

HP XC System Software

Release Notes

Version 3.2



© Copyright 2007, 2008 Hewlett-Packard Development Company, L.P.

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

AMD and AMD Opteron are trademarks or registered trademarks of Advanced Micro Devices, Inc.

FLEXlm is a trademark of Macrovision Corporation.

InfiniBand is a registered trademark and service mark of the InfiniBand Trade Association.

Intel, Itanium, and Xeon are trademarks or registered trademarks of Intel Corporation in the United States and other countries.

Linux is a U.S. registered trademark of Linus Torvalds.

LSF and Platform Computing are trademarks or registered trademarks of Platform Computing Corporation.

Myrinet and Myricom are registered trademarks of Myricom, Inc.

Nagios is a registered trademark of Ethan Galstad.

The Portland Group and PGI are trademarks or registered trademarks of The Portland Group Compiler Technology, STMicroelectronics, Inc.

Quadrics and QsNet^{II} are registered trademarks of Quadrics, Ltd.

Red Hat and RPM are registered trademarks of Red Hat, Inc.

syslog-ng is copyrighted by BalaBit IT Security.

SystemImager is a registered trademark of Brian Finley.

TotalView is a registered trademark of Etnus, Inc.

UNIX is a registered trademark of The Open Group.

Table of Contents

About This Document.....	7
Intended Audience.....	7
Typographic Conventions.....	7
HP XC and Related HP Products Information.....	8
Related Information.....	9
Manpages.....	12
HP Encourages Your Comments.....	13
1 New and Changed Features.....	15
1.1 Base Distribution and Kernel.....	15
1.2 Support for Additional Hardware Models.....	15
1.3 OpenFabrics Enterprise Distribution for InfiniBand.....	15
1.4 HP Scalable Visualization Array.....	15
1.5 Partition Size Limits on Installation Disk.....	16
1.6 More Flexibility in Customizing Client Node Disk Partitions.....	16
1.7 Enhancements to the discover Command.....	16
1.8 Enhancements to the cluster_config Utility.....	16
1.9 System Management and Monitoring Enhancements.....	17
1.10 Enhancements to the OVP.....	17
1.11 Installing and Upgrading HP XC System Software On Red Hat Enterprise Linux.....	17
1.12 Support For HP Unified Parallel C.....	17
1.13 Documentation Changes.....	18
2 Important Release Information.....	19
2.1 Firmware Versions.....	19
2.2 Patches.....	19
3 Hardware Preparation.....	21
3.1 Upgrading BMC Firmware On HP ProLiant DL140 G2 and DL145 G2 Nodes.....	21
4 Software Installation On The Head Node.....	23
4.1 Manual Installation Required For NC510F Driver.....	23
5 System Discovery, Configuration, and Imaging.....	25
5.1 Notes That Apply Before You Invoke The cluster_prep Utility.....	25
5.1.1 Required Task for Some NIC Adapter Models: Verify Correct NIC Device Driver Mapping.....	25
5.2 Notes That Apply To The Discover Process.....	26
5.2.1 Discovery of HP ProLiant DL140 G3 and DL145 G3 Nodes Fails When Graphics Cards Are Present.....	26
5.3 Notes That Apply Before You Invoke The cluster_config Utility.....	26
5.3.1 Adhere To Role Assignment Guidelines for Improved Availability.....	26
5.4 Benign Message From C52xcgraph During cluster_config.....	26
5.5 Processing Time For cluster_config Might Take Longer On A Head Node With Improved Availability.....	27
5.6 Notes That Apply To Imaging.....	27

5.6.1 HP ProLiant DL140 G3 and DL145 G3 Node Imaging Fails When Graphics Cards Are Present.....	27
6 Software Upgrades.....	29
6.1 Do Not Upgrade If You Want Or Require The Voltaire InfiniBand Software Stack.....	29
7 System Administration, Management, and Monitoring.....	31
7.1 Perform A Dry Run Before Using The si_updateclient Utility To Update Nodes.....	31
7.2 Possible Problem With ext3 File Systems On SAN Storage.....	31
8 HP XC System Software On Red Hat Enterprise Linux.....	33
8.1 Enabling 32-bit Applications To Compile and Run	33
9 Programming and User Environment.....	35
9.1 MPI and OFED InfiniBand Stack Fork Restrictions.....	35
9.2 InfiniBand Multiple Rail Support.....	35
9.3 Benign Messages From HP-MPI Version 2.2.5.1.....	35
10 Cluster Platform 3000.....	37
11 Cluster Platform 4000.....	39
12 Cluster Platform 6000.....	41
12.1 Network Boot Operation and Imaging Failures on HP Integrity rx2600 Systems.....	41
12.2 Notes That Apply To The Management Processor.....	41
12.2.1 Required Task: Change MP Settings on Console Switches.....	41
12.2.2 MP Disables DHCP Automatically.....	41
12.2.3 Finding the IP Address of an MP.....	41
13 Integrated Lights Out Console Management Devices.....	43
13.1 iLO2 Devices In Server Blades Can Hang.....	43
14 Interconnects.....	45
14.1 InfiniBand Interconnect.....	45
14.1.1 enable Password Problem With Voltaire Switch Version 4.1.....	45
14.2 Myrinet Interconnect.....	45
14.2.1 Myrinet Monitoring Line Card Can Become Unresponsive.....	45
14.2.2 The clear_counters Command Does Not Work On The 256 Port Switch.....	45
14.3 QsNet ^{II} Interconnect.....	45
14.3.1 Possible Conflict With Use of SIGUSR2.....	46
14.3.2 The qsnet Database Might Contain Entries To Nonexistent Switch Modules.....	46
15 Documentation.....	47
15.1 Documentation CD Search Option.....	47
15.2 HP XC Manpages.....	47
15.2.1 New device_config.8.....	47
15.2.2 Changes to ovp.8.....	47

15.2.3 New preupgradesys-lxc.8.....47
15.2.4 New upgradesys-lxc.8.....48

Index.....51

About This Document

This document contains release notes for HP XC System Software Version 3.2. This document contains important information about firmware, software, or hardware that might affect the system.

An HP XC system is integrated with several open source software components. Some open source software components are being used for underlying technology, and their deployment is transparent. Some open source software components require user-level documentation specific to HP XC systems, and that kind of information is included in this document when required.

HP relies on the documentation provided by the open source developers to supply the information you need to use their product. For links to open source software documentation for products that are integrated with the HP XC system, see “Supplementary Software Products” (page 9).

Documentation for third-party hardware and software components that are supported on the HP XC system is supplied by the third-party vendor. However, information about the operation of third-party software is included in this document if the functionality of the third-party component differs from standard behavior when used in the XC environment. In this case, HP XC documentation supersedes information supplied by the third-party vendor. For links to related third-party Web sites, see “Supplementary Software Products” (page 9).

Standard Linux® administrative tasks or the functions provided by standard Linux tools and commands are documented in commercially available Linux reference manuals and on various Web sites. For more information about obtaining documentation for standard Linux administrative tasks and associated topics, see the list of Web sites and additional publications provided in “Related Software Products and Additional Publications” (page 11).

Intended Audience

The release notes are intended for anyone who installs and configures an HP XC system, for system administrators who maintain the system, for programmers who write applications to run on the system, and for general users who log in to the system to run jobs.

The information in this document assumes that you have knowledge of the Linux operating system.

Typographic Conventions

This document uses the following typographical conventions:

<code>%</code> , <code>\$</code> , or <code>#</code>	A percent sign represents the C shell system prompt. A dollar sign represents the system prompt for the Korn, POSIX, and Bourne shells. A number sign represents the superuser prompt.
<code>audit(5)</code>	A manpage. The manpage name is <code>audit</code> , and it is located in Section 5.
Command	A command name or qualified command phrase.
Computer output	Text displayed by the computer.
Ctrl+x	A key sequence. A sequence such as Ctrl+x indicates that you must hold down the key labeled Ctrl while you press another key or mouse button.
ENVIRONMENT VARIABLE	The name of an environment variable, for example, <code>PATH</code> .
[ERROR NAME]	The name of an error, usually returned in the <code>errno</code> variable.
Key	The name of a keyboard key. Return and Enter both refer to the same key.
Term	The defined use of an important word or phrase.
User input	Commands and other text that you type.

<i>Variable</i>	The name of a placeholder in a command, function, or other syntax display that you replace with an actual value.
[]	The contents are optional in syntax. If the contents are a list separated by , you can choose one of the items.
{ }	The contents are required in syntax. If the contents are a list separated by , you must choose one of the items.
...	The preceding element can be repeated an arbitrary number of times.
	Separates items in a list of choices.
WARNING	A warning calls attention to important information that if not understood or followed will result in personal injury or nonrecoverable system problems.
CAUTION	A caution calls attention to important information that if not understood or followed will result in data loss, data corruption, or damage to hardware or software.
IMPORTANT	This alert provides essential information to explain a concept or to complete a task.
NOTE	A note contains additional information to emphasize or supplement important points of the main text.

HP XC and Related HP Products Information

The HP XC System Software Documentation Set, the Master Firmware List, and HP XC HowTo documents are available at this HP Technical Documentation Web site:

<http://www.docs.hp.com/en/linuxhpc.html>

The HP XC System Software Documentation Set includes the following core documents:

<i>HP XC System Software Release Notes</i>	Describes important, last-minute information about firmware, software, or hardware that might affect the system. This document is not shipped on the HP XC documentation CD. It is available only on line.
<i>HP XC Hardware Preparation Guide</i>	Describes hardware preparation tasks specific to HP XC that are required to prepare each supported hardware model for installation and configuration, including required node and switch connections.
<i>HP XC System Software Installation Guide</i>	Provides step-by-step instructions for installing the HP XC System Software on the head node and configuring the system.
<i>HP XC System Software Administration Guide</i>	Provides an overview of the HP XC system administrative environment, cluster administration tasks, node maintenance tasks, LSF® administration tasks, and troubleshooting procedures.
<i>HP XC System Software User's Guide</i>	Provides an overview of managing the HP XC user environment with modules, managing jobs with LSF, and describes how to build, run, debug, and troubleshoot serial and parallel applications on an HP XC system.
QuickSpecs for HP XC System Software	Provides a product overview, hardware requirements, software requirements, software licensing information, ordering information, and information about commercially available software that has been qualified to interoperate with the HP XC System Software. The QuickSpecs are located on line: http://www.hp.com/go/clusters

See the following sources for information about related HP products.

HP XC Program Development Environment

The Program Development Environment home page provide pointers to tools that have been tested in the HP XC program development environment (for example, TotalView® and other debuggers, compilers, and so on).

<http://h20311.www2.hp.com/HPC/cache/276321-0-0-0-121.html>

HP Message Passing Interface

HP Message Passing Interface (HP-MPI) is an implementation of the MPI standard that has been integrated in HP XC systems. The home page and documentation is located at the following Web site:

<http://www.hp.com/go/mpi>

HP Serviceguard

HP Serviceguard is a service availability tool supported on an HP XC system. HP Serviceguard enables some system services to continue if a hardware or software failure occurs. The HP Serviceguard documentation is available at the following Web site:

<http://www.docs.hp.com/en/ha.html>

HP Scalable Visualization Array

The HP Scalable Visualization Array (SVA) is a scalable visualization solution that is integrated with the HP XC System Software. The SVA documentation is available at the following Web site:

<http://www.docs.hp.com/en/linuxhpc.html>

HP Cluster Platform

The cluster platform documentation describes site requirements, shows you how to set up the servers and additional devices, and provides procedures to operate and manage the hardware. These documents are available at the following Web site:

<http://www.docs.hp.com/en/linuxhpc.html>

HP Integrity and HP ProLiant Servers

Documentation for HP Integrity and HP ProLiant servers is available at the following Web site:

<http://www.docs.hp.com/en/hw.html>

Related Information

This section provides useful links to third-party, open source, and other related software products.

Supplementary Software Products This section provides links to third-party and open source software products that are integrated into the HP XC System Software core technology. In the HP XC documentation, except where necessary, references to third-party and open source software components are generic, and the HP XC adjective is not added to any reference to a third-party or open source command or product name. For example, the SLURM `srun` command is simply referred to as the `srun` command.

The location of each Web site or link to a particular topic listed in this section is subject to change without notice by the site provider.

- <http://www.platform.com>

Home page for Platform Computing Corporation, the developer of the Load Sharing Facility (LSF). LSF-HPC with SLURM, the batch system resource manager used on an HP XC system, is tightly integrated with the HP XC and SLURM software. Documentation specific to LSF-HPC with SLURM is provided in the HP XC documentation set.

Standard LSF is also available as an alternative resource management system (instead of LSF-HPC with SLURM) for HP XC. This is the version of LSF that is widely discussed on the Platform Web site.

For your convenience, the following Platform Computing Corporation LSF documents are shipped on the HP XC documentation CD in PDF format:

- *Administering Platform LSF*
- *Administration Primer*
- *Platform LSF Reference*
- *Quick Reference Card*
- *Running Jobs with Platform LSF*

LSF procedures and information supplied in the HP XC documentation, particularly the documentation relating to the LSF-HPC integration with SLURM, supersedes the information supplied in the LSF manuals from Platform Computing Corporation.

The Platform Computing Corporation LSF manpages are installed by default. *lsf_diff(7)* supplied by HP describes LSF command differences when using LSF-HPC with SLURM on an HP XC system

The following documents in the HP XC System Software Documentation Set provide information about administering and using LSF on an HP XC system:

- *HP XC System Software Administration Guide*
- *HP XC System Software User's Guide*

- <http://www.llnl.gov/LCdocs/slurm/>

Documentation for the Simple Linux Utility for Resource Management (SLURM), which is integrated with LSF to manage job and compute resources on an HP XC system.

- <http://www.nagios.org/>

Home page for Nagios®, a system and network monitoring application that is integrated into an HP XC system to provide monitoring capabilities. Nagios watches specified hosts and services and issues alerts when problems occur and when problems are resolved.

- <http://oss.oetiker.ch/rrdtool>

Home page of RRDtool, a round-robin database tool and graphing system. In the HP XC system, RRDtool is used with Nagios to provide a graphical view of system status.

- <http://supermon.sourceforge.net/>

Home page for Supermon, a high-speed cluster monitoring system that emphasizes low perturbation, high sampling rates, and an extensible data protocol and programming interface. Supermon works in conjunction with Nagios to provide HP XC system monitoring.

- <http://www.llnl.gov/linux/pdsh/>

Home page for the parallel distributed shell (pdsh), which executes commands across HP XC client nodes in parallel.

- http://www.balabit.com/products/syslog_ng/

Home page for syslog-ng, a logging tool that replaces the traditional syslog functionality. The syslog-ng tool is a flexible and scalable audit trail processing tool. It provides a centralized, securely stored log of all devices on the network.

- <http://systemimager.org>

Home page for SystemImager®, which is the underlying technology that distributes the golden image to all nodes and distributes configuration changes throughout the system.

- <http://linuxvirtualserver.org>
Home page for the Linux Virtual Server (LVS), the load balancer running on the Linux operating system that distributes login requests on the HP XC system.
- <http://www.macrovision.com>
Home page for Macrovision®, developer of the FLEXlm™ license management utility, which is used for HP XC license management.
- <http://sourceforge.net/projects/modules/>
Web site for Modules, which provide for easy dynamic modification of a user's environment through modulefiles, which typically instruct the `module` command to alter or set shell environment variables.
- <http://dev.mysql.com/>
Home page for MySQL AB, developer of the MySQL database. This Web site contains a link to the MySQL documentation, particularly the *MySQL Reference Manual*.

Related Software Products and Additional Publications This section provides pointers to Web sites for related software products and provides references to useful third-party publications. The location of each Web site or link to a particular topic is subject to change without notice by the site provider.

Linux Web Sites

- <http://www.redhat.com>
Home page for Red Hat®, distributors of Red Hat Enterprise Linux Advanced Server, a Linux distribution with which the HP XC operating environment is compatible.
- <http://www.linux.org/docs/index.html>
This Web site for the Linux Documentation Project (LDP) contains guides that describe aspects of working with Linux, from creating your own Linux system from scratch to bash script writing. This site also includes links to Linux HowTo documents, frequently asked questions (FAQs), and manpages.
- <http://www.linuxheadquarters.com>
Web site providing documents and tutorials for the Linux user. Documents contain instructions for installing and using applications for Linux, configuring hardware, and a variety of other topics.
- <http://www.gnu.org>
Home page for the GNU Project. This site provides online software and information for many programs and utilities that are commonly used on GNU/Linux systems. Online information include guides for using the bash shell, emacs, make, cc, gdb, and more.

MPI Web Sites

- <http://www.mpi-forum.org>
Contains the official MPI standards documents, errata, and archives of the MPI Forum. The MPI Forum is an open group with representatives from many organizations that define and maintain the MPI standard.
- <http://www-unix.mcs.anl.gov/mpi/>
A comprehensive site containing general information, such as the specification and FAQs, and pointers to other resources, including tutorials, implementations, and other MPI-related sites.

Compiler Web Sites

- <http://www.intel.com/software/products/compilers/index.htm>
Web site for Intel® compilers.
- <http://support.intel.com/support/performance/tools/>
Web site for general Intel software development information.
- <http://www.pgroup.com/>
Home page for The Portland Group™, supplier of the PGI® compiler.

Debugger Web Site

<http://www.etnus.com>

Home page for Etnus, Inc., maker of the TotalView® parallel debugger.

Software RAID Web Sites

- <http://www.tldp.org/HOWTO/Software-RAID-HOWTO.html> and <http://www.ibiblio.org/pub/Linux/docs/HOWTO/other-formats/pdf/Software-RAID-HOWTO.pdf>
A document (in two formats: HTML and PDF) that describes how to use software RAID under a Linux operating system.
- <http://www.linuxdevcenter.com/pub/a/linux/2002/12/05/RAID.html>
Provides information about how to use the mdadm RAID management utility.

Additional Publications

For more information about standard Linux system administration or other related software topics, consider using one of the following publications, which must be purchased separately:

- *Linux Administration Unleashed*, by Thomas Schenk, et al.
- *Linux Administration Handbook*, by Evi Nemeth, Garth Snyder, Trent R. Hein, et al.
- *Managing NFS and NIS*, by Hal Stern, Mike Eisler, and Ricardo Labiaga (O'Reilly)
- *MySQL*, by Paul Debois
- *MySQL Cookbook*, by Paul Debois
- *High Performance MySQL*, by Jeremy Zawodny and Derek J. Balling (O'Reilly)
- *Perl Cookbook, Second Edition*, by Tom Christiansen and Nathan Torkington
- *Perl in A Nutshell: A Desktop Quick Reference*, by Ellen Siever, et al.

Manpages

Manpages provide online reference and command information from the command line. Manpages are supplied with the HP XC system for standard HP XC components, Linux user commands, LSF commands, and other software components that are distributed with the HP XC system.

Manpages for third-party software components might be provided as a part of the deliverables for that component.

Using *discover(8)* as an example, you can use either one of the following commands to display a manpage:

```
$ man discover
$ man 8 discover
```

If you are not sure about a command you need to use, enter the man command with the `-k` option to obtain a list of commands that are related to a keyword. For example:

```
$ man -k keyword
```

HP Encourages Your Comments

HP encourages comments concerning this document. We are committed to providing documentation that meets your needs. Send any errors found, suggestions for improvement, or compliments to:

feedback@fc.hp.com

Include the document title, manufacturing part number, and any comment, error found, or suggestion for improvement you have concerning this document.

1 New and Changed Features

This chapter describes the new and changed features delivered in HP XC System Software Version 3.2.

1.1 Base Distribution and Kernel

The following table lists information about the base distribution and kernel for this release as compared to the last HP XC release.

HP XC Version 3.2	HP XC Version 3.1
Enterprise Linux 4 Update 4	Enterprise Linux 4 Update 3
HP XC kernel version 2.6.9-42.9hp.XC	HP XC kernel version 2.6.9-34.7hp.XC
Based on Red Hat kernel version 2.6.9-42.0.8.EL	Based on Red Hat kernel version 2.6.9-34.0.2.EL

1.2 Support for Additional Hardware Models

In this release, the following additional hardware models and hardware components are supported in an HP XC hardware configuration.

- HP ProLiant servers:
 - HP ProLiant DL360 G5
 - HP ProLiant DL380 G5
 - HP ProLiant DL580 G4
 - HP ProLiant DL145 G3
 - HP ProLiant DL385 G2
 - HP ProLiant DL585 G2
- HP Integrity servers and workstations:
 - HP Integrity rx2660
 - HP Integrity rx4640
 - HP xw9400 workstation

1.3 OpenFabrics Enterprise Distribution for InfiniBand

Starting with this release, the HP XC System Software uses the OpenFabrics Enterprise Distribution (OFED) InfiniBand software stack.

OFED is an open software stack supported by the major InfiniBand vendors as the future of InfiniBand support. OFED offers improved support of multiple HCAs per node. The OFED stack has a different structure and different commands from the InfiniBand stack that was used in previous HP XC releases.

See the following web page for more information about OFED:

<http://www.openfabrics.org/>

The *HP XC System Software Administration Guide* provides OFED troubleshooting information.

1.4 HP Scalable Visualization Array

HP Scalable Visualization Array (SVA) software is now included on the HP XC System Software DVD distribution media. SVA provides a comprehensive set of services for deployment of visualization applications, allowing them to be conveniently run in a Linux clustering environment.

The following are the key features of SVA:

- Capturing and managing visualization-specific cluster information
- Managing visualization resources and providing facilities for requesting and allocating resources for a job in a multi-user, multi-session environment
- Providing display surface configuration tools to allow easy configuration of multi-panel displays
- Providing launch tools, both generic and tailored to a specific application, that launch applications with appropriate environments and display surface configurations
- Providing tools that extend serial applications to run in a clustered, multi-display environment

See the HP XC *QuickSpecs* and the SVA documentation set for more information about SVA features. The SVA documentation set is included on the HP XC Documentation CD.

Because the SVA RPMs are included on the HP XC distribution media, the SVA installation process has been integrated with the HP XC installation process. The *HP XC System Software Installation Guide* was revised where appropriate to accommodate SVA installation and configuration procedures.

1.5 Partition Size Limits on Installation Disk

Because the installation disk size can vary, partition sizes are calculated as a percentage of total disk size. However, using a fixed percentage of the total disk size to calculate the size of each disk partition can result in needlessly large partition sizes when the installation disk is larger than 36 GB. Thus, for this release, limits have been set on the default partition sizes to leave space on the disk for other user-defined file systems and partitions.

1.6 More Flexibility in Customizing Client Node Disk Partitions

You can configure client node disks on a per-image and per-node basis to create an optional scratch partition to maximize file system performance. Partition sizes can be fixed or they can be based on a percentage of total disk size. To do so, you set the appropriate variables in the `/opt/hptc/systemimager/etc/make_partitions.sh` file or set the variables in user-defined files with a `.part` extension.

The procedure that describes how to customize client node disk partitions is documented in the *HP XC System Software Installation Guide*.

1.7 Enhancements to the `discover` Command

. The following options were added to the `discover` command:

- The `--nodesonly` option reads in the database and discover all nodes if the hardware configuration contains HP server blades and enclosures. This option is valid only when the `--enclosurebased` option is also used
- The `--nothreads` option runs the node discovery process without threads if the hardware configuration contains HP server blades and enclosures. This option is valid only when the `--enclosurebased` option is also used.

1.8 Enhancements to the `cluster_config` Utility

The `cluster_config` utility prompts you to specify whether you want to configure the Linux virtual server (LVS) director to act as a real server, that is, a node that accepts login sessions.

If you answer **yes**, the LVS director is configured to act as a login session server in addition to arbitrating and dispersing the login session connections.

If you answer **no**, the LVS director does not participate as a login session server; its only function is to arbitrate and disperse login sessions to other nodes. This gives you the flexibility to place

the `login` role on the head node yet keep the head node load to a minimum because login sessions are not being spawned.

This configuration choice is documented in the *HP XC System Software Installation Guide*.

1.9 System Management and Monitoring Enhancements

System management and monitoring utilities have been enhanced as follows:

- A new resource monitoring tool, `resmon`, has been added. `resmon` is a job-centric resource monitoring Web page initially inspired by the open-source `clumon` product. `resmon` invokes useful commands to collect and present data in a scalable and intuitive fashion. The `resmon` Web pages update automatically at a preconfigured interval (120 seconds by default). See `resmon(1)` for more information.
- The HP Graph Web interface has been enhanced to include a `cpu temperature` graph. To access this new graph, select **temperature** from the **Metrics** pull-down menu at the top of the Web page.

1.10 Enhancements to the OVP

The operation verification program (OVP) performance health tests were updated to accept an option to specify an LSF queue. In addition, you can run two performance health tests, `network_stress` and `network_bidirectional`, on systems that are configured with standard LSF or configured with LSF-HPC with SLURM.

1.11 Installing and Upgrading HP XC System Software On Red Hat Enterprise Linux

The *HP XC System Software Installation Guide* contains two new chapters that describes the following topics:

- Installing HP XC System Software Version 3.2 on Red Hat Enterprise Linux
- Upgrading HP XC System Software Version 3.1 on Red Hat Enterprise Linux to HP XC System Software Version 3.2 on Red Hat Enterprise Linux

1.12 Support For HP Unified Parallel C

This release provides support for the HP Unified Parallel C (UPC) application development environment.

HP UPC is a parallel extension of the C programming language, which runs on both common types of multiprocessor systems: those with a common global address space (such as SMP) and those with distributed memory. UPC provides a simple shared memory model for parallel programming, allowing data to be shared or distributed among a number of communicating processors. Constructs are provided in the language to permit simple declaration of shared data, distribute shared data across threads, and synchronize access to shared data across threads. This model promises significantly easier coding of parallel applications and maximum performance across shared memory, distributed memory, and hybrid systems.

See the following Web page for more information about HP UPC:

<http://www.hp.com/go/upc>

1.13 Documentation Changes

The following changes were made to the HP XC System Software Documentation Set

- The following manuals have been affected by the new functionality delivered in this release and have been revised accordingly:
 - *HP XC Hardware Preparation Guide*
 - *HP XC System Software Installation Guide*
 - *HP XC System Software Administration Guide*
 - *HP XC System Software User's Guide*
- The information in the *Configuring HP XC Systems With HP Server Blades and Enclosures - Edition 9* HowTo was merged into the *HP XC Hardware Preparation Guide* and *HP XC System Software Installation Guide*, reducing the number of documents you have to read to install and configure an HP XC system that contains HP server blades and enclosures.
- The *HP XC System Software Release Notes* are updated periodically. Therefore, HP recommends that you go to <http://www.docs.hp.com/en/linuxhpc.html> and make sure you have the latest version of this document because the version you are reading now might have been updated since the last time you downloaded it.

HP XC HowTos On the Worldwide Web

HP XC information that is published between releases is issued in HowTo documents at the following Web site:

<http://www.docs.hp.com/en/linuxhpc.html>

2 Important Release Information

This chapter contains information that is important to know for this release.

2.1 Firmware Versions

The HP XC System Software is tested against specific minimum firmware versions. Follow the instructions in the accompanying hardware documentation to ensure that all hardware components are installed with the latest firmware version.

The master firmware tables for this release are available at the following Web site:

<http://www.docs.hp.com/en/linuxhpc.html>

The master firmware tables list the minimum firmware versions on which the Version 3.2 HP XC System Software has been qualified. At a minimum, the HP XC system components must be installed with these versions of the firmware.

Read the following guidelines before upgrading the firmware on any component in the hardware configuration:

- **Never** downgrade to an older version of firmware unless you are specifically instructed to do so by the HP XC Support Team.
- The master firmware tables clearly indicate newer versions of the firmware that are known to be incompatible with the HP XC software. Incompatible versions are highlighted in **bold** font. **Do not install these known incompatible firmware versions because unexpected system behavior might occur.**
- There is always the possibility that a regression in functionality is introduced in a firmware version. It is possible that the regression could cause anomalies in HP XC operation. Report regressions in HP XC operation that result from firmware upgrades to the HP XC Support Team:
xc_support@hp.com
- Contact the HP XC Support Team if you are not sure what to do regarding firmware versions.

2.2 Patches

Software patches might be available for this release. Because network connectivity is not established during a new installation until the `cluster_prep` utility has finished preparing the system, you are instructed to download the patches when you reach that point in the installation and configuration process. The *HP XC System Software Installation Guide* provides more information about where to access and download software patches.

3 Hardware Preparation

Hardware preparation tasks are documented in the *HP XC Hardware Preparation Guide*. This chapter contains information that was not included in that document at the time of publication.

3.1 Upgrading BMC Firmware On HP ProLiant DL140 G2 and DL145 G2 Nodes

This note applies only if the hardware configuration contains HP ProLiant DL140 G2 or DL145 G2 nodes and you are upgrading an existing HP XC system from Version 2.1 or Version 3.0 to Version 3.2.

The HP ProLiant DL140 G2 (G2) and DL145 G2 series of hardware models must be installed with BMC firmware version 1.25 or greater. However, the BMC version 1.25 firmware was not supported by HP XC Version 3.0 or earlier. As a result, you must update the BMC firmware on these nodes **after** you upgrade the system to HP XC Version 3.2, which is contrary to the upgrade instructions for a typical upgrade.

Before upgrading an HP XC system to Version 3.2, contact the HP XC Support Team and request the procedure to upgrade the BMC firmware on HP ProLiant DL140 G2 and DL145 G2 nodes:

xc_support@hp.com

4 Software Installation On The Head Node

This chapter contains notes that apply to the HP XC System Software Kickstart installation session.

4.1 Manual Installation Required For NC510F Driver

The `unm_nic` driver is provided with the HP XC software distribution, however, it does not load correctly.

If your system has a NC510F 10 GB Ethernet card, run the following commands to load the driver:

```
# depmod -a
# modprobe -v unm_nic
```

Then, edit the `/etc/modprobe.conf` file and specify `unm` as the driver for the `eth` device assigned to the NC510F driver.

5 System Discovery, Configuration, and Imaging

This chapter contains information about configuring the system. Notes that describe additional configuration tasks are mandatory and have been organized chronologically. Perform these tasks in the sequence presented in this chapter.

The HP XC system configuration procedure is documented in the *HP XC System Software Installation Guide*.



IMPORTANT: Before you begin, depending upon the cluster platform type, see Chapter 10 (page 37), Chapter 11 (page 39), or Chapter 12 (page 41) to determine if additional platform-specific notes apply to the system discovery, configuration, or imaging process.

5.1 Notes That Apply Before You Invoke The `cluster_prep` Utility

Read the notes in this section before you invoke the `cluster_prep` utility.

5.1.1 Required Task for Some NIC Adapter Models: Verify Correct NIC Device Driver Mapping

On head nodes installed with dual-fiber NIC server adapter models NC6170 or NC7170, Ethernet ports might be reordered between the Kickstart kernel and the subsequent HP XC kernel reboot. Use the procedure described in this section to correct the mapping if a re-ordering has occurred.

At the time of the Kickstart installation, the fiber ports are identified as `eth0` and `eth1`, and the onboard ports are identified as `eth2` and `eth3`.

The `/etc/modprobe.conf` file is written as follows:

- `alias eth0 e1000`
- `alias eth1 e1000`
- `alias eth2 tg3`
- `alias eth3 tg3`

You must correct this mapping if you find that upon the HP XC kernel reboot, `eth0` and `eth1` are the `tg3` devices, and `eth2` and `eth3` are the `e1000` devices. To get the external network connection working, perform this procedure from a locally-connected terminal before invoking the `cluster_prep` utility:

1. Unload the `tg3` and `e1000` drivers:

```
# rmmmod e1000
# rmmmod tg3
```

2. Use the text editor of your choice to edit the `/etc/modprobe.conf` file to correct the mapping of drivers to devices. The section of this file should look like this when you are finished:

```
alias eth0 tg3
alias eth1 tg3
alias eth2 e1000
alias eth3 e1000
```

3. Save your changes and exit the text editor.
4. Use the text editor of your choice to edit the `/etc/sysconfig/network-scripts/ifcfg-eth[0,1,2,3]` files, and remove the `HWADDR` line from each file if it is present.
5. If you made changes, save your changes and exit each file.
6. Reload the modules:

```
# modprobe tg3
# modprobe e1000
```

7. Follow the instructions in the *HP XC System Software Installation Guide* to complete the cluster configuration process (beginning with the `cluster_prep` command).

5.2 Notes That Apply To The Discover Process

The notes in this section apply to the `discover` utility.

5.2.1 Discovery of HP ProLiant DL140 G3 and DL145 G3 Nodes Fails When Graphics Cards Are Present

When an HP ProLiant DL140 G3 or DL145 G3 node contains a graphics card, the nodes often fail to PXE boot. Even when the BIOS boot settings are configured to include a PXE boot, these settings are often reset to the factory defaults when the BIOS restarts after saving the changes. This action causes the discovery and imaging processes to fail.

Follow this procedure to work around the discovery failure:

1. Begin the discovery process as usual by issuing the appropriate `discover` command.
2. When the discovery process turns on power to the nodes of the cluster, manually turn off the DL140 G3 and DL145 G3 servers that contain graphics cards.
3. Manually turn on power to each DL140 G3 and DL145 G3 server one at a time, and use the cluster's console to force each node to PXE boot. Do this by pressing the **F12** key at the appropriate time during the BIOS start up.

After you complete this task for each DL140 G3 and DL145 G3 server containing a graphics card, the discovery process continues and completes successfully.

The work around for the imaging failure on these servers is described in “HP ProLiant DL140 G3 and DL145 G3 Node Imaging Fails When Graphics Cards Are Present” (page 27), which is the appropriate place to perform the task.

5.3 Notes That Apply Before You Invoke The `cluster_config` Utility

Read the notes in this section before you invoke the `cluster_config` utility.

5.3.1 Adhere To Role Assignment Guidelines for Improved Availability

When you are configuring services for improved availability, you must adhere to the role assignment guidelines in Table 1-2 in the *HP XC System Software Installation Guide*. Role assignments for a traditional HP XC system without improved availability of services is slightly different, thus it is important that you follow the guidelines in Table 1-2.

5.4 Benign Message From `C52xcgraph` During `cluster_config`

You might see the following message when you run the `cluster_config` utility on a cluster with an InfiniBand interconnect:

```
.
.
.
Executing C52xcgraph gconfigure
Found no adapter info on IR0N00
Failed to find any Infiniband ports
Executing C54httpd gconfigure
.
.
.
```

This message is displayed because the C52xcgraph configuration script is probing the InfiniBand switch to determine how many HCAs with an IP address are present. Because the HCAs have not yet been assigned an IP address, C52xcgraph does not find any HCAs with an IP address and prints the message. This message does not prevent the `cluster_config` utility from completing.

To work around this issue, after the cluster is installed and configured, run `/opt/hptc/hpcgraph/sbin/hpcgraph-setup` with no options.

5.5 Processing Time For `cluster_config` Might Take Longer On A Head Node With Improved Availability

The `cluster_config` utility processing time can take approximately ten minutes longer if it is run on a head node that is configured for improved availability with Serviceguard when the remaining nodes of the cluster are up and running.

After the entire system has been imaged and booted, you might need to re-run the `cluster_config` procedure to modify the node configuration. If the other node in the availability set with the head node is up and running, the Serviceguard daemons attempt to establish Serviceguard related communication with the node when they are restarted. Because the other node in the availability set is not actively participating in a Serviceguard cluster, it will not respond to the Serviceguard communication.

The Serviceguard software on the head node retries this communication until the communication times out. On a system running with the default Serviceguard availability configuration, the timeout is approximately ten minutes.

5.6 Notes That Apply To Imaging

The notes in this section apply to propagating the golden image to all nodes, which is accomplished when you invoke the `startsys` command.

5.6.1 HP ProLiant DL140 G3 and DL145 G3 Node Imaging Fails When Graphics Cards Are Present

As described in “Discovery of HP ProLiant DL140 G3 and DL145 G3 Nodes Fails When Graphics Cards Are Present” (page 26), the discovery and imaging processes might fail on HP ProLiant DL140 G3 or DL145 G3 servers containing graphics cards.

The work around for the discovery failure is described in “Discovery of HP ProLiant DL140 G3 and DL145 G3 Nodes Fails When Graphics Cards Are Present” (page 26), and the work around for the imaging process described in this section assumes that all nodes were discovered.

Follow this procedure to propagate the golden image to DL140 G3 and DL145 G3 servers containing a graphics card:

1. Issue the appropriate `startsys` command and specify one of the DL140 G3 or DL145 G3 nodes with a graphics card in the `[nodelist]` option of the `startsys` command.
2. When power to the node is turned on, use the cluster console to connect to the node and force it to PXE boot by pressing the **F12** key at the appropriate time during the BIOS start up.
3. When the node is successfully imaged, repeat this process for the remaining nodes containing graphics cards.
4. When all nodes containing graphics cards are imaged, issue the `startsys` command without the `[nodelist]` option to image all remaining nodes of the cluster in parallel.

6 Software Upgrades

This chapter contains notes about upgrading the HP XC System Software from a previous release to this release.

Installation release notes described in Chapter 4 (page 23) and system configuration release notes described in Chapter 5 (page 25) also apply when you upgrade the HP XC System Software from a previous release to this release. Therefore, when performing an upgrade, make sure you also read and follow the instructions in those chapters.

6.1 Do Not Upgrade If You Want Or Require The Voltaire InfiniBand Software Stack

HP XC System Software Version 3.2 installs and uses the OFED InfiniBand software stack by default. Previous HP XC releases installed the Voltaire InfiniBand software stack. If you want to continue using the Voltaire InfiniBand software stack, do not upgrade to HP XC System Software Version 3.2.

7 System Administration, Management, and Monitoring

This chapter contains notes about system administration, management, and monitoring.

7.1 Perform A Dry Run Before Using The `si_updateclient` Utility To Update Nodes

The `si_updateclient` utility can leave nodes in an unbootable state in certain situations. You can still use `si_updateclient` to deploy image changes to nodes. However, before you update any nodes, HP recommends that you perform a dry run first to ensure that files in the `/boot` directory are not updated. Updating files in `/boot` can result in nodes being unable to boot.

You can retrieve a list of files that will be updated by `si_updateclient` by specifying `--dry-run` on the command line.

7.2 Possible Problem With ext3 File Systems On SAN Storage

Issues have been reported when an ext3 file system fills up to the point where `ENOSPC` is returned to write requests for a long period of time, and the file system is subsequently unmounted. A forced check is initiated (`fsck -fy`) before the next mount. It appears that the `fsck` checks might corrupt the file system inode information.

This problem has been seen only on fibre channel (SAN) storage; it has not been seen with directly attached storage or NFS storage.

For more information about details and work arounds, consult Bugzilla number 175877 at the following URL:

https://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=175877

8 HP XC System Software On Red Hat Enterprise Linux

The notes in this chapter apply when the HP XC System Software is installed on Red Hat Enterprise Linux.

8.1 Enabling 32-bit Applications To Compile and Run

To compile and run 32-bit applications on a system running HP XC System Software on Red Hat Enterprise Linux 4 on HP Integrity platforms, use the following commands to install the `glibc-2.3.4-2.25.i686.rpm` from the HP XC distribution media DVD:

```
# mount /dev/cdrom
# cd /mnt/cdrom/LNXHPC/RPMS
# rpm -ivh glibc-2.3.4-2.25.i686.rpm
```


9 Programming and User Environment

This chapter contains information that applies to the programming and user environment.

9.1 MPI and OFED InfiniBand Stack Fork Restrictions

With the introduction of the OFED InfiniBand stack in this release, MPI applications cannot call `fork()`, `popen()`, and `system()` between `MPI_Init` and `MPI_Finalize`. This is known to affect some applications like NWChem.

9.2 InfiniBand Multiple Rail Support

HP-MPI provides multiple rail support on OpenFabric through the `MPI_IB_MULTIRAIL` environment variable. This environment variable is ignored by all other interconnects. In multi-rail mode, a rank can use up to all cards on its node, but it is limited to the number of cards on the node to which it is connecting.

For example, if rank A has three cards, rank B has two cards, and rank C has three cards, then connection A--B uses two cards, connection B--C uses two cards, and connection A--C uses three cards. Long messages are striped among all the cards on that connection to improve bandwidth.

By default, multi-card message striping is off. To turn it on, specify `-e MPI_IB_MULTIRAIL=N` where N is the number of cards used by a rank:

- If $N \leq 1$, message striping is not used.
- If N is greater than the maximum number of cards M on that node, all M cards are used.
- If $1 < N \leq M$, message striping is used on N cards or less.

If you specify `-e MPI_IB_MULTIRAIL`, the maximum possible cards are used.

On a host, all the ranks select all the cards in a series. For example: given 4 cards and 4 ranks per host:

- rank 0 will use cards 0, 1, 2, 3
- rank 1 will use cards 1, 2, 3, 0
- rank 2 will use cards 2, 3, 0, 1
- rank 4 will use cards 3, 0, 1, 2

The order is important in SRQ mode because only the first card is used for short messages. The selection approach allows short RDMA messages to use all the cards in a balanced way.

For HP-MPI 2.2.5.1 and older, all cards must be on the same fabric.

9.3 Benign Messages From HP-MPI Version 2.2.5.1

When running jobs with XC Version 3.2, OFED InfiniBand, and HP-MPI Version 2.2.5.1 the following messages are printed once for each rank:

```
libibverbs: Warning: fork()-safety requested but init failed
```

HP-MPI Version 2.2.5.1 has support for `fork()` using OFED 1.2, but only for kernels more recent than version 2.6.12. HP XC Version 3.2 is currently based on kernel version 2.6.9. This message is a reminder that `fork()` is not supported in this release.

You can suppress this message by defining the `MPI_IBV_NO_FORK_SAFE` environment variable, as follows:

```
% /opt/hpmpi/bin/mpirun -np 4 -prot -e MPI_IBV_NO_FORK_SAFE=1 -hostlist nodea,nodeb,nodec,noded /my/dir/hello_world
```

10 Cluster Platform 3000

At the time of publication, no release notes are specific to Cluster Platform 3000 systems.

11 Cluster Platform 4000

At the time of publication, no release notes are specific to Cluster Platform 4000 systems.

12 Cluster Platform 6000

This chapter contains information that applies only to Cluster Platform 6000 systems.

12.1 Network Boot Operation and Imaging Failures on HP Integrity rx2600 Systems

An underlying issue in the kernel is causing MAC addresses on HP Integrity rx2600 systems to be set to all zeros (for example, 00.00.00.00.00), which results in network boot and imaging failures.

To work around this issue, enter the following commands on the head node to network boot and image an rx2600 system:

1. Prepare the node to network boot:
`# setnode --resync node_name`
2. Turn off power to the node:
`# stopsys --hard node_name`
3. Start the imaging and boot process:
`# startsys --image_and_boot node_name`

12.2 Notes That Apply To The Management Processor

This section describes limitations with the management processor (MP) that are expected to be resolved when a new firmware version is available.

12.2.1 Required Task: Change MP Settings on Console Switches

Perform this task before invoking the `discover` command.

In order for the discovery process to work correctly using the MP in DHCP mode, you must increase the amount of time the console switches hold MAC addresses. Increase this value from the default of 300 seconds to 1200 seconds. Make this change only on the console switches in the system, typically the ProCurve 26xx series.

From the ProCurve prompt, enter the configuration mode and set the `mac-age-time` parameter, as follows:

```
# config
(config)# mac-age-time 1200
```

12.2.2 MP Disables DHCP Automatically

A known limitation exists with the MP firmware that causes the MP to disable DHCP automatically.

To work around this issue, the HP XC software performs the discovery phase with DHCP enabled. You must then perform a procedure to change the addresses on all MPs in the system to use the address received from DHCP as a static address.

For more information on how to perform this procedure, contact the HP XC Support Team at xc_support@hp.com.

12.2.3 Finding the IP Address of an MP

Because the IP addresses for the MPs are being set statically for this release, if a node must be replaced, you must set the IP address for the MP manually when the node is replaced.

To find the IP address, look up the entry for the MP in the `/etc/dhcpd.conf` file. The MP naming convention for the node is `cp-node_name`.

13 Integrated Lights Out Console Management Devices

This chapter contains information that applies to the integrated lights out (iLO and iLO2) console management device.

13.1 iLO2 Devices In Server Blades Can Hang

There is a known problem with the iLO2 console management devices that causes the iLO2 devices to hang. This particular problem has very specific characteristics:

- This problem is typically seen within one or two days of the initial cluster installation.
- Most of the time, but not always, all iLO2 devices in a particular enclosure hang at the same time.
- The problem usually affects multiple enclosures.

The work around for this problem is to completely power cycle the entire cluster (or at least all enclosures) after the initial cluster installation is complete or if the problem is encountered. This problem has never been reported after the power has been cycled and the cluster is in its normal running state.

This problem is targeted for resolution in iLO2 firmware Version 1.28, but at the time of publication, had not been tested yet.

14 Interconnects

This chapter contains information that applies to the supported interconnect types:

- InfiniBand Interconnect (page 45)
- Myrinet Interconnect (page 45)
- QsNet^{II} Interconnect (page 45)

14.1 InfiniBand Interconnect

The notes in this section apply to the InfiniBand interconnect.

14.1.1 enable Password Problem With Voltaire Switch Version 4.1

The instructions for configuring Voltaire InfiniBand switch controller cards requires you to change the factory default passwords for the `admin` and `enable` accounts, as follows:

```
Insert new (up to 8 characters) Enter password :
```

An issue exists where you must enter a password with exactly eight characters for the `enable` account. The `admin` account is not affected.

If the new password does not contain exactly eight characters, the following message appears when you try to log in with the new password:

```
Unauthorized mode for this user, wrong password or illegal mode in the first word.
```

This problem has been reported to Voltaire. As a work around, choose a password that is exactly eight characters.

14.2 Myrinet Interconnect

The following release notes are specific to the Myrinet interconnect.

14.2.1 Myrinet Monitoring Line Card Can Become Unresponsive

A Myrinet monitoring line card can become unresponsive some period of time after it has been set up with an IP address with DHCP. This is a problem known to Myricom. For more information, see the following:

<http://www.myri.com/fom-serve/cache/321.html>

If the line card becomes unresponsive, re-seat the line card by sliding it out of its chassis slot and then slide it back in. You can do this while the system is up; doing so does not interfere with Myrinet traffic.

14.2.2 The `clear_counters` Command Does Not Work On The 256 Port Switch

The `/opt/gm/sbin/clear_counters` command does not clear the counters on the Myrinet 256 port switch. The web interface to the Myrinet 256 port switch has changed from the earlier, smaller switches.

To clear the switch counters, you must open an interactive Web connection to the switch and clear the counters using the menu commands. The `gm_prodmode_mon` script, which uses the `clear_counters` command, will not clear the counters periodically, as it does on the smaller switches.

This problem will be resolved in a future software update from Myricom.

14.3 QsNet^{II} Interconnect

The following release notes are specific to the QsNet^{II}® interconnect.

14.3.1 Possible Conflict With Use of SIGUSR2

The Quadrics QsNet^{II} software internally uses SIGUSR2 to manage the interconnect. This can conflict with any user applications that use SIGUSR2, including for debugger use.

To work around this conflict, set the environment variable LIBELAN4_TRAPSIG for the application to a different signal number other than the default value 12 that corresponds to SIGUSR2. Doing this instructs the Quadrics software to use the new signal number, and SIGUSR2 can be once again used by the application. Signal numbers are define in the `/usr/include/asm/signal.h` file.

14.3.2 The qsnet Database Might Contain Entries To Nonexistent Switch Modules

Depending on the system topology, the qsnet diagnostics database might contain entries to nonexistent switches.

This issue is manifested as errors reported by the `/usr/bin/qsctrl` utility similar to the following:

```
# qsctrl
qsctrl: failed to initialise module QR0N03: no such module (-7)
.
.
.
```

In the previous example, the `switch_modules` table in the qsnet database is populated with QR0N03 even though the QR0N03 module is not physically present. This problem has been reported to Quadrics, Ltd.

To work around this problem, delete the QR0N03 entry (and any other nonexistent switch entries) from the `switch_modules` table, and restart the `swmlogger` service:

```
# mysql -u root -p qsnet
mysql> delete from switch_modules where name="QR0N03";
mysql> quit
# service swm restart
```

In addition to the previous problem, the IP address of a switch module might be incorrectly populated in the `switch_modules` table, and you might see the following message:

```
# qsctrl
qsctrl: failed to parse module name 172.20.66.2
.
.
.
```

Resolve this issue by deleting the IP address from the `switch_modules` table and restarting the `swmlogger` service:

```
# mysql -u root -p qsnet
mysql> delete from switch_modules where name="172.20.66.2";
mysql> quit
# service swm restart
```



NOTE: You must repeat the previous procedure if you invoke the `cluster_config` utility again and you choose to re-create the qsnet database during the `cluster_config` operation.

15 Documentation

This chapter describes known issues with the HP XC documentation.

15.1 Documentation CD Search Option

If you are viewing the main page of the HP XC Documentation CD, you cannot perform a literature search from the **Search:** option box at the top of the page.

To search <http://www.docs.hp.com> or to search all of HP's global Web service, click on the link for **More options**. The Advanced search options page is displayed, and you can perform the search from the advanced page.

15.2 HP XC Manpages

The notes in this section apply to the HP XC manpages.

15.2.1 New `device_config.8`

A manpage is available for the `device_config` command. The `device_config` command enables you to modify the device configuration information in the HP XC command and management database (CMDB). Uses for this command include configuring a range of default external network interface cards (NICs) across multiple nodes and configuring one or two additional, external NICs on the same node.

15.2.2 Changes to `ovp.8`

Note the following two changes to the `ovp(8)` manpage:

1. Under `-o options`, `--opts_for_test [=] options`, add the following before
`--user=username:`
`--queue LSF_queue`
Specifies the LSF queue for the performance health tests.
2. Change the following portion of the `-v component`, `--verify [=] component` as follows:

OLD:

For all users:

This option takes the form `--verify=perf_health/test`

`cpu` Tests CPU core performance using the Linpack benchmark

NEW:

For all users:

This option takes the form `--verify=perf_health/test`

NOTE: Except for the `network_stress` and `network_bidirectional` tests, these tests only apply to systems that install LSF-HPC incorporated with SLURM. The `network_stress` and `network_bidirectional` tests also function under Standard LSF.

`cpu` Tests CPU core performance using the Linpack benchmark.

15.2.3 New `preupgradesys-lxc.8`

The `preupgradesys-lxc(8)` manpage was not included in the HP XC Version 3.2 distribution.

preupgradesys-lxc(8)

NAME

preupgradesys-lxc - Prepares a system for an XC software upgrade

SYNOPSIS

Path: /opt/hptc/lxc-upgrade/sbin/preupgradesys-lxc

DESCRIPTION

Running the `preupgradesys-lxc` command is one of several commands that are part of the process to upgrade HP XC System Software on Red Hat Enterprise Linux to the next release of HP XC System Software on Red Hat Enterprise Linux. The software upgrade process is documented in the HP XC System Software Installation Guide. This command is never run for any reason other than during a software upgrade.

The `preupgradesys-lxc` command prepares your system for a XC software upgrade by modifying release-specific files, recreating links where required, and making backup copies of important files. It also removes specific XC RPMs that do not upgrade properly. Running `preupgradesys-lxc` is a required task before beginning a software upgrade.

The `preupgradesys-lxc` command does not prepare your system for upgrading Red Hat Enterprise Linux RPMs.

OPTIONS

The `preupgradesys-lxc` command does not have any options.

FILES

/var/log/preupgradesys-lxc/preupgradesys-lxc.log
Contains command output and results

SEE ALSO

`upgradesys-lxc(8)`
HP XC System Software Installation Guide

15.2.4 New `upgradesys-lxc.8`

The `upgradesys-lxc(8)` manpage was not included in the HP XC Version 3.2 distribution.

NAME

`upgradesys-lxc` - For XC software upgrades, this command upgrades and migrates configuration data to the new release format

SYNOPSIS

Path: /opt/hptc/lxc-upgrade/sbin/upgradesys-lxc

DESCRIPTION

Running the `upgradesys-lxc` command is one of several commands that are part of the process to upgrade HP XC System Software on Red Hat Enterprise Linux to the next release of HP XC System Software on Red Hat Enterprise Linux. The software upgrade process is documented in the HP XC System Software Installation Guide. This command is never run for any reason other than during a software upgrade.

The `upgradesys-lxc` utility is run immediately after the head node is upgraded with the new XC release software and any other required third-party software products. The `upgradesys-lxc` utility performs the following tasks to upgrade your system:

- o Makes a backup copy of the database from the previous release.
- o Modifies attributes in the database to signify that the system has been upgraded.
- o Removes RPMs from the previous release that are no longer supported in the new release.
- o Executes internal migration scripts to migrate system configuration data to the new release format.

OPTIONS

The `upgradesys-lxc` command does not have any options.

FILES

`/opt/hptc/lxc-upgrade/etc/gupdate.d`
Location of migration scripts

`/opt/hptc/etc/sysconfig/upgrade/upgradesys.dbbackup-date_time_stamp`
Location of database backup

`/var/log/upgradesys-lxc/upgradesys-lxc.log`
Contains the results of the RPM upgrade process and lists customized configuration files

SEE ALSO

`preupgradesys-lxc(8)`
HP XC System Software Installation Guide

Index

B

base operating system, 15

C

C52xcgraph error, 26
clear_counters command, 45
client node disk partition, 16
cluster_config utility, 26
 C52xcgraph error message, 26
 new features, 16
CP3000 system, 37
CP4000 system, 39
CP6000 system, 41
 SIGUSR2 signal, 46

D

data corruption on ext3 file systems, 31
discover command
 new features, 16
discover utility, 26
documentation, 47
 additional publications, 12
 changed in this release, 18
 compilers, 12
 FlexLM, 11
 HowTo, 8
 HP XC System Software, 8
 Linux, 11
 LSF, 10
 manpages, 12
 master firmware list, 8
 Modules, 11
 MPI, 11
 MySQL, 11
 Nagios, 10
 pdsh, 10
 reporting errors in, 13
 rrdtool, 10
 SLURM, 10
 software RAID, 12
 Supermon, 10
 syslog-ng, 10
 SystemImager, 10
 TotalView, 12

E

ext3 file system, 31

F

failed to find InfiniBand ports, 26
feedback
 e-mail address for documentation, 13
firmware version, 19
found no adapter info on IR0N00, 26

H

hardware preparation tasks, 21
hardware support, 15
HowTo, 18
 Web site, 8
HP documentation
 providing feedback for, 13
HP Scalable Visualization Array (*see* SVA)
HP-MPI
 fork restrictions with kernel version, 35
 fork restrictions with OFED, 35
 init failed, 35
 multiple rail support, 35

I

iLO, 43
iLO2
 hang, 43
InfiniBand
 multiple rail support, 35
InfiniBand interconnect
 failed to find ports, 26
inode information, 31
installation notes, 23
integrated lights out console management device (*see* iLO) (*see* iLO2)
interconnect, 45

K

kernel version, 15
Kickstart installation, 23

L

Linux operating system, 15
LSF
 documentation, 10

M

management processor (*see* MP)
manpages, 12
mdadm utility, 12
MP, 41
MPI (*see* HP-MPI)
multiple rail support, 35
Myrinet interconnect, 45

N

NC6170 NIC adapter, 25
NC7170 NIC adapter, 25
new features, 15
NIC device driver mapping, 25

O

OFED, 15
 fork restrictions with HP-MPI, 35

OVP
enhancements, 17

P
partition size limit, 16
patches, 19

Q
qsnet diagnostics database, 46
QsNet interconnect, 45

R
reporting documentation errors
feedback e-mail address for, 13
resmon utility, 17

S
si_updateclient utility, 31
signal
Quadrics QsNet, 46
software RAID
documentation, 12
mdadm utility, 12
SVA, 15
system administration
notes, 31
system configuration, 25
system management
enhancements, 17
notes, 31
system monitoring, 17

T
temperature graph, 17

U
unified parallel C, 17
UPC, 17
upgrade, 29
upgrade installation, 29

W
Web site
HP XC System Software documentation, 8

