A Community Resource

Lustre Users Group 2012

Austin Texas

4/23/2012

Prologue to OpenSFS

2010: The Lustre community moves to support Lustre on its own.

- Two companies – Xyratex & Whamcloud – hire Lustre developers and offer Level 3 support to organizations.
- Whamcloud steps up to support the Lustre tree
- Three community groups form to offer venues for collaboration and support for end-users.

2011: The Community consolidates to align efforts to create a unified, global effort to support Lustre.

- OpenSFS gathers community requirements, solicits proposals, then funds a 2-year development contract to improve Lustre metadata performance and scalability.

What a Difference a Year Makes!

- Since LUG 2011 OpenSFS has established itself as the primary venue for collaborative Lustre maintenance and development
- As a result, the future of Lustre is bright!
  - As of yesterday we have 215 attendees here at LUG
  - Over 70 different organizations
  - Participants from every corner of the world

What is OpenSFS?

OpenSFS is a vendor neutral, member supported non-profit organization bringing together the open source file system community for the high performance computing sector.

Our mission is to aggregate community resources and be the center of collaborative activities to ensure efficient coordination of technology advancement, development and education.

The end goal is the continued evolution of robust open source file systems for the HPC community.

OpenSFS Provides a Community Structure for Open Source File Systems

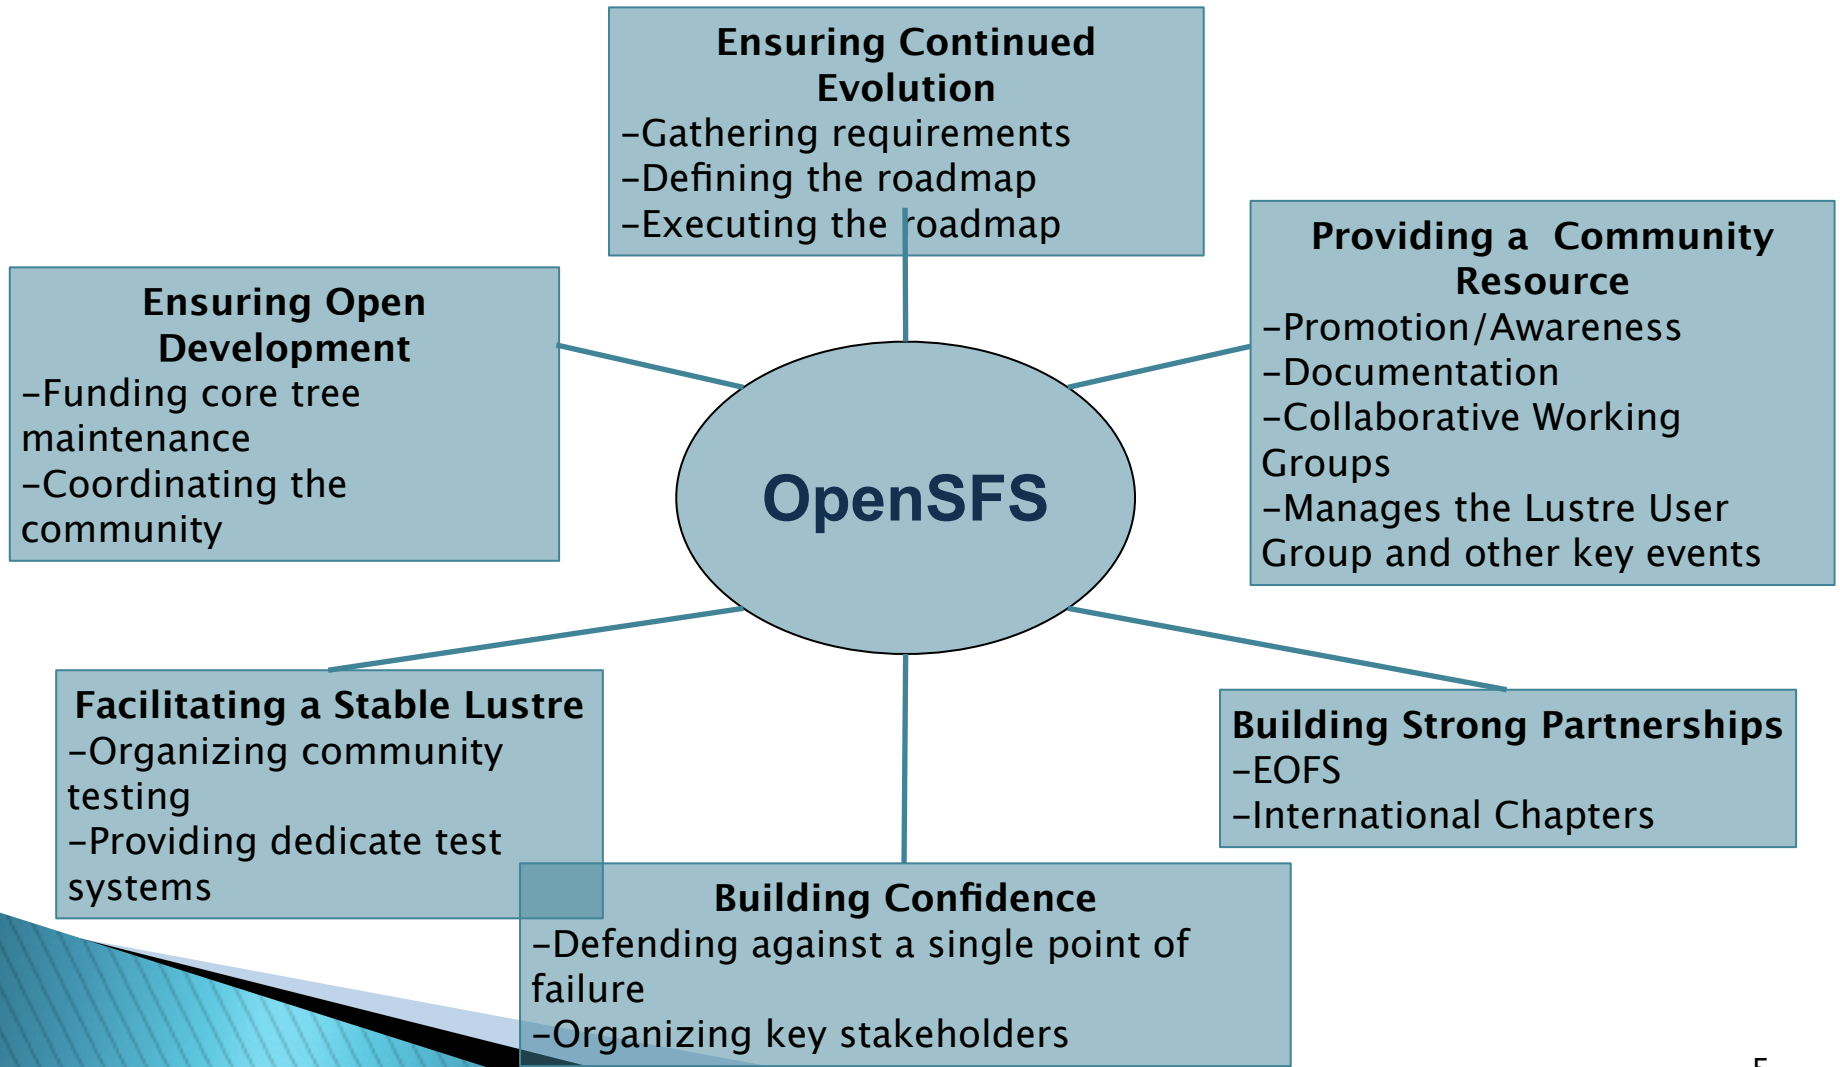

Who Are We?

Current Members

Pending Members:

A Community Working Group Model

- Technical Working Group
  - Gather requirements from community
  - Propose and manage development projects
  - Generate Lustre feature roadmap
- Community Development Working Group
  - Manage Lustre releases
  - Coordinate release roadmap for new features
- Benchmarking Working Group
  - Investigate HPC user workloads
  - Define tests to evaluate system performance and scalability
- WAN Working Group
  - Coordinate use cases and features for wide area Lustre

OpenSFS ensures open source file

OpenSFS	=	Prevents single point of failure for open source file systems	=	Understands and consolidates community issues, efforts, projects & future direction
Collaborative open source file system community				

Outcomes

- Stabilization, support and maturation of Lustre as open source.
- New technology development projects.
- Shared support and development costs.
- Community engagement.
- Promote awareness and expand technical knowledge base including subject matter experts supporting HPC open source file systems.

Three Components for the Continued Maturation of Open Source File

A. Support and maintenance of Lustre
B. Development programs to meet future requirements
C. Community support related activities

There are real costs to each of these components.

The HPC community must share these costs if there is to be production capable open source file systems

There are Real Costs to Open

A. Annual cost of support and maintenance of Lustre

- Regular releases for new features and maintenance
  - Gatekeeper, current and next release (1-2 FTE)
  - Release Testing (2 FTE)
  - Release management (1 FTE)
  - Bug fixing (4 FTE)
  - Admin support (0.5 FTE)
  - Documentation (0.25 FTE)
  - Dedicated testing equipment (~$150k/yr)

- Total annual cost: $3.0M
  - OpenSFS funds via support contract: $1.5M/yr
  - Whamcloud currently contributes an additional:

There are Real Costs to Open

B. Cost of development Programs

Currently Funded

- Whamcloud $2.3M total (2011– 2013)
  - Single Server Metadata Performance Improvements
  - Distributed Namespace: Remote Directories and Striped Directories
  - Lustre File System Checker
- Indiana University $250k total (2012 – 2013)
  - Lustre Wide Area Network
- Open test system for Lustre development ~$275K + Operating costs (ongoing)
  - To be co-located at Indiana University

There are Real Costs to Open

C. Community support related costs

- Community outreach (website, wiki, media, etc.)
- Training (content development, developer tools, tutorials, resource guidance, etc.)
- Documentation (content development, archiving, etc.)
- Community events (users events, related conference presence, etc.)
- Business related (legal, etc.)
- These costs can be substantial

Community Lustre Roadmap

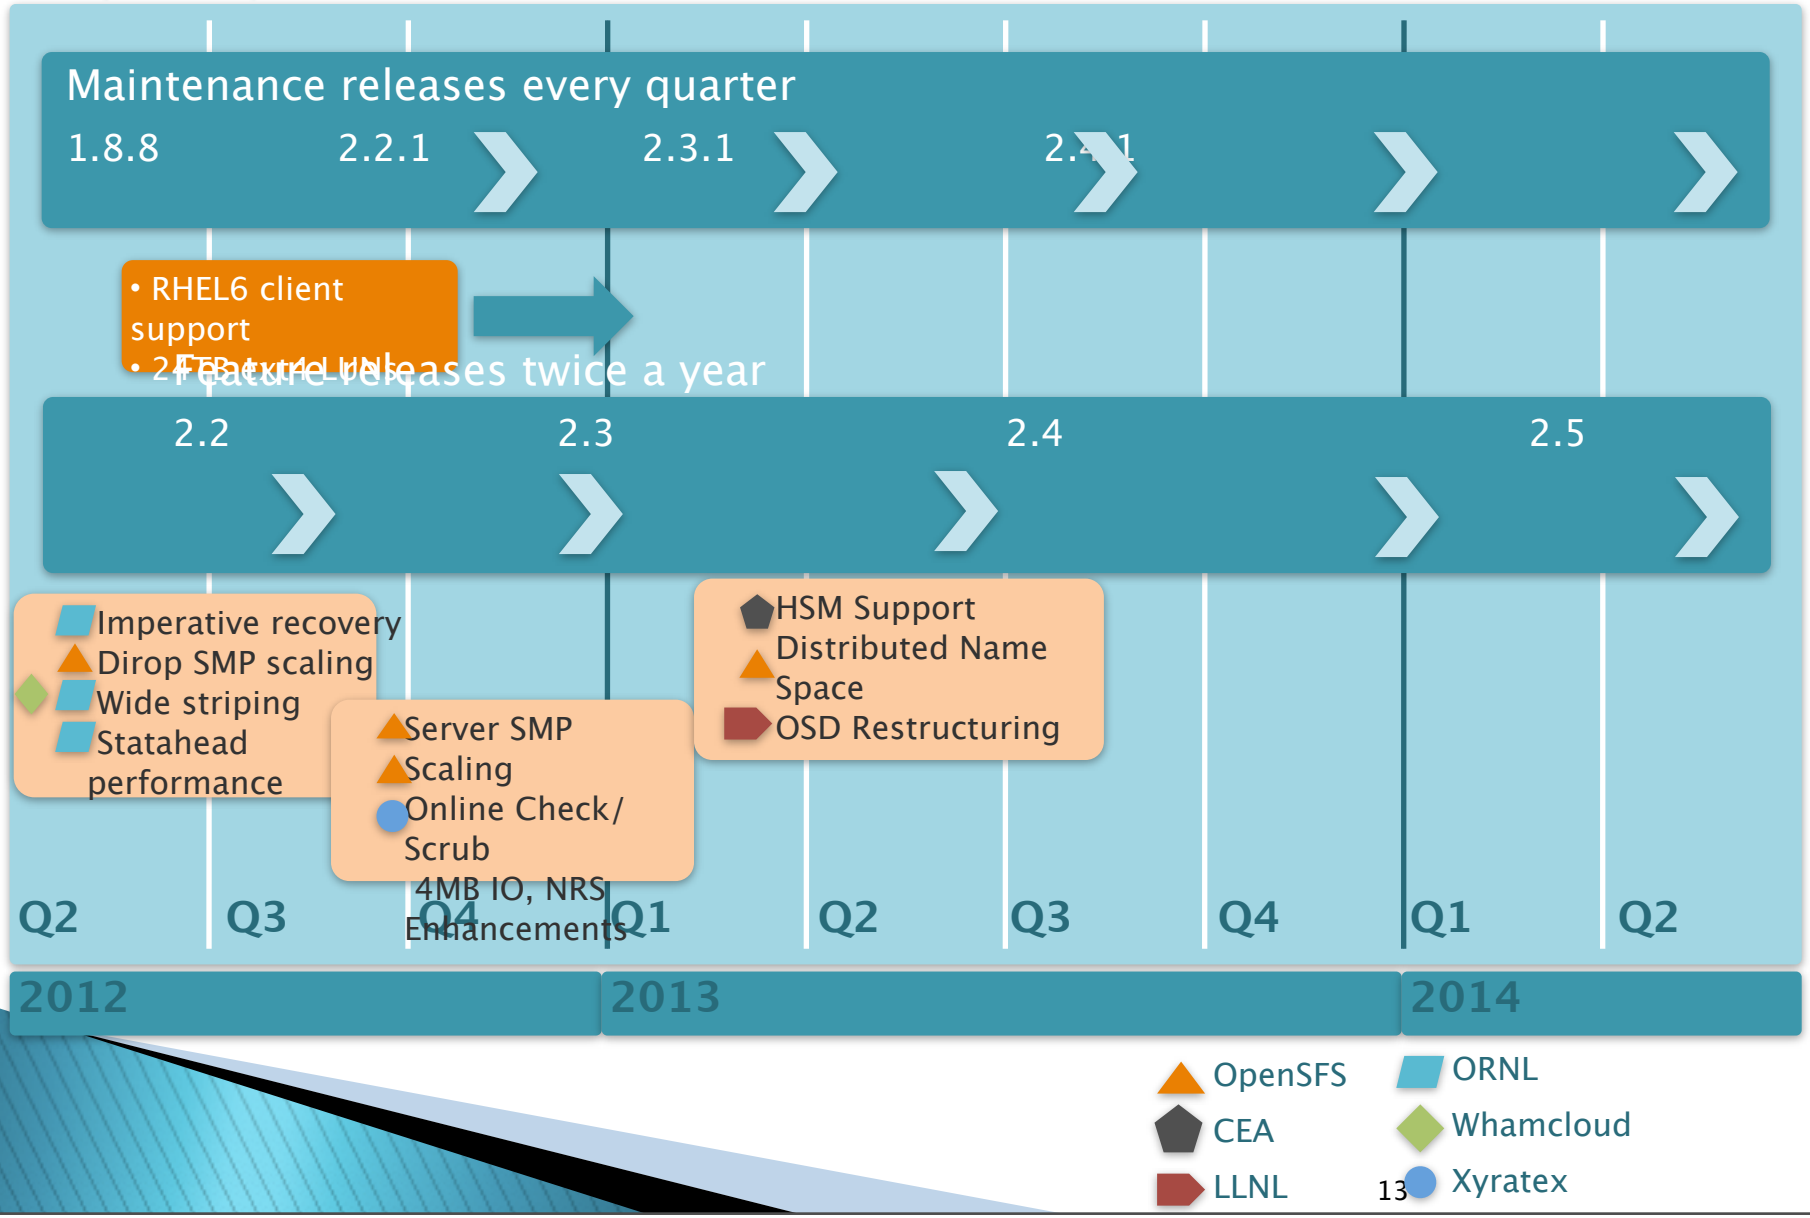

Lustre Roadmap Feature Summary

Feature	Impact	Sponsor	Developer	Release
Wide Striping	File Size	ORNL/WC	ORNL/WC	2.2
Imperative Recovery	Usability	ORNL	WC	2.2
Statahead	MDS	ORNL	WC	2.2
Parallel Directory Ops	MDS	OSFS	WC	2.2
SMP Affinity	MDS	OSFS	WC	2.3
Online File System	Reliability	OSFS	WC	2.3
DNE Remote	MDS	OSFS	WC	2.4
HSM	Usability	CEA	CEA	2.4
Object Storage	Backing Store	LLNL	WC	2.4
DNE Striped	MDS	OSFS	WC	TBD
OSD Quotas	Backing Store	LLNL	WC	TBD

Currently Funded Development Programs

Feature	Purpose	Project
Single Server Metadata Performance Improvements	Scale-up strategy to remove MDS processing	SMP Node Affinity
		Parallel Directory Operations
Distributed Namespace	Scale-up strategy to enable multiple MDS per file system	Remote Directories
		Striped Directories
Lustre File System Checker	Monitor, validate, and repair file system state on-line	Inode Iterator & OI
		MDT-OST Consistency
		MDT-MDT Consistency
WAN Security	Improve support for Lustre as wide-area file system	UID Mapping
		GSSAPI Extensions

A Snapshot of Our Next Steps

Finalizing requirements for the next phase of development

- Latest requirements document is published
- The Technical working group will determine the highest priority items for development

Determining a path forward for a vendor neutral tree

- Towards a shared repository, bug tracker

Defining and develop a community benchmark suite

But What's Missing From OpenSFS?

But What's Missing From OpenSFS?

YOU!

Ensuring the Future of Lustre

What does OpenSFS do for you?

- Funds the lion-share of the Lustre roadmap, without which Lustre would atrophy
- Funds half of the Lustre tree support, ensuring a free and open Lustre tree
- Provides a collaborative environment for ensuring the continued success of the Lustre file system
- Manages and hosts major community events including the LUG

Why you should join OpenSFS?

- Your dues contribute to the continued success of Lustre
- Participation in OpenSFS put you in the drivers seat of the future of Lustre