

Lustre* Releases Update

Peter Jones, HPDD

October 2016

Copyright © Intel Corporation 2016. All rights reserved. Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

* Other names and brands may be claimed as the property of others.

Legal Information

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit <http://www.intel.com/performance>.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at <http://www.intel.com/content/www/us/en/software/intel-solutions-for-lustre-software.html>.

Intel technologies may require enabled hardware, specific software, or services activation. Check with your system manufacturer or retailer.

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

The products described 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.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

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 EXPENSES 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.

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

* Other names and brands may be claimed as the property of others.

© 2016 Intel Corporation

Community Survey

Has been running for 5 years


OpenSFS Lustre Working Group
devises questions

Useful for tracking trends in Lustre
usage


Highest number of respondents yet for
2016 survey (97)

Full details available at
http://wiki.opensfs.org/Lustre_Community_Survey

Respondents



Community Survey - Versions



Lustre 2.5.x remains the dominant version in production usage

Strong adoption of more current releases (even 2.8 before it was `GA)

Lustre 1.8.x usage finally dwindling

Lustre Quality

Well-established release validation practices

- Automated functional regression tests across test matrix
- Stress testing runs on Hyperion
- Execution of feature test plans

Continuing to evolve testing practices

- Soak testing
- Static code analysis tools

Drawing on expertise from outside Lustre community

- Influence from upstream Linux community
- Input from Intel software experts

Lustre 2.8

GA March 2016

•Strong feature set

- RHEL* 7.2 server support (LU-5022)
- DNE Phase 2 (LU-3534)
- LFSCK Phase 4 (LU-6361)
- Client Metadata RPC Scaling (LU-5319)

Number of production sites on this release

- Lustre on ZFS* perf and newer kernel support drivers

http://wiki.lustre.org/Release_2.8.0

Lustre 2.9

Converging towards release; hopefully out soon

- http://wiki.lustre.org/Release_2.9.0

Will support RHEL 7.2 servers/clients and SLES*12 SP1 clients

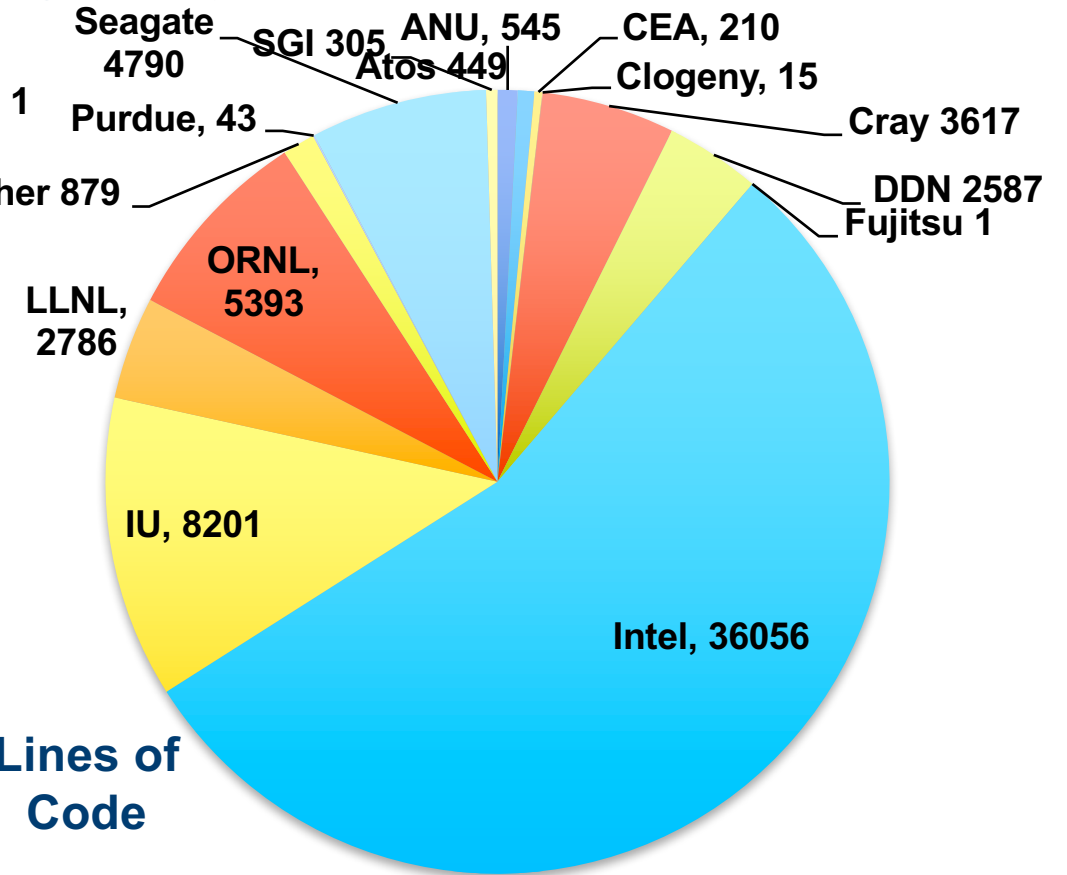
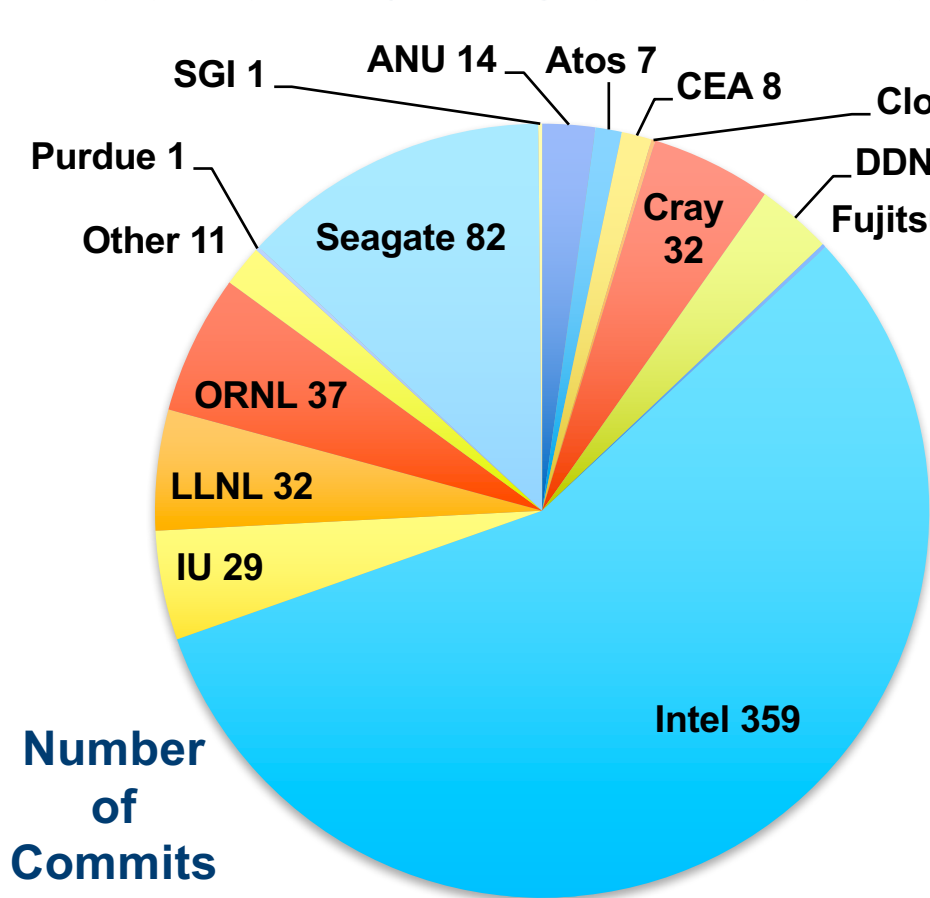
- Dropping support for RHEL 6 and SLES11

Interop/upgrades from Lustre 2.8 servers/clients

A number of new features targeted for this release:

- UID/GID Mapping (LU-3291)
- Subdirectory Mounts (LU-28)
- Large Block IO (16MB IO for ldiskfs) (LU-7990)
- ZFS Metadata Performance Improvements (LU-7895)
- Shared Key Crypto (LU-3289)
- Weak updates (LU-5614)
- Server side advise and hinting (LU-4931)

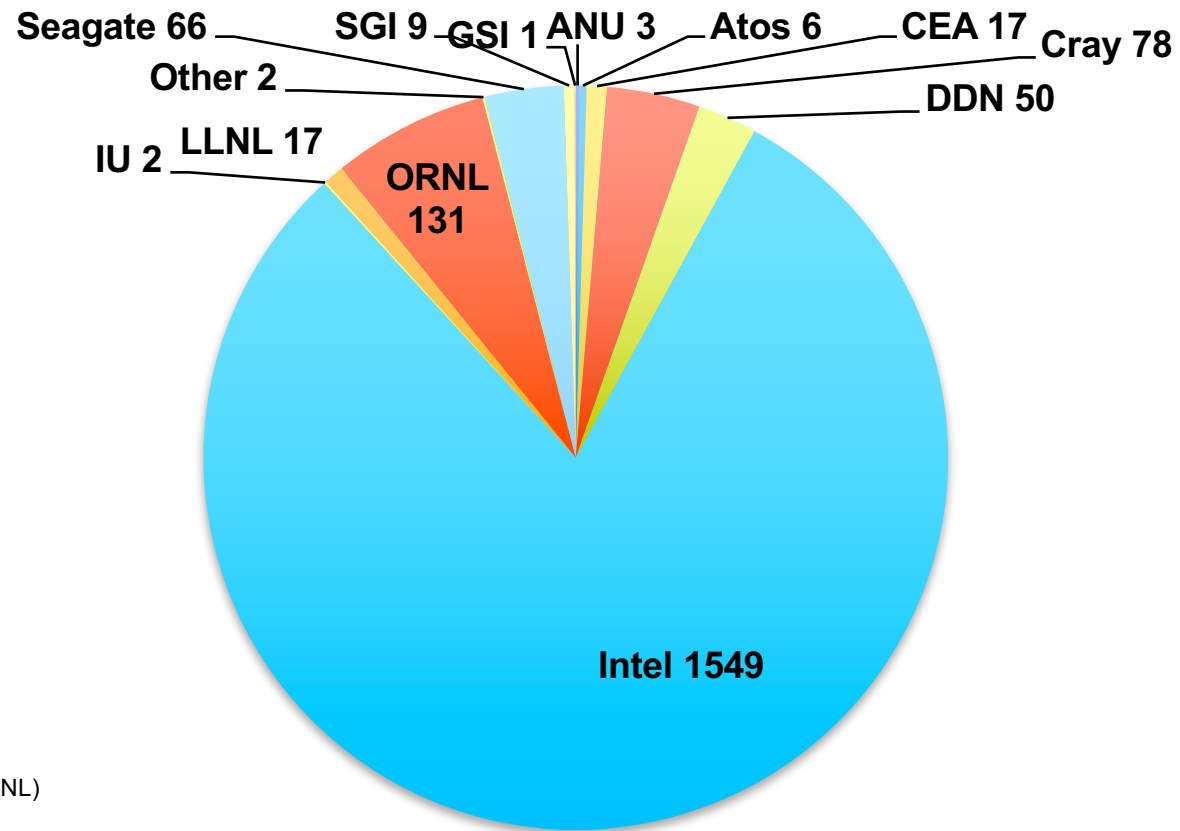
Lustre 2.9 – Contributions So Far



Statistics courtesy of Dustin Leverman(ORNL)
 Source: <http://git.whamcloud.com/fs/lustre-release.git>

Aggregated data by organization between 2.8.50 and 2.8.59 tags

Lustre 2.9 – Reviews So Far



Statistics courtesy of Dustin Leverman(ORNL)

Source: <http://git.whamcloud.com/fs/lustre-release.git>

Aggregated data by organization between 2.8.50 and 2.8.59 tags

Lustre Version Statistics

Version	Commits	LOC	Developers	Organizations
1.8.0	997	291K	41	1
2.1.0	752	92K	55	7
2.2.0	329	58K	42	10
2.3.0	586	87K	52	13
2.4.0	1123	348K	69	19
2.5.0	471	102K	70	15
2.6.0	885	147K	76	14
2.7.0	741	193K	65	15
2.8.0	992	147K	92	17
2.9.0**	635	66K	95	16

Statistics courtesy of Dustin Leverman(ORNL)

Source: <http://git.whamcloud.com/fs/lustre-release.git>

** Data until 2.8.59

Lustre 2.10 and Beyond

Many interesting long term features in development

- Multi-rail LNET
- Progressive File Layouts
- Project quotas

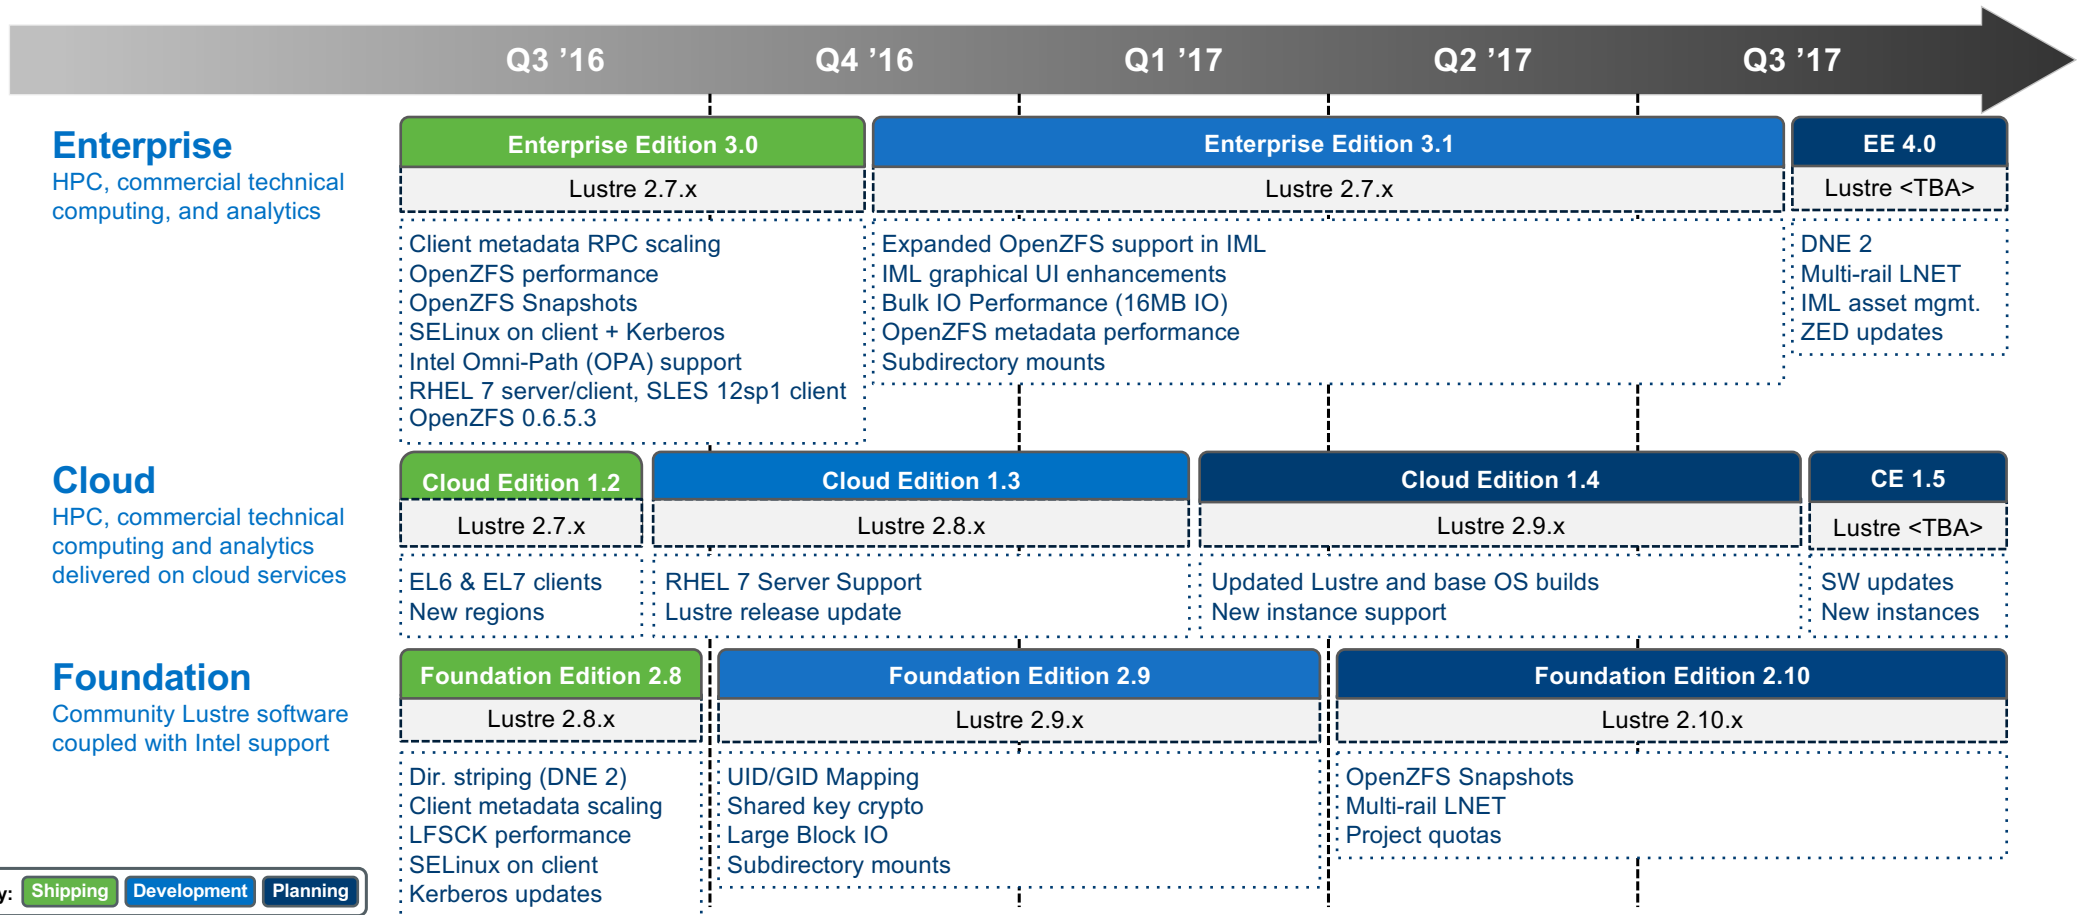
Full list maintained on lustre.org wiki

- <http://wiki.lustre.org/Projects>

Highlights and timelines shown on community roadmap

- <http://lustre.org/roadmap/>

Intel® Solutions for Lustre* Roadmap



Releases – Intel® Enterprise Edition for Lustre* Software

Intel® EE for Lustre* Software 2.4.x

- Based on community 2.5.x release
- Most widely deployed codeline today

Intel® EE for Lustre* Software 3.0.x

- Based on community 2.7 release
- RHEL 7.2 server support
- RHEL 7.2 / RHEL 6.8 / SLES12 SP1 / Ubuntu 14.04 client support
- Intel* Omni-Path Architecture (Intel® OPA) Gen 1 fabric support
- ZFS Snapshots
- ZFS Streaming IO performance improvements
- Single client metadata concurrency
- SE Linux for Lustre clients

Releases – Intel® EE for Lustre* Software (cont)

Intel® EE for Lustre* Software 3.1.x

- Based on community 2.7 release
- Targeted for Q4 2016
- Latest RHEL 7.x server/client support
- IML managed mode for OpenZFS
- Metadata performance improvements for Lustre on OpenZFS
- Bulk IO Performance for Idiskfs (16MB IO)
- Subdirectory mounts

Releases – Cloud Edition

Cloud Edition 1.3

- Maintenance update
- Lustre 2.8.x baseline
- RHEL 7.2 servers and RHEL 7.x/6.x clients

Releases – Foundation Edition

Maintenance releases based on community releases

- Available to HPDD customers

Lustre 2.8.1 FE

- Finalizing; will issue once ORNL/LLNL are confirmed in full production

Lustre 2.7.2 FE

- Used by multiple larger sites; 2.7.3 FE possible

Lustre 2.5.x FE

- No current plans for another release; sites still running are stable

Summary

EE 3.1 due out this quarter

Strong community adoption of recent releases

Community 2.9 release should be GA in the near future

Some interesting developments in progress for 2.10 and beyond

