

HP Cluster Platform

Microsoft Windows Compute Clusters

HP Part Number: A-CPWCS-1C
Published: September 2007



© Copyright 2007 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.

Voltaire™ and GridStack™ are trademarks of Voltaire Inc.

InfiniBand™ is a trademark the InfiniBand® Trade Association

AMD, the AMD Arrow logo and combinations thereof, and AMD Opteron™ are trademarks of Advanced Micro Devices, Inc.

Intel® and Xeon™ are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Microsoft®, Windows®, the Windows NT operating system, and Windows Server 2003 operating system are U.S. registered trademarks of Microsoft Corporation. Compute Cluster Server 2003 is a U.S. trademark of Microsoft Corporation.

Table of Contents

About This Document.....	9
Intended Audience.....	9
Document Organization.....	9
Typographic Conventions.....	9
Related Information.....	10
HP Cluster Platform Documentation.....	10
HP Cluster Platform Links.....	10
HP Documentation for Microsoft Windows Compute Cluster 2003.....	10
HP Encourages Your Comments.....	10
1 Understanding Windows Compute Cluster 2003 on Cluster Platform.....	11
2 Understanding Supported Cluster Configurations.....	13
2.1 Understanding Configuration Constraints.....	13
2.2 Identifying Supported Topologies.....	13
2.3 Understanding Maximum System Limits.....	16
2.4 Identifying Host Channel Adapters and Firmware.....	16
2.5 Identifying the Servers Models Supported as Nodes.....	16
2.6 Upgrading or Modifying Clusters.....	17
3 Setting up your HP Cluster Platform for Microsoft Windows Compute Cluster 2003.....	19
4 Verifying the InfiniBand Stack Software.....	21
4.1 Verifying the Stack Installation.....	21
4.2 Installing the Stack.....	21
4.3 Configuring the Stack.....	22
5 HP Cluster Platform Terminology.....	23
Index.....	27

List of Figures

2-1	Two-Network Topology.....	14
2-2	Three-Network Topology.....	15

List of Tables

2-1	Supported HCA Cards.....	16
2-2	Supported Servers.....	16
5-1	HP Cluster Platform Terminology.....	23

About This Document

This manual describes the implementation of the Microsoft® Windows® Compute Cluster Server 2003 operating system on HP Cluster Platform models, including HP Cluster Platform Express.

Intended Audience

This document is intended for HP Cluster Platform administrators and end users. Certain operations described in this document, if performed incorrectly, might cause system crashes and loss of data. Ensure that you have the appropriate expertise before you begin any operation.

This guide is also intended for HP service representatives and other persons trained to install and service hardware options in the HP 10000-series racks. Such persons are expected to understand the hazards of working in this environment and to take suitable precautions to minimize danger to themselves and others.

Document Organization

This document is organized as follows:

Chapter 1 (page 11)	Describes the implementation of Microsoft Windows Compute Cluster 2003 on HP Cluster Platform.
Chapter 2 (page 13)	Microsoft Windows Compute Cluster 2003 supports several different cluster topologies. This chapter explains which topologies map to HP Cluster Platform.
Chapter 3 (page 19)	Describes the basic steps for configuring your cluster to install Microsoft Windows Compute Cluster 2003.
Chapter 4 (page 21)	This chapter describes the required GridStack™ software that must be installed on HP Cluster Platform to support Microsoft Windows Compute Cluster 2003.
Chapter 5 (page 23)	Maps HP Cluster Platform terminology to Microsoft Windows Compute Cluster 2003 terms and concepts.

Typographic Conventions

This document uses the following typographic conventions:

Reader Notes



NOTE: Content of the note.

Notes provide additional information, supplementing the adjoining content or emphasizing points of information.

Reader Warnings



WARNING! Content of the warning.

A warning calls attention to important information that if not understood or followed will result in personal injury or nonrecoverable system problems.

Reader Cautions



CAUTION: Content of the 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: Content of the important information.

This alert provides essential information to explain a concept or to complete a task.

Related Information

The following sections describe where to find more information about the Cluster Platform products.

HP Cluster Platform Documentation

HP Cluster Platform ships with a hardware documentation CD-ROM containing all books for all models of cluster. To mount the CD-ROM, and identify the books that you need for your specific model of HP Cluster Platform, read the printed *Customer Letter*.



IMPORTANT:

It is important that you visit <http://www.docs.hp.com/en/highperfcomp.html> and download the latest versions of the HP Cluster Platform documentation. You can identify updated books by their part number, such as *A-CPWCS-1A* and the date of publication. HP updates part numbers by incrementing the last two characters. For example; *A-CPWCS-1B* and *A-CPWCS-2G*

HP Cluster Platform Links

HP's Unified Cluster Portfolio offers comprehensive modular solutions utilizing industry-standard servers, with a choice of operating environments. For more information, go to the following URL:

<http://h20311.www2.hp.com/HPC/cache/274276-0-0-0-121.html>

You can find the latest versions of HP Cluster Platform guides at the following URL:

<http://docs.hp.com/>

You can also find Voltaire InfiniBand and Microsoft Windows Gridstack documentation at the following URL:

http://www.voltaire.com/partner_login.htm

See the section on grid computing. You will require the appropriate login credentials, which are available from Voltaire.

HP Documentation for Microsoft Windows Compute Cluster 2003

Your preinstalled cluster platform solution employing Microsoft Windows Compute Cluster 2003 ships with a CD of documentation, containing the following HP publications:

- *HP Installed Microsoft® Windows Compute Cluster 2003, User Guide*
- *HP Installed Microsoft® Windows Compute Cluster Server 2003 Getting Started*

Full Microsoft documentation for Microsoft Windows Compute Cluster 2003 is also provided on the CD.

HP Encourages Your Comments

HP encourages your 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 Understanding Windows Compute Cluster 2003 on Cluster Platform

Microsoft Windows Compute Cluster 2003 is preinstalled on your HP Cluster Platform during factory integration of the cluster. After you complete any on-site cabling (described in the HP Cluster Platform hardware documentation), you need only complete a few software installation steps to enable the cluster for use. The sequence of installation tasks is as follows:

1. Unpack and align the racks, using the information provided in the *HP Cluster Platform Site Preparation Guide*.
2. Cable the cluster together using the information provided in the following documents, which you can determine by using the guidelines provided on the hardware documentation CD.
 - Use the *HP Cluster Platform Server and Workstation Overview* to identify and cable cluster nodes (servers).
 - Use the interconnect guide for your model of HP Cluster Platform to cable the system interconnect. In the current release, Microsoft Windows Compute Cluster 2003 supports only:
 - InfiniBand protocol, using Voltaire system interconnects. See:
 - *HP Cluster Platform Voltaire InfiniBand Interconnect Installation and User's Guide*
 - *HP Cluster Platform InfiniBand Fabric Management and Diagnostic Guide*
 - Gigabit Ethernet protocol, using HP ProCurve switches as system interconnects. See the *HP Cluster Platform Gigabit Ethernet Interconnect Guide*.

The system interconnect is the network fabric that carries all the high-speed message passing interface (MPI) traffic. Depending on your model of HP Cluster Platform, this fabric might be dedicated to MPI traffic, or it might be shared with other cluster traffic. See Chapter 2.

 - The cabling tables for your type of system interconnect, cluster model and specific implementation (such as the number of nodes):
 - *HP Cluster Platform Cabling Guide for Gigabit Ethernet Solutions*
 - *HP Cluster Platform Infiniband Cabling Tables*
3. Power up the cluster, as described in the the interconnect guide for your model of HP Cluster Platform or the *HP Cluster Platform Cluster Platform Express Installation and User's Guide*, if you have a Cluster Platform Express model.
4. Perform any hardware post-installation verification procedures, to ensure that all cabling is functioning. For InfiniBand clusters, see the *HP Cluster Platform InfiniBand Interconnect*. Gigabit Ethernet clusters have no specific post-installation verification procedures, other than those documented in the HP ProCurve hardware documentation shipped with your cluster.

2 Understanding Supported Cluster Configurations

This release of Microsoft Windows Compute Cluster 2003 is constrained to certain models and specific configurations of HP Cluster Platform. Supported configurations are automatically verified during the order and build process for factory-installed clusters. The information in this chapter defines the supported configurations and components.

The following components and topics are discussed:

- For information about the limitations on cluster designs, see [Section 2.1](#)
- For definitions of supported cluster topologies, see [Section 2.2](#)
- The maximum size of a cluster configuration, see [Section 2.3](#)
- To identify the supported InfiniBand HCAs, see [Section 2.4](#)
- To identify the server models supported as nodes,, see [Section 2.5](#)
- For information about adding or removing components, see [Section 2.6](#)

2.1 Understanding Configuration Constraints

The configuration of an HP Cluster Platform is flexible, so that it might support several different operating environments. The cluster architecture is explained in *HP Cluster Platform, Cluster Platform Overview*.

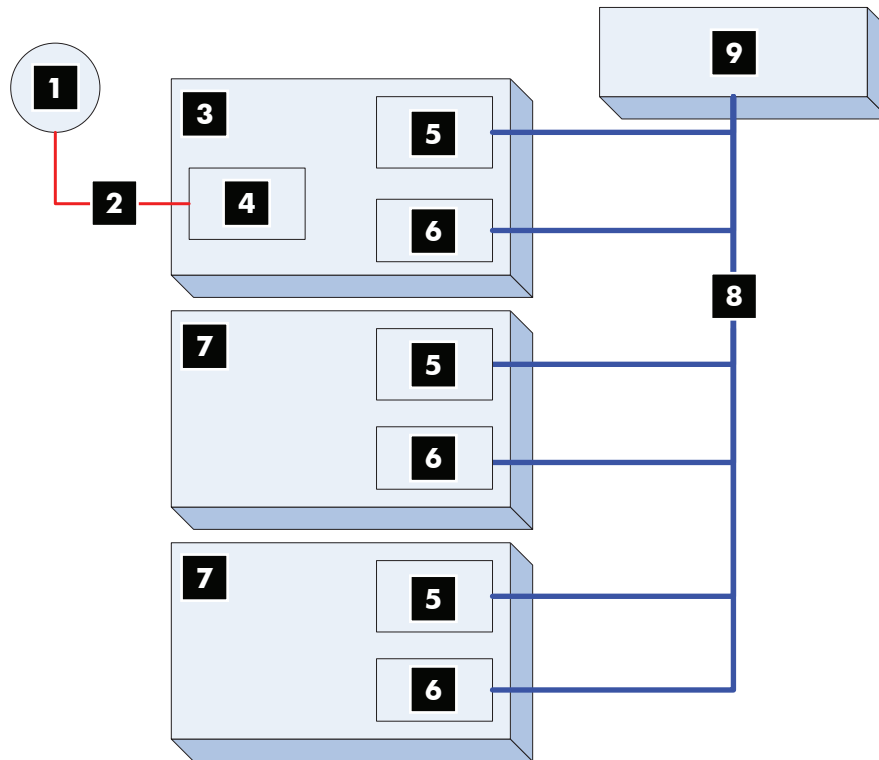
To support Microsoft Windows Compute Cluster 2003, a cluster must have the following specific characteristics:

- The cluster consists of a single head node and a number of compute nodes. The HP Cluster platform concept of a utility node is not supported.
- Only servers (not workstations) are supported as the head node or compute nodes. The supported server types are defined in [Section 2.5](#).
- The head node requires an optical drive (CD reader).
- Nodes may differ by model, providing the models have identical CPU type and memory and I/O configurations, compliant with the basic system requirements for HP Cluster Platform. Clusters configured this way will support the concept of grouping servers into pools of up to four pools (pool 1 through pool 4).
- Only the following types of system interconnect are supported as the high-speed Message Passing Interface (MPI) network:
 - InfiniBand, which is typically employed in clusters that require a dedicated high-speed MPI network, segregating job traffic from cluster administration and job management traffic.
 - Gigabit Ethernet, which might be employed for a dedicated MPI network, or as an in-band network where cluster administration and job management traffic passes over the same network used for the MPI traffic. (Typical for implementations of Microsoft Windows Compute Cluster 2003 on HP Cluster Platform Express.)
- The HP Cluster Platform concept of separate console and administrative networks does not apply where Microsoft Windows Compute Cluster 2003 is employed as the operating environment.

2.2 Identifying Supported Topologies

Microsoft Windows Compute Cluster 2003 supports five different cluster topologies. Only two of these topologies map to HP Cluster Platform topologies. [Figure 2-1](#) shows the simplest topology, based on in-band (shared) use of the Cluster's HP ProCurve switch as the MPI fabric. The same network provides routing for cluster administrative and job management traffic. A second network is provided via the site LAN connection.

Figure 2-1 Two-Network Topology

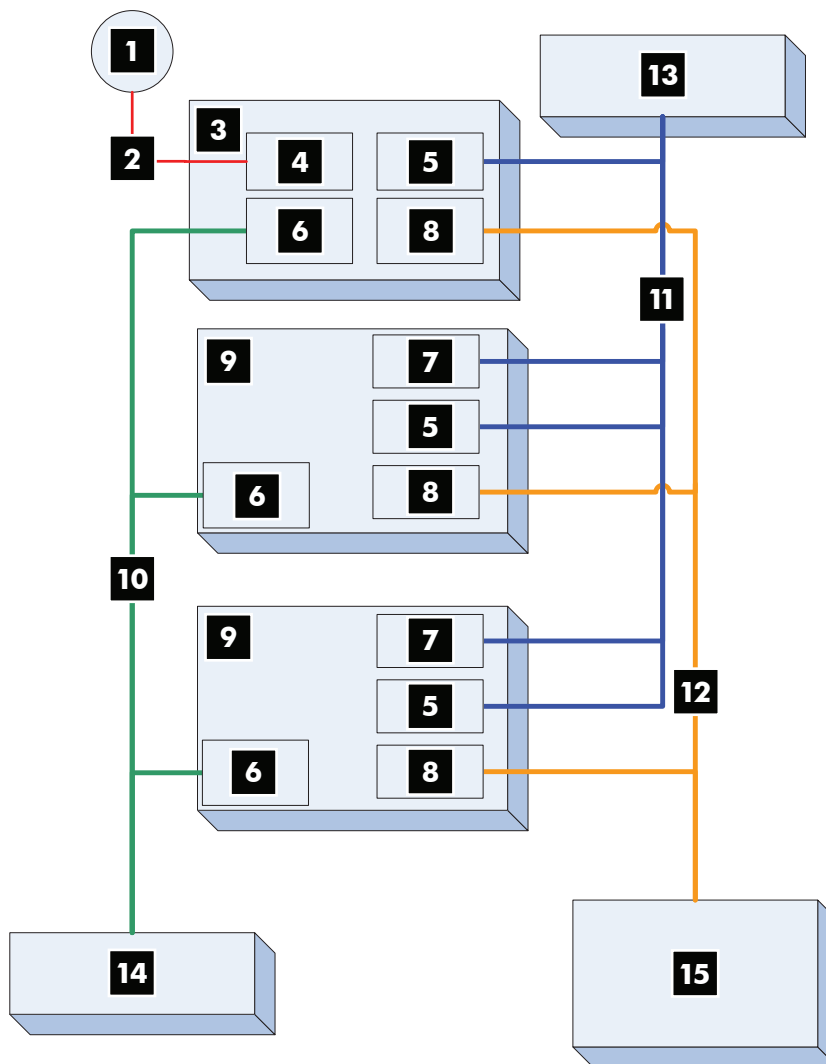


The following cluster features are shown in Figure 2-1:

- 1 Site local area network (LAN) connection.
- 2 LAN network, connected to a network interface card (NIC) installed (or embedded) in the cluster's head node.
- 3 The server designated as the head node (control node in HP Cluster Platform terminology).
- 4 NIC dedicated to the LAN network.
- 5 Connection to the server's management interface, a dedicated hardware interface that listens for, and executes commands received via the Ethernet network. This card is also known as the management processor (MP), or integrated lights-out (iLO) card, depending on the model of server. Where supported by the operating environment, this connection is used for server hardware management functions, such as power-off or boot.
- 6 NIC dedicated to the Message Passing Interface (MPI).
- 7 Compute nodes.
- 8 Common Gigabit Ethernet network, used both for MPI traffic, and for in-band job management, and cluster administrative functions.
- 9 Gigabit Ethernet switch, used as the system interconnect.

Figure 2-2 shows a topology based on a dedicated MPI fabric. This might be built around Gigabit Ethernet or InfiniBand. A second Ethernet network provides routing for cluster administrative, and job management traffic. A third network is provided via the site LAN connection.

Figure 2-2 Three-Network Topology



The following cluster features are shown in Figure 2-2:

- 1** Site local area network (LAN) connection.
- 2** LAN network, connected to a network interface card (NIC) installed (or embedded) in the cluster's head node.
- 3** The server designated as the head node (control node in HP Cluster Platform terminology).
- 4** NIC dedicated to the LAN network.
- 5** NIC connected to a dedicated Gigabit Ethernet MPI network, if present.
- 6** Connection to the server's management interface, a dedicated hardware interface that listens for, and executes, commands received via the Ethernet network. This card is also known as the management processor (MP), or integrated lights-out (ILO) card, depending on the model of server. Where supported by the operating environment, this connection is used for server hardware management functions, such as power-off or boot.
- 7** NIC connected to a shared Gigabit Ethernet MPI network, if present.
- 8** Host Channel Adapter (HCA) card, connected to an InfiniBand MPI network, if present.
- 9** Compute nodes.
- 10** Ethernet network for cluster administration.
- 11** Gigabit Ethernet MPI network, if present.
- 12** InfiniBand MPI network, if present.

- 13 Gigabit Ethernet system interconnect, if present.
- 14 Gigabit Ethernet switch for cluster administrative network, if present.
- 15 InfiniBand system interconnect, if present.

2.3 Understanding Maximum System Limits

In this release of Microsoft Windows Compute Cluster 2003, supported HP Cluster Platform configurations are constrained as follows:

- HP Cluster Platform Express is constrained to clusters ranging from 5 to 33 nodes.
- HP Cluster Platform is constrained to clusters ranging from 5 to 512 nodes for CP3000 and CP4000 models. Models employing blade and c-Class blade servers are constrained to 1024 nodes.
- InfiniBand interconnect single data rate (SDR) and double data rate (DDR).

2.4 Identifying Host Channel Adapters and Firmware

Every cluster node that is part of the high-speed InfiniBand fabric requires a host channel adapter (HCA) installed in its PCI bus. The HCA is always installed in a specific PCI slot on an unshared bus of the appropriate speed, ensuring system performance. For this release, the supported HCA models are specified in Table 2-1. Single-port InfiniBand cards are also supported.



Note:

For instructions on upgrading the HCA firmware on Microsoft Windows Compute Clusters, refer to the *Voltaire HCA 4X0 User Manual for Windows GridStack*.

Table 2-1 Supported HCA Cards

Model	PCI Bus Type	Description
380299-B21: InfiniBand 4X	PCI-X	Dual-port card with 128 MB attached memory
380298-B21: InfiniBand 4X	PCI-Express	Dual-port with 128 MB attached memory
409377-B21: InfiniBand 4X	PCI-Express	HP HPC 4X DDR IB Mezzanine HCA for HP BladeSystem c-Class

For information on supported firmware and stack revisions, see the *HP Cluster Platform InfiniBand Fabric Management and Diagnostic Guide*. You can download the latest information about supported firmware revisions from <http://www.docs.hp.com/en/highperfcomp.html>.

For mor information, see Chapter 4, which describes the GridStack™ InfiniBand stack.

2.5 Identifying the Servers Models Supported as Nodes

HP Cluster Platform hardware utilizes several CPU architectures and supports many models of HP server. Table Table 2-2 defines the servers that are supported in cluster platform under Microsoft Windows Compute Cluster 2003.

Table 2-2 Supported Servers

Model	CPU Type	Description
ProLiant DL140	Intel Xeon	1U rack-mount server
ProLiant DL145	AMD Opteron	1U rack-mount server
ProLiant DL360	Intel Xeon	1U rack-mount server
ProLiant DL365	AMD Opteron	1U rack-mount server

Table 2-2 Supported Servers (continued)

Model	CPU Type	Description
ProLiant DL380	Intel Xeon	2U rack-mount server
ProLiant DL385	AMD Opteron	2U rack-mount server
ProLiant DL585	AMD Opteron	4U rack-mount server
ProLiant BL20	Intel Xeon	Server blade
ProLiant BL25	AMD Opteron	Server blade
ProLiant BL30	Intel Xeon	Server blade
ProLiant BL35	AMD Opteron	Server blade
ProLiant BL460c	Intel Xeon	c-Class server blade
ProLiant BL465c	AMD Opteron	c-Class server blade
ProLiant BL480c	Intel Xeon	c-Class server blade
ProLiant BL685c	AMD Opteron	c-Class server blade

2.6 Upgrading or Modifying Clusters

Factory-integrated clusters are not designed for field upgrades using preinstalled servers, (otherwise known as Microsoft Windows Compute Cluster 2003 **standalone** servers). To find out whether an upgrade is available for your cluster, contact your HP sales and service representative, or authorized HP reseller.



CAUTION: Never add components to an HP Cluster Platform rack, even if available space exists in the rack. Your HP Cluster Platform is engineered to prevent thermal problems and to ensure proper power load balancing. Adding components, or otherwise modifying the rack might result in degraded performance.

3 Setting up your HP Cluster Platform for Microsoft Windows Compute Cluster 2003

To cable your cluster system, follow these steps:

1. Locate the appropriate cabling guides in your hardware documentation. HP Cluster Platform for Microsoft Windows Compute Cluster 2003 uses the following sets of cabling tables, available on the hardware documentation CD:
 - *HP Cluster Platform Cabling Guide for Gigabit Ethernet Solutions*
 - *HP Cluster Platform Infiniband Cabling Tables*
 - *Ethernet Network Cabling Tables*

Cables are labeled with the origin and destination ports. You need only reference the cabling tables to ensure that you understand the port naming syntax. In the event of missing or damaged cable labels (or if replacing damaged cables), use the cabling tables for your model and type of cluster to determine the correct point-to-point cable routing.

2. When cabling the cluster according to the *Ethernet Network Cabling Tables*, you might see cabling table footnotes that are operating-environment specific. Follow all cabling instructions that apply to the HP XC operating environment, which applies the same cabling rules as Microsoft Windows Compute Cluster 2003.
3. Clusters are delivered preconfigured. For single-rack solutions, you need only connect the cluster to your local area network, using the Ethernet switch port or NIC card designated in the cabling tables.

Go to [Step 4](#).

4. For clusters of more than one rack, you must do the following:
 - a. Complete the intrarack cabling, as described in the appropriate hardware documentation for InfiniBand or Gigabit Ethernet.



CAUTION: InfiniBand cables and interconnect ports are highly sensitive and prone to damage if not handled correctly. Before you begin to connect cables, read the *HP Cluster Platform InfiniBand Interconnect Installation and User's Guide*.

- b. Connect the cluster to your local area network, using the Ethernet switch port or NIC card designated in the cabling tables.
5. Perform a visual check of the fabric, as described in the cluster hardware documentation, referring to the user guides for the specific interconnect and Ethernet network switches used in your cluster. (This step involves only verifying that the link status and activity LEDs at each end of a cable link indicate that the link is good.)
 6. If not already installed, install the GridStack stack software. The installation procedure is described in [Chapter 4](#).
 7. Proceed to install or configure the Microsoft Windows Compute Cluster 2003 operating environment, as described in the operating environment documentation.



IMPORTANT: The operating environment configuration procedure requires that all cabling is complete and all hardware devices are in an appropriate state.

4 Verifying the InfiniBand Stack Software

Read this chapter only if your cluster employs an InfiniBand interconnect. The Voltaire GridStack stack is provided with the Microsoft Windows Compute Cluster 2003 operating environment in the supplied HP software kit on CD-ROM. Depending on how you order your configuration, the operating environment might be fully installed and partly configured. An overview of the installation procedure is provided in this chapter, if you need to install the stack on a replacement node. For the full procedure, and any current release notes, read the *Voltaire HCA 4X0 User Manual for Windows GridStack*.

4.1 Verifying the Stack Installation

To verify whether the stack is already installed, follow these steps:

1. Open the Windows **Control Panel** folder.
2. Click on the option to **Add or Remove Programs**.
3. A list titled **Currently Installed Programs** is displayed. Scroll down the list. If the GridStack software is installed, it is identified as *Voltaire WinIB64* or *Voltaire WinIB32* or a similar title.
4. Click on the option titled **Click here for support information** to determine the installed version.

4.2 Installing the Stack

If the stack is not installed, or after deinstalling the existing stack, follow these steps to install a new version while logged in system administrator (root user):

1. Ensure that the cluster configuration is complete, and you have verified the operation of HCA cards and links, as described in [Chapter 3](#). The operating environment must be installed and configured first.
2. Obtain the Microsoft software installer (MSI) package and copy it to a temporary location on the operating environment host disk. Potential sources for this package are:
 - In your Microsoft Windows Compute Cluster 2003 distribution kit.
 - From HP Global Services, as part of a recommended update.
 - By download from Voltaire.

A typical package name is: `Voltaire_WinIB64_2.0.363.201.msi`, and the actual file name corresponds to the architecture of the CPUs used in your cluster. The Voltaire .msi installer can also be found in `<CD drive volume>\HP-CCS\Voltaire` (for factory pre-installed operating systems) on the “HP Supplement for Microsoft Windows Computer Cluster Server 2003” CD.

3. Ensure that the MSI package version you receive is specified for your cluster hardware, as specified in either of the following sources:
 - *HP Cluster Platform InfiniBand Supported Firmware Matrix and Qualified Solution Tables*, available online and updated frequently at: <http://docs.hp.com/>, under the section titled *High Performance Computing*.
 - *HP Cluster Platform InfiniBand Fabric Management and Diagnostic Guide*.
4. Using the guides and tables specified in the preceding step, identify and record the correct version of HCA firmware to match the stack revision. (You might need this information later if the installed firmware is out of date.)
5. Navigate to the folder where you stored the MSI package, and double-click on its icon to begin the installation. An installation wizard leads you through the following steps:
 - a. Accept the software license.
 - b. Setup type: Choose the **Complete** installation.
 - c. Press the **Install** button to begin the installation.

6. Check for patches and updates to Microsoft Windows Server 2003. At the time of publication, the following patch is mandatory: <http://support.microsoft.com/?kbid=910481>. The KB910481 patch has already been installed with the factory-installed operating system. If not, you can also find this patch in <CD drive volume>:\HP-CCS\Hotfixes on the “HP Supplement for Microsoft Windows Computer Cluster Server 2003” CD, and it can also be automatically installed by using the provided Pre-CN.bat file.

4.3 Configuring the Stack

After you install the stack, follow these steps to configure it:

1. Internet Protocol over InfiniBand (IPoIB) addresses are already pre-configured by HP for factory-installed versions of the operating system. See the *HP Cluster Platform InfiniBand Fabric Management and Diagnostic Guide* for more information about the configuration interface.
2. Configure the **Windows Network Connection Order**, as described in the operating environment documentation, setting **Ethernet** to be the first priority.
3. Manually restart the GridStack as follows:
 - a. Open the **Control Panel** and choose **Administrative Tools**.
 - b. Expand the **Computer Management** options list and choose the **Device Manager**.
 - c. Under the **InfiniBand Host Channel Adapters**, Choose **InfiniBand HCA** and click the right mouse button to display the options menu.
 - d. From the options menu, choose **disable**.
 - e. Choose **InfiniBand HCA**, and click the right mouse button to display the options menu a second time. enable from the drop-down menu. The GridStack is now reset.
 - f. From the options menu, choose **enable**.
4. Determine the version of HCA firmware and, if necessary, upgrade it to a newer version using the procedure described in the *HP Cluster Platform InfiniBand Fabric Management and Diagnostic Guide*, which is on the Cluster Platform documentation CD.

5 HP Cluster Platform Terminology

Table 5-1 lists and defines generic terms specific to the HP Cluster Platform. Terms in bold in the Definition column are defined in the list. Where a concept is defined differently, or does not apply to clusters running Microsoft Windows Server 2003, and appropriate definition is included.

Table 5-1 HP Cluster Platform Terminology

HP Cluster Platform Term	Windows Compute Cluster Term	Definition
administrative network	Windows clusters do not distinguish a separate administrative network and console network	A cluster has two internal networks used for low-speed communication between nodes . The administrative network is used to manage the running cluster and for tasks such as job management.
AMD Opteron		The ProLiant DL series servers used in the CP4000cluster platform employ the AMD Opteron 32/64-bit CPU.
application node	compute node	A server that runs end-user programs. A bundle of application nodes requires one control node . Other than the high-speed network connection, an application node does not have significant I/O or storage capabilities.
bounded cluster		Any cluster with less than 128 nodes . Clusters with 128 nodes or fewer do not require their interconnects to be organized into a federated topology. Bounded clusters cannot easily be upgraded.
CBB		See compute building block .
chassis		See interconnect chassis.
cluster		A set of industry standard servers joined together using various interconnect technologies to form a supercomputer.
cluster configuration	Windows clusters do not support utility building blocks (UBB) or utility expansion racks (UXR)	The modular makeup of the cluster using CBBs , IBBs , UBBs , and a UXR . The configuration can be dense or modular.
cluster interconnect	The term switch might be used to describe this component	See interconnect .
cluster model		The HP Cluster Platform has three models based on the type of server used in the cluster . See CP3000 , CP4000 , and CP6000 .
component		Any discrete device that is used to build the cluster, such as a server, network switch or high-speed interconnect .
Compute Building Block (CBB)		A module containing a number of nodes (compute nodes and possibly also control nodes) management network switches, and other supporting components.
console network	Windows clusters do not distinguish a separate administrative network and console network	A cluster has two internal networks used for low-speed communication between nodes . The console network is used to access nodes for hardware management tasks, such as power on, shut down, and boot.
console NIC		See network interface card .
control node	head node	A server that provides management functions for the cluster . One control node exists in the cluster.
control processor		See management port .
CP3000		An HP Cluster Platform based on ProLiant DL series servers that employ the Intel Xeon 32-bit CPU.

Table 5-1 HP Cluster Platform Terminology *(continued)*

HP Cluster Platform Term	Windows Compute Cluster Term	Definition
CP4000		An HP Cluster Platform based on ProLiant DL-series servers that employ the AMD Opteron 32/64-bit CPU.
dense cluster		A preconfigured bounded cluster that cannot be expanded.
Ethernet switch		Devices that connect the nodes into administrative networks , such as the HP ProCurve Switch 2650. These networks facilitates management traffic between nodes and is used to perform functions such as booting.
external network		The customer's local area network (LAN) or wide area network (WAN), to which the cluster can be connected via its internal administrative network .
factory integration		A method of assembling cluster components into modules and temporarily connecting the modules to form a running, testable cluster. The cluster is partially disassembled into its modules for shipping.
fat tree		A network topology used by the interconnect . The topology is characterized by a link capacity that is progressively larger closest to the root of the tree than it is at the leaves, making the network nonblocking.
federated cluster		A cluster with more than 128 nodes , requiring that its interconnects are organized into a node-level and top-level hierarchy to increase the number of links . See bounded cluster .
flat tree		A network topology consisting of a single root device with several individual connections to multiple branch devices. The cluster's internal networks are organized as a flat tree.
Gigabit Ethernet	Windows clusters do not support Myrinet.	An interconnect protocol and connection type supported by the hardware that is based on HP ProCurve and Myricom Myrinet technologies.
high-speed network	message passing interface (MPI)	The fast data connection between nodes , using links through the interconnect . This term is used only to differentiate the interconnect network from the (low-speed) Ethernet networks.
HCA		See host channel adapter .
HP Cluster Platform		The generic product name for cluster models . This product name might be followed by a model designation, such as CP3000 .
host channel adapter		PCI interconnect card used in InfiniBand configurations
IBB		See Interconnect Building Block .
InfiniBand		An interconnect protocol and connection type supported by interconnect technology from Voltaire that complies with the specification defined by the InfiniBand trade group.
Intel Xeon		ProLiant DL series servers used in the CP3000 cluster platform use the Intel Xeon 32-bit CPU.
interconnect	Only InfiniBand and Gigabit Ethernet are supported.	A high-speed networking device that connects nodes together though an option card installed in the node's PCI bus. The interconnect is not referred to as a switch.
Interconnect Building Block (IBB)		A module containing one or more interconnects .
interconnect port		See port .

Table 5-1 HP Cluster Platform Terminology *(continued)*

HP Cluster Platform Term	Windows Compute Cluster Term	Definition
interconnect technology	Only InfiniBand and Gigabit Ethernet are supported.	The type of communications protocol and link implemented in an interconnect . See Quadrics , Myricom , Infiniband , and Gigabit Ethernet .
internal network	Windows clusters do not distinguish a separate administrative network and console network	A cluster has two internal networks used for low-speed communication between nodes . See administrative network and console network .
KVM		A rack-integrated workstation (keyboard, video, and mouse) enabling you to administer the cluster locally when necessary for tasks such as upgrading firmware and installing software. (You perform most run-time administration tasks by using a remote login from a client.)
KVM switch		An optional multiport (4 or 8) switch that enables you to switch the connection between the KVM and a utility node .
LAN		See external network .
link		A physical connection between a port on the interconnect and a port on an HBA that is located in the node's PCI bus.
local workstation		See KVM .
low-speed network	Windows clusters do not distinguish a separate administrative network and console network	The internal Ethernet networks are called “low speed” only to differentiate them from the interconnect's nonblocking high-speed network . This term is relative, because the Ethernet network switches used might be fast Gigabit Ethernet.
management network	Windows clusters do not distinguish a separate administrative network and console network	See administrative network .
management port		A connection or device that is part of the interconnect , enabling you to connect to it and administer it. You use the control processor for tasks such as assigning TCP/IP addresses, loading firmware or running interconnect diagnostics.
module (*BB)	Windows clusters do not support utility building blocks (UBB) or utility expansion racks (UXR)	A number of components of similar function connected into a functional unit and installed in a single rack. A module can be a CBB , IBB , UBB , or UXR .
modular cluster	Windows clusters do not support utility building blocks (UBB) or utility expansion racks (UXR)	A configuration designed to customer specifications by using CBBs , IBBs , UBBs , and a single UXR . Modular clusters can be upgraded (expanded).
network interface card		Each node has at least two integrated network interface cards (NICs). One NIC might be dedicated to the server's console. These NICs are used for connections to the internal networks .
network topology	HP Cluster Platform supports only a subset of the topologies supported by Windows Clusters.	The architecture of the cluster networks. In HP Cluster Platform there can be two low-speed networks and one high-speed network . The connections between any node and these networks determine their function.
NIC		See network interface card .
node	Head node and compute node. (Utility nodes are not supported).	An individual server, which can be a control node \, an application node or a utility node .

Table 5-1 HP Cluster Platform Terminology (*continued*)

HP Cluster Platform Term	Windows Compute Cluster Term	Definition
node-level interconnect (NLI)		In federated clusters, an interconnect that has connections to nodes . See also top-level interconnect (TLI) .
PCI card		See host channel adapter (HCA) .
PCI-X		Peripheral Component Interconnect Extended.
pool	Windows clusters support only pools 1 through 4	A related group of nodes in one or more CBB that share identical attributes. A cluster can have several pools of nodes, providing that all nodes in a pool (or the CBBs comprising a pool) are configured identically.
port		A physical outlet or connector. There are multiple ports on the switch cards installed in the interconnect and one or two ports on the HCA installed in the nodes .
ProCurve switches		The HP ProCurve networking switches are used in the internal Ethernet networks.
ProCurve interconnect		The HP ProCurve networking switches are used as Gigabit Ethernet interconnects in some HP Cluster Platform models.
ProLiant DL series	Windows clusters supports a subset of servers supported by HP Cluster Platform	A server type used as a node in the CP3000 and CP4000 platforms. These servers employ both the Intel Xeon 32-bit CPU and AMD Opteron 32/64-bit CPU.
server	Windows clusters supports a subset of servers supported by HP Cluster Platform	See node .
top-level interconnect (TLI)		In federated clusters , an interconnect that provides links to other interconnects, multiplying the number of possible pathways between nodes to create a full-bandwidth fat tree topology. See node-level interconnect (NLI) .
topology		The organization of a network. HP Cluster Platforms contain networks based on flat tree and fat tree topologies.
Utility Building Block (UBB)	Windows clusters do not support utility nodes	A module containing one or more SAN controllers, optional JBOD arrays, a terminal, a keyboard, and other supporting components. For dense clusters , this building block might also contain an interconnect .
Utility Expansion Rack (UXR)	Windows clusters do not support utility nodes	A module containing only utility nodes .
utility node	Windows clusters do not support utility nodes	A server that provides administrative and I/O functions.

Index

A

- about this document, 9
- adapter, PCI, 16
- administrative network, 13
- AMD Opteron, 16
- audience, 9

B

- blade server, 17

C

- c-Class blade server, 17
- cable
 - labels, 19
- cabling rules, 19
- cabling tables, 19
- CD-ROM, 10
- CD-ROM reader, 13
- cluster
 - terminology specific to, 23
- comments, reader, 10
- compute node, 14, 15
- configuratiob
 - supported servers, 13
- configuration, 9
 - constraints, 13
 - interconnects supported, 13
 - pools, 13
 - supported, 13
 - understanding, 11
- console network, 13
- constraints on configuring, 13
- conventions
 - typographic, 9
- CPU, 16

D

- DDR, 16
- dedicated fabric, 14
- document
 - InfiniBand, 10
 - organization, 9
 - reader comments, 10
 - related information, 10
- documentation
 - cabling tables, 19
 - CD-ROM, 10
 - HP cluster platform, 10
 - online, 10
 - stack, 21
 - Windows Compute Cluster 2003, 10
- double data rate, 16

E

- Ethernet

- cabling tables, 19

F

- fabric
 - dedicated, 14
 - dedicated Ethernet , 15
 - dedicated InfiniBand , 15
 - shared, 14
 - shared Ethernet , 15
 - verifying hardware, 19
- feedback, 10
- firmware, 16
 - identifying, 21
 - matrix, 21
 - support matrix, 16
 - upgrading, 22

G

- gigabit Ethernet, 13, 15
 - cabling tables, 19
 - switch, 14
- GridStack, 16
 - configuring, 22
 - installing, 21
 - restarting, 22
 - software kit, 21
 - verifuiong, 21

H

- HCA, 16
 - restarting the stack, 22
 - supported, 16
 - upgrading the firmware, 22
- head node, 13, 14, 15
- host channel adapter, 16
- HP Cluster Platform
 - documentation, 10
 - Web links, 10
 - Web site, 10

I

- ILO, 14, 15
- in-band, 13
- InfiniBand, 13, 15
 - cabling tables, 19
 - firmware, 16
 - host channel adapter, 16
 - management documents, 10
 - single data rate, 16
 - stack, 16
- information, related, 10
- installation, 19
 - documentation, 19
 - GridStack software, 19
 - patching, 22
 - stack, 21

- verifying fabric, 19
 - verifying stack, 21
- inteconnect
 - gigabit Ethernet , 16
 - InfiniBand, 16
- Intel Xeon, 16
- intended audience, 9
- interconnect, 13
 - gigabit Ethernet, 13
 - InfiniBand, 13
- IPoIB, 22

L

- labels, 19
- legal notices and trademarks, 2
- local network, 13, 14, 15
- log-in
 - Web site, 10

M

- mangement networ, 14, 15
- maximum system limits, 16
- Microsoft software installer, 21
- MP, 14, 15
- MPI, 13
 - dedicated Ethernet fabric , 15
 - dedicated fabric, 14
 - dedicated InfiniBand fabric , 15
 - gigabit Ethernet, 14, 15
 - InfiniBand, 15
 - NIC card, 14
 - shared gigabit Ethernet, 15
- MSI, 21

N

- network
 - administrative, 13
 - console, 13
 - in-band, 13
 - local, 14, 15
 - management, 14, 15
- network connection order, 22
- network topology
 - three networks, 15
 - two networks, 14
- NIC, 14
- node
 - compute, 14, 15
 - head, 13, 14, 15
 - standalone servers as, 17
- nodes
 - maximum, 16
 - ProLiant, 17
 - servers used as, 16

O

- online documentation, 10
- Opteron, 16
- organization of this book, 9

P

- patches, 22
- PCI, 16
- pools, 13
- ProLiant
 - BL20, 17
 - BL25, 17
 - BL30, 17
 - BL35, 17
 - BL460c, 17
 - BL465c, 17
 - BL480c, 17
 - BL685c, 17
 - DL145, 16
 - DL360, 16
 - DL365, 16
 - DL380, 17
 - DL385, 17
 - DL585, 17
- ProLiant DL140
 - DL140, 16
- ProLiant servers, 17
- providing feedback, 10

R

- reader
 - comments, 10
 - Web site, 10
- related information, 10
- restarting the stack, 22

S

- SDR, 16
- server
 - blade, 17
 - c-Class blade, 17
- servers, 13
 - ProLiant, 17
 - standalone, 17
 - supported variants, 16
- setup, 9, 19
- shared fabric, 14
- site lan, 13
- stack, 9, 16
 - configuring, 22
 - documentation, 21
 - GridStack, 16
 - installing, 21
 - installing software, 19
 - restarting, 22
 - software kit, 21
 - verifying, 21
- standalone servers, 17
- support matrix, 16
- supported servers, 13
- switch, 13, 14
 - (*see also* interconnect)
- system interconnect, 13
- system limits, 16

T

- terminology, 9, 23
- topology, 13
 - three networks, 15
 - two networks, 14
- typographic conventions, 9

U

- understanding configuration , 11
- upgrades, 17
- user guides, 10

V

- Voltaire
 - Web site, 10

W

- Web links, 10
- Web site
 - HP Cluster Platform, 10
 - log-in, Voltaire, 10
 - reader feedback, 10
 - Voltaire, 10
- windows cluster
 - cabling rules, 19
 - configurations, 9
 - hardware installation , 19
 - implementation, 9
 - setup, 9
 - supported HCA, 16
 - supported topologies, 13
 - system limits, 16
 - terminology, 9
 - upgrading, 17
 - Windows network connection order,, 22
- Windows Compute Cluster 2003
 - configuration , 11
 - configuration constraints, 13
 - documentation, 10
 - supported configuration , 13
- wiondows cluster
 - configuring stack, 22
- wondows cluster
 - stack, 9

X

- Xeon, 16