



Cray Linux Environment (CLE) Software Release Overview Supplement

S-2497-4202

© 2010–2013 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.

Lustre is a trademark of Xyratex and/or its affiliates. UNIX is a trademark of The Open Group.

Contents

	<i>Page</i>
Software Enhancements [1]	5
1.1 Software Enhancements in CLE 4.2.UP02	5
1.1.1 Resource Utilization Reporting (RUR)	5
1.1.1.1 Comprehensive System Accounting (CSA), Mazama, and Application Completion Reporting (ACR) Deprecated in CLE and SMW Releases	5
1.1.2 High-speed Network Metrics	6
1.2 Software Enhancements in CLE 4.2.UP01	6
1.2.1 Batch Process and Memory Limits (Torque Only)	6
1.2.2 Fast Compute Node Reboot	6
1.2.3 ALPS aprun Client Exclusion Lists	7
1.3 Bugs Addressed Since the Last Release	7
1.4 Compatibilities and Differences	7
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
3.2 Changed Cray Man Pages	12
Tables	
Table 1. Books Provided with This Release	11

Software Enhancements [1]

Cray Linux Environment (CLE) 4.2 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 4.2 release are described in *Cray Linux Environment (CLE) Software Release Overview*.

1.1 Software Enhancements in CLE 4.2.UP02

1.1.1 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.

Initially, RUR will be coordinated by ALPS but will also be made available to be used with third-party workload management systems. The extensible RUR infrastructure allows plugins to be easily written to collect data uniquely interesting to each site administrator. Cray supplied plugins collect process accounting data, energy usage data, and GPU accounting data.

1.1.1.1 Comprehensive System Accounting (CSA), Mazama, and Application Completion Reporting (ACR) Deprecated in CLE and SMW Releases

With the release of RUR, Comprehensive System Accounting (CSA), Mazama (also known as *Cray Management Services (CMS)*), and Application Completion Reporting (ACR) are deprecated as the resource reporting performed by RUR is intended to replace the functionality previously found in those software features. These features will remain in Cray software but will be removed in future SMW and CLE software release packages. 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.

1.1.2 High-speed Network Metrics

The Cray-provided kernel(s) are modified to provide Gemini network metrics on a per-node and per-NIC basis. Userspace programs are also provided to facilitate modification, collection and aggregation of network metrics through third-party tools such as OVIS.

The userspace programs extend the Gemini Performance Counters Daemon (`gpcdr`) functionality and are called: `gpcdr-init` and `gpcdr-ctl`. `gpcdr-init` configures the metrics as specified in the corresponding configuration file, `gpcdr-init.conf`. `gpcdr-ctl` allows administrators to experiment with metrics in an *ad hoc* fashion post-boot. `rca-helper` also has a new `-O` option which produces the tile-to-link mapping used by the `l0sysd` daemon on each node.

For more information, see *Managing System Software for the Cray Linux Environment* and the `gpcdr-ctl(8)` man page included with this release package as a PDF.

1.2 Software Enhancements in CLE 4.2.UP01

1.2.1 Batch Process and Memory Limits (Torque Only)

CLE provides a new daemon for Torque MOM nodes called `batchlimitd`. It is used in order to avoid a potentially harmful proliferation of processes (e.g., *fork bomb*) that some jobs may create, resulting in OOM conditions and unstable or crashed nodes. `batchlimitd` allows administrators to tune these limits in the configuration file.

For more information, see the `batchlimitd` man page included with this release (PDF is included) and the default configuration file located at `/etc/opt/cray/batchlimitd/batchlimitd.conf` on the boot and SDB nodes.

1.2.2 Fast Compute Node Reboot

Fast Compute Node Reboot (FCNR) is an option of the `xtbootsys` command that reboots compute nodes without bouncing them, reducing the time needed for reboots. FCNR addresses a specific problem: when memory fragmentation reaches a level that prevents allocation of huge pages. Do not use FCNR as a general method of rebooting nodes.

The intended means of calling FCNR is through the configurations of Node Health Checker and the SMW daemon dumpd. The `xtbootsys` command syntax for use of FCNR specifies the `--fast` option:

```
% xtbootsys --reboot --fast [node_list]
```

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

1.2.3 ALPS aprun Client Exclusion Lists

The `aprun` client command is modified to allow users to exclude nodes on a per-application instance basis using the `-E`, `--exclude-node-list` or `--exclude-node-list-file` options. In very rare cases, it is useful to have the option to exclude a certain node or list of nodes from application placement. For example, a node has not yet been marked into a DOWN state but otherwise known as "bad." See the `aprun(1)` man page for usage information.

1.3 Bugs Addressed Since the Last Release

The list of customer-filed bug reports that were closed with CLE 4.2 releases is included in the *CLE 4.2 Errata* specific to this release package.

1.4 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 4.2 update release on a Cray system that was running an earlier version of the CLE 4.2 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 4.2.UP02 update release supports Cray XE and Cray XK systems.

2.2 Supported Software Upgrade Path

The CLE 4.2.UP02 release supports initial system installations and upgrade installations from CLE 4.2.UP01, 4.2.UP00, and CLE 4.1 release packages.

2.2.1 System Management Workstation (SMW) Requirements

You must be running the SMW 7.1 release or later before you install the CLE 4.2.UP02 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 XE and Cray XK Programming Environments Release is recommended as Perftools 6.1.2 and `libsci_acc 3.0.00` is required for this CLE release package. 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

Cray Linux Environment (CLE) Software Release Overview includes a section that lists third-party software requirements for the CLE 4.2 release; this information applies to CLE 4.2 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 4.2.UP02 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-2425-42), 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-2497-4202	Yes
<i>Cray Linux Environment (CLE) Software Release Overview</i>	S-2425-42	No
<i>Installing and Configuring Cray Linux Environment (CLE) Software</i>	S-2444-4202	Yes
<i>Managing System Software for the Cray Linux Environment</i>	S-2393-4202	Yes
<i>Managing Lustre for the Cray Linux Environment (CLE)</i>	S-0010-4201	No
<i>Introduction to Cray Data Virtualization Service</i>	S-0005-51	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-2496-42	No
<i>Using the GNI and DMAPP APIs</i>	S-2446-42	No
<i>Using Compute Unit Affinity on Cray Systems</i>	S-0030-5002	No
<i>Using Balanced Injection in Cray Systems</i>	S-0040-A	No
<i>Installing CLE Support Package on a Cray Development and Login (CDL) Node</i>	S-2528-4202	Yes

3.2 Changed Cray Man Pages

- `intro_LLM_logfiles(5)`: updated to reflect that some messages from the state manager are included in LLM log files.
- `xtlog(8)`: Corrected example error.