

RAID Introduction

RAID – Redundant Array of Independent Disks


- combines multiple drives into a single large storage
- widely used in NASes and servers

Different RAID types are called *levels*

- Most common RAID 1, RAID 0, RAID 10, RAID 5
- Levels trade between speed, size, and fault tolerance

RAID Triangle

- speed
 - size
 - fault tolerance
- } Choose any two
- when criteria are in vertices, sides represent RAID levels



RAID Introduction

RAID can be created

- with hardware controllers (in higher-end servers)
- with software drivers (in NASes and low-end servers)

RAIDs distribute data to disks

- dividing data into blocks (block-based RAID) – RAID 0, RAID 5
- in contiguous streams (non-block RAID) – JBOD, RAID 1

Drawing block RAID layouts

- columns for disks
- numbers for data blocks in their order
- *P* for parity, *Q* for more parity

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

RAID Introduction

RAID metadata

- describes RAID parameters
- stored on the RAID itself, either before or after RAID data

RAID parameters

- set of parameters describes a particular RAID
- different RAID levels have different parameter sets

Common parameters – block RAIDs

- block size
- start offset
- disk order

