# Cray SHMEM and OpenSHMEM

David Knaak
Cray Inc.
OpenSHMEM Workshop 2016

# About Cray

## Seymour Cray founded Cray Research in 1972

- 1972-1996, Cray Research grew to leadership in Supercomputing
- 1996-2000, Cray was subsidiary of SGI
- 2000- present, Cray Inc. growing to $724.7M in revenue in 2015
- Cray Inc. formed in April 2000

## Cray Inc.

- NASDAQ: CRAY
- Over 1,200 employees across 30 countries
- Headquartered in Seattle, WA

## Three Focus Areas

- Computation
- Storage
- Analytics

## Seven Major Development Sites:

- Austin, TX
- Chippewa Falls, WI
- Pleasanton, CA
- St. Paul, MN
- San Jose, CA
- Seattle, WA
- Bristol, UK

# Cray's Vision:
# The Fusion of Supercomputing and Big & Fast Data

## Modeling The World

# Computational Tools for Your Most Challenging Problems

Compute

Store

Analyze









A Workflow Approach for a Data-Intensive World

COMPUTE | STORE | ANALYZE

# Cray® XC40™ Supercomputer

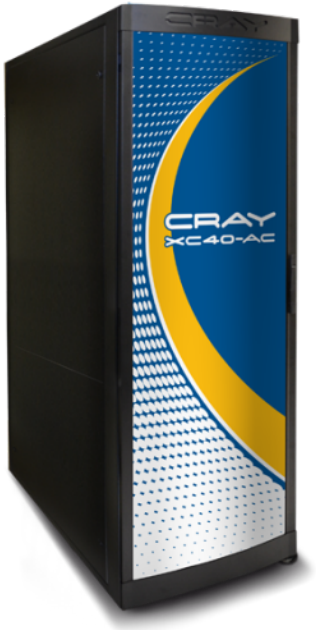


**Adaptive**

- **Flexibility** of x86 processors, coprocessors and accelerators
- **Investment protection,** upgradable by design

**Integrated**

- **Integrated** HPC software environment and storage

**High Performing**

- **Sustained, scalable** application performance
- Upgradable to **100 petaflops**
- **Extreme** performance interconnects, packaging, cooling and more

**Scalable Performance**

# Cray® CS400™ Cluster Supercomputers

**Configurable**

- **Designed for broad range of workloads**
- **Industry standards**-based
- Blades or rackmount server; multiple **interconnect** and **storage options**

**Manageable**

- Customizable **HPC cluster software stack**
- **Advanced Cluster Engine** system management software

**Reliable & Efficient**

- Built-in and optional **energy efficiencies**
- **Designed for reliability**
- Multiple levels of redundancy

**Flexible Performance**

# Cray® Sonexion® Scale-Out Lustre® Storage


**Simplified Management**

- **Fully integrated** and **preconfigured** design; fewer components to manage

**Precision Performance**

- Scale I/O from **7.5 GB/s more than 1.7 TB/s** in a single file system

**Balanced Scalability**

- Scale in modular increments; **reduce capital costs** as capacity grows

**Quality Design**

- **Reduce hardware footprint by 50%** over component-based solutions
- Cray-ensured reliability and stability at scale

**Get Results — Up to 24x Faster**

# Urika®-GX Agile Analytics Platform

**Serious Agility that's enterprise accessible**

- Run analytics workloads concurrently: Hadoop®, Spark™, graph, and HPC*
- Dynamic resource repurposing
- Simple, standards-based tools for straightforward management

**Pervasive Speed based on Cray's Supercomputing experience**

- Pre-integrated and validated software + hardware
- Cray Aries fabric with high I/O throughput and low latency

**High-Frequency Insights**

- Make decisions with near immediacy using the Cray Graph Engine

**Unprecedented versatility and speed**

* Validated now, additional support to follow

# Cray Customers by Segment



**Earth Sciences**

Met Office
DWD
ECMWF
NCEP
Australian Government Bureau of Meteorology
DMI
KMA Korea Meteorological Administration
Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra
FINNISH METEOROLOGICAL INSTITUTE

**Govt. & Defense**

NERSC
Los Alamos NATIONAL LABORATORY
Sandia National Laboratories
Argonne NATIONAL LABORATORY
OAK RIDGE National Laboratory
National Nuclear Security Administration
University of California Lawrence Livermore National Laboratory

**Higher Education**

HLRS
EPSRC
PAWSEY supercomputing centre
KAUST
INDIANA UNIVERSITY
NCSA
CSCS
KTH VETENSKAP OCH KONST
PDC Center for High Performance Computing
UNIVERSITY OF HAWAI'I
STATE

**Life Sciences**

Mount Sinai
MSSM MOUNT SINAI SCHOOL OF MEDICINE
MAYO CLINIC
THE UNIVERSITY OF CHICAGO MEDICAL CENTER & BIOLOGICAL SCIENCES
Duke
養和醫院 Hong Kong Sanatorium & Hospital
FORTUNE 100
Undisclosed

**Energy**

PGS
ExxonMobil
Chevron
TOTAL COMMITTED TO BETTER ENERGY
PETROBRAS
Undisclosed

**Manufacturing**

FORTUNE 100
Formula 1
Undisclosed
Undisclosed
Stalprodukt S.A.
Undisclosed

**Financial Services**

FORTUNE 100
Undisclosed

COMPUTE | STORE | ANALYZE

# Cray Technology Innovations

**Combining advanced technologies into productive, scalable systems**

### System Interconnect
Interconnect and optimization software to address the data transfer bottleneck at large scale

### Systems Management & Performance Software
Software to productively manage and extract performance out of thousands of processors used as a single system

### Packaging
Greenest x86 supercomputers with innovative cooling and upgradability to improve TCO

# Cray SHMEM and OpenSHMEM

Cray is a leader in HPC systems, in SHMEM design and implementation, and in OpenSHMEM standardization.

Cray is committed to supporting its HPC customers with a high performance and high quality implementation of the OpenSHMEM API. Our customers' needs often require that we pioneer extensions to the API. We usually design and implement the extensions as soon as possible to meet those needs but we also work with the OpenSHMEM Committee to define these APIs so that they also meet the needs of the broader set of OpenSHMEM users.

# OpenSHMEM Extensions

**Why do we need more than current API?**
**Primarily due to trends in system architectures as we move towards exascale:**

- Increased complexity
- Increasing number of cores in multi-core processors
- Memory hierarchies, including high bandwidth memory
- Processor accelerators
- Increased network capabilities to offload communication work from compute processors
- Other new concepts that help programmability and performance

# OpenSHMEM Extensions

## What makes a desirable OpenSHMEM extension?

- Improves ease of OpenSHMEM programming
- Improves performance of OpenSHMEM programs
- Aids portability by hiding system differences within the specific implementations
- Has a user friendly API
- Is consistent with existing OpenSHMEM API

# Cray SHMEM extensions to OpenSHMEM

| Proposed Features | Cray SHMEM | OpenSHMEM |
|---|---|---|
| Non-Blocking Put | 2007 | V1.3 |
| Non-Blocking Get | 2011 | V1.3 |
| Alltoall Collectives | 2013 | V1.3 |
| Alltoallv Collectives | 2013 | #183 |
| Global Exit | 2014 | V1.2 |
| Flexible PE Subsets,aka Teams | 2014 | #193 |
| Thread-Safety | 2014 | #186 |
| Put With Signal | 2014 | #77 |
| Memory Hierarchy, HWB memory | 2016 | N/A |
| Contexts-Domains | prototype | #177 |

# Summary – Cray SHMEM and OpenSHMEM

- We are committed to giving our customers a high performance and high quality implementation of the OpenSHMEM API
- We are committed to working with our customers to advance OpenSHMEM with new features
- We evaluate requests in light of resources available and other priorities
- We are committed to working with OpenSHMEM Committee
- We have a proven track record of doing these

# Discussion

- Cray extensions you have  used?
- Other vendors' extensions you have used?
- Proposed extensions you hope to use?
- Additional extensions you would like to see?