Nov-3

Note Title 03-11-2012

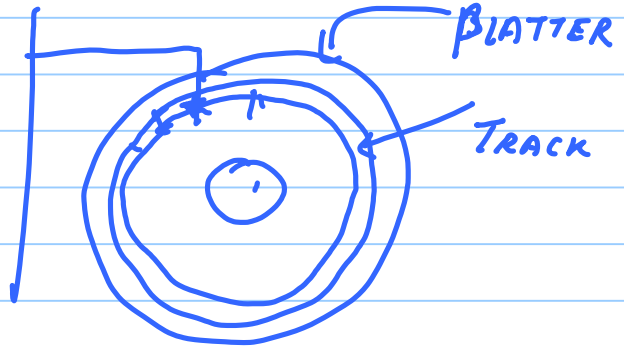

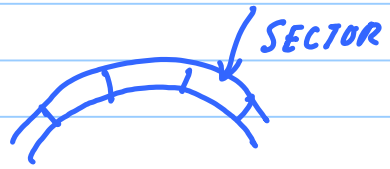

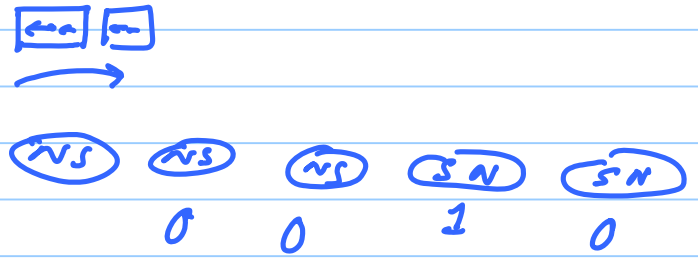

1) Position the head on the right track.
(seek time)

2) Rotational Latency.

3) Transfer time: $\frac{\text{Size}}{\text{BW}}$

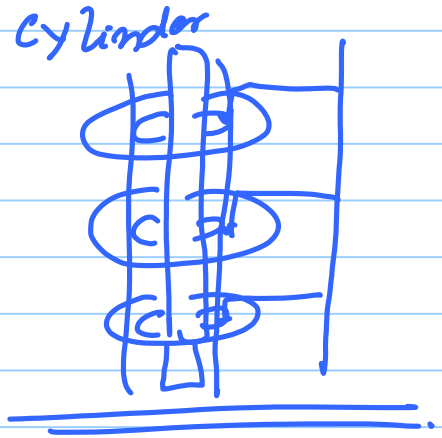

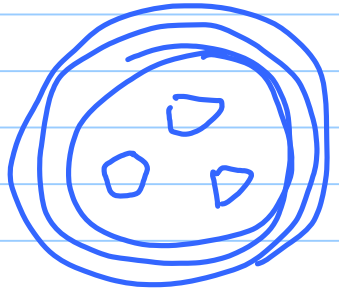

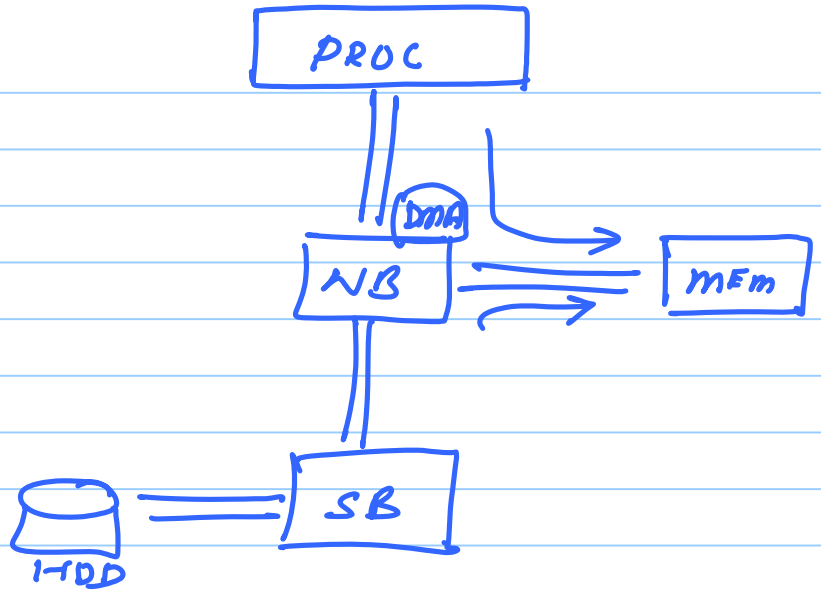

PROC — High Prio
DMA — Low Prio
} Cycle stealing mode

PROC — Low
DMA — High
} Burst Mode

Bank: 10^7 customers
1 MB of data/customer.

$= [10\,TB]$

Facebook: 30PB

RAID: Redundant Array of Inexpensive Disks.

RAID-0

No fault tol.

RAID-1

Mirroring.

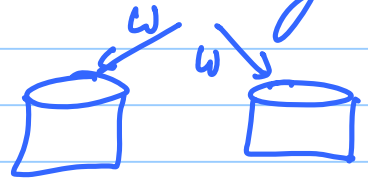

RAID-2

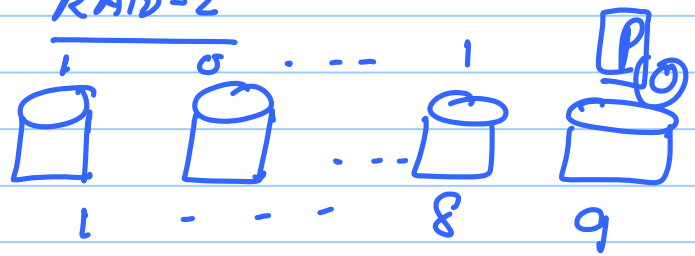

1000 ... 001

Failure Models:

FAIL-STOP: Immediately detect & stop

FAIL-SILENT: Silent Data Corruption.

RAID 3.

Byte Level (RAID 2)

RAID 4.

Block Level (RAID 2)

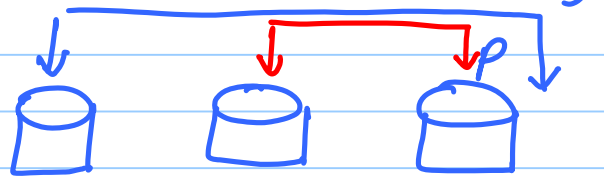

RAID-5

Distribute Parity

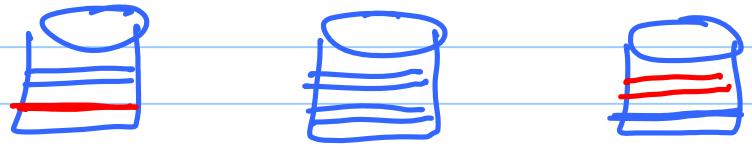

Example.

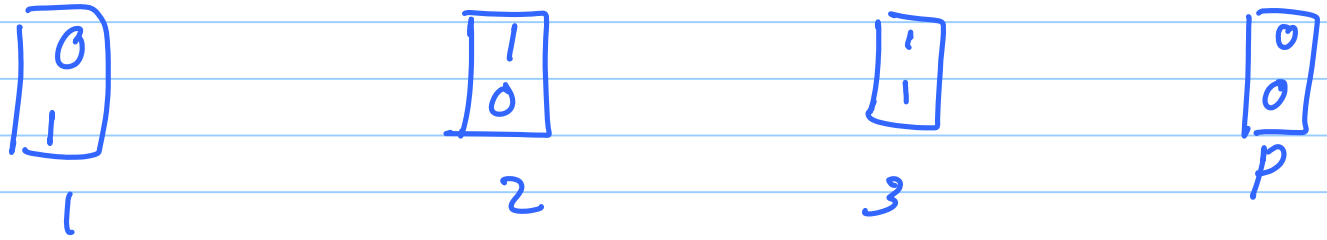

V_{blk}^{old} (old value) V_{blk}^{new} (new value)

p^{old} (old value of parity)

$$p^{new} = p^{old} \oplus V_{blk}^{old} \oplus V_{blk}^{new}$$

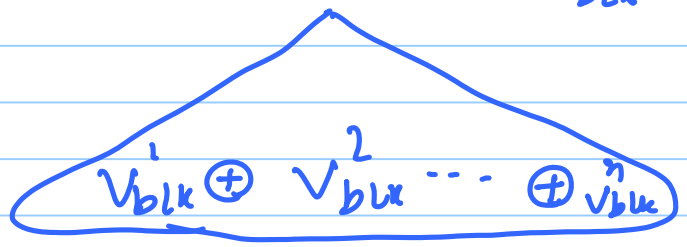

Any write: 2 Reads + 2 Writes.

RAID 6: Tolerate 2 disk failures.

$(N+2)$ disks [2 disks distributed parity]

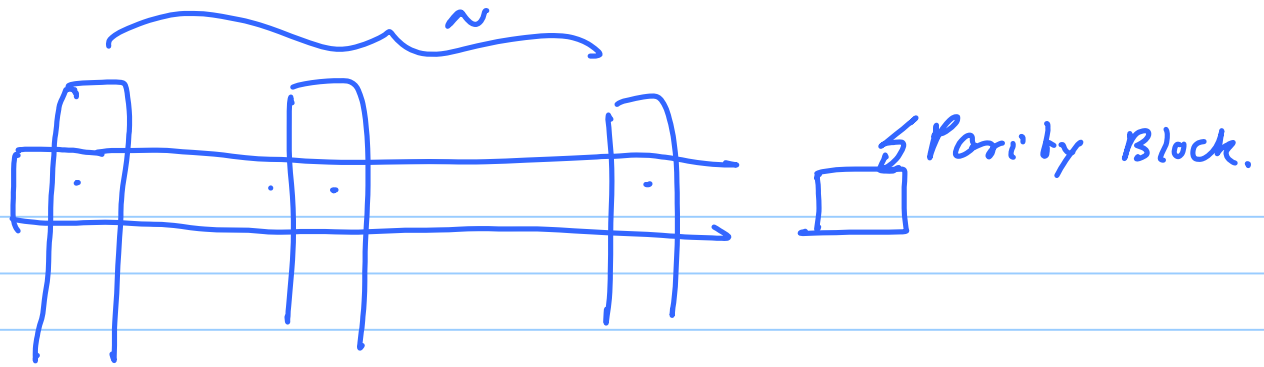
Parity Block.