Technical Working Group (TWG) Face-to-Face Meeting

Dave Dillow
John Carrier
TWG Co-chairs

TWG Mission Statement

- Work with the Lustre community to ensure that Lustre continues to support the stability, performance, and management requirements of the OpenSFS members as HPC compute platforms continue to scale
- Responsible for creating and managing the OpenSFS development roadmap
  - Gather requirements from the Lustre HPC community
  - Prioritize and recommend development projects to the OpenSFS Board
  - Initiate RFPs for important features
  - Work with contractors to meet these requirements

2012 TWG Participants

Diego Moreno	Bull	Bob Ciotti	NASA Ames
J-C Lafoucriere	CEA	Mahmoud Hanafi	NASA Ames
Cory Spitz	Cray	Ken Hornstein	NRL
John Carrier	Cray	Norm Morse	OpenSFS
David Vasil	DDN	Dave Dillow	ORNL
Kit Westeneat	DDN	Galen Shipman	ORNL
Alex Kulyavtsev	FNAL	Sarp Oral	ORNL
Gene Oleynik	FNAL	Evan Felix	PNNL
Andreas Dilger	Intel HPDD	Roger Spellman	Terascala
Bryon Neitzel	Intel HPDD	Alexander Lezhoev	Xyratex
Eric Barton	Intel HPDD	Cheng Shao	Xyratex
Richard Henwood	Intel HPDD	Colin Faber	Xyratex
Robert Read	Intel HPDD	Denis Kondratenko	Xyratex
Andrew Korty	IU	John Forgan	Xyratex
Joshua Walgenbach	IU	Nathan Rutman	Xyratex
Justin Miller	IU	Nic Henke	Xyratex
Steve Simms	IU	Peter Bojanic	Xyratex
Chris Morrone	LLNL	Roman Grigoryev	Xyratex
Marc Stearman	LLNL		

Community Lustre Roadmap

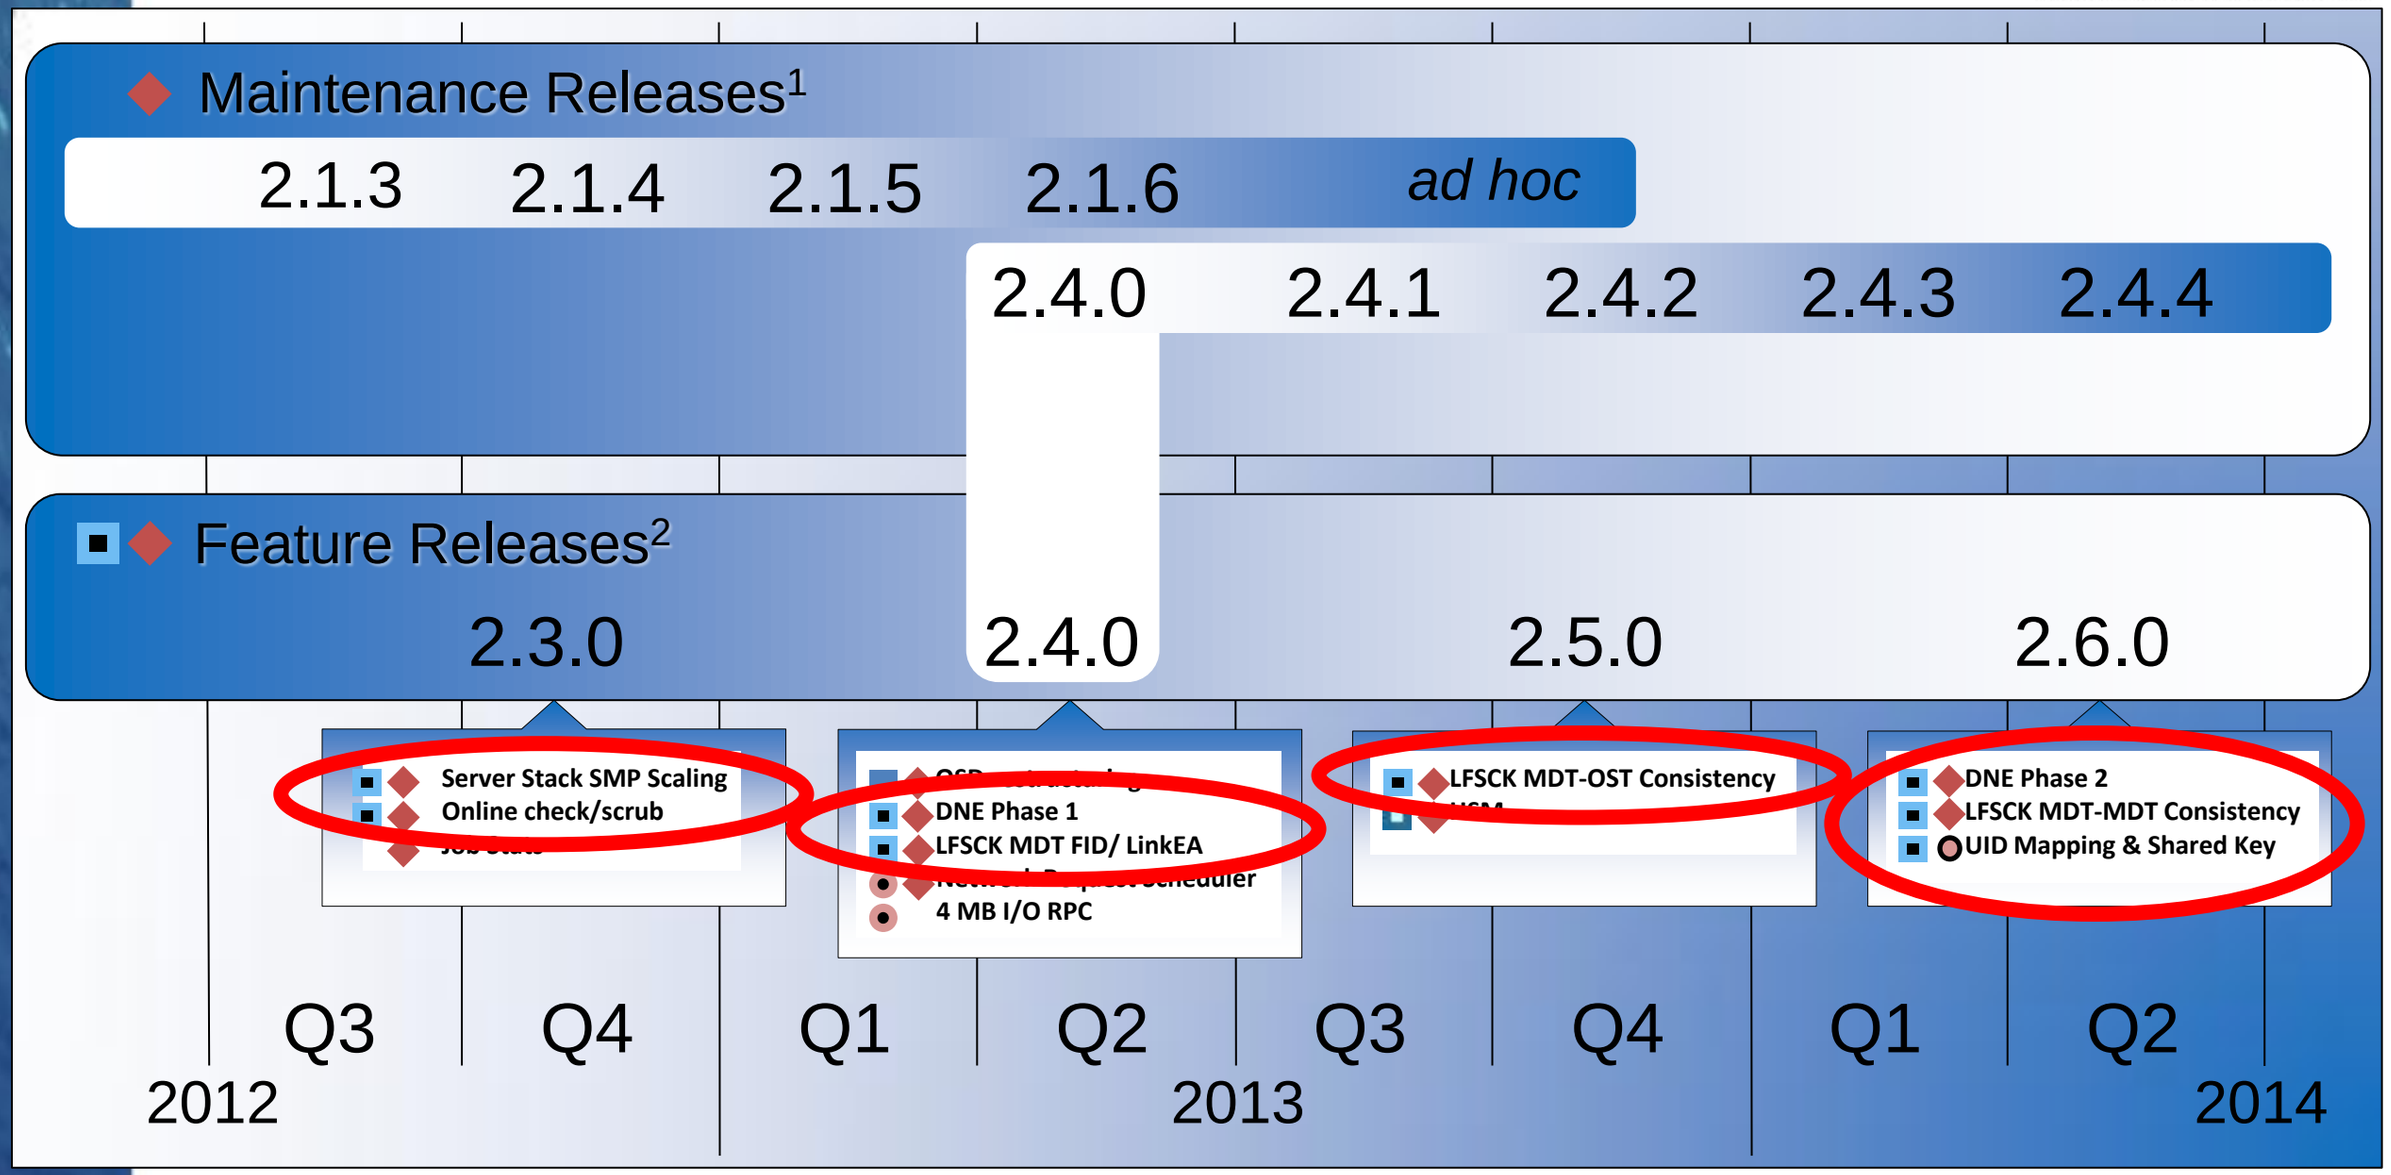

Sponsors for Development and Releases: ORNL, OpenSFS, LLNL, Intel, CEA, Xyratex, Indiana University

[1] Maintenance releases focus on bug fixes and stability. Updates to the current version are made at 3 month intervals. Updates to past versions will be made on an ad hoc basis.

[2] Feature releases focus on introducing new features. New release versions are expected at 6 month intervals. New maintenance versions from the feature release stream are anticipated at 18 month intervals.

SFS DEV-001

- Contract awarded to Intel HPDD (Whamcloud)
  - Dave Dillow, OpenSFS tech. rep.
- Projects
  - Single Server Metadata Performance Improvements
    - SMP Node Affinity
    - Parallel Directory Operations
  - Distributed Namespace
    - Remote Directories
    - Striped Directories
  - Lustre File System Checker
    - Inode Iterator and OI Scrub
    - FID-in-dirent and linkEA consistency
    - MDT-OST Consistency
    - MDT-MDT Consistency
    - Performance

SFS DEV-001 (cont.)

- Status
  - Single Server Metadata Performance Improvements
    - SMP Node Affinity -- complete
    - Parallel Directory Operations -- complete
  - Distributed Namespace
    - Remote Directories -- in demonstration
    - Striped Directories – in scope phase
  - Lustre File System Checker
    - Inode Iterator and OI Scrub -- complete
    - FID-in-dirent and linkEA consistency – in demonstration
    - MDT-OST Consistency – in scope phase
    - MDT-MDT Consistency – in scope phase
- Contract Wiki
  - http://wiki.opensfs.org/Contract_SFS-DEV-001

SFS DEV-002

- Contract awarded to Indiana University
  - Nathan Rutman, OpenSFS tech. rep.
- Projects
  - UID/GID Mapping
  - Shared Key Authentication and Encryption
- Status
  - Design documentation -- complete
  - Implementations -- in progress
- Contract Wiki
  - http://wiki.opensfs.org/Contract_SFS-DEV-002

2012 Recommended Development Areas

Category (prioritized)	**Requirements** (not prioritized)
File system availability and robustness	• Avoid RPC timeouts • Scalable fault management
Storage management	• HSM and storage management infrastructure • OST migration/rebalancing
Performance	• Single client IO performance • File create performance • Directory traversal and attribute retrieval
Lustre networking (LNET)	• LNET channel bonding • Improved LNET robustness • Dynamic LNET configuration

http://wiki.opensfs.org/images/f/f9/OpenSFSTWGRequirements2012.pdf

RFP W4570

- OpenSFS announced the new RFP on 2/21/2013
  - http://www.opensfs.org/rfp-w4570/
- RFP sought proposals to:
  - Further the Lustre roadmap to meet the highest priority requirements defined by the community
  - Develop production quality tools to ease administration and use of open source scalable file systems
  - Address Lustre technical debt to improve the code base and documentation thereof
  - Encourage new efforts in open source scalable file systems for high performance and data intensive computing to broaden the set of solutions available to the community

RFP W4570 (cont)

- OpenSFS received a well intentioned number of responses
- TWG review committee created to review the responses
  - We are seeking volunteers from OpenSFS to help review and prioritize the proposals
  - Our goal is to provide initial feedback to the OpenSFS board by the end of April to confirm the budget
- We will open another opportunity for proposals after completing the evaluation of the responses

Meeting Schedule

- Previously planned to be weekly placeholder
  - –$DAYJOBS and slow progress on the RFP interfered
  - –Fairly good about getting cancellations to the list
- New schedule
  - –Bi-weekly meeting placeholder
  - –Thursdays, 9:30am Pacific
  - –Will confirm or cancel by preceding Tuesday
  - –Restarting meetings May 16
    - RFP progress first!

Next Steps

- After we complete the current RFP, we will begin updating our annual requirements report for the OpenSFS board
- The quality of this information depends on the quality of the input – without your participation, the process will keep following the advice of those that show up
- If you want to change the outcome, you must make your voice heard

Open Discussion

- Gather topics for future TWG meetings
- ….?

Thank You!