

OpenSFS Update

Presented by

Galen Shipman
Technology Integration Group



OpenSFS in a Nutshell

- **A non-profit, technical organization**
 - **California Mutual Benefit Corporation (501.C6)**
- **Focused on high-end, commodity-based, open-source scalable file system technologies**
- **An organization to support Petascale storage requirements and development**
- **A driver and focus for future requirements and development**

Organization

Participants:
Promoters (\$500k)
Adopters (\$50k)
Supporters (\$5k)

Board of Directors:
1 per Promoter
1+ elected by Adopters & Supporters

Officers:
President
Secretary
Treasurer

Technical Working
Group

Gather requirements, Generate SOWs, Roadmap,
Work with contractors on technical requirements

Community Development
Working Group

Facilitate collaborative development model working
with the entire Lustre community

Lustre WAN Working
Group

Improve Lustre's performance and security when
used across the wide area

Other Working Groups
(as necessary)



Who is OpenSFS?



A number of other organizations are planning to announce participation at SC!

OpenSFS – (Linux Foundation model)

- Promoters (Board Members) \$500k/yr
 - Cray, DDN, LLNL, ORNL, Xyratex, Community Board Member
- Adopters \$50k/yr
 - Indiana University, Whamcloud, SNL, +1*
- Supporters \$5k/yr
 - RAID, SGI, Terascale
- Peer Organization
 - European Open File Systems SCE
- Transitioning Supporters
 - TACC, CSC, PSC, SFW, Fermi, SLAC

* will be announced during SC11

Community Development Working Group

Pam Hamilton

Ryan Weiss

Working Group Co-Chairs

Community Development Working Group Mission

- Collaborate with the broader Lustre community on regular releases of the Lustre file system
- Provide manpower from member organizations to facilitate the testing process
- Provide hardware resources from member organizations to facilitate the testing process
- Contribute code which fixes bugs found during testing

CDWG – Current Status

- **The OpenSFS CDWG combined efforts to support the Lustre 2.1 release**
- **Goals of this effort included:**
 - **Collaborate with the entire Lustre community on version 2.1 of the Lustre file system**
 - **Provide manpower from the OpenSFS community to facilitate the testing and bug-fix process**
 - **Provide hardware resources from member organizations to facilitate the testing process**
 - **Contribute code which fixes bugs found during testing**
- **Lustre 2.1 was released on Sept 30th!**
- **Up next is Lustre 2.2!**

OpenSFS and Whamcloud to provide Lustre Gatekeeping and Releases

- **OpenSFS and Whamcloud have announced an agreement to provide Lustre gatekeeping and releases**
- **This agreement will provide the necessary funding and development resources to care for the Lustre tree**
- **This is just one of many ways that OpenSFS is helping to ensure the future of Lustre**

Get Involved

- **CDWG has open meetings most Tuesdays**
 - 9:30am Pacific Time
 - 866-914-3976 or 925-424-8105
 - ID & Password 534986
- **Mailing lists**
 - cdwg@lists.opensfs.org
 - Mailman: <http://lists.opensfs.org/listinfo.cgi>
 - Archives available through link above

Technical Working Group

David Dillow

John Carrier

Working Group Co-chairs

Technical Working Group Mission

- **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 roadmap for the OpenSFS community**
 - **Gather requirements from the Lustre HPC community**
 - **Prioritize and recommend development projects to the Board**
 - **Initiate RFPs for important features**
 - **Work with contractors to meet these requirements.**

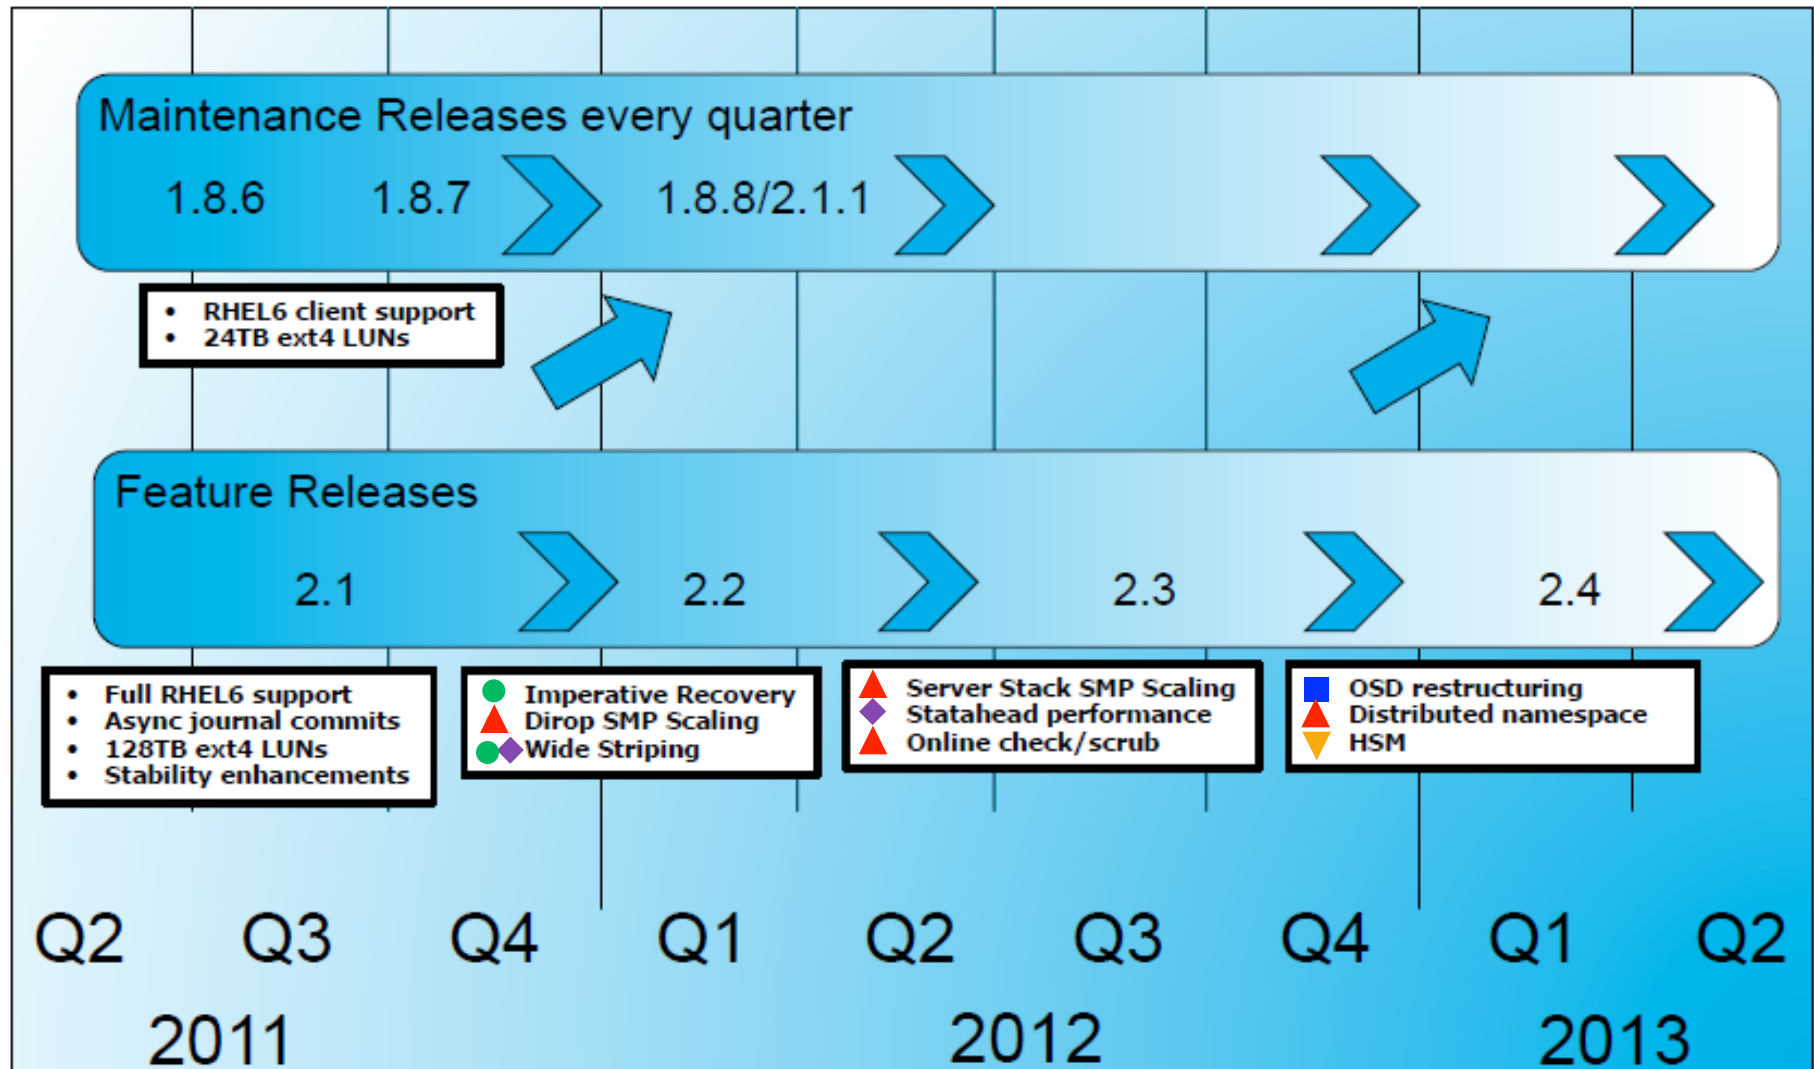
Why is this important?

- **The needs of the HPC community are focused on file system scalability and stability**
- **HPC requires Lustre to run on Linux server platforms and to include features critical to its continued growth in scale**
- **OpenSFS requires technical leadership to guide the funding of a Lustre roadmap that meets those requirements**

TWG – Contract Development

- **Contract development underway with Whamcloud**
- **Single MDS Performance**
 - **Parallel Directory Operations**
- **Distributed Namespace**
 - **Remote Directories**
 - **Phase 1: Multiple MDTs - No striping of directories**
 - **Phase 2: Multiple MDTs and Striped directories**
- **Lustre File System Checker**
 - **Inode Iterator and OI scrub**
 - **Verifies inode->FID mapping in object index table**
 - **Allows for file-level backups of MDT**

Community Lustre Roadmap



Contributor/Sponsor : ◆ Whamcloud ● ORNL ▲ OpenSFS ■ LLNL ▼ CEA

Get Involved!

- **TWG has open meetings most Thursdays**
 - 9:30am Pacific Time
 - 866-304-8294 or 715-726-4994
 - ID & Password 72090
- **Mailing lists**
 - discuss@lists.opensfs.org
 - twg@lists.opensfs.org
 - Mailman: <http://lists.opensfs.org/listinfo.cgi>
 - Archives available through link above

Wide Area Networking Working Group

**Steven Simms
Working Group Lead**

WAN Working Group Mission

- **There are currently gaps in Lustre's functionality that makes its use as a site-wide and wide area filesystem difficult.**
- **The purpose of this group is to gather requirements from others who would like to improve Lustre's performance and security when used across the wide area.**
- **As a secondary mission, for the benefit of a larger community, this group will gather, disseminate, and promote best practices for wide area use of the Lustre filesystem.**

WWG – Development Contract

- **There are currently gaps in Lustre's functionality that makes its use as a site-wide and wide area filesystem difficult.**
- **The purpose of this group is to gather requirements from others who would like to improve Lustre's performance and security when used across the wide area.**
- **As a secondary mission, for the benefit of a larger community, this group will gather, disseminate, and promote best practices for wide area use of the Lustre filesystem.**

Lustre WAN Contract

- **UID Mapping**
 - Provides consistent ownership and permissions across heterogeneous namespaces
 - Current versions break quotas
 - No version exists for 2.x
- **Shared key security mechanism**
 - Will no alter Lustre internals (will use GSSAPI)
 - Alternative to kerberos which will make bridging administrative domains easier

Contact

Galen Shipman

gshipman@ornl.gov

865-576-2672



21 Managed by UT-Battelle
for the U.S. Department of Energy

Presentation Name

