




Leveraging Ceph on Next Gen Storage

Jamon Bowen

Seagate

Director, Customer Management



Ceph Deployment

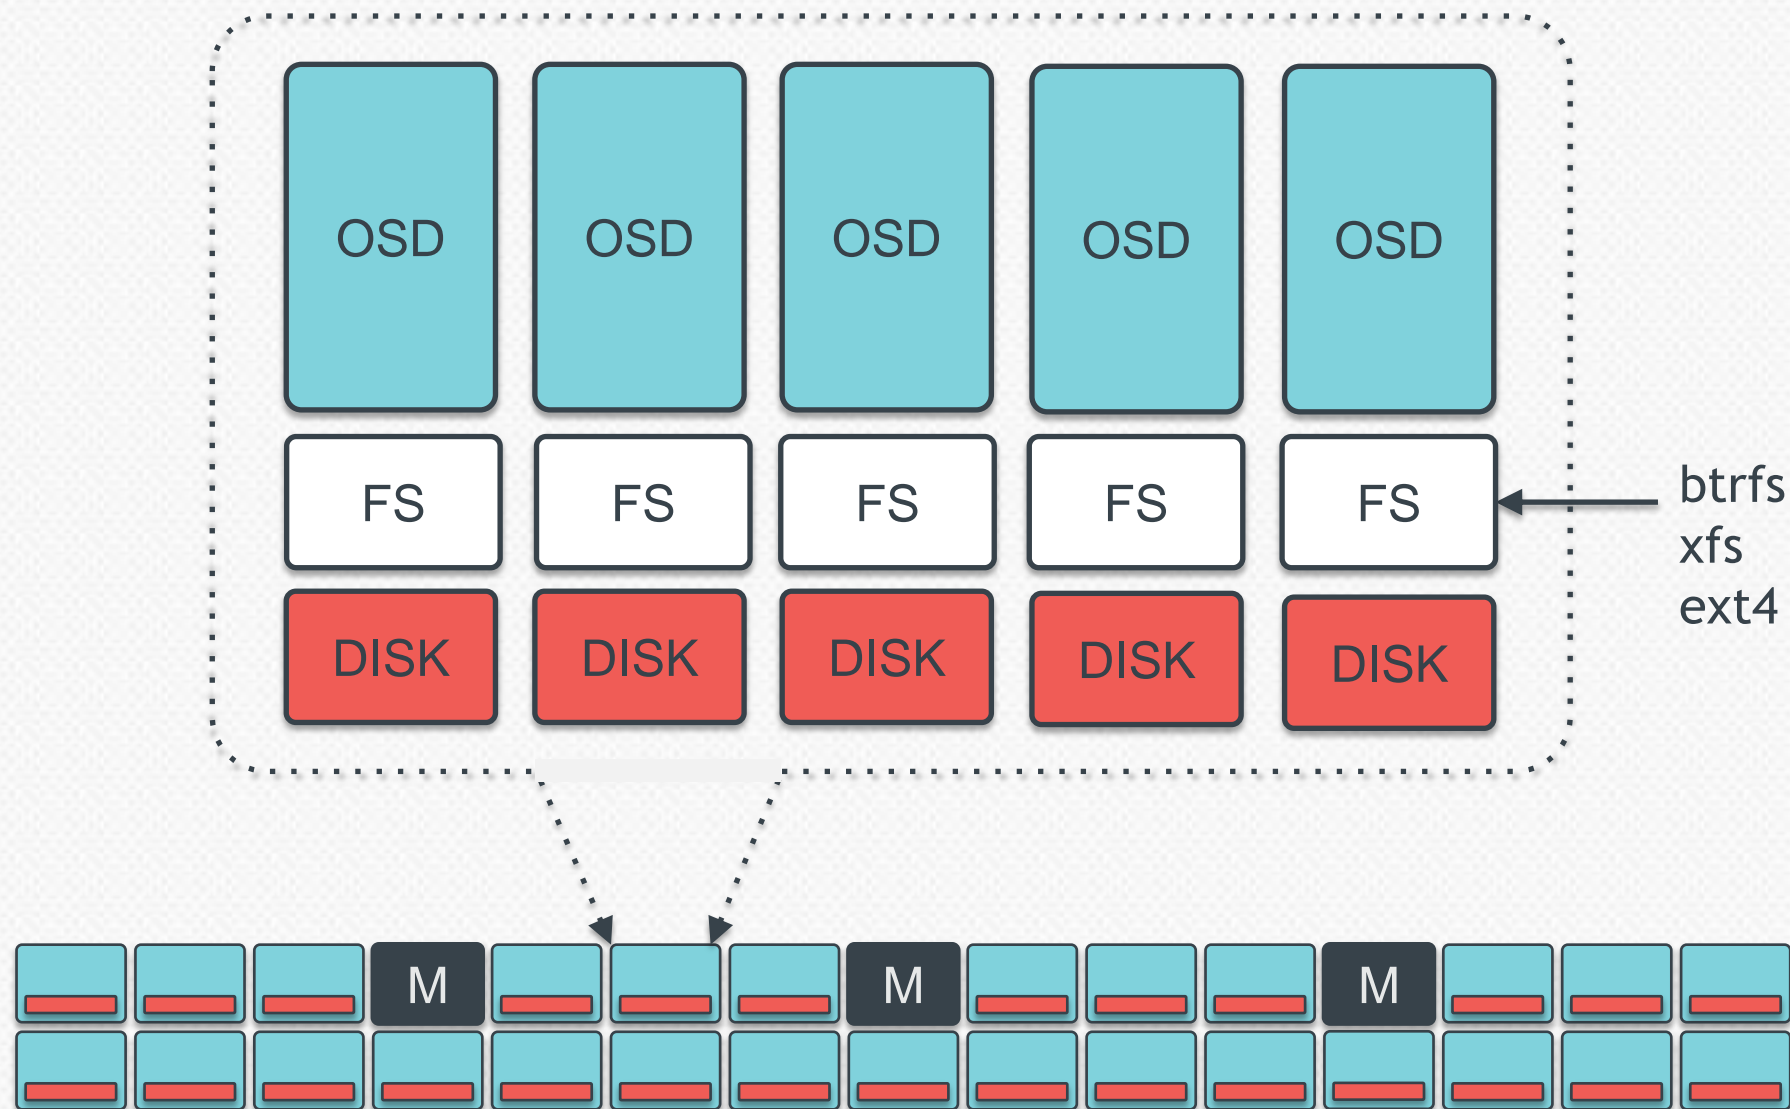


Object Storage Daemons (OSDs):

- 10s to 10000s of OSDs in a cluster
- One per disk or SSD
- Serves stored objects to clients
- Intelligently peers to perform replication and recovery tasks
- Support for erasure encoding of objects
- Support for strong consistency leveraging write journals

Advanced Storage:

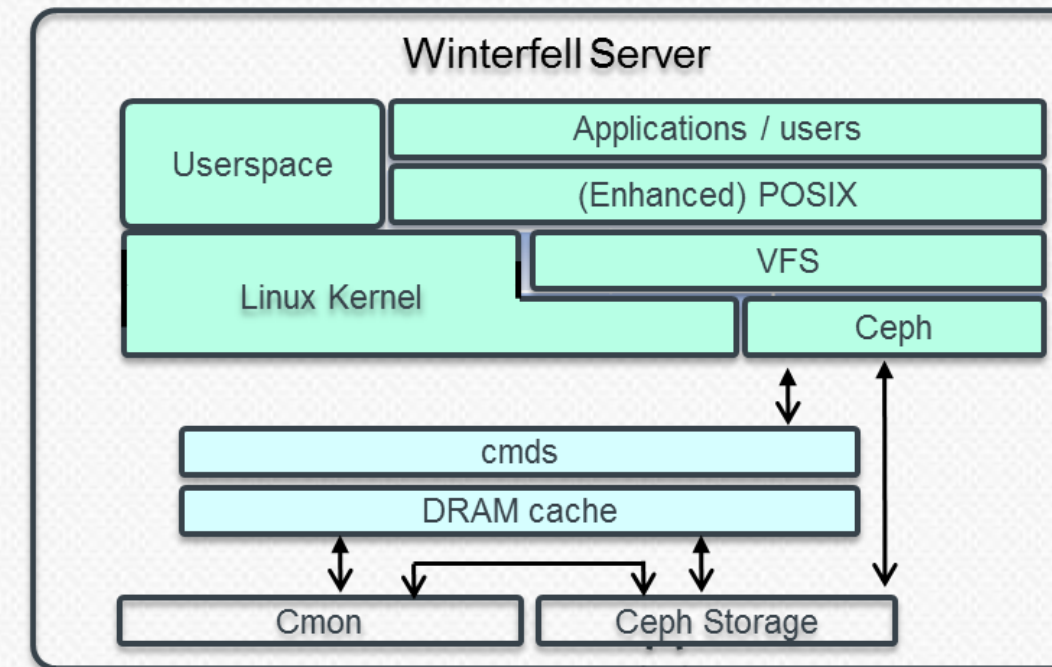
- Shingled Magnetic Recording (SMR) Drives (drive managed)
 - Excellent density, cost, read performance, and sequential write streaming.
 - Poor random write performance.
- PCIe Flash
 - Excellent performance, in server integration
 - Poor cost per GB.



Storage Configuration

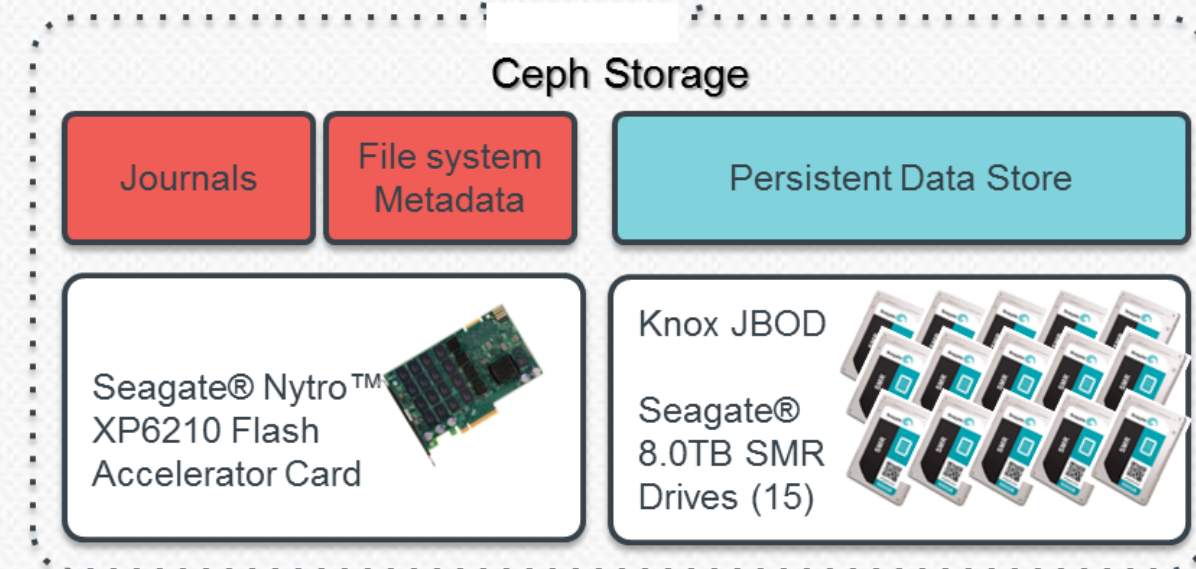
Futures:

- Filter driver small IO cache to flash integration
- Ceph cache tiering
- Host aware SMR plus flash



Seagate storage technology


- Accelerate Ceph journals and Metadata with Nytro PCIe flash
- Achieves 120TB of storage per U Seagate SMR drives




Scaled up Hardware Configuration

Per Open Rack:

- 18 OSDs
- 33.5 TBs of Flash
- 2.2 PBs of SMR Storage





OPEN

Compute Engineering Workshop

March 9, 2015

San Jose