

LUSTRE ON ZFS: BUILDING BLOCK ARCHITECTURE

New in 2019!

Highly-Available NVMe Flash Array. Accelerate metadata intensive workloads or unleash all-flash Lustre with NVMe. Up to 20 GB/s and 140TB in 2U.



Highlights

- Scalable storage capacity and I/O
- Free and open source software
- Low total cost of ownership
- Maximum storage density
- High availability
- Data integrity
- Standard enterprise administrative tools & security
- 24/7 enterprise support

Unleashing Performance & Scalability

Today, HPC is a key technology for enterprises of all sizes. Getting maximum performance from HPC and data-intensive applications is a necessity that requires storage infrastructure that scales endlessly and delivers unmatched I/O levels while pushing down total cost of ownership. In leveraging free open source software with Lustre 2.10 and ZFS on Linux 0.7, RAID Inc unleashes the performance and scalability of the Lustre parallel file system for HPC workloads with higher density and lower TCO. Enterprise level support available from RAID Inc. and backed by core Lustre developers at Whamcloud®.

Building Block Architecture

Lustre building blocks are built from RAID Inc's Ability® line of enterprise class JBODs (EBODs) and commodity server hardware. All building blocks are configured for high-availability with no single point of failure. Building block architecture allows seamless scale-out of capacity and performance for new and existing systems.

Metadata

MDS building blocks leverage RAID Inc's 24-Bay Ability® EBOD. SFF 2.5" drive bays allow for SSDs or high-RPM spinning disks. Up to 100 TB of usable metadata capacity with enclosure protection in a single building block.

Object

OSS building blocks leveraging RAID Inc's Ability® Series of EBODs. LFF 3.5" drive bays allow high capacity drives for up to 2.8 PB in a single building block. With up to 10 GB/s throughput per building block, scale-out to meet your performance and capacity needs is easier than ever.

Support

Through our experience, we have identified and developed best practices for delivering levels of technical support that fit our customers' ever-changing needs for enterprises of all sizes. RAID Inc. delivers high-quality, customized technical support for all stages of the customer lifecycle.

Benefits

- Support for ZFS on Linux 0.7. Large dnodes for greatly increased metadata performance. Integrated and simplified JBOD management. Multiple import protection for HA clusters. Improved disk I/O monitoring.
- Snapshots. Fast, mountable, read-only snapshot of a Lustre filesystem in a separate namespace. Only available if all storage targets are using ZFS.
- Multi-rail LNET. Support for load balancing network traffic across multiple physical interfaces. Decreased configuration complexity for servers and clients that have large bandwidth requirements.
- NEW FEATURE: Lustre 2.11 support for Data on MDT. Greatly increase small file performance in Lustre by allocating them directly on SSD-backed MDTs.

ScaleAbility LUSTRE ON ZFS Solution


New in 2019!

Highly-Available NVMe Flash Array. Accelerate metadata intensive workloads or unleash all-flash Lustre with NVMe. Up to 20 GB/s and 140TB in 2U.



Extreme-Scale Storage for High Performance

Industry first, highly-available NVMe flash array with dual ported drives and integrated Skylake server modules. Accelerate filesystem workloads with the highest performance disks on the market while retaining five nines uptime. Deploy all flash OSTs for performance up to 20 GB/s in 2U. For capacity needs, daisy chain dense HDD enclosures for spinning disk-based OSTs.



All Flash Solution

140 GB/s
1 PB
16U



Capacity Solution

10 GB/s
3 PB
20U

Highlights

- Highly-Available clusters with no single point of failure
- Scale for capacity and performance of both data and metadata
- 24 bay enclosure SSD or HDD for metadata
- 90-bay enclosure for capacity object storage
- Infiniband, Omni-Path and Ethernet Lustre networking

Benefits

- Support for ZFS on Linux 0.7. Large dnodes for greatly increased metadata performance. Integrated and simplified JBOD management. Multiple import protection for HA clusters. Improved disk I/O monitoring.
- Snapshots. Fast, mountable, read-only snapshot of a Lustre filesystem in a separate namespace. Only available if all storage targets are using ZFS.
- Multi-rail LNET. Support for load balancing network traffic across multiple physical interfaces. Decreased configuration complexity for servers and clients that have large bandwidth requirements.
- NEW FEATURE: Lustre 2.11 support for Data on MDT. Greatly increase small file performance in Lustre by allocating them directly on SSD-backed MDTs.