



Community Release Update

LUG 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

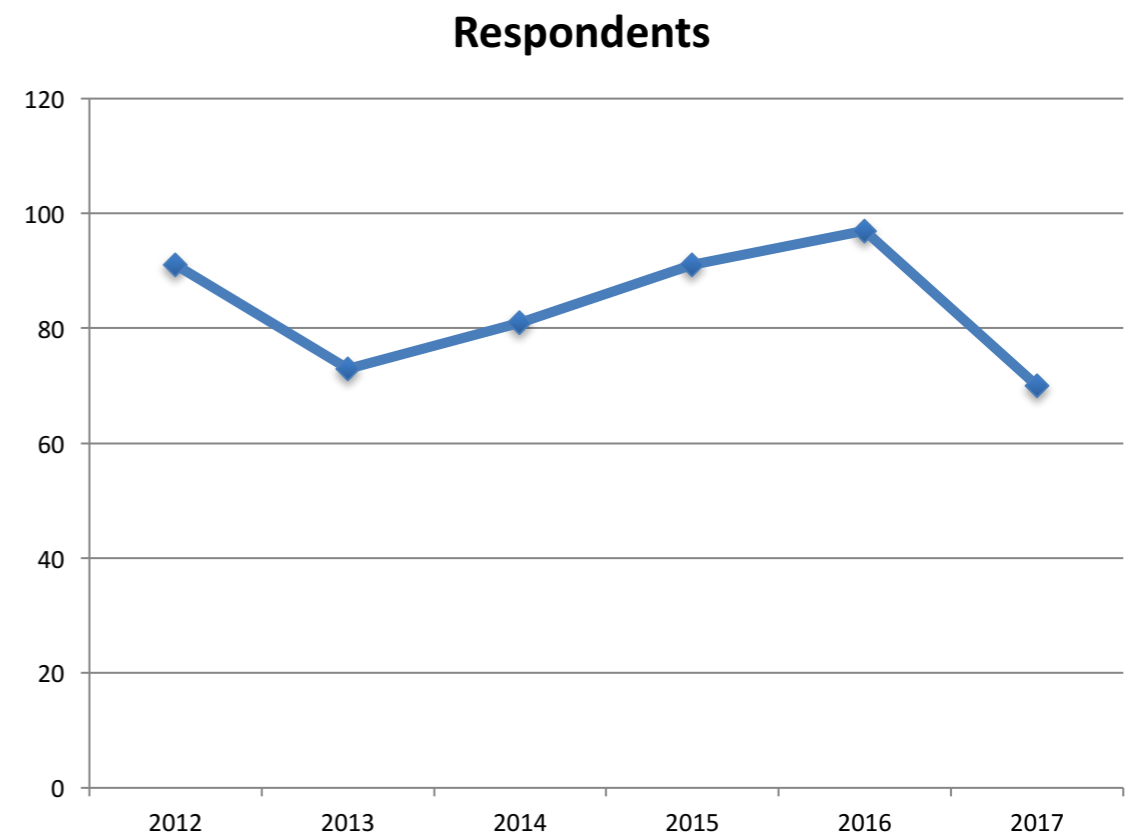
- Oversees entire Lustre development cycle
- Maintains the roadmap
- Plans major releases
- Collects requirements for future Lustre features
- Sets priorities for test matrix

For more information visit the wiki

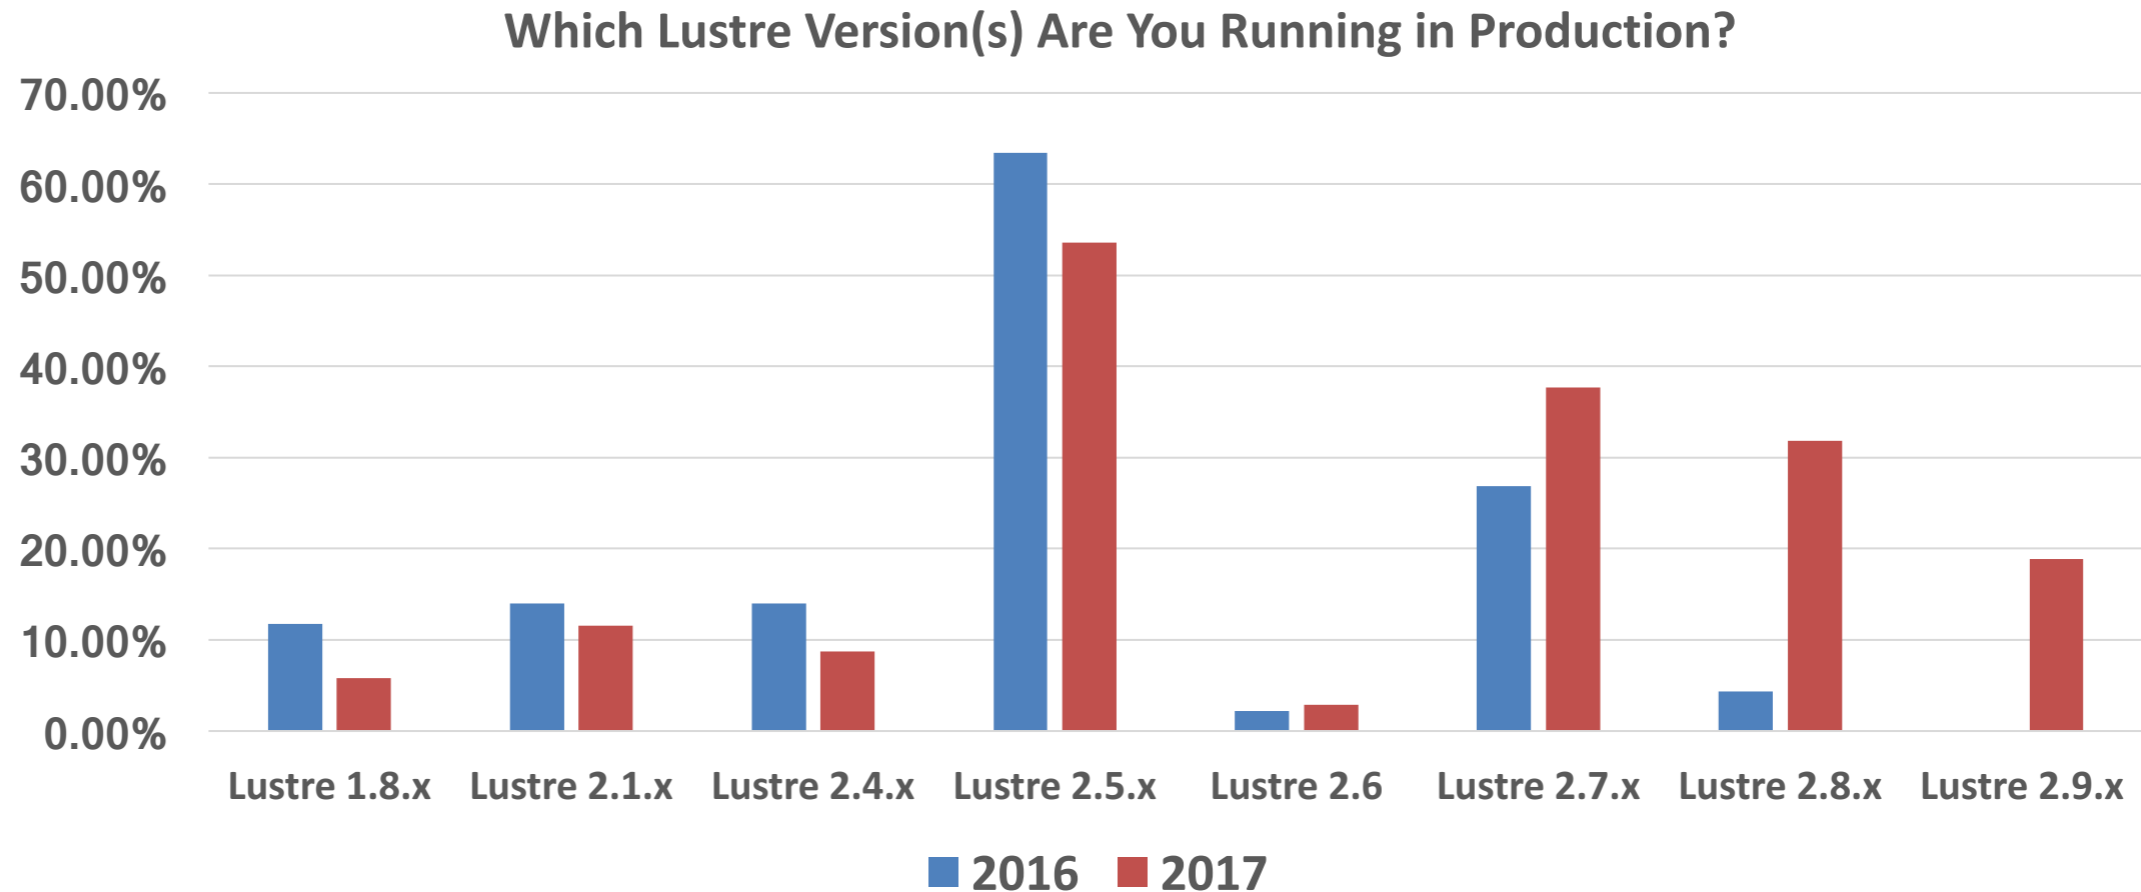
http://wiki.opensfs.org/Lustre_Working_Group

Community Survey

- Has been running for 6 years
- LWG devises questions
- Useful for tracking trends in Lustre usage
- Drop in respondents compared to 2016 survey (97 to 70)
- Full details available at http://wiki.opensfs.org/Lustre_Community_Survey

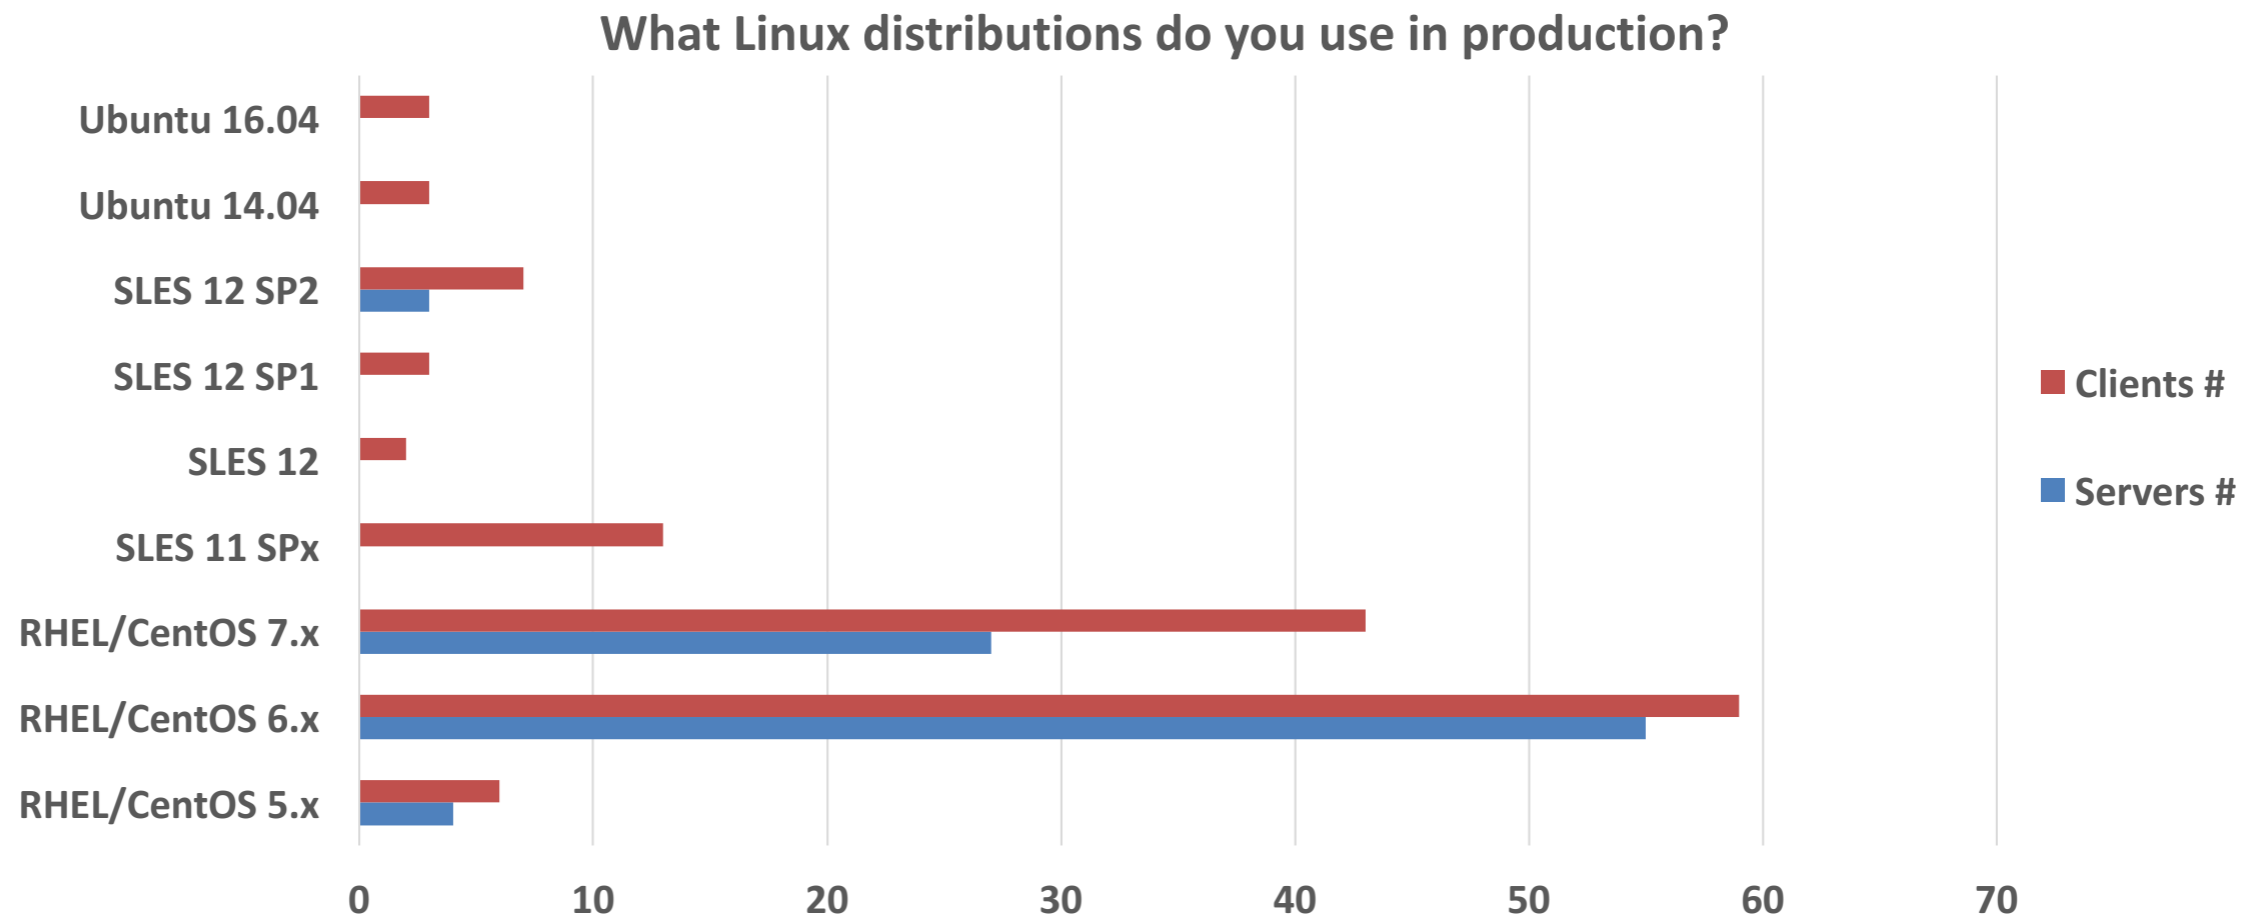


Community Survey - Versions



- Lustre 2.5.x remains the most commonly-used version in production
- Strong adoption of more current releases

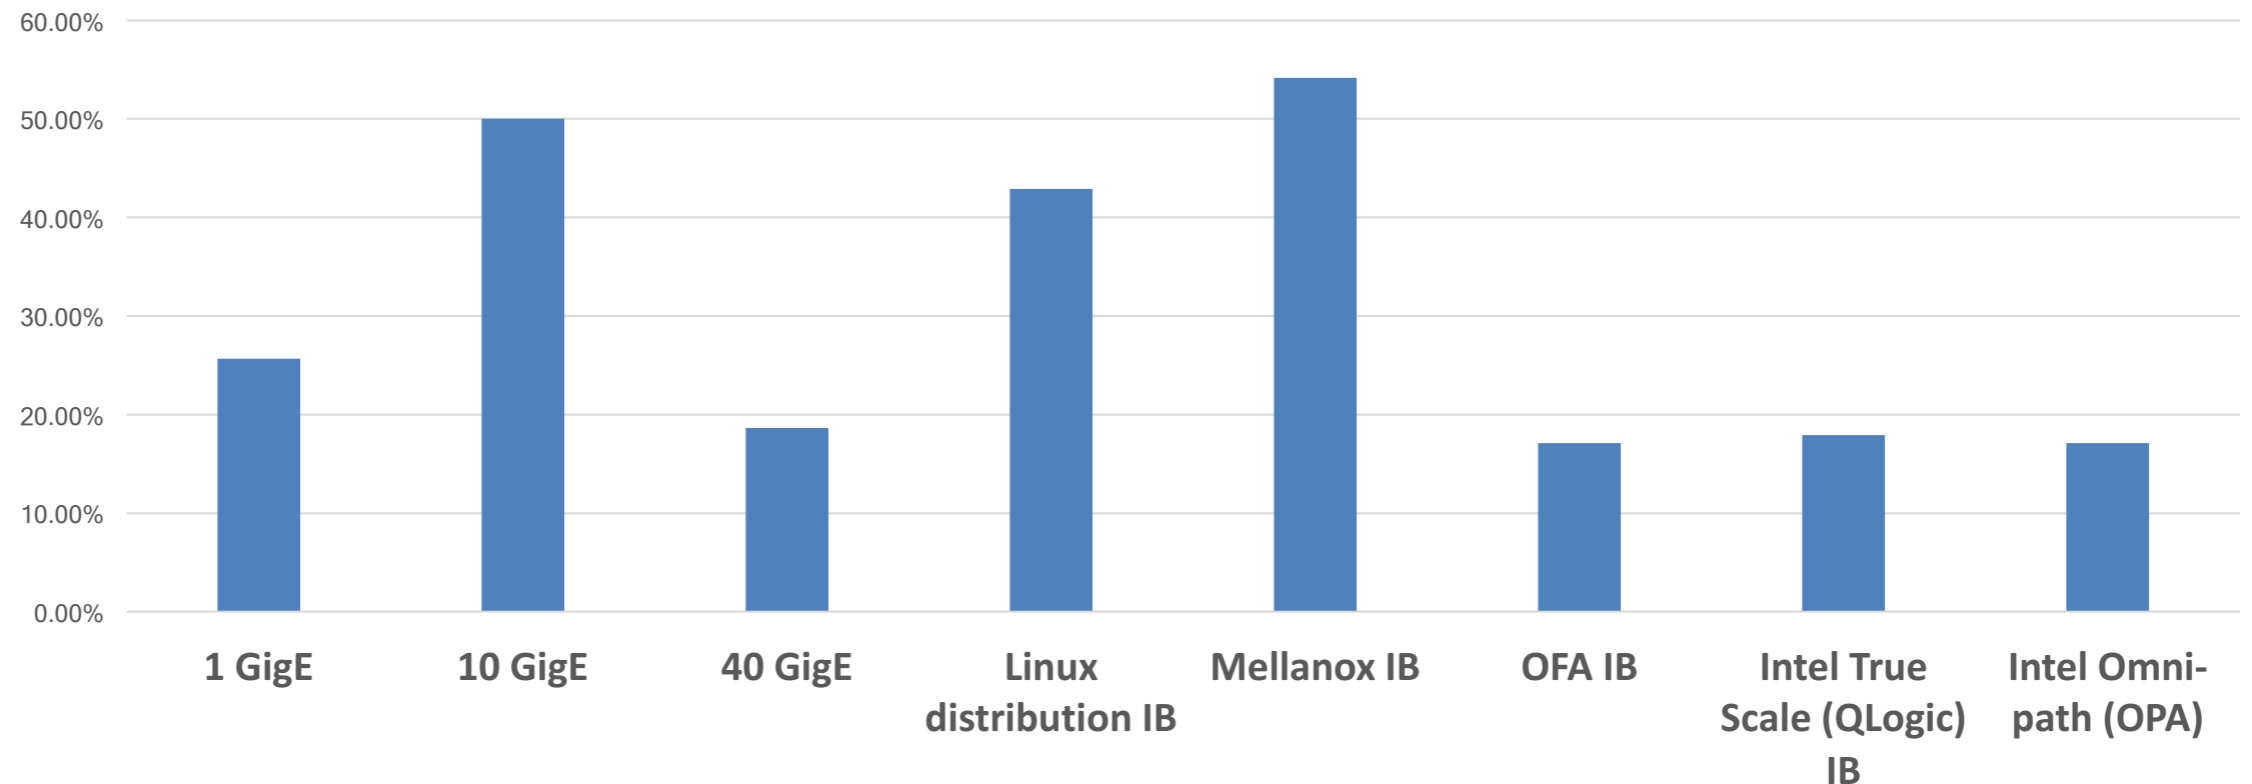
Community Survey - Linux distros



- RHEL/CentOS 6.x still the most widely used distro
- RHEL/CentOS 7.x now widely used

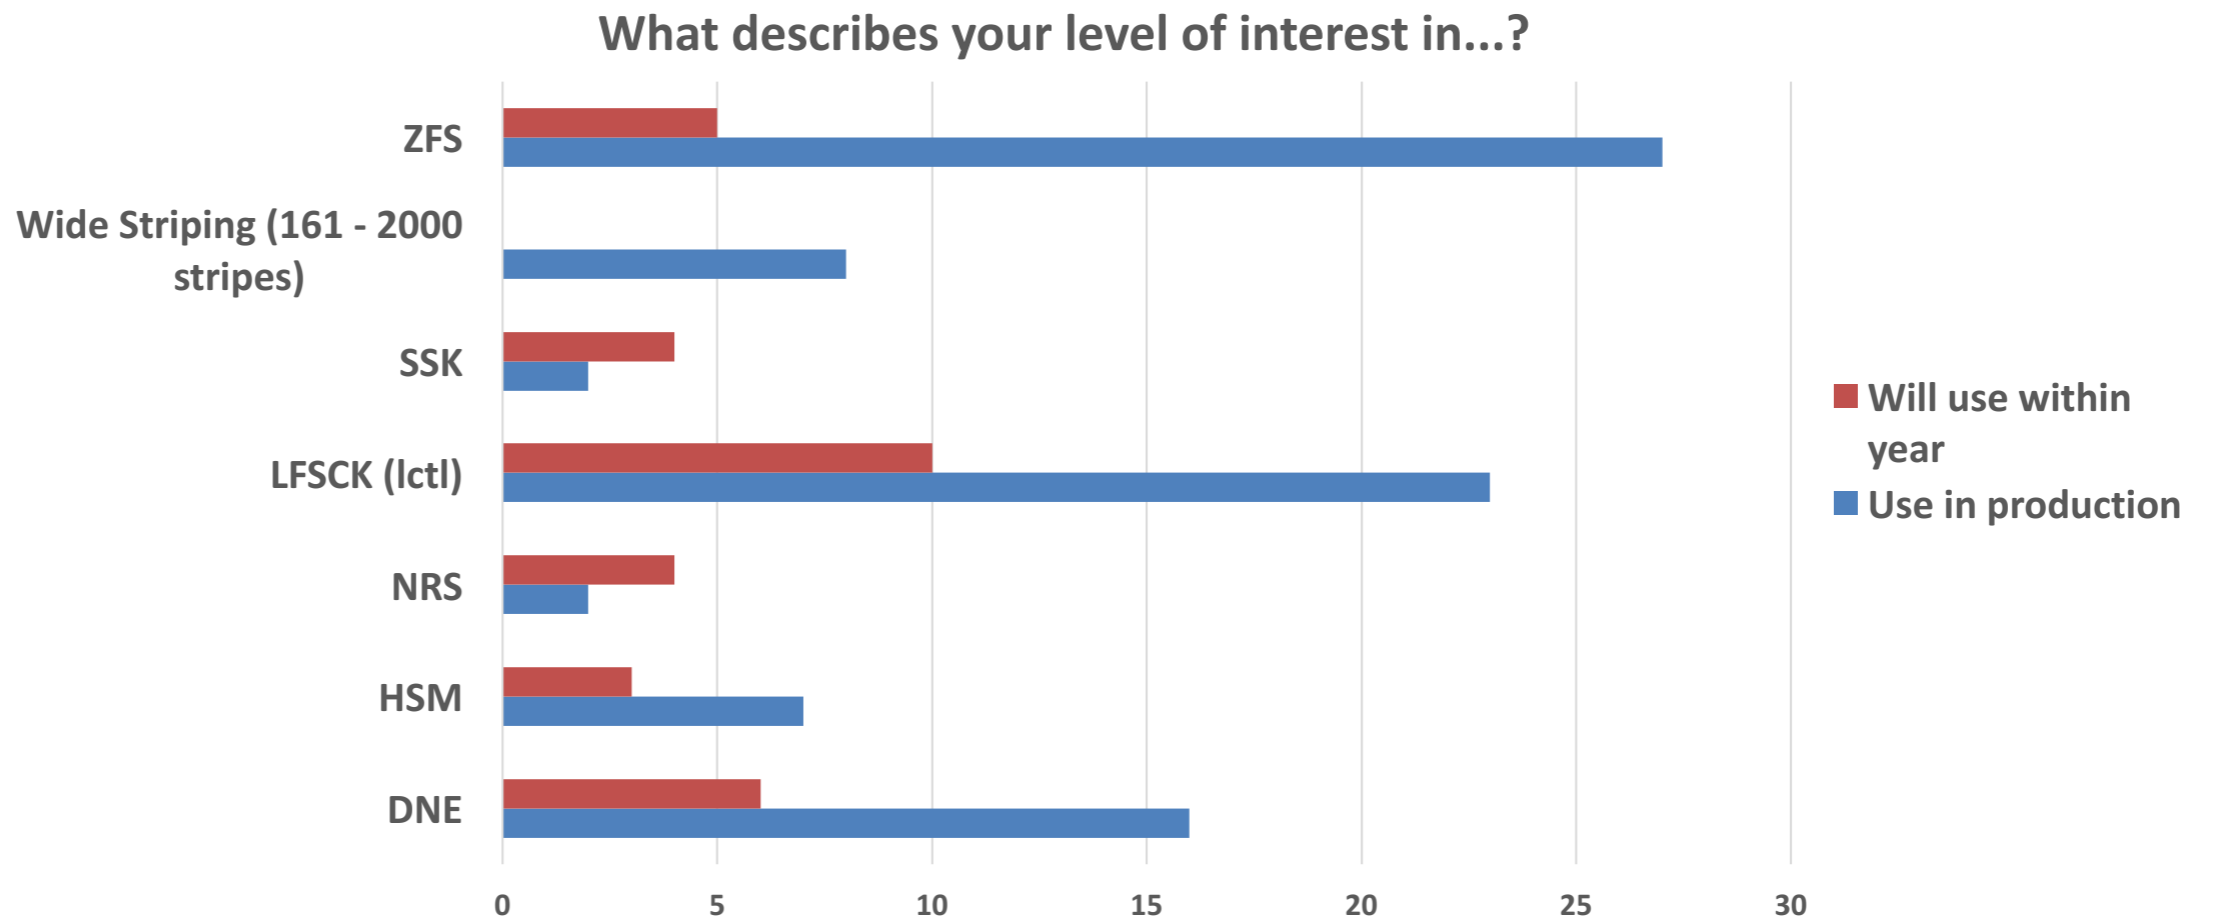
Community Survey - Networks

Which kind of networks do you use with Lustre within your organization?



- Many options used but Mellanox and 10 GigE remain the most common

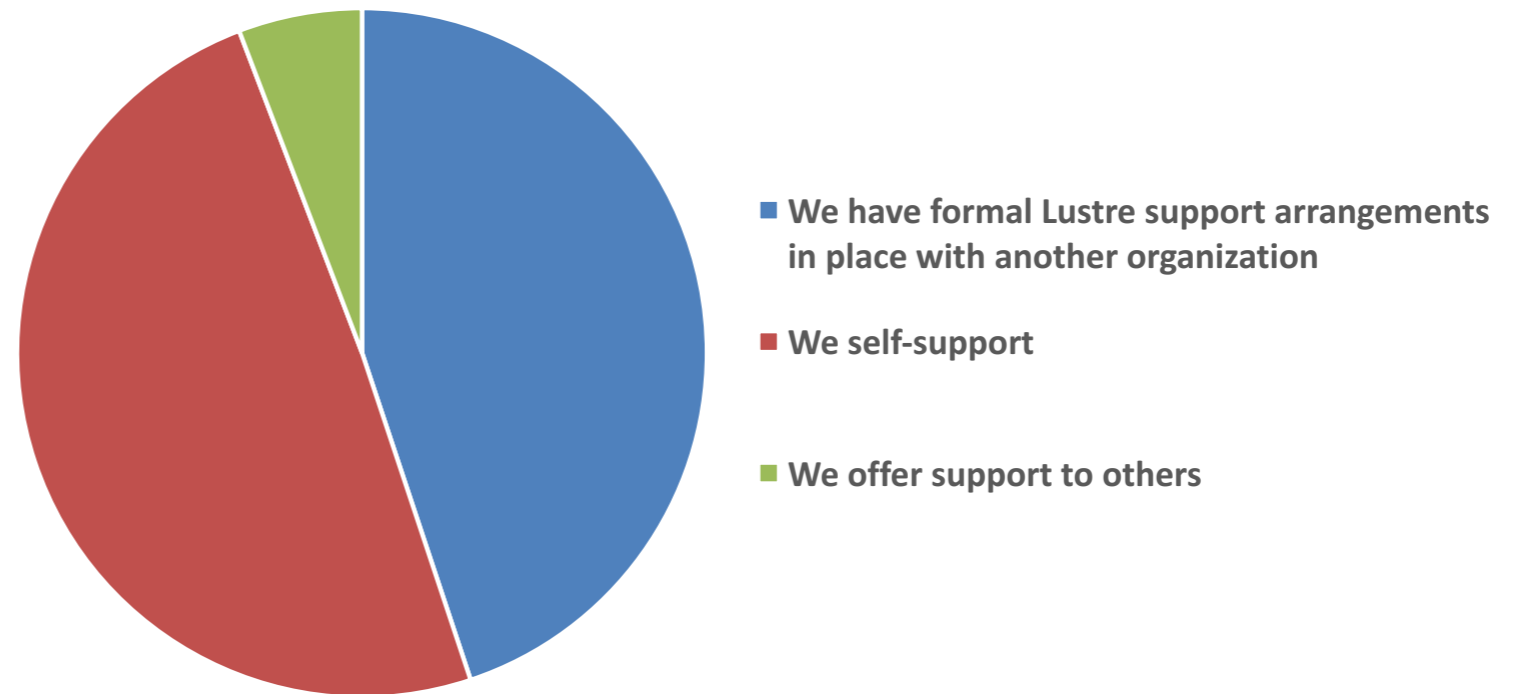
Community Survey – New Features



- DNE and LFSCK usage grew since last year

Community Survey - Support

What best describes your Lustre support situation?



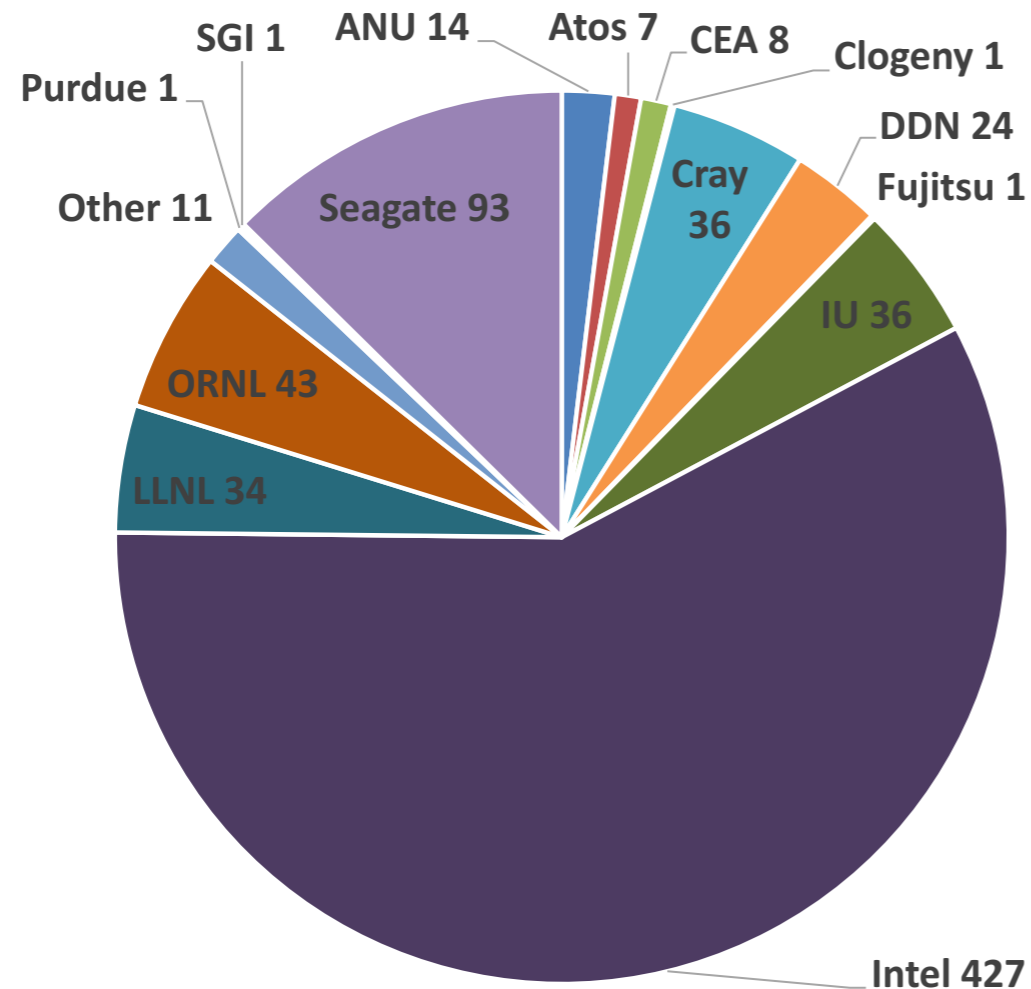
- Even split between those self-supporting and those using/providing vendor support

Lustre 2.9

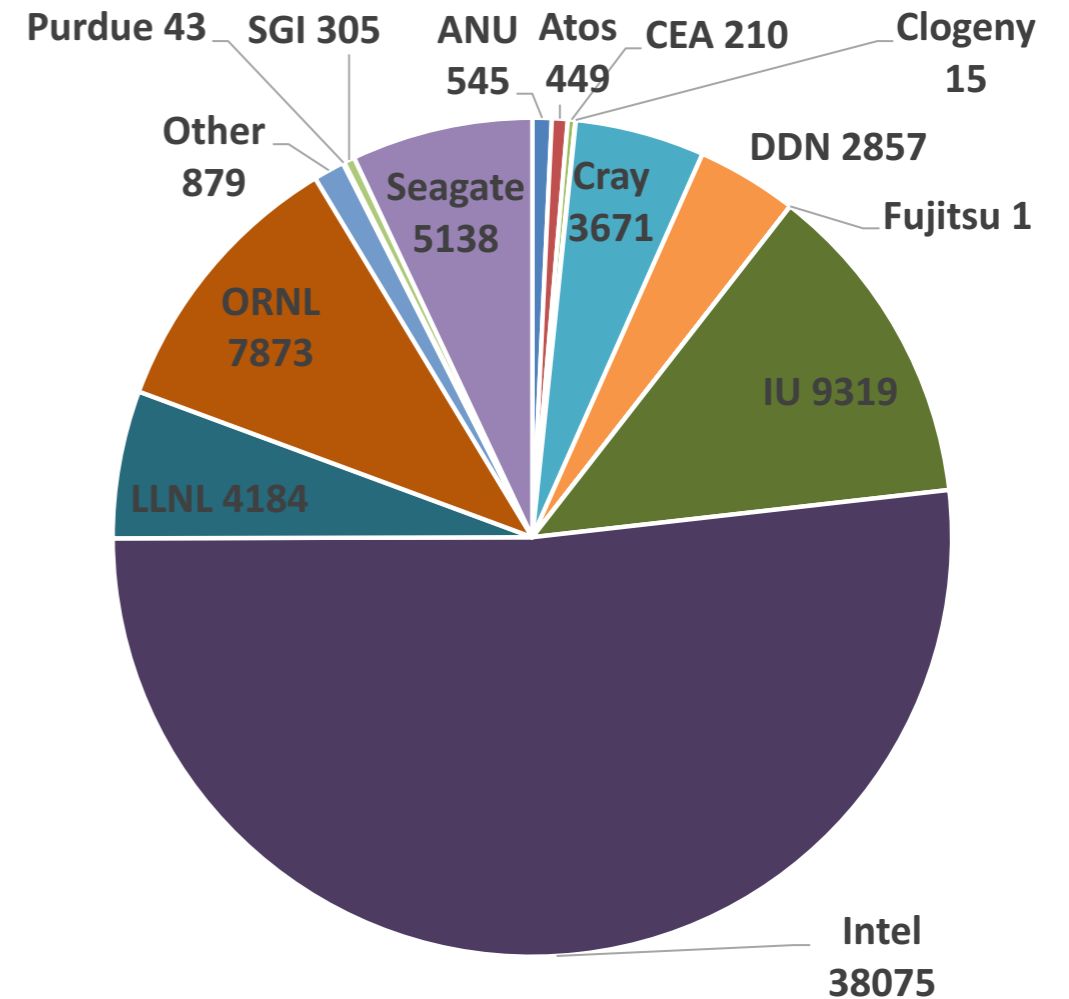
- GA December 2016
- Supports RHEL 7.3 servers/clients and SLES12 SP1 clients
- Varied features
 - UID/GID mapping (LU-3291)
 - Shared Secret Key Crypto (LU-3289)
 - Subdirectory Mounts (LU-28)
 - ZFS Metadata Improvements (LU-7895)
 - RPM Kernel Module Weak Updates (LU-5614)
 - Server Side I/O Advise (ladvice) (LU-4931)
 - Large Bulk IO (LU-7990)
- http://wiki.lustre.org/Release_2.9.0

Lustre 2.9 - Contributions

Number of Commits



Lines of Code Changed

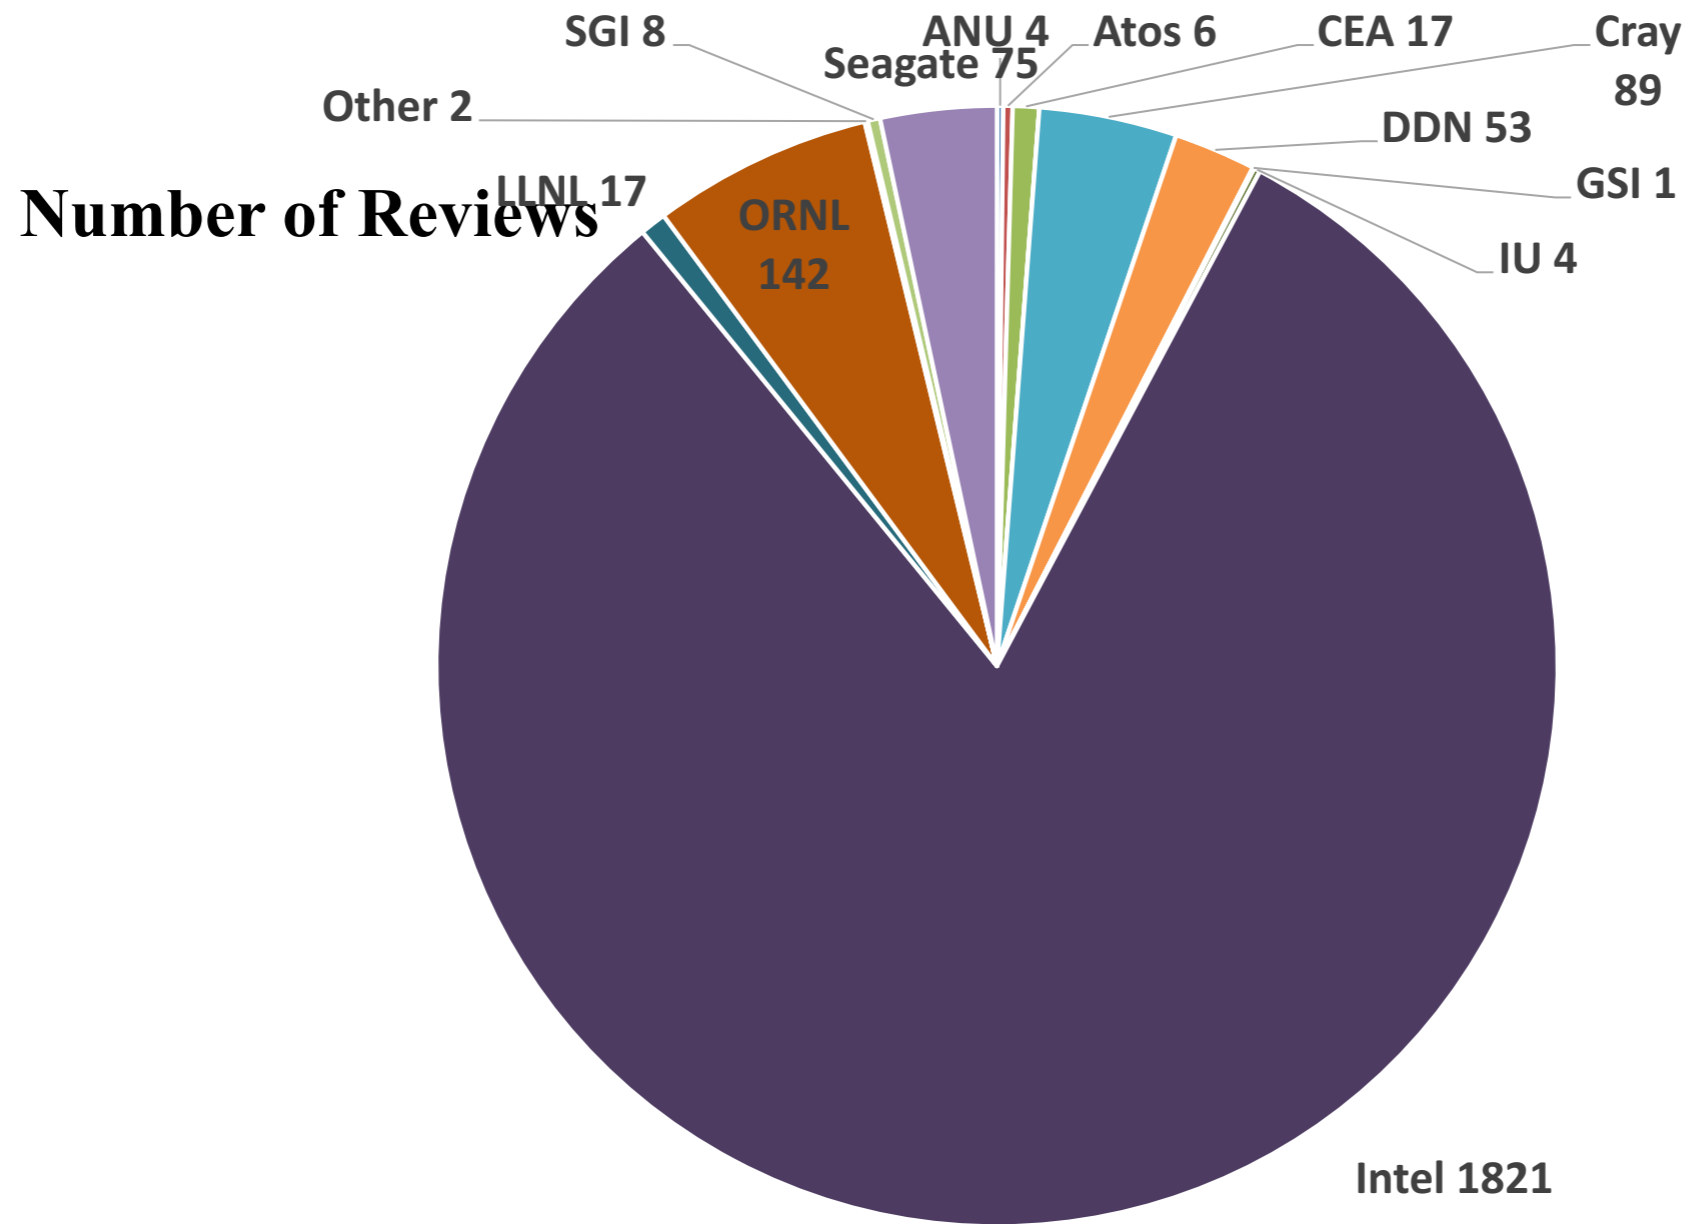


Statistics courtesy of Dustin Leverman (ORNL)

Source: http://git.whamcloud.com/fs/lustre-release.git/shortlog/refs/heads/b2_9

Aggregated data by organization between 2.8.50 and 2.9.0 tags

Lustre 2.9 - Reviews



Statistics courtesy of Dustin Leverman (ORNL)

Aggregated data by organization between 2.8.50 and 2.9.0 tags

Lustre 2.10

- Targeted GA June 2017
 - Latest status on how this is trending...
- Will support RHEL 7.3 servers/clients and SLES12 SP2 clients
- Interop/upgrades from Lustre 2.9 servers/clients
- Will be designated an LTS Release and have 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)
- http://wiki.lustre.org/Release_2.10.0

Lustre 2.10 – Progressive File Layouts

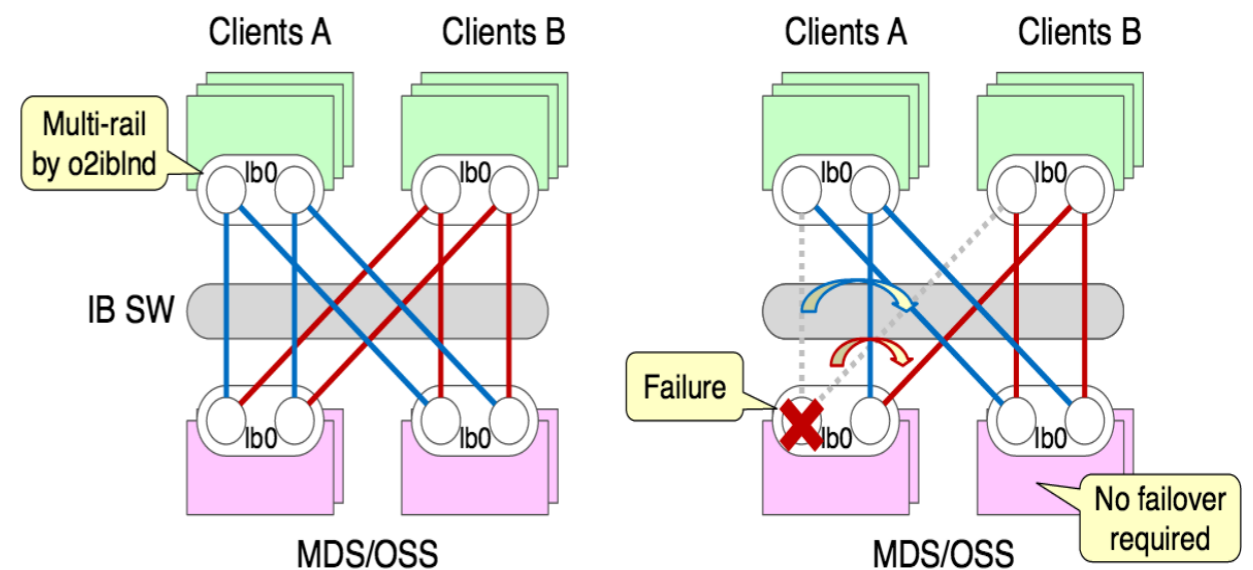
- Progressive File Layout (PFL) simplifies usage for users and admins
 - Optimize performance for diverse users/applications
 - One PFL layout could be used for all files
 - Low stat overhead for small files
 - High IO bandwidth for large files
- Collaboration between Intel and ORNL

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.x – Additional Content

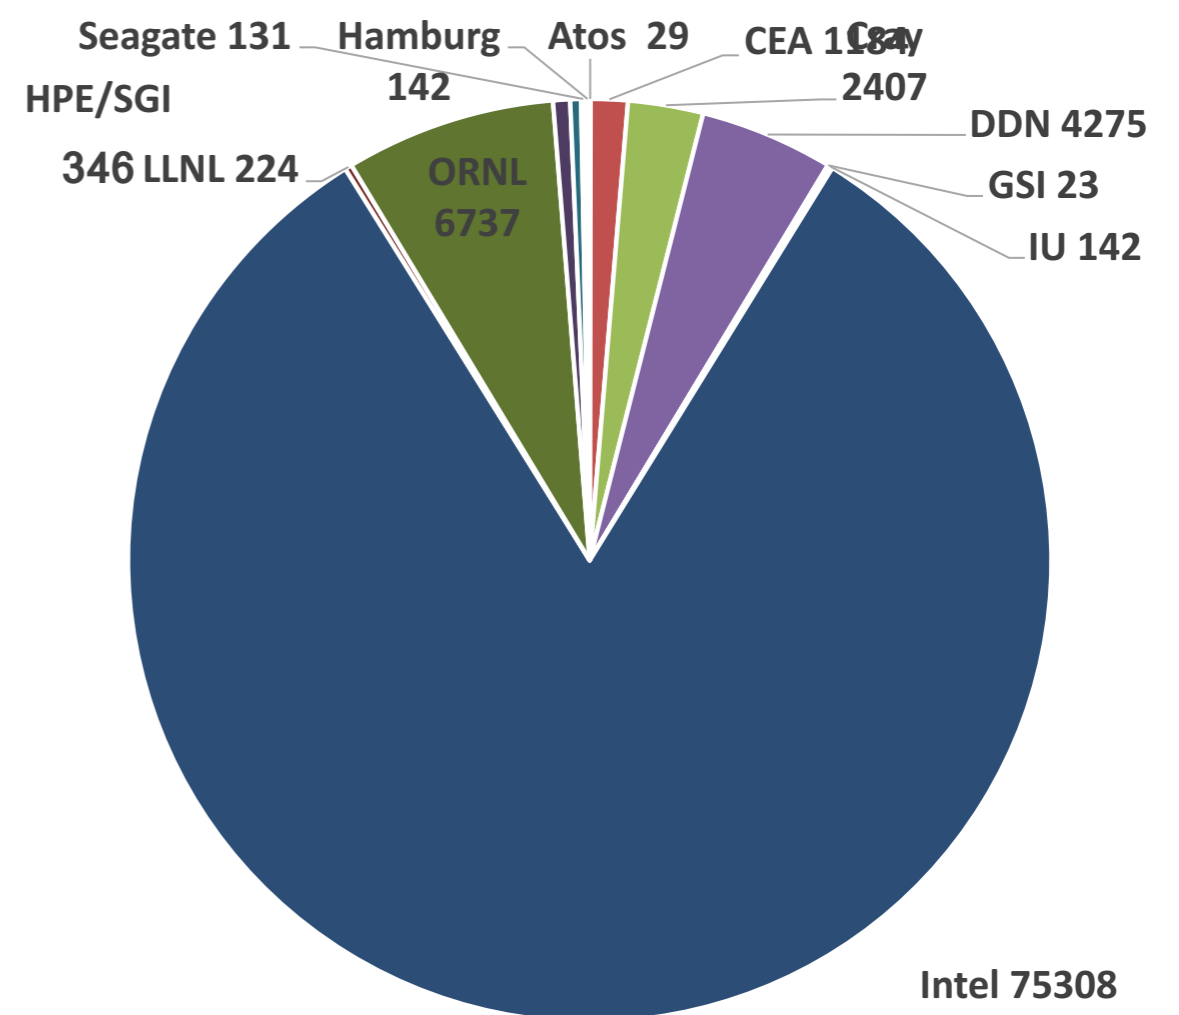
- Confirmed in 2.10.0
 - ZFS Metadata Improvements (LU-7895)
 - OPA Performance improvements (LU-8943)
 - Pacemaker scripts (LU-8457/8458)
 - Feature/bugfix parity with latest EE 3.x Lustre releases
- Coming in 2.10.x maintenance release
 - Patchless servers (LU-20)
 - Support for 4.9 kernel Lustre clients (LU-9183)
 - Single thread performance improvements (LU-8964)
 - SLES12 SP2 server support
 - Ubuntu 16.04 LTS Lustre client support
 - MOFED 4.x support

Lustre 2.10 - Contributions

Number of Commits



Lines of Code Changed

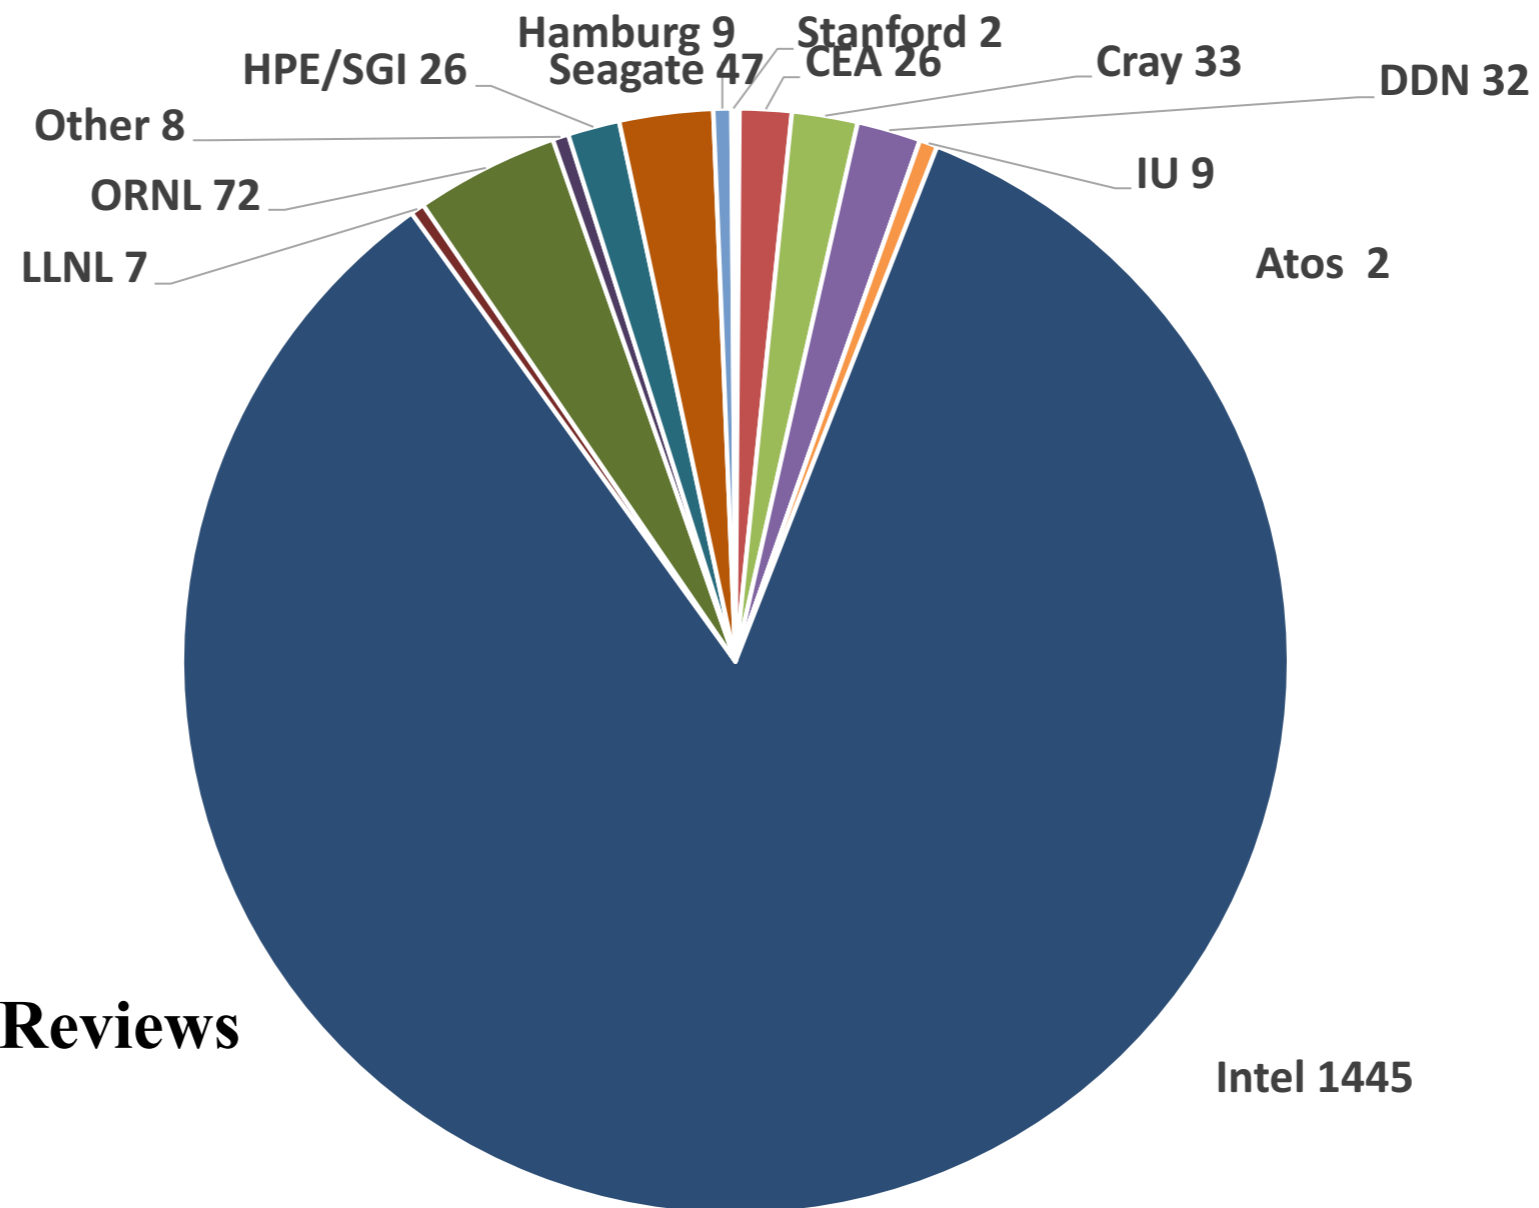


Statistics courtesy of Dustin Leverman (ORNL)

Aggregated data by organization between 2.9.50 and 2.9.56 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 2.9.58 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*	605	93K	81	13

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

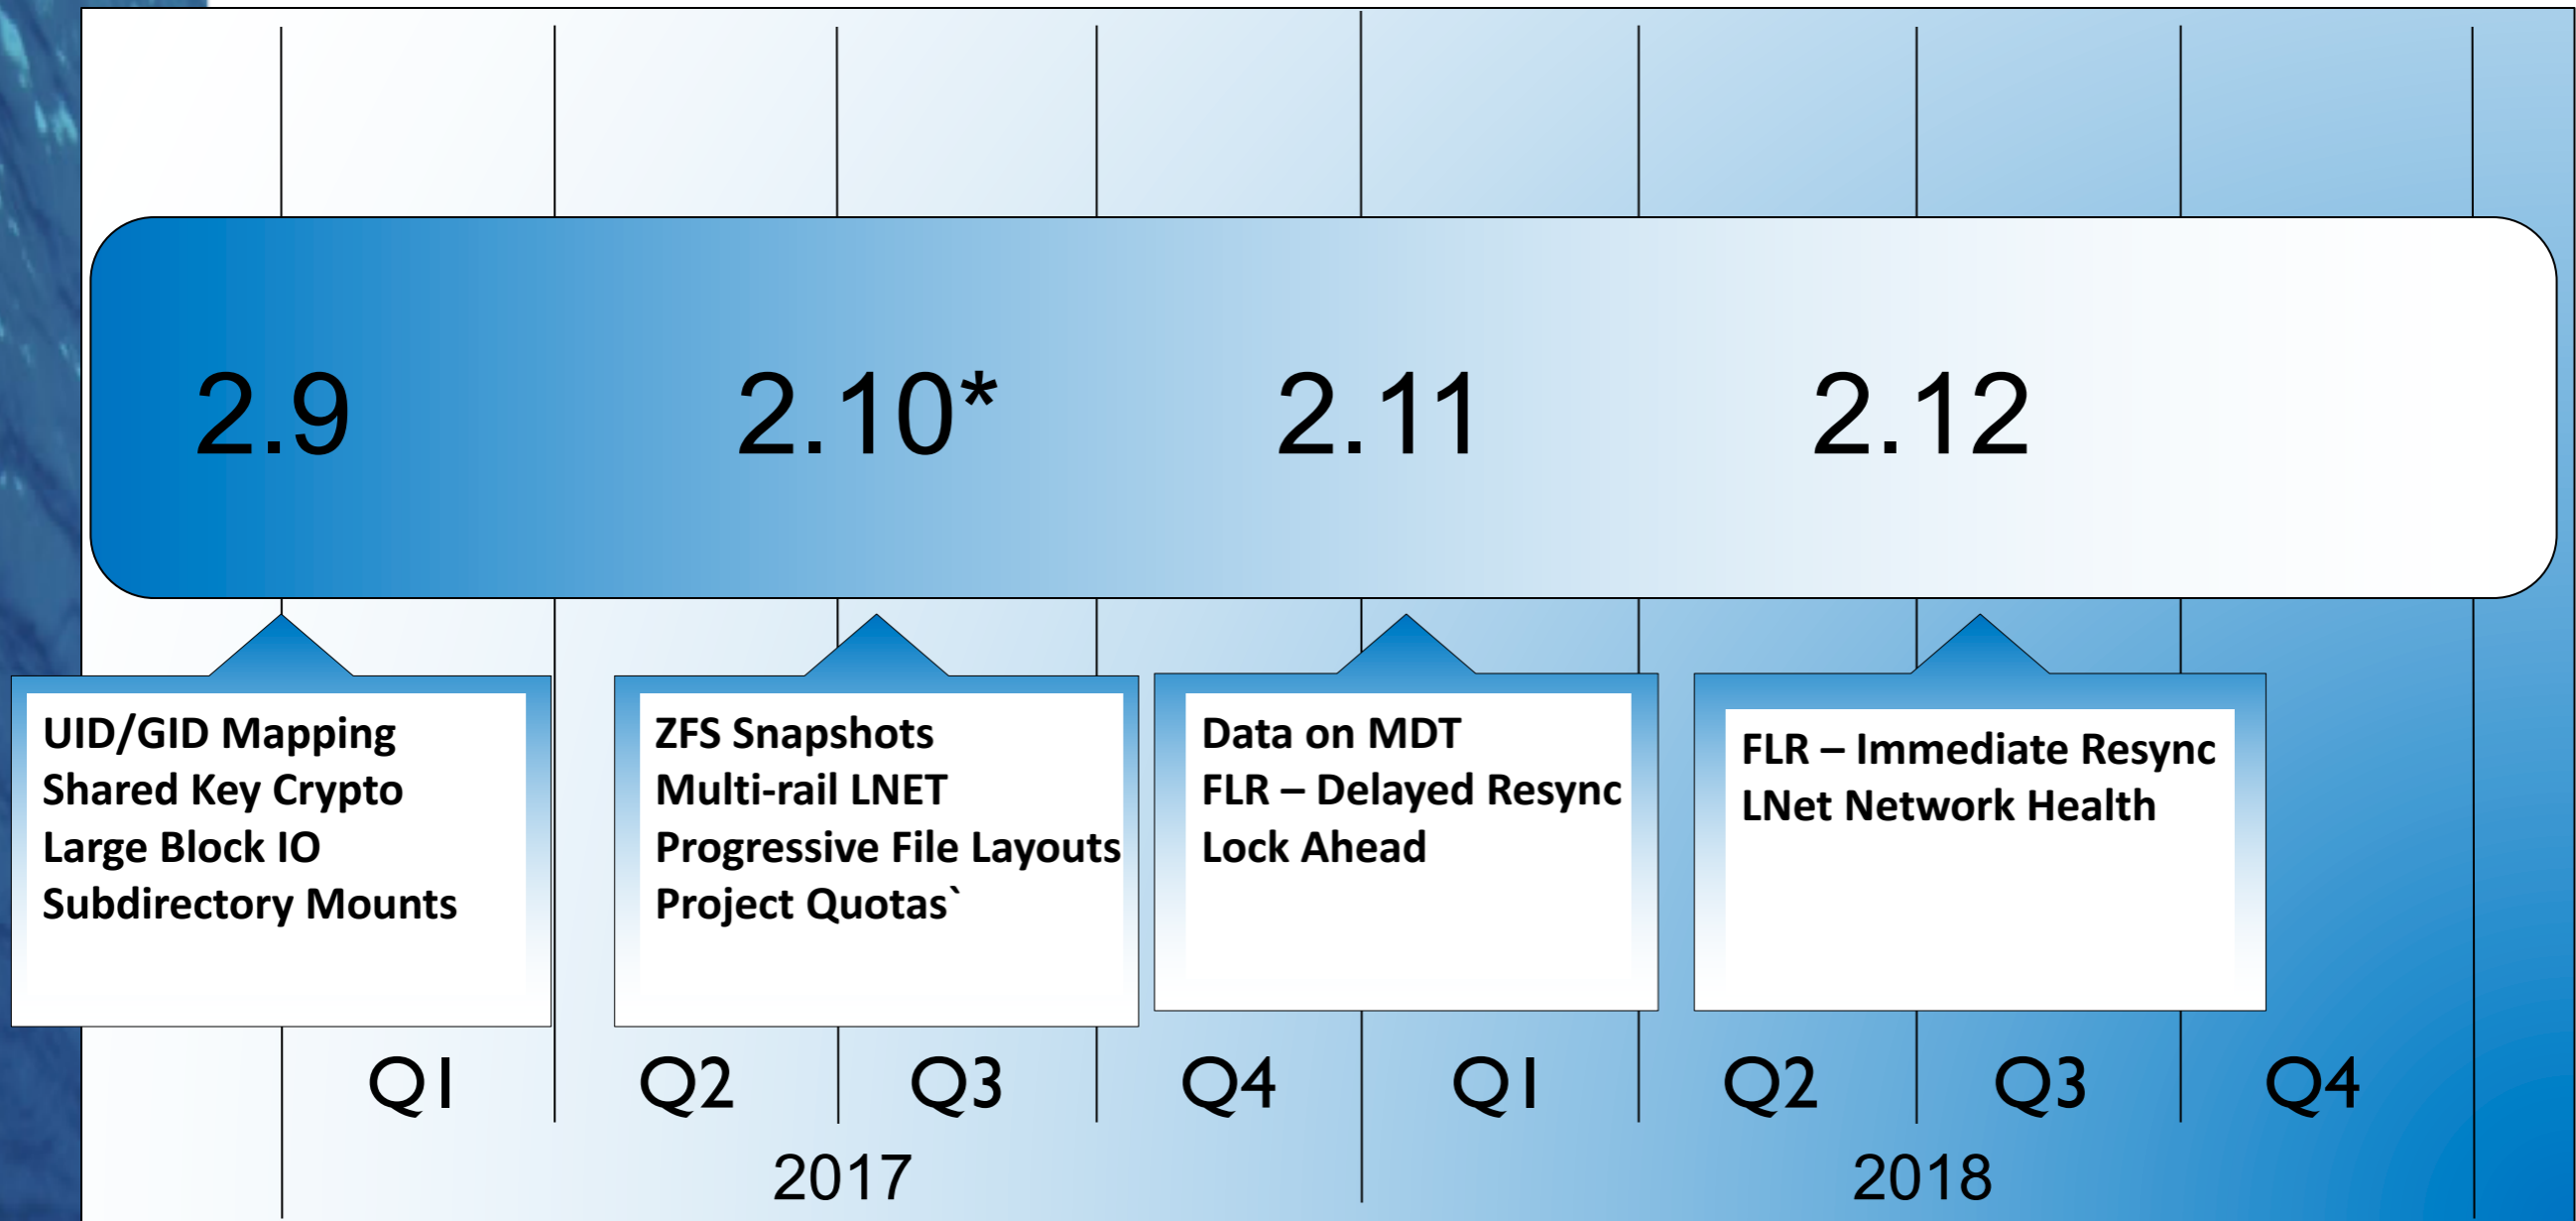
*Data for 2.10 is incomplete and only runs up until 2.9.58

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

Lustre 2.11 and Beyond

- Many interesting long term features in development
 - File Level Redundancy
 - Data on MDT
 - LNet Network Health
- Full list maintained on lustre.org wiki
 - <http://wiki.lustre.org/Projects>
- Roadmap kept up to date on lustre.org
 - <http://lustre.org/roadmap/>

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

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 updates for 2.n when 2.n+1 has been released
 - Will depend upon levels of adoption

Lustre Release Documentation

- Latest version of user documentation dynamically available to download
 - <http://lustre.org/documentation/>
 - Also links to details on 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
- More dynamic content on <http://wiki.lustre.org>

Summary

- Lustre 2.10.0 is expected shortly
- Freely available 2.10.x maintenance releases planned
- Plenty of options for those interest in contributing to Lustre

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