Cray Linux Environment (CLE) 5.2 Software Release Overview

S–2425–52xx

Contents

Tables

Introduction [1]

This document provides an overview of the Cray Linux Environment (CLE) 5.2 operating system software release and highlights new functionality and changes from previous CLE releases. This document supports Cray XE6, Cray XK6, and Cray XK7 systems.

Chapter 2, Software Enhancements on page 9 and describe changes made since the last release package of the CLE operating system software.

This document does **not** describe hardware, software, or installation of related products, such as the Cray Compiling Environment or products that Cray does not provide. To determine the release levels of other software products that are compatible with CLE 5.2, see Software Requirements on page 17.

1.1 Emphasis for the CLE 5.2 Release

The CLE 5.2 release provides the following key enhancements:

- **ALPS Standardization** — With this release, ALPS adopts Cray standard conventions for building and packaging ALPS. Some file locations and directory paths are changed.
- **Reservation-level Cleanup** — The interaction between ALPS and NHC (Node Health Checker) is enhanced. Cleanup is now a two-stage process in which nodes are checked both after application exit and after batch system node reservations are released.
- **Memory Compaction** — Move physical pages of memory in order to defragment or compact physical memory. This helps in the prevention of and recovery from memory fragmentation.
- **SLES 11 SP3** — The base operating system for CLE 5.2 is now SLES 11 SP3.
- **DAL Upgrade, Lustre 2.5** — Cray supports Direct-Attached Lustre (DAL) on Cray XE6, Cray XK6, and Cray XK7 systems. Service nodes that support DAL use a CentOS™ operating system running on ramdisk (i.e., not the traditional shared root file systems) to support Lustre 2.x service daemons. This will allow sites with legacy internal Lustre servers and Lustre 1.8.x file systems to upgrade to DAL and Lustre 2.5, which includes significant advances in Lustre's underlying design and several new features. DAL nodes will still be able to

communicate with the Cray high-speed and hardware supervisory networks. For more information, see *Installing and Configuring Cray Linux Environment (CLE) Software*.

- **Image Management and Provisioning System (IMPS)** — *IMPS* is a new set of features that change how software is installed, managed, provisioned, booted, and configured. It brings together industry-standard methods used for image management and Cray proprietary technologies to produce a coherent image management and provisioning system. In future release IMPS will support the installation of other Cray products, but for this release only DAL is installed using IMPS. For more information, see *IMPS Guide for DAL Installation*.
- **Cray Data Virtualization Service (Cray DVS)** — The Cray Data Virtualization Service (Cray DVS) is a distributed network service that provides transparent access to file systems residing on the service I/O nodes and remote servers in the data center. Cray DVS provides a service analogous to NFS™ and projects local file systems resident on I/O nodes or remote file servers to compute and service nodes within the Cray system. For more information, see *Introduction to Cray Data Virtualization Service*.
- **Resource Utilization Reporting (RUR)** — An extensible utility introduced to support the collection of system-wide and per-application resource statistics in a simple unified way for users, system administrators, and accounting tools for use with workload managers, RUR supersedes CSA (Comprehensive System Accounting). Mazama and ACR (Application Completion Reporting) are removed in this release.

The System Management Workstation (SMW) 7.2 release also provides several key enhancements that directly impact the operation of your Cray system running the CLE 5.2 release. For more information, see the *README* provided with the SMW Release Software.

1.2 Supported System Configurations

This release (CLE 5.2) supports initial installations on Cray XE6, Cray XK6, and Cray XK7 systems. It also supports an upgrade from the previous CLE 4.2.UP02 release. Upgrade tests were also done from the CLE 4.0.UP03 and CLE 4.1.UP01 releases.

1.3 Description of the CLE 5.2 Release Software

CLE is a Linux-based operating system that runs on Cray systems. The CLE 5.2 release includes Cray's customized version of the SLES 11 SP3 (Service Pack 3) operating system. All software is installed by means of scripts and RPM files. RPMs include related security fixes.

For complete information about the release package, including detailed information about prerequisites for other Cray software products, see Chapter 5, Release Contents on page 17.

1.4 CLE 5.2 Support Policy

Cray continually enhances the Cray Linux Environment (CLE) with new releases and periodically discontinues support for older releases. Our current policy is to support the latest major release of CLE and the previous major release.

The previous major release for Cray XE6 and Cray XK7 systems was CLE 4.2.

- During the 12 months after initial release, CLE software is supported with update packages at approximately three- to six-month intervals, depending on need.
- Cray will provide patches for available critical and urgent bug fixes for a period of 18 months following an initial (generally available) CLE release.
- Beyond 18 months, support is limited to critical fixes on a best-effort basis.

All applicable recommended and security-related SUSE Linux updates released by Novell are included in the CLE releases and update packages. Security-related patches are also available through Field Notices (FNs).

Contact your Cray representative for information about current software availability and release schedules.

Software Enhancements [2]

This chapter describes the software enhancements made to the Cray Linux Environment (CLE) since the CLE 4.2.UP02 release package.

For information about issues that you may encounter when using, installing or maintaining CLE 5.2 (when compared to previous CLE releases), see Chapter 3, Compatibilities and Differences on page 11.

In addition to the documentation noted in each feature description, see Cray-developed Books Provided with This Release on page 14.

2.1 Software Enhancements in CLE 5.2

2.1.1 ALPS Enhancements

ALPS now supports:

- `aprun` node exclusion: Provides the ability to exclude a node or node lists
- 32-bit Reservation IDs: Widens the range of number the of Reservation IDs from 4096 to 2^32.
- Reservation-level cleanup: ALPS supports reservation data tracking and removal, and works with Node Health Checker to guarantee that nodes with orphaned reservations are not made available to future reservations and applications

2.1.2 `xt-asyncpe` Superseded

With the release of CLE 5.2.UP00, the compiler driver script package module `xt-asyncpe` is superseded by the module `craype`, release 2.1.0 or later. For more information about how this change affects compiler behavior, enter the following command.

```
> module help craype
```

2.2 DVS Enhancements

Cray Data Virtualization Service (DVS) now features improved I/O, additional error recovery, better failover support, improved performance, and decreased application launch time. Aside from the improved performance, this change is transparent to users and requires no administrative or configuration changes.

2.3 Resource Utilization Reporting (RUR)

Resource Utilization Reporting (RUR) is an administrator tool for gathering statistics on how system resources are being used by applications. RUR is a low-noise, scalable infrastructure that collects compute node statistics before an application runs and again after it completes.

The RUR *taskstats* plugin includes the option to collect extended process accounting information similar to that which had been collected by, the now deprecated, Cray System Accounting (CSA). All data elements within CSA's *pacct* record, with the exception of system billing units, are collected when the extended process accounting feature of *taskstats* is enabled.

For more information, see *Managing System Software for the Cray Linux Environment*.

2.4 Audit Support

This release includes an updated audit 1.8.0 set of RPMs. These RPMs replace audit-1.2.8 provided in previous versions of CLE.

Note: Audit 1.8.0 is installed on the boot node and shared root, but not into the compute node image. To add auditing support to a compute node (CNL) image, contact Cray for assistance.

The `cray-audit` RPMs are removed from this release. The `cray-audit` RPMs provided tools specific to the 2.6.16 kernel. No replacement is intended at the current time.

Compatibilities and Differences [3]

This chapter compares Cray Linux Environment (CLE) 5.2 to CLE 4.2 and lists compatibility issues and functionality changes. The *CLE 5.2 Limitations* document describes temporary limitations of the release. The *CLE 5.2 Errata* document describes any installation and configuration changes identified after documentation for this release was packaged. It also includes a list of customer-filed critical and urgent bug reports that are closed with this release. These documents are included with this release package and also available from your Cray representative.

3.1 Binary Compatibility

Applications and binaries compiled for platforms other than Cray XE6, Cray XK6, and Cray XK7 systems may need to be recompiled.

Binaries compiled for Cray XE6, Cray XK6, and Cray XK7 systems running earlier releases of CLE should run on Cray XE6, Cray XK6, and Cray XK7 systems running CLE 5.2 without needing to be recompiled, provided the binaries are dynamically linked.

However, statically linked binaries which are direct or indirect consumers of the network interface libraries (uGNI/DMAPP) must be relinked. This is because the DMAPP and uGNI libraries are tied to specific kernel versions and no backward or forward compatibility is provided. This includes applications that use the MPI or SHMEM libraries, as well as applications using the PGAS (Partitioned Global Address Space) languages such as UPC, Fortran with Coarrays, and Chapel.

3.2 Installation and Configuration Functionality Changes for System Administrators

For detailed initial and upgrade installation procedures, see *Installing and Configuring Cray Linux Environment (CLE) Software*.

For temporary limitations of this release and changes identified after the documentation for this release was packaged, see the *CLE 5.2 Limitations* document provided with the release package. Additional information may be included in the *CLE 5.2 Errata* and *CLE 5.2 README* document provided with release packages.

3.2.1 Supported Installation Path

The CLE 5.2 release supports initial and upgrade installations.

3.2.2 System Management Workstation (SMW) Upgrade Requirements

You **must** install or upgrade the System Management Workstation (SMW) to the SMW 7.2 release before you install or upgrade to CLE 5.2.

For information about the content of the SMW 7.2 release, see the *SMW README* document included in the SMW 7.2 release package.

3.2.3 Mazama and ACR Software No Longer Provided

With the release of RUR, Mazama (also known as *Cray Management Services (CMS)*) and Application Completion Reporting (ACR) are no longer provided, as the resource reporting performed by RUR replaces this functionality. Users of these features should migrate to RUR. See *Managing System Software for the Cray Linux Environment* for more information on how to configure and interact with RUR.

Comprehensive System Accounting (CSA) is deprecated and will be removed in a future release.

Documentation [4]

This chapter describes the documentation that supports the Cray Linux Environment (CLE) 5.2 release.

4.1 Accessing Product Documentation

With each software release, Cray provides books and man pages, and in some cases, third-party documentation. These documents are provided in the following ways:

CrayPort

CrayPort is the external Cray website for registered users that offers documentation for each product. CrayPort has portal pages for each product that contains links to all of the documents that are associated to that product. CrayPort enables you to quickly access and search Cray books, man pages, and in some cases, third-party documentation. You access CrayPort by using the following URL:

http://crayport.cray.com

CrayDoc

CrayDoc is the Cray documentation delivery system. CrayDoc enables you to quickly access and search Cray books, man pages, and in some cases, third-party documentation. Access the HTML and PDF documentation via CrayDoc at the following locations.

- The local network location defined by your system administrator
- The CrayDoc public website: http://docs.cray.com

Man pages

Man pages are textual help files available from the command line on Cray machines. To access man pages, enter the `man` command followed by the name of the man page. For more information about man pages, see the `man`(1) man page by entering:

```
% man man
```

Third-party documentation

Third-party documentation that is not provided through CrayPort or CrayDoc is included with the third-party product.

4.2 Cray-developed Books Provided with This Release

The books provided with this release are listed in Table 1, which also indicates whether each book was updated. Books are provided in HTML and PDF formats.

Table 1. Books Provided with This Release

Book Title	Most Recent Document	Updated
Cray Linux Environment (CLE) Software Release Overview (this document)	S–2425–52	Yes
Installing and Configuring Cray Linux Environment (CLE) Software	S–2444–52	Yes
Managing System Software for the Cray Linux Environment	S–2393–52xx	Yes
Managing Lustre for the Cray Linux Environment (CLE)	S–0010–52	Yes
Introduction to Cray Data Virtualization Service	S–0005–52	Yes
Workload Management and Application Placement for the Cray Linux Environment	S–2496–52	Yes
Using the GNI and DMAPP APIs	S–2446–52	Yes
IMPS Guide for DAL Installation	S–0049–52	Yes
Writing a Node Health Checker (NHC) Plugin Test	S–0023–5002	No
Using Compute Unit Affinity on Cray Systems	S–0030–5002	No
Installing CLE Support Package on a Cray Development and Login (CDL) Node	S–2528–52	Yes
Network Resiliency for Cray XE and Cray XK Systems	S–0032–C	No
Modifying Your Application to Avoid Gemini Network Congestion	S–0031–A	No

4.2.1 Additional Cray-developed Release Documents

Two additional documents are provided with the CLE 5.2 release package. These documents are also available from your Cray representative.

CLE 5.2 Limitations

Describes temporary limitations of the release.

CLE 5.2 Errata

Describes any installation and configuration changes that were identified after documentation for this release was packaged; also includes a list of customer-filed critical and urgent bug reports closed with this release.

Contact your Cray representative about CLE-related information addressed in Field Notices (FNs).

4.3 Changes to Man Pages

Updated Linux man pages are included with CLE 5.2. To access Linux man pages, use the `man` command on a login node.

The following Cray man pages are either new or updated:

- `impscli`(8) — New: documents the Image Management and Provisioning System command line interface
- `apbridge`(8), `apinit`(8), `apres`(8), `apsched`(8), `apsys`(8), `apwatch`(8) — Updated: new file paths to configuration and log files resulting from ALPS standardization
- `lustre_control`(8) — Updated: new `-l` option added to the reformat command
- `xthotbackup`(8) — Updated: new `-i` and `-L` options

4.4 Other Related Documents Available

The following publications contain additional information that may be helpful in setting up your Cray system; they are not provided with this release, but are supplied with other products purchased from Cray. You can access these publications from the CrayPort website. You can also order the printed form of release overviews and installation guides from Cray.

Table 2. Other Related Documents Available

Book Title	Number
Installing Cray System Management Workstation (SMW) Software	S–2480–72
Using and Configuring System Environment Data Collections (SEDC)	S–2491–71
Cray Programming Environments Installation Guide	S–2372
Cray Compiling Environment Release Overview and Installation Guide	S–5212

4.5 Additional Documentation Resources

Table 3 lists additional resources for obtaining documentation not included with this release package.

Table 3. Additional Documentation Resources

Product	Documentation Source
Linux	Documentation for SLES and Linux is at http://www.novell.com/linux, and documentation for the Linux Documentation Project is at http://www.tldp.org.
Lustre	Additional Lustre documentation is available at http://wiki.whamcloud.com/display/PUB/Documentation.
RPM	RPM documentation is available at http://www.rpm.org.
PBS Professional	Documentation for the PBS Professional Workload Manager software is available from Altair Engineering, Inc. at http://www.altair.com.
Moab with TORQUE	Documentation for Moab Workload Manager and TORQUE Resource Manager software is available from Adaptive Computing: http://www.adaptivecomputing.com/.
SLURM	Documentation for SLURM job scheduling and resource management software is available from SchedMD: http://slurm.schedmd.com/cray.html.
IBM Platform LSF	Documentation for Platform LSF workload management system software is available at: http://www-947.ibm.com/support/entry/portal/documentation_expanded_list/software/platform_computing/platform_lsf.

Release Contents [5]

5.1 Hardware Requirements

The Cray Linux Environment (CLE) 5.2 release supports initial and upgrade installations on Cray XE6, Cray XK6, and Cray XK7 systems.

5.2 Software Requirements

The following sections list the required or recommended release levels for products that run on Cray systems but are released separately from CLE 5.2.

5.2.1 Release Level Requirements for Other Cray Software Products

The product versions listed in Table 4 are the minimum release level required for verified compatibility with CLE 5.2. Support for these products is provided in the form of updates to the latest released version only. Unless otherwise noted in the associated release documentation, Cray recommends that you continue to upgrade these releases as updates become available.

Table 4. Minimum Release Level Requirements for Other Cray Software Products with CLE 5.2

Product	Minimum Release Level	Release Information
System Management Workstation (SMW)	Release 7.2	*SMW README*
Cray Application Developer's Environment (CADE)	Release 6.30 or later	Cray XE and Cray XK systems
Cray Message Passing Toolkit (MPT)	Release 6.2.0 or later	Cray XE and Cray XK systems
Cray Performance, Measurement, and Analysis Tools (CPMAT)	Release 6.1.4	*Cray Performance Measurement and Analysis Tools Release Overview and Installation Guide* (S–2474)
Cray Compiling Environment (CCE)	Release 8.2.5 or later	*Cray Compiling Environment Release Overview and Installation Guide* (S–5212)

5.2.2 Third-party Software Requirements

Third-party compiler products are available for Cray systems as noted in Table 5. The release level indicated has been tested with CLE 5.2. Cray recommends that you continue to upgrade these products as updates become available.

Table 5. Minimum Release Level Requirements for Third-party Compilers with CLE 5.2

Product	Minimum Release Level	Release/Ordering Information
PGI Compiler	Release 13.3 or 13.1 or later.	Contact your Cray representative for licensing/purchasing information. For product information see The Portland Group, Inc.: http://www.pgroup.com
Intel Compiler	Release 14.0 or later.	Intel Corporation. See: http://software.intel.com
GCC (GNU Compiler Collection)	Release 4.8.0 or later.	The GNU Project. See: http://gcc.gnu.org

Batch system software products are available for Cray systems as indicated in Table 6. Information regarding supported and certified batch system software release levels is available on the CrayPort website at http://crayport.cray.com. Click on **3rd Party Batch SW** in the menu bar.

Table 6. Third-party Batch System Software Products Available for Cray Systems

Product	Minimum Release Level	Release/Ordering Information
Moab and TORQUE:	Moab Version 7.2.4.1 or later. TORQUE 4.3.4 or later.	Contact your Cray representative for licensing/purchasing information. For product information see Adaptive Computing: http://www.adaptivecomputing.com/
PBS Professional:	Release 12.1 or later.	Contact your Cray representative for licensing/purchasing information. For product information see Altair Engineering, Inc.: http://www.altair.com/

5.3 Supported Upgrade Path

The CLE 5.2 .UP00 release supports initial installations. It also supports an upgrade from the previous CLE 4.2.UP02 releases. Upgrade tests were also done from the CLE 4.0.UP03 and CLE 4.1.UP01 releases.

5.4 Contents of the Release Package

The release package includes:

- All necessary RPMs and installation utilities for the components listed in CLE 5.2 Software Components on page 19
- A printed copy of this release overview
- A printed copy of *Installing and Configuring Cray Linux Environment (CLE) Software*
- A printed copy of *CLE 5.2 Limitations*
- A printed copy of *CLE 5.2 Errata*
- A printed copy of *CLE 5.2 README*

5.4.1 CLE 5.2 Software Components

The CLE 5.2 release includes, but is not limited to, the following system software products:

- Cray's customized version of the SLES 11 SP3 operating system
- CNL compute node operating system
- CentOS base (for DAL)
- Lustre file system (Client Based on Version 2.5.0 from Whamcloud)
- Application Level Placement Scheduler (ALPS)
- Cray Data Virtualization Service (Cray DVS)
- Cluster Compatibility Mode (CCM)
- Resource Utilization Reporting (RUR)
- Cray Lightweight Log Management (LLM) System
- Dynamic Shared Objects and Libraries (DSL)
- Linux `ldump` and `lcrash` utilities
- Node Health Checker (NHC)
- OpenFabrics InfiniBand
- Realm-Specific Internet Protocol (RSIP)

5.5 Licensing

The CLE release is covered under a software license agreement for Cray software. Upgrades to this product are provided only when a software support agreement for this Cray software is in place.

Cray licenses the following as separate products for Cray systems under a Cray license agreement:

- Cray system OS binary (which provides rights to the CLE operating system and its components)

  Note: Source Code Option: The Cray OS licenses for Cray systems is binary by default. Certain U.S. customers may be eligible to obtain a buildable OS source license on Cray systems for an additional fee. For more information regarding source code, contact your sales representative.

For more information about licensing and pricing, contact your Cray sales representative, or send E-mail to crayinfo@cray.com.