

HP Cluster Management Utility

Simple, scalable, flexible management for Linux clusters

Data sheet



HP Cluster Management Utility (HP CMU) simplifies the management and the rapid deployment of large Linux clusters and groups of clusters. HP CMU's flexibility and efficient interface enables optimal workload performance and reduced administrative costs.

HP Cluster Management Utility

A simple and affordable cluster toolkit, HP CMU is used to configure and install cluster operating environments, monitor cluster and node metrics, and remotely manage resources. HP CMU serves as a powerful tool for installing Linux software images, including middleware such as Message-Passing Interface (MPI) and job schedulers. You can use HP CMU anywhere you need to manage a number of standalone systems that are similar in hardware and software configuration.

Designed to manage high-performance computing (HPC) clusters or CPU farms, HP CMU operates independently of the underlying operating system, interconnects and software applications. This flexible design enables CMU to operate across multiple customized Linux software stacks, and manage groups of nodes running in parallel with different Linux distributions.

From an historical perspective

The emergence of Linux as the most popular HPC operating system, the rise in performance of x84-64 processors with Intel® Xeon™ and AMD Opteron chips, and the high performance now achievable with Gigabit Ethernet and InfiniBand have created an enormous demand for industry-standard clusters.

While the HPC cluster architecture has proven its efficiency in terms of performance and price, managing configurations of hundreds or thousands of systems is challenging. To support a range of customers—each with a unique cluster implementations—HP CMU was designed to manage key tasks, independent of Linux variant, MPI, or other software components. HP CMU is an ideal tool for customers implementing a cluster solution utilizing multiple open source and/or third-party software components, and customers who need a simple interface for installation and administration of HPC clusters.

First developed in the late 1990s for AlphaServers, HP CMU now supports HP ProLiant servers—including the latest HP ProLiant SL series and HP BladeSystem infrastructures—and includes recent enhancements to enable GPU-based computation. Companies around the world, including many organizations on the Top500¹, use HP CMU to manage a range of clusters.

¹ Semi-annual listing of the Top500 performing systems, as listed at Top500.org




HP CMU is supported on Red Hat Linux, Novell SUSE Linux Enterprise Server (SLES) and multiple community distributions. The toolkit has also been deployed with popular workload managers such as Altair Engineering PBS Pro, Adaptive Computing Moab, and Platform LSF. HP works with its partners to support dynamic provisioning based on workload, and resource scheduling to minimize power consumption.

HP CMU architecture

HP CMU manages a cluster from a single management or “head” node. Making use of DHCP, NFS, and tftp services, HP CMU can PXE-boot the cluster nodes during installation and any subsequent software updates. Using the features of the Baseboard Management Card (BMC) of each compute node server (for example, Integrated Lights-Out or Lights-Out 100) HP CMU enables:

- Remote text-based console availability during all server states (set-up, boot, OS, or halted)
- Remote control of server power, regardless of the state of the server (even if the server is off)
- Remote BIOS set-up

Figure 1: GUI sample screen shot displaying global monitoring information of a cluster



Some customers deploy and support open source cluster stacks, using HP CMU to simplify installation and management. A description of a stack deployed by the HP Competency Center for Asia Pacific Japan is available at www.hp.com/go/cmu.

CMU features

Designed to streamline management of many compute nodes (4000+), HP CMU ships with a complete Java™ graphical user interface (GUI) for remote management of the cluster, and a command-line interface (CLI) for simple interactive and/or scriptable cluster management from the head node.

CMU provides central access to all nodes in the cluster in several ways—from individual access to the OS or BMC of a node to a parallel distributed shell over all nodes. The GUI provides a multi-window broadcast capability that displays OS or serial console activity from a selected group of nodes, and broadcasts commands from a single keyboard session.

With a few mouse clicks in the GUI or a single command in the CLI, HP CMU can shut down, boot up, reboot, or power off any selection of nodes.


You can use HP CMU to deploy software to the cluster hardware in several ways—the most efficient and scalable method being the HP CMU backup and cloning mechanism. You configure one compute node, and HP CMU will copy and distribute that disk image to some or all cluster nodes using an intelligent hierarchical process. HP CMU can store multiple disk images, which means you can deploy different software stacks, including different Linux OS distributions, within a single cluster. CMU also includes tools to support installation of GPU drivers and CUDA software.

HP CMU also supports kick-starting compute nodes, as well as deploying Linux in a diskless environment.

The HP CMU GUI displays real-time monitoring of selected metrics from some or all nodes in an intuitive and uniquely scalable design. This capability allows the system manager to view the status of the entire cluster at a glance or to focus on a group of nodes and change the displayed metrics as necessary.

CMU provides a rich set of pre-configured metrics and the flexibility to add customized metrics. System administrators can also customize the grouping of nodes to enhance monitoring. CMU makes available a script that integrates with various batch schedulers to dynamically create and delete groups of nodes in CMU that correspond with the workloads currently active on the cluster.

Step 1: Cloning one node per network entity



Cloning feature

HP CMU offers the ability to clone the disk contents of one node to other cluster nodes. This feature reduces the complexity of scale and ensures a consistent configuration of compute nodes within the cluster.

A cluster is often split into logical groups, which may be owned by different parts of the IT organization. HP CMU can manage these logical groups by associating one or more disk images with each group. Once one node in a group has been installed and configured, HP CMU builds a compressed image (called a “clone image”) of this master disk. The clone image is then ready to be propagated to other members of the group using the HP CMU cloning mechanism.


Cloning chronology

A node that is ready for cloning is first rebooted to free the local disk of any operating system activity. Then the diskless HP CMU imaging environment is loaded on the node, and a backup image of the system disk is captured and stored.

Once a successful backup has occurred, the cloning process can be invoked. Using a tree propagation approach, this process is divided in two steps to protect the network topology.

In the first step, one node from each network entity is given a copy of the backup image. A CMU network entity is a group of nodes that share a common local switch. For example, all blades in a blade enclosure belong to a network entity. Once that node receives the backup image, it begins to propagate the image to the other nodes within the network entity.

Step 2: Cloning inside the network entities



Once a node has received a completed image, the node tries to download the image to another node within the group. HP CMU manages the group node list waiting for an image. This algorithm parallelizes the propagation process and takes advantage of the available network bandwidth. This high performance tree propagation algorithm can quickly clone a full operating system and its applications.

The propagation mechanism makes efficient use of network bandwidth and is specifically designed to take advantage of switched network configurations.

Monitoring feature

The monitoring feature was designed to synthesize environmental, performance and administration information from the cluster for the IT manager. In this way, HP CMU helps decrease total cost of solution ownership (TCO). Visibility to the cluster is presented to the administrator in a single system view, as opposed to a compilation of data collected on the compute engines distributed over the cluster.

This feature, similar to all other features of the HP CMU, is designed for scalability. No additional resources are needed to ensure high-performance, highly efficient monitoring—even as a cluster grows to include thousands of nodes. To enhance monitoring even further, HP CMU includes a popular monitoring package called collectl, comprised of a large number of lightweight sensors. Collectl can be configured to store monitoring data for later analysis. The data can be viewed directly, and historical displays can be viewed using the colplot software, also included with CMU.

CMU also monitors GPU health and utilization, displaying metrics such as GPU and GPU memory load, and GPU and IOH temperature.

Table 1: CMU support matrix

	Processor	Linux distribution	Cluster nodes supported	Cluster interconnects
HP Cluster Platform	Intel Xeon	<ul style="list-style-type: none">• Red Hat• Novell SUSE	<ul style="list-style-type: none">• ProLiant DL160, DL380, DL170h• Blades BL460c, BL280c, BL2x220c	<ul style="list-style-type: none">• InfiniBand• Gigabit Ethernet
HP Cluster Platform 4000	AMD Opteron	<ul style="list-style-type: none">• Red Hat• Novell SUSE	<ul style="list-style-type: none">• ProLiant DL165, DL385, DL585• Blades BL465c	<ul style="list-style-type: none">• InfiniBand• Gigabit Ethernet
HP Cluster Platform 3000SL/4000SL	Intel Xeon	<ul style="list-style-type: none">• Red Hat• Novell SUSE 1	<ul style="list-style-type: none">• ProLiant SL170s, SL390s, SL2x170z, SL165z	<ul style="list-style-type: none">• InfiniBand• Gigabit Ethernet
Custom HP cluster	x86 Intel Itanium®	<ul style="list-style-type: none">• Debian• Scientific Linux• CentOS	<ul style="list-style-type: none">• HP hardware• Non-HP hardware	<ul style="list-style-type: none">• Quadrics• InfiniBand• Myrinet• Gigabit Ethernet

Figure 2: GUI sample screen shot displaying global monitoring information of a cluster



The monitoring data are accessible through the monitoring graphic user interface, which is available locally on the HP CMU Management Node or remotely from any workstations running Java that can have network access to the HP CMU Management Node.

The primary advantage of HP CMU's monitoring GUI is to provide a single system view of the cluster. The monitored data of all individual compute nodes displays in a single pie view, providing the administrator with a comprehensive picture of the entire cluster's behavior.

Comprehensive support and services

Available factory installed, HP CMU software is a featured cluster management option for HP Cluster Platforms. CMU software is supported worldwide by the service professionals at HP. CMU software can be complemented by a full range of services from design to factory pre-installation and configuration to acceptance and training. A public forum for CMU users is also hosted at HP's IT Resource Center.

To understand how the HP CMU tool and the HP Unified Cluster portfolio can help you drive business performance, visit: www.hp.com/go/cmu and www.hp.com/go/clusters

Share with colleagues



Get connected

www.hp.com/go/getconnected

Current HP driver, support, and security alerts delivered directly to your desktop

© Copyright 2010 Hewlett-Packard Development Company, L.P. 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 is a trademark of Advanced Micro Devices, Inc. Java is a U.S. trademark of Sun Microsystems, Inc. Intel, Xeon, and Itanium are trademarks of Intel Corporation in the U.S. and other countries.

