



B6545: TECHNOLOGY FOUNDATIONS  
FOR E-COMMERCE

STORAGE TECHNOLOGIES  
SUMMARY



Figure 1: 40 GB Hard Disk Drive with EIDE Interface (Maxtor)

---

**HARD DISK DRIVES**


---

Class	Length	Width	Height	
3.5"	5.75"	4.00"	1.00"	The device that is most widely used for data storage is the hard disk drive (see Figure 1). Disk drives which come in a variety of brands and capacity can be classified according to their physical dimensions (see Table 1) which refers to the diameter of the platters, their revolution speed or their data handling interface. 3.5" hard disks are mostly used in desktop and server systems, and they are also the choice for the composition of disk arrays. 2.5" drives are used in laptops, and 1.0" drives are usually found in PDAs or in digital cameras.
2.5"	4.00"	2.75"	0.50"	
1.0"	1.70"	1.40"	0.20"	

Table 1: Physical Dimensions of Disk Drives

When hard disks are classified according to their data interface, we usually encounter drives with IDE<sup>1</sup> interface (nowadays EIDE<sup>2</sup>) in the consumer and low-end server market and disk drives with SCSI<sup>3</sup> interface in the workstation and high-end server market. SCSI drives are also used for high-performance disk arrays whereas budget disk arrays are often equipped with IDE drives. IDE drives are cheaper than their SCSI counterparts. Although the realistically attainable streaming data transfer capacity is similar in modern IDE and SCSI drives, the SCSI drives are definitively faster when multiple files are accessed at the same time, and therefore, they are the preferred choice for high-end applications. The technical reason behind this is that disk drives need a certain time to reposition the read-write head on the platters. This time is called the *access time* and ranges from 3ms in very fast times to 8-10ms in typical consumer drives. During this time, the IDE drives block the IDE bus whereas SCSI drives can release the bus and free it up for a data transfer with another drive or simply let the CPU on the mainboard of a computer continue working on other tasks. Recently, a new interface names Serial-ATA has come up, and the first drives have already been introduced in the market [2]. The reason behind Serial-ATA is to ease cabling in the PCs and to be able to daisy-chain drives. In high-end disk arrays, Fibre Channel is another interface for disk drives although it is usually mostly used for interfacing complete disk arrays with servers rather than individual disk drives.

Hard disks also can be classified according to their revolution speed, and a faster speed means a smaller access time and a faster data throughput. Typical values are 5,400 rpm for consumer drives and 7,200 rpm for newer consumer and low-end SCSI drives. High-end SCSI drives work at 10,000 or even at 15,000 rpm. While faster spinning drives have lower access time, they generate more noise and heat and often require additional cooling measures. Laptop drives have a slightly lower speed of 4,800 or 5,400 rpm

Recently, technological advances like AFC<sup>4</sup> coating have stretched the data density of disk driver platters to the range of 100 Gbit per inch<sup>2</sup> which allows drive capacities of up to 400 GB in the near future.

The major manufacturers of hard disk drives are Seagate (27%), IBM (19.3%), Maxtor (18.5%), Western Digital (9.8%), and Fujitsu (9.5%). Individual hard disks have become a commodity and the prices are expected to decrease further [1].

In desktop computers and small servers, hard disk are usually either directly connected by IDE or SCSI or sometimes by a RAID<sup>5</sup> controller, in order to increase the data security in case of a hardware failure.

---

**DISK ARRAYS**


---

Disk Arrays are sets of combined hard disk drives that are used when large amounts of data has to be stored centrally. This is typically the case for big-scale applications like databases, of when data shall be held centrally for purposes of integrity or easy backup. One can distinguish three different systems as Figure 2 shows.

---

<sup>1</sup> IDE = Integrated Device Electronics, a term used because the early disk drives did not have a complete set of electronics on board, but only basic electronics for stepping the RW head and elementary signal processing.

<sup>2</sup> EIDE = Enhanced IDE. EIDE is also known as Advanced Technology Attachment (ATA).

<sup>3</sup> SCSI = Small Computer System Interface

<sup>4</sup> AFC = antiferromagnetic coupling

<sup>5</sup> RAID = redundant array of independant disks, usually RAID-0 or RAID-1 or RAID-5 (best solution) is used.




Figure 2: Storage Systems in an Enterprise

Direct Attached Storage (DAS) denominates a disk array that is connected to one or two servers usually by means of a SCSI interface. Typical applications include WAP servers in a clustered environment which store their data on a clustered DAS system for redundancy. In general, DAS can be used in single host & single OS<sup>6</sup>, multiple host & single OS or in multiple host & multiple OS scenarios. The major vendors of DAS systems are Compaq (18.2%), EMC (15.6%), IBM (14.2%), and HP (8.1%). The DAS market represents the largest market share in disk array systems, but it is expected to decrease in the future [1].

Storage Area Networks (SAN) are defined as systems that supply data to one or more servers and connect through a hub or switch. SAN do not contain filesystems or similar applications, in fact they are very similar to DAS. However, they have richer capabilities in that data can be transferred over IP or fibre channel rather than over SCSI. Therefore, SAN can be located physically apart from the application servers which they supply with data. The major advantage of SAN is their scalability which is much greater than the one of DAS. Also, backup operations can typically be performed “in the background” without affecting the network traffic of the LAN. SAN are typically high-performance and fast bandwidth storage systems that are ideally used when heavy duty application servers like transaction servers or databases or heavily loaded web servers have to be connected to disk arrays [1]. The SAN market is dominated by EMC (37.9%), IBM (18.2%), and Compaq (17.9%).

Network Attached Storage (NAS) systems are external disk storage systems that are directly attached to a LAN and include higher level protocols like SMB<sup>7</sup> or NFS<sup>8</sup> for file sharing. Consequently, they include their own small specialized computer, an OS<sup>6</sup> and file sharing and authentication applications. NAS can easily be attached into any LAN and are therefore an easy means of increasing storage capacity. However, all data transfer between computers like Office PCs and the NAS add up to the data traffic in the LAN. Therefore, NAS are not an ideal means for heavy-duty application servers unless the LAN is designated exclusively for the data transfer between the servers and the NAS. The NAS market is dominated by EMC (42.3%) and Network Appliance (32.5%).

<sup>6</sup> OS = operating system, like Windows 2000, Windows XP, Windows Server 2003, Linux, Unix, Solaris, AIX, etc.

<sup>7</sup> SMB = server message block, this is essentially the Windows file sharing method, also used by Samba servers under Linux.

<sup>8</sup> NFS = network file system, a file sharing method under Unix and Linux, see RFC 1094, RFC 1813 and RFC 3010.

---

**TAPE DRIVES**

---



Figure 3: 15 GB Tape Cartridge

A storage media which is still used for backups of large amounts of data are magnetic tapes which are often packaged in cartridges like shown in Figure 3. Although disk arrays offer a high degree of reliability due to technologies like RAID, hot-swapping and clustering, data on disk arrays is not safe against manipulation or unintentional deletion or destruction by viruses. Often, it is also necessary to store several version of data sets in order to be able to resort to old data if that is necessary for some reason. Sometimes, data achieving is a legal issue like for banks or insurance companies who are forced to keep data sets of many years. Although a backup on tape is not as fast as mirroring a disk array to another disk array, tapes are a relatively cheap medium for the storage of large amounts

of data.

---

**LITERATURE AND INTERNET SOURCES**

---

- [1] PriceWaterhouseCoopers, "Technology Forecast: 2002-2004 – Volume 2: Emerging Patterns of Internet Computing", PriceWaterhouseCoopers, Menlo Park (CA), October 2002.
- [2] Bremer, L., "Aktuelle Festplatten im Vergleich", *c't*, vol. 14/2003, pp. 136 *ff.*, Verlag Heinz Heise, Zorneding, July 2003.