



Community Release Update

LAD 2017

Peter Jones HPDD, Intel
OpenSFS Lustre Working Group

OpenSFS Lustre Working Group

Lead by Peter Jones (Intel) and Dustin Leverman (ORNL)


Single forum for all Lustre development matters

- Oversees entire Lustre development cycle
- Maintains the roadmap
- Plans major releases
- Collects requirements for future Lustre features
- Sets priorities for test matrix

For more information visit the wiki

http://wiki.opensfs.org/Lustre_Working_Group

Community Survey - Versions



- Survey conducted March 2017 (so before 2.10 was available)
- Lustre 2.5.x remains the most commonly-used version in production
- Strong adoption of more current releases
- Full results at http://wiki.opensfs.org/Lustre_Community_Survey

Lustre 2.10

- Went GA July 13th 2017
- Supports RHEL 7.3 servers/clients and SLES12 SP2 clients
- Interop/upgrades from Lustre 2.9 servers/clients
- Designated LTS Release - freely available maintenance releases
 - Lustre 2.10.1 targeted for Q3 release
 - Lustre 2.10.x will remain LTS branch for at least 18 months
- http://wiki.lustre.org/Release_2.10.0


Lustre 2.10 - Features

- A number of new features included in this release
 - Progressive File Layouts (LU-8998)
 - Multi-Rail LNET (LU-7734)
 - Project Quotas (LU-4017)
 - NRS Delay Policy (LU-6283)
 - ZFS Snapshots (LU-8900)
- Some useful enhancements
 - ZFS Metadata Improvements (LU-7895)
 - OPA Performance improvements (LU-8943)
 - Pacemaker scripts (LU-8457/8458)
 - Feature/bugfix parity with latest EE 3.x Lustre releases

Lustre 2.10 – Progressive File Layouts

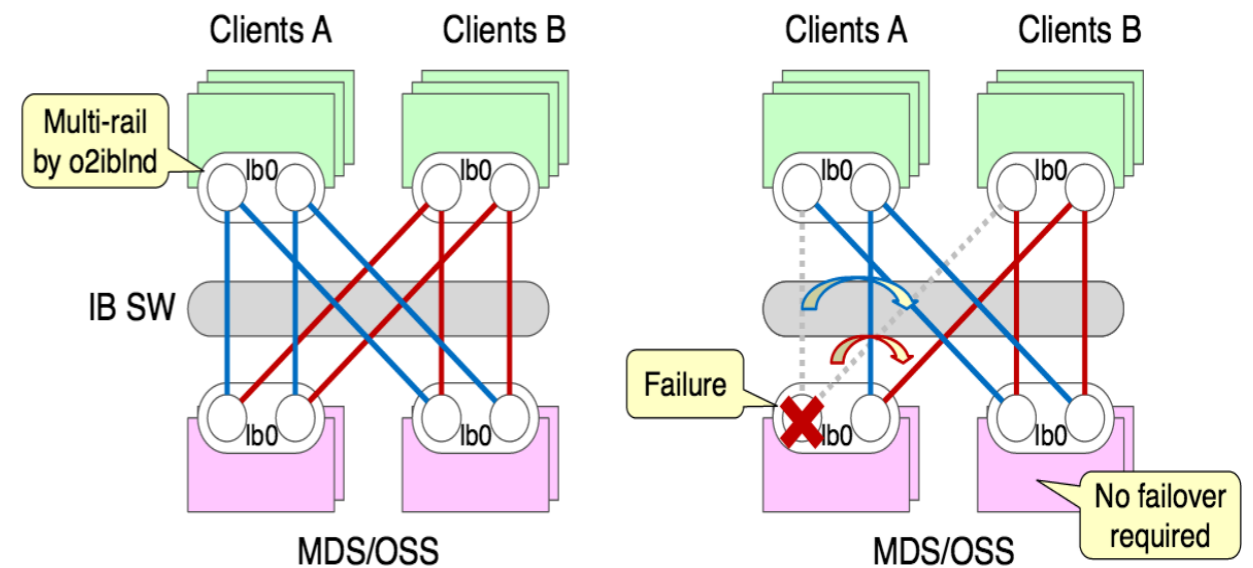
- Progressive File Layout (PFL) simplifies usage for users and admins
 - Optimize performance for diverse users/applications
 - One PFL layout could be used for all files
 - Low stat overhead for small files
 - High IO bandwidth for large files
- Collaboration between Intel and ORNL

Example progressive file layout with 3 components




Lustre 2.10 – Multi-Rail LNet

- Allow LNet across multiple network interfaces
 - Supports all LNet networks – LNet layer instead of LND layer
 - Allows concurrent use of different LNDs (e.g. both TCP & IB at once)
- Scales performance significantly
- Improves reliability
 - Active-active network links between peers
- Collaboration between Intel and HPE/SGI




Lustre 2.10 - Contributions

Number of Commits



Lines of Code Changed

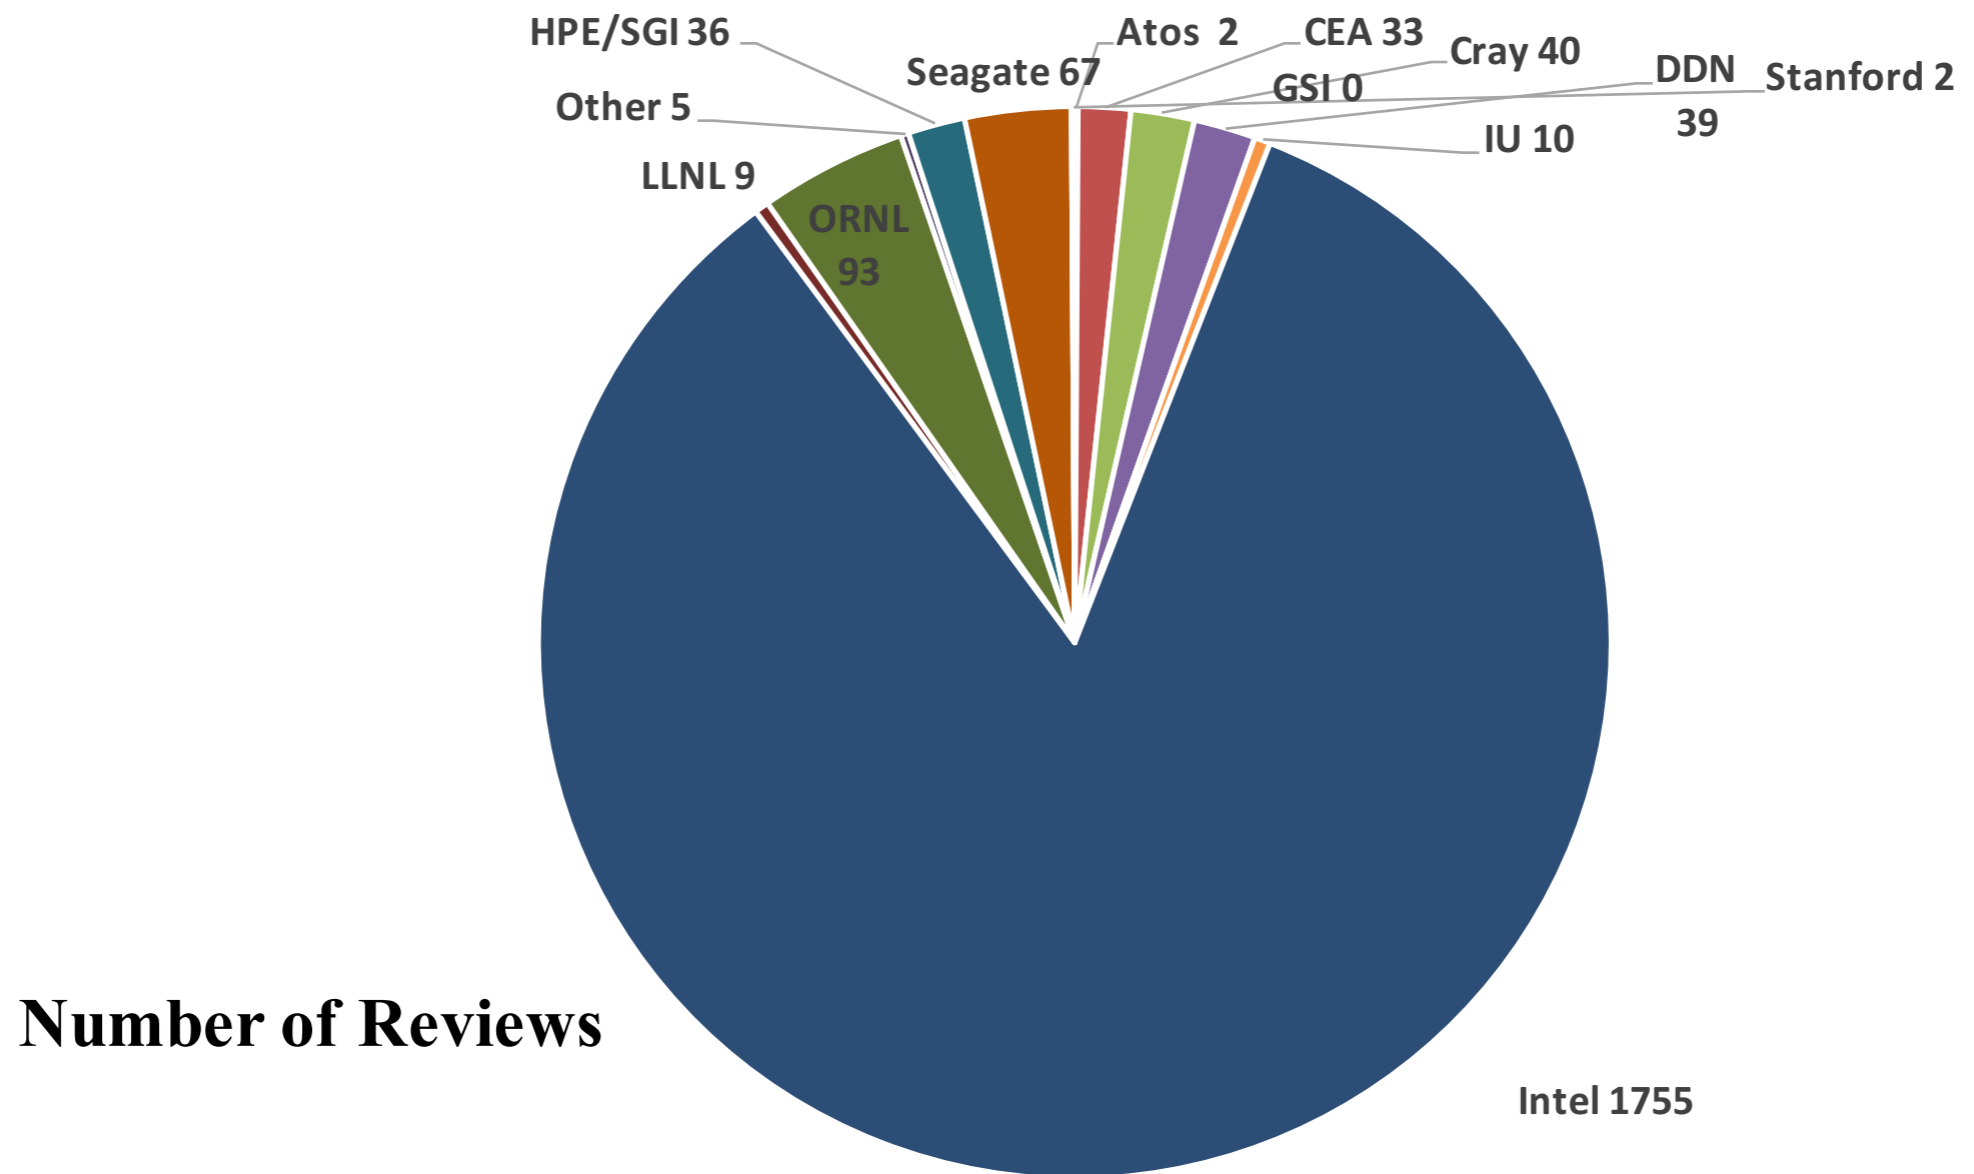


Statistics courtesy of Dustin Leverman (ORNL)

Aggregated data by organization between 2.9.50 and 2.10.0 tags

Source: <http://git.whamcloud.com/fs/lustre-release.git/shortlog/refs/heads/master>

Lustre 2.10 - Reviews



Number of Reviews

Source: <http://git.whamcloud.com/fs/lustre-release.git/shortlog/refs/heads/master>

Aggregated data by organization between 2.9.50 and 2.10.0 tags

Statistics courtesy of Dustin Leverman (ORNL)

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	742	201K	65	15
2.8.0	995	147K	92	17
2.9.0	737	74K	121	16
2.10.0	732	108K	85	14

Statistics courtesy of Chris Morrone (LLNL)/ Dustin Leverman (ORNL)

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

LTS Releases

- Long Term Stable (LTS) releases will remain active for at least 18 months
 - Similar idea to Whamcloud maintenance release streams
 - Updates for bugfixes and to support newer Linux distro releases
 - Will announce details of next release as they become available
- Maintenance releases may also be produced for other community releases but only until next major release
 - i.e. no further updates for 2.n when 2.n+1 has been released
 - Will depend upon levels of adoption


Lustre 2.10.x Maintenance Releases

- Lustre 2.10.1 maintenance release approaching GA
 - RHEL 7.4 server and client support
 - ZFS 0.7.1 default version of ZFS used
 - MOFED 4.1 support
 - Support for 4.9 kernel Lustre clients (LU-9183)
 - RHEL 6.9 Lustre client support
 - Ubuntu 16.04 LTS Lustre client support
- Coming in Future 2.10.x maintenance releases
 - Patchless servers (LU-20)
 - Support for 4.12 kernel Lustre clients (LU-9558)
 - Single thread performance improvements (LU-8964)
 - SLES12 SPx server support

Lustre 2.11

- Targeted for Feb 2018 release
- Will support
 - RHEL 7.4 servers/clients
 - SLES12 SP3 clients
 - Ubuntu 16.04 clients
- Interop/upgrades from latest Lustre 2.10.x servers/clients
- Several features targeted for this release
 - Lock-ahead (LU-6179)
 - Data on MDT (LU-3285)
 - File Level Redundancy – Delayed Resync (LU-9771)
- http://wiki.lustre.org/Release_2.11.0

Community Release Roadmap



*LTS Release with maintenance releases provided

Estimates are not commitments and are provided for informational purposes only

Fuller details of features in development are available at <http://wiki.lustre.org/Projects>

Last updated: April 20th 2017

IML 4.0

- IML 4.0 approaching GA
- <https://github.com/intel-hpdd/intel-manager-for-lustre/releases>
- First open source release of IML
 - Provides intuitive browser-based administration of Lustre filesystems
 - Distributed under an MIT license
 - Compatible with Lustre 2.10.1
- Possible for sites running Intel EE 2.x and 3.x to upgrade
- IML releases are now decoupled from Lustre releases
 - This will mean more flexibility for release timing
- Simple for those interested to setup for demo or development
 - <https://github.com/intel-hpdd/intel-manager-for-lustre/wiki/Installing-IML-on-HPC-Storage-Sandbox>

IML 4.1

- Planned for Q1 2018; content being finalized
 - <https://github.com/intel-hpdd/intel-manager-for-lustre/issues>
 - **Your input into this process is welcomed!**
 - Mechanisms in place for community contributions
- IML is in plans for upcoming extreme scale deployments
 - This will influence roadmap over coming releases
 - Event-driven architectures vs polling
- IML to become more modular
 - Allows users to pick and choose which parts to utilize
- ZED provides interesting opportunities in both scale and scope
 - Leverage to discover pools, datasets and properties
 - Monitoring and alerts

Lustre Release Documentation

- Latest version of manual dynamically available to download
 - <http://lustre.org/documentation/>
 - Also links for how to contribute
- If you know of gaps then please open an LUDOC ticket
 - If you have not got time to work out the correct format to submit then unformatted text will provide a starting point for someone else to complete
- Large amount of content being added on lustre.org
 - http://wiki.lustre.org/Category:Lustre_Systems_Administration
 - Lustre Internals content being refreshed

Summary

- Lustre 2.10.1 approaching GA; Lustre 2.10.2 targeted for Q4
- Feature freeze for Lustre 2.11 is approaching
- IML 4.0 approaching GA
- There is still time to influence content for IML 4.1
- Plenty of options for contributing to Lustre/IML

Thank you

Open Scalable File Systems, Inc.

3855 SW 153rd Drive

Beaverton, OR 97006

Ph: 503-619-0561

Fax: 503-644-6708

admin@opensfs.org



www.opensfs.org