

Cray Operating System Plans and Status

Charlie Carroll
May 2012



Cray Operating Systems and I/O

Compute node

- Compute Node Linux
- NVIDIA GPU driver

Service node OS

File systems: Lustre

Networking

- HSN: Gemini and Aries (GNI, DMAPP)
- IBGNI (IB verbs -> Gemini)
- TCP/IP

Third-party extensions

- GPFS
- Panasas
- Batch systems

Operating system services

- Core specialization
- Dynamic Shared Library (DSL) support
- Cluster Compatibility Mode
- DVS (Data Virtualization Service)

System management

- ALPS (Application-Level Placement Scheduler)
- Node Health Checker (NHC)
- CMS (Cray Management Services)
- Command interface

Hardware Supervisory System

- Handling errors, resiliency
- Event routing
- Booting the system

Cray Operating Systems Focus

Performance

- Maximize compute cycles delivered to applications while also providing necessary services
 - Lightweight operating system on compute node
 - Standard Linux environment on service nodes
- Optimize network performance through close interaction with hardware
- GPU infrastructure to support high performance

Stability and Resiliency

- Correct defects which impact stability
- Implement features to increase system and application robustness

Scalability

- Scale to large system sizes without sacrificing stability
- Provide system management tools to manage complicated systems

Accomplishments of the Past Year

CLE 4.0, UP01, UP01A, UP02, UP03

- Support for AMD's Interlagos, used in our currently shipping systems
- Nvidia Fermi GPU support
- Cray Sonexion introduction
- DVS performance improvements
- Memory Control Groups reduced OOM issues

Cluster Compatibility Mode and ISV Application Acceleration

Quality improvements

- Resiliency features improving system reliability
 - Moving to fewer PMs (preventive maintenance) and less down-time
 - Steve Johnson's talk follows this one

Aries bring-up

Programming Environments coordination

- Resiliency features, GPUs, DSLs (dynamic shared libraries)

External servers

Hiring

Demographics (April 2012)

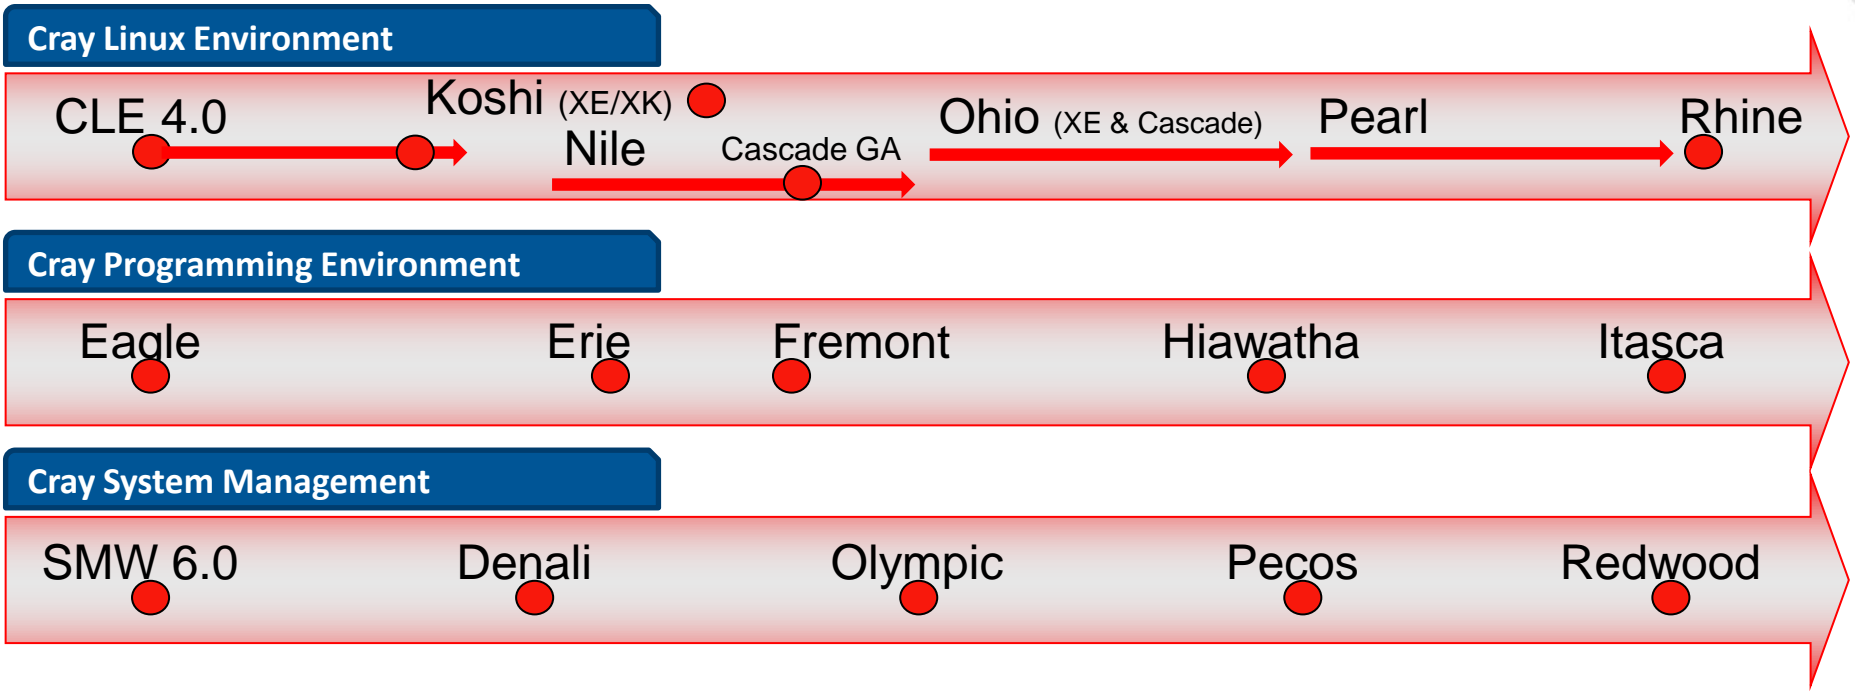
Year	Release	Systems	Cabinets
2007	CLE 2.0	3	4
2008	CLE 2.1	6	118
2009	CLE 2.2	37	262
2010	CLE 3.1	27	276
2011	CLE 4.0	48	720



Cray Software



2011				2012				2013				2014				2015			
Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4



XE - Gemini

XK - Gemini

Cascade - Aries

CLE Koshi Features

Supports Cray XE/XK systems

- Releases in December 2012

Kernel features

- Compute Unit Affinity
- Compute Node Cleanup (CNCU)
- Faster warm boots

Lustre 1.8.7; Lustre 2.2 client

Application resiliency features

Lightweight Log Manager (LLM)

Accelerator features

- Nvidia Kepler support
- Soft GPU reset
- GPU memory scrub

CCM/IAA improvements

CLE Nile Features

Supports Cray Cascade systems

- General Availability (GA) in March 2013
- Based on Intel processors and Cray's new Aries interconnect

Kernel features

- Compute Unit Affinity (Intel Hyperthreads)
- Compute Node Cleanup (CNCU)

Lustre 2.2 client

Application resiliency features

Lightweight Log Manager (LLM)

Aries HSN (high-speed network) features

- Deadlock Avoidance (DLA)
- Aries collective engine

CCM/IAA improvements

Power management features

Application Resiliency: Two Steps

Application Relaunch

- Current behavior: node dies -> application dies
- New behavior
 - Node dies
 - Application torn down
 - If flag set, relaunch the job
 - App restarts from application's checkpoint file
- No need to wait again in the job queue

ALPS Reconnect

- Current behavior: node dies -> application dies
- New behavior
 - Node dies
 - If flag set, ALPS rebuilds its communication tree
 - Passes new info to PMI, which rebuilds its communication tree
 - PMI passes failure info to the programming model
 - CHARM++ will be the initial programming model

Cascade Power Management

Enhanced power monitoring

- Take advantage of power features in Intel and Aries chips

Job power profiling

- Tie power usage data to job data

Static system power capping

- Use case: data center has a hard power limit

P-state control at job launch

- Run a particular job in a particular p-state

Idle node power conservation

- If not needed for a while, turn it off
- Requires involvement of the batch system software

Lustre Road Map (in brief)

Cray XE/XK systems with direct-attached Lustre

- Customers today are running a variety of CLE levels
- Can stay or move to CLE 4.0 UP03 (with Lustre 1.8.6)
- In December 2012 can move to Koshi (with Lustre 1.8.7)
- Patch support for CLE 4.0 is available through mid-2013
- Patch support for Koshi is available through mid-2014

Cray XE/XK systems with external Lustre file systems (esFS)

- Can run any CLE release through and including Koshi in December 2012
- Can run later versions of Lustre by:
 - Upgrading their esFS to ESF running Lustre 2.x
 - Upgrading their Cray XE/XK to Koshi and running its Lustre 2.x client

Cray XE/XK systems with Sonexion devices

- Can run any release from CLE 4.0 UP03 through the last Ohio update in 2014

External Servers

Three new products

- EsLogin
- EsFS (external Lustre servers)
- EsMS (management server for Bright Cluster Manager)

Synchronicity

- EsLogin releases are tied to CLE (through OS levels)
- EsFS releases are tied to Lustre releases
- EsMS releases are tied to BCM releases

Transition to new products

- EsLogin and EsFS releases in December with Koshi
- New shipments in 2H12 will use new products

Upcoming CUG Events

Mon 1:30 in Hamburg

- Reliability and Resiliency of Cray XE6 and XK6 Systems

Mon 3:00 in Köln

- Getting Up and Running with Cluster Compatibility Mode (CCM)

Tue 1:00 in Köln

- Entire track is devoted to CCM

Tue 1:00 in Hamburg

- Cray's Lustre Support Model and Road Map

Tue 3:30 in Hamburg

- Minimizing Lustre Ping Effects at Scale

Thu 1:00 in Köln

- The Year in Review in Cray Security

Thu 2:00 in Hamburg

- Node Health Checker

OSIO Road Map Summary

Koshi will release in December 2012

- XE/XK systems

Nile will release in March 2013

- Cascade systems

Three new products added for external servers

- ESL, ESF and ESM

Ohio is planned for 2H13 and 1H14

Features: application resiliency and power management

Thank You



Dynamic Linking

Make it possible for Cray customers to link dynamically (done)

Improve our implementation of dynamic linking (2011)

- Name and install files in standard Linux ways
 - Improve application startup time
- Static or dynamic can be configured as default site-by-site

Make dynamic linking the default in CCE in mid-2013

No longer ship static libraries (TBD)

Power Management Ideas

Power down unused node resources at job startup

- Example: job not using all the cores, then turn off unused cores

In-band power monitoring

Job-based power consumption accounting

Job-based power capping

Power profiling tools

API for low-latency user space performance/power tradeoffs

Auto-tuning for power optimization

Network power scaling and control

Cray I/O Models

