

Open Source Storage

Does Use Case Matter?

Michael Letschin
Product Manager, Nexenta
[@mletschin](#)

Drastic Growth....

- Storage is growing at 2.7 zettabytes/year¹
- HDD Shipments growth 30%/year¹
- HDD Areal Density growth 20%/year¹
- New drive tech improves areal density by 25%¹

¹<http://www.storagenewsletter.com/news/disk/seagate-shingled-magnetic-recording>



Is your application special?



Application Centric Storage


- General Purpose Storage is DEAD
- Applications have custom storage requirements
- Storage should change with application load
- Augments “virtual first” mentalities
- Availability is integrated with infrastructure



How does that change the hardware?



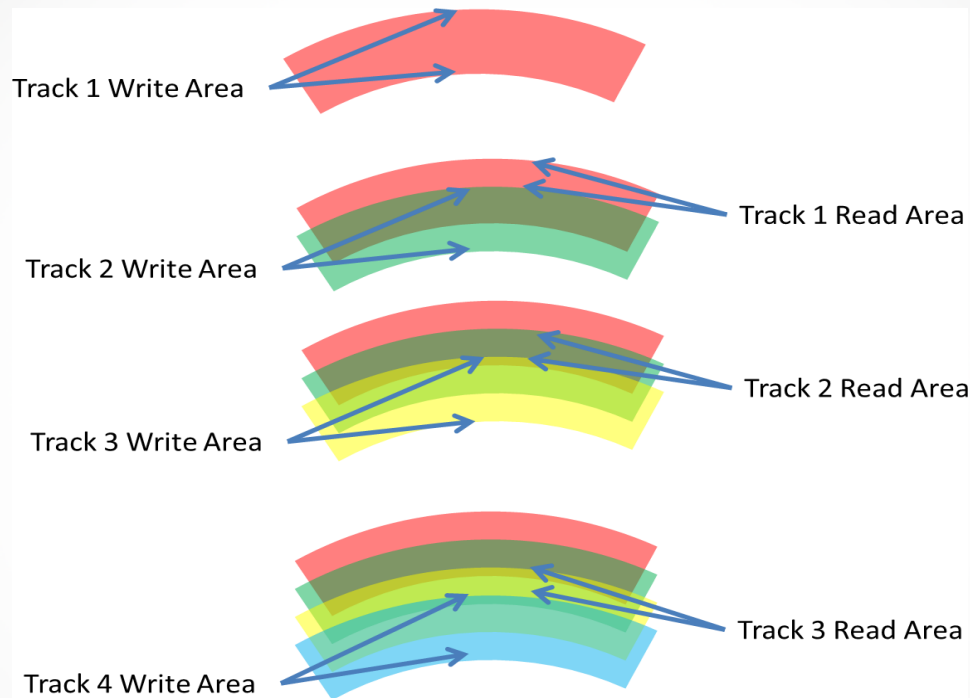
Traditional Drives



- Data areas – white
- Guarded by gaps –gray
- Can randomly read/write any data area




LAYING DOWN SHINGLES (LIKE A ROOF)



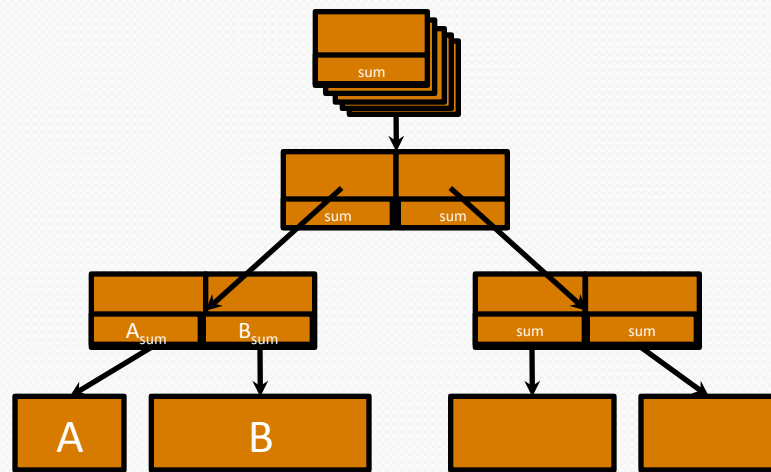
Shingled Magnetic Recording

- Shingles are grouped in bands
- Copy from one band to another to consolidate/delete
- Random I/O in outer zone and inner zone
- Perfect organization for ZFS Uberblocks

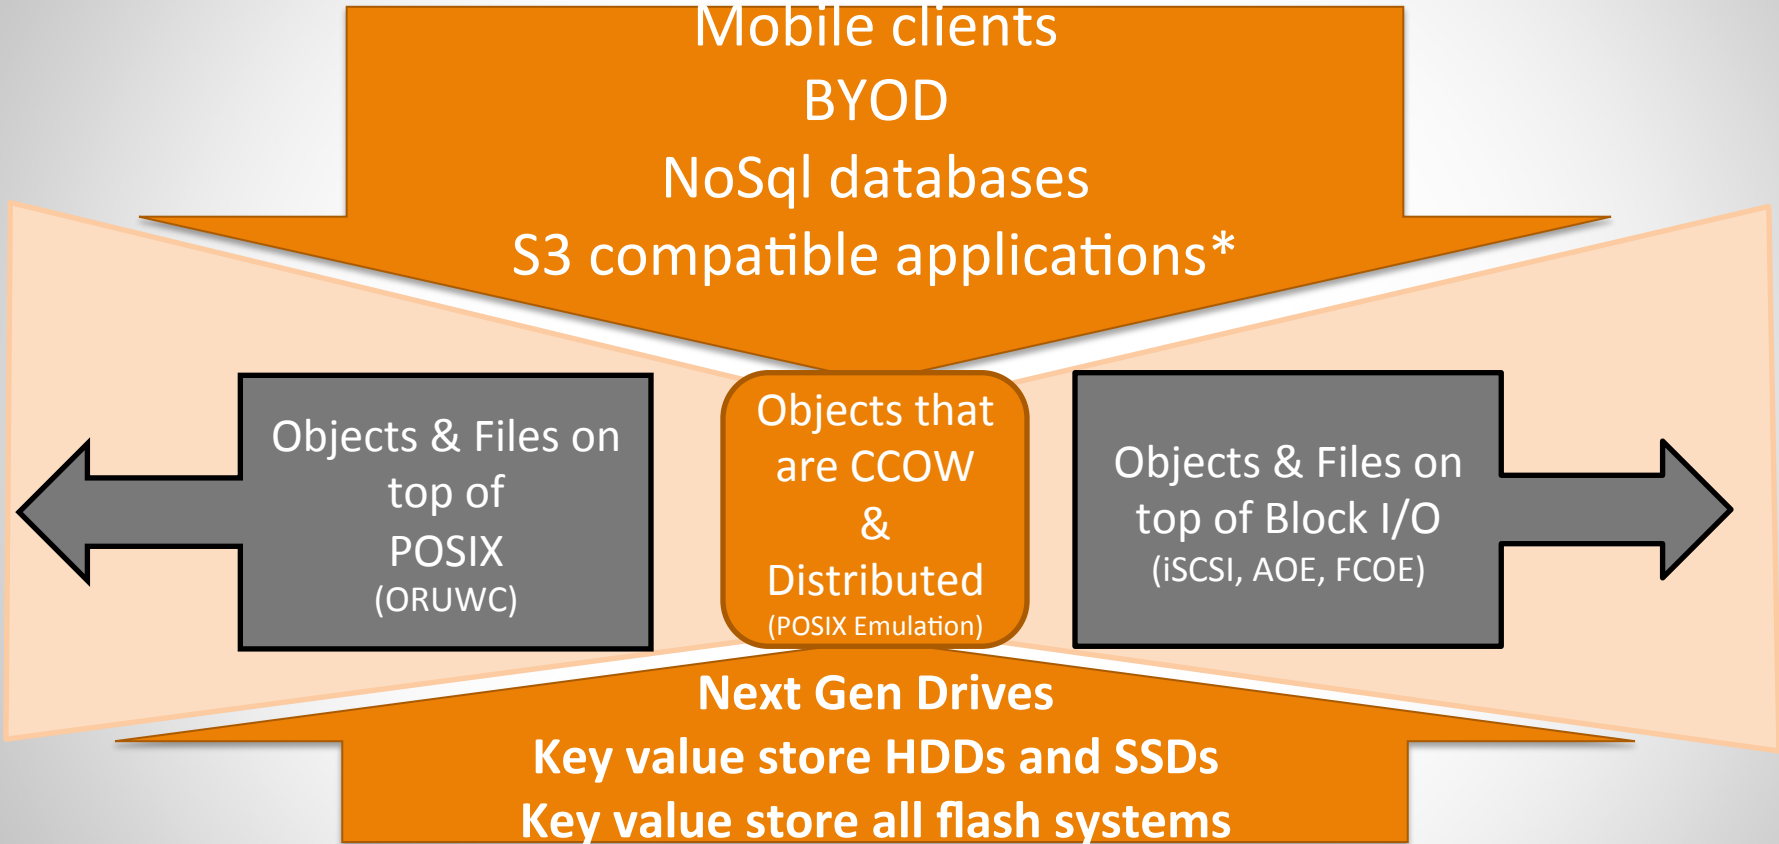


ZFS Highlights

- 🌀 **HUGE Capacity: 128-bit file system**
 - 🌀 16 billion-billion times that of a 64-bit file system
- 🌀 **Hybrid Pooled Storage: Automatic & Intelligent use of DRAM, SSD, & HDD**
 - 🌀 File system + volume manager + RAID built together
- 🌀 **Always Consistent on Disk**
 - 🌀 Copy on write; no journaling; no fsck ever!
- 🌀 **Unmatched Data Integrity**
 - 🌀 Data checksummed at every point
 - 🌀 Detects & **corrects** silent data corruption
- 🌀 **Simple & Efficient Snapshots, Clones, & Replication**
- 🌀 **Simple & Fast Initialization and Resilvering**
 - 🌀 Increasingly important with growing HDDs



TREND TO OBJECT STORAGE



Future of Open Source Archival

- HDD will give way to SMR as the dominant media
- Archive is not big metal storage – Needs Application Centric Storage
- Data Integrity over time is the key
- Self Healing Filesystems dominate the archival space





Michael Letschin

mletschin@nexenta.com

@mletschin

- **3.1** What do you see as the role of open source software in the current and future options for inter-operability and portability in storage architectures?
- **3.2** What do you think large archive users and planners should be doing if they are interested in expanding the use of open source software for these architectures?
- **3.3** What interfaces do you think will be used in archival storage in 1, 5, 10 years?
 - a. REST/SOAP (AWS, object)
 - b. POSIX (NFS, CIFS, file system)

