Cray Lustre Update
CUG 2014

Cory Spitz
Tom Sherman

Safe Harbor Statement

This presentation may contain forward-looking statements that are based on our current expectations. Forward looking statements may include statements about our financial guidance and expected operating results, our opportunities and future potential, our product development and new product introduction plans, our ability to expand and penetrate our addressable markets and other statements that are not historical facts. These statements are only predictions and actual results may materially vary from those projected. Please refer to Cray's documents filed with the SEC from time to time concerning factors that could affect the Company and these forward-looking statements.

The Role of Lustre at Cray

- **Cray delivers maximum system performance**
  - Implies a **fast file-system**
  - **Repeatable maximum performance** often is the key
  - Balance throughput by **eliminating bottlenecks**
  - **Lustre is best** for delivering performance at scale
- **Multiple delivery platforms**
  - Sonexion
  - CLFS – Lustre File System by Cray w/DDN & NetApp
  - DAL – Direct-attached Lustre w/DDN & NetApp
  - C3 – Cray Cluster Connect (client)

Cray's Internal Lustre Strategy

- Enhance and expand a strong cross-discipline storage team
- Augment with support vendor
- Generate and gather system requirements
- Prioritize engineering efforts
- Create a strong test regimen and release process
- Keep close customer involvement

Cray's engineering priorities change as Lustre capabilities evolve

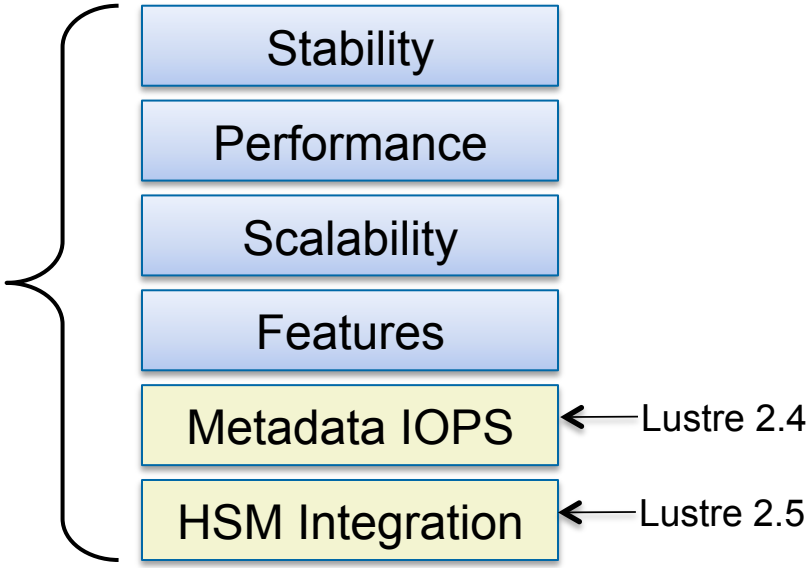

Cray's Community Lustre Strategy

OpenSFS – Original Founder and promoter-level member

- Founded by Cray, DDN, LLNL, and ORNL
- Non-profit technical organization focused on high-end open-source scalable file system technologies, primarily Lustre

Goals

- Foster collaboration among entities developing and deploying Lustre
- Drive technologies for future requirements into open, scalable file systems
- Deliver Lustre file system releases and roadmap designed to meet those goals

OpenSFS partnership created → **Lustre development process reestablished** → **Multi-stage roadmap in place**

Cray's Lustre Release Strategy

Demonstrate excellence in HPC

- **Quality** – continuously improve dependability and reliability
- **Performance** – measure and monitor performance without regressions
- **Scale** – maintain Cray-level performance & quality at scale

Tactics

- Work with community at head of development (master)
- Continuously and automatically integrate and test upstream changes
- Provide comprehensive test feedback of release candidates
- Leverage both feature and community maintenance branches

Plan added enhancements independently

- Lustre development is moving rapidly
- Watch for regressions; new features shouldn't destabilize functionality
- Deliver when ready

We've Reached an Inflection Point

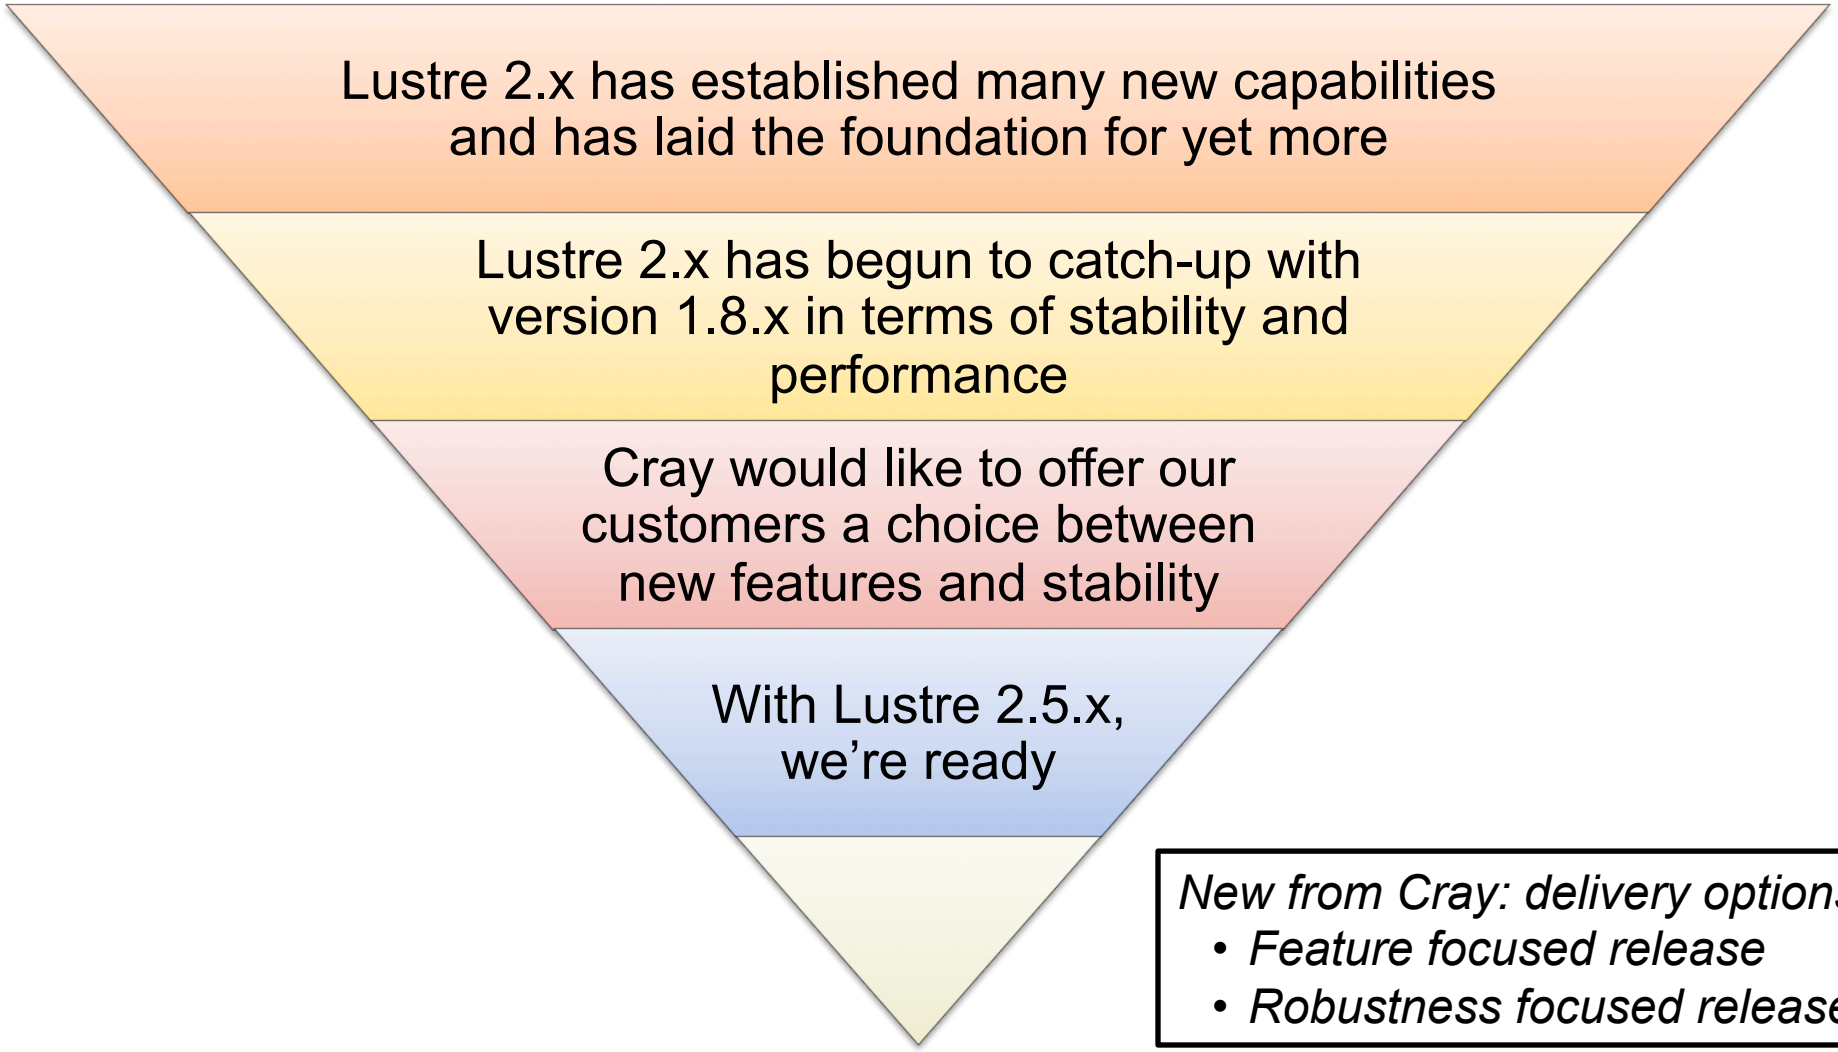

New from Cray: delivery options

- *Feature focused release*
- *Robustness focused release*

SDM Feature and Maintenance Streams

Cray SDM will begin to offer two release streams

- Feature release stream based on upstream community feature release
- Maintenance release stream based on Intel maintenance branch

Cray HPCS roadmap to include:

- Regular maintenance release updates
- Feature releases as needed
  - (e.g. CLE Rhine on SLES 12 may require newer Lustre clients)

Community Lustre Roadmap

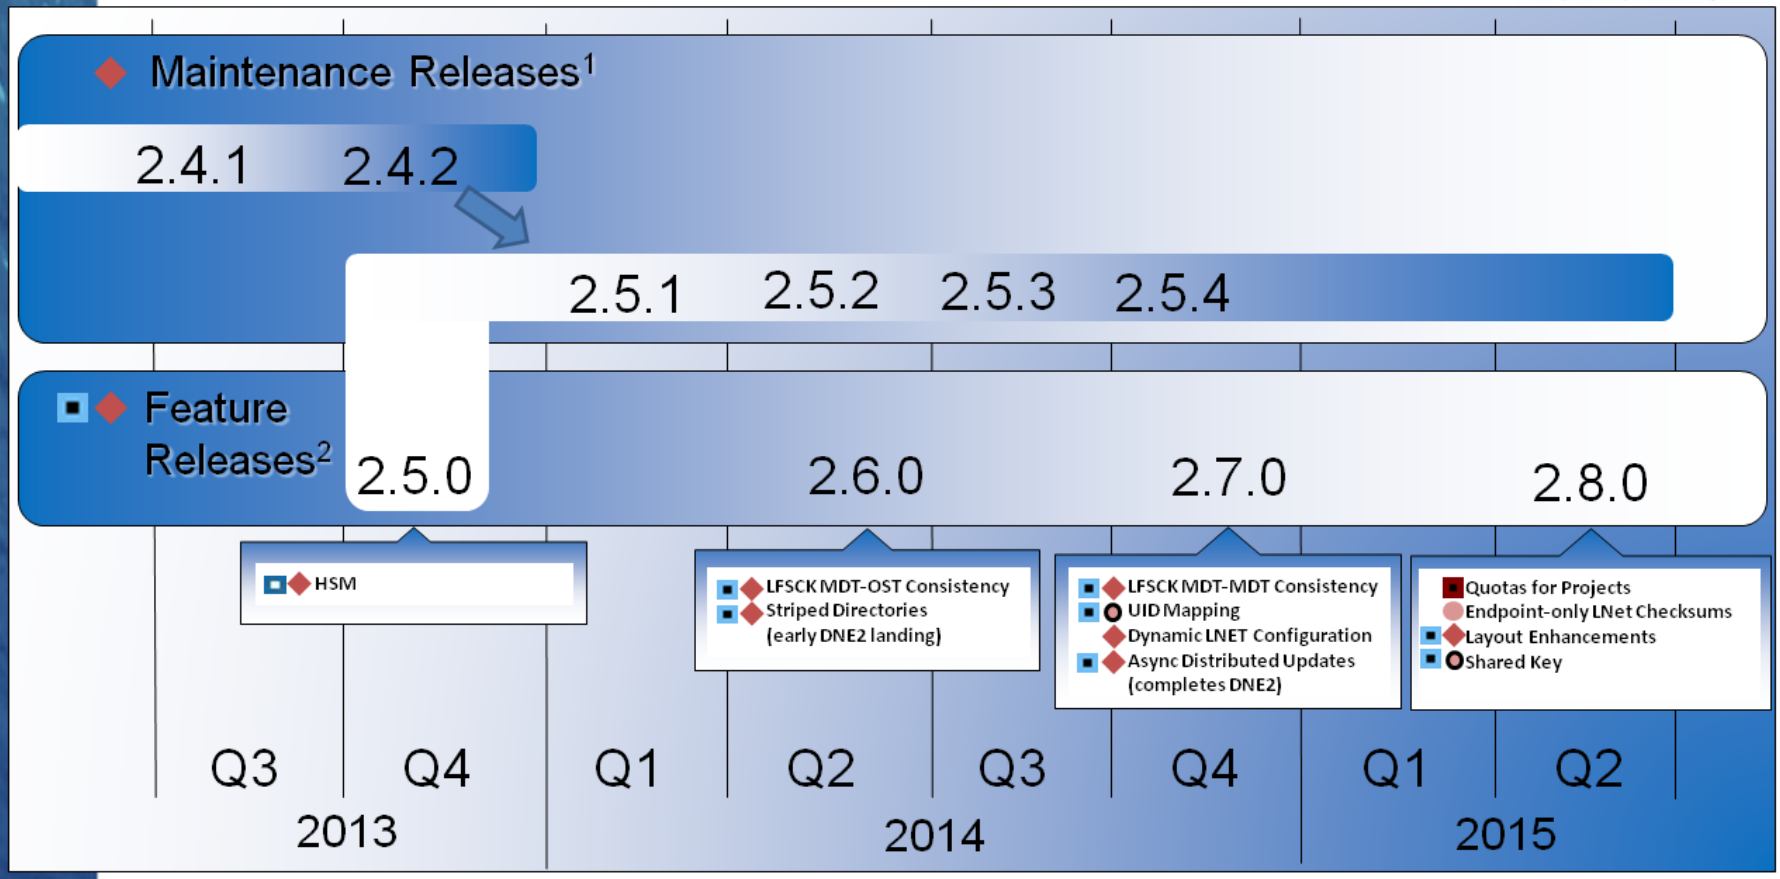

Sponsors for Development and Releases: ORNL, OpenSFS, Intel, CEA, Indiana University, DDN

[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.

Cray Sonexion 1600 Software Roadmap

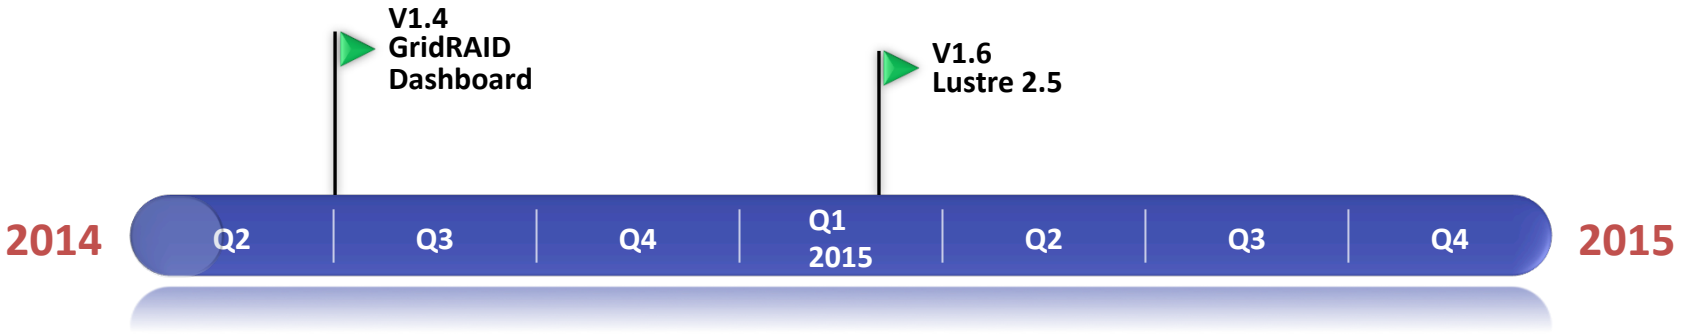

Sonexion 1600

- **Features**
  - V1.4 offers GridRAID declustered parity
    - 300% faster rebuilds
  - V1.4 includes an enhanced dashboard
  - V1.6 provides Lustre 2.5
    - HSM support

CLFS Software Roadmap

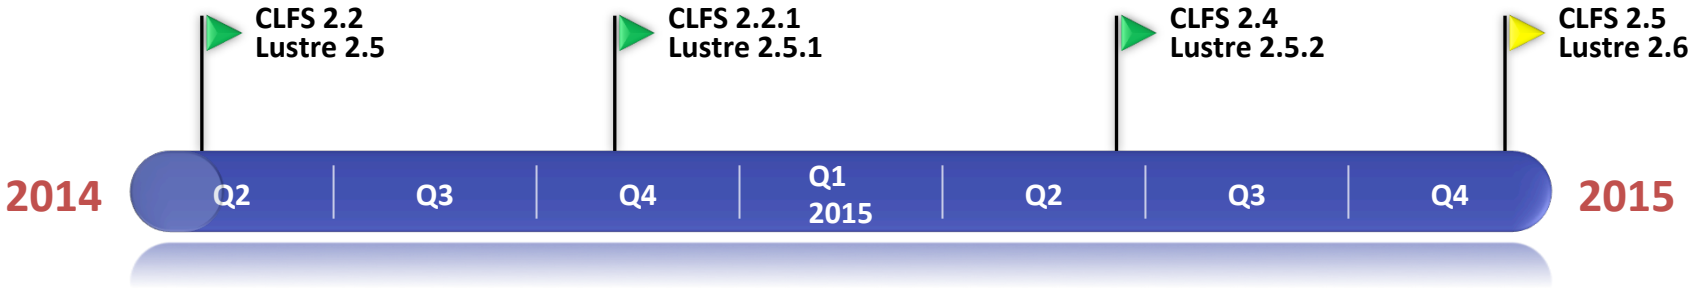

- Features
  - Lustre 2.5: HSM, DNE Phase 1
  - Lustre 2.6: DNE Phase 2, part 1
- Notes
  - Release versions and dates are subject to change

Under promise / Over deliver

- **Old plan:**
  - Last release for XE in 2013
- **New plan:**
  - Additional XE release: 5.2.UP00
    - SLES 11SP3
    - Lustre 2.5.0 server
    - Lustre 2.4.1 or 2.5.0 client
  - Additional XE release: 5.2.UP02
    - Lustre 2.5.1

Loads of new features in Lustre 2.x

- **2.0/2.1**
  - Commit on share
  - **Changelogs**
  - lustre_rsync
- **2.2**
  - **Statahead**
  - **Asynchronous Glimpse Lock (AGL)**
  - Client parallel checksums
  - **Imperative recovery**
  - Large xattrs (aka wide striping)
  - Multi threaded ptlrpcd
  - OSD API
  - Parallel directory operations (pDirOps)
  - Quota protocol and space accounting
- **2.3**
  - Server stack SMP scaling
  - **Online check/scrub**
  - **Job Stats**
  - OFD (using OSD API)
  - CRC in memory checksum
  - **Quota accounting**
- **2.4**
  - OSD restructuring
  - **DNE Phase I**
  - **LFSCK MDT FID/LinkEA**
  - **Network Request Scheduler**
  - **4 MB I/O RPC**
- **2.5**
  - HSM *[stay for Scott's talk next]*
  - Quota improvements

Coming soon from Cray: pingless clients, improved LNET RAS & re-routing

Helpful Cray Docs

- **Lustre Operations Manual**
- **S-2444 – Installing and Configuring Cray Linux Environment (CLE) Software**
  - Includes IMPS + DAL
- **S-2520F – CDL installation**
- **S-2521C – CLFS installation**
- **S-2522E – CIMS installation**
- **S-2327C – DMP admin guide**
  - includes esfsmon and DNE configuration
- **S-2508A – C3 Lustre**
- **Field support have access to helpful TOI**
  - CIMS upgrade/migration
  - CLFS upgrade/migration
  - IMPS
  - DAL migration
  - DNE

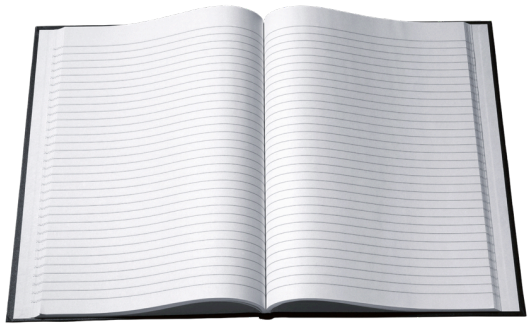

Legal Disclaimer

Information in this document is provided in connection with Cray Inc. products. No license, express or implied, to any intellectual property rights is granted by this document.

Cray Inc. may make changes to specifications and product descriptions at any time, without notice.

All products, dates and figures specified are preliminary based on current expectations, and are subject to change without notice.

Cray hardware and software products may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Cray uses codenames internally to identify products that are in development and not yet publically announced for release. Customers and other third parties are not authorized by Cray Inc. to use codenames in advertising, promotion or marketing and any use of Cray Inc. internal codenames is at the sole risk of the user.

Performance tests and ratings are measured using specific systems and/or components and reflect the approximate performance of Cray Inc. products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance.

The following are trademarks of Cray Inc. and are registered in the United States and other countries: CRAY and design, SONEXION, URIKA, and YARCDATA. The following are trademarks of Cray Inc.: ACE, APPRENTICE2, CHAPEL, CLUSTER CONNECT, CRAYPAT, CRAYPORT, ECOPHLEX, LIBSCI, NODEKARE, THREADSTORM. The following system family marks, and associated model number marks, are trademarks of Cray Inc.: CS, CX, XC, XE, XK, XMT, and XT. The registered trademark LINUX is used pursuant to a sublicense from LMI, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis. Other trademarks used in this document are the property of their respective owners.