RAID levels

- RAID – Redundant Array of Independent Disks
  - combines multiple drives into a single large storage

- Data distributed in one of the two ways
  - non-block – contiguously over array disks one by one
  - block – data is cut into the blocks, blocks are written to all the disks in pattern

- RAID Triangle
  - criteria – size, speed, and fault tolerance (F/T) – in vertices
  - RAID levels – on the sides

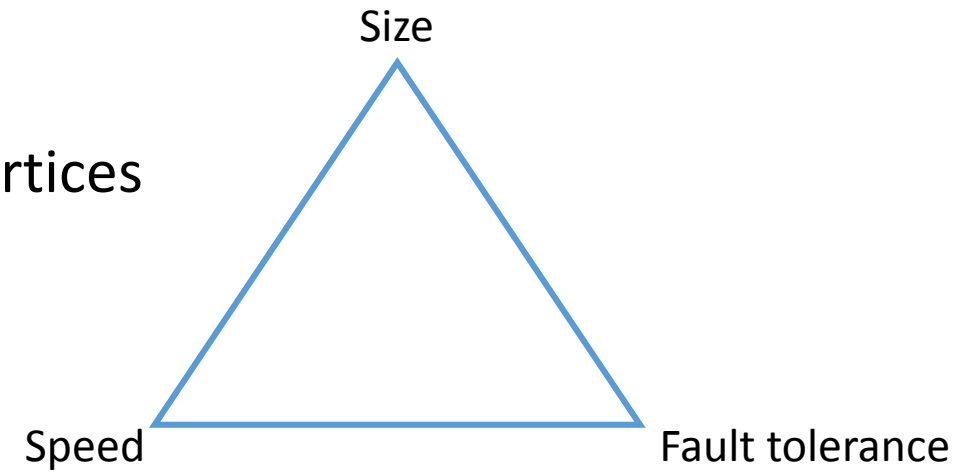

Non-block RAIDs

JBOD – Just a Bunch Of Disks (also called *span*)

- minimum two disks
- provides only capacity
- data is written to the first disk until it is full, then to the next disk

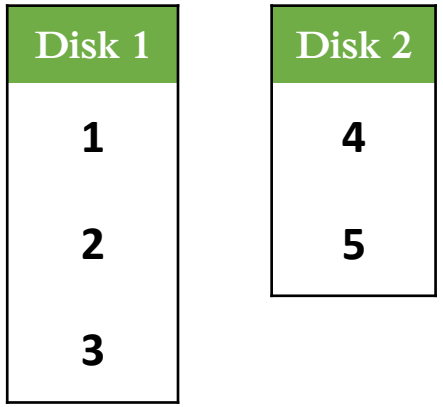

RAID 1 – also called *mirror*

- minimum two disks
- provides fault tolerance and faster reads
- two copies of data are stored on two disks

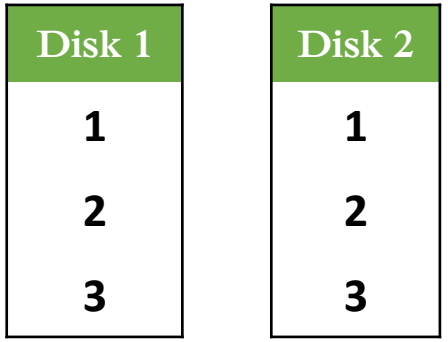

Block RAIDs

Block arrays can be

- non-redundant (RAID 0)
- redundant (RAID 10, RAID 5, RAID 6)

Redundancy can be provided by

- mirroring (writing multiple copies) – RAID 10, RAID 1E
- parity (storing computed extra blocks) – RAID 5, RAID 6

Mirror technique

Disk 1	Disk 2	Disk 3	Disk 4
1	2	1	2
3	4	3	4
5	6	5	6

Parity technique

Disk 1	Disk 2	Disk 3
1	2	P
3	P	4
P	5	6

RAID 0 – Non-redundant block array

Requirements

- at least two disks
- no CPU overhead

Disk 1	Disk 2
1	2
3	4
5	6

Redundancy

- none – any single drive failure destroys all data

Pros and cons

- Faster reads and writes
- No redundancy

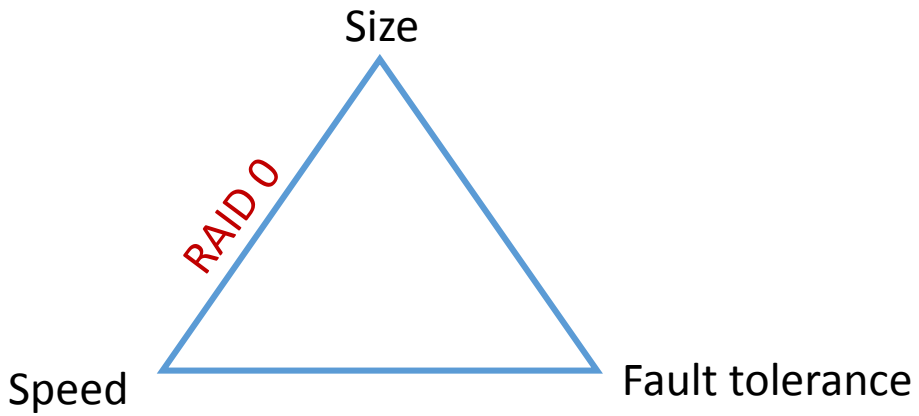

Block, redundant RAIDs - mirroring

RAID 10 (also called RAID 01, 0+1, 1+0)

- minimum four disks
- faster reads and writes
- survives any single drive failure and some double failures

Disk 1	Disk 2	Disk 3	Disk 4
1	2	1	2
3	4	3	4
5	6	5	6

RAID 1E – similar to RAID 10 over odd number of disks

- at least three disks
- faster reads and writes
- survives any single drive failure and some double failures
- multiple variations of the layout

Disk 1	Disk 2	Disk 3
1	1	2
2	3	3
4	4	5
5	6	6

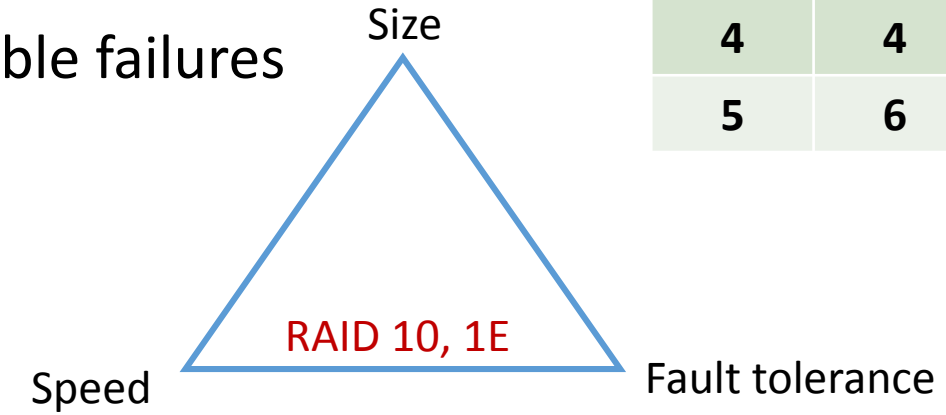

Block, redundant RAIDs – RAID 5

Standard single parity array

- minimum three disks
- faster reads, slower writes
- survives any single drive failure
- uses P-parity (XOR function)

Layout variations

- left and right parity ordering
- symmetric and asymmetric data ordering

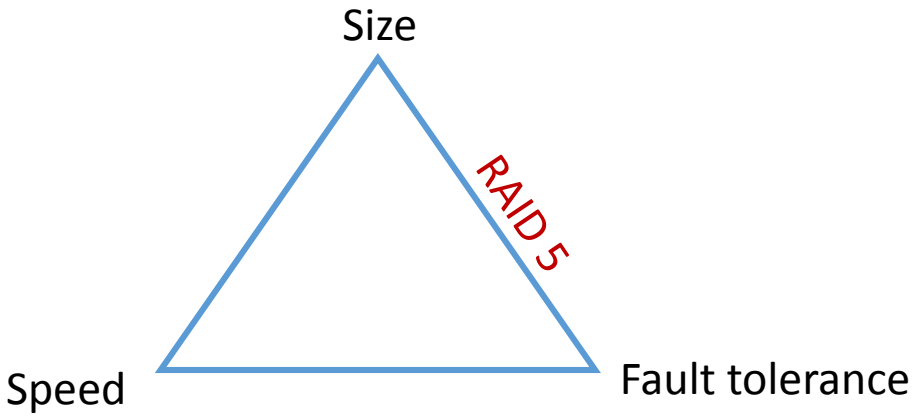

Disk 1	Disk 2	Disk 3
1	2	P
3	P	4
P	5	6

Left asymmetric

Disk 1	Disk 2	Disk 3
1	2	P
4	P	3
P	5	6

Left symmetric

Disk 1	Disk 2	Disk 3
P	1	2
3	P	4
5	6	P

Right asymmetric

Disk 1	Disk 2	Disk 3
P	1	2
4	P	3
5	6	P

Right symmetric

Block, redundant RAIDs – RAID 6

Double parity array

- minimum four disks
- faster reads, much slower writes
- survives any double drive failure
- uses P- and Q- parities (XOR and Reed-Solomon)

Layout variations

- left/right, symmetric/asymmetric same as RAID5
- Q- or P-parity on top in layout
- Calculation order: Q=Q(1, 2, P) or P=P(1, 2, Q)
- Wide pace (*Promise* controllers)

Disk 1	Disk 2	Disk 3	Disk 4
Q			P
		P	Q
	P	Q	
P	Q		

Disk 1	Disk 2	Disk 3	Disk 4
P			Q
		Q	P
	Q	P	
Q	P		

Disk 1	Disk 2	Disk 3	Disk 4
		P	Q
P	Q		
		P	Q
P	Q		

Block, redundant RAIDs – Exotics

Dedicated parity disk – RAID 4

- spare disk becomes a bottleneck

Disk 1	Disk 2	Disk 3	Disk 4
			P
			P
			P

Hot spares

- regular – standalone hot spare drive
- RAID 5E – hot spare blocks at the end
- RAID 5EE – striped hot spare blocks

Disk 1	Disk 2	Disk 3	Disk 4
			P
		P	
	P		
S	S	S	S

Disk 1	Disk 2	Disk 3	Disk 4
		P	S
	P	S	
P	S		
S			P

Delayed parity

- parity blocks larger than data blocks
- *HP SmartArray* controllers

Disk 1	Disk 2	Disk 3
		P
		P
	P	
	P	
P		
P		