OpenSFS®
Lustre Matters.

LUG 2015

LUSTRE USER GROUP

April 13 - 15 | Denver, CO

$5,000 OFF 1st MEMBERSHIP

Is Big Data and HPC key to your business? OpenSFS has the perfect opportunity for you and your company to get involved—at zero cost.

OpenSFS is providing a limited time offer to subsidize your first year membership—up to $5,000 dollars! You can apply the money to any membership level. Our Supporter level is $5000 per year, so that means your first year membership is FREE!

This gives your technical and business teams the ability to ask questions directly to Lustre experts and test out ideas.

CONTACT US NOW!

http://opensfs.org/get-involved-with-lustre-for-free/

CRAY

Lustre Matters.

Welcome!

Welcome to the 13th Annual Lustre User Group Conference. As key stakeholders in the open source file system community, your participation and input is invaluable. This conference is our primary venue for discussions and seminars on open source parallel file system technologies.

While you are here, we invite you to:

Learn, Share, Develop Best Practices, Provide Feedback to Developers, and Provide Feedback to LUG Organizers

Agenda

LUG will feature three days of sessions, as well as a number of ways to engage with attendees, customers, and industry representatives.

All attendees are invited to join the OpenSFS Board of Directors on Tuesday, April 14th for a discussion on the OpenSFS community and industry topics. This discussion will be held from 5-6pm in the LUG general session room but is not part of the official LUG agenda.

Presentations

- All presentations will be held in Colorado Ballroom A-E (Lower Level 2)
- All sponsor exhibit tables will be located in the Colorado Ballroom Foyer

Technical Poster Exhibitions • April 13-15

- All posters will be located in the back of the Colorado Ballroom A-E

Meals

Breakfast and lunch will be served in the Penrose Ballroom (Lower Level 1). The networking reception on Day 1 will be held in the Colorado Ballroom Foyer (Lower Level 2).

- Breakfast will be served on Days 1, 2 and 3, and lunch will be served on Days 1 and 2
- Snacks will be provided during morning and afternoon breaks in the Colorado Ballroom Foyer

Work Group Meetings

All OpenSFS and EOFS members are invited to participate in the OpenSFS Work Group meetings being held in conjunction with LUG.

- **Benchmarking Working Group**: Tuesday, April 14 from 6:00-7:00pm in the Independence Room (Lower Level 1)
- **Lustre Working Group**: Wednesday, April 15 from 2:00-3:30pm in the Independence Room (Lower Level 1)

OpenSFS is dedicated to providing a harassment-free experience for all Lustre User Group participants. Please review our anti-harassment policy statement online at: http://www.opensfs.org/policies/

Monday, April 13 • Day 1

7:30am - 8:45am	**Registration & Breakfast**
8:45am - 9:10am	**Welcome Remarks and Lustre Overview for Newcomers** Stephen Simms, OpenSFS
9:10am - 9:40am	**OpenSFS and EOFS Welcome and Update** Charlie Carroll, OpenSFS Hugo Falter, EOFS
9:40am - 10:10am	**Lustre 101** Stephen Simms, OpenSFS Brian Andrus, Firstspot, Inc.
10:10am - 10:40am	**Break**
10:40am - 11:30am	**Lustre Development Update** Christopher Morrone, OpenSFS/ Lawrence Livermore National Laboratory
11:30am - 12:00pm	**Status of the Linux Upstream Lustre Client** James Simmons, Oak Ridge National Laboratory
12:00pm - 12:15pm	**Cray's Storage History and Outlook – Lustre+** Jason Goodman, Cray
12:15pm - 1:15pm	**Lunch**
1:15pm - 1:45pm	**From Lab to Enterprise: Growing the Lustre Ecosystem** Malcolm Cowe, Intel
1:45pm - 2:15pm	**SDSC's Data Oasis Gen II: ZFS, 40GbE, and Replication** Rick Wagner, San Diego Supercomputer Center
2:15pm - 2:45pm	**There and Back Again: The Battle of Lustre at LANL** Kyle Lamb, Michael Masson, and Susan Coulter Los Alamos National Laboratory
2:45pm - 3:00pm	**Seagate Lustre Update** Peter Bojanic, Seagate Technology
3:00pm - 3:30pm	**Break**
3:30pm - 4:00pm	**Lustre & Kerberos: In Theory and in Practice** Sebastien Buisson, Bull
4:00pm - 4:30pm	**GSS Shared Key Update and Using UID Mapping in Lustre 2.7** Stephen Simms, Indiana University
4:30pm - 5:00pm	**ZFS *as Backend File System for Lustre* the Current Status, How to Optimize, Where to Improve** Gabriele Paciucci, Intel
5:00pm - 5:30pm	**Per User Lustre File Systems Leveraging ZFS Pool** Marc Stearman, Lawrence Livermore National Laboratory
5:30pm - 7:30pm	**Networking Reception**

Tuesday, April 14 • Day 2

8:00am - 9:00am	**Registration & Breakfast**
9:00am - 9:30am	**Lustre 2.8 and Beyond** Andreas Dilger, Intel
9:30am - 10:00am	**Progressive File Layouts** John Hammond, Intel Jason Hill, Oak Ridge National Laboratory
10:00am - 10:30am	**Break**
10:30am - 11:00am	**Intelligent Cache Hinting in Lustre** Micah Bhakti, Intel
11:00am - 11:15am	**Intel and Lustre: Growth of Lustre Adoption and Intel's Continued Commitment** Brent Gorda, Intel
11:15am - 11:45am	**Lustre HSM in the Cloud** Robert Read, Intel
11:45am - 12:00pm	**Application-optimized Lustre Solutions for Big-Data Workflows** Robert Triendl, DataDirect Networks
12:00pm - 1:00pm	**Lunch**
1:00pm - 1:30pm	**Metadata Access Reduction and Analysis and Elimination of Client Evictions on a Large Scale Lustre Based File System** Keiji Yamamoto, RIKEN AICS Shinji Sumimoto, Fujitsu
1:30pm - 2:00pm	**Lustre Network (LNET) Router Configuration and Tuning** John Fragalla, Seagate Technology
2:00pm - 2:30pm	**Towards Continuous 24x7x365 Production Operation on a Router-less 20 Petabyte, 11,200 Node Cluster** **– OR –** **Defending the Planet with Lustre: Your Life Could Depend on it!** Bob Ciotti, NASA
2:30pm - 3:00pm	**Break**
3:00pm - 3:30pm	**Lustre Metrics: New Techniques for Monitoring Lustre** Scott Nolin and Andrew Wagner UW Space Science and Engineering Center
3:30pm - 4:00pm	**New and Improved Lustre Performance Monitoring Tool** Torben Kling Petersen, Seagate Technology
4:00pm - 4:30pm	**Correlating Multiple TB of I/O Performance Data to User Jobs** Michael Kluge, ZIH
4:30pm - 5:00pm	**Monitoring a Heterogeneous Lustre Environment** Frédérick Lefebvre, Calcul Quebec - Universite Laval
Optional 5:00pm - 6:00pm	**OpenSFS Board Meeting and Community Discussion** OpenSFS Board Members

Wednesday, April 15 • Day 3

8:00am - 9:00am	**Registration & Breakfast**
9:00am - 9:30am	**Reliability of NVM Devices for I/O Acceleration on Supercomputing Systems** Hitoshi Sato, Tokyo Institute of Technology
9:30am - 10:00am	**Scalability Testing of DNE2 in Lustre 2.7** Stephen Simms, Indiana University
10:00am - 10:30am	**Break**
10:30am - 11:00am	**Shared File Performance in Lustre: Challenges and Solutions** Patrick Farrell, Cray
11:00am - 11:30am	**OSD-Btrfs, a Novel Lustre OSD Based on Btrfs** Shuichi Ihara and Li Xi, DataDirect Networks
11:30am - 12:00pm	**Deploying Hadoop on Lustre Storage: Lessons Learned and Best Practices** J.Mario Gallegos, Dell Zhiqi Tao, Intel
12:00pm - 12:30pm	**Understanding Hadoop Performance on Lustre** Stephen Skory, Seagate Technology
12:30pm - 12:40pm	**Closing**
12:40pm	**End of LUG 2015**

Technical Poster Exhibitions • April 13-15

Bottom-up Performance Estimation for a Cray Sonexion 2000

- Brett Kettering, Lawrence Livermore National Laboratory
- Alfred Torrez, Lawrence Livermore National Laboratory
- Bradley Settlemyer, Lawrence Livermore National Laboratory
- Ruth Klundt, Lawrence Livermore National Laboratory

Early Test Results for Trinity's Sonexion 2000 Lustre PFS

- Brett Kettering, Lawrence Livermore National Laboratory
- Alfred Torrez, Lawrence Livermore National Laboratory
- Bradley Settlemyer, Lawrence Livermore National Laboratory
- Ruth Klundt, Lawrence Livermore National Laboratory

FastForward Project

- John Bent, EMC

Lustre Beyond 4-Walls

- Keith Mannthey, Intel
- Gerry Jankauskas, Bay Microsystems
- Mark Rodriguez, Bay Microsystems

Updating the HPC Platform at the Bank of Italy: Intel EE Lustre as Multipurpose Filesystem

- Giuseppe Bruno, Bank of Italy

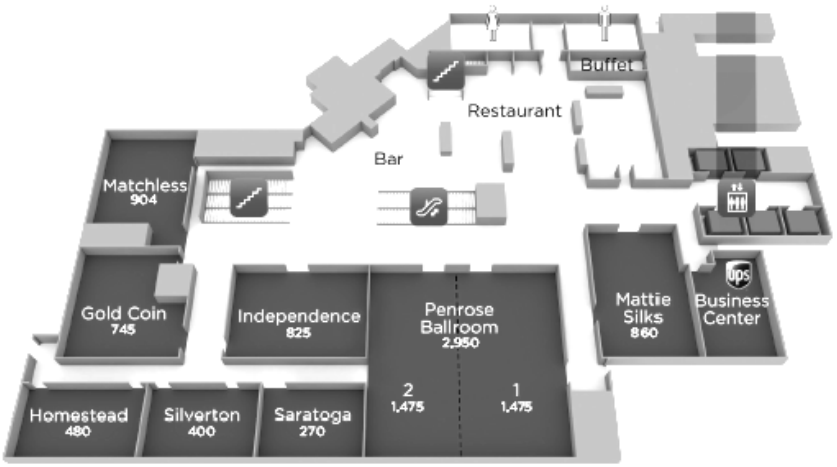

Lower Level 1

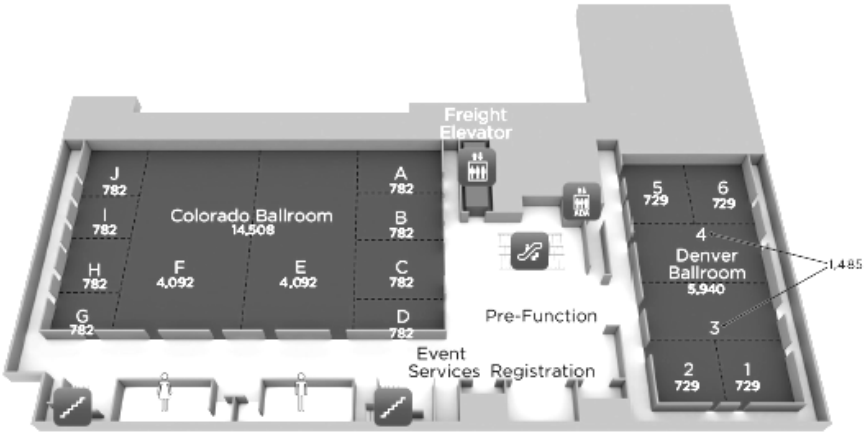

Lower Level 2

About OpenSFS: OpenSFS (Open Scalable File Systems) is a strong and growing nonprofit organization dedicated to the success of the Lustre® file system. OpenSFS was founded in 2010 to advance Lustre development, ensuring it remains vendor-neutral, open, and free. Since its inception, OpenSFS has been responsible for advancing the Lustre file system and delivering new releases on behalf of the open source community. Through working groups, events, and ongoing funding initiatives, OpenSFS harnesses the power of collaborative development to fuel innovation and growth of the Lustre file system worldwide.

We hope you enjoyed LUG 2015 and welcome your input!

Complete the LUG Conference survey for your chance to win a *Kindle Fire HD Tablet, Bose® Acoustic Noise Cancelling Headphones,* or *a GoPro Hero 3 Camera.*

LUG 2015 is Made Possible by our Generous Sponsors

Gold Sponsors

Silver Sponsors

LUG 2015 Planning Committee

Sebastien Buisson – Bull
Aurélien Degrémont - CEA
Stephane Thiell - CEA
Stephen Simms - Indiana University
Peter Jones – Intel
Mark Gary - Lawrence Livermore National Laboratory
Parks Fields - Los Alamos National Laboratory
Feiyi Wang - Oak Ridge National Laboratory
Sarp Oral - Oak Ridge National Laboratory
Jesse Casman - Oppkey
Steve Monk - Sandia National Laboratories
Meghan McClelland – Seagate Technology
Michael Connolly – Seagate Technology
Jen Franklin – VTM Group
Jordan Kees – VTM Group
Molly Nelson – VTM Group

www.opensfs.org