that extends across the width of a track. Recording is in the longitudinal direction [that is, along the circular track], but they also generate fringe fields that extend radially. If the tracks are spaced too closely, this effect can cause information on adjacent tracks to be overwritten and lost.

One solution is to fabricate the recording head more precisely to smaller dimensions. "You can use a focused ion beam to trim the write head and to narrow the width of the track that a writer writes," one researcher says. But the read head, which is a complex sandwich of elements, poses a harder manufacturing problem. Furthermore, for 150,000 tracks or more per inch to be squeezed in, the tracks will have to be less than about 170 manometers wide. Such microscopically narrow tracks will be difficult for the heads to follow, and thus each head will need a secondary actuator for precise positioning. (In current products, just one actuator controls the entire assembly of heads.)

Last, smaller bits in thinner tracks will generate weaker signals. To separate those signals from background noise, researchers need to develop new algorithms that can retrieve the information accurately. Today's software requires a signal-to-noise ratio of at least 20 decibels. According to some analyists, current industry is at least six decibels short of being able to work with the signal-to-noise ratio that would apply when dealing with the bit sizes entailed in disks with areal densities of 100 to 150 gigabits per square inch.

Nevertheless, such problems are well understood, many industry experts concur. In fact, analyists assert that the improvements in materials, fabrication techniques and signal processing already being studied at IBM and elsewhere will, over the next few years, enable the manufacture of disk drives with areal densities in the range of 100 to 150 gigabits per square inch.

The introduction of thin-film heads took nearly 10 years. The transition from that to magnetoresistive technology required six more years because of various technical demands, including separate read and write elements for the head, a manufacturing technique called sputter deposition and different servo controls.

But the switch to giant magnetoresistive drives is occurring much faster, taking just between 12 and 18 months. In fact, IBM and Toshiba began shipping such products before the rest of the industry had fully converted to magnetoresistive heads.

The quick transition was possible because giant magnetoresistive heads have required relatively few changes in the surrounding disk-drive components. According to researchers, the progression to drive capacities of 100 gigabits per square inch will likewise be evolutionary, not revolutionary, requiring only incremental steps.

The Giant Magnetoresistive Head: An important discovery from IBM

After intense research and development, GMR-Giant Magnetoresistance makes its mass-market debut in IBM's16.8-gigabyte hard disk drive for desktop computers [19].

The History of GMR

The "giant magnetoresistive" (GMR) effect was discovered in the late 1980s by two European scientists working independently: Peter Gruenberg of the KFA research institute in Julich, Germany, and Albert Fert of the University of Paris-Sud . They saw very large resistance changes -- 6 percent and 50 percent, respectively -- in materials comprised of alternating very thin layers of various metallic elements. This discovery took the scientific community by surprise; physicists did not widely believe that such an effect was physically possible. These experiments were performed at low temperatures and in the presence of very high magnetic fields and used laboriously grown materials that cannot be mass-produced, but the magnitude of this discovery sent scientists around the world on a mission to see how they might be able to harness the power of the Giant Magnetoresistive effect [19].

Stuart Parkin and two groups of colleagues at IBM's Almaden Research Center, San Jose, Calif, quickly recognized its potential, both as an important new scientific discovery in magnetic materials and one that might be used in sensors even more sensitive than MR heads.

Parkin first wanted to reproduce the Europeans' results. But he did not want to wait to use the expensive machine that could make multilayers in the same slow-and-perfect way that Gruenberg and Fert had. So Parkin and his colleague, Kevin P. Roche, tried a faster and less-precise process common in disk-drive manufacturing: sputtering. To their astonishment and delight, it worked! Parkin's team saw GMR in the first multilayers they made. This demonstration meant that they could make enough variations of the multilayers to help discover how GMR worked, and it gave Almaden's Bruce Gurney and co-workers hope that a room-temperature, low-field version could work as a super-sensitive sensor for disk drives [19].

The key structure in GMR materials is a spacer layer of a non-magnetic metal between two magnetic metals. Magnetic materials tend to align themselves in the same direction. So if the spacer layer is thin enough, changing the orientation of one of the magnetic layers can cause the next one to align itself in the same direction. Increase the spacer layer thickness and you'd expect the strength of such "coupling" of the magnetic layers to decrease. But as Parkin's team made and tested some 30,000 different multilayer combinations of different elements and layer dimensions, they demonstrated the generality of GMR for all transition metal elements and invented the structures that still hold the world records for GMR at low temperature, room temperature and useful fields. In addition, they discovered oscillations in the coupling strength: the magnetic alignment of the magnetic layers periodically swung back and forth from being aligned in the same magnetic direction (parallel alignment) to being aligned in opposite magnetic directions (anti-parallel alignment). The overall resistance is relatively low when the layers were in parallel alignment and relatively high when in anti-parallel alignment. For his pioneering work in GMR, Parkin won the European Physical Society's prestigious 1997 Hewlett-Packard Europhysics Prize along with Gruenberg and Fert [19].

Current Hard Drives, LaCie FireWire HDDs

LaCie, a Boston Manufacturer, has begun shipping the 3.5 inch 7200-rpm desktop hard drive, which provides up to 75GB of capacity -- enough to store more than five hours of DV video, more than 100 CD-Audio images or more than a mile-high stack of text documents. By enhancing the external 75GB drives with a built-in FireWire controller, LaCie provides digital video, music and digital imaging professionals with a secure storage capability. The compact high-speed drive also allows for exchanging large files such as high-resolution photos and digital videos with no risk of quality loss. The LaCie 75GB FireWire HDD can be connected to the native FireWire ports that are standard with Macintosh G3, G4 and iMac DV systems as well as PCs incorporating Windows(R) 98 Second Edition or Windows 2000. LaCie's built-in FireWire controller is a 400 Mbps interface that is based on a 1394 to IDE/ATAPI bridge chipset.

LTO (Linear Tape-Open) Ultrium Tape Drives

IBM last week announced that its Linear Tape-Open (LTO) Ultrium tape drive was awarded the XChange Xcellence Award for "Best New Product, Hardware Division" at CMP's Solution Provider XChange held in New Orleans. The LTO Ultrium tape drive was voted the best new hardware innovation of the year, over competitive products from more than 70 vendors attending the show [19].

The LTO family of products consists of: Ultrium Tape Drive with a capacity of up to 200 gigabytes compressed Ultrium Tape Autoloader can hold up to seven tape cartridges (1.4 terabytes compressed) Ultrium Scalable Tape Library with a capacity of up to 72 cartridges (14.4 terabytes compressed) UltraScalable Tape Library can hold up to 2,481 cartridges (496 terabytes compressed) [19].

IBM's comprehensive family of LTO-based solutions include tape drives and small, medium and large tape-automation products that leverage the Ultrium ultra high-capacity recording format. This family is designed to deliver exceptional performance, widespread availability and a cost-effective data storage solution [19].

Using Ultrium LTO products, customers can transfer information at 30 megabytes per second compressed, which equates to an astounding 172 miles per second. That rate is more than 8 megabytes/second faster than competitor Quantum's SDLT 220. IBM LTO products were specifically designed for use in automated tape storage libraries that employ sophisticated robotic arms to automatically retrieve tapes. The process of selecting and retrieving the correct tape is reduced to seconds [19].

IBM was the first to market with LTO products less than a year ago. The LTO Ultrium products are the culmination of an industry initiative by IBM, Hewlett-Packard and Seagate to create an open industry-standard format for digital tape. It serves as an alternative to Quantum's proprietary DLT and Super DLT formats. According to a recent analyst report, LTO Ultrium products are expected to make inroads into market share currently held by Quantum, the dominant player in the midrange market [19].

The LTO specification was jointly developed by three of the world's leading storage technology providers: IBM, Hewlett-Packard and Seagate [19].

The CD-ROM Technology

The Compact Disc-CD was invented in 1982 by two well-known companies, Philips and Sony. This was an Audio CD, digital audio. Soon the computer industry recognized that the large amount of digital audio data could be replaced by digital computer data [5].

CD-ROM (Compact Disc Read-Only Memory) represents a new important tool for the storage, retrieval, and the distribution of information. CD-ROMs can store and play back audio, video, graphic images, digital data, and digitized text. CDs contain computer data in the same way as the hard disk. Compact Disks are usually in the ISO 9660 format. This is a data format that was introduced in 1984 by the International Standards Organization (ISO). Since then it has succeeded in becoming a widely accepted cross-platform [5].

Because of the digital storage techniques used, CD-ROM can be used to retrieve all types of digitally encoded data. The diameter of a CD disc is 4.7 inches (120 mm or 12 centimeters) and its thickness is 0.047 inches. CDs rotate at a constant linear velocity of 200 to 530 revolutions per minute [5].

A single Compact Disk can store up to 600-700MB of data and this is equivalent of:

- 1500 5.25-inch floppy disks
- 450 3.5-inch floppy disks
- 200 books each containing 1,000 pages
- 10 computer magnetic tapes
- 275,000 pages of text

Types of Compact Disks

(1)	CD-DA	:	Audio CD, Digital Audio by Philips and Sony
(2)	Photo CD	:	Developed by Kodak and Philips
(3)	Video CD	:	Can hold around 70 minute of video footage
(4)	CD Extra	:	For Multimedia and developed in 1996
(5)	CD-R	:	Well known technology for WORM disks
(6)	CD-RW	:	Rewritable (erasable)-CD, developed in 1995

The Evolution of CD-ROM & DVD-ROM

1980 / CD-Audio Philips and Sony create standards for optical disk-based digital audio format [18].

1983 / CD Players Sony introduces the first audio CD player, priced at $1,000. The first CD title is Billy Joel's 52nd Street.

1985 / CD-ROM Philips and Sony announce the standard for compact disk-read-only memory for computer data.

1987 / CD-ROM Drive The CD-ROM format enters in the PC arena. Drives cost $1,000; the read-only disks hold 650MB of data.

1994 / 4X CD-ROM Drive Throughput doubles to 600 KBps. Prices for 4X drives are initially $1,000.

1995 / 6X CD-ROM Drive Throughput rises to 900 KBps, with 6X drives initially costing $600, but these models are about to be usurped by 8X drives.

1996 / 8X CD-ROM Drive Throughput is up to 1,200 KBps, with the price of an 8X drive initially $400.

10X and 12X CD-ROM Drives These drives have a claimed throughput of up to 1,800 KBps. They cost around $250.

CD-R Prices for compact disk-recordable drives drop below $500. This technology lets mainstream PC users create their own 650MB CD-ROMs for data archiving or distribution. The user can write to each blank disk only once.

1997 / CD-RW Compact disk-rewritable (also called compact disk-erasable) drives and media let users overwrite files on CDs they have created. CD-RW disks are backward-compatible, letting any standard CD-ROM drive read them.

DVD-ROM The paradigm shift to DVD disks begins. These read-only disks hold 4.7GB of data, and the format is standard to both the PC and the consumer electronics markets. The drives can read legacy CD-ROM disks [18].

1998 / High-Capacity DVD-ROM Subsequent iterations of the DVD format increase capacities to 8.5GB for dual-layer designs; to 9.7GB for double -sided, single-layer implementations; and to 17GB for double-sided, dual-layer designs [18]

DVD-RAM DVD-random access memory drives let users create their own 2.6GB DVDs. The industry will likely skip the write-once format and go straight to designing rewritable disks [18].

Digital Video Disc (DVD)

Digital Video Disc or Digital Versatile Disc (DVD), an optical storage device that looks the same as a compact disc but is able to hold about 15 times as much information and transfer it to the computer about 20 times as fast as a CD-ROM. A DVD, also called a Super Density disc (SD), can hold 8.5 gigabytes of data or four hours of movies on a side; double -sided and rewriteable DVD discs are available now [30].

DVDs come in two formats: the DVD-Video format and the DVD-ROM (DVD-Read Only Memory) format. The DVD-Video format is used for home movie entertainment through a DVD player. DVD players are backward compatible to existing technologies, so they can also play Audio CD and CD-ROM formats. The DVD-ROM stores computer data. DVD-ROM uses include interactive games, video file storage, and photographic storage; it is called the "digital versatile disc" because it can be used in a variety of ways. Recently, DVDs are also used to record data on a DVD-RAM (DVD-Random Access Memory) or DVD-R (DVD-Recordable) disc. When compared to CD-ROM technology, DVD allows for better graphics, greater resolution, and increased storage capacity [31].

Creating your own DVD

DVD is now the simplest way to share the Movies with family, friends and co-workers. Thanks to DVD technology, for the first time ever there is a seamless solution from creation to distribution of movies. Now one can *burn* his/her own DVD-Video discs for playback on most standard DVD players. Power Mac G4 Computers with the DVD-R/CD-RW SuperDrive are the first computers to ship with everything needed to record DVD titles for playback on most standard DVD players [31].

The new DVD-R/CD-RW SuperDrive reads DVD titles at 6x (7.8 megabytes per second), and writes to 4.7-gigabyte DVD-R discs at 2x (2.6 megabytes per second). The SuperDrive also reads CDs at 24x, writes to CD-R at 8x, and writes to CD-RW at 4x. It supports DVD-Video, DVD-ROM and DVD-R, as well as CD-ROM, CD-Audio, CD-R, CD-RW, CDI, CD Bridge, CD Extended, CD Mixed Mode and Photo CD media. That's why it's called the SuperDrive [31].

Power Mac G4 models with the built-in SuperDrive not only come pre-loaded with iMovie, they also give you Apple's revolutionary new iDVD software for creating your own DVD titles. Professionals will want to augment that with Apple's "Final Cut Pro" software (for sophisticated editing, compositing and special effects) and DVD Studio Pro (a complete set of interactive authoring and production tools for producing professional-level DVDs from start to finish) [31].

Microdrive Technology

In April 2001, IBM unveiled the 1 GB Microdrive. They claim this microdrive is the world's smallest hard disk drive and it has successfully completed two NASA shuttle missions. The one-inch Microdrive was used to store hundreds of digital images taken by astronauts on the recent Atlantis and Discovery shuttle missions. The drive was first subjected to a series of tests including high doses of radiation and surviving in a weightless environment. IBM designed the Microdrive for reliability, high performance and huge capacity - features that are essential to meet NASA's rigorous standards for critical missions. The IBM Microdrive boasts capacities ranging from 340 MB to 1 GB, and serves portable electronic devices including digital music players, digital cameras and personal digital assistants. The Microdrive also supports multiple data types including MP3, text, JPEG [19].

Magnetic Tapes

Electromagnetism underlies technologies of great importance to the computer industry. The history of magnetic recording begins with Danish physicist Valdemar Poulsen's patent application for a "*Method of Recording Sounds or Signals*" in 1899. September 1900 issue of Scientific American reported that Poulsen, invented a magnetic-wire speech recorder called the telegraphone. This instrument used several hundred feet of fine steel piano wire wrapped around a brass drum of 5-inch diameter [9].

Magnetic recorders using steel wire or steel tape were used in some specialized areas, but the real break-through came in the years immediately after the second World War with the replacement of wire, as a recording medium, by a thin flexible layer of plastic tape coated with a magnetizable material. Then innovations of the 1950s had made possible commercially viable magnetic tape and disk storage products. Before that time engineers throughout the industry pursued many approaches to storage that failed to find long term market acceptance either because of the cost of converting customer files from one medium to another or because of technological limitations. Three of such unsuccessful developments were hypertape, magnetic strip file, and photo storage [9].

The Removable Disk Pack

In October 1962 IBM unveiled two disk packs, IBM 1316 and IBM 1311. These packs contained an array of six disks and with its protective covering weighed less than 10 pounds. The ten recording surfaces provided a storage capacity of 2 million characters (2 MB). A disk pack was the equivalent in storage capacity of 25,000 punched cards of a fifth of a tape reel, served very well as auxiliary storage in many system environments during 1960s [1].

Modular Storage Systems / Giga Screamers

In April 2001, Amdahl Corporation unveiled its modular storage system GSS 4900 known as the "Giga-Screamer." The Giga Streamer GSS 4900 is powerful enough to run demanding data-warehousing and media-streaming applications, and is configured with as many as four storage arrays, which can handle up to 240,000 I/Os per second or achieve data rates well over a gigabyte per second. The range of configurations allows customers to start small and grow their capacity, performance, and connectivity to huge levels. A capacity of 100 gigabytes, for example, can be expanded to a massive 64 terabytes-over 500 times growth [24].

The design of the GSS 4900 includes its own integrated storage area network (SAN), allowing connectivity to many open systems servers and enabling advanced storage management solutions. Up to 48 100MB/second Fibre Channel ports can be configured. The integrated SAN gives customers immediate access to the benefits of a SAN-improved performance, non-disruptive growth, a centralized storage pool, centralized management, and LANless and serverless backups-without the challenges of designing, installing, and integrating their own SAN [24].

The GSS 4900 integrates storage management software within the storage system. Advanced Function Engines can carry out centralized backups to directly attached tape drives. This offloads the customer's application servers and networks, and avoids contention on the LAN and SAN. The Giga Screamer integrates industry-leading components from several major partners including servers from Fujitsu, storage arrays from LSI Logic Storage Systems, and I-SAN switches from BROCADE Systems [24].

Storage Virtualization

Storage Virtualization is a powerful new tool which can be applied to both hard disk and automated tape storage devices. In storage virtualization the user sees a single interface, which provides a logical view, rather than a physical configuration, of the storage devices. With virtualization techniques implemented in software and in hardware, the user doesn't need to know how storage devices are configured, where they're located, their physical geometry, or what their capacity limit is [23].

SAN-STORAGE AREA NETWORKS

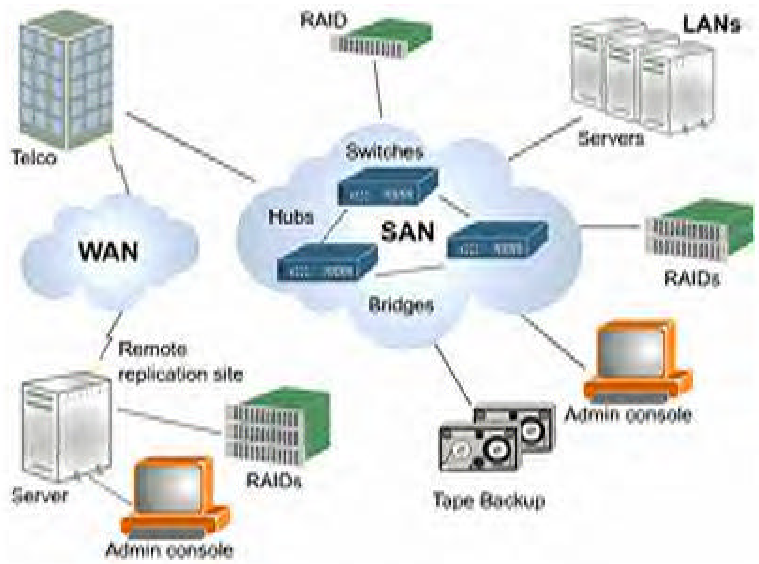

A SAN, or storage area network, is a dedicated network that is separate from LANs and WANs. It generally serves to interconnect the storage-related resources that are connected to one or more servers. It is often characterized by its high interconnection data rates (Gigabits/sec) between member storage peripherals and by its highly scalable architecture. Though typically spoken of in terms of hardware, SANs very often include specialized software for their management, monitoring and configuration [25].

SANs can provide many benefits. Centralizing data storage operations and their management is certainly one of the chief reasons that SANs are being specified and deployed today. Administrating all of the storage resources in high-growth and mission-critical environments can be daunting and very expensive. SANs can dramatically reduce the management costs and complexity of these environments while providing significant technical advantages [25].

SANs can be based upon several different types of high-speed interfaces. In fact, many SANs today use a combination of different interfaces. Currently, Fibre Channel serves as the de facto standard being used in most SANs. Fibre Channel is an industry-standard interconnect and high-performance serial I/O protocol that is media independent and supports simultaneous transfer of many different protocols. Additionally, SCSI interfaces are frequently used as sub-interfaces between internal components of SAN members, such as between raw storage disks and a RAID controller.

Providing large increases in storage performance, state-of-the-art reliability and scalability are primary SAN benefits. Storage performance of a SAN can be much higher than traditional direct attached storage, largely because of the very high data transfer rates of the electrical interfaces used to connect devices in a SAN (such as Fibre Channel). Additionally, performance gains can come from opportunities provided by a SAN's flexible architecture, such as load balancing and LAN-free backup. Even storage reliability can be greatly enhanced by special features made possible within a SAN. Options like redundant I/O paths, server clustering, and run-time data replication (local and/or remote) can ensure data and application availability. Adding storage capacity and other storage resources can be accomplished easily within a SAN, often without the need to shut down or even quiese the server(s) or their client networks. These features can quickly add up to large cost savings, fewer network outages, painless storage expansion, and reduced network loading [25].

By providing these dedicated and "very high speed" networks for storage and backup operations, SANs can quickly justify their implementation. Offloading tasks, such as backup, from LANs and WANs is vital in today's IT environments where network loads and bandwidth availability are critical metrics by which organizations measure their own performance and even profits. Backup windows have shrunken dramatically and some environments have no backup windows at all since entire data networks and applications often require 24x365 availability.

As with many IT technologies, Sans depend on new and developing standards to ensure seamless interoperability between their member components. SAN hardware components such as Fibre Channel hubs, switches, host bus adapters, bridges and RAID storage systems rely on many adopted standards for their connectivity. SAN software, every bit as important as its hardware, often provides many of the features and benefits that Sans have come to be known for. SAN software can provide or enable foundation features and capabilities, including:

- SAN Management
- SAN Monitoring (including "phone home" notification features)
- SAN Configuration
- Redundant I/O Path Management
- LUN Masking and Assignment
- Serverless Backup
- Data Replication (both local and remote)

- Shared Storage (including support for heterogeneous platform environments)

SAN HARDWARE

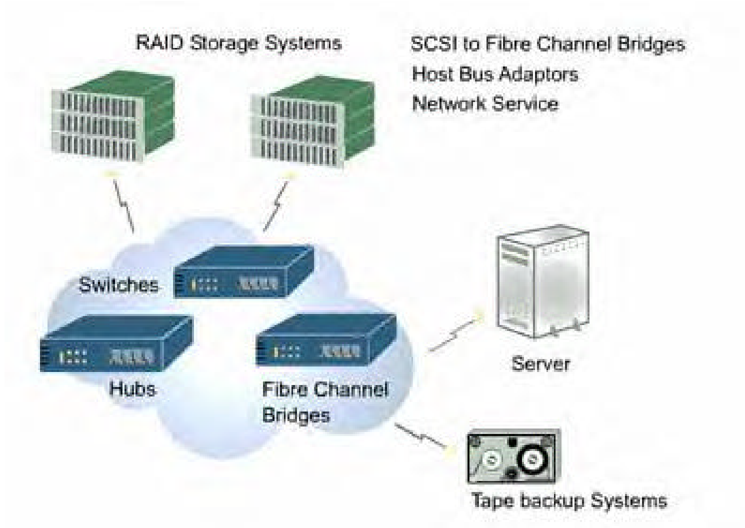

SANs are built up from unique hardware components. These components are configured together to form the physical SAN itself and usually include a variety of equipment. RAID storage systems, hubs, switches, bridges, servers, backup devices, interface cards and cabling all come together to form a storage system that provides the resources that facilitate the policies of an IT organization [25].

It is very important to select the hardware devices (and their configuration) for a SAN with care and consideration. Many of the "standards" that are involved with SANs are concerned with interoperability. Some of these standards are still evolving and haven't been equally adopted by all manufacturers of equipment used in SANs. This can lead to difficulties when matching up devices from different vendors and suppliers. Since SANs are typically just as dependent upon software for their proper operation, it can be vital to secure the latest version information about software (and firmware) and potential compatibility issues [25].

Working with companies that specialize in the design, integration and implementation of SAN systems can provide great benefits. Firms that specialize in SANs are often familiar with the latest software and hardware and can speed the process of successfully deploying SAN technology. By working with other vendors, manufacturers and standards bodies, these SAN specialists can help ensure that the promised benefits are realized and successfully integrated into new or existing IT infrastructures [25].

RAID SYSTEMS

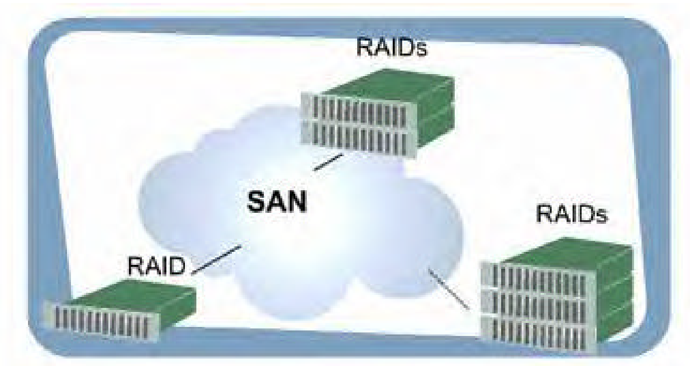

Most contemporary SANs include RAID systems as their primary data storage devices. These systems have become highly complex. They offer the foundation principles that have come to be the basic modern SAN. First, RAID systems offer data protection, or fault tolerance, in the event of a component or I/O path failure. This is true even if fundamental elements, such as disk drives, fail in the system. Additionally, by way of numerous data striping techniques (described below), and controller configurations, today's RAID systems offer very high performance, storage capacity, scalability, and survivability. Other reliability features available in today's RAID systems include redundant cooling systems, power supplies, controllers and even monitoring circuitry. These, and other features and characteristics, contribute dramatically to high data availability in a SAN. Modern RAID systems can even permit the direct connection of backup equipment, thus facilitating LAN-free and even serverless data backup and replication [17].

The roots of RAID technology can be traced back to Patterson, Gibson and Katz at the University of California at Berkeley. The ideas presented and explained in the paper involved combining multiple small, inexpensive disk drives into arrays in order to provide features that single drives alone couldn't supply. These new features centered on improving I/O performance and automatically preserving the contents of drives during, and after, drive or component failures [17].

These drive arrays are presented to a computer as a single logical storage unit (LUN) or drive. Additional benefits of drive arrays include the ability to make these arrays fault-tolerant by redundantly storing data in various ways. Five of the array architectures, RAID levels1 through 5, were defined by the Berkeley paper as providing disk fault-tolerance with each offering various trade-offs in features and performance. Overall, the idea was to improve the reliability of the storage system by significantly increasing the Mean Time Between Failure (MTBF) for the array and to dramatically improve the storage system's performance. A sixth common type of RAID architecture, RAID 0, has subsequently been defined that can substantially improve the I/O performance of an array but it provides no data protection should a hardware component fail. The performance gains possible with RAID 0 arrays can be very dramatic. RAID 0 arrays are ideal for applications that demand the highest possible data throughput. Note that these applications must be able to tolerate possible data loss, and service interruption, if a drive or other component in the array fails [17].

Data Striping in RAID

Fundamental to the RAID concept is "striping", a method of dividing and distributing data among the arrayed drives and effectively joining multiple drives into one logical storage unit. Striping involves partitioning each drive's storage space into stripes that may be as small as one block (512 bytes) or as large as several megabytes. These stripes are then interleaved in a round robin fashion, so that the combined space is composed of joined stripes from each drive. In most instances, the application environment determines the suitability of larger vs. smaller stripe sizes.

Most contemporary multi-user operating systems like UNIX, Solaris, NT and Netware support overlapping disk I/O operations across multiple drives. However, in order to maximize throughput for a combined disk subsystem, its I/O load must be balanced between all of its member drives so that each drive can be kept as active as possible. High parallelism during I/O operations generally translates into much greater performance [25].

In multiple drive systems without striping, the disk I/O load is almost never perfectly balanced. Some drives will contain data that is frequently accessed while other drives will only occasionally be accessed. During intense I/O operations, performance can be optimized by striping the drives in the array with stripes large enough so that each record potentially falls entirely within one stripe segment.

This helps insure that data and I/O operations are evenly distributed across the arrayed drives, thus allowing each drive to work on separate I/O operations at the same time, thereby maximizing the number of simultaneous I/O operations that can be serviced by the array. By contrast, in data-intensive applications that access large records, smaller stripe sizes can be used so that each record will span across many, or all, of the drives in an array with each drive storing only part of a record's data. This can allow long record accesses to be performed faster, since the data transfers can occur in parallel on multiple drives in the array. Applications such as digital video editing, audio/video on demand, imaging and data acquisition that employ long record accesses are examples of applications that often achieve optimum performance with smaller stripe sizes. Unfortunately, smaller stripe sizes typically rule out multiple overlapping I/O operations since each I/O will typically involve all of the drives [25].

SWITCHES, HUBS AND BRIDGES

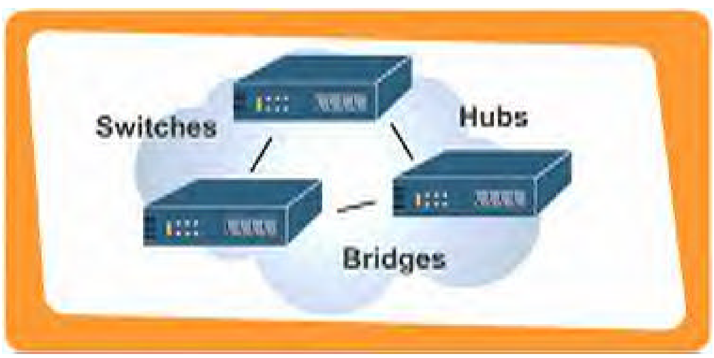

More and more, the design and deployment of SAN technology involves incorporating specialized interconnection equipment. This category of devices often includes Fibre Channel Hubs, Switches and Bridges. This hardware is generally responsible for linking together the data storage peripherals, such as RAID systems, tape backup units and servers within a SAN.

These interconnection devices are somewhat analogous to their LAN-related counterparts. They perform functions such as data frame routing, media and interface conversion (i.e. copper to optical, Fibre Channel to SCSI), network expansion, bandwidth enhancement, zoning, and they allow concurrent data traffic. Just as customers today are more involved in the design and implementation of their LANs and WANs, they are also looking at these building blocks of SANs to create their own SAN solutions [25].

Fibre Channel HBAs, hubs, switches, and FC/SCSI bridges are some of the building block components with which IT administrators can develop SAN-based backup solutions, server clusters, enhanced bandwidth, extended distance and other application driven challenges. Selecting the appropriate pieces to address these issues requires an understanding of what each component can do. When, for example, is a fabric switch a better solution than a hub? When should hubs and switches be used in combination? There are no universal answers to these questions, but understanding the architecture and capabilities of switches, hubs and bridges provides a basis for making appropriate choices for SAN designs.

Backup Solutions

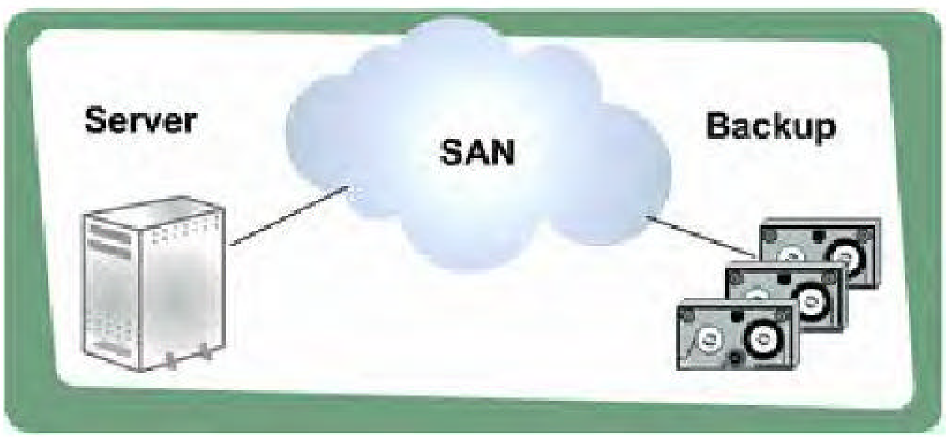

One of the most valuable time- and cost-saving features of a SAN architecture is its ability to offload backup operations from a LAN and/or backup servers. This capability can significantly increase the amount of LAN bandwidth available to network clients and end users during backup operations. When backup servers are relieved from the "data mover" role, they become more available for other productive tasks.

LAN-free and serverless backup solutions optimize backup operations by offloading backup data traffic from a LAN, thereby increasing the amount of LAN bandwidth available to end users. Serverless backup extends these performance gains by removing more than 90 percent of the backup administration overhead that is usually placed on a backup server as backups are performed. This is achieved by incorporating some of the backup intelligence into the data storage or connectivity peripherals themselves. This can significantly free up backup servers by releasing them from large portions of a backup operation's administration and data moving chores. Using these SAN based backup solutions lets administrators optimize network and server utilization.

SANS' Background

Traditional backup operations place the application server, the backup server and the LAN all in the data path. Consequently, as the amount of storage grows, the amount of time and network resources needed to back it up grows. Now that businesses and organizations have moved toward 24 x 365 operation, backup tasks are competing with critical business applications for server time and network resources. Invariably, this causes network congestion and can result in business slowdowns.

For "serverless" backup operations, host computers (servers) do not "handle" or touch the backup data itself. Instead, these hosts merely direct and monitor the backup without actually moving the data. The backup data is copied directly from disk to tape, or disk to disk, by the storage peripherals themselves using intelligence that is incorporated into them. Optionally, this intelligence can even be placed inside of other SAN components, such as Fibre Channel switches or hubs. Freed from the routine data transport burden, server resources can be put back to more productive uses. Or, in other words, the backup or tape server is delegated the role of "backup coordinator," rather than data mover. Serverless backup takes LAN-free backup a step further since it removes backup traffic from both the LAN and the backup server.

By contrast, with simple "LAN-free" backup operations, the backup and restore data (traffic) is removed from the LAN but still flows through the administrating server as it moves between data storage and backup devices. The benefit here is still valuable, since backup traffic is taken off of the LAN, reducing LAN congestion. While both serverless and LAN-free backup keep backup data off of the LAN, only the serverless backup frees up the administrating server as well - placing the data movement tasks onto the smart peripherals.

Specifically, smarter peripherals can now perform much of their own backup by supporting newer technologies and APIs - such as the "extended copy command," a Storage Networking Industry Association specification that lets data be moved between storage devices on different buses. The backup server issues the command to a data mover in the SAN, and then removes itself from the data path. This way only the source, destination and SAN devices are involved. The constraints related to the memory, I/O and CPU performance of the backup server itself are eliminated as the data moves through a high-performance copy device or agent that is optimized for data movement. The logic here is fairly obvious since this frees up the backup server for other business-critical applications, and supports server consolidation. The dedicated backup server is no longer needed. Additionally, backups can complete much more quickly over higher speed networks - such as Fibre Channel.

Serverless backup systems can also provide additional cost savings by eliminating expensive, high-end servers. Another advantage unique to the serverless backup architecture is its ability to stream the same data to several tape libraries or other targets simultaneously, even if they are geographically separated, without the need for copying and moving the actual tapes - an important advantage in disaster recovery plans.

Backup Hardware in SANS

Mechanically, backup equipment used in SANs is typically the same as that used in conventional configurations. What is different, however, is how these devices are interfaced to their host servers and client storage systems. Since most contemporary SANs are connected together using Fibre Channel, and since many backup devices use SCSI interfaces, some type of bridge is often required.

These bridges perform the electrical, and any protocol, conversions required between the disparate buses or channels. There are many bridge manufacturers that supply these units, but it is vital to confirm compatibility with the selected backup device(s) before attempting to configure the units together or specifying units for purchase. This SAN topic has many caveats and is often an area that benefits from direct experience. SAN consultants, equipment vendors, and SAN solutions providers can be excellent sources for this type of compatibility information.

If a serverless backup solution is being considered or designed, it is important to note that some of these bridge manufacturers offer "smart" units that include built-in copy functions. As mentioned above, this set of features is generally required for true serverless backup. In practice, small applications running on a selected server activate and instruct these copy agents remotely, then sit back and monitor progress while the smart bridge, or similar device, moves the data.

These software, or firmware, copy agents can even be found in certain Fibre Channel switches and hubs. There are some software companies that have even developed special programs that can be 'downloaded' into these units to give them these independent copy capabilities. With all of these options available, one can see the importance of selecting components very carefully. First time SAN adopters may want to consider consulting with SAN specialists before purchasing this type of equipment for deployment in a backup solution.

Though SAN backup solutions typically employ a tape server, tape library, and disk-based storage attached together with a Fibre Channel infrastructure, it is becoming fairly common for backup solutions to include disk to disk copies. With today's backup windows shrinking and IT policies calling for remote site replication, backup can now mean much more than just making tapes. Backup can now include data replication to local or remote disk sites via WANs, disk to disk backups that accommodate offline data processing and short backup windows, or all of these at the same time.

Patents on Storage Related Devices in 2000

IBM was awarded with over 400 storage-related U.S. patents in 2000 [18].

References

[1] Charles J. Bashe, Lyle R. Johnson, John H. Palmer, Emerson W. Pugh, *IBM's Early Computers,* The MIT Press, Bernard Cohen, Editor, Cambridge Massachusetts, 1986

[2] Richard E. Matick, *Computer Storage Systems and Technology,* John Wiley & Sons, A Wiley-Interscience Publication, 1977

[3] Russel L.A., et al., *Ferrite Memory Systems*, IEEE Trans. Magazine., MAG-4, No 2, June 1968.

[4] Smith, C.V., *Electronic Digital Computers*, McGraw Hill, New York, 1959.

[5] Judith Paris Roth (Ed.) 1986. *Essential Guide to CD-ROM.* Rockville, Md., Computer Science Press, 3d ed.

[6] Stifler, W.W. (Ed.)., *High Speed Computing Devices*, Engineering Research Associates Staff, McGraw Hill, New York, 1950.

[7] Eckert, W.J. *Punched Card Methods in Scientific Calculations,* Columbia University, New York, January 1940.

[8] Hoagland A.S., *Digital Magnetic Recording*, John Wiley & Sons, New York, 1963.

[9] Sanjay Ranade, *Mass Storage Technologies*, Meckler Publishing, the division of Meckler Corporation, 11 Ferry Lane London, 1991.

[10] Sebestyen, L.G., *Digital Magnetic Tape Recording for Computer Applications*, Chapman and Hall Publishing, New Fetter Lane, EC4P 4EE London, 1973.

[11] John J. Mallinson, *Magneto-Resistive Heads: Fundamentals and Applications,* Electromagnetism Series, Academic Press, 1995

[12] Kanu G. Ashar, *Magnetic Disk Drive Technology: Heads, Media, Channel, Interfaces, and Integration*, IEEE Press, 1997

[13] Shan X. Wang, A. M. Taratorin, *Magnetic Information Storage Technology*, Academic Press, 1999

[14] Eric D. Daniel, C. Dennis Mee, *Magnetic Recording: The First 100 years,* IEEE Press, 1999

[15] Hans J. Coufal, Demetri Saltis, *Holographic Data Storage,* Springer-Verlag, 1999
TELOS, 2000

[16] Judith Paris Roth, Berth Gal, Ashok Mathur, Donald Rodgers, *Essential Guide to CD-ROM,* Meckler Publishing, 11 Ferry Lane, Westport CT, 1986

[17] Patterson, Gibson, Katz, *A Case for Redundant Arrays of Inexpensive Disks (RAID),* University of California at Berkeley, 1987

[18] http://www.yamahacdrwinfo.com

[19] http://www.ibm.com

[20] http://www.sony.com

[21] http://www.philips.com

[22] http://www.hp.com

[23] http://www.compaq.com

[24] http://www.maxell.com

[25] http://www.gartner.com/storage/usa
Gartner, Dallas, Texas 75391

[26] http://www.101com.com/solutions/storage/

[27] http://www.amdahl.com/doc/products/storage/

[28] http://www.dothill.com/tutorial/

[29] http://www.Internetwire.com

[30] http://www.cd-info.com

[31] http://www.apple.com