

The DASD and Evolution of Storage Devices

Rule 1: “Data processing requires data storage”. – I just made that up.

Data were originally stored on paper media, first as written documents but fairly soon (Hollerith, late 19th century) the storage medium was machine-readable.

In the 1950’s, New York Life Insurance Company was devoting an entire floor of its main building to the storage of punched cards. Something had to change.

IBM quickly came out with two magnetic media for storing data

the magnetic tape

the DASD (disk)


The acronym “DASD” stands for **D**irect **A**ccess **S**torage **D**evice.

Until recently, the standard disk drive was the only commercially viable example.

We now have another very popular example, these USB “flash drives”. While different from standard disk drives, these are managed as if they were disk drives and are considered disk drives.

Structure of a Large Disk Drive

The typical large-capacity (and physically small) disk drive has a number of glass platters with magnetic coating. These spin at a high rate (7,200 rpm or 120 / second)




This drawing shows a disk with three platters and six surfaces. In general, a disk drive with N platters will have $2 \cdot N$ surfaces, the top and bottom of each platter.

On early disk drives, before the introduction of sealed drives, the top and bottom surfaces would not be used because they would become dirty.

More on Disk Drive Structure

Each surface is divided into a number of concentric tracks.

Each track has a number of sectors.



A sector usually contains 512 bytes of data, along with a header and trailer part.

Seek Time and Rotational Latency

In order to read from a disk track, the read/write heads must be moved to the track.

This is a mechanical action, as the read/write heads are physical devices.

There are two seek times typically quoted for a disk.

Track-to-track: the time to move the heads to the next track over

Average: the average time to move the heads to any track.

The **rotational delay** is due to the fact that the disk is spinning at a fixed high speed.

It takes a certain time for a specific sector to rotate under the read/write heads.

Suppose a disk rotating at 12,000 RPM. That is 200 revolutions per second.

Each sector moves under the read/write heads 200 times a second,


once every 0.005 second or every 5 milliseconds.

The **rotational latency**, or average rotational delay, is one half of the time for a complete revolution of the disk. Here it would be 2.50 milliseconds.

The Idea of a Cylinder

Fixed head disks have one head per track. The last time I heard of such a device was 1977, when working with a 1 MB fixed head disk on a PDP-11/45.

I claim that fixed head disks are obsolete. Revisit the picture of a typical disk.



Question: How many tracks can be read before the read/write heads must be moved?

Answer: One track per surface can be read without moving the heads. Here it is 6.

Definition: A cylinder is that set of tracks that can be read without moving the disk read write heads. A disk has as many cylinders as a surface has tracks.
A cylinder has as many tracks as the disk has surfaces.