

OS/R SC15

Larry Kaplan
Chief Software Architect

What Cray Has Today

- OpenMP 4.0 moving to 4.5
- MPI 3.0/3.1
- Core/Thread Specialization
- DataWarp
- **ALPS placement and resiliency**
 - Binding of threads and memory
 - Control tree reconstruction on node failure



Workload Management

- **Cray provides significant integration with third party workload management systems (WLMs)**
- **Power management**
 - P-state control, power capping (system and job, with variation), idle node power off, ramp/band management, accounting
- **DataWarp**
 - Storage instance management, pre-/post-run operations, swap
- **Processor reconfiguration**
 - Today: NUMA modes, MCDRAM config



COMPUTING

COMPUTE

STORE

ANALYZE

What Cray is Planning

- **Containers**

- Shifter work from NERSC

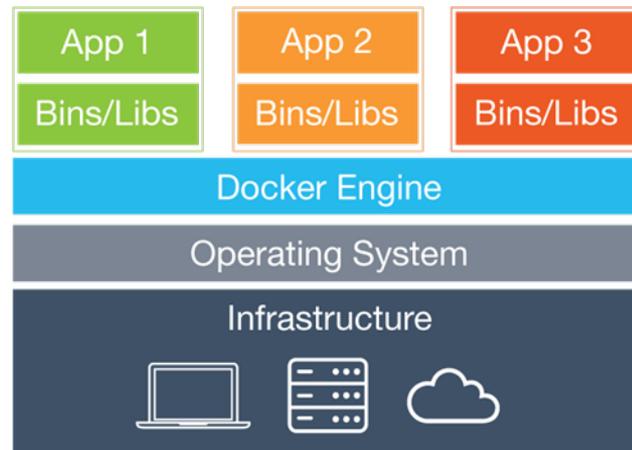
- **OpenMP extensions**

- E.g. memory hierarchy handling

- **“Public” pub/sub bus**

- **Considering fine-grained power management**

- Providing infrastructure



Looking to the Community for

- **Task-based runtimes**

- Which are solving the important issues, are usable, integratable?

- **Global resource management and meta-scheduling**

- How should multiple different resource types be coordinated for efficient execution of workflows?
- Cray must provide appropriate management infrastructure

- **Programming model resiliency**

- Not clear it will be needed but don't want to be caught flat-footed
- Cray participating and planning to adopt

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 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 names and brands may be claimed as the property of others. Other product and service names mentioned herein are the trademarks of their respective owners.

Copyright 2015 Cray Inc.

