Shingled Disk Drives

File System Vs. Autonomous Block Device

ROBERT E. NOVAK | Director of Systems Architecture

JUNE 2014

Three Versions of SMR Firmware

- Drive Managed Firmware
  - Drive will emulate a random I/O drive
  - Response times will be have more variability than traditional Random Drives
  - For Enterprise transactional profiles - Challenging to develop & QA from a drive manufacturer perspective
- Host Aware Firmware
  - Drive prevents accidental overwrite
  - Drive performs limited random I/O remapping
  - Less challenging than full random I/O
- Host Managed Firmware
  - All aspects of reads/writes are managed by Host
  - Conventional zones possible on I.D. and O.D. of drive
  - Drive prevents accidental overwrite
  - Simplest firmware

File Systems today are using 1890 technology

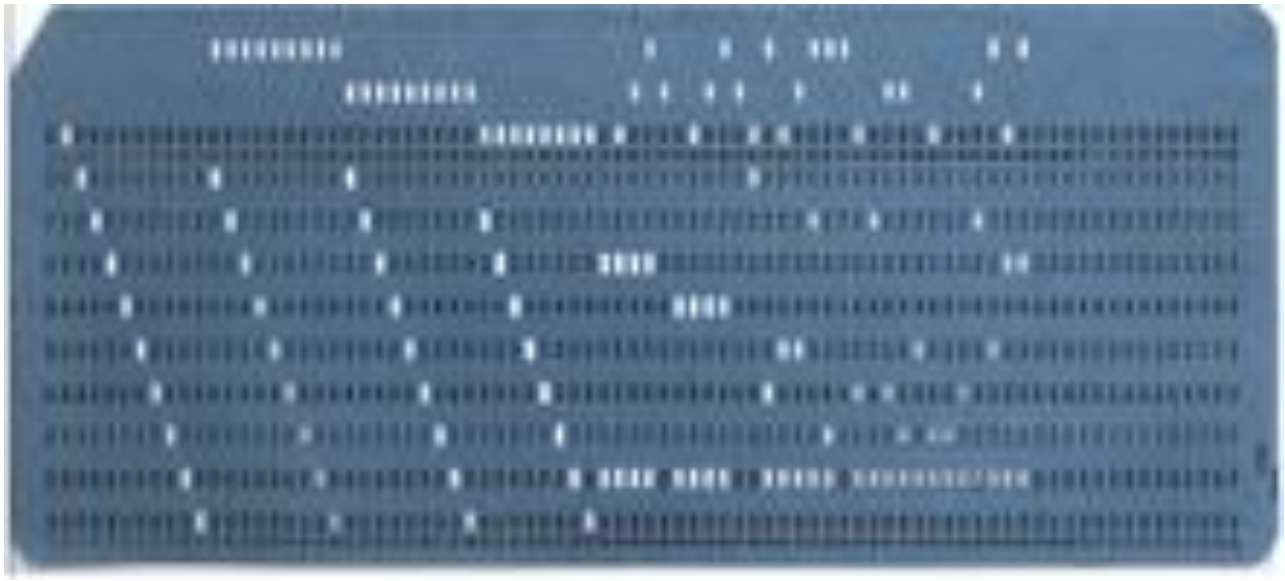

What do Blocks have to do with Objects/Files?

- Blocks are for convenience in addressing
- Why are blocks visible outside of the drive?
- Block addresses lead to problems (e.g., FSCK, CHKDSK, BP-Rewrite for ZFS)
- Blocks mean data can’t move

*Omnes relinquite spes, o vos intrantes

*Abandon hope all ye who enter – Dante

"Inferno" – The inscription on the gate to hell – It. *Lasciate ogne speranza, voi ch'intrate*

- Abandon Blocks outside the device AKA Key/Value
- Put enough intelligence in the device to worry about internal and external fragmentation of data as well as garbage collection
- Keep reorganization of data internal to the drive
- Give the host hints on when reorganization is needed
- This requires that the host can live without the drive for some period of time
- In replication clusters, this should be OK, even if the host does not direct the drive when to "be unresponsive for awhile"
- This works best with Ethernet attached drives
- SAS/SATA can emulate a Key/Value drive