

Standard Features

The HP Cluster Test software has been used for several years by Hewlett Packard's manufacturing facilities as a diagnostic utility for cluster verification prior to customer shipments. This latest version, HP Cluster Test V4.0 (HP CT) is the first version of that software which has been licensed to HP customers so that they can perform their own cluster diagnostics.

HP CT software is designed to verify the proper operation of an HPC cluster configuration; to test the functionality of the cluster as a whole with emphasis on testing the cluster interconnect, including the interconnect switch, cabling, and interface cards; and to provide additional stress testing on the cluster nodes.

Furthermore, HP CT software is capable of detecting failures or inefficiencies of hardware and network connections prior to installation or upgrade of a complete cluster management stack. That way, failed components can be replaced or corrected before the cluster software installation is started. This makes cluster software integration, particularly on a complex solution with several hundred nodes, much easier and faster.

HP CT software supports HP ProLiant servers with optional GPGPUs, running a Linux operating system and connected with and InfiniBand and/or Ethernet network.

Interface

HP Cluster Test software comes with a graphical user interface (GUI) that includes two test control tabs comprising the Control Interface (CI): the Cluster Test Control Interface and the Accelerator Test Control Interface. The entire test suite is invoked via the GUI, and all output is automatically captured and saved to log files.

Cluster Tests

The Cluster Test Control Interface normally includes 8 test buttons: CrissCross, Test4, Pallas, Stream, Disk Test, Linpack, Netperf, and InfiniBand verbs tests. For systems using a QLogic InfiniBand interconnect, the CT installation procedure enables a 9th test button, Deviation test. Parameters, if needed, for each test are also given. For each of the tests, the run time can be specified.

- CrissCross: In this test, nodes take turn sending data packets to every other process.
- Test4: In this test, the nodes send small packets to all other nodes at the same time. This test basically floods the network with data and in the process also stresses the CPU(s).
- Pallas: This is the industry standard test suite formerly known as Pallas now known as IMB (Intel MPI Benchmark). Only a few of the tests are selected out of the original test suite: Exchange, SendRecv, Bcast, AllGather, AllReduce, and PingPong. Individual tests can be run by selecting the check boxes.
- Stream: is a standard memory benchmark.
- Disk Test: This test uses xdd to measure disk performance on all disks in a cluster. Parameters include the Disk device (or all to test every disk), the Percentage (%) of the disk size to be tested, and the Queue Depth to use, to simulate a more realistic test with multiple processes accessing the disks.
- Linpack: The test is a well-known matrix algorithm implemented to run separately on each node using all CPU cores and about 80% or less of memory. This test will not give the cluster wide performance but it exercises CPU and memory at the same time and it provides another set of numbers used to compare the relative performance among the cluster nodes
- Netperf: is a benchmark that is used to measure network performance (e.g. GigE and 10 GigE). It provides tests for unidirectional throughput and takes 2 parameters: UDP or TCP protocol.
- InfiniBand verbs tests: are a collection of InfiniBand verbs bandwidth/latency performance tests provided with OFED.
- Deviation test: A test to measure the consistency of bandwidth and latency between nodes using



Standard Features

MPI over QLogic InfiniBand.

Accelerator Tests

The Accelerator Test Control Interface includes 5 test buttons: Verify, SGEMM, Bandwidth, Memory Test, and Thermal Test. Parameters, if needed, for each test are also given.

- Verify: Each node selected is tested for the presence of one or more GPUs. If the device is detected (using lspci) then it is queried for information. If the query operation is successful the GPU is deemed to be functional.
 - SGEMM: Single Precision General Matrix Multiply Test for testing expected computational performance of the GPU
 - Bandwidth: GPU Bandwidth Test for testing expected data motion performance.
 - Memory Test: Writes then reads a pattern to memory and tests for errors
 - Thermal Test: Reports GPU temperatures for 5 minutes while a benchmark runs in the background. Temperature is compared to the documented expected temperature.
-

Reports

The GUI provides a rich set of reporting options. One option is a Nodes Monitoring Window in which is shown the set of nodes selected according to whether or not they are active, and if active, what is their network connectivity. For each node is a notation indicating whether it has active GPUs. Nodes can be activated or deactivated, powered off or rebooted, etc. using simple mouse clicks.

Another option of the GUI, is the Test Output Window. By tab selection, the user can choose to see live results, or logged results that can be saved to files. There are various display-options for each of the tests. All tests offer options to look at standard error, disk error and memory error output. There is also an analyze option for most tests, in which certain parameters are compared amongst nodes, and statistically anomalous results are reported. For some of those tests, there is also a graph option used to examine some test results as a function of the number of nodes.

Performance Monitor

The Performance Monitor is the graphic utility that enables monitoring the nodes in the cluster simultaneously. This utility is invoked from the Tools menu of the Control Interface GUI. The Performance Monitor displays an array of icons. Each icon represents a node in the cluster and provides the following information.

- Ethernet activity
- Core utilization
- Memory utilization
- Interconnect I/O activity
- Disk I/O activity

The data is color-coded with respect to the percent of utilization.

From the Performance Monitor window, a mouse click will invoke the xperf utility, which displays a dynamic graph showing the performance metrics for the node. The statistics are shown in a color-coded histogram. Pull-down menus describe the color coding for each metric.



Standard Features

Hardware Inventory

The Hardware Inventory graphical user interface is invoked from the Control Interface. This interface is designed to perform hardware inventory on all active nodes in the cluster. The hardware information includes server models, serial numbers, BIOS versions, BMC firmware versions, number of CPU's, NIC's status, and memory in MBytes.

Firmware Summary

The Firmware Summary graphical user interface is invoked from the Control Interface. This interface is designed to summarize firmware versions on all nodes in the cluster. Each line in the Firmware Summary display represents a group of nodes of the same model, BIOS firmware versions, ILO/MP firmware versions, Interconnect HCA firmware versions, and sorted based on node type[Hardware model] so it's easier to check for correct firmware versions.

Installation

The HP Cluster Test software, in this external release, must be installed on nodes that are already running a Linux operating system. (Internally at HP, the manufacturing facilities are able to use a variant of HP CT, in which the installation can be done on bare metal, but that version of HP CT is not currently provided external to HP, due to licensing restrictions.) A single HP CT RPM is installed on each of the nodes to be tested with the HP CT software. The Cluster Test RPM does not include an operating system, drivers, or message passing interface (MPI). These software packages must already be installed and configured on each node of the cluster prior to installing the Cluster Test RPM. Cluster Test RPM supports both HP-MPI and Open MPI. Each node being tested must have the appropriate MPI version installed and configured. In addition, ssh and a common node naming convention are required for Cluster Test RPM cluster communication. Details are provided in the CT documentation.

The CT documentation also provides information about Linux distribution RPM's that may need to be installed prior to installation of the CT RPM, and about drivers and other software tools that may be required depending on the cluster configuration and on the version of the operating system. After installation of the Cluster Test RPM, the remainder of the installation consists of running a script to either install the pre-built Cluster Test executables or build the executables from the Cluster Test sources. This must be done on each node. Tools such as HP Cluster Management Utility can be used to streamline this installation process.

When running HP Cluster Test the first time, a warning message will appear, notifying the administrator to set some configuration settings including, for example, the range of nodes to be tested by CT. Once this configuration is completed, HP CT will be ready to be used.



Standard Features

- Hardware Requirements** HP Cluster Test is supported on each node of a valid configuration such as the HP Cluster Platform 3000/3000BL/3000SL or HP Cluster Platform 4000/4000BL/4000SL.
- The supported Cluster Platform 3000 systems are based on HP ProLiant DL140, DL160, DL170h, DL360, and DL380 each of which uses Intel Xeon processors.
 - The supported Cluster Platform 3000BL systems are based on HP ProLiant blade BL2x220c, BL260c, BL280c, BL460c, and xw460c servers each of which use Intel Xeon processors.
 - The supported Cluster Platform 3000SL systems are based on HP ProLiant SL2x170z servers and HP ProLiant SL390s servers.
 - The supported Cluster Platform 4000 systems are based on HP ProLiant DL145, DL165, DL385, and DL585 servers each of which use AMD Opteron processors.
 - The supported Cluster Platform 4000BL systems are based on HP ProLiant blade BL465c, BL495c and BL685c servers each of which use AMD Opteron processors.
 - The supported CP4000SL systems are based on the HP ProLiant SL165z server.

As new HP servers are released, the above list of CT-supported platforms will be augmented. Please see the Cluster Test documentation for more information.

Each node may be connected to the others with a valid system interconnect that could be Gigabit Ethernet, 10 Gigabit Ethernet (Mellanox ConnectX EN), or InfiniBand (Mellanox, Voltaire and QLogic). For ProLiant SL390s servers, HP Cluster Test supports the HP Tesla M1060, M2050 and M2070 GPU Modules.

-
- Operating Systems** HP Cluster Test is supported on Red Hat Enterprise Linux V5 and SUSE Linux Enterprise Server V11 and on CentOS V5.

-
- Related High Performance Cluster Software**
- | | |
|-------------------------------|---|
| HP Cluster Platforms | HP Cluster Platforms are specifically engineered, factory-integrated large-scale ProLiant clusters optimized for High Performance Computing, with a choice of servers, networks and software. Operating system options include specially priced offerings for Red Hat Enterprise Linux and Novell SLES, as well as Microsoft Windows HPC Server. A Cluster Platform Configurator simplifies ordering. [http://www.hp.com/go/clusters] |
| HP HPC Interconnects | High Performance Computing (HPC) interconnect technologies are available for this server as part of the HP Cluster Platform portfolio. These high-speed InfiniBand and Gigabit interconnects are fully supported by HP when integrated within an HP cluster. Flexible, validated solutions can be defined with the help of configuration tools. [http://www.hp.com/techservers/clusters/ucp/index.html] |
| HP Cluster Management Utility | HP Cluster Management Utility (CMU) is an HP-licensed and HP-supported suite of tools that are used to manage large-scale Linux ProLiant systems. CMU includes software for the centralized provisioning, management and monitoring of nodes. CMU makes the administration of clusters user friendly, efficient, and effective. [http://www.hp.com/go/cmu] |
| HP HPC Linux Value Pack | HP HPC Linux Value Pack (Value Pack) is an HP-licensed and HP-supported specially priced software bundle for the development and deployment of applications on HPC Cluster Platforms. Value Pack includes the Platform LSF workload scheduler, the HP-MPI parallelization library, the HP Unified Parallel C compiler and the HP Shmem library, as well as the execution environments for the libraries and compiler. [HP HPC Linux Value Pack]. |



Ordering Information

HP Cluster Test software is available "as is" at no charge. HP does not provide technical support for this product, but users may send questions and comments to ClusterTestSupport@hp.com. There is no guarantee that questions will be addressed by HP, although as time permits, some responses may be provided.

The HP Cluster Test software will be available as part of the HP Cluster Management Utility distribution starting after CMU V4.2, and is available with CMU V4.2 as downloadable software from the CMU support area at the HP IT Resource Center support portal <http://www.itrc.hp.com/>.

The HP Cluster Test software can also be downloaded from HP Software Depot at: <http://www.hp.com/go/ct-download>. Documentation for HP Cluster Test software can be found at: <http://www.hp.com/go/ct-docs>.

Related High Performance Cluster Software	HP Cluster Management Utility Compute Node License NOTE: This part number can be used to purchase one certificate for multiple licenses with a single activation key. Each license is for one node (server). Customer will receive a printed end user license agreement and license entitlement certificate via physical shipment. The license entitlement certificate must be redeemed online in order to obtain a license key. NOTE: For additional license kits please see the QuickSpecs at: http://h18004.www1.hp.com/products/quickspecs/12612_div/12612_div.html	436284-B21
	HP Cluster Management Utility Media and Management License NOTE: Order a minimum of one license per cluster to purchase media including software and documentation, which will be delivered to the customer, and also licenses CMU management. No license key is delivered or required. NOTE: For additional license kits please see the QuickSpecs at: http://h18004.www1.hp.com/products/quickspecs/12612_div/12612_div.html	433257-B21
	HP HPC Linux Value Pack One Processor License NOTE: This part number can be used to purchase one certificate for multiple licenses with a single activation key. Each license is for one socket (a.k.a. processor). Customer will receive a printed end user license agreement and license entitlement certificate via physical shipment. The license entitlement certificate must be redeemed online in order to obtain a license key. NOTE: For additional license kits please see the QuickSpecs at: http://h18004.www1.hp.com/products/quickspecs/13485_div/13485_div.html	TC293A
	HP HPC Linux Value Pack Media NOTE: This part number can be used to purchase media including software and documentation, which will be delivered to the customer. NOTE: For additional license kits please see the QuickSpecs at: http://h18004.www1.hp.com/products/quickspecs/13485_div/13485_div.html	TC294A

Software Licensing Information

The following is taken from the End User License Agreement which comes with the software when delivered. Notwithstanding references to "purchase" of the Software, these terms and conditions apply to any use of the software, whether it was purchased or acquired in some other way.

a. Subject to the terms and conditions of this Agreement and the payment of any applicable license fee, HP grants You a non-exclusive, non-transferable license to Use (as defined below) in object code form one copy of the Software on one device at a time for Your internal business purposes, unless otherwise indicated above or in applicable Transaction Document(s). "Use" means to install, store, load, execute and display the Software in accordance with the Specifications. Your Use of the Software is subject to these license terms and to the other restrictions specified by HP in any other tangible or electronic documentation delivered or otherwise made available to You with or at the time of purchase of the



Ordering Information

Software, including license terms, warranty statements, Specifications, and "readme" or other informational files included in the Software itself. Such restrictions are hereby incorporated in this Agreement by reference. Some Software may require license keys or contain other technical protection measures. You acknowledge that HP may monitor your compliance with Use restrictions remotely or otherwise. If HP makes a license management program available which records and reports license usage information, You agree to appropriately install, configure and execute such license management program beginning no later than one hundred and eighty (180) days from the date it is made available to You and continuing for the period that the Software is Used.

b. This Agreement confers no title or ownership and is not a sale of any rights in the Software. Third-party suppliers are intended beneficiaries under this Agreement and independently may protect their rights in the Software in the event of any infringement. All rights not expressly granted to You are reserved solely to HP or its suppliers. Nothing herein should be construed as granting You, by implication, estoppel or otherwise, a license relating to Software other than as expressly stated above in this section 2.

c. Unless otherwise permitted by HP, You (a) may only make copies or adaptations of the Software for archival purposes or when copying or adaptation is an essential step in the authorized Use of the Software on a backup device, provided that copies and adaptations are used in no other manner and provided further that the Use on the backup device is discontinued when the original or replacement device becomes operable, and (b) may not copy the Software onto or otherwise Use or make it available on, to, or through any public or external distributed network.

d. To Use Software identified as an update or upgrade, You must first be licensed for the original Software identified by HP as eligible for the update or upgrade. If the update or upgrade is intended to substantially replace the original Software, after updating or upgrading, You may no longer Use the original Software that formed the basis for Your update or upgrade eligibility unless otherwise provided by HP in writing. Nothing in this Agreement grants You any right to purchase or receive Software updates, upgrades, or support, and HP is under no obligation to make such support available to you. Updates, upgrades, enhancements, or other Support may only be available under separate HP support agreements. You may contact HP to learn more about any support offerings HP may make available. HP reserves the right to require additional licenses and fees for Software upgrades or other enhancements, or for Use of the Software on upgraded devices.

e. You must reproduce all copyright notices that appear in or on the Software (including documentation) on all permitted copies or adaptations. Copies of documentation are limited to internal use.

f. Notwithstanding anything to the contrary herein, if the Transaction Document(s) identifies that the Software may be utilized on another Designated System(s) (as defined below), Your license to Use the Software may be transferred to another Designated System(s).
A "Designated System" means a computer system owned, controlled, or operated by or solely on behalf of You and may be further identified by HP by the combination of a unique number and a specific system type. Such license will terminate in the event of a change in either the system number or system type, an unauthorized relocation, or if the Designated System ceases to be within Your possession or control.

g. Operating system Software may only be Used when operating the associated hardware in configurations as approved, sold, or subsequently upgraded by HP or an HP authorized reseller.

h. Software is not specifically designed, manufactured, or intended for use as parts, components, or assemblies for the planning, construction, maintenance, or direct operation of a nuclear facility. You are solely liable if Software is Used for these applications and will indemnify and hold HP harmless from all loss, damage, expense, or liability in connection with such Use.



Ordering Information

i. You will not modify, reverse engineer, disassemble, decrypt, decompile, or make derivative works of the Software. Where You have other rights mandated under statute, You will provide HP with reasonably detailed information regarding any intended modifications, reverse engineering, disassembly, decryption, or decompilation and the purposes therefore.

j. Extending the Use of Software to any person or entity other than You as a function of providing services, (i.e.; making the Software available through a commercial timesharing or service bureau) must be authorized in writing by HP prior to such Use and may require additional licenses and fees. You may not distribute, resell, or sublicense the Software.

k. Notwithstanding anything in this Agreement to the contrary, all or any portion of the Software which constitutes Ancillary Software is licensed to You subject to the terms and conditions of the Software license agreement accompanying such Ancillary Software, whether in the form of a separate agreement, shrink wrap license or electronic license terms accepted at time of download. Use of the Ancillary Software by You shall be governed entirely by the terms and conditions of such license and, with respect to HP, by the limitations and disclaimers of sections 3 and 5 hereof. HP has identified any Ancillary Software by either noting the Ancillary Software provider's ownership within each Ancillary Software program file and/or by providing information in the "ancillary.txt" or "readme" file that is provided as part of the installation of the Software. The Ancillary Software licenses are also set forth in the "ancillary.txt" or "readme" file. By accepting the terms and conditions of this Agreement, You are also accepting the terms and conditions of each Ancillary Software license in the ancillary.txt or "readme" file. If the Software includes Ancillary Software licensed under the GNU General Public License and/or under the GNU Lesser General Public License ("GPL Software"), a complete machine-readable copy of the GPL Software Source Code ("GPL Source Code") is either:

(i) included with the Software that is delivered to You; or (ii) upon your written request, HP will provide to You, for a fee covering the cost of distribution, a complete machine-readable copy of the GPL Source Code, by mail, or (iii) if You obtained the Software by downloading it from a HP website and neither of the preceding options are available, you may download the GPL Source Code from the same website. Information about how to make a written request for GPL Source Code may be found in the ancillary.txt file or, if an address is not listed in that file, at the following website: www.hp.com.

Software and Services Warranty

DISCLAIMER OF WARRANTIES:

TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, HP AND ITS SUPPLIERS PROVIDE THE SOFTWARE "AS IS" AND WITH ALL FAULTS, AND HEREBY DISCLAIM ALL INDEMNITIES, WARRANTIES AND CONDITIONS, EITHER EXPRESS, IMPLIED, WHETHER BY STATUTE, COMMON LAW, CUSTOM OR OTHERWISE, INCLUDING, BUT NOT LIMITED TO, WARRANTIES OF TITLE AND NON-INFRINGEMENT, ANY IMPLIED WARRANTIES, DUTIES OR CONDITIONS OF MERCHANTABILITY, OF FITNESS FOR A PARTICULAR PURPOSE, AND OF LACK OF VIRUSES. HP does not warrant that the operation of Software will be uninterrupted or error free or that the Software will meet Your requirements. Some states/jurisdictions do not allow exclusion of implied warranties or limitations on the duration of implied warranties, so the above disclaimer may not apply to You in its entirety.



Ordering Information

© Copyright 2003-2010 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.

Intel® and Itanium® and Xeon® are trademarks or registered trademarks of Intel Corporation in the U.S. and other countries and are used under license. Pentium® is a U.S. registered trademark of Intel Corporation. UNIX® is a registered trademark of The Open Group. AMD and AMD Opteron are trademarks or registered trademarks of Advanced Micro Devices, Inc

NVIDIA® is a registered trademark of NVIDIA Corporation.

Red Hat® is a registered trademarks of Red Hat, Inc.

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

Red Hat® is a registered trademark of Red Hat, Inc. in the United States and other countries.

Linux® is a registered trademark of Linus Torvalds.

Mellanox is a registered trademark of Mellanox Technologies, Inc. and ConnectX is a trademark of Mellanox Technologies, Inc.

QLogic is a registered trademark of QLogic Corporation.

This product includes software developed at the University of Tennessee, Knoxville, Innovative Computing Laboratory.

Restricted Rights Legend

Use, duplication or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c) (1) (ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 for DOD agencies, and subparagraphs (c) (1) and (c) (2) of the Commercial Computer Software Restricted Rights clause at FAR 52.227-19 for other agencies.

HEWLETT-PACKARD COMPANY
3000 Hanover Street
Palo Alto, California 94304 U.S.A.

Use of this QuickSpecs and media is restricted to this product only. Additional copies of the programs may be made for security and back-up purposes only. Resale of the programs, in their present form or with alterations, is expressly prohibited. A variety of customer service options are available from Hewlett-Packard for CMU. For more information, contact your local Hewlett-Packard office.

Amongst the service options are software factory installation services, where CMU is installed and configured on Cluster Platforms at the factory. There are also on-site Consulting and Integration Services available for CMU. Contact your local Hewlett-Packard office or your Hewlett-Packard services representative for more information.

