



**Hewlett Packard**  
Enterprise

# HPE Cluster Test Installation Guide

## **Abstract**

This guide describes the procedures for installing HPE Cluster Test RPM. For information on HPE Cluster Test Image and CTlite, see their respective documents.

Part Number: 5900-3747R  
Published: November 2015  
Edition: 6

© Copyright 2011, 2015 Hewlett Packard Enterprise Development LP

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise 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. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise 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.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

#### **Acknowledgments**

Microsoft® and Windows® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

#### **Revision history**

<b>Manufacturing part number</b>	<b>Edition number</b>	<b>Publication date</b>
5900-3747	5	<i>January 2014</i>
5900-3311	4	<i>November 2013</i>
5900-2458	3	<i>September 2012</i>
5900-1968	2	<i>October 2011</i>
5070-6697	1	<i>May 2011</i>

# Contents

1 Cluster Test overview.....	4
2 Requirements and installation of CT RPM.....	5
Requirements for Cluster Test RPM.....	5
Additional requirements for accelerator support.....	15
Getting the CT RPM kit.....	15
Procedure for installing CT RPM.....	16
3 Uninstalling Cluster Test.....	20
Uninstalling CT RPM.....	20
4 Support and other resources.....	21
Accessing Hewlett Packard Enterprise Support.....	21
Accessing updates.....	21
Intended audience.....	22
Websites.....	22
Customer self repair.....	22
Remote support.....	22
Documentation feedback.....	23
Index.....	24

# 1 Cluster Test overview

Cluster Test is designed to:

- verify the validity of a cluster configuration.
- test the functionality of the cluster as a whole, with emphasis on testing interconnect, including the interconnect switch, cabling, and interface cards.
- provide stress testing on the cluster nodes.

Cluster Test detects failures of hardware and network connections to the node level early, so the failed components can be replaced or corrected before cluster software (HPE Insight Cluster Management Utility (CMU), XC, Scali, Scyld, etc.) installation is started. This makes cluster software integration much easier and faster, particularly on a complex solution with several hundred nodes.

Cluster Test is available in the following versions:

- Image
- RPM
- CTlite

---

**NOTE:** This installation guide applies only to the Cluster Test RPM version. For information on the Cluster Test Image and CTlite versions, see their respective documents.

---

The Cluster Test Image and RPM versions have a common GUI and suite of tests.

CT Image allows you to test and diagnose the hardware on a cluster that does not have an operating system installed. CT Image contains the Red Hat Enterprise Linux (RHEL) operating system, libraries, software packages, scripts, and GUIs that allow you to set up an HPC cluster in a very little time and run tests immediately afterward. The cluster's compute nodes are set up in a diskless cluster environment, which allows the tests to be performed on the compute nodes with pre-installed OS and software. CT Image has two distribution media options:

- DVD ISO image
- USB key

Cluster Test Image can also be set up to do a network installation instead of using physical media. All versions of Cluster Test Image install nothing on compute nodes; compute nodes are not altered any way. Cluster Test Image does require installation on the head node's hard drive.

Cluster Test RPM provides the ability to diagnose and test the hardware and software on a cluster that is already installed and running. CT RPM contains the software tools, scripts, and GUIs that allow you to set up an HPC cluster in very little time and run tests immediately afterward. Cluster Test RPM requires certain software packages on the cluster, such as the operating system, drivers, etc., but allows the flexibility of testing combinations of these specific components.

Supported servers and component hardware are listed in the Cluster Test Release Notes document, available at <http://www.hpe.com/info/ct-docs>.

## 2 Requirements and installation of CT RPM

Cluster Test RPM is installed on an existing cluster as an additional software package. The current cluster settings are not modified.

### Requirements for Cluster Test RPM

- Currently, Cluster Test RPM supports only X86–64 systems.
- Cluster Test RPM does not include an operating system, drivers, or message passing interface (MPI). These software packages must be installed and configured on each node of the cluster prior to installing the Cluster Test RPM. If the cluster uses an interconnect, the interconnect interface must be enabled and present.
- Currently, Cluster Test RPM only supports Platform MPI and Open MPI. Each node being tested must have the appropriate MPI version installed and configured.
- For Open MPI:
  - The `PATH` environment variable must include an entry pointing to the `mpirun` and `mpicc` commands.
  - The `LD_LIBRARY_PATH` must contain the path to the Open MPI libraries as its first entry. The `mpi-selector` command can be used to set the environment appropriately.
- The `ssh` network protocol must be available and configured to allow `ssh` between all nodes without being challenged for a password.
- A common node-naming convention is required for cluster communication. The node-naming scheme should encompass a common base name followed by a unique number. If the cluster has an interconnect, its network alias should follow the same naming convention, with a common prefix added to the base name. As an example, a common base name of “node” could be used for the administrative network of the cluster and adding a prefix of “i” would identify the interconnect network as “inode”.
- Cluster Test RPM requires the installation of some Perl modules to run the user interface. These Perl modules are included with RPM, but the individual modules have the following specific package requirements (additional packages might be required when building Cluster Test executables). These packages should be installed prior to installing the Cluster Test RPM and are generally available from the Linux installation media.
  - `fontconfig-devel`
  - `gd`
  - `libpng-devel`
  - `libXp`
  - `openmotif22`
  - `perftest`
  - `tcl-devel` (required for rebuild only)
  - `tk-devel` (required for rebuild only)
  - `gd-devel` (required for rebuild only)

- libaio-devel (required for rebuild only)
- libjpeg-devel (required for rebuild only)
- You should have access to the appropriate RPM package. Information on how to download RPMs is available in the section [“Getting the CT RPM kit” \(page 15\)](#).
- Prerequisite OS RPMs:
  - xinetd
  - openmpi (test binaries are built against openmpi-1.6.4)
  - hphealth (Optional. Required to run Health Check.)
  - hponcfg (Optional. Required to run Hardware and Firmware Summaries.)
- You should install the following RPMs from your OFED distribution if they are not already installed:
  - infiniband-diags
  - libibmad1
- **RHEL 6.1 CT RPM Kit Dependent packages**
  - tk:
    - tk-8.5.7-5.el6.x86\_64.rpm
  - OpenIPMI, ipmitool, net-snmp-libs, lm\_sensors-libs (needed to start ipmi service):
    - OpenIPMI-2.0.16-12.el6.x86\_64.rpm
    - OpenIPMI-libs-2.0.16-12.el6.x86\_64.rpm
    - ipmitool-1.8.11-7.el6.x86\_64.rpm
    - net-snmp-libs-5.5-31.el6.x86\_64.rpm
    - lm\_sensors-libs-3.1.1-10.el6.x86\_64.rpm
  - libxml, libstdc++, zlib (32-Bit [686]RPMs needed for conrep):
    - libstdc++-4.4.5-6.el6.i686.rpm
    - libxml2-2.7.6-1.el6.i686.rpm
    - zlib-1.2.3-25.el6.i686.rpm
  - Mesa-libGL-devel required to install nbody after cuda tool kit installation:
    - mesa-libGL-devel-7.10-1.el6.x86\_64.rpm
  - mcelog package to decode kernel machine check log on x86 machines used by CT memory error check:
    - mcelog-1.0pre3\_20110718-0.14.el6.x86\_64.rpm
  - Install the following two font-related RPM packages from the respective DVD ISO distribution.
    - urw-fonts-2.4-10.el6.noarch.rpm
    - ghostscript-fonts-5.50-23.1.el6.noarch.rpm

- **RHEL 6.2 CT RPM Kit Dependent packages**

- tk:
  - tk-8.5.7-5.el6.x86\_64.rpm
- OpenIPMI, ipmitool, net-snmp-libs, lm\_sensors-libs (needed to start ipmi service):
  - OpenIPMI-2.0.16-12.el6.x86\_64.rpm
  - OpenIPMI-libs-2.0.16-12.el6.x86\_64.rpm
  - ipmitool-1.8.11-12.el6.x86\_64
  - net-snmp-libs-5.5-37.el6.x86\_64.rpm
  - lm\_sensors-libs-3.1.1-10.el6.x86\_64.rpm
- libxml, libstdc++, zlib (32-Bit [i686]RPMs needed for conrep):
  - libstdc++-4.4.6-4.el6.i686.rpm
  - libxml2-2.7.6-4.el6.i686.rpm
  - zlib-1.2.3-27.el6.i686.rpm
- Mesa-libGL-devel required to install nbody after cuda tool kit installation:
  - mesa-libGL-devel-7.11-3.el6.x86\_64.rpm
- mcelog package to decode kernel machine check log on x86 machines used by CT memory error check:
  - mcelog-1.0pre3\_20110718-0.14.el6.x86\_64.rpm
- Install the following two font-related RPM packages from the respective DVD ISO distribution.
  - urw-fonts-2.4-10.el6.noarch.rpm
  - ghostscript-fonts-5.50-23.1.el6.noarch.rpm

- **RHEL 6.3 CT RPM Kit Dependent packages**

- tk:
  - tk-8.5.7-5.el6.x86\_64.rpm
- OpenIPMI, ipmitool, net-snmp-libs, lm\_sensors-libs (needed to start ipmi service):
  - OpenIPMI-2.0.16-12.el6.x86\_64.rpm
  - OpenIPMI-libs-2.0.16-12.el6.x86\_64.rpm
  - ipmitool-1.8.11-13.el6.x86\_64.rpm
  - net-snmp-libs-5.5-41.el6.x86\_64.rpm
  - lm\_sensors-libs-3.1.1-10.el6.x86\_64.rpm
- libxml, libstdc++, zlib (32-Bit [i686]RPMs needed for conrep):
  - libstdc++-4.4.5-6.el6.i686.rpm
  - libxml2-2.7.6-4.el6\_2.4.i686.rpm
  - zlib-1.2.3-27.el6.i686.rpm

- Mesa-libGL-devel required to install nbody after cuda tool kit installation:
  - mesa-libGL-devel-7.11-5.el6.x86\_64.rpm
- mcelog package to decode kernel machine check log on x86 machines used by CT memory error check:
  - mcelog-1.0pre3\_20110718-0.14.el6.x86\_64.rpm
- Install the following two font-related RPM packages from the respective DVD ISO distribution.
  - urw-fonts-2.4-10.el6.noarch.rpm
  - ghostscript-fonts-5.50-23.1.el6.noarch.rpm
- **RHEL 6.4 CT RPM Kit Dependent packages**
  - tk:
    - tk-8.5.7-5.el6.x86\_64.rpm
  - OpenIPMI, ipmitool, net-snmp-libs, lm\_sensors-libs (needed to start ipmi service):
    - OpenIPMI-2.0.16-14.el6.x86\_64.rpm
    - OpenIPMI-libs-2.0.16-14.el6.x86\_64.rpm
    - ipmitool-1.8.11-13.el6.x86\_64.rpm
    - net-snmp-libs-5.5-44.el6.x86\_64.rpm
    - lm\_sensors-libs-3.1.1-17.el6.x86\_64.rpm
  - libxml, libstdc++, zlib (32-bit [i686] RPMs needed for conrep):
    - libstdc++-4.4.7-3.el6.i686.rpm
    - libstdc++-devel-4.4.7-3.el6.x86\_64.rpm
    - libxml2-2.7.6-8.el6\_3.4.i686.rpm
    - zlib-1.2.3-29.el6.i686.rpm
  - Mesa-libGL-devel required to install nbody after cuda tool kit installation:
    - mesa-libGL-9.0-0.7.el6.x86\_64.rpm
    - mesa-libGL-devel-9.0-0.7.el6.x86\_64.rpm
  - mcelog package to decode kernel machine check log on x86 machines used by CT memory error check:
    - mcelog-1.0pre3\_20120814\_2-0.6.el6.x86\_64.rpm
  - Install the following two font-related RPM packages from the respective DVD ISO distribution.
    - urw-fonts-2.4-10.el6.noarch.rpm
    - ghostscript-fonts-5.50-23.1.el6.noarch.rpm
  - Other dependent packages that might be required (if not already installed):
    - expect-5.44.1.15-4.el6.x86\_64.rpm
    - oddjob-0.30-5.el6.x86\_64.rpm
    - python-imaging-1.1.6-19.el6.x86\_64.rpm



- telnet-0.17-47.el6\_3.1.x86\_64
- nfs-utils-1.2.3-36.el6.x86\_64.rpm
- nfs-utils-lib-1.1.5-6.el6.x86\_64.rpm
- make-3.81-20.el6.x86\_64.rpm
- openssl-devel-1.0.0-27.el6.x86\_64rpm
- gcc-c++-4.4.7-3.el6.x86\_64.rpm
- gcc-gfortran-4.4.7-3.el6.x86\_64.rpm
- gdb-7.2-60.el6.x86\_64.rpm
- tigervnc-server-1.1.0-5.el6.x86\_64.rpm
- rpm-build-4.8.0-32.el6.x86\_64.rpm
- gcc-4.4.7-3.el6.x86\_64.rpm
- libgcc-4.4.7-3.el6.x86\_64.rpm
- libgcc-4.4.7-3.el6.i686.rpm
- glibc-2.12-1.107.el6.i686.rpm
- libXp-1.0.0-15.1.el6.x86\_64.rpm
- libXpm-devel-3.5.10-2.el6.x86\_64.rpm
- libXpm-3.5.10-2.el6.x86\_64.rpm
- openmotif22-2.2.3-19.el6.x86\_64.rpm
- xorg-x11-fonts-misc-7.2-9.1.el6.noarch.rpm

- **SLES 11.1 CT RPM Kit Dependent packages**

- Install first **before** ofed and mpi installation:
  - gcc-4.3-62.198.x86\_64.rpm
  - gcc-c++-4.3-62.198.x86\_64.rpm
  - gcc-fortran-4.3-62.198.x86\_64.rpm
  - gcc43-4.3.4\_20091019-0.7.35.x86\_64.rpm
  - gcc43-c++-4.3.4\_20091019-0.7.35.x86\_64.rpm
  - gcc43-fortran-4.3.4\_20091019-0.7.35.x86\_64.rpm
  - glibc-devel-2.11.1-0.17.4.x86\_64.rpm
  - glibc-devel-32bit-2.11.1-0.17.4.x86\_64.rpm
  - libgfortran43-4.3.3\_20081022-11.18.x86\_64.rpm
  - libstdc++43-devel-4.3.4\_20091019-0.7.35.x86\_64.rpm
  - linux-kernel-headers-2.6.32-1.4.13.noarch.rpm
- Install **after** ofed and mpi installation:
  - ImageMagick-6.4.3.6-7.20.1.x86\_64.rpm
  - ImageMagick-devel-6.4.3.6-7.20.1.x86\_64.rpm
  - boost-devel-1.36.0-11.17.x86\_64.rpm
  - fontconfig-devel-2.6.0-10.6.x86\_64.rpm

- freetype2-devel-2.3.7-25.8.x86\_64.rpm
- gd-2.0.36.RC1-52.18.x86\_64.rpm
- gd-devel-2.0.36.RC1-52.18.x86\_64.rpm
- glib-1.2.10-737.22.x86\_64.rpm
- glib-devel-1.2.10-737.22.x86\_64.rpm
- glibc-info-2.9-13.2.x86\_64.rpm
- libMagickWand1-6.4.3.6-7.18.x86\_64.rpm
- libboost\_date\_time1\_36\_0-1.36.0-11.17.x86\_64.rpm
- libboost\_filesystem1\_36\_0-1.36.0-11.17.x86\_64.rpm
- libboost\_graph1\_36\_0-1.36.0-11.17.x86\_64.rpm
- libboost\_iostreams1\_36\_0-1.36.0-11.17.x86\_64.rpm
- libboost\_math1\_36\_0-1.36.0-11.17.x86\_64.rpm
- libboost\_mpi1\_36\_0-1.36.0-11.17.x86\_64.rpm
- libboost\_program\_options1\_36\_0-1.36.0-11.17.x86\_64.rpm
- libboost\_python1\_36\_0-1.36.0-11.17.x86\_64.rpm
- libboost\_serialization1\_36\_0-1.36.0-11.17.x86\_64.rpm
- libboost\_signals1\_36\_0-1.36.0-11.17.x86\_64.rpm
- libboost\_system1\_36\_0-1.36.0-11.17.x86\_64.rpm
- libboost\_test1\_36\_0-1.36.0-11.17.x86\_64.rpm
- libboost\_thread1\_36\_0-1.36.0-11.17.x86\_64.rpm
- libboost\_wave1\_36\_0-1.36.0-11.17.x86\_64.rpm
- libjpeg-devel-6.2.0-879.10.x86\_64.rpm
- liblcms-devel-1.17-77.14.19.x86\_64.rpm
- libpciaccess0-devel-7.4-8.24.2.x86\_64.rpm
- libpixman-1-0-devel-0.16.0-1.2.22.x86\_64.rpm
- libpng-devel-1.2.31-5.10.x86\_64.rpm
- libstdc++-devel-4.3-62.198.x86\_64.rpm
- libtiff-devel-3.8.2-141.6.x86\_64.rpm
- libuuid-devel-2.16-6.8.2.x86\_64.rpm
- libwmf-0.2.8.4-206.26.x86\_64.rpm
- libwmf-devel-0.2.8.4-206.26.x86\_64.rpm
- libwmf-gnome-0.2.8.4-206.26.x86\_64.rpm
- libxml2-devel-2.7.6-0.1.37.x86\_64.rpm
- libxml2-devel-32bit-2.7.6-0.1.37.x86\_64.rpm
- ncurses-devel-5.6-90.55.x86\_64.rpm
- openmotif-libs-2.3.1-3.13.x86\_64.rpm
- readline-devel-5.2-147.9.13.x86\_64.rpm

- tack-5.6-90.55.x86\_64.rpm
  - tcl-devel-8.5.5-2.81.x86\_64.rpm
  - tk-devel-8.5.5-3.12.x86\_64.rpm
  - xorg-x11-devel-7.4-8.24.2.x86\_64.rpm
  - xorg-x11-fonts-devel-7.4-1.15.x86\_64.rpm
  - xorg-x11-libICE-devel-7.4-1.15.x86\_64.rpm
  - xorg-x11-libSM-devel-7.4-1.18.x86\_64.rpm
  - xorg-x11-libX11-devel-7.4-5.5.x86\_64.rpm
  - xorg-x11-libXau-devel-7.4-1.15.x86\_64.rpm
  - xorg-x11-libXdmcp-32bit-7.4-1.15.x86\_64.rpm
  - xorg-x11-libXdmcp-devel-32bit-7.4-1.15.x86\_64.rpm
  - xorg-x11-libXdmcp-devel-7.4-1.15.x86\_64.rpm
  - xorg-x11-libXext-devel-32bit-7.4-1.14.x86\_64.rpm
  - xorg-x11-libXext-devel-7.4-1.14.x86\_64.rpm
  - xorg-x11-libXfixes-devel-7.4-1.14.x86\_64.rpm
  - xorg-x11-libXmu-devel-7.4-1.17.x86\_64.rpm
  - xorg-x11-libXp-devel-32bit-7.4-1.14.x86\_64.rpm
  - xorg-x11-libXp-devel-7.4-1.14.x86\_64.rpm
  - xorg-x11-libXpm-devel-32bit-7.4-1.17.x86\_64.rpm
  - xorg-x11-libXpm-devel-7.4-1.17.x86\_64.rpm
  - xorg-x11-libXprintUtil-devel-7.4-1.17.x86\_64.rpm
  - xorg-x11-libXrender-devel-7.4-1.14.x86\_64.rpm
  - xorg-x11-libXt-devel-32bit-7.4-1.17.x86\_64.rpm
  - xorg-x11-libXt-devel-7.4-1.17.x86\_64.rpm
  - xorg-x11-libXv-devel-7.4-1.14.x86\_64.rpm
  - xorg-x11-libfontenc-devel-7.4-1.15.x86\_64.rpm
  - xorg-x11-libxcb-devel-32bit-7.4-1.15.x86\_64.rpm
  - xorg-x11-libxcb-devel-7.4-1.15.x86\_64.rpm
  - xorg-x11-libxkbfile-devel-7.4-1.14.x86\_64.rpm
  - xorg-x11-proto-devel-7.4-1.21.x86\_64.rpm
  - xorg-x11-util-devel-7.4-1.15.x86\_64.rpm
  - xorg-x11-xtrans-devel-7.4-4.22.x86\_64.rpm
  - zlib-devel-1.2.3-106.34.x86\_64.rpm
- OpenIPMI, ipmitool (needed to start ipmi service):
    - ipmitool-1.8.11-0.1.48
    - OpenIPMI-2.0.16-0.3.29

- Mesa-libGL-devel required to install nbody after cuda tool kit installation:
  - Mesa-libGL-devel-8.0.4-20.11.1.x86\_64.rpm
- mcelog package to decode kernel machine check log on x86 machines used by CT memory error check:
  - mcelog-1.0.2011.06.08-0.9.12
- **SLES 11.2 CT RPM Kit Dependent packages**
  - Install first **before** ofed and mpi installation:
    - gcc-4.3-62.198.x86\_64.rpm
    - gcc-c++-4.3-62.198.x86\_64.rpm
    - gcc-fortran-4.3-62.198.x86\_64.rpm
    - gcc43-4.3.4\_20091019-0.7.35.x86\_64.rpm
    - gcc43-c++-4.3.4\_20091019-0.7.35.x86\_64.rpm
    - gcc43-fortran-4.3.4\_20091019-0.7.35.x86\_64.rpm
    - glibc-devel-2.11.1-0.17.4.x86\_64.rpm
    - glibc-devel-32bit-2.11.1-0.17.4.x86\_64.rpm
    - libgfortran43-4.3.3\_20081022-11.18.x86\_64.rpm
    - libstdc++43-devel-4.3.4\_20091019-0.7.35.x86\_64.rpm
    - linux-kernel-headers-2.6.32-1.4.13.noarch.rpm
  - Install **after** ofed and mpi installation
    - ImageMagick-6.4.3.6-7.20.1.x86\_64.rpm
    - ImageMagick-devel-6.4.3.6-7.20.1.x86\_64.rpm
    - boost-devel-1.36.0-11.17.x86\_64.rpm
    - fontconfig-devel-2.6.0-10.6.x86\_64.rpm
    - freetype2-devel-2.3.7-25.8.x86\_64.rpm
    - gd-2.0.36.RC1-52.18.x86\_64.rpm
    - gd-devel-2.0.36.RC1-52.18.x86\_64.rpm
    - glib-1.2.10-737.22.x86\_64.rpm
    - glib-devel-1.2.10-737.22.x86\_64.rpm
    - glibc-info-2.9-13.2.x86\_64.rpm
    - libMagickWand1-6.4.3.6-7.18.x86\_64.rpm
    - libboost\_date\_time1\_36\_0-1.36.0-11.17.x86\_64.rpm
    - libboost\_filesystem1\_36\_0-1.36.0-11.17.x86\_64.rpm
    - libboost\_graph1\_36\_0-1.36.0-11.17.x86\_64.rpm
    - libboost\_iostreams1\_36\_0-1.36.0-11.17.x86\_64.rpm
    - libboost\_math1\_36\_0-1.36.0-11.17.x86\_64.rpm
    - libboost\_mpi1\_36\_0-1.36.0-11.17.x86\_64.rpm
    - libboost\_program\_options1\_36\_0-1.36.0-11.17.x86\_64.rpm

- libboost\_python1\_36\_0-1.36.0-11.17.x86\_64.rpm
- libboost\_serialization1\_36\_0-1.36.0-11.17.x86\_64.rpm
- libboost\_signals1\_36\_0-1.36.0-11.17.x86\_64.rpm
- libboost\_system1\_36\_0-1.36.0-11.17.x86\_64.rpm
- libboost\_test1\_36\_0-1.36.0-11.17.x86\_64.rpm
- libboost\_thread1\_36\_0-1.36.0-11.17.x86\_64.rpm
- libboost\_wave1\_36\_0-1.36.0-11.17.x86\_64.rpm
- libjpeg-devel-6.2.0-879.10.x86\_64.rpm
- liblcms-devel-1.17-77.14.19.x86\_64.rpm
- libpciaccess0-devel-7.4-8.24.2.x86\_64.rpm
- libpixman-1-0-devel-0.16.0-1.2.22.x86\_64.rpm
- libpng-devel-1.2.31-5.10.x86\_64.rpm
- libstdc++-devel-4.3-62.198.x86\_64.rpm
- libtiff-devel-3.8.2-141.6.x86\_64.rpm
- libuuid-devel-2.16-6.8.2.x86\_64.rpm
- libwmf-0.2.8.4-206.26.x86\_64.rpm
- libwmf-devel-0.2.8.4-206.26.x86\_64.rpm
- libwmf-gnome-0.2.8.4-206.26.x86\_64.rpm
- libxml2-devel-2.7.6-0.1.37.x86\_64.rpm
- libxml2-devel-32bit-2.7.6-0.1.37.x86\_64.rpm
- ncurses-devel-5.6-90.55.x86\_64.rpm
- openmotif-libs-2.3.1-3.13.x86\_64.rpm
- readline-devel-5.2-147.9.13.x86\_64.rpm
- tack-5.6-90.55.x86\_64.rpm
- tcl-devel-8.5.5-2.81.x86\_64.rpm
- tk-devel-8.5.5-3.12.x86\_64.rpm
- xorg-x11-devel-7.4-8.24.2.x86\_64.rpm
- xorg-x11-fonts-devel-7.4-1.15.x86\_64.rpm
- xorg-x11-libICE-devel-7.4-1.15.x86\_64.rpm
- xorg-x11-libSM-devel-7.4-1.18.x86\_64.rpm
- xorg-x11-libX11-devel-7.4-5.5.x86\_64.rpm
- xorg-x11-libXau-devel-7.4-1.15.x86\_64.rpm
- xorg-x11-libXdmcp-32bit-7.4-1.15.x86\_64.rpm
- xorg-x11-libXdmcp-devel-32bit-7.4-1.15.x86\_64.rpm
- xorg-x11-libXdmcp-devel-7.4-1.15.x86\_64.rpm
- xorg-x11-libXext-devel-32bit-7.4-1.14.x86\_64.rpm
- xorg-x11-libXext-devel-7.4-1.14.x86\_64.rpm

- xorg-x11-libXfixes-devel-7.4-1.14.x86\_64.rpm
  - xorg-x11-libXmu-devel-7.4-1.17.x86\_64.rpm
  - xorg-x11-libXp-devel-32bit-7.4-1.14.x86\_64.rpm
  - xorg-x11-libXp-devel-7.4-1.14.x86\_64.rpm
  - xorg-x11-libXpm-devel-32bit-7.4-1.17.x86\_64.rpm
  - xorg-x11-libXpm-devel-7.4-1.17.x86\_64.rpm
  - xorg-x11-libXprintUtil-devel-7.4-1.17.x86\_64.rpm
  - xorg-x11-libXrender-devel-7.4-1.14.x86\_64.rpm
  - xorg-x11-libXt-devel-32bit-7.4-1.17.x86\_64.rpm
  - xorg-x11-libXt-devel-7.4-1.17.x86\_64.rpm
  - xorg-x11-libXv-devel-7.4-1.14.x86\_64.rpm
  - xorg-x11-libfontenc-devel-7.4-1.15.x86\_64.rpm
  - xorg-x11-libxcb-devel-32bit-7.4-1.15.x86\_64.rpm
  - xorg-x11-libxcb-devel-7.4-1.15.x86\_64.rpm
  - xorg-x11-libxkbfile-devel-7.4-1.14.x86\_64.rpm
  - xorg-x11-proto-devel-7.4-1.21.x86\_64.rpm
  - xorg-x11-util-devel-7.4-1.15.x86\_64.rpm
  - xorg-x11-xtrans-devel-7.4-4.22.x86\_64.rpm
  - zlib-devel-1.2.3-106.34.x86\_64.rpm
- o OpenIPMI, ipmitool (needed to start ipmi service):
    - ipmitool-1.8.11-0.1.48
    - OpenIPMI-2.0.16-0.3.29

- Mesa-libGL-devel required to install nbody after cuda tool kit installation:
  - Mesa-libGL-devel-8.0.4-20.11.1.x86\_64.rpm
- mcelog package to decode kernel machine check log on x86 machines used by CT memory error check:
  - mcelog-1.0.2011.06.08-0.9.12

## Additional requirements for accelerator support

In addition to those listed in “Requirements for Cluster Test RPM” (page 5), the following are requirements for clusters with accelerators:

- To work with NVIDIA accelerators, the following NVIDIA software packages must be installed and functional. They should be installed in the order listed:

1. NVIDIA-Linux-x86\_64-319.32.run (Or newer)

The installation of NVIDIA drivers must be done with the X server shut down. This can be done by putting the node into `runlevel 3` using the `telinit` command.

2. `cuda_5.0.35_linux_64_<OS_used>.run`

Install the toolkit and the samples. The `nbody` sample must be compiled as it is used by Nvidia Accelerator tests.

Once the `cuda` toolkit is installed, the user environment must include `PATHs` to the new directories. One way of doing this is to create the file `/etc/profile.d/cuda.sh` with the following content (assuming `cuda` is installed in `/usr/local/cuda`).

```
PATH=/usr/local/cuda/bin:${PATH}
LD_LIBRARY_PATH=/usr/local/cuda/lib64:${LD_LIBRARY_PATH}
```

3. `gpucomputingsdk_4.2.9_linux.run`

Once installed, the applications must be compiled as follows.

```
# cd NVIDIA_GPU_Computing_SDK_C
# make
```

If compile errors related to missing libraries are reported, you will have to create `symlinks` from existing libraries to the missing names. For example:

```
# ln -s /usr/lib64/libGLU.so.1.3.070700 /usr/lib64/libGLU.so
```

Create the missing `symlinks` as needed until all the applications have been built.

Cluster Test requires the `nbody` and `deviceQuery` binaries present in the SDK directory `C/bin/linux/release`.

Add the path to the newly built binaries to your user environment.

To verify the NVIDIA components are working, run the `nvidia-smi -a` command and the `deviceQuery` command.

## Getting the CT RPM kit

1. Download the Cluster Test file self-extracting binary file `HP_Cluster_Test_Software_Vx.x_Clusterx64_rpm_vx.x_XXXX.bin` from <http://www.hpe.com/info/ct-download>.
2. Set the binary file as executable with the `chmod` command.

- Execute `HP_Cluster_Test_Software_Vx.x_Clusterx64_rpm_vx.x_XXXX.bin`. You will be presented with the HPE End User License Agreement. Use the space bar to scroll through the agreement. At the end, the following prompt is displayed:  
To accept the license please type `accept`:  
After typing `accept`, the `Clusterx64_rpm-vx.x-XXXX.noarch.rpm` file is extracted into the current directory.

## Procedure for installing CT RPM

The Cluster Test RPM installation is a two step process. The first step is to install the RPM package on each node in the cluster; the second step is to run a script to either install the prebuilt Cluster Test executables or build the executables from the Cluster Test sources.

Repeat the following steps on each node in the cluster.

- Install the Cluster Test sources on the node.

```
# rpm -ivh Clusterx64_rpm-vx.x-XXXX.noarch.rpm
Preparing... ##### [100%]
1:clustertest ##### [100%]
```

Run `'/opt/clustertest/setup/ct_setup.sh -i'` to complete the install

- Cluster Test RPM kits for external download and use are available for RHEL 6.4, RHEL 6.3, RHEL 6.2, RHEL 6.1, SLES 11.2, and SLES 11.1. You do not need to rebuild the RPMs if you downloaded the correct Cluster Test RPM kit for your OS and compatible MPI. The prebuilt binaries can be installed directly. Skip to [Step 3](#).


- To rebuild the RPMs for RedHat, SLES, or CentOS:

```
# /opt/clustertest/setup/ct_setup.sh -r
```

---

**NOTE:** The normal rebuild command will not rebuild the `ct_atlas` RPM. The atlas rebuild can take several hours and the included binaries have been shown to work across operating systems. Should you need to rebuild `ct_atlas`, you can do so by using the `-ra` option to `ct_setup` in place of the `-r` option shown above.

---

 **TIP:** The rebuild process takes some time to complete. If you have multiple nodes of the same configuration, you can run the rebuild command on one node and then copy the built packages located in `/opt/clustertest/setup/build/RPMS/x86_64/` to `/opt/clustertest/setup/` on the other nodes, then run `/opt/clustertest/setup/ct_setup.sh -i` on those nodes.

---

- Install the Cluster Test executables and scripts on the node.

```
# /opt/clustertest/setup/ct_setup.sh -i
```

Monitor the installation and note any errors that are reported. Errors indicate an installation problem. Two types of errors may be seen.

- An RPM is already installed. It is possible that an RPM being installed as part of the Cluster Test executable kit is already present.

For Red Hat or CentOS: Move the offending RPM from `/opt/clustertest/setup` to another directory. For example:

```
# mv /opt/clustertest/setup/perl-DateTime-0.47-1* /home/not-needed
```

For SLES: The list of Cluster Test Executable RPMs is contained in the text file `/opt/clustertest/setup/rpmlist.sles`. Make a backup copy of the file, then edit it removing the offending RPM.





**TIP:** In the event of an installation error, Hewlett Packard Enterprise recommends you uninstall the Cluster Test executables, correct the reported problem, then repeat the installation. To uninstall the Cluster Test Executable use:

For Red Hat, SLES, and CentOS: `/opt/clustertest/setup/ct_setup.sh -u`

- A prerequisite RPM is not installed. Although the installation checks and reports on the most common RPM prerequisites, it is possible that one of more required RPMs is not installed. In this case the installation will abort with a dependency error message specific to the missing component. The component must be installed before the Cluster Test RPMs will install.
4. Log out and then log in again to the node to update the `PATH` variable and load the Cluster Test aliases.
  5. Edit `/opt/clustertest/logs/.ct_config` and add:

```
INFINIBAND=generic
```
  6. If you are going to use OpenMPI, run `mpi-selector` to select it as the default MPI, then log out and log in again to pick up the changes.

```
# mpi-selector --list
mvapich2-1.9
openmpi-1.6.5
openmpi_gcc-1.6.4
#
```
  7. Perform the following procedures to be sure the `ib_send_bw` and Linpack tests run correctly.

- **ib\_send\_bw test**

The `ib_send_bw` test might not run successfully, generate an empty `ib_send_bw.analysis` file, and report errors in the `ib_send_bw.err` file such as `Conflicting CPU frequency values detected: 2000.000000 != 1200.000000`.

To be sure the test runs successfully, change the content of the `/sys/devices/system/cpu/cpu*/cpufreq/scaling_governor` files from `userspace` or `ondemand` to `performance` by copying and running the following scripts on all nodes. Use the `pdcp` and `pdsh` commands to copy to the same location and run the scripts on all cluster nodes.

```
#cat substitute_performance_in_Filescaling_governor.sh for CPUFREQ in
/sys/devices/system/cpu/cpu*/cpufreq/scaling_governor; do [ -f $CPUFREQ ] || continue; echo -n
performance > $CPUFREQ; done
#
```

After you run the script, verify your changes are reflected on all nodes.

```
#cat /sys/devices/system/cpu/cpu*/cpufreq/scaling_governor
performance
performance
performance
performance
performance
performance
....
....
#
```

- **Linpack test**

The node-wise Linpack test might generate the following STDERR messages:

```
Node0: sh: hponcfg: command not found
```

```
Node0: sh: hponcfg: command not found
```

or

```
node0: ERROR: Could not open input file
/opt/clustertest/bin/smartscripts/Get_FW_Version.xml
node0: Error processing the XML file:
/opt/clustertest/bin/smartscripts/Get_FW_Version.xml
```

To ensure the Linpack test runs successfully, install the `hponcfg` package and the `linux-LOsamplescripts4.11.0.tgz` `smartscripts` on all nodes.

1. Install the `hponcfg` RPM package on all nodes.

```
# rpm -ivh /home/hponcfg-4.2.0-0.x86_64.rpm
```

2. Create the directory `/opt/clustertest/bin/smartscripts/`.

```
# mkdir -p /opt/clustertest/bin/smartscripts/
```

3. Extract the `linux-LOsamplescripts4.11.0.tgz` archive file to `/opt/clustertest/bin/smartscripts/`.

```
# tar -xzf linux-LOsamplescripts4.11.0.tgz --directory=/opt/clustertest/bin/smartscripts/
```

---

**NOTE:** Download the Lights-Out XML PERL Scripting Sample for Linux `linux-LOsamplescripts4.11.0.tgz` archive from either of the following locations.

<http://www.hpe.com/support/iLO-XMLPerlScript-Linux>

<http://www.hpe.com/info/ftp-linux-LOsamples>

---

8. Update the Cluster Test configuration file to define the cluster specifics. Do this by invoking Cluster Test:

```
# testnodes.pl
```

If this is the first time you have started the GUI, you will see a message window:

```
5 of 12 config file entries are missing/invalid
```

After clicking **OK**, the Cluster Test configuration settings window is displayed. Check each entry and modify as necessary:

- Cluster Test base directory – The location of the Cluster Test installation. This should always be `/opt`.
- Admin network alias name – The node prefix used in `/etc/hosts` for each node. The default is "node".
- Interconnect type – Must be one of the following. Note the entry is case sensitive.
  - Admin – When there is no separate interconnect and the Admin network is to be used as the interconnect.
  - GigE – When the interconnect is GigE.
  - 10GigE – When the interconnect is 10GigE.
  - InfiniBand – When the interconnect is InfiniBand.
- Interconnect alias name – The interconnect alias prefix used in `/etc/hosts`. The default is "inode".
- Node number range managed by this head node – For large clusters there can be several head nodes, sometimes configured on a per rack basis. This entry specifies the nodes to be managed by this head node. For clusters with one head node, the entry is for the entire cluster, for example "0-63".
- Node number range of the entire cluster – Range of node numbers that describe the entire cluster, for example "0-63".

---

**NOTE:** The node number range managed by this head node and the node number range of the entire cluster are typically the same.

---

- Node number list of all head nodes – If there is only one head node this entry, its node number, for example “0”. If there are multiple head nodes, this is a comma-delimited list of all head nodes.
  - Alternate interconnect type – Must be one of the following. Entries are case sensitive.
    - None
    - GigE – When the interconnect is GigE.
    - 10GigE – When the interconnect is 10GigE.
    - InfiniBand: When the interconnect is InfiniBand.
  - Alternate interconnect alias name – The alternate interconnect alias prefix used in `/etc/hosts`. The default is "anode".
  - mpirun command – Full path to the `mpirun` command. The default is `/opt/platform_mpi/bin/mpirun`.
  - pdsh command – Full path to the Cluster Test-provided `pdsh` command. Default is `/opt/clustertest/bin/pdsh`.
  - Performance monitor command – Full path to the Cluster Test provided performance monitoring tools. Default is `/opt/clustertest/bin/xclus`.
  - light command – Command used to turn node UID light on and off. Default is `/opt/clustertest/bin/light`.
9. You are now ready to run Cluster Test.

## 3 Uninstalling Cluster Test

The RPM version of Cluster Test installs bits on the compute nodes, and has an uninstall script to remove those bits.

### Uninstalling CT RPM

Uninstalling Cluster Test RPM is a two step process. The first step involves removing the Cluster Test executables; the second step involves removing the Cluster Test sources. These steps must be executed on each node that has Cluster Test installed.

1. Run the uninstall script to remove the Cluster Test executables.

For Red Hat, SLES, and CentOS: `/opt/clustertest/setup/ct_setup.sh -u`

2. Remove the Cluster Test RPM.

```
# rpm -e Clusterx64_rpm-vx.x-XXXX.noarch.rpm
```

The uninstall process will remove all the files and directories used under normal Cluster Test operation, with the exception of the `logs` directory. If any files or directories are created or moved by the user in the `/opt/clustertest` directory, they might remain after the Cluster Test RPM removal. If Perl modules were installed along with the Cluster Test RPM, they will remain after uninstalling the Cluster Test RPM.

# 4 Support and other resources

## Accessing Hewlett Packard Enterprise Support

- For live assistance, go to the Contact Hewlett Packard Enterprise Worldwide website:  
[www.hpe.com/assistance](http://www.hpe.com/assistance)
- To access documentation and support services, go to the Hewlett Packard Enterprise Support Center website:  
[www.hpe.com/support/hpesc](http://www.hpe.com/support/hpesc)

### Information to collect

- Technical support registration number (if applicable)
- Product name, model or version, and serial number
- Operating system name and version
- Firmware version
- Error messages
- Product-specific reports and logs
- Add-on products or components
- Third-party products or components

## Accessing updates

- Some software products provide a mechanism for accessing software updates through the product interface. Review your product documentation to identify the recommended software update method.
  - To download product updates, go to either of the following:
    - Hewlett Packard Enterprise Support Center **Get connected with updates** page:  
[www.hpe.com/support/e-updates](http://www.hpe.com/support/e-updates)
    - Software Depot website:  
[www.hpe.com/support/softwaredepot](http://www.hpe.com/support/softwaredepot)
  - To view and update your entitlements, and to link your contracts and warranties with your profile, go to the Hewlett Packard Enterprise Support Center **More Information on Access to Support Materials** page:  
[www.hpe.com/support/AccessToSupportMaterials](http://www.hpe.com/support/AccessToSupportMaterials)
- 
- ① **IMPORTANT:** Access to some updates might require product entitlement when accessed through the Hewlett Packard Enterprise Support Center. You must have an HP Passport set up with relevant entitlements.
-

## Intended audience

It is assumed the reader has the experience in the following areas:

- the Linux operating system
- Hewlett Packard Enterprise hardware, including all HPE ProLiant models, Integrity servers, and ProCurve switches
- configuration of BIOS settings, iLO/IPMI, and ProCurve switches

## Websites

Website	Link
Hewlett Packard Enterprise Information Library	<a href="http://www.hpe.com/info/enterprise/docs">www.hpe.com/info/enterprise/docs</a>
Hewlett Packard Enterprise Support Center	<a href="http://www.hpe.com/support/hpesc">www.hpe.com/support/hpesc</a>
Contact Hewlett Packard Enterprise Worldwide	<a href="http://www.hpe.com/assistance">www.hpe.com/assistance</a>
Subscription Service/Support Alerts	<a href="http://www.hpe.com/support/e-updates">www.hpe.com/support/e-updates</a>
Software Depot	<a href="http://www.hpe.com/support/softwaredepot">www.hpe.com/support/softwaredepot</a>
Customer Self Repair	<a href="http://www.hpe.com/support/selfrepair">www.hpe.com/support/selfrepair</a>
Insight Remote Support	<a href="http://www.hpe.com/info/insightremotesupport/docs">www.hpe.com/info/insightremotesupport/docs</a>
Serviceguard Solutions for HP-UX	<a href="http://www.hpe.com/info/hpux-serviceguard-docs">www.hpe.com/info/hpux-serviceguard-docs</a>
Single Point of Connectivity Knowledge (SPOCK) Storage compatibility matrix	<a href="http://www.hpe.com/storage/spock">www.hpe.com/storage/spock</a>
Storage white papers and analyst reports	<a href="http://www.hpe.com/storage/whitepapers">www.hpe.com/storage/whitepapers</a>

## Customer self repair

Hewlett Packard Enterprise customer self repair (CSR) programs allow you to repair your product. If a CSR part needs to be replaced, it will be shipped directly to you so that you can install it at your convenience. Some parts do not qualify for CSR. Your Hewlett Packard Enterprise authorized service provider will determine whether a repair can be accomplished by CSR.

For more information about CSR, contact your local service provider or go to the CSR website:

[www.hpe.com/support/selfrepair](http://www.hpe.com/support/selfrepair)

## Remote support

Remote support is available with supported devices as part of your warranty or contractual support agreement. It provides intelligent event diagnosis, and automatic, secure submission of hardware event notifications to Hewlett Packard Enterprise, which will initiate a fast and accurate resolution based on your product's service level. Hewlett Packard Enterprise strongly recommends that you register your device for remote support.

For more information and device support details, go to the following website:

[www.hpe.com/info/insightremotesupport/docs](http://www.hpe.com/info/insightremotesupport/docs)

## Documentation feedback

Hewlett Packard Enterprise is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback ([docsfeedback@hpe.com](mailto:docsfeedback@hpe.com)). When submitting your feedback, include the document title, part number, edition, and publication date located on the front cover of the document. For online help content, include the product name, product version, help edition, and publication date located on the legal notices page.

# Index

## A

- accelerator
  - install requirements, 15
- accessing
  - updates, 21

## C

- Cluster Test
  - about, 4
  - installing, 5
  - RPM, 5
  - RPM requirements, 5
  - uninstalling, 20
- compute node
  - what Cluster Test installs, 20
- contacting Hewlett Packard Enterprise, 21
- customer self repair, 22

## D

- documentation
  - providing feedback on, 23

## F

- files
  - RPM test kit, 15

## M

- Message Passing Interface (MPI)
  - and Cluster Test RPM, 5

## O

- Open MPI requirements, 5

## P

- Perl modules, 5
- procedures
  - installing Cluster Test RPM, 16

## R

- remote support, 22
- requirements
  - accelerator support, 15
  - Open MPI, 5
  - RHEL 6.1 support, 6
  - RHEL 6.2 support, 7
  - RHEL 6.3 support, 7
  - RHEL 6.4 support, 8
  - RPM version, 5
  - SLES 11.1 support, 9
  - SLES 11.2 support, 12
- RHEL 6.1 requirements, 6
- RHEL 6.2 requirements, 7
- RHEL 6.3 requirements, 7
- RHEL 6.4 requirements, 8
- RPMs
  - OFED, 6

- removing, 20
- RHEL 6.1 requirements, 6
- RHEL 6.2 requirements, 7
- RHEL 6.3 requirements, 7
- RHEL 6.4 requirements, 8
- SLES 11.1 requirements, 9
- SLES 11.2 requirements, 12
- test kit download, 15

## S

- SLES 11.1 requirements, 9
- SLES 11.2 requirements, 12
- software
  - Cluster Test RPM, 5
- ssh, 5
- support
  - Hewlett Packard Enterprise, 21

## U

- updates
  - accessing, 21

## W

- websites, 22
  - customer self repair, 22