

Cray System Software Stack: Plans and Status

**Charlie Carroll
Rose Olson
Keith Stephenson
May 2013**



Cray System Software

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 and Aries)
- TCP/IP

Third-party extensions

- GPFS
- Batch systems

Operating system services

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

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
- Routing the network

Big Picture Goals

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
- Accelerator 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 4.1, 4.2 Releases

- Scaled to 200+ cabinets
- Compute Node Clean-Up
- Zone Movable (prevents memory fragmentation)
- Internal login node OOM protection
- DVS enhancements

XC30 Launch

- Intel processor support
- CLE 5.0 / SMW 7.0 release and updates
- Early power management features

Kepler Support

Application Resiliency

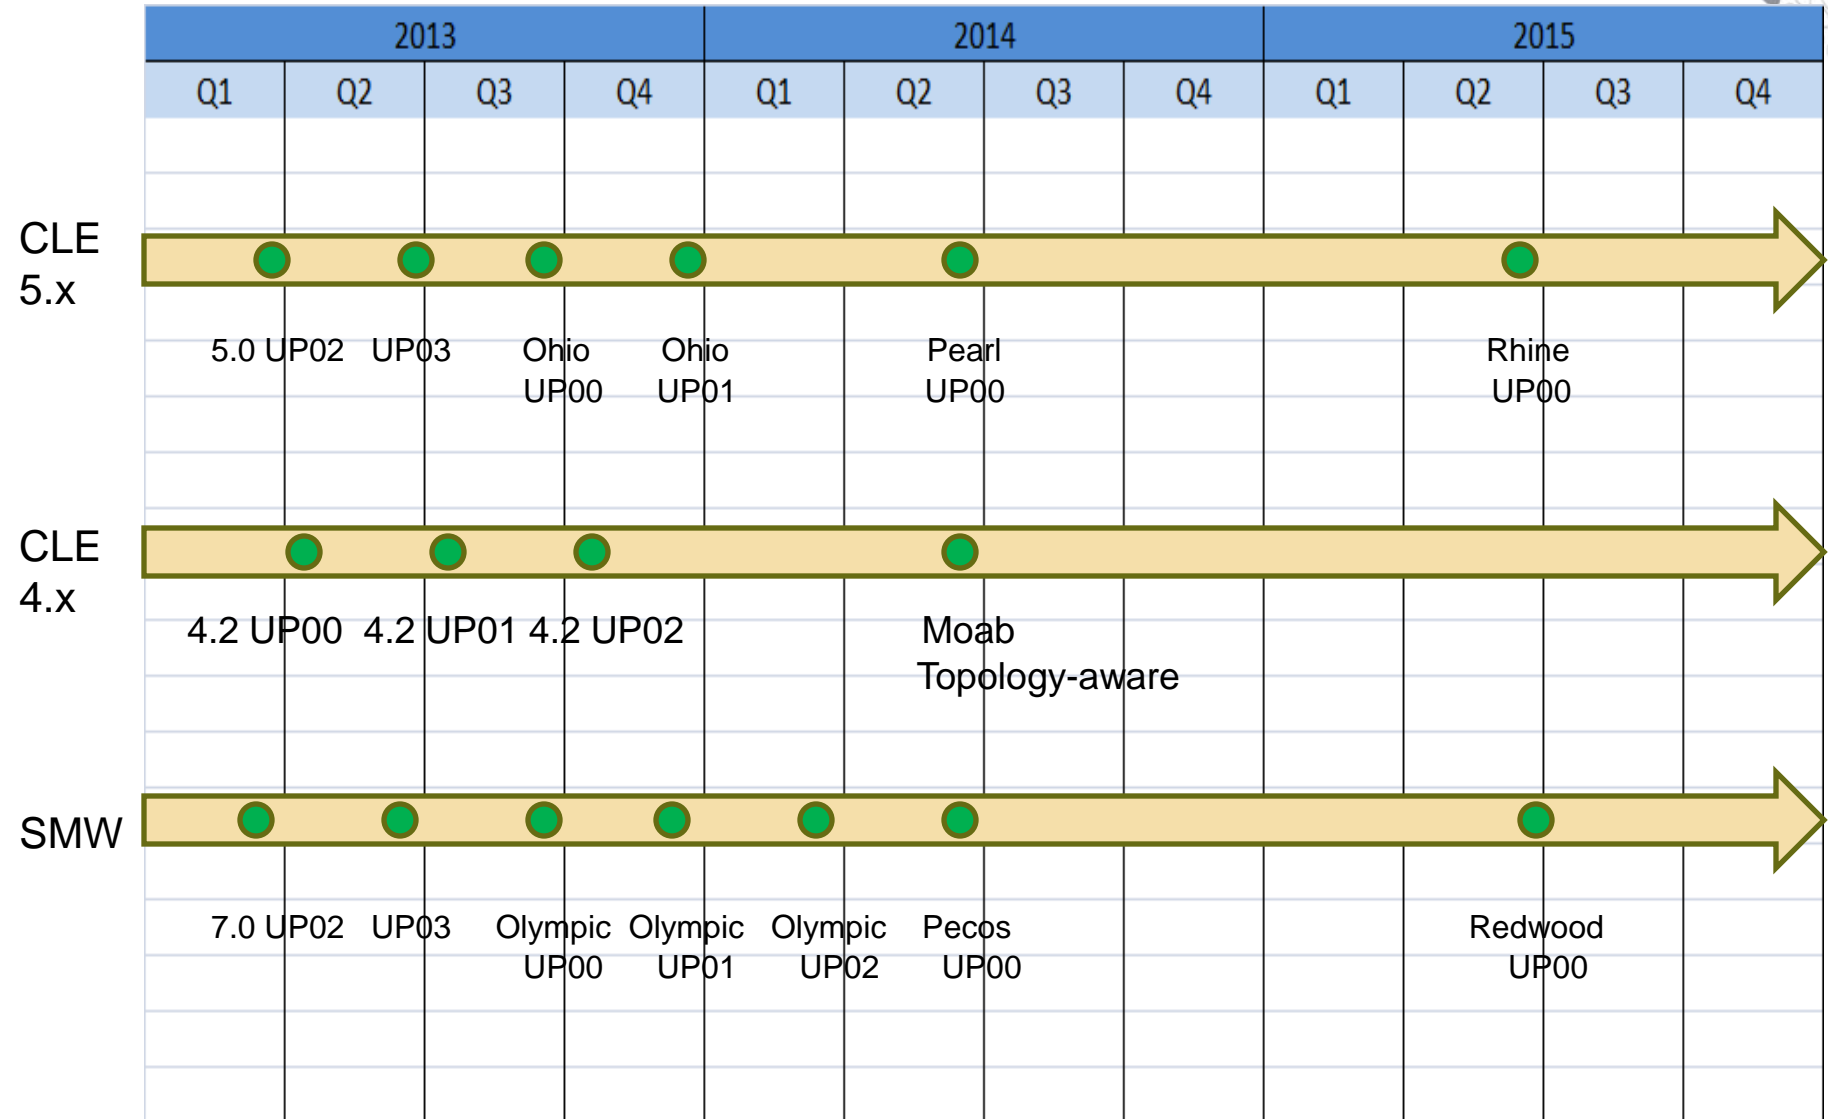
- Application Relaunch
- ALPS Reconnect



Demographics (April 2013)

Year	Release	Systems	Cabinets
2007	CLE 2.0	1	2
2008	CLE 2.1	3	4
2009	CLE 2.2	27	141
2010	CLE 3.1	17	136
2011	CLE 4.0	52	481
2012	CLE 4.1	28	691
2013	CLE 5.0	9	38

System Software Stack Overview





CLE 4.2 UP00 (April 2013) Features

Feature	Benefit
Dynamic SDB	Enables resynching of the SDB and ALPS after a blade warm swap
DVS/DSL performance improvements (Phase 1)	Reduce application start-up times when using dynamic shared libraries
Congestion tools	Diagnose HSN congestion
OFED 1.5.4.1	Support FDR InfiniBand HCAs



CLE 4.2 UP01 (July 2013) Planned Features

Feature	Benefit
Fast Compute Node Reboot	Warmboot nodes without going through Coldstart to clear memory fragmentation
Internal MOM node OOM protection	Imposes limits, kills an app as a set
Controller vitality check*	Monitoring and logging of controller resource (CPU and memory) data
Routing table repair*	Automatic error detection and correction of Gemini routing tables improves system reliability
Support access to L3 cache counters	L3 cache data in Cray performance tools (available in Q3 2013)

* Features in SMW 7.0 UP03, which also releases in July 2013

CLE 4.2 UP02 (October 2013) Planned Features



Feature	Benefit
NVIDIA CUDA 5.5	Driver update with bug fixes and support for GPU Utilization Accounting
GPU Utilization Accounting	Gather and report GPU usage statistics associated with jobs
HSN monitoring	HSN data to the admin; can be correlated with application data
Resource Utilization Reporting (RUR)	Per-application resource utilization statistics, including process accounting and GPU utilization; future releases will include power accounting and replace ACR and CSA
Lustre 2.4	Long-term maintenance release for Lustre; support of new Lustre 2.x features (imperative recovery and wide striping require Lustre 2.2+ server)



Post-CLE 4.2 UP02 Planned Features

Feature	Benefit
Topology-Aware Scheduling (Moab)	Improved application performance by better placement
Compute node metrics	Compute node information such as memory usage

XC30 (Cascade) Systems





CLE 5.0 UP03 (June 2013) Planned Features

Feature	Benefit
Ivy Bridge support	Many XC30 installations in 2013 and beyond
XC30-AC support	Air-cooled counterpart to XC30
DVS / XFS	Lower cost storage option on smaller machines
Lustre 2.3 client features	Support of new Lustre 2.3 features (imperative recovery and wide striping require Lustre 2.2+ server)



Ohio/Olympic (Sep 2013) Planned Features

Feature	Benefit
Kepler support	First Nvidia accelerator supported in XC30
Resource Utilization Reporting (RUR)	Per-application resource utilization statistics, including process accounting and GPU utilization; future releases will include power accounting and replace ACR and CSA
Suspend / Resume, phase 1	Use case: lower-priority is suspended so that higher-priority job can run; both must fit into node memory
Direct-attached Lustre	Another way to connect a Lustre file system to an XC30

Ohio/Olympic UP01 (Dec 2013)

Planned Features



Feature	Benefit
Knights Corner (KNC) support	First Xeon Phi processor in XC30
Atlas support	Atlas is next member of the Kepler family
Suspend / Resume, phase 2	Four jobs per node; better sharing of network resources
Power management features	Monitoring and logging of controller resource (CPU and memory) data
Lustre 2.4.1	Following Lustre road map



Ohio / Olympic: Future Releases

Feature	Benefit
Power management features	See later slides
IMPS: Image Management and Provisioning System	Stick around for John Hesterberg's talk



Cascade Power Management: UP03 (Jun 2013)

Enhanced power monitoring

- Take advantage of power features in Intel and Aries chips

In-band node power/energy/cap performance counters

- Enables tools to have node-level access to power 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

Power database on SMW

- Out-of-band job power profiling using data in power database



Cascade Power Management: Futures

Monitoring and capping for Kepler and KNC

- Apply existing infrastructure to newer blade types

In-band energy reporting

- Provided by RUR (Resource Utilization Reporting)

Off-SMW power database alternative

- Helpful to large sites or anyone wanting to do heavy post-processing

WLM-based idle node power on/off

- Save power when nodes are idle for extended periods

Job-based power capping by WLM

- One use case: charge based on power cap reservation



Upcoming CUG Events in System Software

Tuesday 1:45 in Zinfandel / Cabernet

- Directions in Installation, Image Management, Provisioning & Config (John H)

Tuesday 3:30 in Zinfandel / Cabernet

- The Evolution of Cray Management Services (Tara Fly)

Tuesday 4:30 in Zinfandel / Cabernet

- Cray External Services Systems Overview (Harold Longley)

Tuesday 5:15 in Atlas Peak

- BOF: Cray System Management Futures (John Hesterberg)

Wednesday 12:00 in Atlas Peak

- BOF: Cray External System Services (Jeff Keopp)

Thursday 11:00 in Zinfandel / Cabernet

- Resource Utilization Reporting on Cray Systems (Andrew Barry)

Thursday 12:00 in Zinfandel / Cabernet

- BOF: System Monitoring, Accounting and Metrics (John Hesterberg)

Thursday 3:30 in Zinfandel / Cabernet

- Evaluating Node Orderings for Improved Compactness (Carl Albing)

Thank You

