



Community Release Update

May 15th 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)

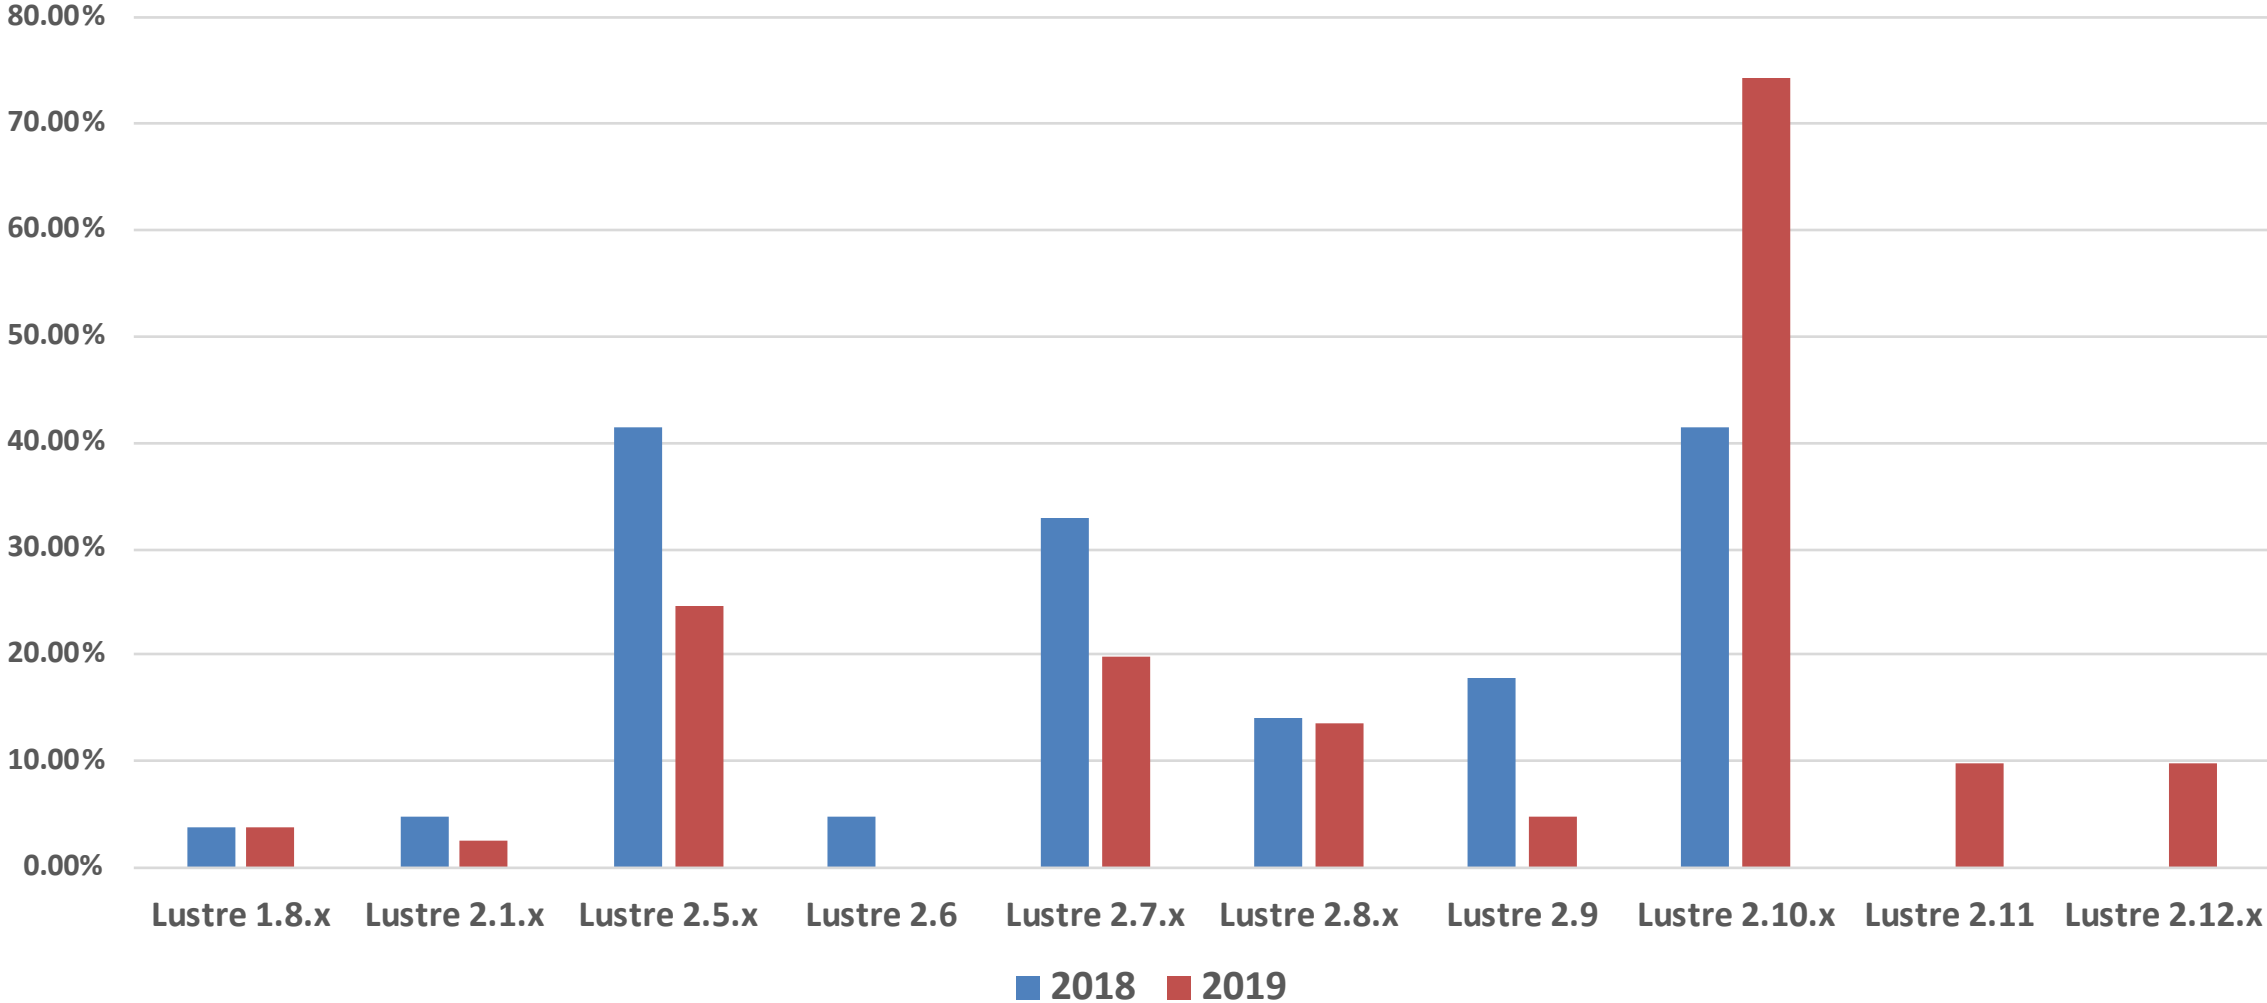
Community Survey

- Has been running for 8 years
- LWG devises questions
- Useful for tracking trends in Lustre usage
- Full details available at [http://wiki.opensfs.org/Lustre Community Survey](http://wiki.opensfs.org/Lustre_Community_Survey)



Lustre Community Survey


Which Lustre versions do you use in production? (select all that apply)



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

Community Survey - Linux distros


What Linux distributions do you use in production?



RHEL/CentOS 7.x remains most widely-used Linux distribution

Community Survey – Processor Architecture


Which processor architecture(s) do you use in production?



X86 is the most common processor architecture

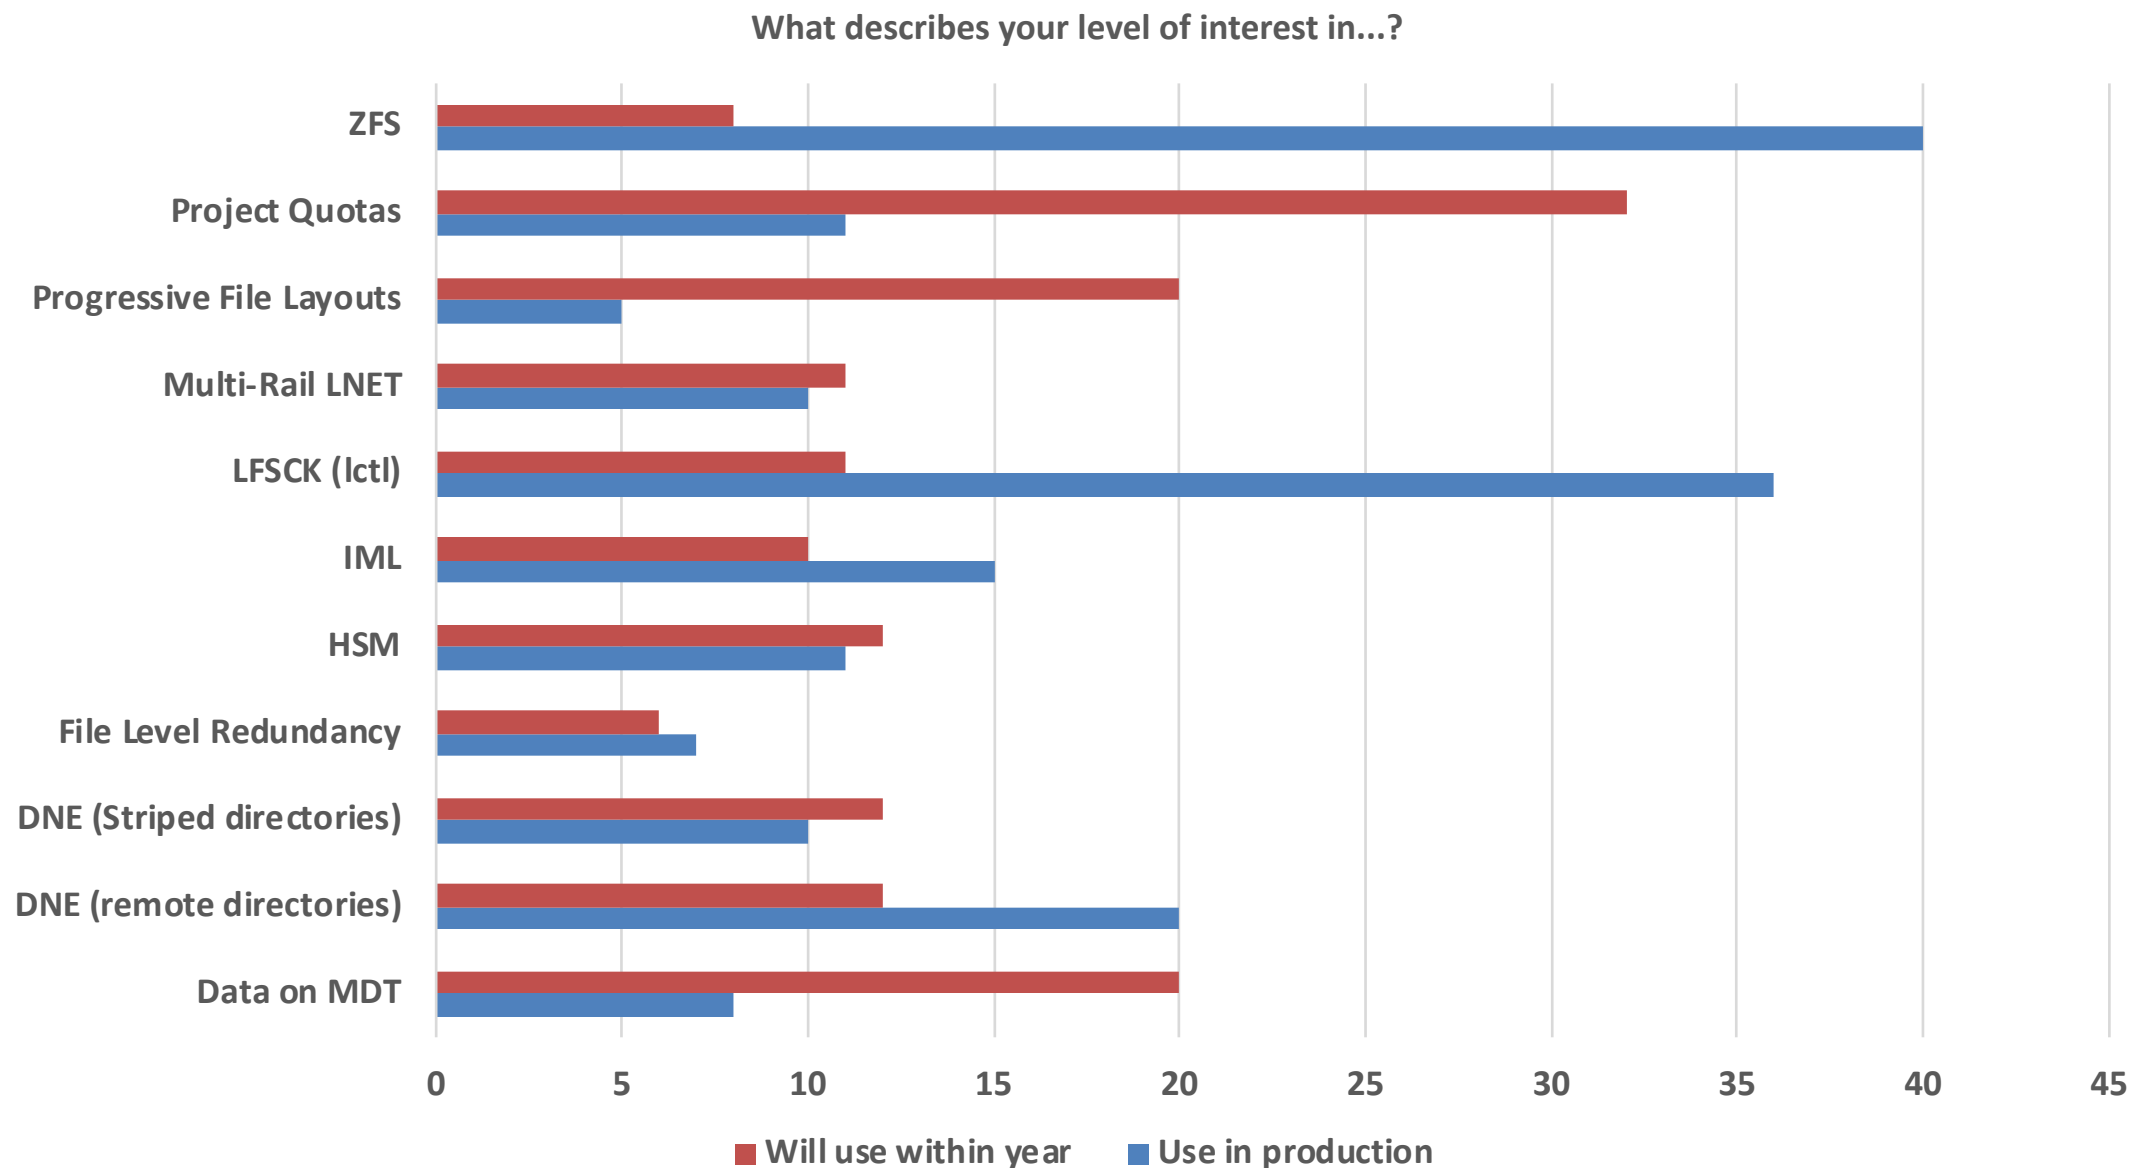
Community Survey - Networks

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




10 GigE most common networking technology used

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

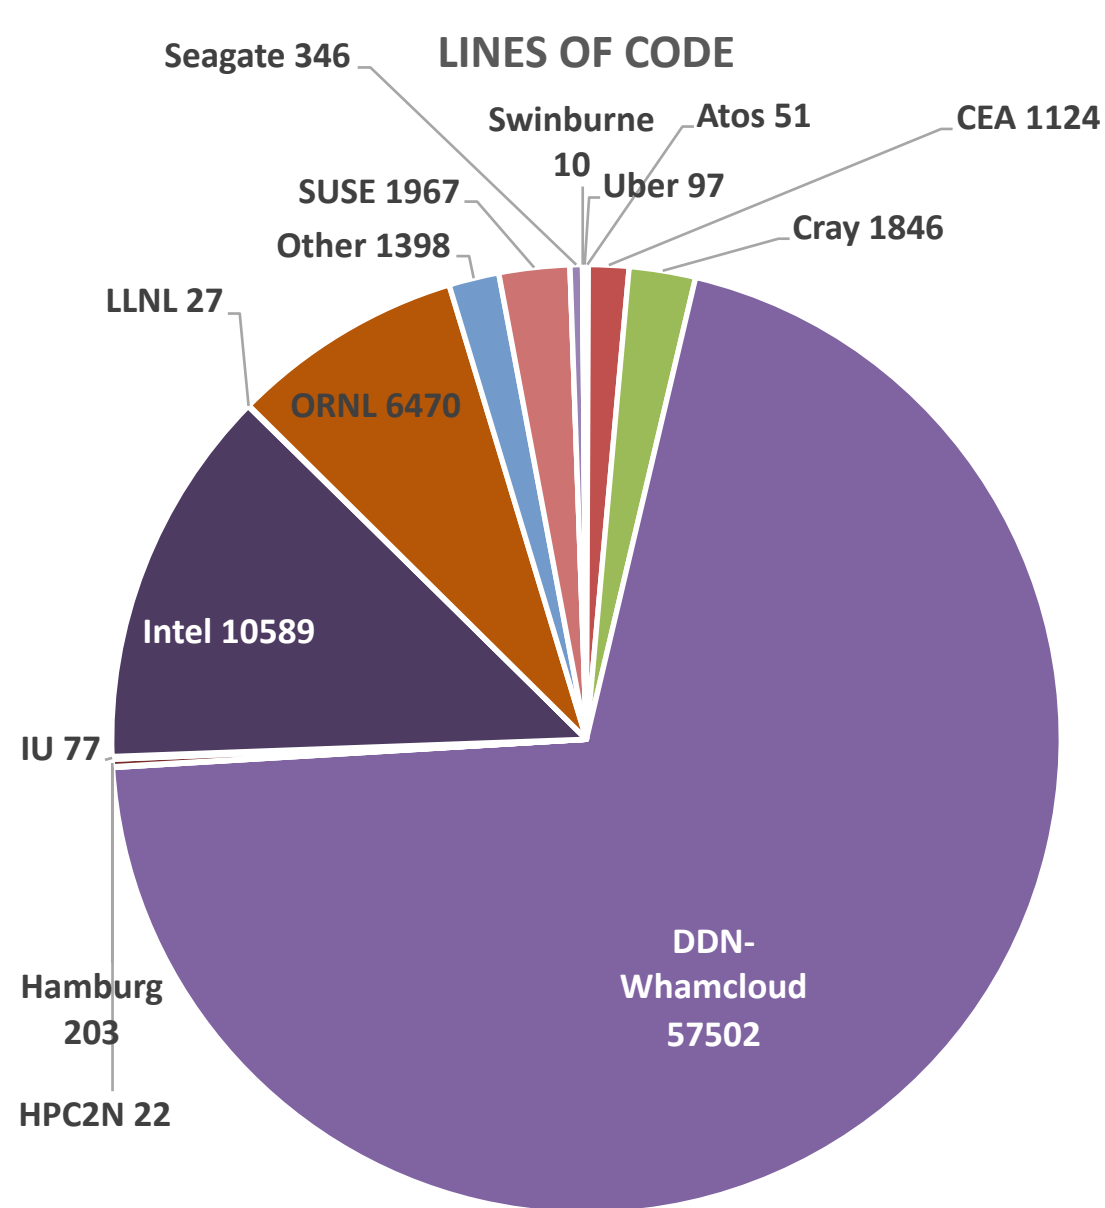
Lustre LTS Releases

- Lustre 2.10.0 went GA July 2017
 - Initial LTS branch with regular updates
 - Lustre 2.10.7 went GA March 25th
- Lustre 2.12.0 went GA Dec 2018
 - Current LTS branch
 - Provides reliable option for newer kernels/features
 - Lustre 2.12.1 went GA April 30th 2019
 - Lustre 2.12.2 coming soon

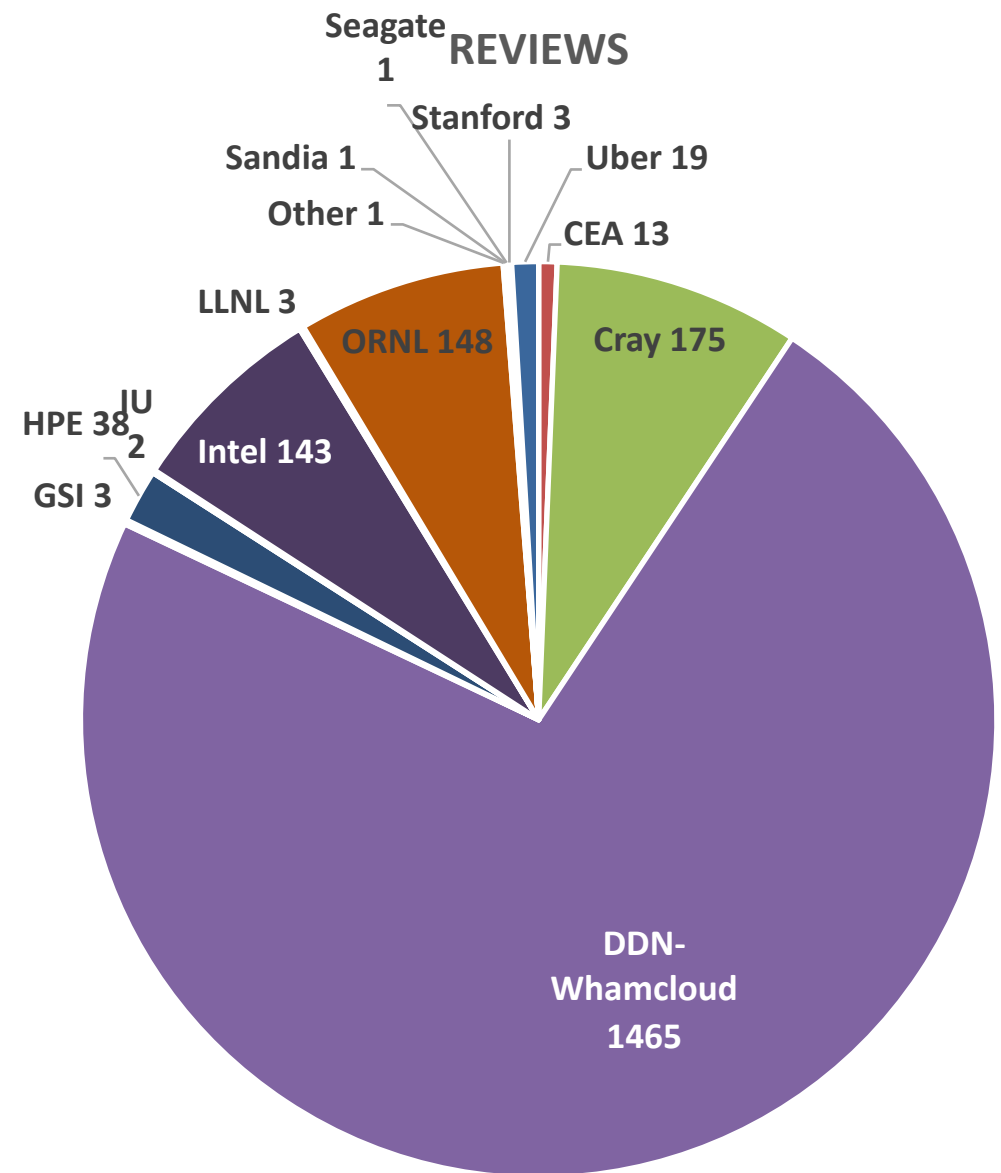
Lustre 2.12

- GA Dec 2018
- OS support
 - RHEL 7.6 servers/clients
 - SLES12 SP3/Ubuntu 18.04 clients
- Interop/upgrades from latest Lustre 2.10.x and 2.11
- Number of useful features
 - DNE Directory Restriping (LU-4684)
 - LNet Network Health (LU-9120)
 - Lazy Size on MDT (LU-9538)
 - T10PI end to end data checksums (LU-10472)
- http://wiki.lustre.org/Release_2.12.0

Lustre 2.12 Contributions



Data courtesy of Dustin Leverman (ORNL)



Analysis of commits on git.whamcloud.com v2.11.50 to 2.12.0

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

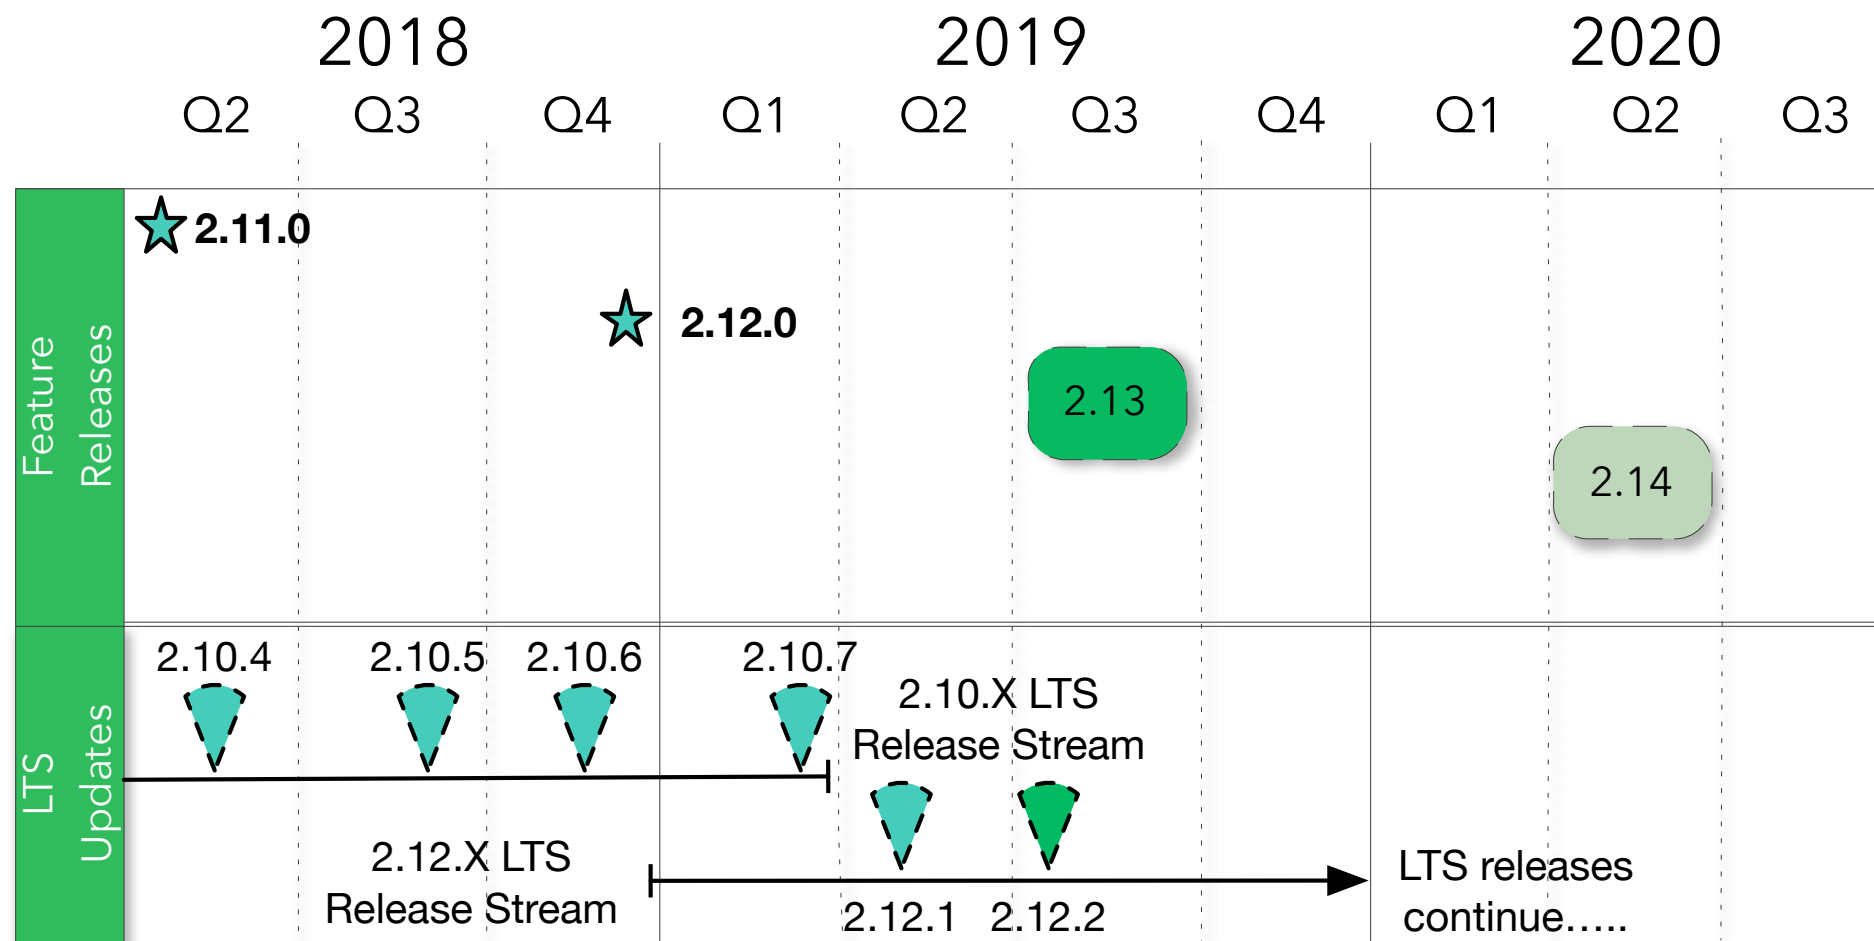
Statistics courtesy of Chris Morrone (LLNL)/ Dustin Leverman (ORNL)

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


Lustre 2.13

- Targeted for Q3 2019 release
- OS support
 - RHEL 7.6 servers/clients
 - SLES12 SP4/Ubuntu 18.04 clients
- Interop/upgrades from latest Lustre 2.12.x
- Ships with latest ZFS 0.7.x but ZFS 0.8.x ready
- Number of useful features targeted; feature freeze May 31st
 - Persistent Client Cache (LU-10092) **IN PROGRESS**
 - User Defined Selection Policy (LU-9121) **IN PROGRESS**
 - Self Extending Layouts (LU-10070) **IN PROGRESS**
- http://wiki.lustre.org/Release_2.13.0

Lustre Community Roadmap




LEGEND:



Expected Timeline



Timeline TBD



Completed

↓ LTS Branch

2.11

- [Data on MDT](#)
- [FLR Delayed Resync](#)
- [Lock Ahead](#)

2.12

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

2.13

- [Persistent Client Cache](#)
- [LNet Selection Policy](#)
- [Self Extending Layouts](#)

2.14

- [FLR Erasure Coding](#)
- [Health Monitoring](#)
- [DNE Auto Restriping](#)

* 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

- Added to kernel staging area in 3.11
 - Basic code reformatting had not been done beforehand
 - Kernel staging rules made subsequent updates challenging
- Removed from staging area in 4.17
- Work is continuing at <https://github.com/neilbrown/lustre>
 - SUSE and ORNL driving this initiative
 - Plan to resubmit when code is ready for acceptance
- Major ldiskfs patches merged into upstream ext4/e2fsprogs
 - Now much easier to keep Lustre e2fsprogs current


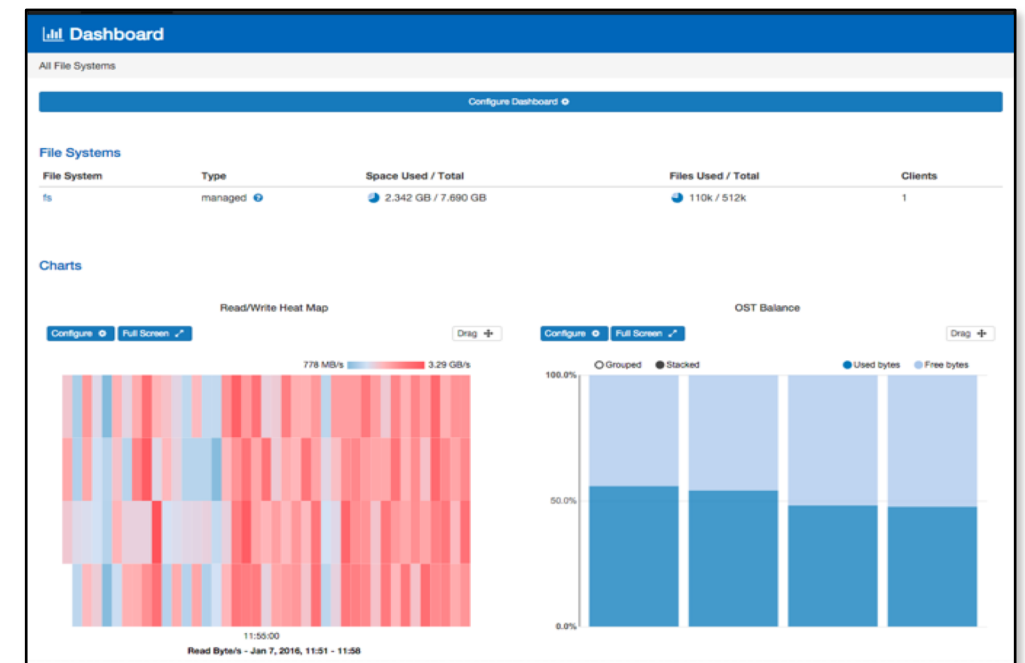
Lustre in the Cloud

- Intel AWS offering transferred to DDN-Whamcloud
 - <https://aws.amazon.com/marketplace/pp/B07FTX9GBV>
 - Used in production environments since 2015
- Amazon announced FSX for Lustre Software in Nov 2018
 - <https://aws.amazon.com/fsx/lustre/>
- Google also interested in Lustre
 - Gave several talks about Lustre at SC18 and entered IO-500
 - Whamcloud GCP offering launched April 2019
 - <https://console.cloud.google.com/marketplace/details/ddnstorage/cloud-edition-for-lustre>
- Multiple offerings on Azure
 - Cray offering has been available since late 2017
 - <https://azure.microsoft.com/en-ca/solutions/high-performance-computing/cray/>
 - Whamcloud offering launched May 2019
 - https://azuremarketplace.microsoft.com/en-us/marketplace/apps/ddn-whamcloud-5345716.lustre_on_azure?tab=Overview&pub_source=email&pub_status=success

Integrated Manager for Lustre

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

- IML 4.0 GA Oct 2017
 - Compatible with Lustre 2.10.x
 - latest maintenance release 4.0.10.1
 - Distributed under an MIT license
 - Possible to upgrade from Intel EE 2.x/3.x
- IML 5.0 coming soon
 - Compatible with Lustre 2.12.x
 - Extending scale of clusters it can manage
 - More flexibility around HA schemes
 - Reduced resource utilization
 - Will require downtime to upgrade



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

Summary

- LTS model has been well adopted; focus switching to 2.12.x LTS
- Lustre 2.13 feature freeze at end of month
- Lustre 2.12.2 coming soon
- Plenty of options for those interested in contributing to Lustre
- LWG [http://wiki.opensfs.org/Lustre Working Group](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