



# Intel® Enterprise Edition for Lustre\* Software

A New Generation of Lustre Software Expands High Performance Computing into the Commercial Enterprise

## KEY FEATURES


- Lustre\* version 2.5 foundation
- POSIX compliant, parallel, distributed scale-out storage
- Infrastructure for policy-driven tiered storage solutions
- Optimized storage for Hadoop\* applications
- Simple but powerful UI lowers management complexity
- Supports larger, more complex configurations
- Native Lustre client for Intel® Xeon Phi™ coprocessor
- Server and client support for all major Linux\* distributions, including SUSE Linux Enterprise\*
- Extensible architecture with REST API and storage plug-in interface
- Global, around the clock technical support

Around the world and across all industries, high-performance computing is being used to solve today’s most important and demanding problems. More than ever, storage solutions that deliver high sustained throughput are vital for powering HPC and Big Data workloads.

**Intel® Enterprise Edition for Lustre\* software** unleashes the performance and scalability of the Lustre parallel file system for enterprises and organizations, both large and small. It allows users and workloads that need large scale, high-bandwidth storage to tap into the power and scalability of Lustre, but with the simplified installation, configuration, and monitoring features of Intel® Manager for Lustre\* software, a management solution purpose-built for the Lustre file system. Intel EE for Lustre software includes proven support from the Lustre experts at Intel, including worldwide 24x7 technical support.

## High Performance Parallel Storage for the Enterprise

Intel EE for Lustre software brings the benefits of Lustre and high performance storage to data-intensive applications. For decision makers seeking to achieve their business goals more quickly and cost effectively, Intel EE for Lustre software offers a parallel file system that enables more powerful computing for improved results and data-driven business and IT intelligence. Intel EE for Lustre software provides open interfaces that promote easy integration and high levels of interoperability with existing infrastructures. Because Lustre is open source software, IT can grow data center storage systems over time and on budget, employing a variety of networking and hardware options.






Figure 1. Typical Enterprise Edition for Lustre\* configuration

## INTEL AND OPEN SOURCE LUSTRE\*

Lustre has been developed in an open, collaborative community and is freely available in open-source software under the GPLv2 license. Intel is committed to maintaining Lustre as an open system while providing expert support and tools to enhance management. Intel takes pride in being a long-standing member of the open-source community, and believes in open source development as a means to create rich business opportunities, advance promising technologies, and bring together top talent from diverse fields to solve computing challenges. Intel's contributions to the open source community include reliable hardware architectures, professional development tools, work on essential open-source components, collaboration and co-engineering with leading companies, investment in academic research and commercial businesses, and helping to build a thriving ecosystem around open source.

## What is Lustre storage software?

The Lustre file system was purpose-built to provide sustained storage performance and stability at large scale for high-performance computing (HPC) systems. Lustre is POSIX-compliant—vital for applications—and capable of handling extremely large amounts of data and huge numbers of files shared concurrently across clustered servers. It offers wide scalability, both in performance and storage capacity. According to IDC, Lustre is the most widely used file system for HPC worldwide, powering 60 percent of the fastest 50 supercomputing sites worldwide.<sup>1</sup>

Traditionally employed in HPC environments that generate and process massive amounts of data for research, governmental, and industrial purposes, storage powered by Lustre software is a breakthrough technology for addressing exascale and emerging “big data” challenges.

HPC is no longer the exclusive province of such data-rich projects as genomics, climate modeling, and seismic analysis. Increasingly, organizations and businesses of all kinds generate high-volume data flows that require scalable, simple to manage storage to deliver very high I/O rates and massive capacity to super-scale applications.

With the release of Intel Enterprise Edition for Lustre software, Intel provides customers with a commercial-grade version of Lustre optimized to address the storage and data throughput challenges impacting HPC-class computing.

## Performance

Intel EE for Lustre software is designed to enable fully parallel I/O throughput across 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. Intel EE for Lustre scales down efficiently to provide fast parallel storage for smaller organizations.

- Many Lustre configurations are running in production at 500 to 750 gigabytes per second, with leading edge installations achieving throughput in excess of 2 terabytes per second
- Massive data flows can utilize a high percentage of underlying storage and network bandwidth, for low latency, high throughput storage performance
- Improves innovation and precision by allowing enterprises to run larger and more complex simulations faster and easier
- Native Lustre client optimized for Intel® Xeon Phi™ coprocessors allows applications to directly connect to Lustre storage

## What's New

- Hierarchical storage management infrastructure to support policy-driven, cost effective, tiered storage
- Streamlined configuration workflow for Intel Manager for Lustre
- Native Lustre client for Intel Xeon Phi coprocessor
- Support for larger configurations
- Intuitive “heat map” chart options integrated with Lustre “jobstats” to easily identify performance constraints
- Improved single client and single stream performance
- Storage server support for SUSE Linux<sup>2</sup> Enterprise

## Storage Infrastructure for Big Data Workflows

### Tackling the Big Data Challenge

From Wall Street to the Great Wall, enterprises and institutions of all sizes are faced with the benefits and challenges promised by Big Data. But before users can take advantage of the nearly limitless potential locked within their data, they must have affordable, scalable, and powerful software tools to analyze and manage their data. Intel EE for Lustre provides the performance at scale to fuel data intensive applications.

### Industry-Leading File Storage

High performance workloads are larger, more complex, and more important than ever before.

Today's HPC users are demanding application frameworks to analyze vast amounts of data created by complex simulations. As the most widely deployed file system for HPC, EE for Lustre software plays a critical role for these data-intensive applications, transforming enormous amounts of data into data-driven decisions.

### Performance Meets Scalability for Big Data Storage

Intel Enterprise Edition for Lustre with the productivity and security of Lustre storage with software “connectors” designed for integration with leading Hadoop\* distributions. This powerful




Figure 2. Intel® Manager for Lustre® Dashboard

combination allows users to run their MapReduce\* applications—without changes—directly on shared, fast, Lustre-powered storage. This optimizes the performance of Hadoop applications, while delivering faster, more scalable, and productive storage.

### Accessing Big Data Insight

When organizations operate both Lustre and Hadoop within a shared HPC infrastructure, they create a powerful foundation for analyzing vast amounts of unstructured data. Hadoop MapReduce applications can access all of their data directly, avoiding the need to replicate data across multiple compute nodes.

Using Intel® EE for Lustre in combination with Hadoop makes storage management simpler—providing a single Lustre file system instance rather than partitioned, hard-to-manage storage.

### Capacity

The object-based storage architecture of Intel EE for Lustre software scales upward to tens of thousands of clients and petabytes of data.

### Affordability

Intel EE for Lustre software is based on open source Lustre software, and is hardware, server, and network fabric neutral, giving enterprises the flexibility to easily scale their storage solutions, yet continue to have simple-to-manage storage.

### Maturity

The performance and scalability of Lustre-powered storage has led it to be the most widely used parallel file system for HPC.<sup>3</sup> Proven in the most demanding data environments, its stable, reliable and backed by Intel, the leading global technical support provider for Lustre solutions of all sizes.

## Intel Manager for Lustre

Intel Manager for Lustre (IML) is a simple but powerful management solution that was purpose-built for Lustre. IML brings together information about your Lustre configuration to present a unified, consistent view about what's going on inside your storage system—while dramatically simplifying the installation, configuration, monitoring and overall management of Lustre.

As storage challenges grow rapidly—and unlocking the value within storage becomes even more important—today's high performance storage solutions have become too large and complex to be managed using software tools that were not designed for today's demanding needs.

Designed to make Lustre-based storage solutions easier to deploy and manage, Intel Manager for Lustre software maximizes return on storage investments. Purpose-built for unleashing the powerful benefits of distributed, parallel storage—from scalability to absolute performance—installing, configuring, and managing Lustre has never been simpler.

## Key Features of Intel Manager for Lustre

### Intuitive, browser-based administration

- Simple but powerful graphical and scriptable command line interfaces
- Point-and-click simplicity for cluster configuration and management
- Centralized definition and management of common administrative tasks

### Real-time system monitoring

- Near real-time monitoring storage health and key performance indicators in real time
- View high level system performance or individual components
- Generate historical and real-time charts and reports

Intel Manager for Lustre software consolidates Lustre information in a central, browser-accessible location for ease of management.

### Advanced troubleshooting tools

- Consolidated view of cluster-wide storage log files
- Intelligent log-scanning for efficient problem isolation and analysis
- Configurable event notifications

Open, documented APIs REST API for integration with management frameworks, and a storage plug-in architecture that enables deep storage hardware reporting.

To learn more about Intel® solutions for Lustre\* software, visit [intel.com/lustre](http://intel.com/lustre)

<sup>1</sup> Intel analysis of Top500 statistics from November 2013 report, see [www.top500.org](http://www.top500.org)

<sup>2</sup> Storage server support only with this release, no support for Intel Manager for Lustre with SUSE

<sup>3</sup> IDC survey research, May 2014

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

A "Mission Critical Application" is any application in which failure of the Intel Product could result, directly or indirectly, in personal injury or death. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL PRODUCT OR ANY OF ITS PARTS.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined". Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or go to: <http://www.intel.com/design/literature.htm>

FTC Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel.

Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

General Performance Disclaimer: For more complete information about performance and benchmark results, visit Performance Test Disclosure <http://www.intel.com/benchmarks>

Copyright © 2014 Intel Corporation. All rights reserved. Intel, the Intel logo, and Xeon Phi are trademarks of Intel Corporation in the U.S. and other countries.

