

# Intel® Enterprise Edition for Lustre Storage Optimized Solutions



## Unleashing the Performance and Scalability of the Lustre Parallel File System for HPC Workloads and Big Data Applications



### The Challenge

Getting maximum performance from HPC applications that requires storage infrastructure that scales endlessly and delivers unmatched I/O levels.

### Highlights

- Enormous storage capacity and I/O
- REST API
- High availability
- Hitless upgrade
- Intuitive graphical user interface and CLI
- Cloud-based services
- Standard enterprise administrative tools & security
- Throughput in excess of 1 terabyte per second
- 24/7 enterprise support

### Use Cases

- HPC Workloads
- Data-intensive Applications
- Technical 'Big Data' Applications

### Overview

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.

In partnering with Intel® Enterprise Edition for Lustre software (Intel® EE for Lustre software), RAID Inc. unleashes the performance and scalability of the Lustre parallel file system for HPC workloads, including technical 'big data' applications.



Figure 1. Intel® Enterprise Edition for Lustre software

### Benefits

**Performance.** Intel EE for Lustre software has been designed to enable fully parallel I/O throughput across thousands of clients, servers, and storage devices. Metadata and data are stored on separate servers to allow optimization of each system for the different workloads they present. Improved metadata scalability using Distributed Namespace (DNE) feature is now integrated in Intel Manager for Lustre. Intel EE for Lustre can also scale down efficiently to provide fast parallel storage for smaller organizations.

**Capacity.** The object-based storage architecture of Intel EE for Lustre software can scale to tens of thousands of clients and petabytes of data.

**Affordability.** Intel EE for Lustre software is based on the community release of Lustre software, and is hardware, server, and network fabric neutral. Enterprises can scale their storage deployments horizontally, yet continue to have simple-to-manage storage.

**Maturity.** Lustre has been in use in the world's largest datacenters for over a decade and hardened in the harshest big data environments. Intel EE for Lustre software is rigorously tested, reliable, and backed by Intel, the leading provider of technical support for Lustre software. Intel EE for Lustre software delivers commercial-ready Lustre in a package that can scale efficiently both up and down to suit your business workloads, with built-in manageability.

## Product Components

**Intel® Manager for Lustre.** Intel Manager for Lustre software includes simple, but powerful, management tools that provide a unified, consistent view of Lustre storage systems and simplify the installation, configuration, monitoring, and overall management of Lustre. The manager consolidates all Lustre information in a central, browser-accessible location for ease of use.

**Integrated Apache Hadoop Adapter.** When organizations operate both Lustre and Apache Hadoop within a shared HPC infrastructure, there is a compelling use case for using Lustre as the file system for Hadoop analytics, as well as HPC storage.

Intel Enterprise Edition for Lustre includes an Intel-developed adapter which allows users to run MapReduce applications directly on Lustre. This optimizes the performance of MapReduce operations while delivering faster, more scalable, and easier to manage storage.

**Table 1.** What's New in EE v3.0

Feature	Benefit
Open ZFS Updates	Several performance and stability improvements after updating to the latest stable version (0.6.5). Improved I/O throughput for large files now that the dataset record size can be set up to 1MB.
Lustre 2.7	The latest stability fixes and performance enhancement from extensive production deployments and comprehensive test frameworks. Provides customers with the latest Intel® EE for Lustre software core.
Online File Consistency (LFSCK)	Online Lustre File System Consistency Checks (LFSCK) allow checks and repairs to be performed without downtime. In this release, LFSCK will verify and repair the Object Index (OI) table, directory FID-in-Dirent and LinkEA, MDT-OST file layout inconsistencies, and inconsistencies between multiple MDTs.
Single Client Metadata Concurrency	An update to metadata communications interface between client and server, which allows multiple metadata RPCs to be in flight in parallel, per-client for both read and write transactions.
Intel® Manager for Lustre	Updates include support for latest RHEL releases, UI refresh with new workflows, and improvements to responsiveness, performance, layout, and navigation.
Enterprise Edition Server and Client Support for RHEL 7.2, 6.7 and SLES 11 SP4 and Client-Only Support for SLES 12 and Ubuntu 14.04	Provides the latest updates and security patches to customers.
Intel® Omni-Path Architecture Support	Intel® OPA is supported by Intel® EE for Lustre software, as are all major HPC fabrics, including InfiniBand. Intel® OPA fabric support is available for systems running RHEL 7.2.
Secure Clients	SELinux support to enforce access control policies, including Multi-Level Security (MLS) on Lustre clients.
Secure Authentication & Encryption	Kerberos functionality can be applied to Intel® EE for Lustre software to establish trust between Lustre servers and clients and to support encrypted network communications.
OpenZFS Updates	Several performance and stability improvements after updating to the latest stable version (0.6.5). Improved I/O throughput for large files now that the dataset record size can be set up to 1MB.
Snapshots	A mechanism that is now capable of leveraging OpenZFS to take a coordinated snapshot of an entire file system (provided that all of the metadata and storage targets are formatted using OpenZFS). With this version, the snapshot is taken across the whole file system, which can be mounted on a Lustre client as a separate name space.
Dynamic Network Configuration	Powerful extension to the Lustre network (LNet) to simplify system administration tasks. Provides tuning and optimization of LNet while Lustre is still running on the target node. Add/remove and update LNet routers without effecting Lustre network traffic.

Intel harnesses the capability of the cloud, the ubiquity of the Internet of Things, the latest advances in memory and programmable solutions, and the promise of always-on 5G connectivity, Intel is disrupting industries and solving global challenges.



RAID Inc. was founded in 1994 to deliver end-to-end performance-driven technical computing and storage solutions. The company has earned industry praise for providing platform agnostic technical guidance in high performance computing (HPC), big data, cloud and software-defined data centers—in the most efficient, reliable and cost effective manner. The world's leading research facilities, government, life science, financial, healthcare, energy, and cloud service providers can leverage the our team of engineers' extensive academic, research lab and commercial expertise that makes RAID Inc. a trusted industry leader.