

Cray Operating Systems Road Map

Charlie Carroll, Cray Inc.
Rose Olson, Cray Inc.
Keith Stephenson, Cray Inc.

ABSTRACT: *This paper discusses the Cray Operating Systems road map. The CLE 4.x and CLE 5.x release streams are discussed. In addition, the rationale for coming changes is discussed.*

KEYWORDS: Operating systems, releases

1. Introduction

The Cray Software Operating Systems and I/O (OSIO) group provides key infrastructure and service components of the software stack.

These components include:

- Compute node kernels
 - XT CNL (Compute Node Linux)
 - NVIDIA GPU drivers
- Service node kernel
 - Supports all compute node types
- File systems
 - Lustre
- Networking
 - GNI and DMAPP
 - IBGNI (IB verbs -> Gemini and Aries)
 - TCP/IP
- Operating system services
 - Core specialization
 - Dynamic shared libraries
 - Cluster Compatibility Mode
 - DVS (Data Virtualization Service)
 - XFS
- System management
 - Cray Management Services
 - Node Health Checker
 - ALPS (Application Level Placement Scheduler)
 - Command interface

- Third-Party Extensions
 - GPFS
 - Batch schedulers

This paper discusses the main themes to be emphasized in upcoming OSIO software releases, followed by specific features to be delivered in these releases.

2. Release Themes

Upcoming OSIO releases will emphasize certain broad themes. Before getting into specifics, we will take a look at the big picture.

2.1. System stability

Stability and robustness are important to any customer system. They are especially important in large supercomputers with millions of separate components.

Cray has invested substantially over the past three years in defect reduction. By our internal measures, we've made substantial progress in reducing customer bugs. By the most important measure—system availability—we've made great progress. System availability is well above 99%.

2.2. Performance

Cray's Compute Node Linux (CNL) implementation performs extremely well, largely because we have limited which services and features run on compute nodes.

In addition, current and upcoming work will focus on GPU and accelerator performance improvements. One example is host-initiated GPU-to-GPU transfers across the Cray high-speed network.

2.3. Hardware Support

Part of OSIO's mission is to support new Cray hardware as it becomes available. CLE 5.0 UP03 and SMW 7.0 UP03 (June 2013) will support Intel's Ivy Bridge processors. Ohio and Olympic (September 2013) will support Kepler GPUs in XC30 systems. The first update to these releases will support Intel's Phi processor known as Knights Corner.

2.4. File Systems

Cray supports a variety of I/O models. Much of our installed XE/XK base uses direct-attached Lustre file systems. Direct-attached Lustre (DAL) will be available on XC30 systems by the end of 2013. Cray Software also ships and supports external Lustre file systems. These integrate white-box servers, Lustre server software and storage. Cray supercomputers connect to these external Lustre servers with Infiniband.

Cray now offers its Sonexion Lustre appliance.

Cray supports other file systems through DVS (Data Virtualization Service). DVS projects the file system from Cray service nodes to the Cray compute nodes. The service nodes, in addition to serving DVS, are also clients to the remote, projected file system. This enables applications running on the compute nodes to access, through DVS, the external file system. To date, Cray has used DVS to interface with Panasas, GPFS, XFS and NFS.

3. Upcoming Releases

Cray maintains two release streams: CLE 4.x for XE/XK systems and CLE 5.x for XC30 systems. The SMW releases support both product and development lines. We discuss each CLE release stream in turn.

3.1. CLE 4.2 UP00

CLE 4.2 UP00, recently released in April 2013, has a number of interesting features: Dynamic SDB, performance improvements in DVS, and congestion tools.

Dynamic SDB is a feature which enables nodes and blades which were not physically present in the system at boot time to be dynamically added.

The congestion tools were developed in conjunction with the acceptance at one of our large sites. The tools have proven quite useful, so Cray R&D has turned them into products which are available for all customers.

3.2. CLE 4.2 UP01

CLE 4.2 UP01, planned for release in July 2013, addresses a number of issues: Fast Compute Node Reboot, Internal MOM Node OOM Protection, Controller Vitality Check, and Routing Table Repair.

Fast Compute Node Reboot introduces a shorter process to warm-boot a node, basically by skipping the hardware initialization steps (known as Coldstart).

Internal MOM Node OOM Protection was implemented to prevent rogue applications from overwhelming the internal MOM node. Limits, such as memory, are imposed on all applications. In addition, the OOM Killer, if invoked, can more reliably completely kill the offending application.

Controller Vitality Check is an HSS feature which periodically checks the health of certain HSS components. If, for example, the L0 controller runs out of memory, Controller Vitality Check will report this to syslog.

Routing Table Repair deals with the fact that occasionally the routing table can become corrupted. This feature will check the routing tables periodically, and reload any which might have been corrupted.

3.3. CLE 4.2 UP02

CLE 4.2 UP02, scheduled for release in October 2013, will offer several new features: CUDA 5.5, GPU Utilization, Resource Utilization Reporting (RUR), HSN Monitoring, and the Lustre 2.4 client.

Over the summer of 2013, Nvidia will release CUDA 5.5. We will incorporate that into our CLE 4.2 UP02 release.

GPU Utilization will build on CUDA 5.5, which has features to give visibility into how the GPUs are being used. This feature is a special case of the more general feature: Resource Utilization Reporting (RUR). RUR will report per-application resource utilization statistics, including process accounting and, as mentioned, GPU statistics. Future releases will add power accounting, eventually replacing ACR and CSA.

HSN Monitoring will pull out some low-level HSN data and make it available on the SMW to administrators.

3.4. CLE 4.2 Later Releases

Later releases of CLE 4.2 will include Topology-Aware Scheduling and Compute Node Metrics.

Topology-Aware Scheduling will allow batch schedulers to use topology information to make intelligent decisions about job placement. It's been shown that job placement can significantly affect job runtime.

Compute Node Metrics will use a similar mechanism to HSN Monitoring to bring compute node metrics such as memory usage to the SMW.

3.5. CLE 5.0 UP03

CLE 5.0 UP03, planned for June 2013, will be the first Generally Available release for the Cray XC30 system. The release will include support for Intel's Ivy Bridge processors, as well as the air-cooled cabinet known as XC30-AC.

XFS projected by DVS to the compute node will be available in CLE 5.0 UP03. The Lustre 2.3 client will also be available in this release. Finally, the first set of power management features, notably system-wide power capping, will be released.

3.6. CLE Ohio and Olympic

The first release of CLE Ohio and Olympic will come out in September 2013. The initial release will include support for Nvidia's Kepler GPUs, Resource Utilization Reporting (as discussed with CLE 4.2 UP02), direct-attached Lustre, and Suspend / Resume phase 1.

Suspend / Resume allows a high-priority job to run on top of a lower-priority job. The weather community provides a typical use case. The lower-priority job (perhaps a research run) is running when the high-priority job (perhaps an emergency weather forecast) needs to run. The high-priority job will load (the lower-priority job remains in memory) and run to completion while the lower-priority job is suspended. When the high-priority job completes, it rolls out and the lower-priority job goes back to running.

The first update to Ohio / Olympic will add support for the Intel Phi (code-named Knights Corner), support for Nvidia's Atlas GPU (an uptick in the Kepler family), and

a Lustre 2.4 client. In addition, enhancements to Suspend / Resume will be released.

Updates to Ohio / Olympic will include a set of features around power management. While the power management road map is still being set, it's expected to give the ability to do in-band power monitoring, turn off hardware which is not expected to be used for some time, and manage power on a job-by-job basis.

4. Conclusion

This paper has presented specific features which will be coming in 2013 and 2014 releases of Cray's operating system. In addition, we have discussed the themes and thought processes behind our plans.

Acknowledgments

The authors would like to thank their colleagues and development teams at Cray. Their commitment to producing the world's best supercomputers makes it a pleasure to come to work every day, as well as making this paper possible.

About the Author

Charlie Carroll is Director, Product Management with Cray Inc. Rose Olson is Director, OS and I/O with Cray Inc. Keith Stephenson is Director, SMW with Cray Inc. If you have comments on our road map, we would love to hear from you at charliec@cray.com, olson@cray.com and rkeiths@cray.com.