



SGI Data Management DMF JBFS

Zsolt Ferenczy/SE
zsolt@sgi.com



Current DMF 6 Environment

HPC Compute Nodes

High-Performance Parallel File Systems
XFS / CXFS & Lustre



SGI DMF
Scalable Data Migration Facility



- Millions of Files
- Focus on Archive
- Tape-Centric

Tape Libraries (hardware)



Lowest Cost & High Durability

Cold Storage
with SGI JBFS



Low Cost & High Performance

Cloud / Object Storage




High Scalability &
Geo-Distribution

DMF 6 JBFS Backend

HPC Compute Nodes

High-Performance Parallel File Systems
XFS / CXFS & Lustre



SGI DMF
Scalable Data Migration Facility




- Millions of Files
- Focus on Archive
- Tape-Centric

Tape Libraries (hardware)



Lowest Cost & High Durability

**Cold Storage
with SGI JBFS**



Low Cost & High Performance

Cloud / Object Storage




High Scalability &
Geo-Distribution

SGI JBFS: High-Performance & Cost Optimized for SGI DMF

- **JBFS** is an acronym for **JBOD File System**
- SGI JBFS provides mounting services and serial access to disk media independent from Linux disk mounts and file systems
- Leverages the significant data management software IP that came with SGI's acquisition of COPAN
- SGI JBFS enables rich capabilities for power management and focuses on management of tiered versus primary data
- Ability to deliver data access and I/O performance significantly beyond alternatives

SGI JBFS

- Any Number of LUNs or Devices
- Full Power Control
- Recoverability
- High-Performance
- Flexible to Many Media Types



DMF Cold Storage

- 90-Bay 12Gb/sec JBOD
- Delivered fully populated with matching firmware drives (6TB or 8TB)
- 540TB or 720TB capacity
- Individual drive power control
- Integrated with SGI DMF and SGI JBFS
- Delivers 'cost correct' archive storage with very low operating costs
- Planned release as a service addition to ISSP 3.5 in March-April 2016 timeframe

SGI IS496J Data Sheet

SGI InfiniteStorage IS496J Cold Storage JBOD

Robust High-Density Cold Storage & Archival Solution

Key Features

- Support for 96 SATA Drives in 4U Enclosure
- Daisy-chain Expansion & Multi-Zone Support
- Fully Integrated & Certified with SGI DMF



A Scalable Solution for Cold Storage

The IS496J Cold Storage JBOD is a robust and scalable storage platform for the data center. Targeted for use cases that require the securing and management of petabyte-scale quantities of less active data - also known as "cold data" - the IS496J delivers a very dense arrangement of high-capacity drives within an enclosure that supports advanced power management capabilities.

The result is a cold storage solution that enables extremely low storage cost per gigabyte of data under management - and lower operating costs through reduced power and cooling requirements.

Based on industry standard SAS 3.0 technology, the IS496J is compatible with a broad array of storage applications that require large capacity JBOD-based storage.

When combined with SGI's DMF tiered data management platform, the IS496J creates an on-premise solution based on a petascale design that is ideal for data centers serving markets including science & technology, healthcare, government, manufacturing and media & entertainment.

SGI DMF with the IS496J is a complete, low cost per GB disk-based solution; designed, integrated, installed and supported by a single vendor.

Product Highlights

- High density cold storage and archival enclosure
- Fully integrated and optimized for use with SGI DMF tiered data management and the SGI JBFS storage management interface
- Robust and full-featured 12 Gb/sec SAS 3.0 storage enclosure
- Carrier-less, hot swappable hard drive installation enhances ease of drive installation and servicing
- Delivered fully-populated with certified 3.5" SATA drives - configured with matching firmware levels - for trouble-free implementation
- Ability to zone the enclosure for maximum flexibility
- Includes a single JBOD SAS 3.0 I/O controller module
- N+1 Hot-pluggable power modules
- Redundant, hot swappable system cooling fans
- Two SFF-8644 12Gb/sec SAS 3.0 host and expansion connections support up to 4,800MB/sec bandwidth each
- SAS daisy-chain expansion to additional IS496J storage arrays or other SFF-8644 SAS compliant host/expansion ports



sgi

©2013-2014 Silicon Graphics International Corp. All rights reserved. SGI, the SGI logo, InfiniteStorage, JBFS and DMF are registered trademarks or trademarks of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries. Intel, Xeon and the Intel Xeon logo are registered trademarks of Intel Corporation. All other trademarks are property of their respective holders. 06122013-4366 0302015

DMF Cold Storage Features

- Hardware specs:
 - 4U 90 3.5" bay dual controller JBOD
 - SAS drive technology
 - 12 Gb dual path SAS interconnect
 - Up to 20 GB/s possible bandwidth
- Reliability, Accessibility, Serviceability
 - 90 hot-swap HDDs & tool-less drive trays
 - 4 hot-swap power supplies
 - 5 hot-swap fan trays
 - 2 hot-swap expander modules
 - Remote reboot, central management and firmware upgrade
 - Hardware monitoring over IPMI

90-Drive ZeroWatt™ JBOD




Hot-Swappable Expander Modules



Tool-less Drive Tray Design



Slide Rails with Cable Management Arm



JBFS Features


- JBFS is implemented in OpenVault
- Can use any block device as a “*virtual tape drive*”
- A collection of block devices become a “*virtual tape library*”
- A block device can hold multiple “*virtual tape cartridges*”
- A JBFS library can have as many virtual drives as there are virtual cartridges

DMF Cold Storage Features


- The SGI product leveraging JBFS technology
- A combined hardware and software solution
- One JBFS “*library*” per JBODs connected to the same set of DMF servers or PDMOs
- One JBFS “*cartridge*” per drive
- Individual drive power management
- Active/Active HA

DMF Cold Storage Configurations


Basic DMF



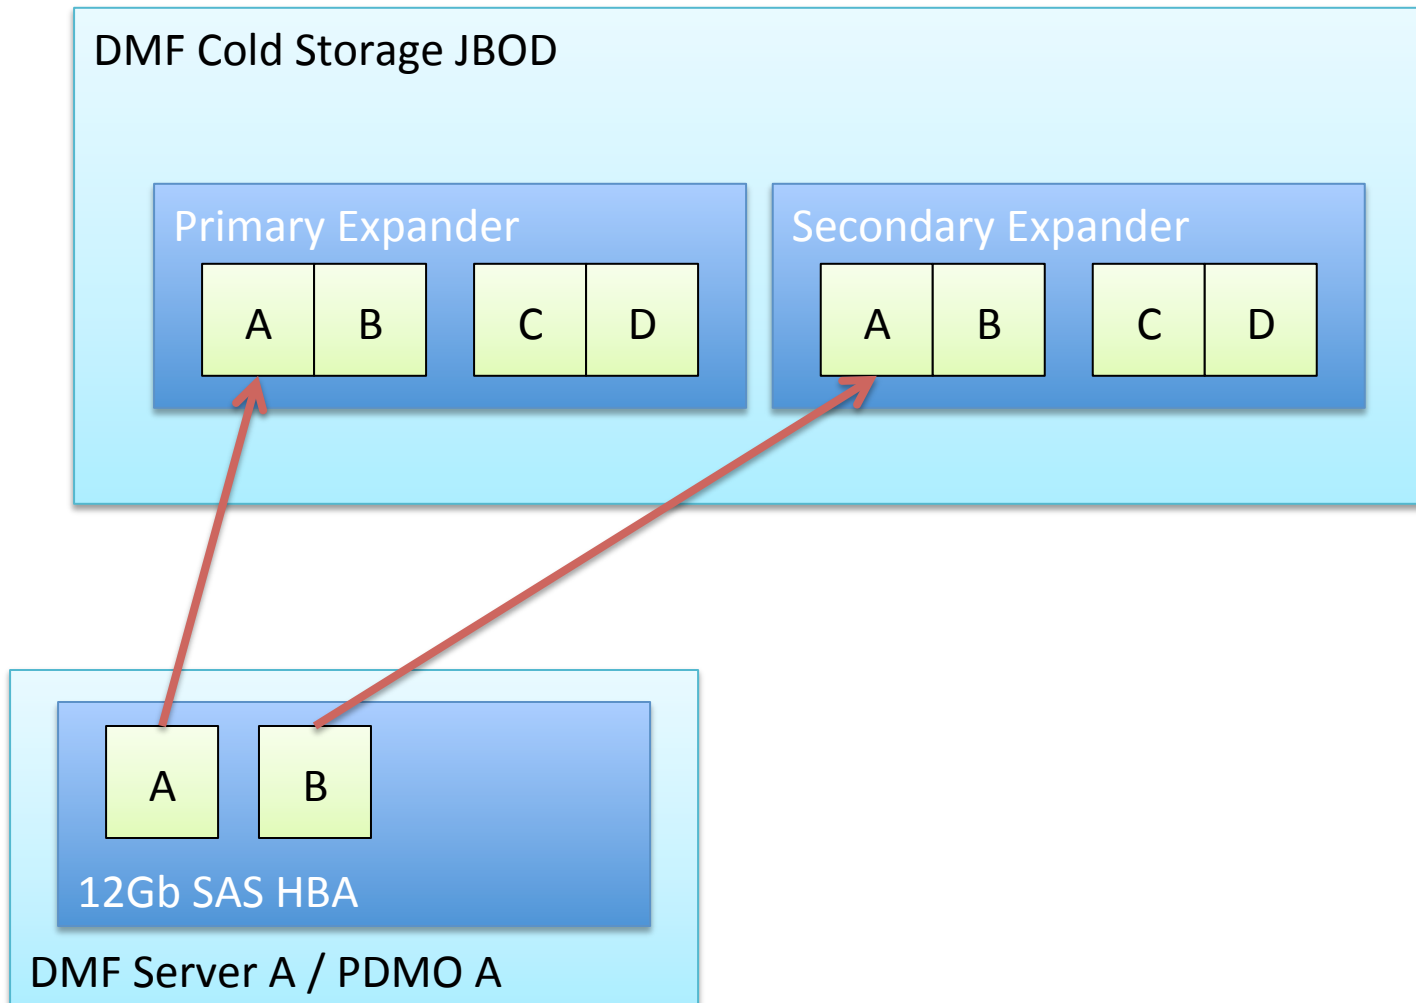
Basic DMF HA




Parallel DMF



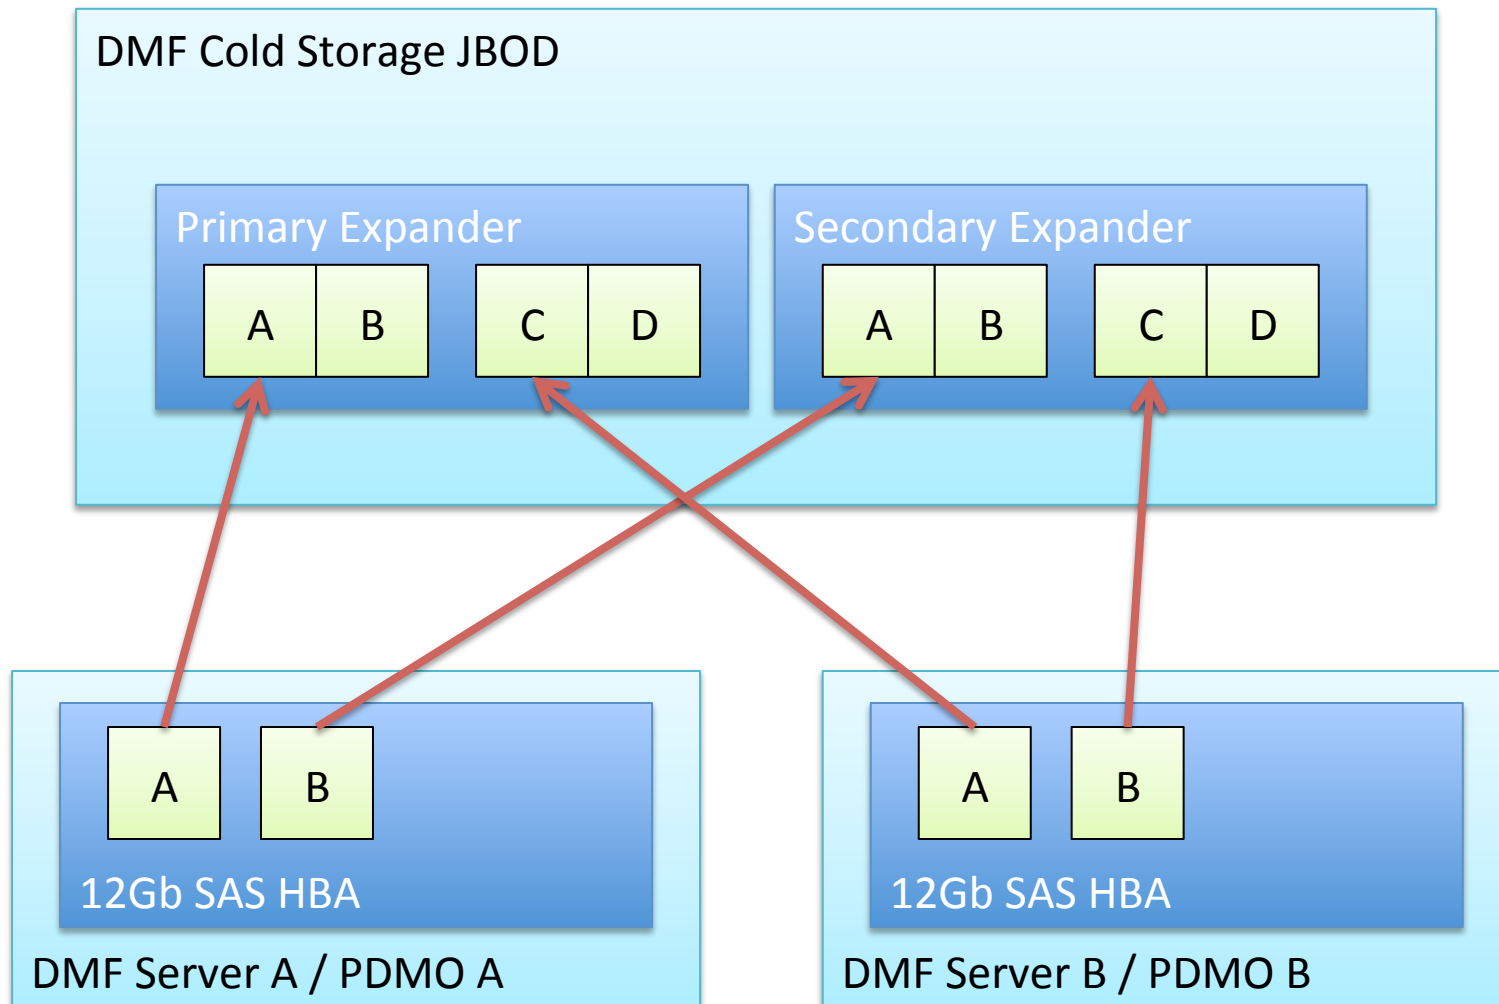
DMF Cold Storage Cabling




DMF Cold Storage Cabling



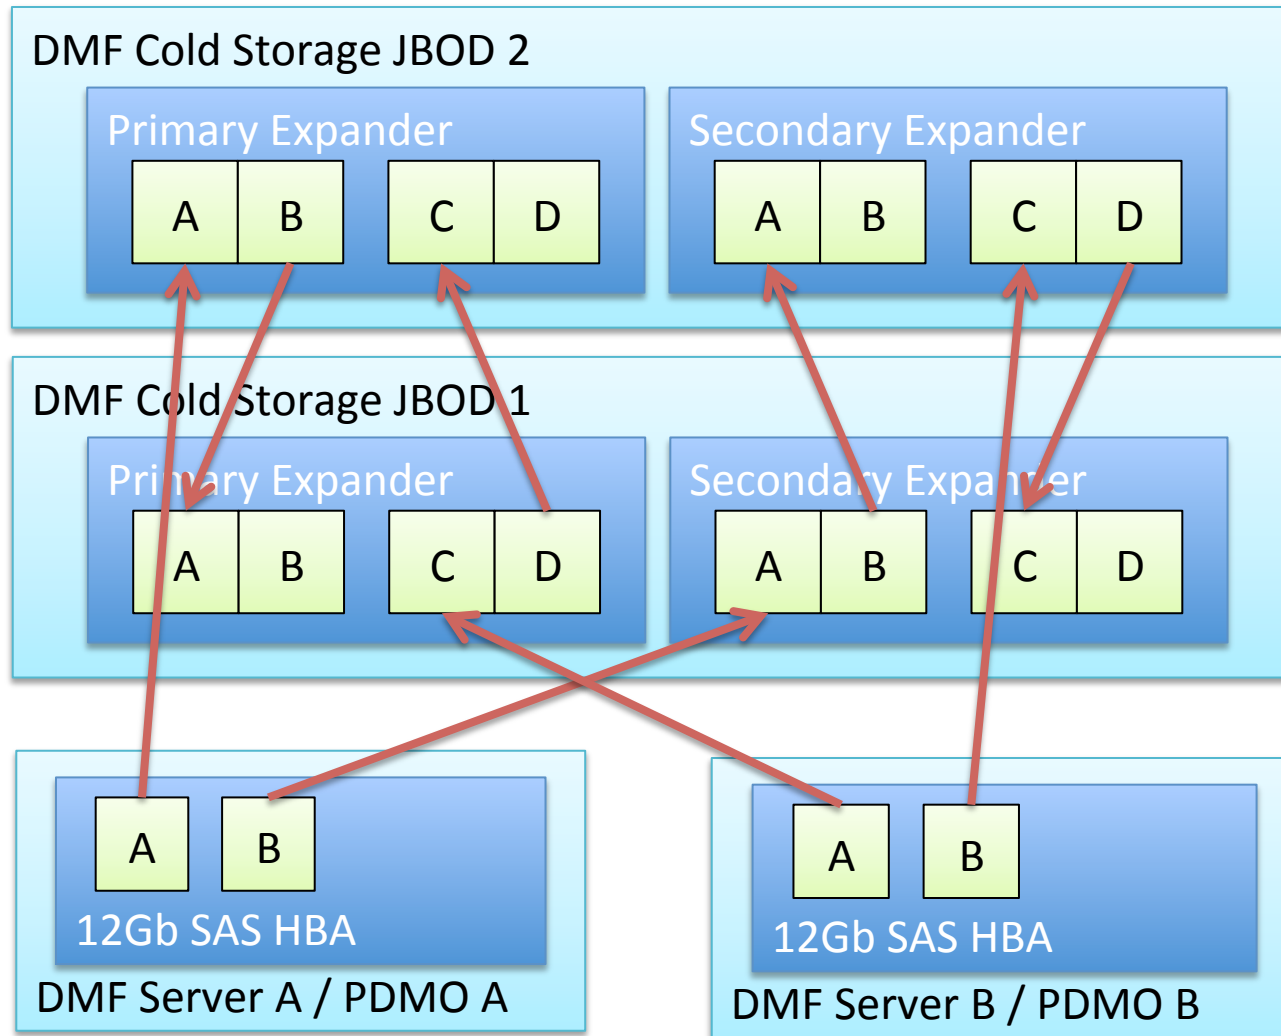
DMF Cold Storage HA Cabling



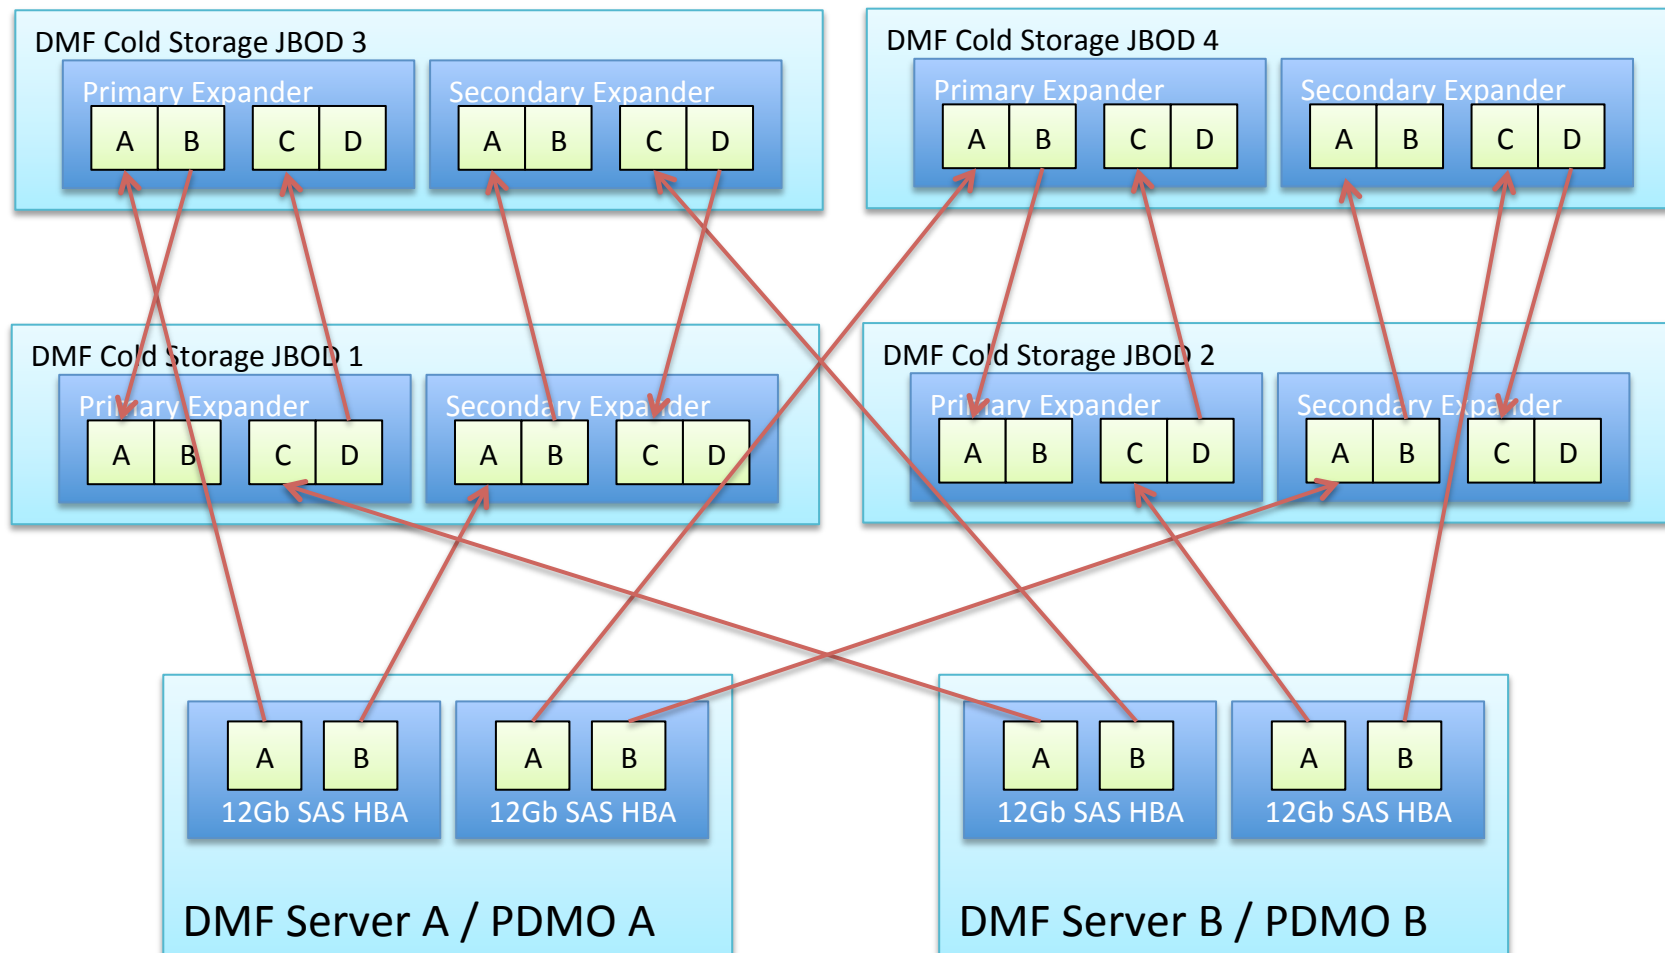
DMF Cold Storage HA Cabling



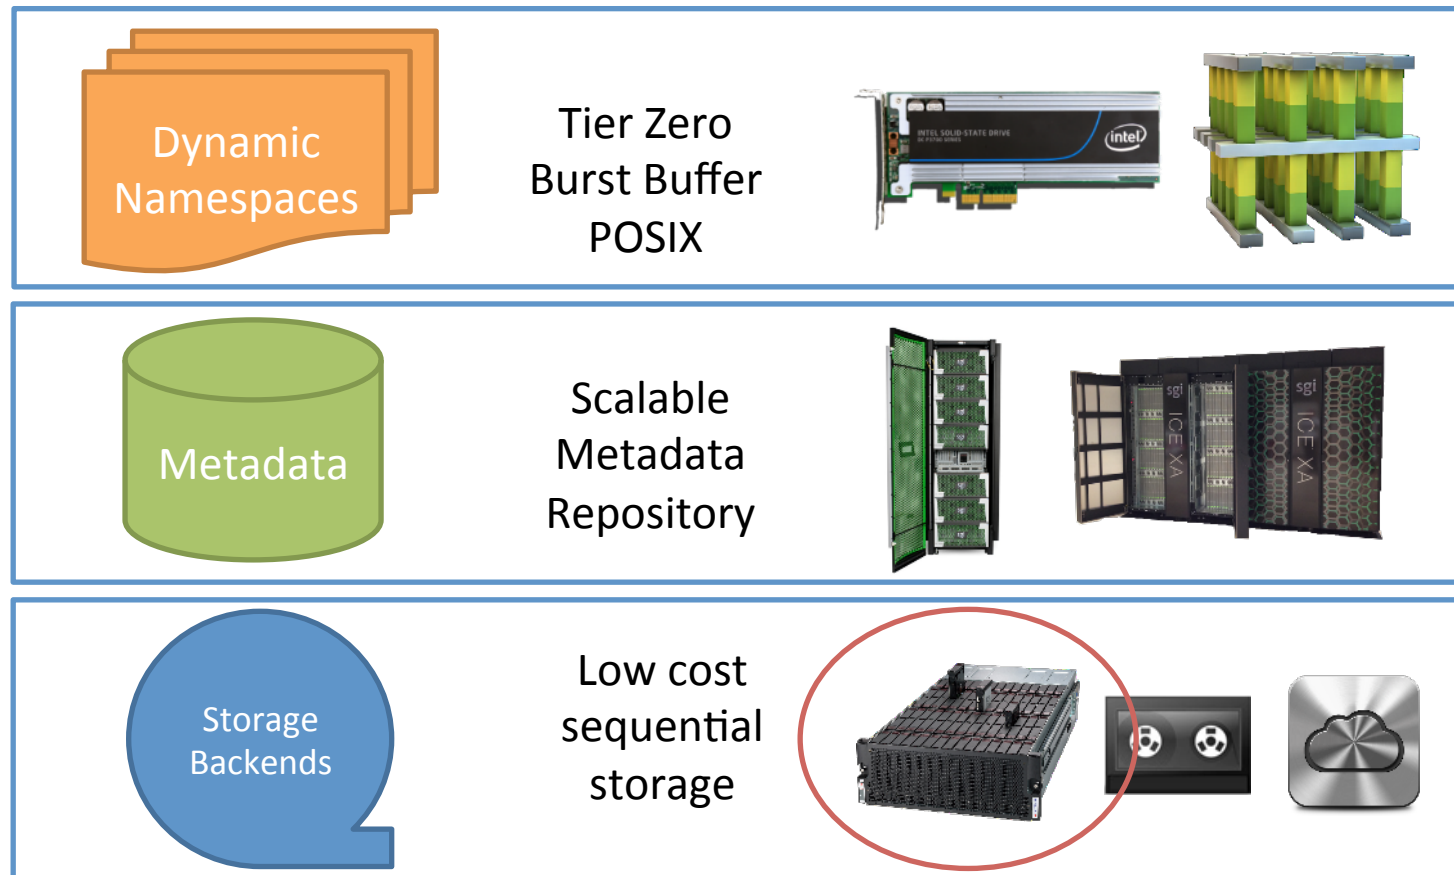
DMF Dual JBOD HA Cabling



DMF Quad JBOD HA Cabling



Data Management Ecosystem



Data Management for HPC with DMF 7: Vision

PBS Professional®  **slurm**
workload manager

HPC and HPDA Compute Nodes



IME

Burst Buffer



NetApp

High-Performance
Parallel File Systems
Lustre, GPFS


DDN
STORAGE



SEAGATE

Tier Zero: Dynamic POSIX Namespaces


On-demand filesystems collocated
with compute and managed by DMF




sgi



**OpenStack
Swift API
Containers**



openstack™



SGI DMF v7

Scalable Data Management Platform



Tape Libraries




SPECTRA

ORACLE

IBM

Lowest Cost & High Durability

**Cold Storage
with SGI JBFS**



sgi

Low Cost & High Performance

Cloud / Object Storage



Private & Public
DDN WOS
NetApp
StorageGRID
Amazon S3

WOS

S3

High Scalability and Geo-Distribution

