



Cray Linux Environment™ (CLE) 3.1 Software Release Overview Supplement

S-2497-3103

© 2010, 2011 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, LibSci, and PathScale are federally registered trademarks and Active Manager, Baker, Cascade, Cray Apprentice2, Cray Apprentice2 Desktop, Cray C++ Compiling System, 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 XD1, Cray XE, Cray XE6, 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, Gemini, Libsci, NodeKARE, RapidArray, SeaStar, SeaStar2, SeaStar2+, Threadstorm, and UNICOS/lc 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. Other names may be trademarks of their respective owners. Moab and TORQUE are trademarks of Adaptive Computing Enterprises, Inc. PBS Professional is a trademark of Altair Grid Technologies. PGI is a trademark of The Portland Group Compiler Technology, STMicroelectronics, Inc. All other trademarks are the property of their respective owners.

RECORD OF REVISION

S-2497-3103 Published March 2011 Supports the 3.1.UP03 update release of the Cray Linux Environment (CLE) operating system running on Cray XT and Cray XE systems.

S-2497-3102 Published December 2010 Supports the 3.1.UP02 update release of the Cray Linux Environment (CLE) operating system running on Cray XT and Cray XE systems.

S-2497-3101a Published October 2010 Supports the 3.1.UP01 update release of the Cray Linux Environment (CLE) operating system running on Cray XT and Cray XE systems.

S-2497-3101 Published September 2010 Supports the 3.1.UP01 update release of the Cray Linux Environment (CLE) operating system running on Cray XT and Cray XE systems.

Contents

	<i>Page</i>
Software Enhancements [1]	5
1.1 Software Enhancements in CLE 3.1.UP03	5
1.1.1 CLE 3.1 Support for Cray XT5 and Cray XT4 Systems	5
1.1.2 dumpd Daemon to Initiate Automatic Dump and Reboot of Nodes	7
1.1.3 CCM (Cluster Compatibility Mode) Enhancements	9
1.1.4 Lustre Upgraded to 1.8.4	10
1.1.5 Configuring a Virtual Local Area Network Interface	10
1.1.6 Enhanced Node Placement Scheme for Gemini Systems	11
1.1.7 RSIP Daemon Log File (syslog) Logging	11
1.1.8 New Cray OpenSM Startup Scripts	12
1.1.9 New Options for the lustre_control.sh Utility	13
1.1.10 CLEinstall Program Enhancements	14
1.1.11 Application Level Placement Scheduler (ALPS) Enhancements	15
1.1.12 DVS bulk_rw Mode and DVS-specific ioctl () Commands	16
1.2 Software Enhancements in CLE 3.1.UP02	17
1.2.1 DVS POSIX Atomicity	17
1.2.2 Changes to GPCD (Gemini Performance Counters Device) to Improve Performance when Accessing Memory Mapped Registers Using CrayPat	17
1.2.3 Repurposed Compute Node Support for PBS Professional	18
1.2.4 New xtverifydefaults Command	19
1.2.5 Core Specialization is NUMA-aware	20
1.3 Software Enhancements in CLE 3.1.UP01	21
1.3.1 Gemini Network Resiliency and Congestion Management	21
1.3.2 Repurposed Compute Nodes	22
1.3.3 Lustre Upgraded to Version 1.8.2	23
1.3.4 Application Completion Reporting	24
1.3.5 Topology and NID Ordering on Cray XE Systems	25
1.4 Software Enhancements in CLE 3.1.UP00	26
1.5 Previously Documented Limitations that are now Resolved	26

	<i>Page</i>
1.5.1 Resolved in CLE 3.1.UP02	26
1.5.1.1 Checkpoint/Restart Supported with Cray XT Systems Running CLE 3.1	26
1.5.2 Resolved in CLE 3.1.UP01	27
1.5.2.1 LSI RDAC Support	27
1.5.2.2 Number of Cores Support for Core Specialization	27
1.5.2.3 Huge Page Support for Both Dynamically and Statically Linked SHMEM, UPC and CAF Applications on Cray XE Systems	27
1.6 Bugs Addressed Since the Last Release	27
1.7 Compatibilities and Differences	28
Support Requirements [2]	29
2.1 Supported Cray System Hardware Platforms	29
2.2 Supported Software Upgrade Path	29
2.2.1 System Management Workstation (SMW) Requirements	29
2.2.2 Supported CLE Software Installations	29
2.3 Binary Compatibility	31
2.4 Additional Software Requirements	32
2.4.1 Release Level Requirements for Other Cray Software Products	32
2.4.1.1 Programming Environments Requirements for Hardware Upgrades	32
2.4.2 Third-party Software Requirements	33
Documentation [3]	35
3.1 Cray-developed Books Provided with This Release	35
3.1.1 CLE Installation and Migration Publications	36
3.2 Changes to Man Pages	37
3.2.1 New Cray Man Pages	37
3.2.1.1 CLE 3.1.UP03	37
3.2.1.2 CLE 3.1.UP02	37
3.2.1.3 CLE 3.1.UP01	37
3.2.2 Changed Cray Man Pages	38
3.2.2.1 CLE 3.1.UP03	38
3.2.2.2 CLE 3.1.UP02	38
3.2.2.3 CLE 3.1.UP01	39
Tables	
Table 1. Books Provided with This Release	35
Figures	
Figure 1. CLE Migration Roadmap	30

Software Enhancements [1]

CLE 3.1 software update releases 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 3.1 release are described in *Cray Linux Environment (CLE) 3.1 Software Release Overview* (S-2425-31), which is also provided with the release package.

1.1 Software Enhancements in CLE 3.1.UP03

1.1.1 CLE 3.1 Support for Cray XT5 and Cray XT4 Systems

Who will use this functionality?

All users and administrators on Cray XT5, Cray XT5m, and Cray XT4 systems.

How can this support help me?

Beginning with the CLE 3.1.UP03, Cray supports the CLE 3.1 operating system release on Cray XT5, Cray XT5m, and Cray XT4 systems. CLE 3.1 release includes a large number of Cray features, fixes and other enhancements for all users and administrators. Because SLES and Lustre are upgraded for this release, additional features and fixes are available.

What prerequisites do I need to know about?

The System Management Workstation (SMW) must be running the SMW 5.1.UP03 release or later. To install the SMW 5.1 release on your SMW, you must run release version CLE 2.2.UP03B or later on your Cray XT system.

To run SMW 5.1 with CLE 2.2.UP02, contact your Cray service representative and request a special CLE 2.2.UP02 release package that includes modifications to the CLE installation software.

Important: Without these changes, you cannot generate a CLE 2.2 boot image for SMW 5.1.UP03 and will not be able to boot your CLE 2.2.UP02 system when your SMW is running 5.1.

What is required to install CLE 3.1 on these systems?

Moving from CLE 2.2 to CLE 3.1 requires a migration from SLES 10 SP2 to SLES 11. The migration process involves preserving some of your local configuration data from the CLE 2.2 system set, performing an initial install of CLE 3.1 on a different system set, and finally restoring the preserved configuration data. Lustre file system data is preserved.

To move from CLE 2.2 to CLE 3.1 on a Cray XT5, Cray XT5m, or Cray XT4 system, refer to *Migrating to Cray Linux Environment (CLE) 3.1 Software for Cray XT Systems* (S–2489).

If you are upgrading your Cray XT5 or Cray XT5m system hardware to a Cray XE system, select either a single phase or a dual phase migration path as follows:

Single phase

Coordinate the hardware replacement and software installation to directly migrate your Cray XT system running CLE 2.2 to a Cray XE system running CLE 3.1.

Dual phase

Schedule the software and hardware changes separately. First move from CLE 2.2 to CLE 3.1 and verify that you have a successful software migration. At some later date, complete the hardware replacement and upgrade from CLE 3.1 for SeaStar systems to CLE 3.1 for Gemini systems.

Where can I find more information?

For CLE 3.1 feature information see *Cray Linux Environment (CLE) 3.1 Software Release Overview* (S–2425–31) and the remaining sections in this document.

For more information on migrating to CLE 3.1, see [Figure 1, Supported CLE Software Installations on page 29](#) and [CLE Installation and Migration Publications on page 36](#).

1.1.2 `dumpd` Daemon to Initiate Automatic Dump and Reboot of Nodes

Who will use this feature?

System administrators and site analysts.

How can this `dumpd` enhancement help me?

Using the new `dumpd` functionality, system administrators can decrease node down time and, if necessary, access evidence that can help identify the problem that caused the node to be unhealthy.

What does this enhancement do?

When a node fails a node health test, Node Health Checker (NHC) can request a dump or reboot of the node depending on the action associated with that test in the node health configuration file. The following are new node health actions, with descriptions of the steps taken for each action:

<code>dump</code>	Sets the compute node's state to <code>admindown</code> and requests a dump from the SMW, in accordance with the <code>maxdumps</code> configuration variable.
<code>reboot</code>	Sets the compute node's state to <code>unavail</code> and requests a reboot from the SMW.
<code>dumpreboot</code>	Sets the compute node's state to <code>unavail</code> and requests a dump and reboot from the SMW.

Requests for dumps and reboots can be made of the `dumpd` daemon in two ways:

- Node Health Checker (NHC) can automatically call `dumpd` (if so configured in the NHC and `dumpd` configuration files).
- System administrators can use the `dumpd-request` script on the SMW or the shared root, provided `dumpd` is enabled in the `/etc/opt/cray-xt-dumpd/dumpd.conf` file. You can define additional actions in this configuration file.

In each case, `dumpd` must be enabled in the `/etc/opt/cray-xt-dumpd/dumpd.conf` file; otherwise, it will not handle any requests.

The `dumpd` binary sits and waits for requests from NHC (or some other entity using the `dumpd-request` tool on the shared root). When `dumpd` gets a request, it creates a database entry in the `mznhc` database for the request, and calls the script `/opt/cray-xt-dumpd/default/bin/executor` to read the configuration file found at `/etc/opt/cray-xt-dumpd/dumpd.conf` and perform the requested actions.

System administrators can use the `dumpd-dbadmin` script to view or delete entries in the `mznhc` database in a convenient manner. The `dumpd-dbadmin` tool can be found in `/opt/cray-xt-dumpd/default/bin`.

The NHC and `dumpd` configuration files enable administrators to completely control when and how `dumpd` responds to nodes set to `admindown` by NHC, as well as how many nodes are dumped.

Where can I find more information?

For more information, see *Managing System Software for Cray XE and Cray XT Systems* and the `intro_NHC(8)`, `dumpd(8)`, `dumpd-request(8)`, and `dumpd-dbadmin(8)` man pages.

In addition, see the configuration files described in the `intro_NHC(8)` man page for specific system variables and examples.

1.1.3 CCM (Cluster Compatibility Mode) Enhancements

Who will use CCM enhancements?

End users, application developers, and system administrators.

How can the CCM enhancements help me?

You can use new command line options to override default behavior when you launch your cluster-based job. Additional enhancements in CCM improve performance and functionality for ISV applications.

What do the CCM enhancements do?

- The `ccmrun` command includes several new options to enable and disable the following functionality: `SSH` daemon, `portmap` and `xinetd` daemons for `rsh`, name service caching daemon, and RSIP (Realm Specific Internet Protocol). Most options describe a default behavior for CCM, but allow you to override that behavior with non-default configurations or environments.
- CCM makes the head node of a job the first processing element (PE[0]) in the associated node list for the reservation. This is a requirement of many ISV applications.
- `ccmlogin` propagates the exact login node environment settings to the CCM interactive environment via SSH.
- CCM handles application termination and job cleanup more efficiently.
- CCM and installation configuration files include parameters to customize your Workload Management (WLM).

Where can I find more information?

The changes visible to the end user are documented in man pages `ccmrun(1)` and `ccmlogin(1)`. Administrators should refer to *Installing and Configuring Cray Linux Environment (CLE) Software* and *Managing System Software for Cray XE and Cray XT Systems* for information on how to install and complete the setup of CCM.

1.1.4 Lustre Upgraded to 1.8.4

Who will use this feature?

Users and administrators of the Lustre file system.

How does this new version of Lustre help me?

It resolves a large number of existing bugs.

What does upgrading to Lustre 1.8.4 do?

When you update or upgrade your system to CLE 3.1.UP03, your Lustre file system software is automatically upgraded to version 1.8.4. This change is transparent to both users and administrators. You can view version-specific change logs from Oracle at http://wiki.lustre.org/index.php/Change_Log_1.8. Additionally, you can view a list of CLE bug fixes in the README file included with the release package

Where can I find more information?

For more information on Lustre, visit <http://wiki.lustre.org/> or see the *Lustre Operations Manual* from Oracle which was updated for this release; a copy is included with your CLE 3.1 release package. For information about upgrades to Lustre 1.8 from Lustre 1.6, see *Cray Linux Environment (CLE) 3.1 Software Release Overview* (S-2425-31).

1.1.5 Configuring a Virtual Local Area Network Interface

Who will use this feature?

System administrators and site analysts.

How does this feature help me?

A Virtual Local Area Network (VLAN) enables you to set up a LAN to have different broadcast domains, which often improves performance by reducing unnecessary network traffic. VLANs enable you to set up virtual workgroups, making it easier to move users from one network to another. VLANs may also provide better network security.

What does this feature do?

CLE 3.1 now supports the 802.1Q VLAN standard. A new procedure is provided to help you configure a 802.1Q VLAN interface for your Cray system.

Where can I find more information?

The new procedure is located in *Managing System Software for Cray XE and Cray XT Systems* (S-2393-3103).

1.1.6 Enhanced Node Placement Scheme for Gemini Systems

Who will use this feature?

All users and administrators on Cray XE systems.

How does this feature help me?

A wide variety of applications perform better as compared to their performance on earlier versions of CLE 3.1.

What does this feature do?

The script that orders the nodes for allocation by ALPS improves node placement on a Gemini based system interconnection network. This script is invoked only when the `ALPS_NIDORDER` configuration variable is set to `-O2`. Most sites with Cray XE systems should use this configuration option.

The CLE 3.1.UP03 release also includes modifications to the node ordering script to support Cray XE6m and Cray XE5m systems larger than three cabinets.

Where can I find more information?

The `apbridge(8)` man page and *Managing System Software for Cray XE and Cray XT Systems* describe `ALPS_NIDORDER` options.

1.1.7 RSIP Daemon Log File (`syslog`) Logging

Who will use this feature?

System administrators and site analysts.

How can the `rsipd` log file help me?

Log file (`syslog`) logging of RSIP daemon (`rsipd`) messages enables the system administrator to more easily access this information and reduces the default log level (as described in the `rsipd.conf` file) by entering fewer `rsipd` messages in the `syslog`, thus eliminating considerable "spam".

What does this feature do?

During the RSIP daemon (`rsipd`) startup process (before disassociating from a controlling terminal), messages are logged to both the `syslog` and to `stderr`. This action now includes log file (`syslog`) logging. While the `rsipd` daemon is running, you can raise or lower the log level by sending a `SIGUSR1` or a `SIGUSR2` signal, respectively. The former use of the `SIGUSR1` signal to have the `rsipd` daemon dump the server state to the default status file `/var/run/rsipd.stat` is replaced by the `SIGWINCH` signal.

Where can I find more information?

The `rsipd(8)` man page is updated in the CLE 3.1.UP03 release.

1.1.8 New Cray OpenSM Startup Scripts

Deferred Implementation: This feature is planned for a future release.

Who will use this feature?

System administrators.

How can this feature help me?

A new package, `cray-opensm-init`, supports starting two subnet managers (`opensm`) for dual-port Host Channel Adapter (HCA) direct-attached storage configurations. Site modifications of the base OpenFabrics Enterprise Distribution (OFED) Session Manager script (`opensmd`) are no longer required.

Note: This package is only intended for dual-port direct-attached storage configurations. It is not recommended as a general solution for starting two subnet managers on the same host.

What does this feature do?

The new package provides two startup scripts, `opensmd-port1` and `opensmd-port2`, which are pre-configured with the appropriate log file, cache directory, and temporary directory locations. These scripts can be modified as needed by changing their configuration files: `/etc/sysconfig/opensm-port1` and `/etc/sysconfig/opensm-port2`. The original `/etc/init.d/opensmd` and `/etc/sysconfig/opensm` are not changed by the new package.

To use this new feature, system administrators must disable `opensmd` and then enable `opensmd-port1` and `opensmd-port2` by using the `chkconfig` command.

Where can I find more information?

For more information, contact your Cray service representative.

1.1.9 New Options for the `lustre_control.sh` Utility

Who will use this feature?

System administrators.

How can this feature help me?

By using new `lustre_control.sh` options, you can now mount or unmount Lustre across all compute nodes, specify a list of nodes, and specify a mount point.

What does this feature do?

The `-c`, `-m`, and `-n` options to `lustre_control.sh` add functionality to the `mount_clients` and `umount_clients` actions.

The `-c` option sends mount and unmount commands strictly to compute node Lustre clients. However, this option only works if passwordless `ssh` is enabled on your system. The `-m` option allows you to specify a mount point other than the `MOUNT_POINT` variable from the `filesystem.fs_defs` file. The `-n` option enables you to specify a list of nodes that receive the `mount_clients` and `umount_clients` actions.

Where can I find more information?

For more information, see the `lustre_control.sh(8)` man page or *Managing Lustre for the Cray Linux Environment (CLE)*

1.1.10 CLEinstall Program Enhancements

Who will use this feature?

System administrators.

How can these enhancements to CLEinstall help me?

For all types of CLE software installations and upgrades, CLEinstall provides additional verification of configuration information and automation for some types of configuration changes, including new features and some hardware configuration changes.

What does this feature do?

- The CLEinstall program verifies and updates the `/etc/hosts` file for Cray hostnames and aliases. Host name aliases in `/etc/hosts` are assigned based on the class name and **order** of NIDs that are specified by each `node_class[idx]` parameter.

Note: CLEinstall modifies Cray system entries in `/etc/hosts` each time you update or upgrade your CLE software.
- For initial installations, CLEinstall creates the `/etc/opt/cray/sdb/node_classes` file based on the `node_class[idx]` parameters you specified in `CLEinstall.conf`. For update or upgrade installations, CLEinstall verifies that they match.
- Each time you upgrade or update CLE software, CLEinstall checks for new service nodes. Any new nodes are initialized in the shared root so that you can use `xtopview -n NID` to customize the nodes, as needed.
- When you add or remove hardware (service nodes, cabinets, or chassis within cabinets), you can run CLEinstall to modify your software configuration for the changes. Use the `--xthwinv` option with CLEinstall to apply the new hardware component information to the specified system set and modify other relevant configuration files. You can do this without changing the operating system release level.

Note: This is in addition to functionality that already existed to change `CLEinstall.conf` parameters without changing the OS level.

- If your system runs CLE 3.1 and has been hardware upgraded from a Cray XT system to a Cray XE system, the `CLEinstall` program performs a number of additional configuration steps to update some system configuration files for the new hardware.

Where can I find more information?

These changes are documented in the books listed in [CLE Installation and Migration Publications on page 36](#) and in the `CLEinstall.conf(5)` man page.

The procedures for *Adding or removing cabinets or chassis within cabinets* and *Adding or removing a service node* in *Managing System Software for Cray XE and Cray XT Systems* (S-2393–3103) are updated to take advantage of `CLEinstall` enhancements.

1.1.11 Application Level Placement Scheduler (ALPS) Enhancements

Who will use these ALPS Enhancements?

End users and application developers.

How can the ALPS enhancements help me?

ALPS provides an output environment variable, `ALPS_APP_DEPTH`, which may assist application programmers when launching multiple programs within a job.

The `aprun(1)` man page language under the MPMD section was modified to clarify that the `-m` option should be specified in the first program executable segment and this value is inherited for subsequent programming segments.

What do the ALPS enhancements do?

`ALPS_APP_DEPTH` can assist application programmers when using MPMD mode. For each programming segment within the ALPS job, there may be varying values for depth (`-d`). Program executables or libraries linked against these executables may make a `getenv()` call, requesting the value for `ALPS_APP_DEPTH` to determine how many threads were requested locally and modify program behavior appropriately.

Where can I find more information?

The changes are documented in the `aprun(1)` man page.

1.1.12 DVS `bulk_rw` Mode and DVS-specific `ioctl()` Commands

Who will use Cray Data Virtualization Service (Cray DVS) Enhancements?

End users and application developers.

How can the DVS enhancements help me?

DVS provides `ioctl` commands to user space applications which can query DVS configuration data for a specific file that is visible within the name space of a DVS mount point.

DVS now allows a bulk read/write option, `bulk_rw`, providing potential performance improvements.

What do the DVS enhancements do?

DVS `ioctl` commands are defined in `dvs_ioctl.h`. To use these DVS-specific `ioctl()` commands, include `dvs_ioctl.h` in your source and load the DVS module in your compiling environment. The following commands are available to application users:

`DVS_GET_REMOTE_FS_MAGIC`

Returns the magic value of the underlying file system.

`DVS_GET_FILE_BLK_SIZE`

Returns the DVS block size in bytes for a file that is visible on the DVS mount point.

`DVS_GET_FILE_STRIPE_WIDTH`

Returns the DVS stripe width for a file that is visible on the DVS mount point.

`bulk_rw` allows DVS to execute read and write operations and eliminates the need for temporary data transfer buffers to or from DVS servers. `bulk_rw` performs RDMA (Remote Direct Memory Access) operations directly to or from file pages in the Linux kernel. There is also an environment variable, `DVS_BULK_RW=<on/off>` which overrides the behavior of a `-o bulk_rw` or `-o nobulk_rw` mount options as specified in `/etc/fstab` for compute node images.

Where can I find more information?

The changes are documented in the `dvs(5)` man page.

1.2 Software Enhancements in CLE 3.1.UP02

1.2.1 DVS POSIX Atomicity

Who will use DVS atomic stripe parallel mode?

System administrators.

How can atomic stripe parallel mode help me?

These enhancements can prevent interleaving of data, which could occur when an application performs I/O on a shared file without using file locking on a DVS stripe parallel mount point.

What does atomic stripe parallel mode do?

Atomic stripe parallel, a Cray DVS mode, adheres to POSIX read/write atomicity rules while still allowing for possible parallelism within a file. It is similar to stripe parallel mode in that the server used to perform the read, write, or metadata operation is selected using an internal hash involving the underlying file or directory inode number and the offset of data into the file relative to the DVS block size.

Where can I find more information?

For information on Cray DVS, see *Introduction to Cray Data Virtualization Service*. For information on how to set up Cray DVS, see *Installing and Configuring Cray Linux Environment (CLE) Software* and *Managing System Software for Cray XE and Cray XT Systems*.

1.2.2 Changes to GPCD (Gemini Performance Counters Device) to Improve Performance when Accessing Memory Mapped Registers Using CrayPat

Who will use this enhancement?

End users.

How can this enhancement help me?

CLE has been modified to decrease MMR (Memory Mapped Register) access time, thus giving a significant performance improvement when using CrayPat. Users should be able to make traces that involve MMR data throughout their program with less overhead than previous releases.

Another direct end user benefit of this improvement is that Gemini network counter event collection can be done more frequently during application execution. The user can now monitor a half-dozen counters without drastically degrading performance. As a reference, when tracing function enter/return, the overhead to count one Gemini event is approximately equivalent to four processor counter events.

What is modified to increase performance?

The Gemini application specific integrated circuit contains MMRs, also known as network performance counters, that are used within CrayPat to capture network performance data during trace experiments. The performance improvement introduced in this update release decreases the access time experienced when using CrayPat to view Gemini MMR data. This is done via a modification to a kernel-level interface, `gpcd` (Gemini Performance Counters Device); thus, end users do not need to take any further action to take advantage of this benefit.

Where can I find more information?

For more information on Gemini MMRs used as performance counters, see *Using the Cray Gemini Hardware Counters*. For more information on CrayPat, see *Cray Performance Analysis Tools Release Overview and Installation Guide* and *Using Cray Performance Analysis Tools*.

1.2.3 Repurposed Compute Node Support for PBS Professional

Who will use repurposed compute node support for PBS Professional?

System administrators.

How can this help me?

You can potentially improve performance or maintainability of PBS Professional batch system services by moving the MOM node to a dedicated node instead of sharing resources with other services on an existing service node.

What is repurposed compute node support for PBS?

Beginning with CLE 3.1.UP02, Cray tests and supports repurposing a compute node as a PBS Professional MOM node. CLE 3.1.UP01 introduced new functionality to repurpose compute nodes as service nodes. For the initial release of this feature, support for repurposing compute nodes as a PBS Professional MOM nodes was not available. For more information on repurposed compute nodes, see [Repurposed Compute Nodes on page 22](#).

Where can I find more information?

For information on how to setup and configure compute nodes as PBS Professional MOM nodes, see *Repurposing Compute Nodes as Service Nodes on Cray XE and Cray XT Systems* and PBS Professional documentation available from Altair Engineering, Inc. at <http://www.altair.com>.

1.2.4 New `xtverifydefaults` Command

Who will use this command?

System administrators, automatically via the `xtopview` utility.

How can this command help me?

The `xtopview` utility automatically invokes a new command, `xtverifydefaults`, to verify that the default software version for Cray packages is the same on all service nodes. In rare cases, the shared root links to Cray packages can become inconsistent so that not all service nodes point to the same default software version. This situation can be difficult to detect. By default, the `xtopview` utility invokes the `xtverifydefaults` command during exit processing to correct any inconsistent links that may have been generated during the `xtopview` session. System administrators can also run the command manually from within `xtopview`.

What does this command do?

The `xtverifydefaults` command verifies and optionally fixes the default links in the shared root for Cray system software packages. If the version of software a link points to is different from the version that the default view link of the same name points to, the link is considered inconsistent. The `xtverifydefaults` command flags and optionally fixes inconsistent links. All of the `node` and `class` view default links for Cray packages (e.g. those that start with "cray-") are forced to match the default view default links.

Where can I find more information?

See the `xtverifydefaults(8)` man page.

1.2.5 Core Specialization is NUMA-aware

Who will use NUMA-aware core specialization?

End users.

How can NUMA-aware core specialization help me?

This enhancement may help to "even out" the distribution of the application-generated service processes when the `aprun -r` core specialization option is greater than 1.

What is NUMA-aware core specialization?

Core specialization may improve application performance by binding system processes and daemons to a set of cores in each processor. Initially, only one specialized core per node was allowed. In CLE 3.1.UP01, this limitation was removed. This feature provides more support for core specialization for specialized cores greater than 1. ALPS allocates specialized cores round-robin across the NUMA nodes on a node (starting with the highest-numbered core on the highest-numbered NUMA node), unless the user specifies the `aprun -cc cpu_list` option. In that case, specialized cores are allocated from the highest-numbered cores, avoiding the cores in *cpu_list*.

Where can I find more information?

For more information on how to use core specialization, see the `aprun(1)` man page and *Workload Management and Application Placement for the Cray Linux Environment*.

1.3 Software Enhancements in CLE 3.1.UP01

1.3.1 Gemini Network Resiliency and Congestion Management

Who will use these enhancements?

End users, system administrators, and site analysts.

How can these enhancements help me?

Network congestion is inherent in any High Performance Computing (HPC) high speed interconnect. Cray's software and hardware are designed to minimize the effects of network congestion. This release includes system software enhancements, administrative recommendations, and user documentation, which together help to prevent network congestion or to protect user jobs and data if congestion occurs.

What are the network resiliency and congestion management enhancements?

The CLE 3.1.UP01 and SMW 5.1.UP01 releases provide the following software enhancements and documentation to improve network resiliency and manage network congestion on Cray XE systems with a Gemini based system interconnection network.

- New user documentation describes some cases where congestion can occur and how to modify PGAS and SHMEM applications to minimize the potential for inducing network congestion.
- The Hardware Supervisory System (HSS) includes new software that monitors the Gemini network and, if necessary, throttles traffic to alleviate the complications due to congestion.
- System administrators can use new information about Gemini High Speed Network (HSN) routing to identify and potentially modify unrouteable network configurations.
- To improve network resiliency on the HSN, system software on the L0 automatically brings up lanes that have been downed.
- System software changes include selective link speed increases and improvements to the HSN routing algorithm (*zone routing*). These adjustments can improve global HSN communication performance by as much as 15 to 25 percent. For more information, contact your Cray service representative.

1.3.2 Repurposed Compute Nodes

Who will use this feature?

System administrators and site analysts.

How can repurposed compute nodes help me?

You can improve performance or maintainability of some Cray system services by moving the service to a dedicated node instead of sharing resources with other services on an existing service node.

You do not need to make hardware configuration changes (swapping modules) to configure additional nodes with a service node role; you can implement a flexible approach to service node configuration on an as-needed basis.

The addition of service nodes can increase throughput for some functions. For example, multiple DVS servers working in parallel can increase I/O throughput; multiple MOM nodes can allow additional jobs to run in parallel.

For some services, configuring additional service nodes can increase redundancy and minimize the negative impact of a failed service node.

What does this feature do?

The CLE 3.1.UP01 release includes new functionality that supports booting compute node hardware with service node images. By using this functionality to change the role of a compute node, you can add additional service nodes for services that do not require external connectivity.

Some services on Cray systems have resource requirements or limitations (for example, memory, processing power or response time) that you can address by configuring a dedicated service node, such as a Cray Data Virtualization Service (Cray DVS) node or a batch system management (MOM) node. On Cray XE and Cray XT systems, service I/O node hardware (on a service blade) is equipped with Peripheral Component Interconnect (PCI) protocol card slots to support external devices. Compute node hardware (on a compute blade) does not have PCI slots. For services that do not require external connectivity, you can configure the service to run on a single, dedicated compute node and avoid using traditional service I/O node hardware.

When you configure a node on a compute blade to boot a service node image and perform a service node *role*, that node is referred to as a *repurposed compute node*.

Cray tests and supports the following services on repurposed compute nodes: DVS servers, Moab TORQUE MOM nodes, and PBS Professional MOM nodes.

Note: Support for repurposing a compute node as a PBS Professional MOM node was originally deferred but is now supported.

Where can I find more information?

A new Cray publication, *Repurposing Compute Nodes as Service Nodes on Cray XE and Cray XT Systems* (S-0029-3101) describes the feature and provides implementation procedures.

Note: The information in this new document supersedes procedures to manually configure compute nodes as compute node root servers for DSL (using DVS) in the following documentation:

- *Installing and Configuring Cray Linux Environment (CLE) Software* (S-2444-31)
- *Managing System Software for Cray XE and Cray XT Systems* (S-2393-31)

1.3.3 Lustre Upgraded to Version 1.8.2

Who will use this feature?

All users and administrators on Cray systems that use the Lustre file system.

How can this new version of Lustre help me?

A large number of existing problems are fixed with this new version of Lustre.

What does upgrading to Lustre 1.8.2 do?

When you update a CLE 3.0 or CLE 3.1 system to CLE 3.1.UP01, your Lustre file system software is automatically upgraded from version 1.8.1 to version 1.8.2. This change is transparent to both users and administrators.

Where can I find more information?

Refer to <http://wiki.lustre.org>. For information about upgrades to Lustre 1.8 from Lustre 1.6, see *Cray Linux Environment (CLE) 3.1 Software Release Overview* (S-2425-31).

1.3.4 Application Completion Reporting

Who will use this feature?

System administrators and site analysts.

How can ACR help me?

Application Completion Reporting (ACR) provides a way to track, manage, and report application data, which gives system administrators an effortless way to see how system resources are being allocated.

Data provided by ACR reports can make it easier to see where system resources are being used so that recommendations can be made for application improvement, different allocation schemes, or use of different time slices.

What does ACR do?

ACR extends the data persistence schema from Cray Management Services (CMS) and provides three commands to examine and report application data.

- Use the `mz job` command to examine reservation data; you can specify the format of the search output so that you can use it in other programs or scripts.
- Use the `mzreport` command to examine application completion status information; the command supports claims (applications), jobs (reservations), or user completion status.
- Use `mz2attr` to read and display node attributes.

ACR requires CLE 3.1.UP01 or later and SMW 5.1.UP01 or later; support for this feature was deferred in earlier SMW and CLE releases.

Where can I find more information?

Using Cray Management Services (CMS) (S-2484-5101) and the `mz job(8)`, `mzreport(8)` and `mz2attr(8)` man pages.

1.3.5 Topology and NID Ordering on Cray XE Systems

Who will use this feature?

End users, site-analysts, and system administrators.

How can this feature help me?

The Application Level Placement Scheduler (ALPS) "xyz-by2" node ordering feature provides performance improvements for larger applications that run on Cray XE systems with no user interaction. System administrators can implement the new ordering with only a single configuration option change to the `/etc/sysconfig/alps` file on the boot node.

What does this feature do?

The node ordering functionality directs ALPS to use a node ordering option that results in better application performance for larger applications run on Cray XE systems.

Cray's first method for application placement on nodes along the interconnect torus was to use the node ID (NID) number list and place applications on nodes in this numerical order. Later, the simple "xyz" placement method was added, which reordered the sequence of NID numbers, used in assigning placements, in a natural 3D communications sequence (that is, following the torus, not the chassis and cabinet layout). The original numerical ordering method provides good "packing" for small numbers of nodes, but for large applications, there are "clumps" of nodes that spread across the machine, which causes inter-application traffic contention on the interconnect. With the "xyz" placement, available with previous releases of CLE, the applications are better isolated, but smaller applications are spread more thinly, losing some of the advantage.

New with this release, the "xyz-by2" NID reordering method combines the incidental small node packing of the simple NID number method with the inter-application interaction reduction of the "xyz" method.

To implement this node ordering, system administrators must set the `ALPS_NIDORDER="-O2"` parameter in the `/etc/sysconfig/alps` configuration file and make identical changes to the file of the same name on the shared root. If configured, a custom ordering is created when ALPS is started on the boot node, triggered in the `/etc/init.d/alps` script file, and used by the `apbridge` component of ALPS.

You should change the `ALPS_NIDORDER` parameter **only** at system reboot time.

Previously, this feature was supplied to selected customers via a software patch. While the internal mechanism used to implement this feature remains the same, the configuration method has changed. The preferred method is to set the `ALPS_NIDORDER` parameter in the `/etc/sysconfig/alps` file. The configuration instructions provided with the patch are obsolete and should no longer be used.

Where can I find more information?

The `apbridge(8)` man page was updated for the CLE 3.1.UP01 release.

1.4 Software Enhancements in CLE 3.1.UP00

For information about feature content in the initial or base CLE 3.1 release (CLE 3.1.UP00), see *Cray Linux Environment (CLE) 3.1 Software Release Overview* (S–2425–31).

1.5 Previously Documented Limitations that are now Resolved

The following features or capabilities were identified as limitations with an earlier CLE 3.1 release and were described in the *CLE 3.1 Limitations* document. These issues are now resolved in the update release indicated.

1.5.1 Resolved in CLE 3.1.UP02

1.5.1.1 Checkpoint/Restart Supported with Cray XT Systems Running CLE 3.1

The CLE 3.1.UP01 Limitations document states that with CLE 3.1, Checkpoint/Restart is not functional on Cray XT systems. The CLE 3.1.UP02 update release removes this limitation. For additional information, see *Managing System Software for Cray XE and Cray XT Systems* for more information on Checkpoint/Restart.

1.5.2 Resolved in CLE 3.1.UP01

1.5.2.1 LSI RDAC Support

Beginning with the CLE 3.1.UP01 update release, the integrated LSI Redundant Disk Array Controller (RDAC) code is updated to version V09.03.0C00.0214 to support all models and firmware releases of LSI hardware. For additional information about support for multipathing I/O in the CLE 3.1 release, see *Cray Linux Environment (CLE) 3.1 Software Release Overview* (S-2425-31), Section 2.15.1 "Support for Multipathing I/O"; This document is also provided with the release package.

Note: Replace the last paragraph in the referenced section with the following, corrected, statement regarding performance impact: Multipathing functionality does not negatively impact I/O performance during normal system operations. In the event of a failure, Multipathing functionality may impact performance depending on the multipath configuration, failure scenario and I/O workload on your system.

1.5.2.2 Number of Cores Support for Core Specialization

The *Cray Linux Environment (CLE) 3.1 Software Release Overview* (S-2425-31) states that with CLE 3.1, only `r=1` is supported for the Core Specialization feature. The CLE 3.1.UP01 update release removes this limitation. You can specify a value for cores greater than or equal to one for the `aprun -r cores` option to enable core specialization. For additional information, see the *Workload Management and Application Placement for the Cray Linux Environment* and the `aprun(1)` man page.

1.5.2.3 Huge Page Support for Both Dynamically and Statically Linked SHMEM, UPC and CAF Applications on Cray XE Systems

Beginning with CLE 3.1.UP01 release, SHMEM, Unified Parallel C (UPC), and Co-Array Fortran (CAF) applications on Cray XE systems may be either statically or dynamically linked with `libhugetlbfs`. For additional information, see the `intro_hugepages(1)` man page that is released in the `xt-asynce 4.2` package.

1.6 Bugs Addressed Since the Last Release

The primary purpose of each CLE 3.1 update package is to provide fixes for the release. The list of customer-filed bug reports that were closed with CLE 3.1 releases is included in the *CLE 3.1 Errata*; this document is provided with the release package.

1.7 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 after you install a CLE 3.1 update release on a Cray system that was running an earlier version of the CLE 3.1 release; for example, a CLE 3.1.UP02 to CLE 3.1.UP03 update installation.

For information about compatibility issues and functionality changes after you migrate or upgrade from the CLE 2.2 or CLE 3.0 releases, see the *Compatibilities and Differences* chapter in *Cray Linux Environment (CLE) 3.1 Software Release Overview* (S–2425–31).

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 3.1.UP03 update release supports Cray XE6, Cray XE6m, Cray XE5, Cray XE5m, Cray XT6, Cray XT6m, Cray XT5, Cray XT5m, and Cray XT4 systems.

2.2 Supported Software Upgrade Path

2.2.1 System Management Workstation (SMW) Requirements

The System Management Workstation (SMW) must be Restriction of Hazardous Substances (RoHS) compliant and must be running the SMW 5.1.UP01B release or later before you install the CLE 3.1.UP03 update release package. For additional information, see the SMW *README* document included with the SMW release package.

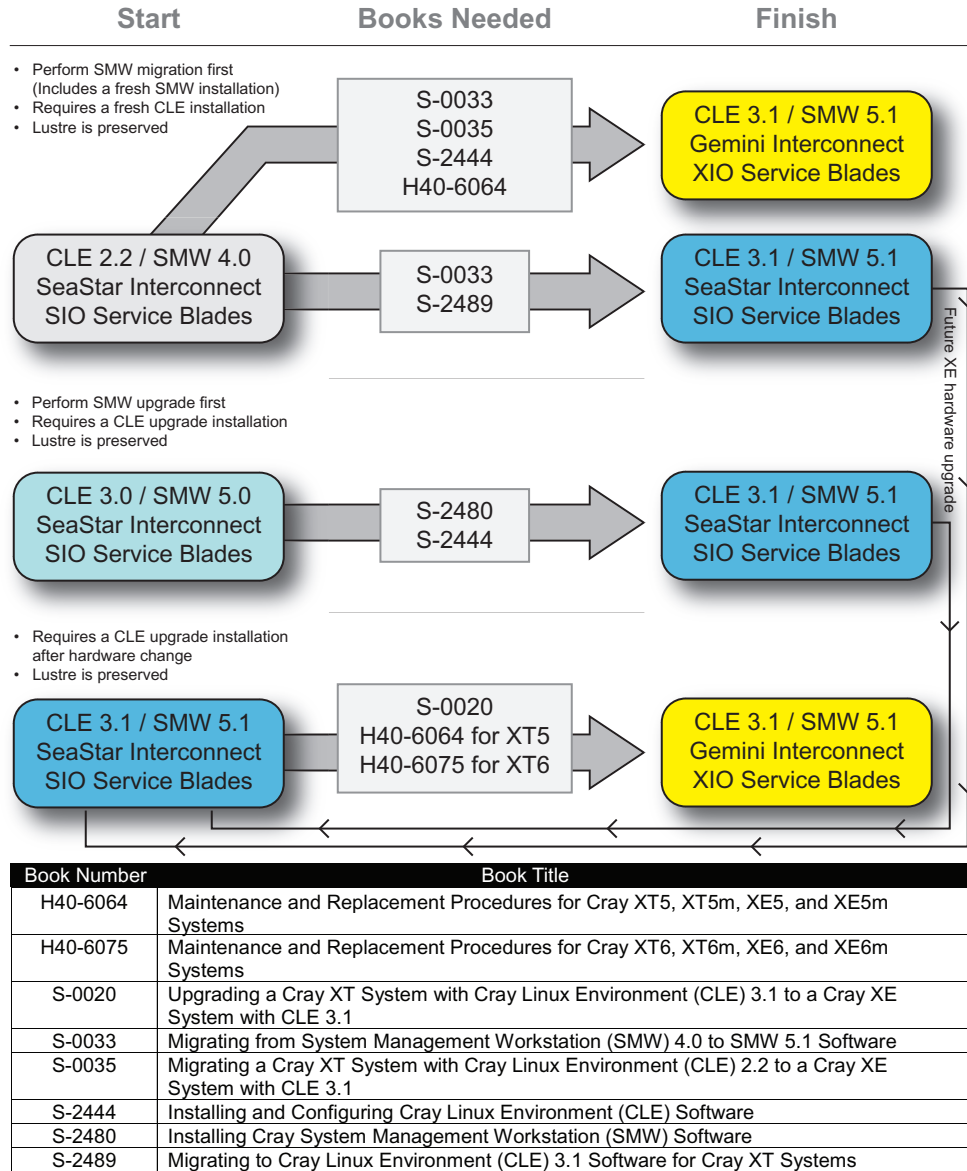
Note: If you have a Cray XT system, you must run release version CLE 2.2.UP03B or later to install the SMW 5.1 release on your SMW. To run SMW 5.1 with CLE 2.2.UP02, contact your Cray service representative and request a special CLE 2.2.UP02 release package that includes modifications to the CLE installation software. Without these changes, you cannot generate a CLE 2.2 boot image for SMW 5.1 and will not be able to boot your CLE 2.2.UP02 system when your SMW is running 5.1.

2.2.2 Supported CLE Software Installations

The CLE 3.1 release supports four types of software installations. CLE software installation types are described in *Cray Linux Environment (CLE) 3.1 Software Release Overview* (S-2425-31).

Figure 1 illustrates the upgrade, update, migration, and installation processes for various Cray hardware paths. The CLE 3.1.UP03 update release supports the following software installations:

Figure 1. CLE Migration Roadmap



- Initial** or new software installations on the following platforms: Cray XE6, Cray XE6m, Cray XE5, Cray XE5m, Cray XT6, Cray XT6m, Cray XT5, Cray XT5m, and Cray XT4 systems.
- Update** installations from CLE 3.1.UP00, CLE 3.1.UP01, and CLE 3.1.UP02.

- **Upgrade** installations from CLE 3.0 on Cray XT6 systems.
- **Upgrade** installations from CLE 3.1 for SeaStar to CLE 3.1 for Gemini, on Cray XT6, Cray XT6m, Cray XT5, or Cray XT5m systems that have been hardware upgraded to Cray XE systems.

Note: Due to hardware changes, this type of installation is a partial upgrade that includes a number of manual migration steps.

- **Migration** installations from CLE 2.2 on Cray XT5, Cray XT5m, and Cray XT4 systems.

Note: This type of installation requires an initial installation with modifications to preserve Lustre file systems.

- **Migration** installations from CLE 2.2 on Cray XT5 or Cray XT5m systems that have been hardware upgraded to upgraded to Cray XE systems.

Note: This type of installation requires an initial installation with modifications to preserve Lustre file systems and reconfigure CLE for Gemini.

System administrators may want to advise users on compatibility of binaries between the various migration, update, and upgrade paths. Please refer to *Cray Linux Environment (CLE) Software Release Overview* for more information on binary compatibility when installing CLE 3.1.

2.3 Binary Compatibility

The *Cray Linux Environment (CLE) 3.1 Software Release Overview* (S-2425-31) states that application binaries built using dynamic libraries do not require recompiling. The language in that statement remains accurate. However, you must recompile applications that are written in C++ or Partitioned Global Address Space (PGAS) programming models when you migrate applications.

2.4 Additional Software Requirements

2.4.1 Release Level Requirements for Other Cray Software Products

Cray Linux Environment (CLE) 3.1 Software Release Overview (S–2425–31) includes a table of other Cray software products supported with CLE 3.1 and the minimum release level required for verified compatibility. The information in this table applies to CLE 3.1 update releases with the following exceptions:

- Cray recommends that you upgrade the System Management Workstation (SMW) to the SMW 5.1.UP03 release before you install the CLE 3.1.UP03 update release package. For more information, see the *README* document included with the SMW release package.
- You must upgrade the Cray Application Developer's Environment (CADE) to the most current version (at time of release, this is 5.08) to ensure compatibility with the CLE 3.1.UP02 release package. For CADE release information, see *Cray Application Developer's Environment Installation Guide* (S–2465).

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.

2.4.1.1 Programming Environments Requirements for Hardware Upgrades

You must uninstall and reinstall the following programming environment software after CLE upgrade/migrations that include a hardware upgrade of a Cray XT system running CLE 3.0 to 3.1 to a Cray XE system running CLE 3.1:

- PGI
- Intel
- Cray Performance Tools
- Cray MPT (Message Passing Toolkit) — Cray MPT 5.0 or later is required for Gemini-based systems
- Chapel

For upgrades to the Gemini-based interconnect, Cray recommends you remove the following:

- SeaStar RPMs for Cray MPT
- Portals RPMs

Cray also recommends that you use the most recent version level of the Cray Compiling Environment as the default.

2.4.2 Third-party Software Requirements

Cray Linux Environment (CLE) 3.1 Software Release Overview (S-2425-31) includes a section that lists third-party software requirements for the CLE 3.1 release; this information applies to CLE 3.1 update releases with the following exceptions:

- You must upgrade the PGI Compiler to version 10.9 or later to ensure compatibility with the CLE 3.1 release. PGI release information is available from The Portland Group, Inc. at <http://www.pgroup.com>.
- The PathScale Compiler Suite version 4.0.7 is compatible with SLES 11. However, Cray Message Passing Toolkit (MPT) and `libsci` libraries built with this compiler are not yet available. PathScale release information is available from PathScale Inc. at <http://www.pathscale.com>. MPT and `libsci` release information is available on the CrayDoc website (docs.cray.com) in the *Programming Environments Release Announcement*.
- 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.

Cray recommends that you continue to upgrade these products as new versions become available.

3.1 Cray-developed Books Provided with This Release

[Table 1](#) list the books provided with the CLE 3.1 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) 3.1 Software Release Overview* (S-2425-31), which is also provided with the release package.

Table 1. Books Provided with This Release

Book Title	Number	Updated
<i>Cray Linux Environment (CLE) Software Release Overview Supplement</i> (this document)	S-2497-3103	Yes
<i>Cray Linux Environment (CLE) Software Release Overview</i>	S-2425-31	No
<i>Installing and Configuring Cray Linux Environment (CLE) Software</i>	S-2444-3103	Yes
<i>Managing System Software for Cray XE and Cray XT Systems</i>	S-2393-3103	Yes
<i>Managing Lustre for the Cray Linux Environment (CLE)</i>	S-0010-3103	Yes
<i>Introduction to Cray Data Virtualization Service</i>	S-0005-3102	No
<i>Writing a Node Health Checker (NHC) Plugin Test</i>	S-0023-3103	Yes
<i>Workload Management and Application Placement for the Cray Linux Environment</i>	S-2496-3103	Yes
<i>Using the GNI and DMAPP APIs</i>	S-2446-3103	Yes
<i>Repurposing Compute Nodes as Service Nodes on Cray XE and Cray XT Systems</i>	S-0029-3101	No
<i>Lustre Operations Manual</i>	S-6540-1815	Yes
<i>CrayDoc Installation and Administration Guide</i>	S-2340-411	No

Book Title	Number	Updated
<i>Upgrading a Cray XT System with Cray Linux Environment (CLE) 3.1 to a Cray XE System with CLE 3.1</i>	S–0020–3103	New
<i>Migrating a Cray XT System with Cray Linux Environment (CLE) 2.2 to a Cray XE System with CLE 3.1</i>	S–0035–3103	New
<i>Migrating to Cray Linux Environment (CLE) 3.1 Software for Cray XT Systems</i>	S–2489–3103	New

3.1.1 CLE Installation and Migration Publications

The following publications provide procedures for installing CLE 3.1 software, including new publications for upgrade and migration scenarios.

- If you have a new system, follow *Installing and Configuring Cray Linux Environment (CLE) Software* (S–2444) to perform an initial installation of the CLE software.
- If you have a system that is already running CLE 3.1, follow *Installing and Configuring Cray Linux Environment (CLE) Software* (S–2444) to install a new CLE 3.1 update release.
- If you have a Cray XT6 system running CLE 3.0, follow *Installing and Configuring Cray Linux Environment (CLE) Software* (S–2444) to upgrade to CLE 3.1.
- If you have a Cray XT5, Cray XT5m, or Cray XT4 system running CLE 2.2, follow *Migrating to Cray Linux Environment (CLE) 3.1 Software for Cray XT Systems* (S–2489) to install CLE 3.1 for SeaStar systems.
- If you have a Cray XT6, Cray XT6m, Cray XT5, or Cray XT5m system running CLE 3.1 and your system is being hardware upgraded to a Cray XE system, follow *Upgrading a Cray XT System with Cray Linux Environment (CLE) 3.1 to a Cray XE System with CLE 3.1* (S–0020) to move to CLE 3.1 for Gemini systems.
- If you have a Cray XT5 or Cray XT5m system running CLE 2.2 and your system is being hardware upgraded to a Cray XE system, follow *Migrating a Cray XT System with Cray Linux Environment (CLE) 2.2 to a Cray XE System with CLE 3.1* (S–0035) and specific sections in *Installing and Configuring Cray Linux Environment (CLE) Software* (S–2444) as directed.

3.2 Changes to Man Pages

3.2.1 New Cray Man Pages

3.2.1.1 CLE 3.1.UP03

- `dumpd(8)`: Initiates automatic dump and reboot of nodes when requested by Node Health Checker (NHC).
- `dumpd-request(8)`: (Deferred implementation) This man page will be provided in a future release and/or in the release package. Use the `-h` option with the command to display a help message.
- `dumpd-dbadmin(8)`: (Deferred implementation) This man page will be provided in a future release and/or in the release package. Use the `-h` option with the command to display a help message.

3.2.1.2 CLE 3.1.UP02

- `xtverifydefaults(8)`: Verifies and fixes inconsistent system default links within the shared root.

3.2.1.3 CLE 3.1.UP01

- `xtorder2db(8)`: Converts a text file to values in the `od_allocator_id` field of the processor table in the Service Database (SDB).
- `xtdb2order(8)`: Converts the processor table `od_allocator_id` field in the SDB to a text file.

3.2.2 Changed Cray Man Pages

3.2.2.1 CLE 3.1.UP03

- `aprun(1)`: Clarified MPMD section regarding use of the `-m` option.
- `ccmrun(1)`: Added `--ssh`, `--nossh`, `--rsh`, `--norsh`, `--ncsd`, `--noncsd`, `--rsip`, `--norsip`, and `--help` options.
- `rsipd(8)`: Updated for logging enhancements.
- `intro_NHC(8)`: Included a description of the `dumpd` interface.
- `lustre_control.sh(8)`: Added new `-c`, `-m`, and `-n` options.
- `dvs(5)`: Added `bulk_rw` option and environment variable; added `ioctl` documentation.
- `xtoprdump(8)`: Corrected definitions for ITEM and FILE specifications; added file specification example.
- `CLEinstall.conf(5)`: Updated text to reflect CLEinstall enhancements; added `node_class[idx]` examples; added batch parameters for CCM; added `CNL_writeable_tmp` parameter.

3.2.2.2 CLE 3.1.UP02

- `apbridge(8)`: Updated to add `-Oy` and other NID ordering clarifications from CLE 3.1.UP01.
- `aprun(1)`: Updated hugepage information; updated `-r` option; added text for `-T`; revised `-t` option text.
- `apstat(1)`: Updated with text to describe unbound and `pendCancel States`.
- `xtopview(8)`: Added `xtverifydefaults` exit processing information.
- `xtcheckhealth(8)`: Updated new and changed exit codes in the RETURN VALUES section.
- `xtcleanup_after(8)`: Updated new and changed exit codes in the RETURN VALUES section.
- `rsipd.conf(5)`: Changed `initial_lease_time` option default setting to 60; system administrators may wish to change the value of `initial_lease_time` in `/etc/opt/cray/rsipd/rsipd.conf` from its default of 60 seconds to a larger value (600 seconds is suggested) to strike a balance between clients renewing their leases too often and clients taking a long time to recover after an RSIP server outage.

3.2.2.3 CLE 3.1.UP01

- `aprun(1)`: Included descriptions of the following environment variables: `APRUN_DEFAULT_MEMORY`, `APRUN_XFER_LIMITS` and `APRUN_SYNC_TTY`.
- `apbridge(8)`: Included the new `-O2` option for the `ALPS_NIDORDER` parameter to provide the "xyz-by2" NID reordering method.
- `basil(7)`: Corrected `ENGINE` query output; marked `-s1`, `-sn`, and `-S` option not for exclusive use; changed `reservation_mode` default to `unset`, use `apsched` default; changed `gemini` in 1.1 to `router_id`.
- `CLEinstall.conf(5)`: Included new default network mask for Cray XE systems; clarified that `sdbnode_failover_interface` must be set even if SDB node failover is not configured.
- `xtcleanup_after(8)`: Included the new `-i` option to allow callers to specify a string in place of the APID when printing to the console.
- `xtnodestat(8)`: Added a detailed description of the legend; included explanation of how repurposed compute nodes are displayed.