

DMF Feature Update

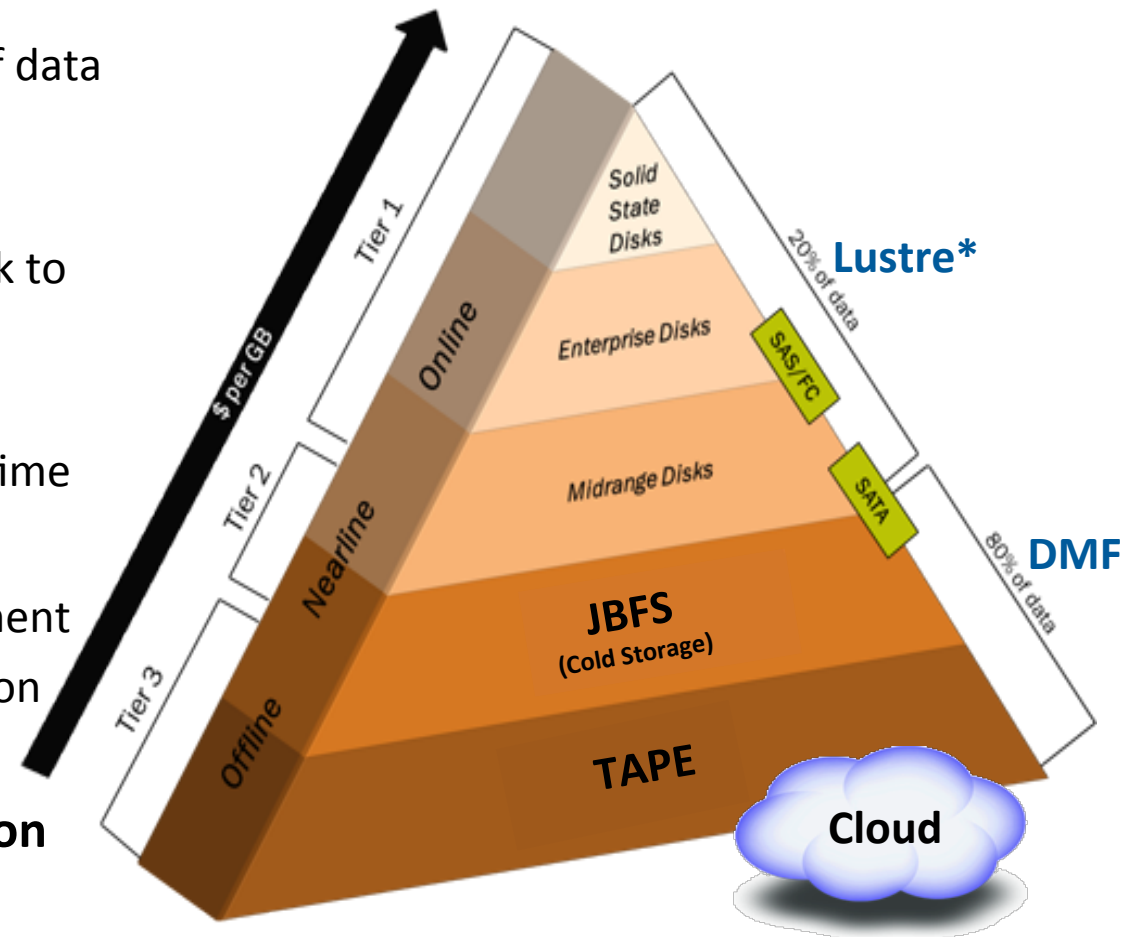
Robert Mollard

Senior Storage Specialist, APAC

HSM | Data Migration Facility (DMF)

Hierarchical Storage Management
Transparently migrate data to Tape, MAID or Cloud

- **Data life cycle management**
 - DMF manages the placement of data within multiple tiers of storage
- **Automated data migration**
 - From expensive, production disk to 2nd or 3rd tier storage
- **Transparent to user**
 - All data appears on line all the time
- **Key Benefits**
 - DMF reduces tier 1 disk investment
 - DMF reduces power consumption
 - DMF protects data long term
- **SGI® DMF™ 25 years in production**



SGI® DMF™ – 25 Years in Production

- **1989** – First DMF Installation at Forschungszentrum Juelich GmbH (KFA), Germany as a Cray Research project.
- **1990** – DMF installed at NOAA GFDL. GFDL is still using DMF to manage over 80PBs of data.
- **1991** – DMF installed at CSIRO Australia – Continuous online data until today, expanding with data and technology.
 - DMF first installed at NASA. Today this has expanded to three different NASA sites, and over 120PB of online data.
- **1996** – SGI acquires Cray Research. DMF is at version 2.5
- **1997** – SGI ports DMF to IRIX with the release of DMF version 2.6
- **1999**– Weta Digital chooses DMF to manage the data for The Lord of the Rings. Nearly 2 billion files and more than 24PB under management today.
- **1999** – SGI announced that XFS will be released as Open Source under the GPL
- **2000** – SGI spins out most of Cray, but keeps DMF and the storage software
- **2002** – DMF 2.7 for IRIX and TRIX is releases with lots of new features
- **2003** – DMF 2.8 is released. First version to support Linux on IA64
- **2003** – DMF 3.0 is released. Adds a third tier of storage with Cache Disk Support
- **2004** – DMF 3.1 is released. Includes file tags & site-defined policies
- **2005** – DMF 3.2 is released. Partial file recall and migrate introduced. Partial file restore used today at the NBA, to accelerate production access to 24PB of online data, growing at 40PB per day.
- **2007** – DMF 3.8 is released. First release of DMF to support the Linux x86_64 platform
- **2008** – DMF 3.9, DMF 3.10 & DMF 3.11 released. Added support for HA with Heartbeat
- **2009** – DMF 4.0 released. Parallel Data Mover option was added. Web based management tool introduced.
- **2010** – DMF 5.0 released. SOAP WS API introduced. Migration to/from non-DMF managed filesystem. Overland library support.
 - V.2 LCP with async mount support and failover.
- **2012** – DMF 5.5 & 5.6 released. MAID support. Fast Mount Cache introduced.
- **2013** – DMF 6.0 released. RedHat server support, Vaulting with pDMF, Recall Priority., Library Auto Config, Integrated Backup,
 - LTFS Import, Logical Block protection,
 - SGI acquires FileTek, to bring Trusted Edge and StorHouse functionality into the DMF solution.

Recent Updates

- ISSP 2.6 Dec 2012
- ISSP 3.0 May 2013
- ISSP 3.1 Feb 2014
- ISSP 3.2 July 2014
- ISSP 3.3 Dec 2014

ISSP 2.6

- Fast-mount cache
- IBM TS1140 tape drive support
- dmcapacity and dmoverview commands
- Ability to specify:
 - How long to wait for a volume to be mounted
 - Size of volumes
 - Size of the alerts and performance databases
 - The number of tape drives in a drive group used for migrations
- Server support for SLES 11 SP2
- DMF recall optimization
 - limit recall of files split across multiple tapes to a designated mover
- Merge protocol to improve performance
- DMF client no longer hang if network goes down
 - commands make use of the CXFS failover network to complete
- Option for DMF client commands to connect outside of a secure port range

ISSP 3.0

- RedHat server support
 - DMF/CXFS/XVM/eSAMBA/eXFS/eNFS
- DMF out of library tape cartridge support
- DMF Priority recall
- specify volume group (VG) from which to recall data
- recall from a second copy if 1st copy in use
- LTO-6

ISSP 3.1

- Oracle T10000 D tape drive support
- Logical Block Protection
- Sles11sp3 server support
- RedHat 6.5 server support

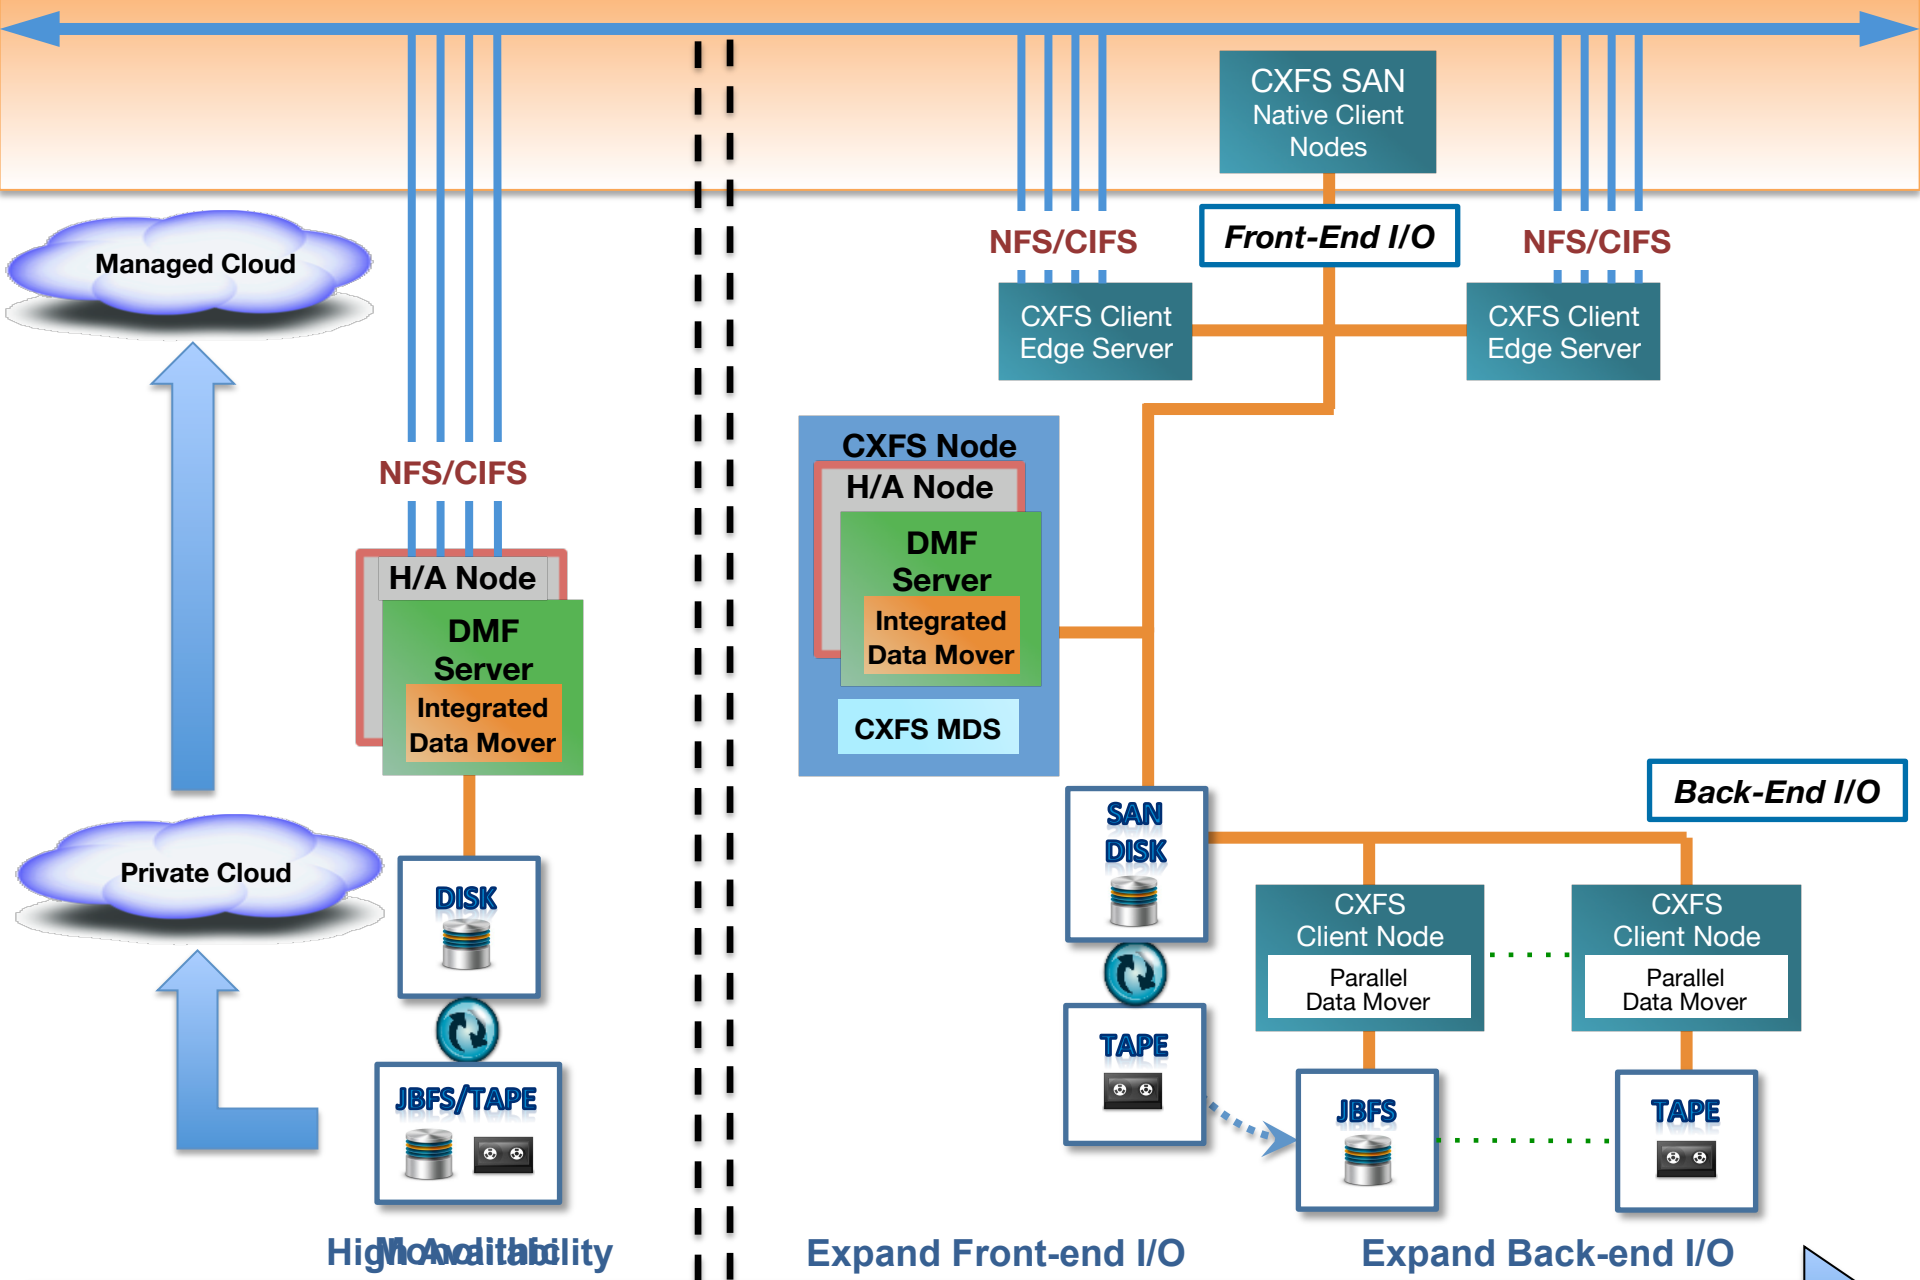
ISSP 3.2

- Cloud support added - S3 and Scalify MSP's
- LTFS capability added
- Project ID support
- libdmfusr.so support to cancel or reprioritize some requests
- persistent drive reservation
- support for 1024 Volume Groups or Migrate Groups

DMF 3.3

- DMF now supports use of the immutable bit to prevent a file from being changed or deleted, even by root.
- New commands:
 - * `dmcancel(8)` cancels some types of DMF requests by request ID
 - * `dmrepri(8)` reprioritizes recall and `dmcopy` request by daemon request ID
- Support for the full line of Overland libraries
- New technology previews: JBFS and DMF queue viewing
- DMF infrastructure includes limited use of the PostgreSQL database for queue viewing capabilities.
- Support for the full line of Overland libraries has been added
- DMF daemon no longer has a `dmdaux` process
- New technology previews: JBFS and DMF queue viewing

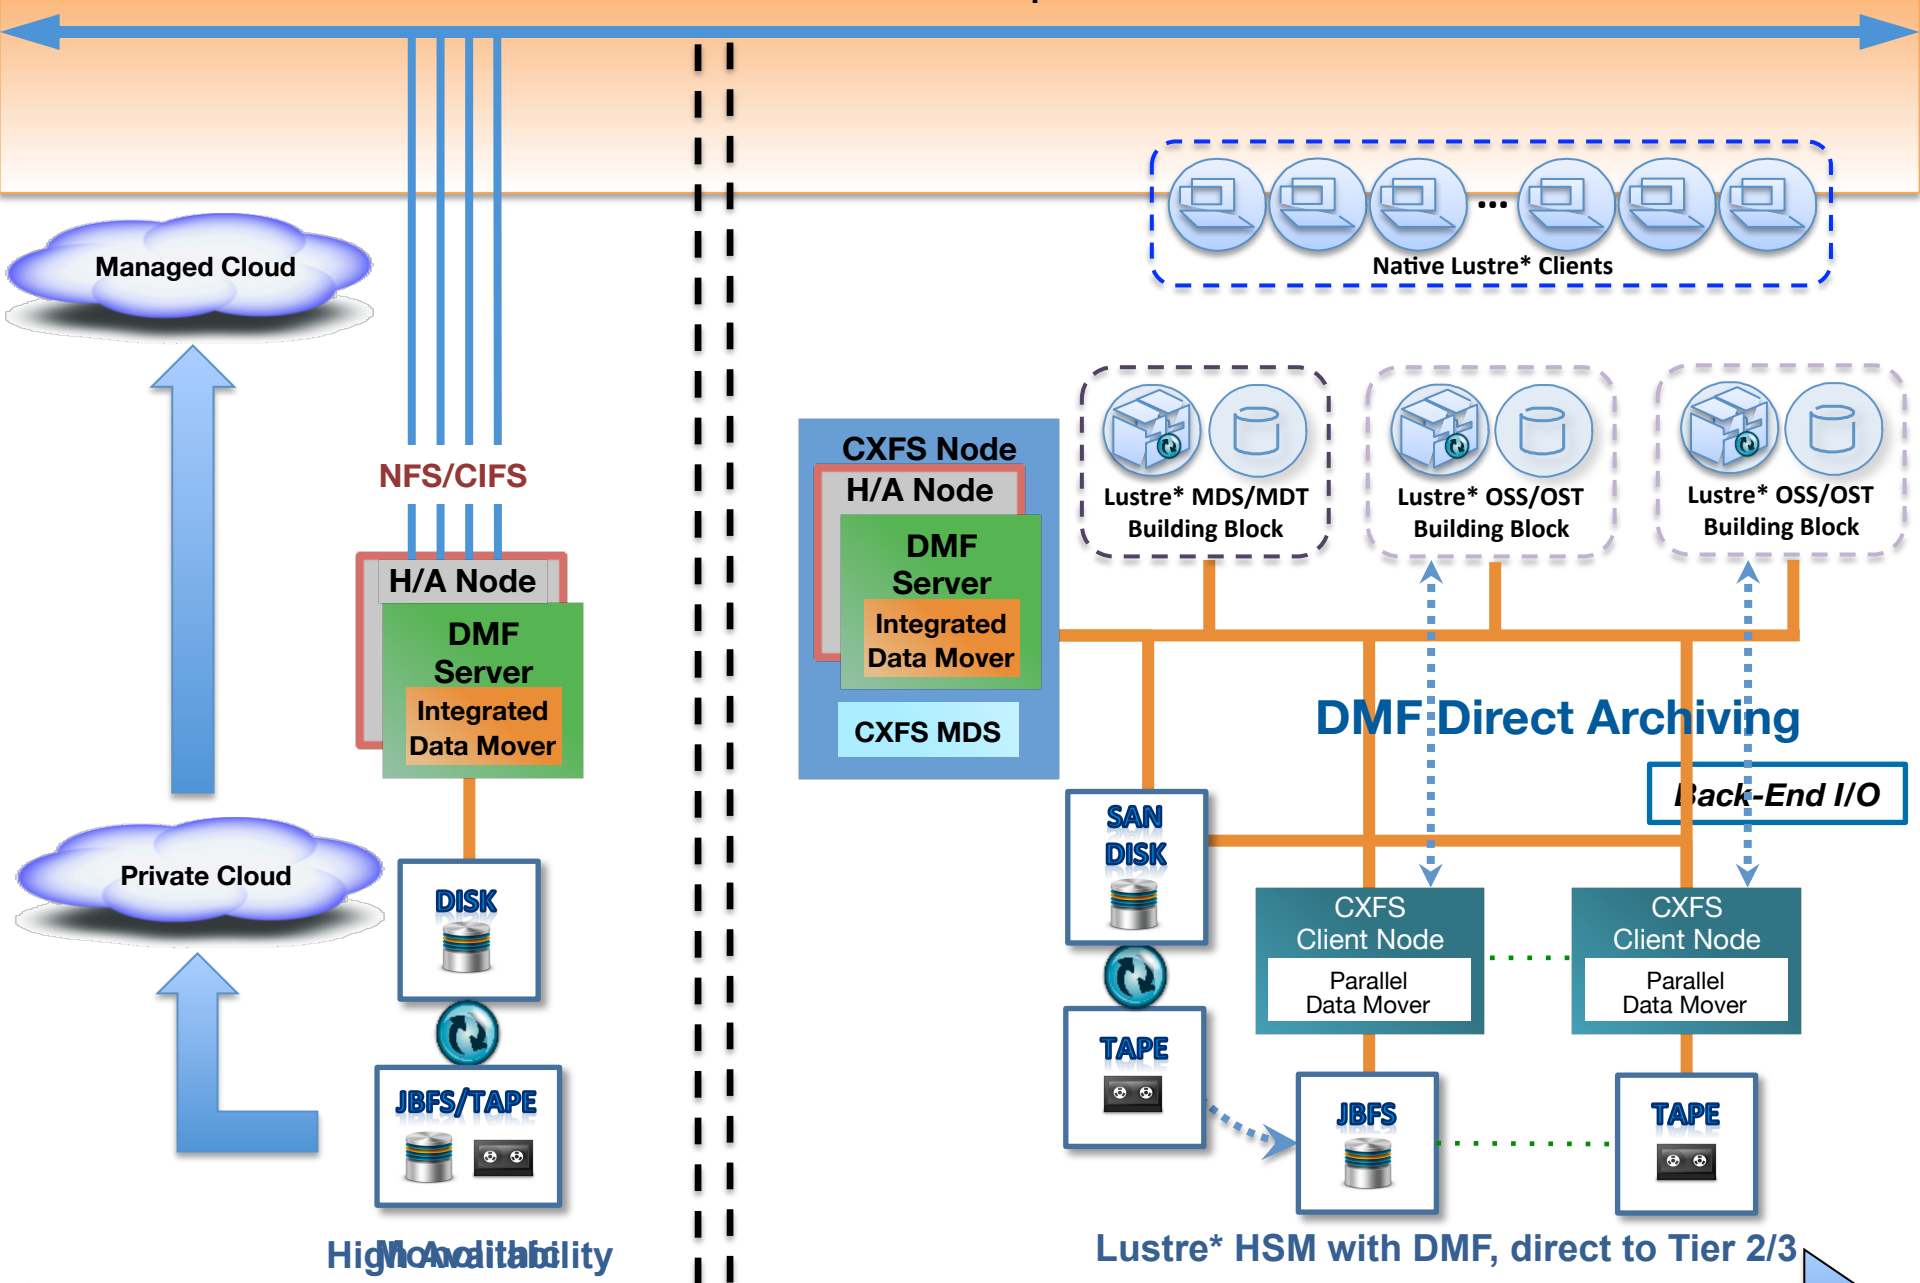
Client Access Space



DMF starts small and grows with you..

* = Some names and brands may be claimed as the property of others

Client Access Space



High Availability

Lustre* HSM with DMF, direct to Tier 2/3

DMF starts small and grows with you..

* = Some names and brands may be claimed as the property of others

DMF Update | Cloud MSP

- Scality RING private cloud
- Amazon Simple Storage Service (S3)
public cloud

Note: Amazon Glacier is not supported.


- Other products that present a service interface that is compatible with S3

DMF Update | LTFS

- Show information about LTFS cartridges/drives and configuration.
- Import new cartridges into the LTFS group.
- Format, mount, or unmount the LTFS on the imported cartridges.
- Eject a cartridge from a library managed by OpenVault.
- Reserve a drive exclusively for LTFS use.
- Borrow the drive from a drive group that belongs to another application, and then return the drive group after the operation completes.
- Purge a cartridge record in the OpenVault catalog.
- Check an LTFS cartridge consistency.

JBFS | The OpenVault VTL for DMF

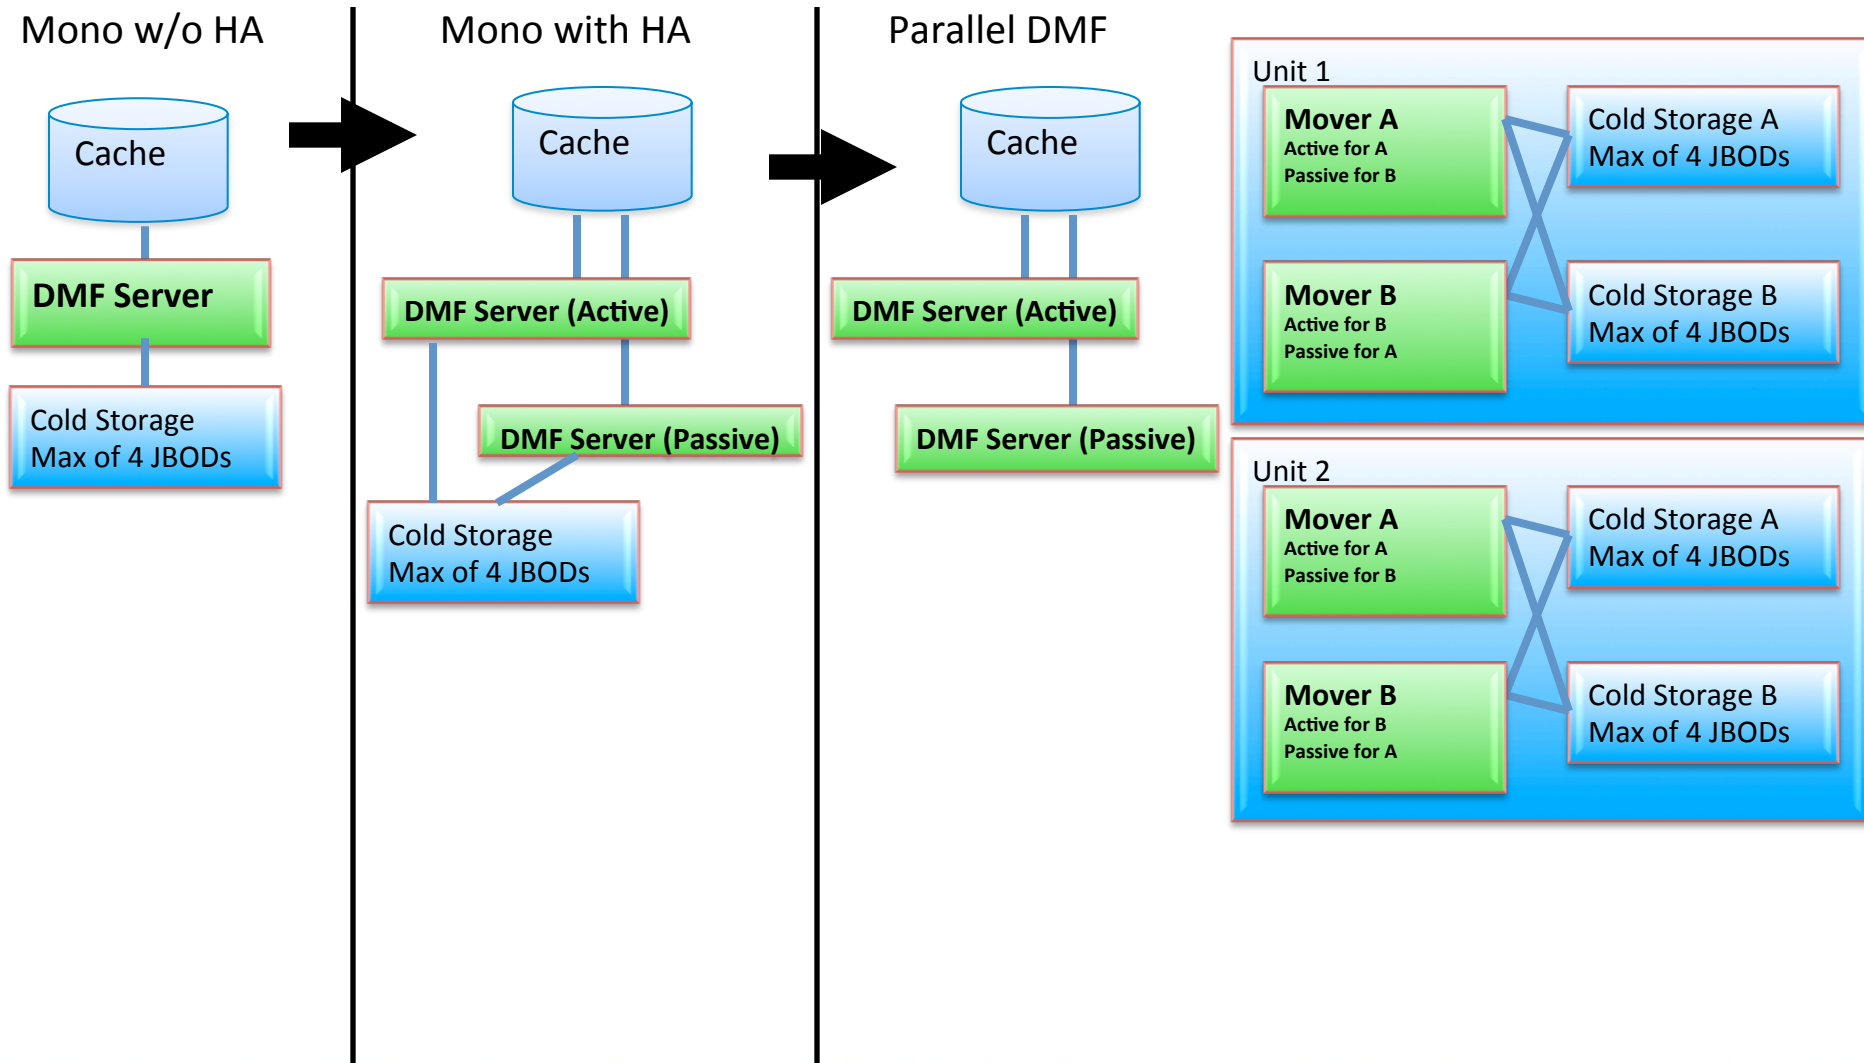
- JBFS is an acronym for JBOD File System
- JBFS provides mounting services
 - Serialised access to disk media
 - Independent from Linux disk mounts and file systems
- Disks treated like tapes mounted in tape drives
- The primary advantages
 - Mount performance
 - Low-cost scalable data throughput performance
 - Power management via ZeroWatt™



DMF Update | JBFS

- Software support for a Virtual Tape Library
- Software that allows DMF to treat disks exactly as it would treat tape cartridges in tape drives.
- The primary advantages
 - mount performance
 - low-cost scalable data throughput performance

Cold Storage Progression



DMF Update | Other Features

- Request reprioritisation and cancellation
- DMF Q-View
- DMF Licensing
- Persistent drive reservation
- Integrated backup for Basic DMF
- Policy by Project ID
- Lustre HSM copytool (pDMF aware and direct to tape capability)
- Integrated-Backup Method for Basic DMF

Cold Storage Hardware


Q2 2015

JBFS JBOD Enclosure

96-bay | Features and Functionality

- Very High Density Cold Storage – 8TB HDD
- 96 hot-swappable 3.5” SATA drives
- 6Gb/s SATA Support
- Individual drive power control
- T-10 SES 3.0 compliant
- No active components in the data paths
- Redundant Hot-swappable Power and Cooling





I/O Module

96 SATA Hard Drives



84-bay Dual host connectivity

- Disk Drives – 6TB HDD
- 84 independent point-to-point connections to each SAS or SATA disk drive with dual-port access and failover by each I/O controller to each drive
- - SATA drives require optional 2:1 Active MUX



84-bay Dual host connectivity



