

A blurred photograph of a crowd of people walking across a street with white zebra crossing stripes. The motion blur suggests a busy, fast-paced environment.

SMR: Moving to Greener Pastures

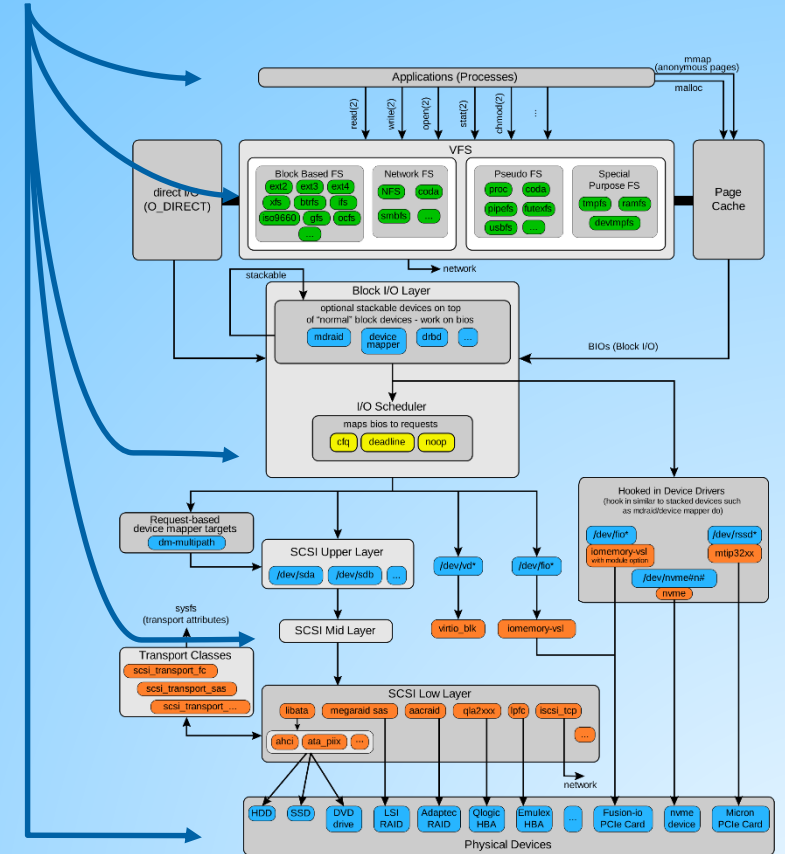
Seagate's effort to vitalize SMR into more generic environments

SMR – Shingled Magnetic Recording

What we know, and what we don't know

- ✓ Mass rotating media storage
- ✓ Forward-Write only
- ? ZBC/ZAC zones partitioning drive
- ? DM vs HA vs HM
- ? Requires new thinking of IO patterns
- ! Requires massive stack changes

Needs work!



Flash, but with Slower speeds

Flash

- Sequential Write
- Erasure Blocks
- ERASE (Blocks)
- FTL – Virtualize LBAs
- TRIM (individual LBAs)

SMR

- Sequential Write
- Zones
- No lengthy Erase operation
- Drive Managed Architecture
- RESET_WRITE_POINTER (Zone)

IO Stack support

Changes **REQUIRED** for generic use


Host responsible for translation layer

In drive: ZAC/ZBC zone information

In kernel: new ATA/SCSI commands

In FS: new allocation and IO management

In RAID: new compositions requirements



Seagate's Open Source Contributions

SMR Friendly File System - SMRFFS

Kernel IO Stack

- ZAC/ZBD aware reference design

FS

- EXT4 reference design

Device Mapper

- ZAC/ZBD aware RAID/JBOD solution



For more information...

ask questions!

Email

- adrian.palmer@seagate.com

Clone

- https://github.com/Seagate/SMR_FS-EXT4

Talk with me

- Adrian Palmer