

ClusterPack V2.5 Release Note



i n v e n t

Manufacturing Part Number: T1843-90014

E0507

U.S.A.

© Copyright 2002-2007 Hewlett-Packard Development Company, L.P.



* T 1 8 4 3 - 9 0 0 1 4 *

Legal Notices

The information in this document is subject to change without notice.

Hewlett-Packard makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be held liable for errors contained herein or direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Warranty. A copy of the specific warranty terms applicable to your Hewlett-Packard product and replacement parts can be obtained from your local Sales and Service Office.

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 manual and flexible disk(s), tape cartridge(s), CD(s), or DVD(s) supplied for this pack 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.

Copyright Notices. ©copyright 2002-2007 Hewlett-Packard Development Company, L.P. all rights reserved.

Reproduction, adaptation, or translation of this document without prior written permission is prohibited, except as allowed under the copyright laws.

©copyright 2002-2007 Regents of the University of California.

This software is based in part on the Fourth Berkeley Software Distribution under license from the Regents of the University of California.

Trademark Notices. UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Limited.

Itanium is a registered trademark of the Intel Corporation or its subsidiaries in the United States and other countries.

ClusterPack V2.5 Release Note

Announcement

This release note describes the release of ClusterPack V2.5.

NOTE ClusterPack V2.5 supports both HP-UX 11iv2 and HP-UX 11iv3. The cluster operating system should be the same for the management server and all Compute Nodes. Mixing HP-UX 11iv2 and HP-UX 11iv3 within the same cluster is not supported.

ClusterPack is a powerful clustering software technology that provides a centralized and comprehensive cluster management solution: cluster setup and installation, system administration workload management, and system inventory/consistency.

The core components of ClusterPack V2.5 are:

- ClusterPack Online Tutorial
- HP Cluster Management Utility Zone V2.5
- HP Application ReStart (AppRS) V2.5

What's in This Version

Benefits

Manageability is one of the key considerations in cluster technology. ClusterPack V2.5 provides several benefits to system management.

- Reduces time and effort spent on repetitive tasks
- Reduces overhead associated with maintaining system configurations, software updates, and user management
- Maintains similar or identical configurations of nodes (or groups of nodes) in a cluster
- Provides proactive problem detection
- Provides flexibility and extensibility

Platforms Supported

ClusterPack V2.5 is available on HP Integrity Servers running either HP-UX 11iv2 Technical Computing Operating Environment (TCOE) or HP-UX 11iv3 TCOE as the management server and compute servers. It supports industry-standard Ethernet and InfiniBand as the cluster interconnect. It also provides console management via the on-board management processor on HP Integrity Servers.

Features

New features for ClusterPack V2.5 include:

- New command line utility for cluster power management (`clpower`)
- New command line utility to manage system file customizations on Compute Nodes (`clsysfile`)
- Integration with HP system management tools
- Flexible configurations
- Utilities to preserve system file customizations during cloning
- Support for non-intrusive Compute Node addition and deletion
- Removal of key configuration restrictions

- Easy upgrade path from supported ClusterPack versions.
- Improvements to the online tutorial

New command line utility for cluster power management (c1power)

This utility performs multiple power operation on Compute Nodes using the Management Processor interface. System administrators will be able to use this utility to turn the power on and off, turn the indicator lights on and off, and inquire about the power status of Compute Nodes.

New command line utility to manage system file customizations on Compute Nodes (c1sysfile)

This utility creates and manages customizations to system files for installation on Compute Nodes. This allows system administrators to preserve customizations across Golden Image cloning operations.

Integration with HP system management tools

ClusterPack configures and performs basic integration with HP Systems Insight Manager (HP SIM) if it is installed on the Management Server.

Flexible configurations

The following services are now optional:

- NIS
- Integration with other HP system management tools
- Secondary DNS servers

Utilities to preserve system file customizations during cloning

This utility packages files for installation on Compute Nodes during a cloning operation. Files are created on the Management Server, and packaged into an SD bundle. The resulting file bundle can be associated with a Golden Image. During a cloning operation, the file bundle will be automatically installed on the Compute Nodes. Multiple file bundles can be maintained, for use with multiple Golden Images. The file bundles can be manually installed on a given Compute Node. Removing the file bundle will restore the original files to the Compute Node.

Support for non-intrusive Compute Node addition and deletion

The Compute Nodes in the cluster that are not candidates for addition or removal will not be touched during these operations. Jobs can continue to run on those nodes while these operations are being performed.

Removal of key configuration restrictions

ClusterPack now has the ability to configure cluster networking with just the management LAN. It also now works in NIS and non-NIS environments.

Easy upgrade path from supported ClusterPack versions

ClusterPack V2.5 supports the upgrade for existing customers using supported versions of ClusterPack V2.3 or V2.4. ClusterPack V2.5 provides a toolset to retain the configuration settings from a ClusterPack V2.3 or ClusterPack V2.4 cluster.

See “Upgrading from V2.4 to V2.5” on page 36 or “Upgrading from V2.3 to V2.5” on page 38 for detailed instructions on how to upgrade from either V2.3 or V2.4 to V2.5.

Improvements to the online tutorial

The ClusterPack online tutorial now includes a QuickStart Installation Guide for experienced HP-UX system administrators. Comprehensive installation instructions are also provided with detailed directions for installation and for the use of Golden Images.

Known Problems and Workarounds

- ClusterPack V2.5 does not support secure shell ssh.
- Giving `clsh` an unknown node name will cause the application to terminate.
- The InfiniBand drivers are not active following a Golden Image installation of a Compute Node. The IB4X-00 driver bundle can be swcopy'd to the ClusterPack depot `/var/opt/clusterpack/depot`. If the InfiniBand drivers are in that depot, `compute_config` will offer the user an option to re-install the drivers on the Compute Nodes. Installation of the InfiniBand drivers will cause the Compute Nodes to reboot.
- If the Management Server IP address is changed, System Images will not install on Compute Nodes. Re-running `sysimage_register` for each image will correct the problem and allow the System Images to be installed on the Compute Nodes.

QuickStart Install Instructions

If you have installed ClusterPack before, follow the instructions in this section as a quick reminder. Refer to “Comprehensive Install Instructions” on page 13 for detailed instructions.

If you have not installed ClusterPack before, use the Comprehensive Install Instructions section instead. This QuickStart Guide does NOT cover the use of Golden Images. If you wish to use Golden Images, use the Comprehensive Install Instructions section instead.

IMPORTANT If you perform these steps out of order or omit steps, your installation will leave your systems in an unknown and non-deterministic state.

Step 1. Fill Out the ClusterPack Installation Worksheet

To help you collect the information needed for the installation, you may print a worksheet from `<DVD mount point>/CPACK-HELP/Tutorials/opt/clusterpack/share/help/ohs/docs/cpack_worksheet.pdf`. Fill out all the information for each node in your cluster.

NOTE You will not be able to complete the following steps if you have not collected all of this information.

Step 2. Install Prerequisites

The installation prerequisites for ClusterPack are important. You should read and understand all the prerequisites for ClusterPack installation before beginning the cluster setup.

ClusterPack v2.5 can be used with either HP-UX 11iv2 or HP-UX 11iv3. The version of the operating system on the cluster must be the same for the management server, and all the Compute Nodes.

For HP-UX 11iv2

Install the following software on the Management Server.

- HP-UX 11iv2 TCOE
- HP-UX 11i Ignite-UX (B5725AA)

Install the following software on each Compute Node.

- HP-UX 11iv2 TCOE
- HP-UX 11i Ignite-UX (B5725AA)

For HP-UX 11iv3

Install the following software on the Management Server.

- HP-UX 11iv3 TCOE
- HP-UX 11i Ignite-UX (IGNITE)

Install the following software on each Compute Node.

- HP-UX 11iv3 TCOE
- HP-UX 11i Ignite-UX (IGNITE)

Allow the default choices to install.

ClusterPack requires a homogeneous operating system environment. That is, all Compute Nodes and the Management Server must have the same release of HP-UX installed as well as the same operating environment.

The Management Server requires at least one LAN connection. The manager must be able to contact all the Compute Nodes using a "management network" that will be configured by ClusterPack. In addition, the management server must be able to connect to all the MP cards on the Compute Nodes. No network connections need to be configured before installing ClusterPack. The console interface can be used for all installation and configuration steps.

The Compute Nodes must have Management Processor (MP) cards.

ClusterPack depends on certain open source software which is normally installed as a part of the operating environment. The minimum release versions required are:

HP-UX 11iv2

- Perl Version 5.8 or higher

HP-UX 11iv3

- Perl Version 5.8 or higher

Step 3. Allocate File System Space

Allocate file system space on the Management Server. Minimum requirements are listed below.

HP-UX 11iv2

- /var - 4GB
- /opt - 4GB

HP-UX 11iv3

- /var - 8GB
- /opt - 4GB

Step 4. Obtain a License File

- Get the Host ID number of the Management Server.
- Contact Hewlett-Packard Licensing Services to redeem your license certificates.
- If you purchased the ClusterPack Base Edition, redeem the Base Edition license certificate.

NOTE It may take up to 24 hours to receive the license file. Plan accordingly.

Step 5. Prepare Hardware Access

Get a serial console cable long enough to reach all the Compute Nodes from the Management Server.

NOTE If you are installing ClusterPack on Compute Nodes for the first time, **DO NOT** power up the systems. ClusterPack will do that for you automatically. If you do accidentally power the compute nodes, **DO NOT** answer the HP-UX boot questions.

Step 6. Power Up the Management Server

Perform a normal boot process for the Management Server.

Step 7. Configure the ProCurve Switch

- Select an IP address from the same IP subnet that will be used for the Compute Nodes.
- Connect a console to the switch.

- Log onto the switch through the console.
- Type **set up**.
- Select **IP Config** and select the **manual** option.
- Select the IP address field and enter the IP address to be used for the switch.

Step 8. Copy the License Files to the Management Server

- Put the files in any convenient directory on the Management Server (e.g. /tmp).

Step 9. Install ClusterPack on the Management Server

- Mount and register the ClusterPack DVD as a software depot.
- Install the ClusterPack Manager software (CPACK-MGR) using `swinstall`.
- Leave the DVD in the DVD drive for the next step.

Step 10. Run `manager_config` on the Management Server

Provide the following information to the `manager_config` program:

- The path to the license file(s)
- The DNS domain and NIS domain for the cluster
- The host name of the manager and the name of the cluster
- The cluster LAN interface on the Management Server
- The IP address(es) of the Compute Node(s)
- Whether to mount a home directory
- Whether to configure HP SIM software

Step 11. Run `mp_register` on the Management Server

Provide the following information to the `mp_register` program about each Management card that is connected to a Compute Node:

- IP address
- Netmask
- Gateway IP address
- username and password to connect

Step 12. Power Up the Compute Nodes

Use the `clbootnodes` program to power up all Compute Nodes that have a connected Management Processor that you specified in the previous step. Provide the following information to the Compute Nodes:

- Language to use
- Host name
- Time and time zone settings
- Network configuration
- Root password

Step 13. Run `compute_config` on the Management Server

The `compute_config` program will register the nodes with various programs.

Step 14. Run `finalize_config` on the Management Server

This program completes the installation and configuration process, verifies the Cluster Management Software, and validates the installation. If it reports diagnostic error messages, repeat the installation process, performing all steps in the order specified.

Comprehensive Install Instructions

ClusterPack uses a two-stage process for setting up an HP-UX cluster.

1. Create a base configuration with a Management Server and one Compute Node.
 - a. Prepare for installation.
 - b. Install and configure the Management Server.
 - c. Install and configure the initial Compute Node and its Management Processor.
 - d. Verify the Management Server and the initial Compute Node.
2. Configure the remaining Compute Nodes with a Golden Image.
 - a. Create a Golden Image.
 - b. Add nodes to the configuration that will receive the Golden Image.
 - c. Distribute the Golden Image to the remaining nodes.
 - d. Install and configure the Compute Nodes that received the Golden Image.
 - e. Verify the final cluster configuration.

These processes are further broken down into a number of detailed steps. Each step contains the following sections:

- Background
- Overview
- Details

The Background section explains why this step is necessary and what will be done for you. The Overview section explains what this step entails in general terms. The Details section gives the exact commands that must be entered.

IMPORTANT The steps in this section must be followed in the specified order to ensure that everything works correctly. Please read all of the following steps **BEFORE** beginning the installation process.

Step 1. Fill Out the ClusterPack Installation Worksheet

Background

ClusterPack simplifies the creation and administration of a cluster of HP Integrity Servers running HP-UX by automating the collection, recording, and distribution of information about the systems in a network. The system administrator must still make decisions about how to identify and secure those network components. All of these decisions can be recorded on this form which is then used as the installation process is performed.

Overview

Print out this form and fill out all the information for each node in your cluster.

`<DVD mount point>/CPACK-HELP/Tutorials/opt/clusterpack/
share/help/ohs/docs/cpack_worksheet.pdf.`

NOTE	You will not be able to complete the following steps if you have not collected all of this information.
-------------	---

Details

At various points during the configuration you will be queried for the following information:

- DNS Domain name (e.g. domain.com)
- NIS Domain name (e.g. hpcluster) Optional
- Network Connectivity:
 - Information on which network cards in each Compute Node connect to the Management Server
 - Information on which network card in the Management Server connects to the Compute Node
- HP SIM Administrator password (You will be asked to set it).

Step 2. Install Prerequisites

Background

ClusterPack works on HP Integrity Servers running HP-UX. In order to install ClusterPack, the Technical Computing Operating Environment (TCOE) version of HP-UX must be installed. You must also have the Ignite-UX software, which is used for installation. Installing Ignite-UX on the Compute Nodes makes it possible to create and distribute 'Golden Images' from the Compute Nodes.

ClusterPack requires a homogeneous operating system environment. That is, all Compute Nodes and the Management Server must have the same release of HP-UX installed as well as the same operating environment.

Overview

HP-UX 11iv2

Install the following software on the Management Server and on one Compute Node:

- HP-UX 11iv2 or TCOE
- HP-UX 11i Ignite-UX (B5725AA)

ClusterPack depends on certain open source software which is normally installed as a part of the operating environment. The minimum release versions required are:

- Perl Version 5.8 or higher

HP-UX 11iv3

Install the following software on the Management Server and on one Compute Node:

- HP-UX 11iv3 or TCOE
- HP-UX 11i Ignite-UX (IGNITE)

ClusterPack depends on certain open source software which is normally installed as a part of the operating environment. The minimum release versions required are:

- Perl Version 5.8 or higher

The Management Server requires a minimum of two LAN connections. One connection must be configured prior to installing ClusterPack.

The Compute Nodes must have Management Processor (MP) cards.

Details

Install these items when you do a fresh install of HP-UX on the Management Server and the Compute Nodes. Or, Ignite-UX can be installed after rebooting using the following method:

- Using the HP-UX 11iv2 or HP-UX 11iv3 TCOE DVD, mount and register the DVD as a software depot.
- Install the Ignite-UX software on the Management Server using `swinstall`.

On the Management Server:

HP-UX 11iv2

```
% /usr/sbin/swinstall -s <source_machine>:/mnt/dvdrom \  
Ignite-UX
```

HP-UX 11iv3

```
% /usr/sbin/swinstall -s <source_machine>:/mnt/dvdrom \  
IGNITE
```

NOTE Allow the default choices to install.

Step 3. Allocate File System Space

Background

ClusterPack installs software in the `/opt` and `/share` file systems. It stores data in the `/var` file system. You must allocate sufficient space in these file systems for correct software operation.

Overview

Allocate file system space on the Management Server. Minimum requirements are listed below.

HP-UX 11iv2

- `/var` - 4GB
- `/opt` - 4GB

HP-UX 11iv3

- /var - 8GB
- /opt - 4GB

Details

Allocate space for these file systems when you do a fresh install of HP-UX on the Management Server.

To resize /opt

1. Go to single user mode.
 - a. **# /usr/sbin/shutdown -r now**
 - b. Interrupt auto boot.
 - c. Select the EFI shell.
 - d. Select the appropriate file system. (Should be fs0; but may be fs1:)
Shell> fs0:
 - e. Boot HP-UX.
fs0:\>hpux
 - f. Interrupt auto boot.
 - g. Boot to single user mode.
HPUX> boot vmunix -is
2. Determine the lvol of /opt.
 - a. **cat /etc/fstab** and look for the lvol that corresponds to /opt.
3. Extend the file system. (Use lvol from Step 2.)
 - a. **# lvextend -L 4096 /dev/vg00/lvol14** (May not be lvol4 or 4096.)
 - b. **# umount /dev/vg00/lvol14** (This should fail.)
 - c. **# extendfs /dev/vg00/lvol14**
 - d. **# mount /dev/vg00/lvol14**
4. Repeat 2 and 3 for /var.

Step 4. Obtain a License File

Background

For the ClusterPack Base Edition, please refer to the Base Edition license certificate for instructions on redeeming your license.

As part of the normal installation and configuration process, you will be asked to provide the license key(s). ClusterPack will install the license files in the correct location(s), and all licensing services will be started.

Overview

- Get the Host ID number of the Management Server.
- Contact Hewlett-Packard Licensing Services to redeem your license certificates.
- Redeem the Base Edition license certificate.

NOTE	It may take up to 24 hours to receive the license file. Plan accordingly.
-------------	---

Details

You will need to contact Hewlett-Packard Licensing Services to redeem your license certificates. You can call, e-mail, or fax your request to Hewlett-Packard Software Licensing Services. Refer to your Software License Certificate for contact information. Prior to installing ClusterPack V2.5, you can request a key by providing the Host ID number of the Management Server. The Host ID can be found using the `uname` command.

```
% /bin/uname -i
```

The number returned by this command must be preceded by a # when making your request. For example, if `uname -i` returns 2005771344, provide the number as #2005771344 in your key request.

Step 5. Prepare Hardware Access

Background

This document does not cover hardware details. It is necessary, however, to make certain hardware preparations in order to run the software.

Overview

Get a serial console cable long enough to reach all of the Compute Nodes from the Management Server.

Details

In order to allow the Management Server to aid in configuring the Management Processors, it is necessary to have a serial console cable to connect the serial port on the Management Server to the console port on the Management Processor that is to be configured. Be sure that the serial cable is long enough to reach all of the Compute Nodes. It is also possible to configure the Management Processors manually by connecting a console to each card.

NOTE If you are installing ClusterPack on Compute Nodes for the first time, **DO NOT** power up the systems, ClusterPack will do that for you automatically. If you do accidentally power the compute nodes, **DO NOT** answer the HP-UX boot questions.

Step 6. Power Up the Management Server

Background

This is the first step in actually configuring your system.

Overview

Perform a normal boot process for the Management Server.

NOTE **DO NOT** boot the Compute Nodes at this time.

Step 7. Configure the ProCurve Switch

Background

The ProCurve switch is used for the management network of the cluster.

Overview

The IP address for the ProCurve switch should be selected from the same IP subnet that will be used for the Compute Nodes.

Details

- Select an IP address from the same IP subnet that will be used for the Compute Nodes.
- Connect a console to the switch.
- Log on to the switch through the console.
- Type **set up**.
- Select **IP Config** and select the **manual** option.
- Select the IP address field and enter the IP address to be used for the switch.

Step 8. Copy the License Files to the Management Server

Background

Copy the license files to the Management Server. The license files can be placed in any convenient directory that is accessible to the Management Server. During the invocation of the `manager_config` tool, you will be asked to provide a path to the license files. As part of `manager_config`, the license files will be installed into the correct locations on the machine, and all licensing services will be started.

Overview

Put the files in any convenient directory on the Management Server.

Details

```
% /usr/bin/ftp your_host  
% > cd your_home  
% > lcd /tmp  
% > get cpack.lic  
% > bye
```

Step 9. Install ClusterPack on the Management Server

Background

The ClusterPack software is delivered on a DVD.

Overview

- Mount and register the ClusterPack DVD as a software depot.
- Install the ClusterPack Manager software (CPACK-MGR) using `swinstall`.
- Leave the DVD in the DVD drive for the next step.

Details

How to mount a DVD on a remote system to a local directory

On the system with the DVD drive (i.e. remote system)

1. Mount the DVD.

```
% mount /dev/dsk/xxx /mnt/dvdrom
```

2. Edit the `/etc/exports` file. DVDs must be mounted read only (`ro`) and, if required, can give root permission to other machines mounting the filesystem. (`root=<machine_foo:machine_bar:machine_baz>`) Add a line to `/etc/exports`:

```
% /mnt/dvdrom -ro,root=<local_system>
```

3. Export the file system using all the directives found in `/etc/exports`.

```
% exportfs -a
```

4. Verify that the line you added is actually exported.

```
% exportfs
```

On the local machine:

5. Mount the DVD to an existing directory.

```
% /etc/mount <remote_system>:/mnt/dvdrom /mnt/dvdrom
```

NOTE

You cannot be in the `/mnt/dvdrom` directory when you try to mount. You will get a file busy error.

6. Unmount the DVD file system.

```
% /etc/umount /mnt/dvdrom
```

On the remote system:

7. Unexport the DVD file system.

```
% exportfs -u -i /mnt/dvdrom
```

8. Unmount the DVD.

```
% /etc/umount /mnt/dvdrom
```

How to enable a DVD as a software depot

During the installation process, two DVDs will be required. Generic instructions for making a DVD accessible as a software depot for installation onto the Management Server are provided here. Please refer to the steps that follow for the specific DVDs that are required.

The steps to mount a DVD for use as a software depot are:

- Insert the DVD into the drive.
- Mount the DVD drive locally on that system.
- Register the depot on the DVD using `swreg`.
- Check the contents of the DVD using `swlist`.

These commands can only be executed as the superuser (i.e. root).

A DVD drive installed in the Management Server can be used for software installations. If the Management Server does not include a DVD drive, use one of the following two methods.

1. Connect a portable DVD drive to the Management Server.
2. Use an HP-UX system with a DVD drive that is network accessible from the Management Server, as a source for installation.

For example, to mount the device `/dev/dvdrom` to the directory `/mnt/dvdrom`, execute the following commands on the “source machine” with the DVD drive.

```
% /sbin/mount -r /dev/dsk/xxx /mnt/dvdrom
```

```
% /usr/sbin/swreg -l depot /mnt/dvdrom
```

```
% /usr/sbin/swlist @ /mnt/dvdrom
```

- Using the ClusterPack DVD, mount and register the DVD as a software depot.
- Install the ClusterPack Manager software (CPACK-MGR) on the Management Server using `swinstall`.

On the Management Server:

```
% /usr/sbin/swinstall -s <source_machine>: /mnt/dvdrom CPACK-MGR
```

- The ClusterPack DVD will be referenced again in the installation process. Please leave it in the DVD drive until the "Run manager_config on the Management Server" step has completed.

Step 10. Run manager_config on the Management Server

Background

This program is the main installation and configuration driver. It should be executed on the Management Server.

Some of the steps are:

- Install the appropriate license files and start the licensing services.
- Assign DNS domain name and NIS domain name based on inputs provided.
- Select and configure the cluster LAN interface on the Management Server that interfaces with the Compute Nodes.
- Specify how many Compute Nodes are in the cluster and the starting IP address of the first Compute Node. This information is used to assign names and IP addresses when the Compute Nodes are brought up. The first five characters of the Management Server's hostname are used for a base for the Compute Nodes. For example, if the starting IP address is 10.1.1.1, and there are 16 Compute Nodes, and the name of the Management Server is hpnod, then the first Compute Node will be called hpnod001 with the address 10.1.1.1. The next Compute Node will be called hpnod002 with the address 10.1.1.2, and so on. (Compute Node names are limited to eight characters.) If the tool is invoked with the `-f` option, the input file will be the source for this information.
- Set up the Management Server as NTP server, NIS server, NFS server, Ignite-UX server, and Web server.
- Install all of the dependent software components from the ClusterPack DVD:
 - This step looks for the source of the CPACK-MGR install and queries for an alternate source, if the source is not found. A local depot is setup. All of the agent components are copied. Other dependent software pieces in the Management Server are validated and installed.
- Modify configuration files on the Management Server to enable auto-startup of the Cluster Management Software components after reboots.

- Configure Cluster Management Software tools. The Management Server components of HP System Management Tools (HP Systems Insight Manager) is also configured if selected.
- Print a PASS diagnostic message if all of the configuration steps are successful.

Overview

Provide the following information to the `manager_config` program:

- The path to the license file(s)
- Whether to store passwords
- The DNS domain and NIS domain for the cluster
- The host name of the manager and the name of the cluster
- The cluster LAN interface on the Management Server
- The count and starting IP address of the Compute Nodes
- Whether to mount a home directory
- The HP-SIM admin password if HP-SIM is configured

Details

This tool can be invoked in two ways, based on your specific requirements.

- (Not recommended) If you want `manager_config` to drive the allocation of hostnames and IP addresses of the Compute Nodes in the cluster (based on some basic questions), `/opt/clusterpack/bin/manager_config` is invoked with no arguments.

```
% /opt/clusterpack/bin/manager_config
```

- If you want `manager_config` to assign specific hostnames and IP addresses to the Compute Nodes in the cluster, supply an input file in the same format as `/etc/hosts`, and invoke the tool as follows:

```
% /opt/clusterpack/bin/manager_config -f <input_file>
```

- The ClusterPack DVD is no longer required during installation. On the source machine, unmount the DVD drive and remove the DVD.

```
% /sbin/umount /mnt/dvdrrom
```

`manager_config` is an interactive tool that configures the Management Server based on some simple queries. Most of the queries have default values assigned and you just need to press **RETURN** to assign those default values.

Step 11. Run `mp_register` on the Management Server

Background

A Management Processor (MP) allows you to remotely monitor and control the state of a Compute Node. By configuring and registering the MP cards for each Compute Node, `clbootnodes` can be used to automatically answer the first boot questions for each Compute Node.

When you telnet to an MP, you will initially access the console of the associated server. Other options such as remote console access, power management, remote reboot operations, and temperature monitoring are available by typing **Ctrl-B** from the console mode. It is also possible to access the MP as a web console. However, before it is possible to access the MP remotely it is first necessary to assign an IP address to each MP. This is normally achieved by connecting a serial console device to the serial port on the MP and performing a series of configuration steps. This can be quite tedious and time consuming for moderate to large clusters. To ease the effort, `mp_register` can perform the configuration for you by issuing the commands via a serial cable.

`mp_register` maintains a database of knowledge about the MP cards in the system. The database is restricted to nodes that have been added to the cluster with `manager_config`. Likewise, nodes removed from the cluster are removed from the MP database. The utility is generally designed for single use when setting up the cluster for the first time. However, it can be run multiple times to make changes to MP designations or when nodes are added to the cluster.

NOTE

It is important to note that the configuration step does not configure accounts for the MP. By default, anyone can access the MP without a password. Leaving the cards without configuring users is a severe security risk. Users can freely access the card and shut down the node or gain root access through the console. The configuration step configures the MP for telnet or web access only to make future modifications, such as adding users, simpler to perform.

`mp_register` will add each MP and associated IP address to the `/etc/hosts` file on the Management Server. This file will later get propagated to the Compute Nodes. Each MP is assigned a name during the configuration step which is also placed in

the `/etc/hosts` file. This name is derived as the name of the associated host appended with `-mp` (for Management Processor). For example, the MP associated with the host `foo` will be named `foo-mp`.

Overview

Provide the following information to the `mp_register` program about each MP card that is connected to a Compute Node. It will configure all of the MPs automatically, instead of requiring you to manually connect the MP to a serial console device.

- IP address
- Netmask
- Gateway IP address

Details

For each node, the utility will ask you if you want to establish an MP for that machine. It will also ask if the MP is already configured. If it is not already configured, you will be prompted to connect a serial cable from the serial port of the Management Node to the serial port of the MP to be configured. The program will then use the information you entered about the card to configure it. Each MP can be configured in turn. MPs which have been previously configured can be added to the database without being configured.

Before invoking `mp_register` to initially configure the MP cards on each Compute Node, obtain a serial cable long enough to connect from the serial console port on the back of the Management Server to the serial console port on the MP card of each Compute Node.

When you are ready to run `mp_register`, use this command:

```
% /opt/clusterpack/bin/mp_register
```

Step 12. Power Up the Compute Nodes

Background

the `clbootnodes` utility is intended to ease the task of booting Compute Nodes for the first time. To use `clbootnodes`, the nodes' MP cards must have been registered and/or configured with `mp_register`

NOTE `clbootnodes` can only be used to boot nodes to the language specification: English

The first time that HP-UX is booted after installation, it asks a series of questions:

- What language to use
- Hostname
- Time and Time zone settings
- Networking configuration
- Root password

Booting each node in a medium to large cluster can be a long and tedious task. `clbootnodes` automates the processes to make it much faster and relatively free of user interaction. It is also possible to boot only specified nodes using `clbootnodes`.

`clbootnodes` will gain console access by using telnet to reach the MP. `clbootnodes` uses a library called Expect to produce the input needed to gain access to the console and step through the boot processes. There are times when manual intervention is necessary. In these cases, a message will be displayed explaining why control is being returned to the user. The user can then interact with the MP/console and then return control to `clbootnodes` by pressing '~'. Control may be given to the user for the following reasons:

- The MP is password protected.
- A LAN card choice was not specified to `clbootnodes`.
- The utility could not determine the state of the console.

`clbootnodes` is intended to boot a node or nodes through the first boot sequence. It can generally be run at any time to ensure that a node is booted and can usually recognize if the console represents a node that is already booted. However, because a user can leave the console in any state, it is not always possible to determine the state of a console. Because of this, it is recommended that `clbootnodes` be used for booting nodes which are known to be in a "first boot" condition.

When booting a node, `clbootnodes` will automatically answer the first boot questions. The questions are answered using the following information:

- Language selection: All language selection options are set to English.

- Keyboard selection: The keyboard selection is US English.
- Time Zone: The time zone information is determined based on the setting of the Management Server.
- Time: The current time is accepted. The time will later be synchronized to the Management Server using NTP.
- Networking: The LAN card specified will be configured to the IP address specified through `manager_config`.
- Hostname: The hostname will be set to the name specified through `manager_config`.
- Root password: The password will be queried before the nodes are booted.

Overview

Use the `clbootnodes` program to power up all Compute Nodes that have a connected MP that you specified in the previous step. It will answer the first boot questions for all the nodes automatically.

Provide the following information to the `clbootnodes` program:

- Language to use
- Hostname
- Time and time zone settings
- Network configuration
- Root password

Details

To run `clbootnodes`, use the following command:

```
% /opt/clusterpack/bin/clbootnodes
```

Before booting the nodes, `clbootnodes` will ask you for the root password to set on the Compute Nodes and the LAN card to configure for networking for each host. The LAN card choice for each host will be set to the IP address specified earlier via `manager_config`.

You can omit the argument list, in which all the nodes in the cluster will be processed. The IP address will be the one that you provided previously. The program will interact with you to obtain the name of the LAN card to use.

Step 13. Run `compute_config` on the Management Server

Background

This tool is the driver that installs and configures appropriate components on every Compute Node.

- Registers Compute Nodes with HP SIM on the Management Server
- Pushes agent components to all Compute Nodes
- Sets up each Compute Node as an NTP client, NIS client, and NFS client
- Starts necessary agents in each of the Compute Nodes
- Modifies configuration files on all Compute Nodes to enable auto-startup of agents after reboots
- Allows for the configuration of additional networks with `clnetworks`
- Prints a PASS diagnostic message if all configuration steps are successful

`clnetworks`

Each Compute Node is known to the Management Server through the IP address and specified to `manager_config`. These interfaces are collectively known as the Cluster Network. This term can be somewhat confusing when a cluster consists of both private nodes and public nodes. This is possible, for example, when an initial set of Compute Nodes is created on a private network and then additional nodes outside the private network are added using `-a`. The IP address of each Compute Node known by the Management Server makes up the Cluster Network.

ClusterPack includes a utility to configure additional networks on all of the Compute Nodes. These networks, like the Cluster Network, refer to a logical collection of interfaces/IP addresses and not to a physical network. However, they must share a common netmask. The concept of a network is defined as:

- A name (for reference only)
- A subset of the nodes in the cluster
- A network interface for each node in the subset
- An IP address for each interface
- A name extension that is added to the hostname of each machine and associated with each host's interface

- A netmask

To define additional networks, use the command `clnetworks`. This tool is also called from `compute_config`.

`clnetworks` provides a text-based interface for selecting nodes, network interfaces and IP addresses. It guides the user through the creation of a network. It is also possible to modify an existing network. When you have finished creating or updating networks, `clnetworks` will ensure that each interface specified is configured correctly and the proper entries exist in each host's `/etc/hosts` file.

Overview

The `compute_config` program will register the nodes with various programs.

Details

Execute the following command.

```
% /opt/clusterpack/bin/compute_config
```

Step 14. Set Up InfiniBand (Optional)

Background

ClusterPack configures IP over InfiniBand (IPoIB) if the appropriate InfiniBand drivers are installed on the systems.

ClusterPack provides a method to re-install the InfiniBand drivers on the Compute Node using `compute_config`.

Overview

If the InfiniBand IPoIB drivers are installed prior to running `compute_config`, the InfiniBand HCA is detected and the administrator is given a choice to configure them.

The administrator can also configure the InfiniBand HCA with IP addresses by invoking `/opt/clusterpack/bin/clnetworks`. See the man pages for `clnetworks` for usage instructions.

Known issues

There is a known issue that IB drivers are not correctly configured following a Golden Image installation of a Compute Node.

`compute_config` can be used to install IB drivers on Compute Nodes following a Golden Image installation. This re-installation of the drivers will allow them to work properly. To use the function, the IB driver bundle (i.e. IB4X-00) must be `swcopy`'d into `/var/opt/clusterpack/depot` on the Management Server:

```
% /usr/sbin/swcopy -x enforce_dependencies=false -s \  
<IB-driver-source> \* @ /var/opt/clusterpack/depot
```

At the end of `compute_config`, if the IB drivers are found in `/var/opt/clusterpack/depot`, an option to install the IB drivers on the Compute Nodes will be given. If you choose to install the IB drivers on the Compute Nodes, a second option will be presented. The IB drivers can be installed on only those Compute Nodes that already have the driver software installed, or the IB drivers can be installed on all the Compute Nodes.

Installing the IB drivers requires the Compute Nodes to reboot. This reboot is done automatically by `compute_config` as part of the installation.

NOTE If the IB drivers are installed on a Compute Node that does not have IB cards installed, the MPI test in `finalize_config` will fail.

Step 15. Run `finalize_config` on the Management Server

Background

This step performs verification checks on the Cluster Management Software, and validates the installation. It prints out diagnostic error messages if the installation is not successful.

NOTE The `finalize_config` tool can be run at any time to validate the cluster configuration and to determine if there are any errors in the ClusterPack software suite.

Overview

This program verifies the Cluster Management Software and validates the installation of the single Compute Node. If it reports diagnostic error messages, repeat the installation process up to this point, performing all the steps in the order specified.

Details

Finalize and validate the installation and configuration of the ClusterPack software.

```
% /opt/clusterpack/bin/finalize_config
```

Step 16. Create a Golden Image of a Compute Node from the Management Server

Background

A system image is an archive of a computer's file system. Capturing the file system of a computer captures the basic state of a computer system. An image does not generally include all files however. By default, `/tmp` and other temporary files, network directories, and host specific configuration files are not included.

A system image may be referred to as a Golden Image or a recovery image. The different names used to refer to the image reflect the different reasons for creating it. Administrators may create a "recovery" image of a node in the event that the node experiences hardware failure or the file system is accidentally removed or corrupted. Administrators may also create a "Golden" Image for the purpose of installing it on other nodes to insure that each node in their cluster is configured exactly the way they want.

Overview

- Clean up anything on the system that shouldn't be in the image.
- Ensure that the system isn't being used.
- Run `sysimage_create` to create the Golden Image.

Details

- Log on to the Compute Node to be archived.
- Perform general file system cleanup and maintenance. For example, it may be desirable to search for and remove core files.

From the Management Server:

- Ensure that the system is not being used. It is advisable that the system stop accepting new LSF jobs while the archive is being made.

```
% badmin hclose <hostname>
```

- In addition, you should either wait until all running jobs complete, or suspend them.

```
% bstop -a -u all -m <hostname>
```

- Execute `sysimage_create` on the Management Server and pass the name of the file from which you would like the image to be made. For example:

```
% /opt/clusterpack/bin/sysimage_create <hostname>
```

- Monitor the output for possible error conditions. The image will be stored in `/var/opt/ignite/archives/<hostname>`

```
% badmin hopen <hostname>
```

Step 17. Add Nodes to the Cluster That Will Receive the Golden Image

Background

This command adds the new node with the specified hostname and IP address to the cluster. It also reconfigures all of the components of ClusterPack to accommodate the newly added node.

Details

Invoke `/opt/clusterpack/bin/manager_config` with the “add node” option (`-a`). You can include multiple `host:ip` pairs if you need to.

```
% /opt/clusterpack/bin/manager_config -a <new_node_name>: \  
  <new_node_ip_addr>
```

Step 18. Distribute the Golden Image to the Remaining Compute Nodes

Background

This is the step that actually installs the Golden Image on the Compute Nodes.

Overview

- Register the image.
- Distribute the image to selected nodes.

Details

To distribute a Golden Image to a set of Compute Nodes, you need to first register the image. To register the image, use the command:

```
% /opt/clusterpack/bin/sysimage_register <full_path_of_image>
```

If the image was created with `sysimage_create`, the full path of the image was displayed by `sysimage_create`. Images are stored in the directory `/var/opt/ignite/archives/<hostname>`

To distribute the Golden Image to the Compute Nodes, use the command:

```
% /opt/clusterpack/bin/sysimage_distribute <full_path_of_image> \  
[hostname|all]
```

The keyword “all” can be used to distribute the image to all of the Compute Nodes in the cluster, or a single hostname can be specified. `sysimage_distribute` will reboot each Compute Node for installation with the specified image.

Step 19. Install and Configure the Remaining Compute Nodes

Background

This tool is the driver that installs and configures appropriate components on every Compute Node.

Overview

Perform this process in the same way as configuring the first Compute Node. Reference Step 13, "Run `compute_config` on the Management Server" for more information.

Details

Use the following command to install and configure a Compute Node that received the Golden Image. Perform this for all nodes. You can specify multiple nodes on the command line. You must place the option `-a` in front of each node name.

```
% /opt/clusterpack/bin/compute_config -a <node_name>
```

Step 20. Verify the Final Cluster Configuration

Background

This step completes the installation and configuration process, performs verification checks on the Cluster Management Software, and validates the installation. It prints out diagnostic error messages if the installation is not successful.

Overview

This program completes the installation and configuration process, verifies the Cluster Management Software, and validates the installation. If it reports diagnostic error messages, repeat the installation process, performing all the steps in the order specified.

Details

Finalize and validate the installation and configuration of the ClusterPack software.

```
% /opt/clusterpack/bin/finalize_config
```

Upgrading from V2.4 to V2.5

ClusterPack V2.5 supports an upgrade path from ClusterPack V2.4. Customers that currently deploy ClusterPack V2.4 on HP Integrity servers use HP-UX 11i Version 2.0 TCOE. ClusterPack V2.5 provides a mechanism for the use of the majority of V2.4 configuration settings for the V2.5 configuration.

Before starting the upgrade, it is important to have all of your Compute Nodes in good working order. All Compute Nodes and MP cards should be accessible. The LSF queues (if in use) should be empty of all jobs, and the nodes should be idle.

NOTE ClusterPack V2.5 does not integrate with Clusterware Pro. If you are using Clusterware Pro, the current setup and functionality will not be removed by ClusterPack V2.5. But, changes made to the cluster configuration (i.e. adding or removing Compute Nodes, creating groups, etc.) will **NOT** be updated automatically in the Clusterware Pro configuration.

Instructions for upgrading from V2.4 to V2.5:

Step 1. Backup the cluster user-level data.

Step 2. Install the V2.5 backup utilities.

```
% swinstall -s <depot_with_V2.5> CPACK-BACKUP
```

Step 3. Take a backup of the cluster information.

```
% /opt/clusterpack/bin/clbackup -f <backup_file_name>
```

Copy the backup file to another system for safe keeping.

Step 4. Install the new ClusterPack manager software.

```
% swinstall -s <depot_with_V2.5> CPACK-MGR
```

Step 5. Run `manager_config` in upgrade mode using the file you created in Step 3.

```
% /opt/clusterpack/bin/manager_config -u <backup_file_name>
```

Step 6. Register your MP cards. (To save time, check out the new `-f` option to `compute_config`.)

```
% /opt/clusterpack/bin/mp_register
```

Step 7. Install the new software on the Compute Nodes. (The `-u` is important.)

```
% /opt/clusterpack/bin/compute_config -u
```

Step 8. Verify that everything is working as expected.

```
% /opt/clusterpack/bin/finalize_config
```

Upgrading from V2.3 to V2.5

ClusterPack V2.5 supports an upgrade path from ClusterPack V2.3. Customers that currently deploy ClusterPack V2.3 on HP Integrity servers use HP-UX 11i Version 2.0 TCOE. ClusterPack V2.5 provides a mechanism for the use of the majority of V2.3 configuration settings for the V2.5 configuration.

Before starting the upgrade, it is important to have all of your Compute Nodes in good working order. All Compute Nodes and MP cards should be accessible. The LSF queues (if in use) should be empty of all jobs, and the nodes should be idle.

NOTE ClusterPack V2.5 does not integrate with Clusterware Pro. If you are using Clusterware Pro, the current setup and functionality will not be removed by ClusterPack V2.5. But, changes made to the cluster configuration (i.e. adding or removing Compute Nodes, creating groups, etc.) will **NOT** be updated automatically in the Clusterware Pro configuration.

Instructions for upgrading from V2.3 to V2.5:

Step 1. Backup the cluster user-level data.

Step 2. Install the V2.5 backup utilities.

```
% swinstall -s <depot_with_V2.5> CPACK-BACKUP
```

Step 3. Take a backup of the cluster information.

```
% /opt/clusterpack/bin/clbackup -f <backup_file_name>
```

Copy the backup file to another system for safe keeping.

Step 4. Install the new ClusterPack manager software.

```
% swinstall -s <depot_with_V2.5> CPACK-MGR
```

Step 5. Run `manager_config` in upgrade mode using the file you created in Step 3.

```
% /opt/clusterpack/bin/manager_config -u <backup_file_name>
```

Step 6. Register your MP cards. (To save time, check out the new `-f` option to `compute_config`.)

```
% /opt/clusterpack/bin/mp_register
```

Step 7. Install the new software on the Compute Nodes. (The `-u` is important.)

```
% /opt/clusterpack/bin/compute_config -u
```

Step 8. Verify that everything is working as expected.

```
% /opt/clusterpack/bin/finalize_config
```

Licensing

ClusterPack V2.5 uses FLEXlm licensing technology. A license is required before the product is installed. One ClusterPack license is required for each CPU in the cluster. Licenses can be redeemed through HP Software Licensing by phone, e-mail, or fax.

Please refer to the license certificate for instructions on redeeming your license.

The license keys are node-locked to the Management Server. You must provide the Host ID of the Management Server when requesting the license keys. The Host ID can be found using the `uname` command.

```
% /bin/uname -i
```

The number returned by this command must be preceded by a # when making your request. For example, If `uname -i` returns 2005771344, provide the Host ID number as #2005771344 in your key request.

Please allow up to 24 hours to receive the license files from HP Software Licensing.

Associated Documentation

You may review online documents from the ClusterPack V2.5 DVD by pointing your browser to **<DVD mount point>/CPACK-HELP/Tutorials/opt/clusterpack/share/help/ohs/index.html**.

- *ClusterPack Tutorial*
http://www.hp.com/techservers/clusterpack_tutorial/
or by pointing your browser to **<DVD mountpoint>/CPACK-HELP/Tutorials/opt/clusterpack/share/help/ohs/index.html**. The tutorial is also available after installation at [http://<management server>](http://<management_server>)
- *ClusterPack V2.5 Release Note*
<http://www.docs.hp.com>
- Additional ClusterPack information is available at
http://www.hp.com/techservers/clusters/hptc_clusterpack.html
- *HP-UX 11i Operating Environments*
<http://www.docs.hp.com/hpux/os/11i/index.html>
- *HP-UX 11i Version 2 Release Notes*
<http://www.docs.hp.com/hpux/onlinedocs/5990-6737/5990-6737.html>
- *HP-UX 11i Version 3 Release Notes* <http://www.docs.hp.com/en/5991-6469/index.html>
- *HP Application ReStart Release Note*
/opt/apprs/doc/releasenote.pdf
- *HP Application ReStart User's Guide*
/opt/apprs/doc/userguide.pdf
- *Software Distributor Administration Guide for HP-UX 11i Ed. 2*
<http://www.docs.hp.com/hpux/onlinedocs/B2355-90979/B2355-90979.html>
- *HP-UX IPFilter Release Note*
<http://www.docs.hp.com/hpux/onlinedocs/B9901-90020/B9901-90020.html>
- *Getting Started Guide HP Integrity rx2600 Server and HP Workstation zx6000*
<http://docs.hp.com/en/A9664-90020/A9664-90020.pdf>
- For information on the Management for the HP Integrity Server rx2600 refer to the Management section on page 31 of the *Getting Started Guide HP Integrity rx2600 Server and HP Workstation zx6000* at
<http://www.docs.hp.com/hpux/onlinedocs/support/A9664-90020/A9664-90020.pdf>

Associated Documentation

- Documentation for HP Integrity Servers is available at <http://docs.hp.com/hpux/hw/index.html>
- MPI documentation is available at <http://www.hp.com/go/mpi>

Software Availability in Native Languages

There is no information on non-English languages for ClusterPack V2.5.

Support Information

Support for ClusterPack V2.5 may be ordered, and service will be provided by the HP Response Center. Please refer to your support contract. Technical support is also available via <http://www.hp.com/techservers> or http://www.hp.com/techservers/clusters/hptc_clusterpack.html

All users can access Hewlett-Packard's Electronic Support Center on the World Wide Web where you can search for bug descriptions, updates, and available patches. The electronic support center is available at <http://us-support.external.hp.com>.