



Community Release Update

September 23rd 2019

Peter Jones, Whamcloud

OpenSFS Lustre Working Group



OpenSFS Lustre Working Group

Lead by Peter Jones (Whamcloud) 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

All welcome to attend and/or join the mailing list

For more information visit the wiki

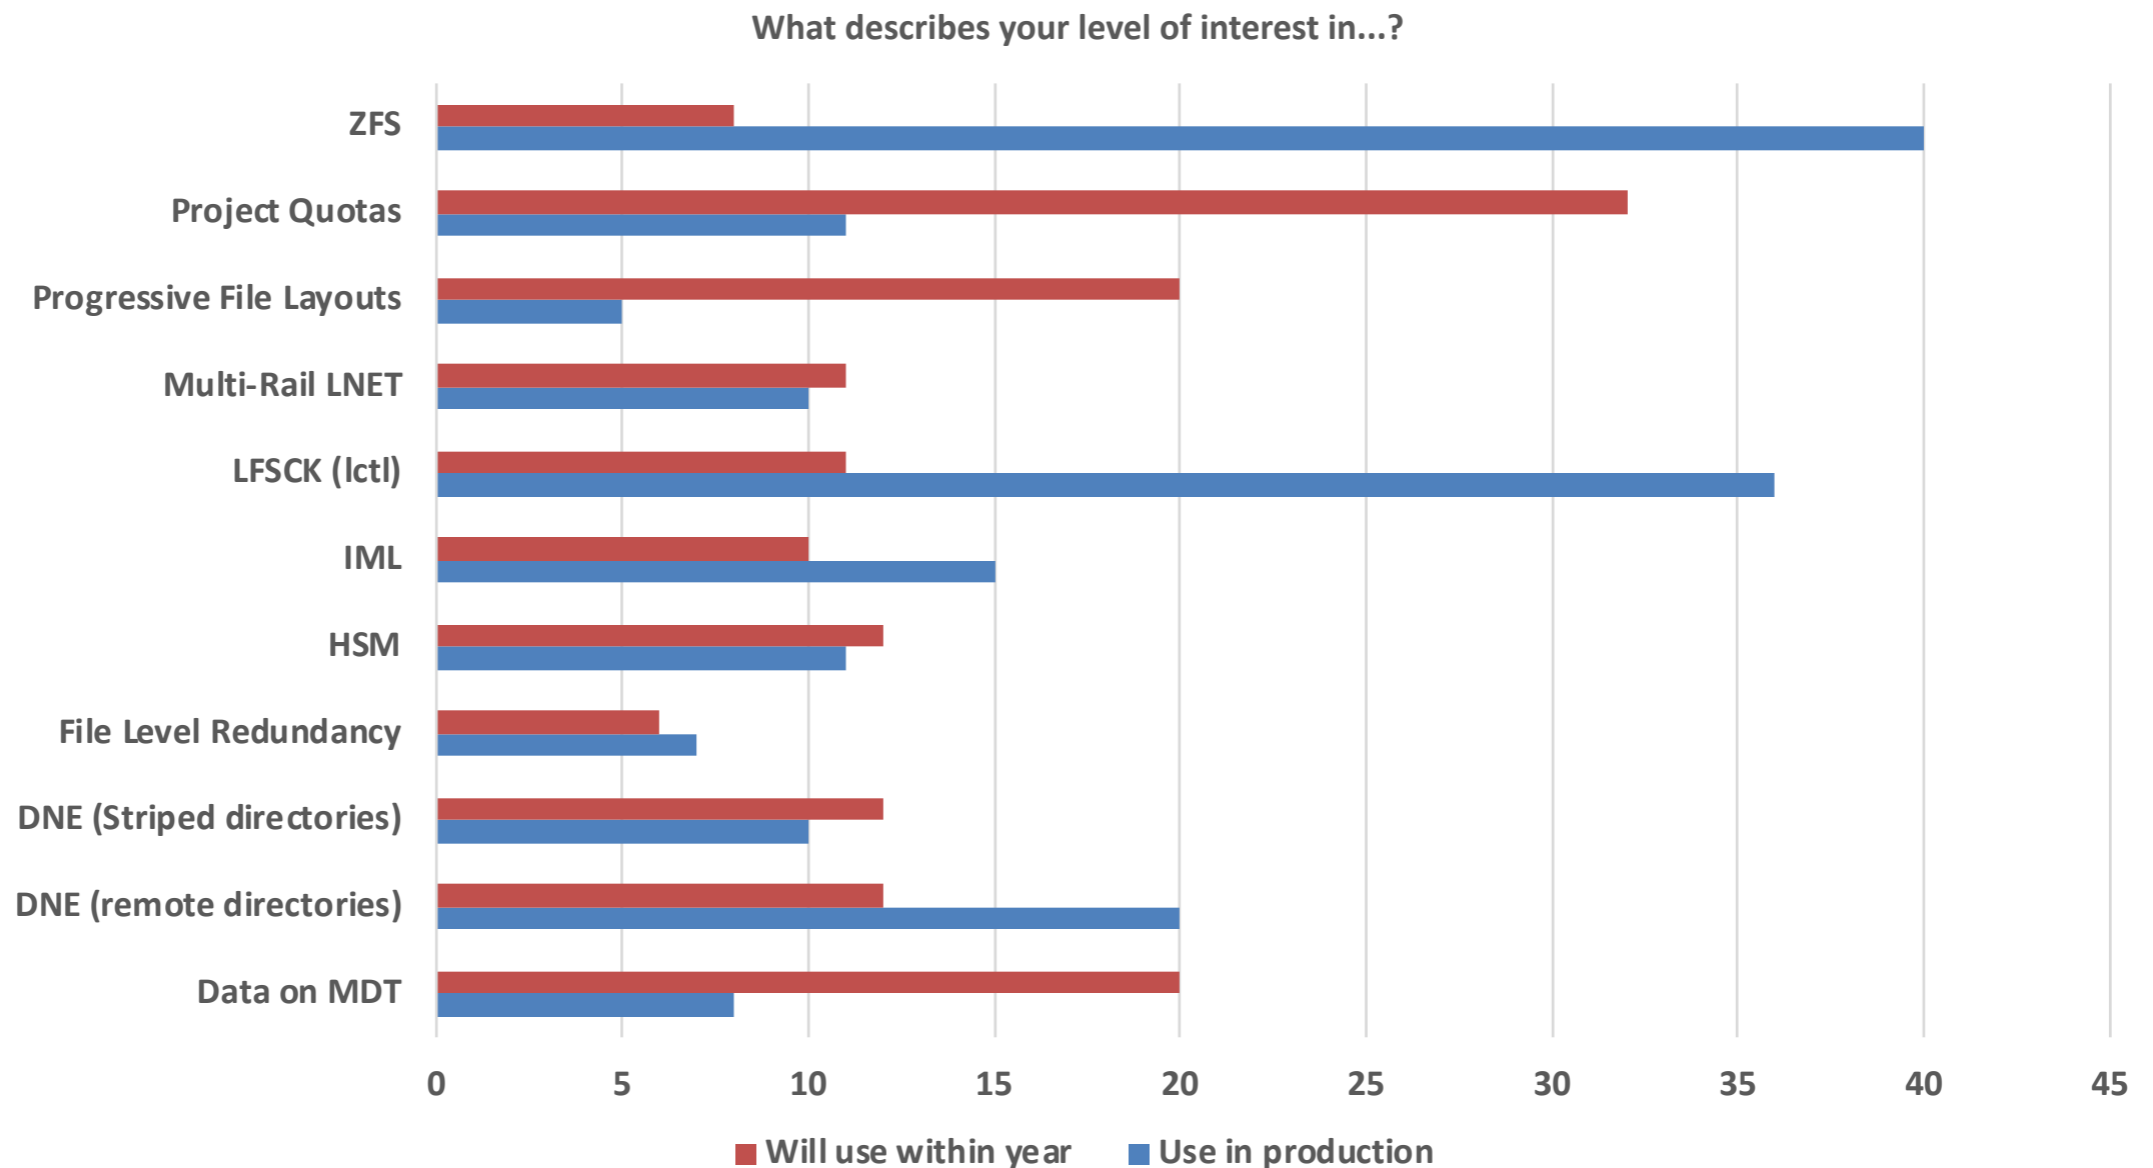
[http://wiki.opensfs.org/Lustre Working Group](http://wiki.opensfs.org/Lustre_Working_Group)

Lustre Community Survey (Mar 2019)



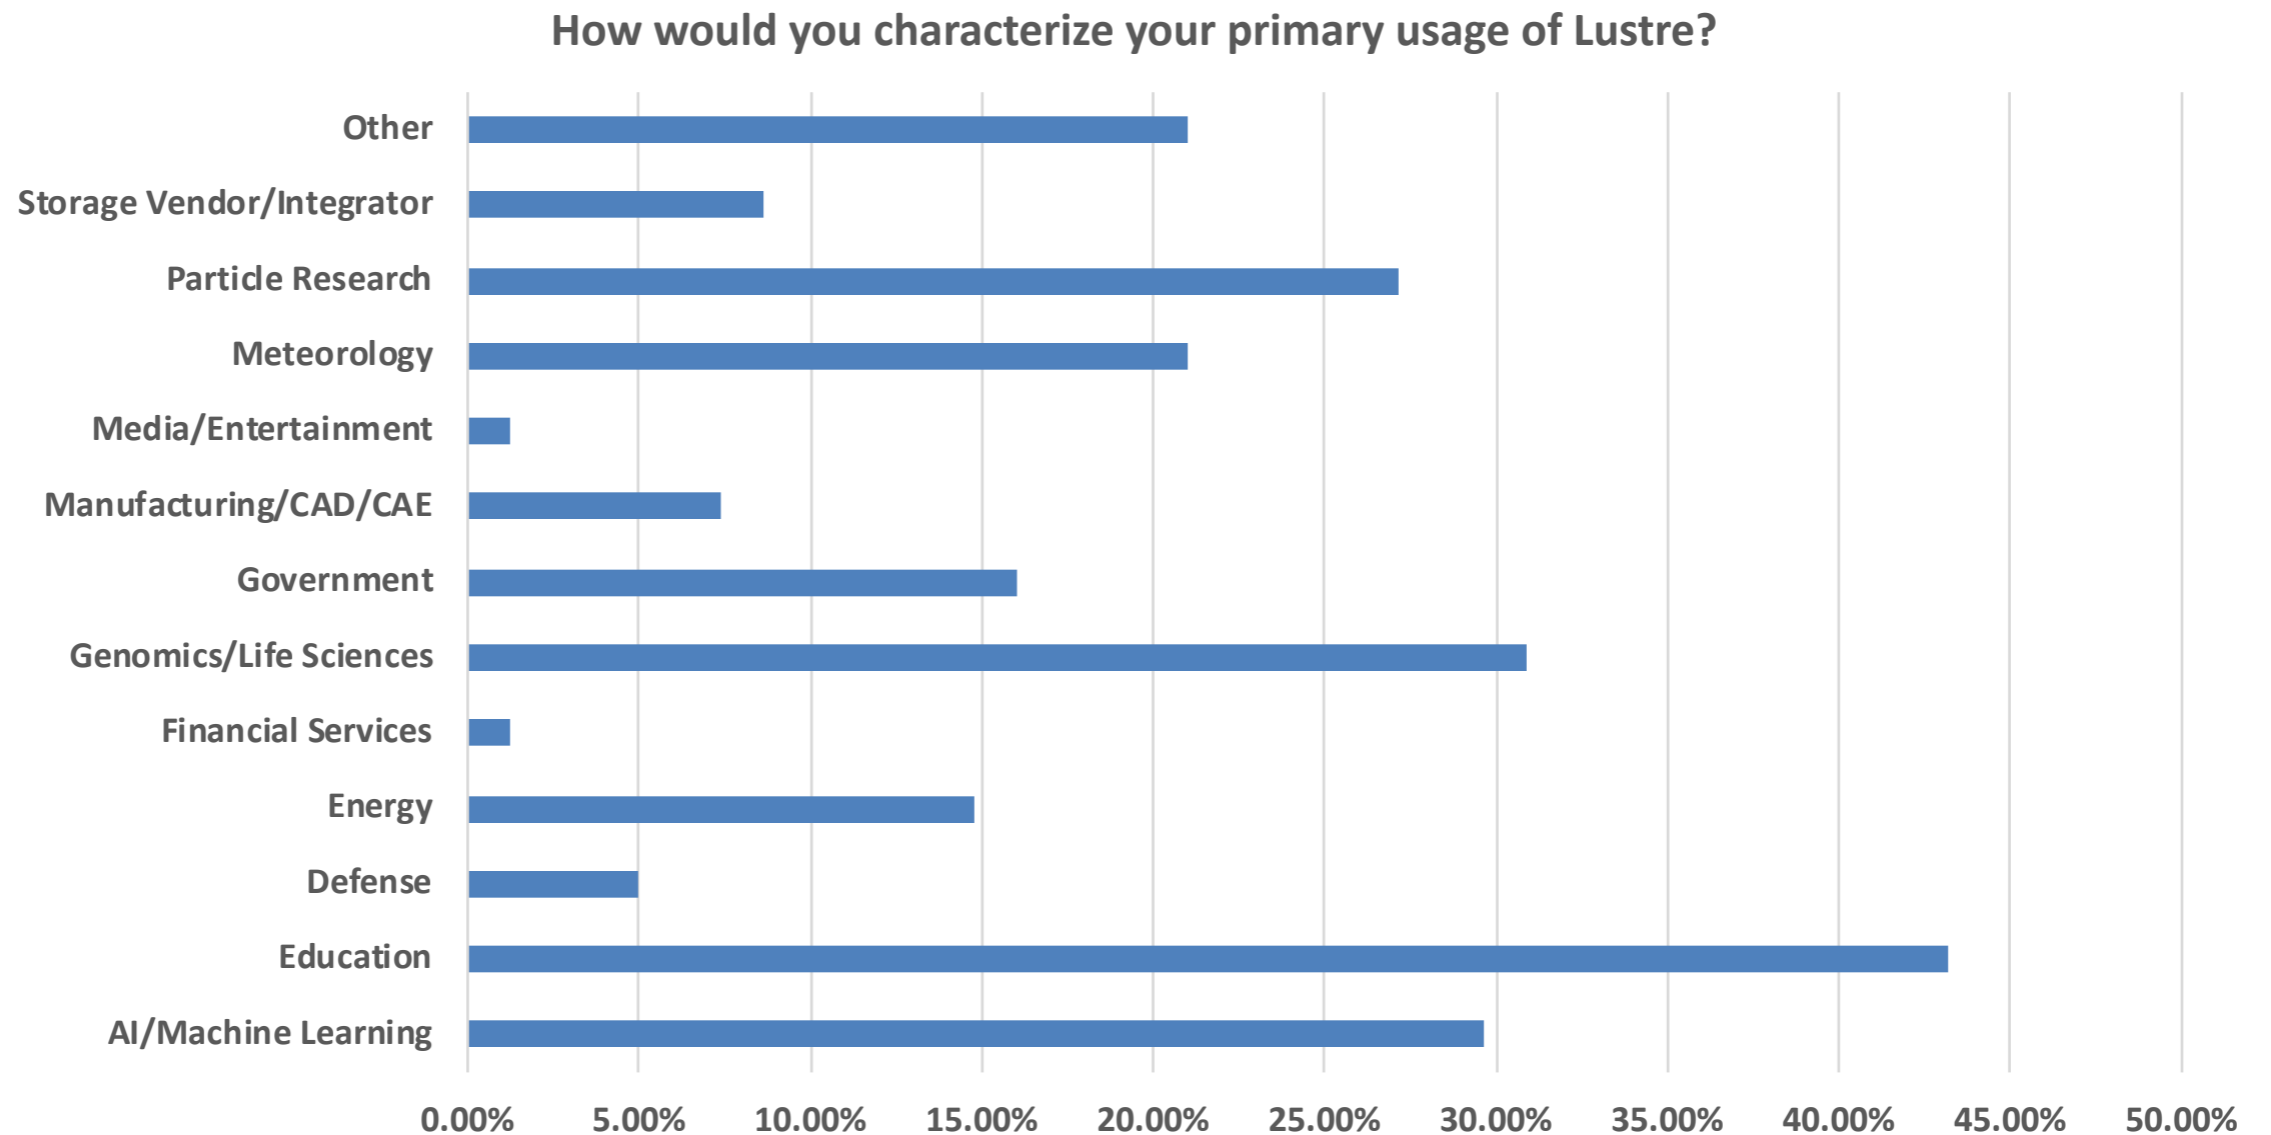
Lustre 2.10.x LTS is the most widely-used production release

Community Survey – Feature Usage



ZFS Production usage continues to grow (~22% 2016; ~36% 2018; 49% 2019)
DNE Striped Directories usage finally starting to get traction
Strong interest in using Project Quotas, PFL and Data on MDT

Community Survey – Primary Usage



- AI/ML, Life Sciences and Particle Research leading categories
- Need to further refine categories based on feedback in Other

Lustre Kernel Policy

- Confusion for a number of years about Linux distro support
 - Questions about omissions from the [Lustre support matrix](#)
 - “Support” is an overloaded term
 - Many more distros than listed work, but these are the ones tested
 - Should developers provide fixes for all kernels in the tree?
- LWG discussed and decided to clarify
 - Developers only need to fix primary distros when making changes
 - Other distros rely on interested parties maintaining them
 - [Lustre/Changelog](#) updated to make clearer distinction
 - Policy on [Lustre.org wiki](#)

Lustre LTS Transition

- LTS change from 2.10.x to 2.12.x announced at SC18
 - Driver was amount of change to support newer kernels
- Some concerns have been raised that two years is not LT 😊
- Large numbers of sites deployed 2.10.x LTS releases
 - Some sites upgraded fully to 2.12.x
 - Some sites using 2.12.x clients
 - Some sites work with vendors who offer longer support versions
 - Some sites using changes ported back to 2.10.x if required
- We will again seek input from the community before making any future changes

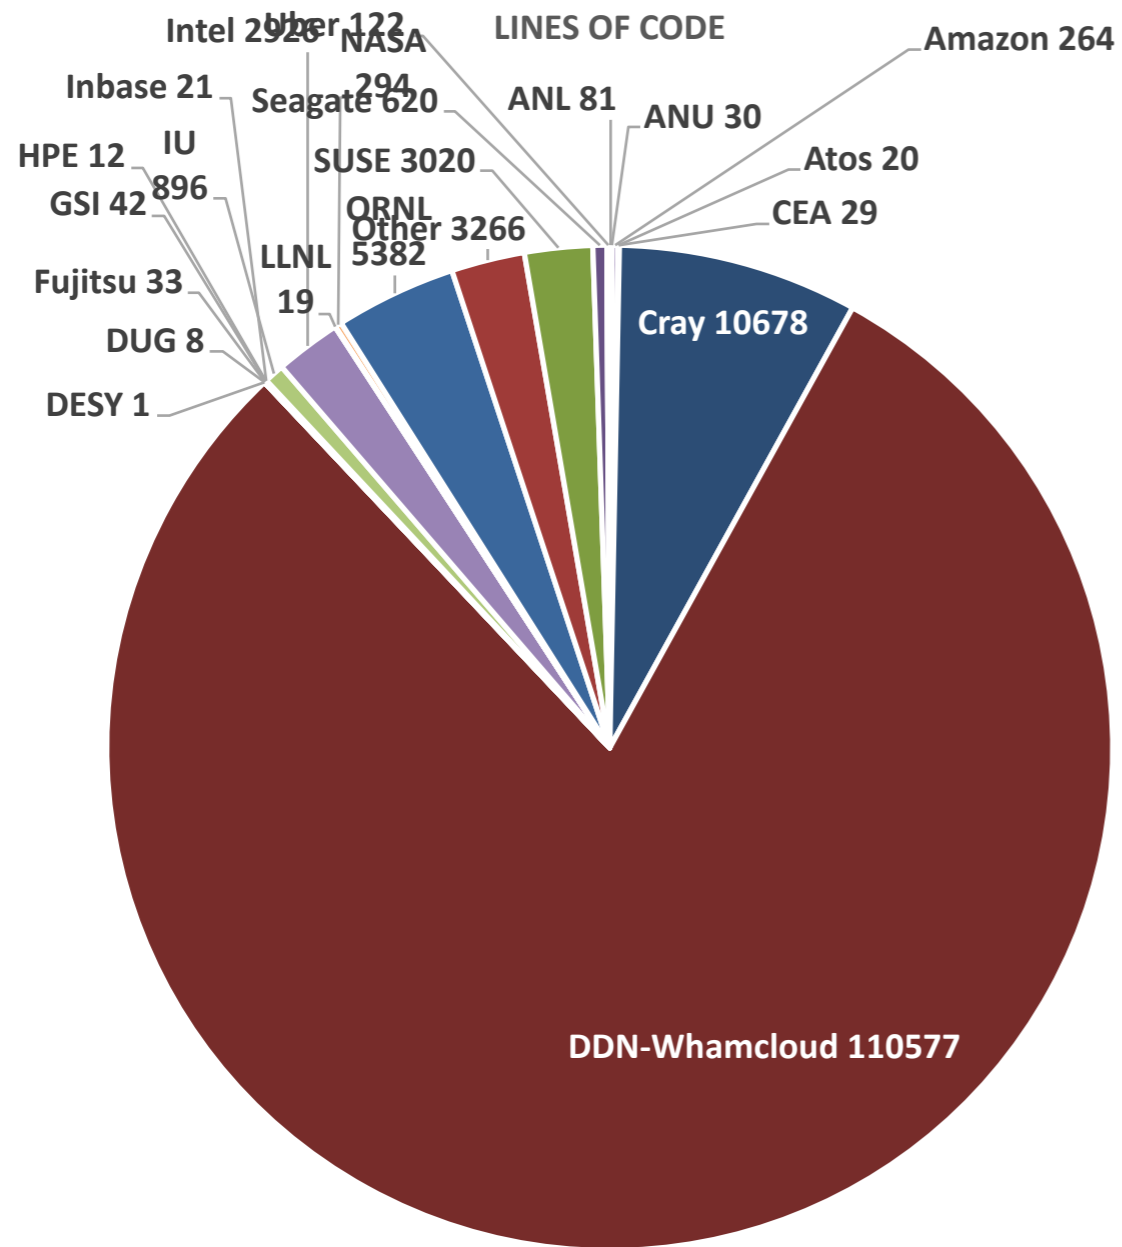
Lustre 2.12.3

- Lustre 2.12.3 targeted for end of month
 - RHEL 7.7 server and client support
 - RHEL 8 client support
 - Bug fixes from early 2.12.x deployments
 - MOFED 4.7
 - ZFS 0.7.13

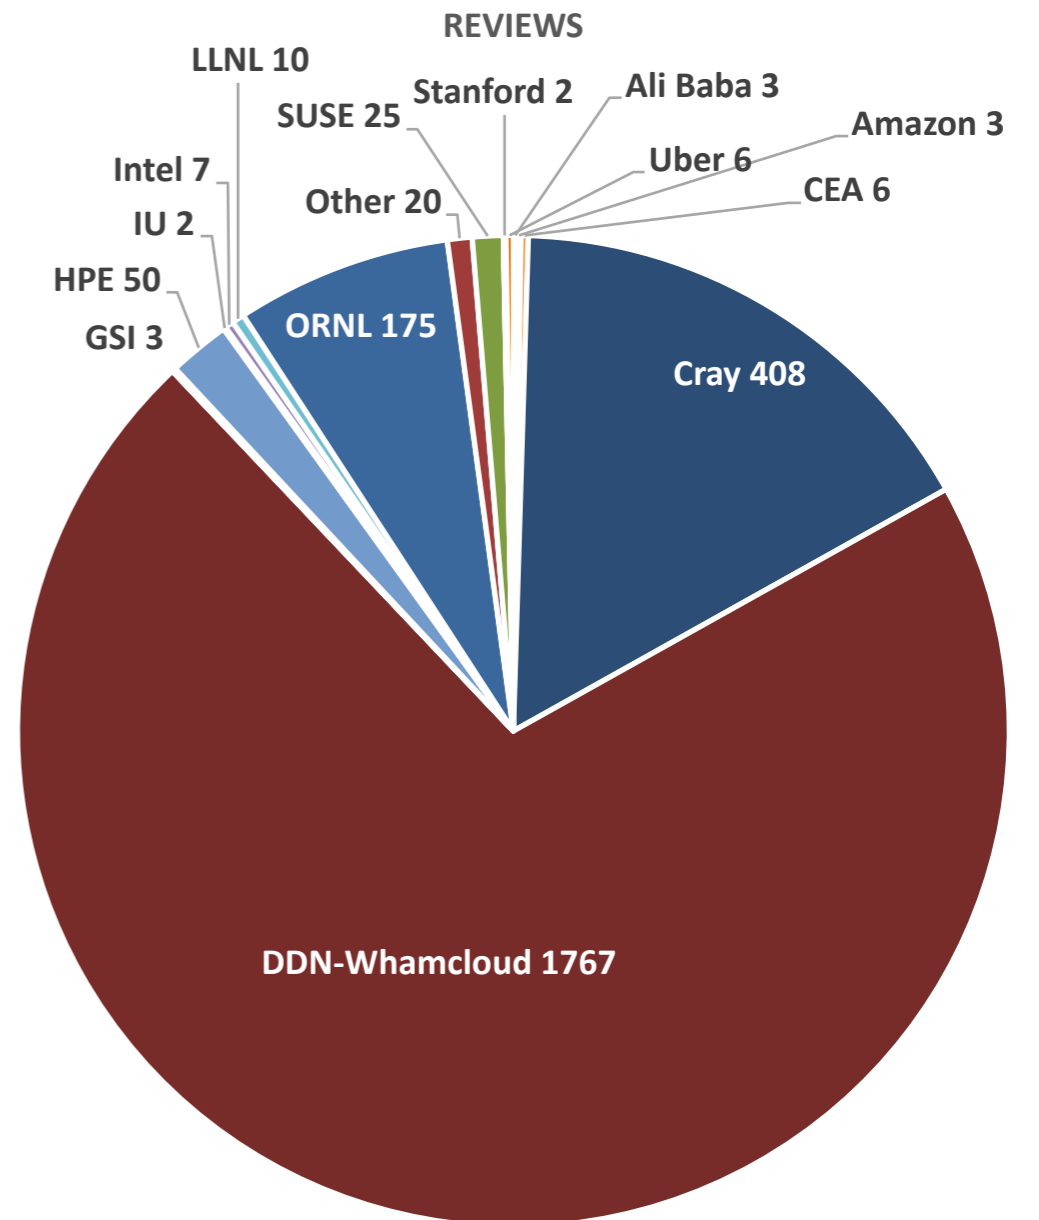
Lustre 2.13

- Trending for Q4 2019 release
- OS support
 - RHEL 7.7 servers/clients
 - RHEL8/SLES12 SP4/Ubuntu 18.04 clients
- Interop/upgrades from latest Lustre 2.12.x
- Ships with ZFS 0.7.13 by default
- Number of useful features
 - Persistent Client Cache (LU-10092)
 - Overstriping (LU-9846)
 - Self Extending Layouts (LU-10070)
- http://wiki.lustre.org/Release_2.13.0

Lustre 2.13 Contributions



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



Aggregated data by organization between 2.12.50 and 2.12.58 tags

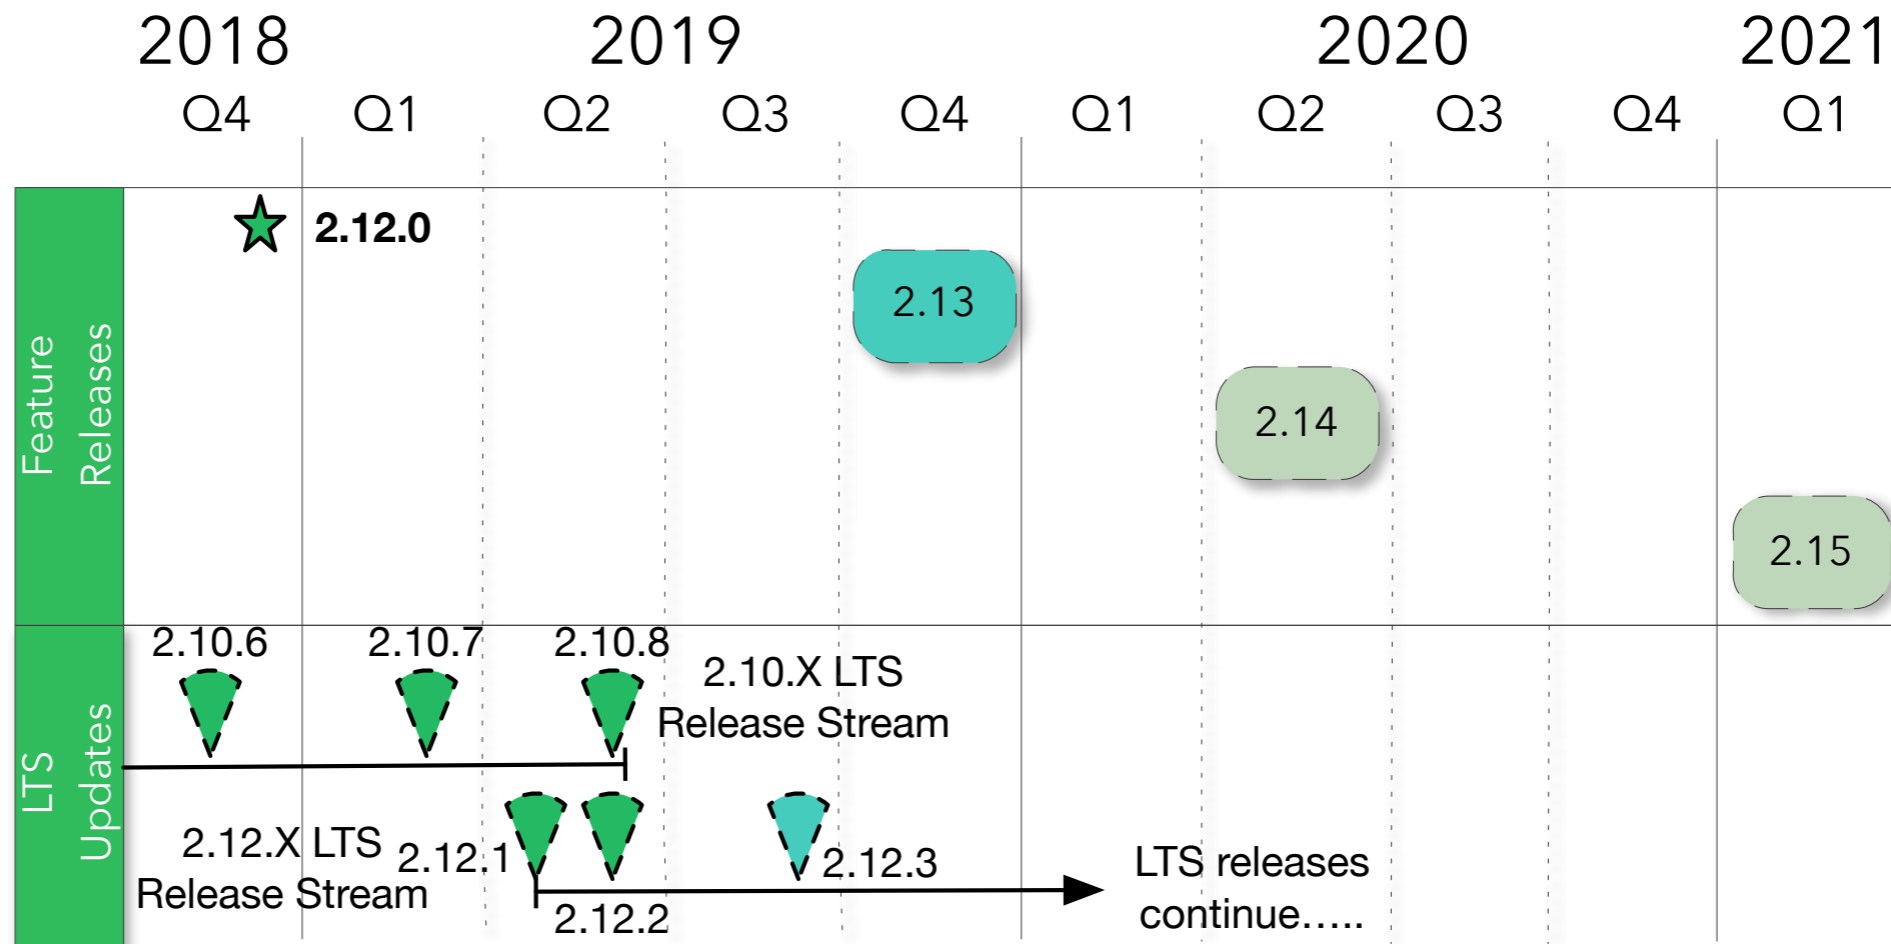
Data courtesy of Dustin Leverman (ORNL)




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
2.11.0	860	134K	87	18
2.12.0	697	82K	90	19
2.13.0 (to 2.12.58)	853	138K	65	24

Lustre Community Roadmap



LEGEND:



2.12

- [Lazy Size on MDT](#)
- [LNet Health](#)
- [DNE Dir Restriping](#)

2.13

- [Persistent Client Cache](#)
- [Multi-Rail Routing](#)
- [Overstriping](#)

2.14

- [FLR Erasure Coding](#)
- [Pool Quota](#)
- [DNE Auto Restriping](#)

2.15

- [Client Encryption](#)
- [Writeback Cache](#)

* 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>

Lustre in Linux Kernel

- SUSE/ORNL working on getting Lustre client into Linux kernel
 - <https://github.com/neilbrown/lustre>
 - Patches pushed for review on lustre-devel mailing list
- Major ldiskfs patches merged into upstream ext4/e2fsprogs
 - Now much easier to keep Lustre e2fsprogs current

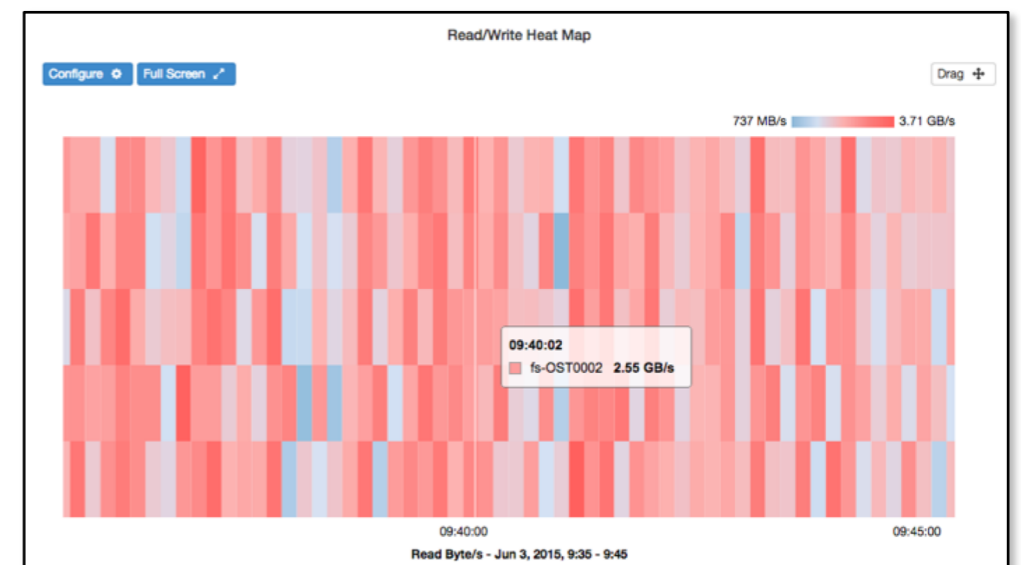
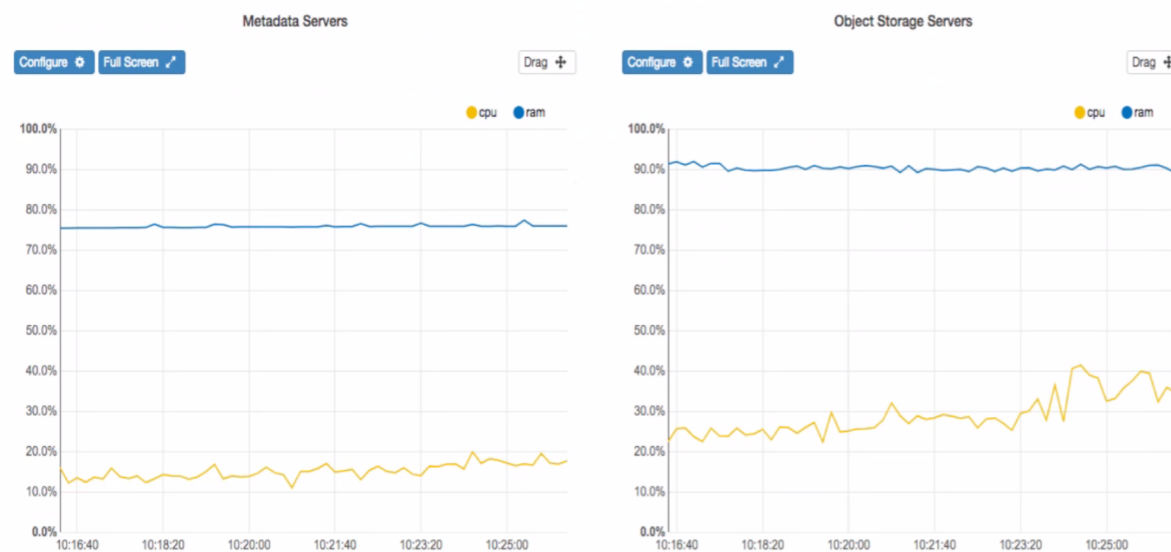
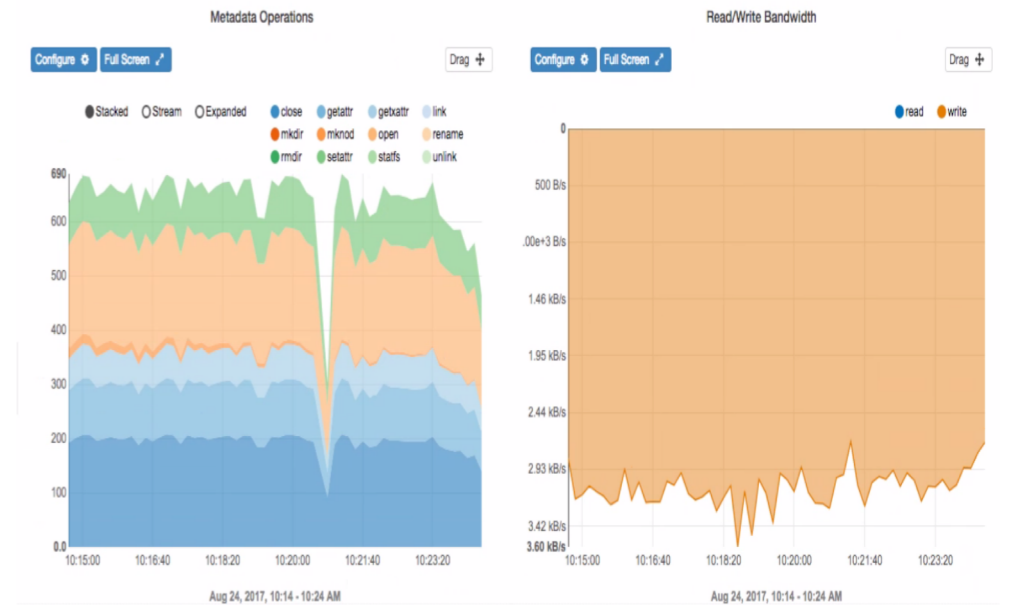
Lustre in the Cloud

- Intel launched AWS offering early 2015
- Cray Azure offering has been available since late 2017
- Amazon announced FSX for Lustre Software in Nov 2018
- DDN-Whamcloud have offerings on GCP, Azure and AWS
- Amazon and Ali Cloud both contributed to Lustre 2.13

Integrated Manager for Lustre

<https://github.com/whamcloud/integrated-manager-for-lustre/releases>

- IML 5.0 GA in May
 - Compatible with Lustre 2.12.x
 - Lower manager/agent resource footprint
- IML 5.1 due out any day
 - <https://github.com/whamcloud/integrated-manager-for-lustre/releases/tag/v5.1.0>



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 exists on lustre.org
 - [http://wiki.lustre.org/Category:Lustre Systems Administration](http://wiki.lustre.org/Category:Lustre_Systems_Administration)
 - Lustre Internals content being refreshed

Lustre's Longevity

- Project active for over 20 years
- Maintaining stability and performance in most demanding operational environments is not an easy problem to solve
- Constant evolution to deal with changing requirements of hardware and usage
- Permissive open source licence has meant that many organizations have been able to collaborate effectively
- For some further Lustre heritage...
 - https://hps.vi4io.org/media/events/2019/hpc-iodc-lustre_next_20_years-dilger.pdf



Summary

- LTS model has been well adopted; focus switching to 2.12.x LTS
- Lustre 2.13 coming soon
- Lustre 2.12.3 coming sooner
- IML 5.1 coming out even sooner 😊
- Plenty of options for those interested in contributing to Lustre
- LWG http://wiki.opensfs.org/Lustre_Working_Group

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