



Cray Linux Environment™ (CLE) Software Release Overview Supplement

S-2497-5003

© 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.

Cray and Sonexion are federally registered trademarks and Active Manager, Cascade, Cray Apprentice2, Cray Apprentice2 Desktop, Cray C++ Compiling System, Cray CS300, Cray CX, Cray CX1, Cray CX1-iWS, Cray CX1-LC, Cray CX1000, Cray CX1000-C, Cray CX1000-G, Cray CX1000-S, Cray CX1000-SC, Cray CX1000-SM, Cray CX1000-HN, Cray Fortran Compiler, Cray Linux Environment, Cray SHMEM, Cray X1, Cray X1E, Cray X2, Cray XC30, Cray XD1, Cray XE, Cray XEm, Cray XE5, Cray XE5m, Cray XE6, Cray XE6m, Cray XK6, Cray XK6m, Cray XK7, Cray XMT, Cray XR1, Cray XT, Cray XTm, Cray XT3, Cray XT4, Cray XT5, Cray XT5_h, Cray XT5m, Cray XT6, Cray XT6m, CrayDoc, CrayPort, CRInform, ECOphlex, LibSci, NodeKARE, RapidArray, The Way to Better Science, Threadstorm, uRiKA, UNICOS/lc, and YarcData are trademarks of Cray Inc.

Intel is a trademark of Intel Corporation or its subsidiaries in the United States and other countries. LSI is a trademark of LSI Corporation. Linux is a trademark of Linus Torvalds. Lustre is a trademark of Oracle and/or its affiliates. Moab and TORQUE are trademarks of Adaptive Computing Enterprises, Inc. NVIDIA, CUDA, Tesla are trademarks of NVIDIA Corporation. PBS Professional is a trademark of Altair Grid Technologies. PGI is a trademark of The Portland Group Compiler Technology, STMicroelectronics, Inc. XFS and SGI are trademarks of Silicon Graphics, Inc. UNIX is a trademark of The Open Group. All other trademarks are the property of their respective owners.

Contents

	<i>Page</i>
Software Enhancements [1]	5
1.1 Software Enhancements in CLE 5.0.UP03	5
1.1.1 Intel Xeon E5-2600 V2 Processor Support	5
1.1.2 Batch Process and Memory Limits (Torque Only)	5
1.1.3 XFS File System Forwarding Over Data Virtualization Service (DVS)	6
1.1.4 ALPS Enhancements	6
1.1.4.1 Application-level Control of Processor Performance States (p-states) and Performance Governors	6
1.1.4.2 Exclusion Lists	7
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
3.2 Changes to Man Pages	12
3.2.1 Changed Cray Man Pages	12
Tables	
Table 1. Books Provided with This Release	11

Software Enhancements [1]

Cray Linux Environment (CLE) 5.0 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 5.0 release are described in *Cray Linux Environment (CLE) Software Release Overview*.

1.1 Software Enhancements in CLE 5.0.UP03

1.1.1 Intel Xeon E5-2600 V2 Processor Support

CLE provides support for the Intel Xeon E5-2600 V2 Series Processors. These processors are based on the Intel microarchitecture code named *IvyBridge*. The primary enhancements that come with the new processor include process improvements resulting in shrunken die area and thus reduced power consumption. Some users may see potential performance improvements.

1.1.2 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.

1.1.3 XFS File System Forwarding Over Data Virtualization Service (DVS)

XFS is a high performance file system developed for UNIX by Silicon Graphics International Corporation and later ported to Linux. The Cray Linux Environment supports XFS file systems with Cray XC30 systems using DVS. The file system will be direct-attached to a service node client, which will in turn be the DVS server for service and compute node clients within the system. Only one XFS over DVS server is supported as mounting more than one XFS client can cause data corruption in the underlying file system.

For instructions on how to configure an XFS file system using DVS see *Introduction to Cray Data Virtualization Service* (S-0005). For XFS documentation, see the project website, <http://oss.sgi.com/projects/xfs/>.

1.1.4 ALPS Enhancements

1.1.4.1 Application-level Control of Processor Performance States (p-states) and Performance Governors

Note: Application level p-state controls for compute nodes are currently provided for expert and research users only. Cray expects to make support for this feature generally available in a future release.

P-states are industry standard mechanisms defined in the Advanced Configuration and Power Interface (ACPI) specification used for CPU dynamic voltage and frequency scaling. They provide a mechanism for power saving and performance when enabled. P-states are defined as discrete levels of frequency while the processors are still executing instructions—they are not in an inactive or sleep state. Users have the option to change the p-state for each application, thus changing the compute node clock frequency for the duration of application execution. When users elect to use a non-default p-state, there is an inherent trade-off between power savings and performance.

Performance governors are supplied by the `CPUFreq` kernel code. For each setting, they supply a certain operating mode and frequency scaling method for the CPU. ALPS will accept available performance governor settings for each `aprun` instance. If the specified performance governor does not match one supported by CLE or the batch reservation, it will be silently discarded. The settings available in the kernel are:

`ondemand` CPU frequency will scale based on usage.

`performance`

Run the CPU at the maximum frequency.

`conservative`

Like `ondemand`, a usage-dependent setting, however the frequency "swings" are not as drastic.

`powersave` Sets the frequency to the lowest frequency available to the kernel.

`userspace` Allows the user to set a specific frequency.

For more information on performance governors, see <https://www.kernel.org/doc/Documentation/cpu-freq>.

Available p-state frequencies are located in:

`/sys/devices/system/cpu/cpu0/cpufreq/scaling_available_frequencies`
on system login nodes.

For more information on using p-states, see the `aprun(1)` man page.

Future versions of workload managers will be able to take advantage of control over p-states for batch reservations. ALPS will accept valid power management requests using the BASIL 1.4 protocol. The `ReserveParam` element will take two new `p-state` and `p-governor` attributes set with the corresponding appropriate values.

1.1.4.2 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."

This change is reflected in the synopsis for the `aprun` man page.

1.2 Bugs Addressed Since the Last Release

The list of customer-filed bug reports that were closed with CLE 5.0 releases is included in the *CLE 5.0 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 5.0 update release on a Cray system that was running an earlier version of the CLE 5.0 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 5.0.UP03 update release supports Cray XC30 systems.

2.2 Supported Software Upgrade Path

The CLE 5.0.UP03 release supports initial system installations and upgrade installations from CLE 5.0.UP00, 5.0.UP01 and 5.0.UP02.

2.2.1 System Management Workstation (SMW) Requirements

You must be running the SMW 7.0.UP03 release or later before you install the CLE 5.0.UP03 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. Applications targeted for Intel Xeon E5-2600 series processors will work without modification on Intel Xeon E5-2600 V2 processors.

2.4 Additional Software Requirements

2.4.1 Release Level Requirements for Other Cray Software Products

Note: You **must** upgrade the Cray Developer Toolkit (CDT) to the most current version (at time of release, this is 1.05) to ensure compatibility with the CLE 5.0.UP03 release package. For CDT release information, see the *Cray Programming Environments Installation Guide*, and the *Cray Programming Environment User's Guide*.

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 (S–2425–5002) includes a section that lists third-party software requirements for the CLE 5.0 release; this information applies to CLE 5.0 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 5.0.UP03 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-2425-5002), 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-2497-5003	Yes
<i>Cray Linux Environment (CLE) Software Release Overview</i>	S-2425-5002	No
<i>Installing and Configuring Cray Linux Environment (CLE) Software</i>	S-2444-5003	Yes
<i>Managing System Software for the Cray Linux Environment</i>	S-2393-5003	Yes
<i>Managing Lustre for the Cray Linux Environment (CLE)</i>	S-0010-5001	No
<i>Introduction to Cray Data Virtualization Service</i>	S-0005-5003	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-2496-5001	No
<i>Using the GNI and DMAPP APIs</i>	S-2446-5002	No
<i>Using Compute Unit Affinity on Cray Systems</i>	S-0030-5002	No
<i>Installing Cray Linux Environment (CLE) Support Package on esLogin Nodes or Linux Workstations</i>	S-2528-5003	Yes

3.2 Changes to Man Pages

3.2.1 Changed Cray Man Pages

- `intro_LLM_logfiles(5)`: power management is no longer deferred implementation.
- `cnselect(1)`: added `-a` option; added `-l` option fields `processor_status`, `alloc_mode`, `x_coord`, and `y_coord` ; removed `-l coremask` argument; added `.rparen.` and `.lparen.` query expression operators for the `-e` option.
- `aprun(1)`: changed `-ss` option text; added `APRUN_XFER_STACK_LIMIT` option; added `-p-state` and `-p-governor` options.
- `apstat(1)`: added `-C` option.