

RAID 5 – Parity Checking

RAID 5 - Block level striping, parity interleaved

Striping unit is 1 block: block1 -> disk1, block2 -> block2, block3 -> block3, etc.

Blocks of stripes written at same time!

Transfer rate = $(N - 1) * \text{transfer rate of one disk}$

Only partial redundancy!

Parity information dispersed round-robin among all disks


Same redundancy overhead as level 3, = $1 / N \%$

Hey! *Block striping can* mean that every disk is NOT involved in a (small) request

Parallel reads and writes *can* occur, depends on which disks store involved blocks

- For error detection, rather than full redundancy
- Each stripe unit has an extra parity stripe
- Parity stripes are distributed

RAID Level 5: Parity Checking



RAID 5 Read/Write

Read: parallel stripes read from multiple disks

- Good performance

Write: 2 reads + 2 writes

- Read old data stripe; read parity stripe (2 reads)
- XOR old data stripe with new data stripe
- XOR result into parity stripe
- Write new data stripe and new parity stripe (2 writes)

RAID 1 vs RAID 5

- o Reads:
- o RAID 1 (mirroring):
 - o always offers parallel reads

RAID 5:

- o Can only sometimes offer parallel reads
- o Depends on where the needed blocks are
- o Two read requests that require blocks on the same disk must be serialized

- o Writes:

RAID 1:

- o Mirroring) must complete two writes before request returns
- o Granularity of serialization can be smaller than a file
- o Can't do parallel writes

RAID 5:

- o Typically does read-modify-write to re-compute parity
- o HP AutoRAID uses combo of read-modify-write and LFS !
- o Can't do parallel writes either

Redundancy Overhead:

- o RAID 1 = full redundancy, storage capacity reduced by 50%
- o RAID 5 = partial redundancy, storage capacity reduced by $1/N\%$