



State of Lustre: Annual Report
Sponsored by OpenSFS

Context

- Lustre Annual Report covers Lustre market trends
- For Lustre technical trends, see

Lustre Community Survey (2012 – 2014)

http://wiki.opensfs.org/Lustre_Community_Survey

I/O Characterization of Large-Scale HPC Centers —
Benchmarking Working Group

<http://lustre.opensfs.org/resources/>

Overview

- Storage growth trends are driving increased opportunity for parallel file systems
- Lustre is the most widely adopted parallel file system in HPC
- However, Lustre is primarily adopted in academic and government research markets, not commercial
- Big Data creates an opportunity for Lustre to expand in commercial markets
- Download full report, free registration required:

<http://www.opensfs.org/2014-i360/>

Big Data Opportunity: Survey

- 269 total respondents
 - 177 “Technical” (HPCwire and HPC500 user group)
 - 92 “Enterprise” (Gabriel Consulting)
 - 150 commercial, 70 academic, 49 government
- Surveys completed summer 2013 – early 2014
- Builds on original survey from early 2012 (306 respondents: 204 Technical, 102 Business)
- End users discuss their environments, challenges, solutions, and “satisfaction gaps” in addressing Big Data challenges

Big Data Challenges

How significant are each of the following barriers, bottlenecks, or constraints? (1 = not a limitation; 5 = severe limitation)	Total	
	Avg.	% 4/5
Compute cycles	2.9	34%
Memory bandwidth	3.1	38%
I/O bandwidth	3.5	55%
I/O latency	3.5	52%
I/O throughput (IOPS)	3.5	51%
Data ingest (receiving or reading data)	3.3	51%
Data output (writing data)	3.2	44%
Storage capacity	3.3	47%
Number of respondents	265 - 267	

Intersect360 Research, 2013-2014

I/O performance (bandwidth, latency, throughput)
rated as top Big Data challenge


Big Data “Satisfaction Gaps”

I/O performance was the #1 “satisfaction gap” for storage, measuring high importance relative to low satisfaction with available technologies

Satisfaction gaps: (Imp - Sat) x 100	Sat Gap
Flash/SSD technology	-3
Disk drive technology	10
Tape technology	-45
InfiniBand	-23
Ethernet	18
Parallel file systems	33
Storage capacity	64
I/O performance	90
Backup and recovery	43
Deduplication	-2
Density	7
Power consumption	36
Reliability, availability, serviceability (RAS)	54
Number of respondents	245 - 255

Intersect360 Research, 2013-2014

Current HPC Adoption of Lustre



Progress, But Still Work to Do

- Lustre is the #1 parallel file system in the market (HPC User Site Census survey data)
- Much lower adoption in commercial than in academia and government
- Big Data presents an opportunity to expand

“Storage, software, and services are all high-growth segments, and Lustre resides at the intersection of these opportunities.”

Intersect360 Research Analysis

- Support and active participation from multiple vendors is critical in continuing to broaden Lustre adoption.
- The roles of OpenSFS and EOFS will be critical. As the leading consortia charged with the ongoing stewardship of Lustre, it will fall to OpenSFS and EOFS to steer the course of development.
- Dual challenge of attracting new organizations into its membership while continuing to represent the interest of the incumbents.

State of Lustre: Annual Report


<http://www.opensfs.org/2014-i360/>

Thanks to OpenSFS Participants

Promoters



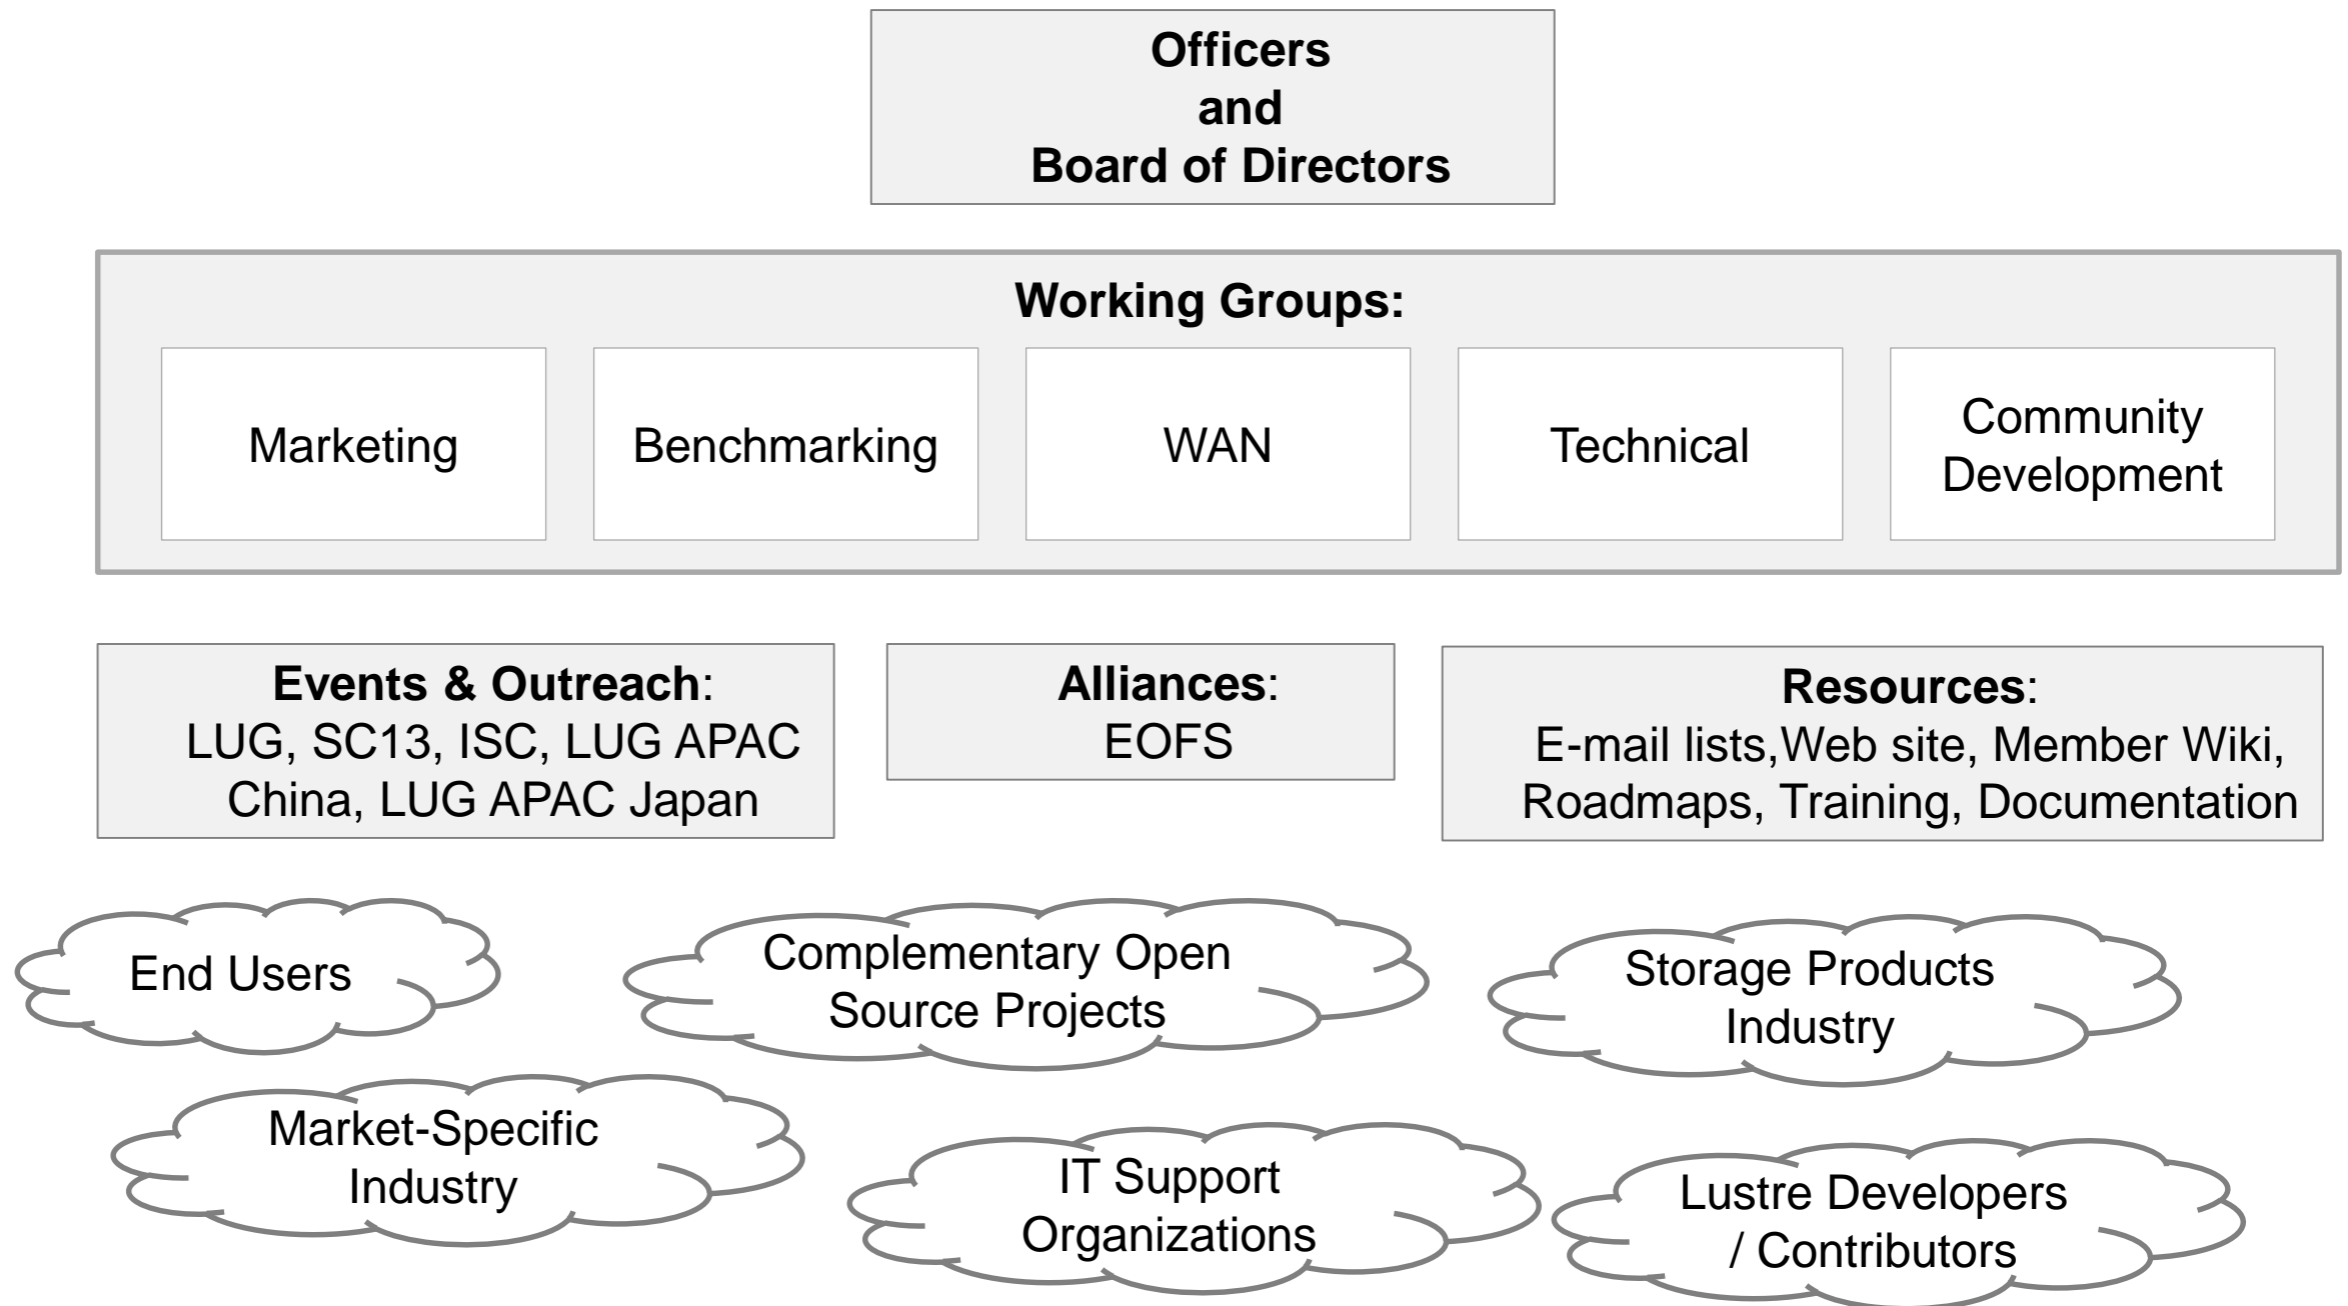
Adopters



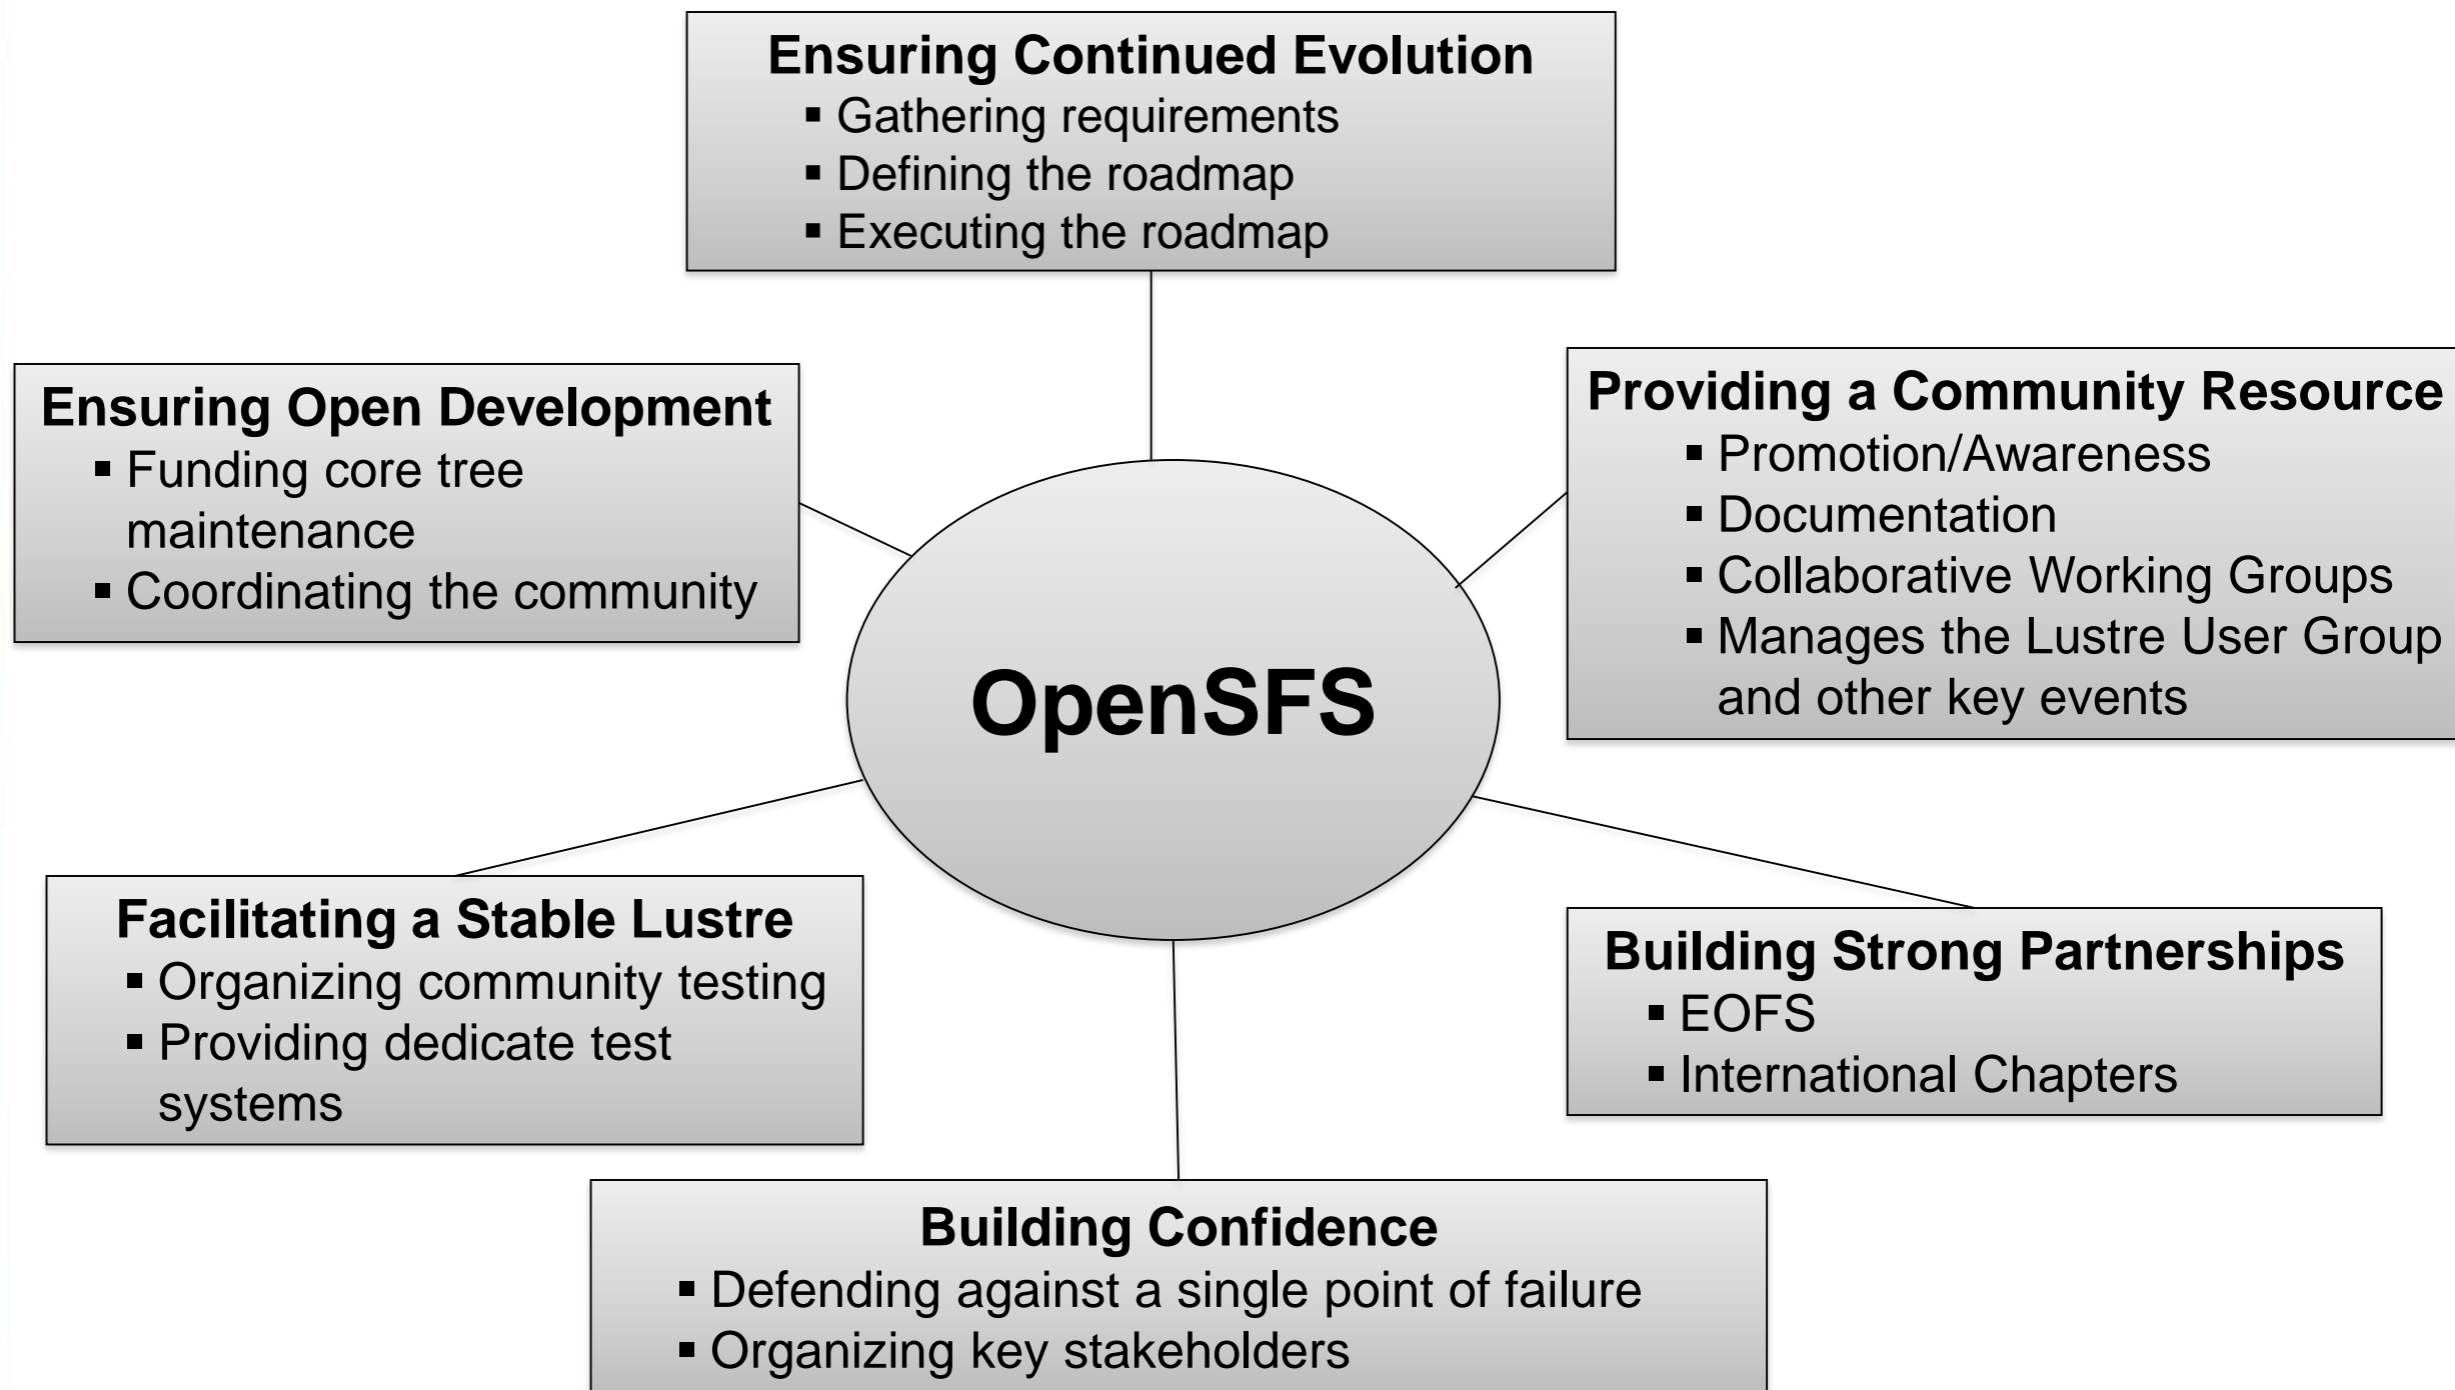
Supporters



How We Are Organized and Who We Serve



Providing Community Structure



2013 Was A Big Year

- Major releases of Lustre 2.4 and 2.5
- Major new Lustre features
- Two new Participants, Fujitsu and NCSA
- Opened Community Lustre Portal at lustre.opensfs.org
- Continued Partnership with EOFS
- First Annual APAC LUG13 in China and Japan
- Largest LUG 2013 to date
- Lustre BoFs
- Enhancements to communications
- More mature management infrastructure

Brought To You By OpenSFS

Feature	Impact	Release
Server Stack SMP Scaling	Performance	2.3
Online check/scrub	Reliability	2.3
Distributed NamespacE (DNE 1)	Performance/Scalability	2.4
LFCK MDT FID/LinkEA	Usability/Reliability	2.4
FSCK MDT-OST Consistency	Usability/Reliability	2.5
DNE Phase 2	Performance/Scalability	2.7
LFCK MDT-MDT Consistency	Usability/Reliability	2.6
UID Mapping & Shared Key	Usability	2.6

New Initiative for 2014



- Fostering broader participation in Lustre community releases
 - Gatekeeping, release management, development, testing
- Laying the groundwork to transition to a volunteer community model
- Leveraging our strong vendor commitment and partnership with EOFS members to make this happen

Fostering a strong and diverse ecosystem
of open source development

Get Involved!

- Encourage your organization to join OpenSFS
- Actively participate in an OpenSFS working group
- Contribute to the Lustre file system (contribute documentation; testing, testing, and more testing; bug fixes, features)

Help shape the future of Lustre!