



Cray Linux Environment (CLE) Software Release Overview Supplement

S-2497-5201

© 2010–2014 Cray Inc. All Rights Reserved. This document or parts thereof may not be reproduced in any form unless permitted by contract or by written permission of Cray Inc.

U.S. GOVERNMENT RESTRICTED RIGHTS NOTICE

The Computer Software is delivered as "Commercial Computer Software" as defined in DFARS 48 CFR 252.227-7014.

All Computer Software and Computer Software Documentation acquired by or for the U.S. Government is provided with Restricted Rights. Use, duplication or disclosure by the U.S. Government is subject to the restrictions described in FAR 48 CFR 52.227-14 or DFARS 48 CFR 252.227-7014, as applicable.

Technical Data acquired by or for the U.S. Government, if any, is provided with Limited Rights. Use, duplication or disclosure by the U.S. Government is subject to the restrictions described in FAR 48 CFR 52.227-14 or DFARS 48 CFR 252.227-7013, as applicable.

The following are trademarks of Cray Inc. and are registered in the United States and other countries: Cray and design, Sonexion, Urika, and YarcData. The following are trademarks of Cray Inc.: ACE, Apprentice2, Chapel, Cluster Connect, CrayDoc, CrayPat, CrayPort, ECOPhlex, LibSci, NodeKARE, Threadstorm. The following system family marks, and associated model number marks, are trademarks of Cray Inc.: CS, CX, XC, XE, XK, XMT, and XT. The registered trademark Linux is used pursuant to a sublicense from LMI, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis. Other trademarks used in this document are the property of their respective owners.

Aries and Intel Xeon Phi are trademarks of Intel Corporation in the United States and/or other countries. CentOS is a trademark of the CentOS project. MySQL is a trademark of Oracle and/or its affiliates. Lustre is a trademark of Xyratex and/or its affiliates.

Contents

	<i>Page</i>
Software Enhancements [1]	5
1.1 Software Enhancements in CLE 5.2.UP01	5
1.1.1 Lustre Monitoring Tool (LMT)	5
1.1.2 Resource Utilization Reporting (RUR) Enhancements	6
1.1.3 DVS Periodic Sync	6
1.1.4 Node Health Checker (NHC) Enhancements	7
1.1.5 High-speed Network (HSN) Enhancements	7
1.1.6 HSS Health Validation	7
1.1.7 Suspend/Resume Enhancements	8
1.2 Bugs Addressed Since the Last Release	8
1.3 Compatibilities and Differences	8
Support Requirements [2]	9
2.1 Supported Cray System Hardware Platforms	9
2.2 Supported Software Upgrade Path	9
2.2.1 System Management Workstation (SMW) Requirements	9
2.3 Binary Compatibility	9
2.4 Additional Software Requirements	9
2.4.1 Release Level Requirements for Other Cray Software Products	9
2.4.2 Third-party Software Requirements	10
Documentation [3]	11
3.1 Cray-developed Books Provided with This Release	11
Tables	
Table 1. Books Provided with This Release	11

Software Enhancements [1]

Cray Linux Environment (CLE) software update release packages provide bug fixes and a limited set of software enhancements or features. This chapter provides an overview of software enhancements that are introduced in each update release.

Software enhancements and features that were introduced with the initial or base CLE 5.2 release are described in *Cray Linux Environment (CLE) Software Release Overview*.

1.1 Software Enhancements in CLE 5.2.UP01

1.1.1 Lustre Monitoring Tool (LMT)

The Lustre Monitoring Tool (LMT), developed by Lawrence Livermore National Lab, is now available on direct-attached Lustre (DAL). LMT monitors Lustre file system servers (MDT, OST, and LNET routers). It collects data using the Cerebro monitoring system and stores it in a MySQL database. Live data is viewed on the LMT server node with a `top`-like utility called `ltop`. The data in the LMT MySQL database can be viewed with `lmtsh` or with scripts that access and present the data. Lustre clients are not monitored.

Configuration and startup of Cerebro and LMT are handled through the IMPS Configurator during installation or update of DAL. The procedure to configure a MySQL database for LMT is described in *Installing and Configuring Cray Linux Environment (CLE) Software*.

Further information can be found in the man pages: `cerebro(7)`, `cerebro.conf(5)`, `cerebro_module(3)`, `cerebro_module_devel(3)`, `cerebro-admin(8)`, `cerebrod(8)`, `cerebro-stat(8)`, `lmt.conf(5)`, `lmt_agg.cron(8)`, `lmtinit(5)`, `lmtsh(8)`, and `ltop(1)`.

1.1.2 Resource Utilization Reporting (RUR) Enhancements

Several enhancements are included in this release of RUR, including new data plugins:

- `kncstats` collects process accounting data from Intel Xeon Phi (KNC) coprocessors
- `taskstats` collects process accounting data
- `memory` collects data from `/proc` and `/sys` that is useful when assessing an application's memory performance

New options:

- `json-dict` option added to the `energy` plugin. This option extends the amount of data reported and formats the outputted data in JSON dictionary format.
- `json-dict` option added to the `taskstats` plugin. This option formats the outputted data in JSON dictionary format.
- When the `user` output plugin is enabled by the administrator, a user can choose whether or not RUR data is written to a file owned by the user (by creating a file that indicates their preference).

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

1.1.3 DVS Periodic Sync

DVS periodic sync improves data resiliency and facilitates a degree of application resiliency so that they may continue executing in the event of a DVS server failure. Prior to periodic sync, when a DVS server failed, DVS on client nodes would unconditionally kill all processes with any open file descriptors that had written data to that server. This was done since the data written to the file system may have been in the page cache on the DVS server rather than on disk, and DVS has no knowledge of exactly when data is written to backing store. Periodic sync improves this by:

- periodically syncing dirty pages.
- tracking each time a server performed the sync and the last time written for each file.
- only killing processes that performed a `write()` after the most recent sync.

You can tune periodic sync behavior using two `/proc` files. Also, DVS `closesync` mount option is obviated by this behavior and periodic sync is more efficient since it is aware of the which files are dirty.

For more information, see *Introduction to Cray Data Virtualization Service* and the `dvs(5)` man page.

1.1.4 Node Health Checker (NHC) Enhancements

The Node Health Checker (NHC) in CLE performs specified tests to determine if compute nodes allocated to an application are healthy enough to support running subsequent applications. These tests are enabled by various plug-in scripts. These scripts run according to the node health configuration, typically after an abnormal termination of a user application, as evidenced by a crash or a non-zero exit code.

A CCM plugin test has been added in this release, to validate the cleanup of a cluster compatibility mode (CCM) environment at the end of a reservation. This test runs on a compute node only when `/var/crayccm` is detected. For more information, see *Managing System Software for the Cray Linux Environment*.

Also in this release, Cray DVS provides a script that outputs information on the oldest outstanding client request in the `/proc/fs/dvs/ipc/requests` file, including the `cname` of the DVS server processing the request. The `dvs_request` plug-in is included by default on systems that have an initial install (as opposed to an update) of CLE 5.2.UP00. If the plug-in is not on your system you can enable it by editing the `/etc/opt/cray/nodehealth/nodehealth.conf` file.

For more information, see *Introduction to Cray Data Virtualization Service* and *Managing System Software for the Cray Linux Environment*.

1.1.5 High-speed Network (HSN) Enhancements

To specify a high-speed network cable to add or remove from service, use the `xtwarmswap --add-cable cable` command or the `xtwarmswap --remove-cable cable` command, respectively. These options provide the ability to replace a single cable without removing blades or shutting down the system.

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

1.1.6 HSS Health Validation

The `xtcheckhss` command initiates a series of tests that validate the health of the HSS. The tests gather and display information supplied by scripts located on blade controllers (BCs) and cabinet controllers (CCs). For a list of the `xtcheckhss` available tests, see the `xtcheckhss(8)` man page.

1.1.7 Suspend/Resume Enhancements

Suspend/resume enables applications to oversubscribe compute node CPUs; BASIL provides a switching mechanism that enables one job to suspend another job on the node. Previously a maximum of two co-resident jobs were permitted; with this release, this limit is raised to four, and suspend/resume is enabled and configured by entries in the `alps.conf` file. For more information, see *Managing System Software for the Cray Linux Environment*.

1.2 Bugs Addressed Since the Last Release

The list of customer-filed bug reports that were closed with CLE 5.2 releases is included in the *CLE 5.2 Errata* specific to this release package.

1.3 Compatibilities and Differences

The *README* document that is included with the release package describes compatibility issues and functionality changes that you should be aware of before you install a CLE 5.2 update release on a Cray system that was running an earlier version of the CLE 5.2 release.

The *README* document also includes additional documentation or changes to the documentation identified after the documentation for this release was packaged.

Support Requirements [2]

2.1 Supported Cray System Hardware Platforms

The CLE 5.2.UP01 update release supports Cray XC series systems.

2.2 Supported Software Upgrade Path

The CLE 5.2.UP01 release supports initial system installations and upgrade installations from CLE 5.1 and CLE 5.2 release packages.

2.2.1 System Management Workstation (SMW) Requirements

You must be running the SMW 7.2.UP01 release or later before you install the CLE 5.2.UP01 update release package. For additional information, see the SMW *README* document included with the SMW release package.

2.3 Binary Compatibility

The language in the binary compatibility statement in *Cray Linux Environment (CLE) Software Release Overview* remains accurate.

2.4 Additional Software Requirements

2.4.1 Release Level Requirements for Other Cray Software Products

Note: Upgrading to the latest Cray XC30 Programming Environments Release is recommended. For release information, see the *Cray Programming Environments Installation Guide*, the *Cray Programming Environment User's Guide*, and the release notes.

Support for other Cray software 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.

2.4.2 Third-party Software Requirements

The *Cray Linux Environment (CLE) Software Release Overview* includes a section that lists third-party software requirements for the CLE 5.2 release. This information applies to CLE 5.2 update packages, with the following exceptions:

- Updated 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.

3.1 Cray-developed Books Provided with This Release

[Table 1](#) lists the books provided with the CLE 5.2.UP01 release and indicates which books are new or revised with this update release. The most recent version of each book is provided with the release package.

For information about additional documentation resources and accessing documentation, see *Cray Linux Environment (CLE) Software Release Overview* (S-2425-52), which is also provided with the release package.

Table 1. Books Provided with This Release

Book Title	Most Recent Document	Updated
<i>Cray Linux Environment (CLE) Software Release Overview Supplement</i> (this document)	S-2497-5201	Yes
<i>Cray Linux Environment (CLE) Software Release Overview</i>	S-2425-52	No
<i>Installing and Configuring Cray Linux Environment (CLE) Software</i>	S-2444-5201	Yes
<i>Managing System Software for the Cray Linux Environment</i>	S-2393-5201	Yes
<i>Managing Lustre for the Cray Linux Environment (CLE)</i>	S-0010-52	No
<i>Introduction to Cray Data Virtualization Service</i>	S-0005-5201	Yes
<i>Writing a Node Health Checker (NHC) Plugin Test</i>	S-0023-5002	No
<i>Workload Management and Application Placement for the Cray Linux Environment</i>	S-2496-52	No
<i>Using the GNI and DMAPP APIs</i>	S-2446-52	No
<i>Using Compute Unit Affinity on Cray Systems</i>	S-0030-5002	No
<i>Installing CLE Support Package on a Cray Development and Login (CDL) Node</i>	S-2528-52	No
<i>Installing Native SLURM on Cray XC Systems</i>	S-2538-5201	Yes
<i>IMPS Guide for DAL Installation</i>	S-0049-52	Yes
<i>Network Resiliency for Cray XC30 Systems</i>	S-0041-B	No
<i>Modifying Your Application to Avoid Aries Network Congestion</i>	S-0048-A	No