

Hard Disk Drives and RAID

Janaka Harambearachchi
(Engineer/Systems Development)


INTERFACES FOR HDD

- A computer interfaces is what allows a computer to send and retrieve information for storage devices such as computer [hard disk drives](#) and [CD-ROM drives](#).


[ATA](#)
[ATAPI](#)
[IDE](#)
[EIDE](#)
[Fast ATA](#)
[Fast IDE](#)
[RAID](#)
[SATA](#)
[SCSI](#)
[Serial ATA](#)
[Ultra ATA](#)

ATA / IDE

- Short for **Integrated Drive Electronics** or **IBM Disc Electronics**, **IDE** is more commonly known as [ATA](#) and is a standard interface for IBM compatible hard drives.
- IDE is different from the Small Computer Systems Interface ([SCSI](#)) and Enhanced Small Device Interface (ESDI) because its controllers are on each drive, meaning the drive can connect directly to the motherboard or controller.
- IDE and its updated successor, Enhanced IDE (EIDE), are the most common drive interfaces found in IBM compatible computers today.



ATA




SATA

- The original ATA interface is now often referred to as **PATA**, which is short for **Parallel AT Attachment** after the introduction of [SATA](#).
- Short for **SerialATA**, **SATA** 1.0 was first released in August [2001](#) and is a replacement for the Parallel ATA
- SerialATA is capable of delivering 1.5[Gbps](#) (150[MBps](#)) of performance to each drive within a disk array,
- Offers a thin, small cable solution , This cable helps make a much easier cable routing and offers better airflow in the computer when compared to the earlier ribbon cables used with ATA drives.

SATA


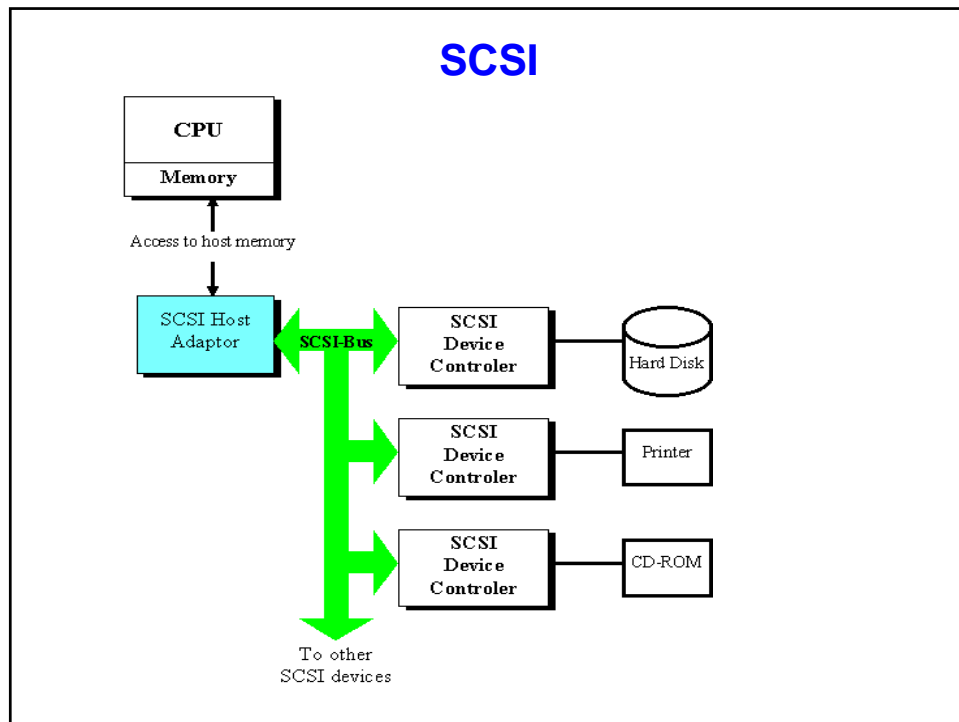


Serial ATA cable vs. Parallel ATA cable.



SCSI

- SCSI, or Small Computer System Interface, pronounced as Scuzzy, was completed in 1982.
- SCSI is a hardware bus similar in function of the IDE controller supporting hard disk drives, CD-ROM drives and other peripherals.
- However, unlike IDE, SCSI allows the support of up to eight devices and uses ID numbers 0 through 07, with the SCSI host adapter located on ID number 07 and boots from the ID 00. This leaves the availability of six device connections.
- Because SCSI devices are connected via chaining the devices together, the SCSI must properly be terminated.



Redundant Array of Inexpensive Disks

- Redundant Array of Inexpensive Disks, or RAID, is a number of hard disks which act as an extremely reliable single disk drive. RAID uses several techniques as explained below.
 1. **Spanning / Software Striping** - Splitting information and writing it across multiple physical disk drives.
 2. **Mirroring** - Duplication of data from one disk drive to another.
 3. **Duplexing** - Duplicates the disk drive as well as the disk controller.
 4. **Deferred** - Data is cached in cache memory and writes to the hard disk drive as the disk drive becomes available.
 5. **Hot Swapping** - Failed disk drives can be replaced and data can be placed back onto the disk drive while the remainder of the system is in operation
 6. **Hot Sparing** - Disk drive is automatically initialized into the array when another fails.
 7. **Spindle Synchronization** - Synchronization of the rotation of all disk drives in the array allowing information to be written all at once.

RAID Levels

RAID 0 - Speed

Level 0 is disk striping only, which interleaves data across multiple disks for better performance. It does not provide safeguards against failure. RAID 0 is widely used in gaming machines for higher speed.

RAID 1 - Fault Tolerance

Uses disk mirroring, which provides 100% duplication of data. Offers highest reliability, but doubles storage cost. RAID 1 is widely used in business applications.

RAID 2 - Speed

Bits (rather than bytes or groups of bytes) are interleaved across multiple disks. The Connection Machine used this technique, but this is a rare method.

RAID 3 - Speed and Fault Tolerance

Data are striped across three or more drives. Used to achieve the highest data transfer, because all drives operate in parallel. Parity bits are stored on separate, dedicated drives.

RAID Levels

- **RAID 4 - Speed and Fault Tolerance**
Similar to Level 3, but manages disks independently rather than in unison. Not often used.

RAID 5 - Speed and Fault Tolerance

Data are striped across three or more drives for performance, and parity bits are used for fault tolerance. The parity bits from two drives are stored on a third drive and are interspersed with user data. RAID 5 is widely used on servers to provide speed and fault tolerance.

RAID 6 - Speed and Fault Tolerance


Highest reliability, but not widely used. Similar to RAID 5, but performs two different parity computations or the same computation on overlapping subsets of the data.

RAID 10 - Speed and Fault Tolerance

A combination of RAID 1 and RAID 0 combined. Raid 0 is used for performance, and RAID 1 is used for fault tolerance.


RAID 0, Striping

- RAID 0 requires at least two hard disk drives. RAID 0 offers no redundancy, it's all about performance.




RAID 1, Mirroring

- RAID 1 requires two hard disk drives. RAID 1 duplicates (mirrors) data from one disk to the other.



RAID 3, Striping with Dedicated Parity

- RAID 3 requires at least three hard disk drives. Data is striped across all disks but one, and that one is a dedicated parity drive containing parity information for data recovery if needed.



RAID 5, Striping with Distributed Parity

- RAID 5 requires at least three hard disk drives. Data is striped across all disks, and a portion of each disk is also used to store parity information for recovery if needed

