



SGI ISLE™ Cluster Manager Installation and Configuration Guide

COPYRIGHT

© 2008, 2009 SGI. All rights reserved; provided portions may be copyright in third parties, as indicated elsewhere herein. No permission is granted to copy, distribute, or create derivative works from the contents of this electronic documentation in any manner, in whole or in part, without the prior written permission of SGI.

LIMITED RIGHTS LEGEND

The software described in this document is “commercial computer software” provided with restricted rights (except as to included open/free source) as specified in the FAR 52.227-19 and/or the DFAR 227.7202, or successive sections. Use beyond license provisions is a violation of worldwide intellectual property laws, treaties and conventions. This document is provided with limited rights as defined in 52.227-14.

The electronic (software) version of this document was developed at private expense; if acquired under an agreement with the USA government or any contractor thereto, it is acquired as “commercial computer software” subject to the provisions of its applicable license agreement, as specified in (a) 48 CFR 12.212 of the FAR; or, if acquired for Department of Defense units, (b) 48 CFR 227-7202 of the DoD FAR Supplement; or sections succeeding thereto. Contractor/manufacturer is SGI, 1140 E. Arques Avenue, Sunnyvale, CA 94085.

TRADEMARKS AND ATTRIBUTIONS

SGI, the SGI logo, and Altix are registered trademarks and SGI ISLE is a trademark of SGI in the United States and/or other countries worldwide.

AMD and AMD Opteron are trademarks or registered trademarks of Advanced Micro Devices, Inc. Intel, Pentium, and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. Java is a registered trademark of Sun Microsystems, Inc. Linux is a registered trademark of Linus Torvalds, used with permission by SGI. Red Hat and all Red Hat-based trademarks are trademarks or registered trademarks of Red Hat, Inc. in the United States and other countries. SUSE LINUX and the SUSE logo are registered trademarks of Novell, Inc. UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company, Ltd. Windows and Windows XP are trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries.

All other trademarks mentioned herein are the property of their respective owners.

Table of Contents

Preface	iii
Audience	iii
Related Documentation	iii
Annotations	iv
Product Support	v
Reader Comments	v
Chapter 1	
Installing ISLE Cluster Manager	1
System Requirements	1
Minimum Hardware Requirements	1
Operating System Requirements	2
Software Requirements	2
Pre-installation Steps	3
Configure the IPMI BMC	3
Configure the Multicast Route	3
Set Up a Cluster Manager Master Host	4
Using the Cluster Manager Interface	7
Starting Cluster Manager	7
Starting and Stopping the Cluster Manager Server	9
Verifying Cluster Manager Services are Running	10
Chapter 2	
Configuring ISLE Cluster Manager	11
Configure DHCP	11
Configure DHCP Settings	11
Configure Network and Email Settings	12
Configure Users, Groups, and Roles	13
Default User Administration Settings	14
Add a User	15
Add a Group	16

Add a Role	17
Privileges	17
Chapter 3	
Set Up Your Cluster	19
Clustered Environments	19
Add a Host	19
Configure Platform Management	22
Import Hosts	25
Add Partitions	26
Create Regions	27
Chapter 4	
Creating Payloads and Images	29
Payload Management	29
Configuring a Payload Source	29
Create a Payload	31
Install Cluster Manager into the Payload	34
Kernel Management	35
Image Management	38
Create an Image	38
Managing Partitions	41
Provisioning	43
Select an Image and Provision	43
Index	47
Glossary	51

Preface

Audience

This guide is intended for system administrators who will install and configure the ISLE Cluster Manager software to manage and control the cluster system.

Related Documentation

The following documents provide additional information relevant to the ISLE Cluster Manager product:

- SGI ISLE Cluster Manager System Administrator's Guide
- SGI Altix XE310 User's Guide
- SGI Altix XE320 User's Guide
- IPMI Management Guide

Note

To access the latter two documents, contact your local sales representative. The following paragraphs describe the general access method for SGI customer documentation.

You can obtain SGI documentation, release notes, or man pages in the following ways:

- Refer to the SGI Technical Publications Library at <http://docs.sgi.com>. Various formats are available. This library contains the most recent and most comprehensive set of online books, release notes, man pages, and other information.
- You can also view man pages by typing `man <title>` on a command line.

SGI systems include a set of Linux man pages, formatted in the standard UNIX “man page” style. Important system configuration files and commands are documented on man pages. These are found online on the internal system disk (or DVD-ROM) and are displayed using the `man` command. For example, to display the man page for the `xscsidisktest` command, type the following on a command line:

```
man xscsidisktest
```

For additional information about displaying man pages using the `man` command, see `man(1)`.

In addition, the `apropos` command locates man pages based on keywords. For example, to display a list of man pages that describe disks, type the following on a command line:

```
apropos disk
```

For information about setting up and using `apropos`, see `apropos(1)`.

 **Note**

SUSE Linux documentation is available at:
<http://www.novell.com/documentation/suse.html>

RHEL documentation is available at:
<https://www.redhat.com/docs/manuals/enterprise/>

Annotations

This guide uses the following annotations throughout the text:

 **Electric Shock!**

Indicates impending danger. Ignoring these messages may result in serious injury or death.

 **Warning!**

Warns users about how to prevent equipment damage and avoid future problems.

 **Note**

Informs users of related information and provides details to enhance or clarify user activities.

 **Tip**

Identifies techniques or approaches that simplify a process or enhance performance.

Product Support

SGI provides a comprehensive product support and maintenance program for its products. SGI also offers services to implement and integrate Linux applications in your environment.

- Refer to <http://www.sgi.com/support/>
- If you are in North America, contact the Technical Assistance Center at +1 800 800 4SGI or contact your authorized service provider.
- If you are outside North America, contact the SGI subsidiary or authorized distributor in your country.

Reader Comments

If you have comments about the technical accuracy, content, or organization of this document, contact SGI. Be sure to include the title and document number of the manual with your comments. (Online, the document number is located in the front matter of the manual. In printed manuals, the document number is located at the bottom of each page.)

You can contact SGI in any of the following ways:

- Send e-mail to the following address: techpubs@sgi.com
- Contact your customer service representative and ask that an incident be filed in the SGI incident tracking system.
- Send mail to the following address:

SGI
Technical Publications
1140 East Arques Avenue
Sunnyvale, CA 94085-4602

SGI values your comments and will respond to them promptly.

Chapter 1

Installing ISLE Cluster Manager

To set up ISLE Cluster Manager in your environment, you must first install the Cluster Manager Server on a Master Host. After your Cluster Manager Server is installed, you can create images to distribute the Cluster Manager Client to the host nodes you want to manage. This lets you monitor and manage compute hosts from a central access point.

System Requirements

Before you attempt to install ISLE Cluster Manager, make sure your master host and compute hosts meet the following *minimum* hardware and software requirements:

Minimum Hardware Requirements

Master Hosts

- 2.2 GHz Intel Xeon or AMD Opteron (64-bit)
- 2 GB of RAM (4 GB or more recommended)
- 4 GB local disk space (minimum) — 50 GB or more is typically used
- 100 Mbps management network (including switches and interface card) — 1000 Mbps recommended

Compute Nodes

- 3.0 GHz Intel Pentium 4 (32-bit) or 2.2 GHz Intel Xeon or AMD Opteron (64-bit)
- 1 GB RAM
- 100 MB local disk typically used, diskless operation is also supported
- 100 Mbps management network (including switches and interface card) — 1000 Mbps recommended

Supported Platforms

- Intelligent Platform Management Interface (IPMI) 2.0 recommended for power control, serial access, and environmental monitoring
- IPMI 1.5, ILO 1.6 (or later), DRAC 3, and DRAC 4 offer power control only

Operating System Requirements

Warning!

Consult SGI before upgrading your Linux distribution or kernel. Upgrading to a distribution or kernel not supported on your system may render Cluster Manager inoperable or impair system functionality. Technical Support is not provided for unapproved system configurations.

Cluster Manager Server

You can run Cluster Manager Server on the following operating systems and architectures:

SUSE LINUX ENTERPRISE SERVER 10 (SP 1-2)

- 64-bit AMD64/EM64T hardware

RED HAT ENTERPRISE LINUX 5 (UPDATES 1-3)

- 64-bit AMD64/EM64T hardware

Cluster Manager Client

You can install and run the Cluster Manager Client on the same operating systems and platforms supported by the Cluster Manager Server as described above.

In addition, you can install the client software on Windows 2000 or Windows XP using 32-bit x86 Intel or AMD hardware.

Cluster Manager Kernel Support

SGI recommends using the kernel that shipped with your version of Linux. If you need to upgrade your kernel, please consult SGI before doing so.

Software Requirements

Cluster Manager requires the following RPM packages:

- Dynamic Host Configuration Protocol (DHCP) — included with your distribution
- Mkelimage — a program that makes an Executable and Linking Format (ELF) boot image for Linux kernel images (available on the Cluster Manager CD)
- Tftp-server (if you require PXE Boot support)
- IPMItool (if using IPMI hosts)

Note

IPMItool is required in the payload for each host controlled via IPMI. OpenIPMI is also required in the payloads for temperature and fan speed monitoring.

Pre-installation Steps

If you are going to use Intelligent Platform Management Interface (IPMI) hosts, you should configure your host's BIOS settings.

Configure the IPMI BMC

The BMC(s) for the nodes should be set up to use networking and serial over LAN. You will also need to know the username and password that will be used for power control and serial with the BMC(s) in order to use power control and serial over LAN with the Cluster Manager. The `ipmitool` utility allows you to set the username and password used to access the BMC on a host. This tool also allows you to set the LAN parameters of the BMC. For more information, consult the user's guide for the SGI Altix XE310 (or XE320), or third-party documentation (in the case of third-party node types).

Configure the Multicast Route

Note

When provisioning nodes, the default multicast configuration may not work properly. Ensure that multicast routing is configured to use the management interface.

Tip

The following examples use a multicast network of 224.0.0.0/4 to provide broad multicast support, but you can also use a more narrow multicast route such as 239.192.0.0/16. By default, the base multicast address in Cluster Manager is 239.192.0.128.

SLES

1. Enter the following from the CLI to temporarily add the route (where `eth0` is the management interface):

```
route add -net 224.0.0.0 netmask 240.0.0.0 dev eth0
```

2. Make the change persistent by entering the following:

```
vi /etc/sysconfig/network/routes
```

Then add:

```
224.0.0.0 0.0.0.0 240.0.0.0 eth0 multicast
```

RHEL

1. Enter the following from the CLI to temporarily add the route (where `eth0` is the management interface):

```
route add -net 224.0.0.0 netmask 240.0.0.0 dev eth0
```

2. Make the change persistent by entering the following:

```
vi /etc/sysconfig/network-scripts/route-eth0
```

Then add:

```
224.0.0.0/4 dev eth0
```

Set Up a Cluster Manager Master Host

After you have installed a Linux distribution and other required software on supported hardware, you are ready to install Cluster Manager Server. (See *Operating System Requirements* on page 2.) You can install the server via RPM(s) or the Cluster Manager installer.

Cluster Manager Server Installation Using RPMs

To install the server via RPMs, do the following:

1. Log on to the system as root.
2. Change directories to the `sgi/<arch>` directory of the Cluster Manager CDROM.
3. Type the following commands, replacing `<version>` and `<arch>` as needed:

```
rpm -ivh islecm-java-<version>.<arch>.rpm
rpm -ivh islecm-<version>.<arch>.rpm
rpm -iv islecm-server-<version>.<arch>.rpm
rpm -ivh lkSGI-<version>.<arch>.rpm
rpm -ivh lkSGI-java-<version>.<arch>.rpm
```

You will need to obtain a license from SGI. For information on the SGI LK licensing mechanism, see the *SGI Foundation 1 Service Pack 3 Start Here* available at <http://docs.sgi.com>. Open the `/etc/lk/keys.dat` file in a text editor. Copy and paste the license string, exactly as given, and save the file.

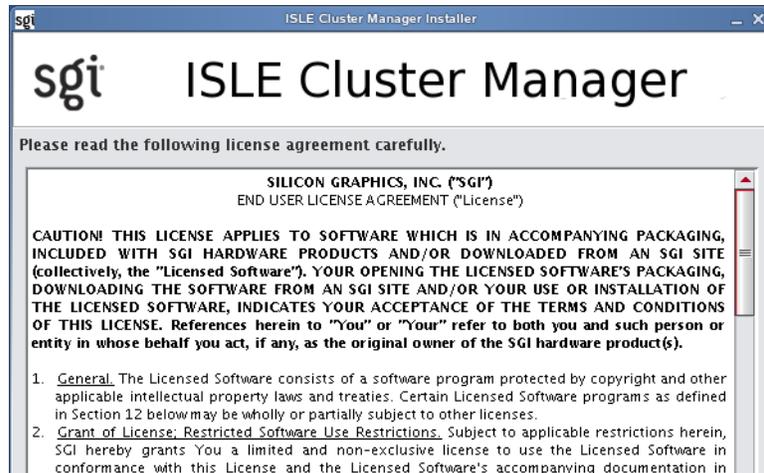
Cluster Manager Server Installation Using the Cluster Manager Installer

To install the server using the installer, do the following:

1. Generate or convert the `/etc/hosts` file for the servers and the compute hosts, then place this file on the Master Host.
The hosts file includes information such as IP address, MAC address, domain, and name of your master host and your compute nodes. If the information in the hosts file is incorrect, Cluster Manager will not communicate properly with your compute nodes.
2. Log in as the root user on the master host and insert the Cluster Manager CD.
If the installation does not begin automatically, enter the following in the command line:

```
run: <cdrom_mount_point>/install.sh
```

- Review and accept the terms of the license agreement and click **Next**.



- Click Cluster Manager Server and enter the installation directory and Master Host name.



 **Note**

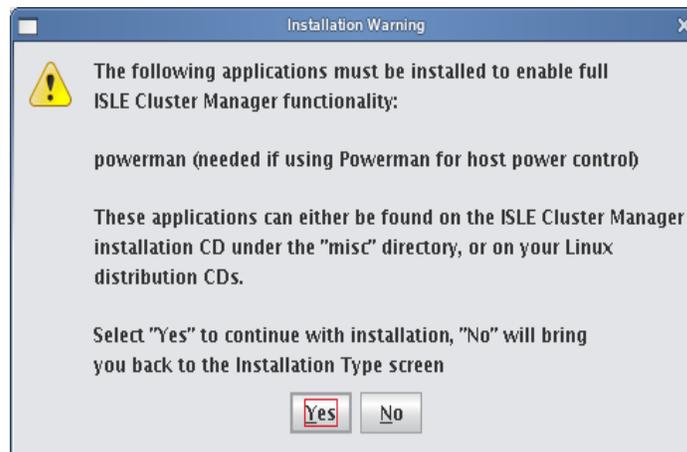
By default, the installation directory is `/opt/sgi/islecm`. The Master Host name (for example, `host`) must be present in the `/etc/hosts` file or you must be able to resolve it through DNS. This host name should resolve to the management interface, such as 10.0.10.1. If it resolves to the public interface, the Cluster Manager will not operate correctly.

! Warning!

To prevent accidentally overwriting Cluster Manager, install on a local file system that is not shared with other hosts.

5. Click **Next**.

If the computer does not have required applications installed, the following warning dialog is displayed.



You can continue the installation, but you will not have full Cluster Manager functionality until you install these required applications. (The applications can be installed after Cluster Manager has been installed.)

6. Click **Yes** to continue the installation.
7. Review the installation settings and click **Install**.

The installation process begins and the progress is displayed in the Cluster Manager Installer.

After the installation files have completed copying to the installation directory, you are prompted to install a Cluster Manager license.



The license is packaged as a text file — for example, *ClusterManager_1.0.txt*.

8. Click **Yes** and select the license file you want from the file browser.
9. After the installation completes, click **Finish**.
10. Restart your X session before starting Cluster Manager.

Using the Cluster Manager Interface

The Cluster Manager interface includes menus, a tool bar, tabbed panels, and frames with navigation trees that allow you to navigate and configure the cluster. From this interface you can add compute hosts and regions to the cluster and create payloads and kernels to provision the hosts.

Starting Cluster Manager

After you have installed the program and have restarted your X session, you can start the Cluster Manager interface from the command line interface.

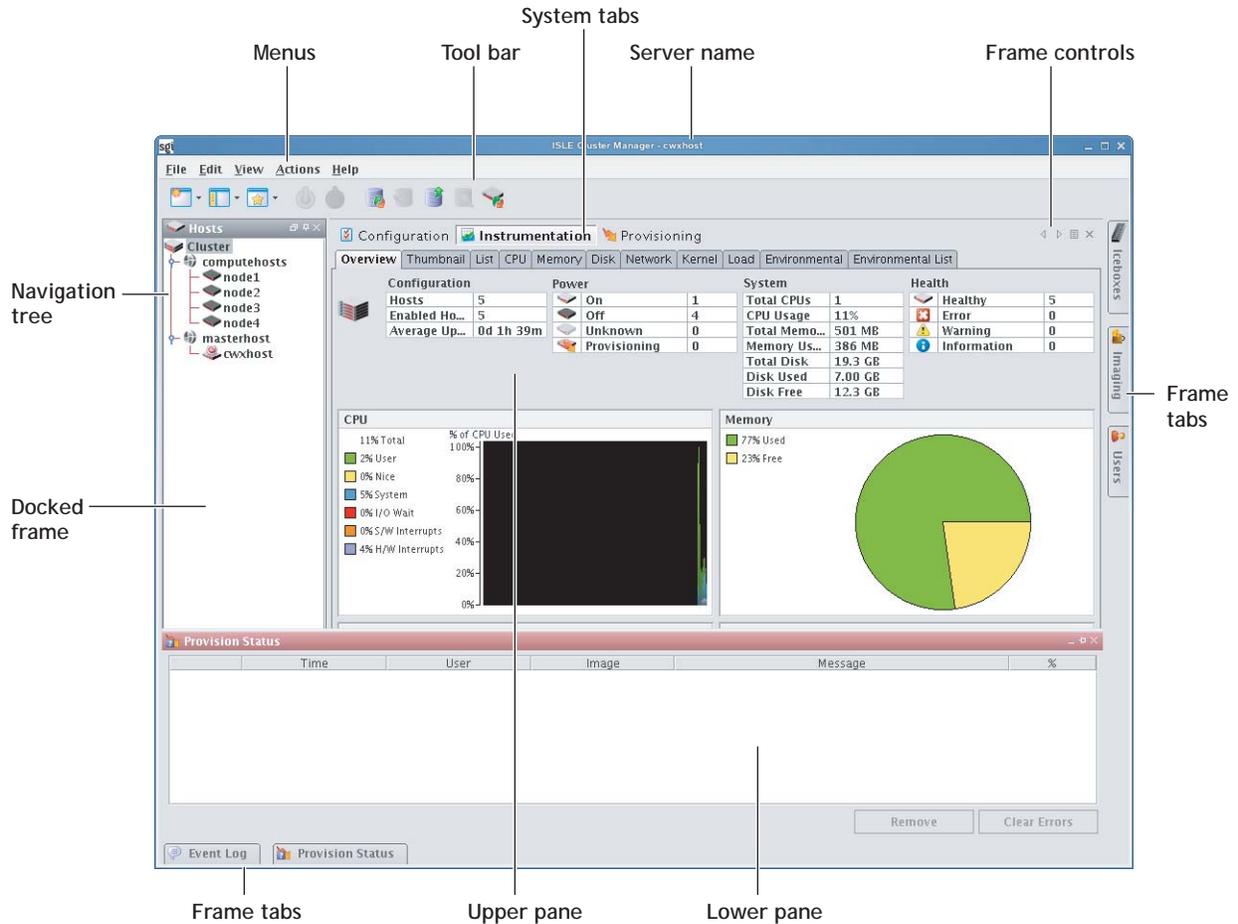
1. Open a command line console.
2. Log in as root.
3. On the command line, enter **cmclient** and press **Enter**.

The Cluster Manager Login is displayed.



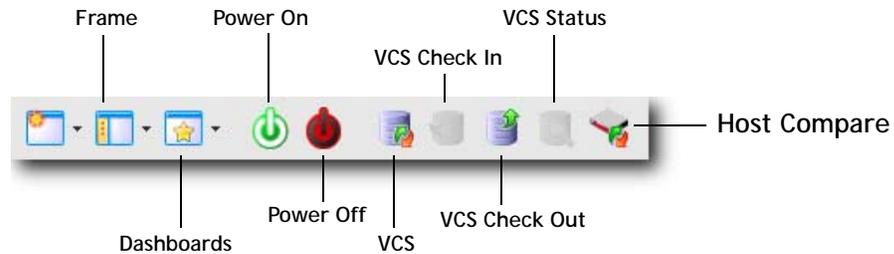
4. Enter a user name (root by default) and password (root by default) and click **OK**.

The Cluster Manager interface is displayed.



Menus — A collection of pull-down menus that provide access to system features and functionality.

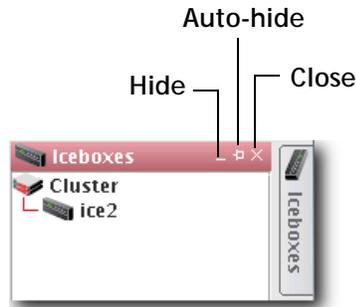
Tool Bar — The tool bar provides quick access to common tasks and features.



Server Name — The name of the server on which Cluster Manager is running.

System Tabs — Allow you to navigate and configure the cluster. Tabs may be opened, closed, and repositioned as needed.

Frame Controls — Lets you dock, un-dock, hide, minimize, and close frames.



Frames — Provide you with specific control over common aspects of cluster systems (for example, imaging and user accounts). Each frame tab opens a frame containing a navigation tree that allows you to manage system components easily. The navigation tree is found in most frames and is used to help organize cluster components. You may dock, close, or relocate frames and frame tabs as needed.

Upper/Lower Panes — These panes allow you to view cluster information in a structured environment.

Starting and Stopping the Cluster Manager Server

Cluster Manager services are started and stopped from scripts in */etc/init.d*. Cluster Manager is controlled by one of these services — this allows you to manage Cluster Manager services using standard Linux tools such as `chkconfig` and `service`. Standard functions for services include `start`, `stop`, `restart`, and `status`.

Example:

```
cd /etc/init.d; ./cmgr status
/etc/init.d/cmgr stop
/etc/init.d/cmgr start
chkconfig --list cmgr
```

Verifying Cluster Manager Services are Running

- Run the `/etc/init.d/cmgr status` command to verify that the following services are running:
 - AuthenticationService
 - DHCPService
 - DNA.<host IP address>
 - DatabaseService
 - DistributionService.provisioning-00
 - DistributionService.provisioning-01

Note

Cluster Manager includes two distribution services for each provisioning channel pair defined in the preferences.

- FileService.<host name>
- HostAdministrationService.<host name>
- IceboxAdministrationService
- ImageAdministrationService
- InstrumentationService
- KernelAdministrationService
- LicenseAdministrationService
- LogMonitoringService
- NotificationService
- PayloadAdministrationService
- PayloadNodeService.<hostname>
- PlatformManagementService
- PowerMonitoringService
- ProvisioningService
- RemoteProcessService.<hostname>
- RNA
- SynchronizationService
- TreeMonitoringService
- VersionService
- VersionService.<host_name>

Chapter 2

Configuring ISLE Cluster Manager

Configure DHCP

If you are using Dynamic Host Configuration Protocol (DHCP) you need to configure it on your master host to ensure proper communication with your compute nodes.

1. In a command line shell, log on as root.
2. Open `/etc/sysconfig/dhcpd`.
3. Look for the **DHCPD_INTERFACE** line and make sure it ends with `=ethx`.
4. Replace “x” with the host interface you use.

Example:

```
DHCPD_INTERFACE="eth0 eth1 eth2"
```

5. Save and close the file.

Configure DHCP Settings

When provisioning occurs, Cluster Manager automatically modifies DHCP settings and restarts the protocol. If you make manual DHCP modifications and want Cluster Manager to stop, start, restart, or reload DHCP, use the controls in the DHCP menu.

Note

When working with DHCP, ensure that the server installation includes DHCP and, if the subnet on which the cluster will run differs from 10.0.0.0, edit the Network sub-class field in the preferences dialog.

The DHCP option of the Actions menu allows you to perform the following operations:

- Stop the DHCP server.
- Start the DHCP server.
- Restart the DHCP server.
- Reload the `dhcpd.conf` file.

Tip

Changes made to `/etc/dhcpd.conf` are overwritten when you provision the host.

Configure Network and Email Settings

1. In the Cluster Manager interface, select **Preferences** from the **Edit** menu.
2. In the Preferences dialog, make sure the **General** button is selected.
3. In the **Email Settings** section, enter the sender, server, and domain information.
Use the email settings to send notifications of cluster events.
 - **Sender** — Used as the “From” address.
 - **Server** — Must be a valid SMTP server and must be configured to receive emails from the authorized domain.
 - **Domain** — The domain used to send email.
4. Configure the network settings.
The network settings must be configured before provisioning the cluster for the first time. The base network sub-class and netmask are mandatory. All other fields are optional.
 - **Base Network Sub-class** — The private network used by the cluster (typically a *192.168.x.x* or *10.x.x.x* network). To set the sub-class, the last octet should be a 0.
 - **Netmask** — The subnet mask used in your cluster.
 - **Log Server, NTP Server, Domain Name Server (DNS), and Default Gateway** — Used to set up DHCP settings. On a small to medium-sized system, these are typically the Master Host (by default, the log and NTP servers are set to use the Master Host). The DNS and default gateway are not set by default, but you should set them if you require all hosts to have external access to the cluster system.
5. Configure Preboot Execution Environment (PXE) Settings.
By default, Cluster Manager uses Etherboot as a scalable multicast boot client. When booting hosts that do not have an Etherboot client, you can use PXE as a boot loader to load Etherboot. Once Etherboot is loaded, the hosts will continue to boot or provision normally.

Note

To configure hosts to use a PXE loader, you must install and configure a TFTP server. After installing and configuring the TFTP server, open the preferences dialog from the Edit menu and browse to locate the `zpxe`-formatted PXE file — Cluster Manager copies the file to the TFTP server’s root directory (typically, `/tftpboot`). Commonly used `zpxe` files are available on the Cluster Manager ISO in the `/misc` directory. To obtain additional Etherboot `zpxe`-formatted files, visit Etherboot’s ROM-O-Matic website, at <http://www.rom-o-matic.net>. When downloading Etherboot files, select **zpxe** as the ROM type.

! Warning!

ROMs will not work with Multicast unless you add support for Scalable Local Area Multicast (SLAM). When downloading the Etherboot files from <http://www.rom-o-matic.net>, do the following:

1. Under the customize ROM configuration step, click **Configure**. You are redirected to the configuration page.
2. Select the **DOWNLOAD_PROTO_SLAM** option. This adds support for SLAM.
3. Click **Back** to return to the download page and finish downloading the Etherboot file.

To Set a Default PXE Boot File for All Hosts

- A. Click **Add** to open the Add PXE Boot Entry dialog.
- B. Select (**default**) from the drop-down list.
- C. Enter the path of the “zpxe” file to use by default or browse to locate the file.
- D. (Optional) Enable or disable the boot entry.
- E. Click **OK**.

To Set a PXE Boot File for a Specific Host

- A. Click **Add** to open the Add PXE Boot Entry dialog appears.
- B. Select one of the configured hosts from the drop-down list.
- C. Enter the path of the “zpxe” file to use by default or browse to locate the file.
- D. (Optional) Enable or disable the boot entry.
- E. Click **OK**.

6. Click **OK** to save the settings and close the Preferences dialog.

Configure Users, Groups, and Roles

Cluster Manager allows you to configure groups, users, roles, and privileges to establish a working environment on the cluster.

A group refers to an organization with shared or similar needs that is structured using specific roles (permissions and privileges) and region access that may be unique to the group or shared with other groups. Members of a group (users) inherit all rights and privileges defined for the groups to which they belong.

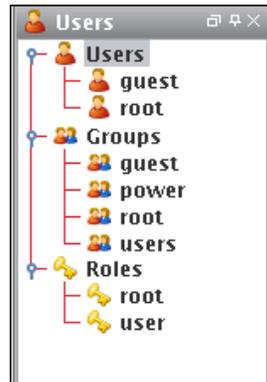
Note

Cluster Manager currently supports adding users and groups to payloads only — it does not support the management of local users and groups on the Master Host. Users with local Unix accounts do not automatically have Cluster Manager accounts, and this information cannot be imported into Cluster Manager.

If you are using local authentication in your payloads and intend to add Cluster Manager users or groups, make sure the user and group IDs (UIDs and GIDs, respectively) match between the accounts on the Master Host and Cluster Manager. Otherwise, NFS may not work properly.

Default User Administration Settings

Cluster Manager creates the following default users, groups, and roles during installation:



After installation, Cluster Manager allows you to create, modify, or delete groups, users, roles, and privileges as needed.

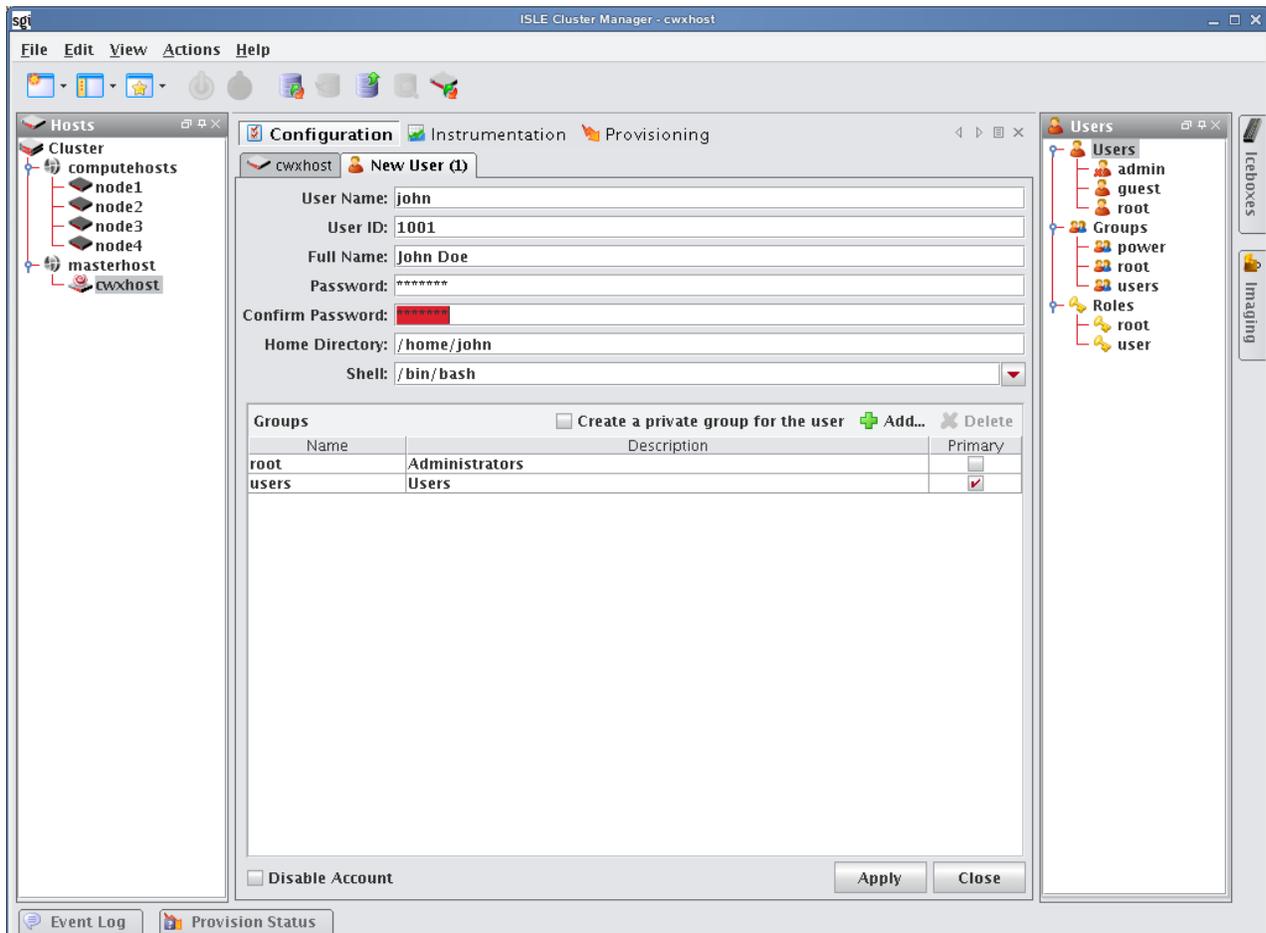
Note

You cannot remove the root user.

Add a User

Adding a user to Cluster Manager creates an account for the user and grants access to the system.

1. Select **New User** from the **File** menu or right-click in the user navigation tree and select **New User** to open a new user pane.



2. In the **User Name** field, enter a login name.
3. (Optional) Enter changes to the system-generated ID in the **User ID** field.
4. In the **Full Name** field, enter the user's first and last name.

Note

The Cluster Manager UID must match the system UID.

5. Enter and confirm a new user password.
6. (Optional) Specify the user's home directory in the **Home Directory** field (for example, `/home/username`).
7. (Optional) Enter a shell for this user or click the drop-down menu to select an existing shell. By default, Cluster Manager uses `/bin/bash`.
8. Click **Apply**.

Define a User Group

The Groups pane allows you to identify the groups to which the user belongs. Users are allowed to be part of any number of groups, but granting access to multiple groups may allow users unnecessary privileges to various parts of the system.

1. To add the user to a group, click **Add** in the **New User** pane.
2. Select the groups to which you want to assign the user (use the **Shift** or **Ctrl** keys to select multiple groups).

Note

Each user must belong to a primary group. If not, Cluster Manager automatically assigns the user to the “users” group. If you are using third-party power controls such as IPMI, the power group must be the primary group for *all* users who will use these controls.

3. Click **OK** to add the user to the groups.
4. (Optional) Click **Create a private group** for the user to create a new group with the same name as the user.
5. (Optional) Click **Disable Account** to prevent users from logging into this account.
Selecting this option also excludes this account from future payloads without requiring you to delete the account from Cluster Manager.

Add a Group

Adding a group creates a collection of users with shared or similar needs (such as an engineering, testing, or administrative group).

1. Select **New Group** from the **File** menu or right-click in the user navigation tree and select **New Group** to open a New Group pane.
2. Enter the group name.
3. (Optional) Enter changes to the system-generated group ID.
4. (Optional) Enter a description.
5. Click **Apply**.

ADD USERS

The Users pane allows you to identify the users that belong to the current group.

1. Click **Add**.
2. Select the users to add to the group.
3. Click **OK**.

ASSIGN ROLES

The Roles pane allows you to assign specific roles to the group.

1. Click **Add**.
2. Select the roles to assign to the group.
3. Click **OK**.

ASSIGN REGIONS

The Regions pane allows you to grant a group access to specific regions of the system.

1. Click **Add**.

2. Select the regions to assign to the group.
3. Click **OK**.

Add a Role

Adding a role to Cluster Manager allows you to define and grant system privileges to groups.

1. Select **New Role** from the **File** menu or right-click in the Users frame and select **New Role** to open a **New Role** pane.
2. Enter a role name.
3. (Optional) Enter a description.
4. Click **Apply**.

Note

Adding or revoking privileges will not affect users that are currently logged into Cluster Manager. Changes take effect only after the users close Cluster Manager and log in again.

ASSIGN GROUPS TO ROLES

The Groups pane allows you to assign a role to multiple groups. This permits users to have varied levels of access throughout the system.

1. Click **Add** in the Groups pane.
2. Select the groups you want to assign to the role.
3. Click **OK**.

GRANT PRIVILEGES

The Privileges pane allows you to assign permissions to a role. Any user with the role will have these permissions in the system.

1. Click **Add** in the Privileges pane.
2. Select the privileges you want to grant the current role.
3. Click **OK**.

Privileges

Privileges are permissions or rights that grant varying levels of access to system users. Cluster Manager allows you to assign privileges as part of a role, then assign the role to specific user groups. Users assigned to multiple groups have different roles and access in each group. This flexibility allows you to establish several types of roles you can assign to users such as, Full Administration, Group Administration, User, or Guest.

Privileges

The following table lists the privileges established for the Cluster Manager module at the function and sub-function levels:

Module	Name	Description
Cluster Manager	Database	The ability to execute database commands from the command line.
	Host	The ability to configure Hosts, Regions, and Partitions.
	Icebox	The ability to configure Iceboxes.
	Image	The ability to configure Images, Payloads, and Kernels.
	Instrumentation	The ability to monitor the system.
	Logging	The ability to view and clear error logs.
	Power	The ability to manage power to hosts.
	Provisioning	The ability to provision hosts.
	Serial	The ability to use the Serial over LAN terminal.
	User	The ability to configure Users, Groups, and Roles.

Chapter 3

Set Up Your Cluster

Clustered Environments

In a clustered environment, there is always at least one host that acts as the master of the remaining hosts (for large systems, multiple masters may be required). The master host is reserved exclusively for managing the cluster and is not typically available to perform tasks assigned to the remaining hosts.

To manage the remaining hosts in the cluster, you can divide them into partitions and regions.

Partitions include a strict set of hosts that may not be shared with other partitions.

Regions are a subset of a partition and may contain any hosts that belong to the same partition. Hosts contained within a partition may belong to a single region or may be shared with multiple regions. Dividing up the system can help simplify cluster management and allows you to enable different privileges on various parts of the system.

Add a Host

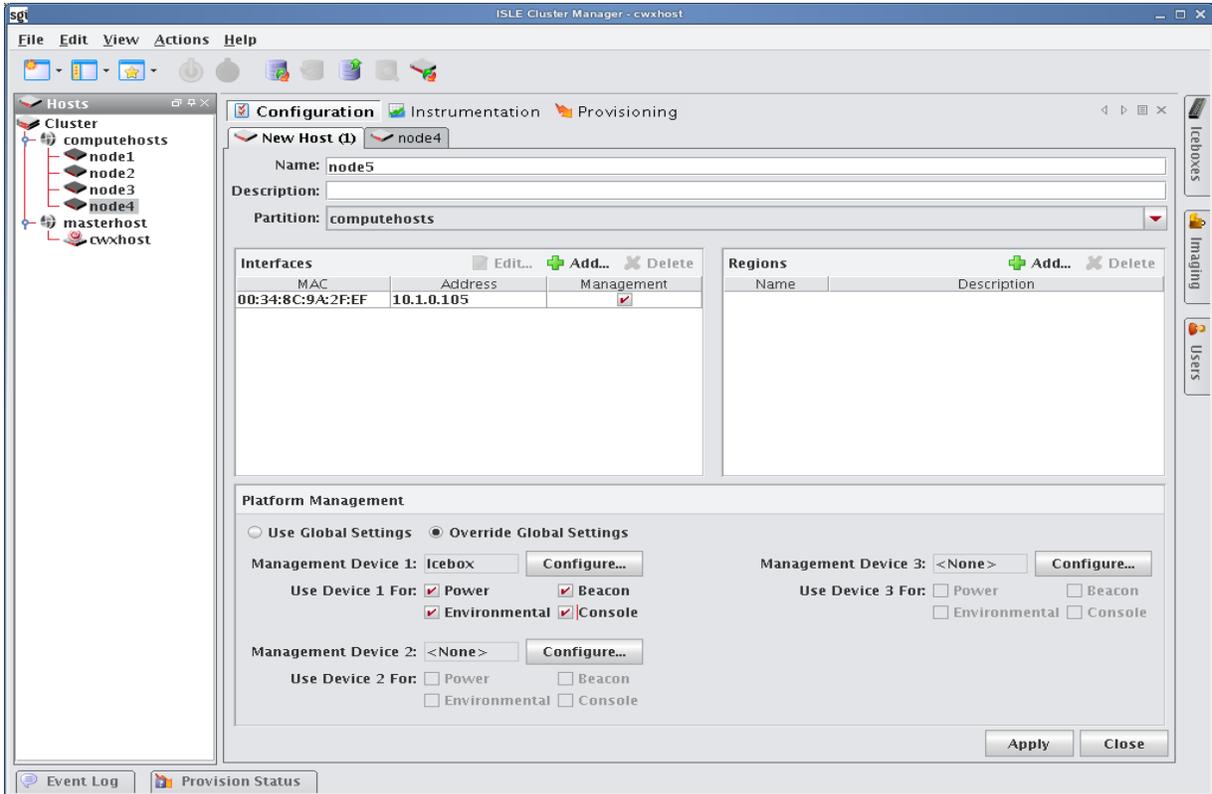
To add a host, you must provide the host name, description, MAC address, IP address, and the partition and region to which the host belongs. Hosts can be added only after you have set up a Master Host.



You can also import a list of existing hosts. See *Import Hosts* on page 25.

1. Select the **Cluster** icon in the **Hosts** frame.

2. Select **New Host** from the **File** menu or right-click in the host navigation tree and select **New Host**. A new host pane appears.



3. Enter the host name.
4. (Optional) Enter a description.
5. (Optional) Select the name of the partition to which this host belongs from the drop-down menu.

 Tip

If you right-click a partition of region in the navigation tree and select **New Host**, the host is automatically assigned to that partition or region.

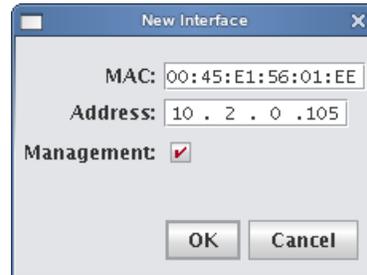
6. Create **Regions** and **Interfaces** assignments as needed.
7. Click **Apply** to create the new host.

Add Interfaces

The Interfaces pane allows you to create new interfaces and assign host management responsibilities.

1. In the Interfaces pane, click **Add**.

The New Interface dialog appears.



2. Enter the host's MAC and IP addresses.

Tip

To find the MAC address of a new, un-provisioned host, you must watch the output from the serial console. Etherboot displays the host's MAC address on the console when the host first boots. For example:

```
Etherboot 5.1.2rc5.eb7 (GPL) Tagged ELF64 ELF (Multiboot) for EEPROM100]
Relocating _text from: [000242d8,00034028) to [17fdc2b0,17fec000)
Boot from (N)etwork (D)isk (F)loppy or from (L)ocal?
Probing net...
Probing pci...Found EEPROM100 ROM address 0x0000
[EEPROM100]Ethernet addr: 00:02:B3:11:03:77
```

```
Searching for server (DHCP)...
```

(*If conman is set up and working, this information is also contained in the conman log file for the host—typically located in `/var/log/conman/console.n[1-x]`)

To find the MAC address on a host that is already running, enter `ifconfig -a` in the CLI and look for the HWaddr of the management interface.

3. Click **Management** to use the Cluster Manager interface to manage the host.
Cluster Manager stores the interface and automatically writes it to `dhcp.conf`.
4. Add any additional interfaces required for this host.
Cluster Manager records the interfaces and writes them to `dhcpd.conf`.

Note

If you are using IPMI or another third-party power controller, you should add the BMC's MAC address and the IP address you are going to assign it. Cluster Manager will set up DHCP to connect to the BMC. In the Platform Management settings, you can select this interface and use it for operations.

5. Click **OK**.

Assign Regions

The Regions pane allows you to identify any regions to which the host belongs.

1. (Optional) In the Regions pane, click **Add**.
The Select Regions dialog appears.
2. Select the region to which the host belongs. (To select multiple regions, use the **Shift** or **Ctrl** keys.)
3. Click **OK**.

Configure Platform Management

Platform management allows you to configure the power and temperature Management Devices you will use for each host.

Note

By default, platform management uses the device specified in your Global preferences settings to control hosts in the cluster. To override this setting, select Override Global Settings.

IPMI

Typically, hosts use one or more Ethernet interfaces. With IPMI, ILO, and DRAC, each host uses at least two interfaces: one management interface and one IPMI/ILO/DRAC interface. The management interface is configured for booting and provisioning, the IPMI/ILO/DRAC interface is used to gather environmental and sensor data (for example, fan speeds) from the host and perform power operations. Additional interfaces are used only for setting up host names and IP addresses.

Note

ILO and DRAC support power control only — they do not support temperature and sensor monitoring.

In order for Platform Management to work correctly, you must first define interfaces for each host (see *Add Interfaces* on page 21). In some cases, you must manually configure an IP address for the Platform Management Controller—in most cases, however, you can use DHCP to configure this address. To view information about each interface, see *dhcpcd.conf*.

The IPMI dialog defines which interface is used for Platform Management. Typically, the Management Device is easily identified because its MAC or IP address is an offset of the host. For example, a host with a MAC address of 00:11:22:33:44:56 and an IP address of 10.0.0.1 might have a Management Device with a MAC address 00:11:22:33:44:59 and set an IP address of 10.0.2.1. In this case, the MAC offset would be 000000000003 (Greater) and the IP offset would be 0.0.2.0 (Greater).

TO CONFIGURE IPMI SETTINGS ON YOUR HOST

1. Select **Override Global Settings**.
2. Select **IPMI** from the **Platform Management Device Type** drop-down list.

3. Select the Management Device IP Address Type:
 - A. **Dynamic** — Enter a hexadecimal MAC offset. For example, if you choose a Greater Than offset of 00:00:00:00:00:04 and the host's MAC address is 00:15:C5:EA:A7:7B, the MAC Address used for power operations will be 00:15:C5:EA:A7:7F (the sum of the original MAC address and the offset).
 - B. **Relative** — Choose an IP address offset and select whether it is Greater Than, the Same As, or Less Than the host's IP address. For example, if you choose a Greater Than offset of 0.0.1.0 and the host's IP address is 10.3.0.14, the host's BMC address will be 10.3.1.14. This is the IP address used for power operations (the sum of the original IP address and the offset).
 - C. **Static** — If you choose Static or if you wish to use different settings for each host, you must configure the IPMI options individually for each host.
4. (Optional) Select the MAC Address vs. Host MAC Address type:
 - A. Not Related
 - B. Greater Than
 - C. Less Than
5. (Optional) Enter the MAC Address Offset.
6. Select the MAC Address to use to manage this host.
7. (Optional) Select the IP Address vs. Host IP Address type:
 - A. Greater Than
 - B. Less Than
 - C. Same As
8. (Optional) Enter the IP address offset from the management interface.
9. (Optional) Enter the IP address for the host.
10. Select a Platform Management User.

 **Note**

Users must belong to Power as their primary group to appear in this list.

DRAC and ILO

1. Select **Override Global Settings**.
2. Select **DRAC** or **ILO** as the Platform Management Device Type.

Management Device 1 Settings

Platform Management Device Type: DRAC

Management Device IP Address Type: Relative

MAC Address vs Host MAC Address: Not Related

MAC Address Offset: : : : : :

MAC Address: [dropdown arrow]

IP Address vs. Host IP Address: Greater Than

IP Address Offset: 0 . 0 . 1 . 0

IP Address: 10 . 1 . 1 . 105

Platform Management User: admin

OK Cancel

3. Select the Management Device IP Address Type:
 - A. **Dynamic** — Enter a hexadecimal MAC offset. For example, if you choose a Greater Than offset of 00:00:00:00:00:04 and the host's MAC address is 00:15:C5:EA:A7:7B, the MAC Address used for power operations will be 00:15:C5:EA:A7:7F (the sum of the original MAC address and the offset).
 - B. **Relative** — Choose an IP address offset and select whether it is Greater Than, the Same as, or Less Than the host's IP address. For example, if you choose a Greater Than offset of 0.0.1.0 and the host's IP address is 10.3.0.14, the host's BMC address will be 10.3.1.14. This is the IP address used for power operations (the sum of the original IP address and the offset).
 - C. **Static** If you choose Static or if you wish to use different settings for each host, you must configure the DRAC and ILO options individually for each host.
4. (Optional) Select the MAC Address vs. Host MAC Address type:
 - A. Not Related
 - B. Greater Than
 - C. Less Than
5. (Optional) Enter the MAC Address Offset.
6. Select the MAC Address to use to manage this host.
7. (Optional) Select the IP Address vs. Host IP Address type:
 - A. Greater Than
 - B. Less Than
 - C. Same As
8. (Optional) Enter the IP address offset from the management interface.
9. (Optional) Enter the IP address for the host.
10. Select a Platform Management User.

 **Note**

Users must be members of the Power group to appear in this list.

Import Hosts

Cluster Manager provides an easy way to import a large group of hosts from a file. When importing a list of hosts, it is important to note that Cluster Manager imports only host information. Cluster Manager accepts the following file types: *nodes.conf*, *dbix*, or *CSV*.

To Import a List of Hosts

1. Obtain or create a host list file for importing. The following examples depict *nodes.conf*, *dbix*, and *CSV* file formats:

- A. *nodes.conf*

SGI *nodes.conf* format lists one host per line with properties being space or tab delimited:

```
MAC HOSTNAME IP_ADDRESS BOOT_MODE UNIQUE_NUM DESCRIPTION
```

Example:

```
0050455C0392 n001 192.168.4.1 boot_mode 1 Node_n001
0050455C03A2 n002 192.168.4.2 boot_mode 2 Node_n002
```

- B. *dbix*

dbix

```
hosts.<hostname>.description: <description>
hosts.<hostname>.enabled:true
hosts.<hostname>.name:<hostname>
hosts.<hostname>.partition:<partition>
interfaces.<MAC_address1>.address:<IP_address1>
interfaces.<MAC_address1>.mac:<MAC_address1>
interfaces.<MAC_address1>.management:true
interfaces.<MAC_address1>.owner:<hostname>
interfaces.<MAC_address2>.address:<IP_address2>
interfaces.<MAC_address2>.mac:<MAC_address2>
interfaces.<MAC_address2>.management:false
interfaces.<MAC_address2>.owner:<hostname>
```

Example:

```
hosts.n1.description:Added automatically by add_hosts.shasd
hosts.n1.enabled:true
hosts.n1.name:n1
hosts.n1.partition:computehosts
interfaces.0030482acc96.address:10.0.1.1
interfaces.0030482acc96.mac:0030482acc96
interfaces.0030482acc96.management:true
interfaces.0030482acc96.owner:n1
interfaces.0030482acc9a.address:10.0.2.1
interfaces.0030482acc9a.mac:0030482acc9a
interfaces.0030482acc9a.management:false
interfaces.0030482acc9a.owner:n1
```

Note

Dbix files are created primarily by obtaining and editing a Cluster Manager database file.

C. CSV

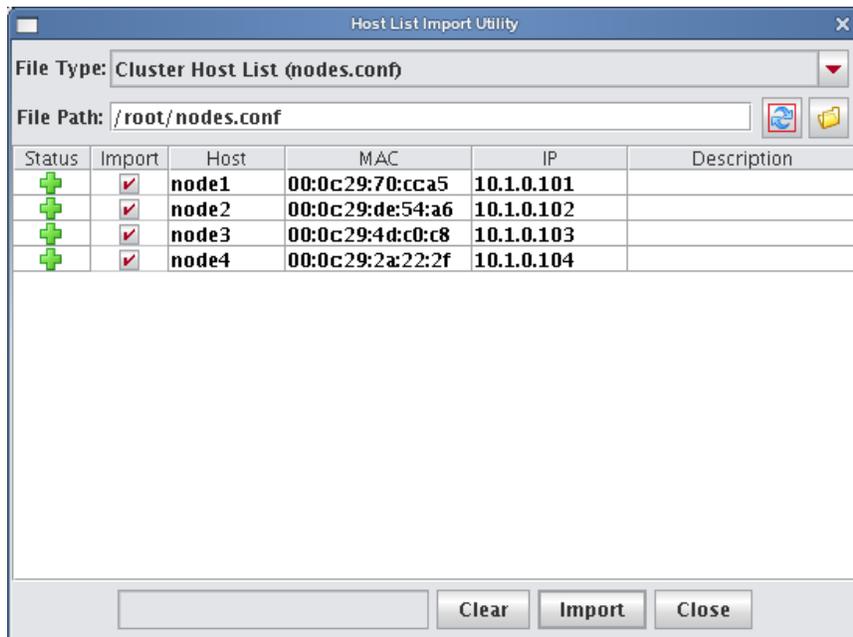
HOSTNAME,MAC_ADDRESS1,IP_ADDRESS1,DESCRIPTION,MAC_ADDRESS2,IP_ADDRESS2

Example:

n14,"0040482acc96,0040482acc9a","10.4.1.1,10.4.2.1",Description

2. Select **Import Host List** from the **File** menu.

The Host List Import Utility dialog appears.



3. Select the host list file type you are importing.
If you change the file type, click **Refresh** to update the dialog.
4. Enter the path for the file you want to import or click **Browse** to locate the file.
5. Review the list of hosts to import and un-check any hosts you do not want.
Errors display for items that cannot be imported.

 **Tip**

To clear the list of selected hosts, click **Clear**.

6. Click **Import** to import the list of hosts.
7. Click **Close**.

Add Partitions

Partitions are used to separate clusters into non-overlapping collections of hosts. Hosts that belong to a partition may not be used by anyone who is not authorized to access the partition. Within a partition, host access may be shared between regions of users.

To Add a Partition

1. Select **New Partition** from the **File** menu or right-click in the host navigation tree and select **New Partition**.
A new partition pane appears.
2. Enter a partition name.
3. (Optional) Enter a description.
4. In the **Regions** pane, click **Add**.
The Select Regions dialog appears.

Note

This dialog displays only regions that are not already assigned to a partition.

5. Select regions you want to include in this partition and click **OK**.
6. In the **Hosts** pane, click **Add** to display the Select Hosts dialog.
7. Select a hosts to add to this partition and click **OK**.
8. Click **Apply**.

Create Regions

A region is a subset of a partition and may share any hosts that belong to the same partition — even if the hosts are currently used by another region. Adding regions allows you to more closely allocate resources to specific groups and users.

To Create a Region

1. Select **New Region** from the **File** menu or right-click in the host navigation tree and select **New Region**.
A new region pane appears.
2. Enter a region name.
3. (Optional) Enter a description.
4. (Optional) From the drop-down list, select the name of the partition to which you want to assign the region.

Note

Regions not assigned to a partition become part of the default or unassigned partition.

5. From the **Hosts** pane, click **Add** to assign a hosts to the region.
The Select Hosts dialog appears.
6. Select the hosts you want to add to the region.
7. Click **OK**.
8. From the **Groups** pane, click **Add**.
The Select Groups dialog appears.
9. Select the groups you want to add to the region.
Adding groups to the region defines which users may access the hosts assigned to the region.
10. Click **OK**.



If you forget to assign groups to the region, the hosts appear to be non-existent.

11. Click **Apply**.

Chapter 4

Creating Payloads and Images

Payload Management

Payloads are stored versions of the operating system and any applications installed on the hosts. Payloads are compressed and transferred to the hosts via multicast during the provisioning process.

Configuring a Payload Source

Before you can build a new payload, you must have a package source available for use. A package source can be the RHEL or SLES physical media, ISO media, ftp or http install, or media copied to your hard drive.

Physical Media

If you are using physical media, you must insert it and mount it for your CDROM:

```
/mnt/cdrom
```

or

```
/media/dvd
```

CD ISOs

If you are using the CD ISOs, you must mount the ISOs one at a time to simulate using the CDROM:

```
mount -o loop <ISO_name> <mount_point>
```

Note

Using either the multiple disks or multiple ISOs may require switching between disks several times.

DVD ISOs

DVD ISOs are perhaps the most convenient because they are simply mounted and do not require changing disks. To use a DVD ISO:

```
mount -o loop <ISO_name> <mount_point>
```

FTP or HTTP

You must follow the operating system vendors recommendations for setting up a network based installation. Some problems have been reported using Apache 2.2.

Copying the Media

If you have CD media or CD ISOs and will be creating multiple payloads or requiring additional packages following payload creation, it is worthwhile to copy the distribution to the hard drive. See *Red Hat Installations* below or *SuSE Linux Enterprise Server Installations* on page 30 for instructions on how to copy the installation disks for your distribution.

RED HAT INSTALLATIONS

Tip

If you choose to copy the *entire* contents of each disc rather than the files described below, you must copy disc1 *LAST*. Failure to copy disks in the correct order may produce payload creation failures (for example, package *aaa_base* may not be found).

1. Mount disk 1 and copy the contents of the entire disk to a location on the hard drive:

```
mount /mnt/cdrom  
  
or  
mount -o loop RHEL-x86_64-WS-disc1.iso /mnt/cdrom  
mkdir /mnt/redhat  
cp -r /mnt/cdrom/* /mnt/redhat
```

2. Mount disk 2 and copy the *.rpm files from the RPMS directory to the RPMS directory on the hard drive:

```
cp /mnt/cdrom/RedHat/RPMS/*.rpm /mnt/redhat/RedHat/RPMS
```

3. Mount each remaining disk and copy the RPMS directory to the RPMS directory on the hard drive.

SUSE LINUX ENTERPRISE SERVER INSTALLATIONS

Tip

If you choose to copy the *entire* contents of each disc rather than the files described below, you must copy disc1 *LAST*. Failure to copy disks in the correct order may produce payload creation failures (for example, package *aaa_base* may not be found).

1. Mount disk 1 and copy the contents of the entire disk to a location on the hard drive:

```
mount /media/cdrom  
  
or  
mount -o loop SLES-9-x86-64-CD1.iso /media/cdrom  
mkdir /mnt/suse  
cp -r /media/cdrom/* /mnt/suse
```

2. Mount disk 2 and copy the RPMs from each architecture subdirectory to the SuSE directory on the hard drive:

```
cp -r /media/cdrom/suse/noarch/* /mnt/suse/suse/noarch  
cp -r /media/cdrom/suse/i586/* /mnt/suse/suse/i586  
cp -r /media/cdrom/suse/i686/* /mnt/suse/suse/i686  
cp -r /media/cdrom/suse/src/* /mnt/suse/suse/src  
cp -r /media/cdrom/suse/nosrc/* /mnt/suse/suse/nosrc  
cp -r /media/cdrom/suse/x86_64/* /mnt/suse/suse/x86_64
```

3. Mount each remaining disk and copy the RPMs from each architecture subdirectory to the SuSE directory.

Create a Payload

Payloads are initially created using a supported Linux distribution installation media (CD-ROM, FTP, NFS) to build a base payload (see *Operating System Requirements* on page 2 for a list of supported distributions) or by importing a payload from a previously provisioned host. Additions and changes are applied by adding or removing packages, or by editing files through the GUI or CLI. Changes to the Payload are managed by the Cluster Manager Version Control System (VCS). Package information and files are stored and may be browsed through Cluster Manager.

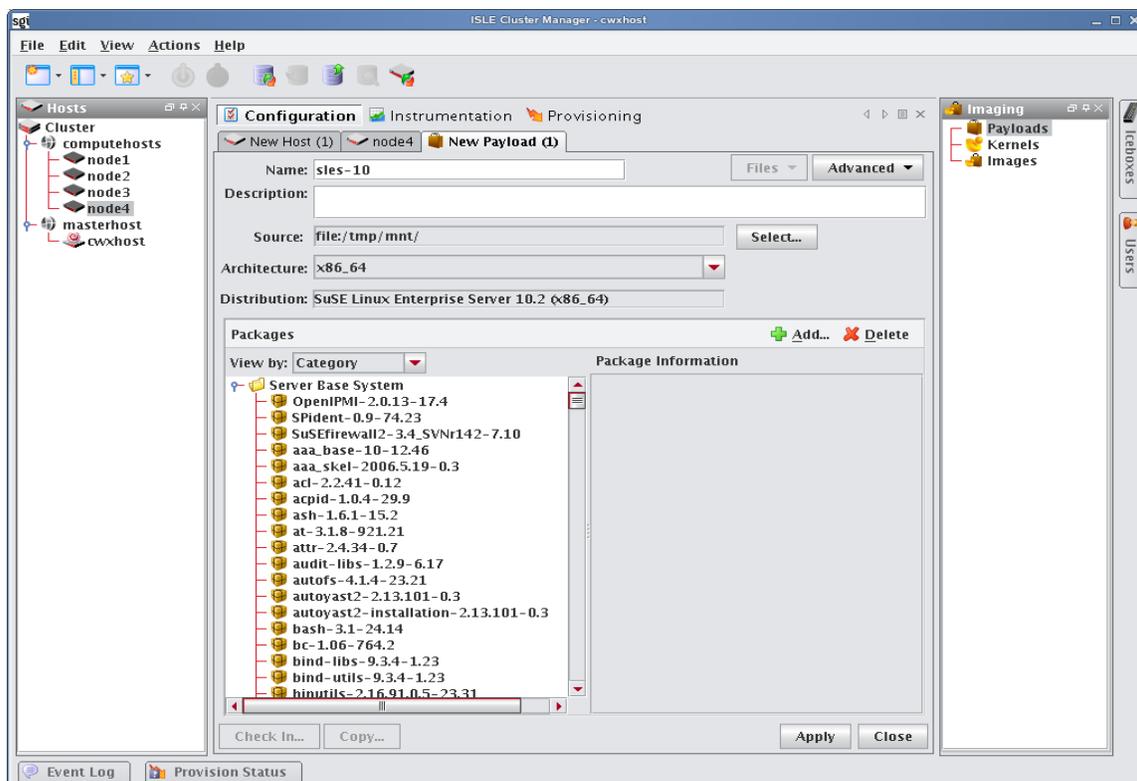
! Warning!

Please consult SGI before upgrading your Linux distribution or kernel. Upgrading to a distribution or kernel not approved for use on your system may render Cluster Manager inoperable or otherwise impair system functionality. Technical Support is not provided for unapproved system configurations.

To Create a New Payload

To create a new payload from a Linux distribution, do the following:

1. Select **New Payload** from the **File** menu or right-click in the imaging navigation tree and select **New Payload**.



A new payload pane appears.

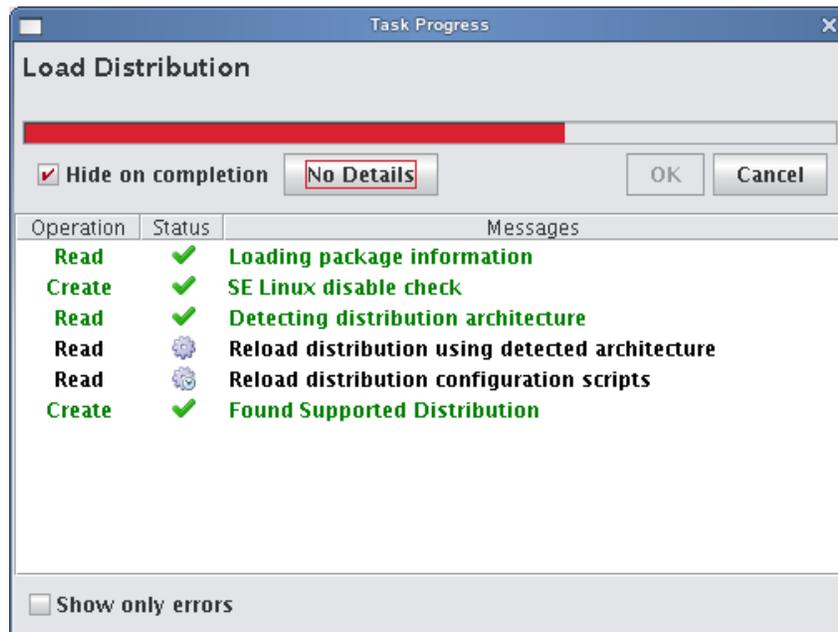
2. Enter the name of the new payload in the **Name** field.
3. (Optional) Enter a description of the new payload in the **Description** field.
4. Click **Select** to display the Package Source dialog.
5. From the **Scheme** drop-down list, select **file:**, **http://**, or **ftp://**.

6. If you select the **file:** scheme, click **Browse** to locate the Linux distribution you want to use. Otherwise, enter the URL or FTP address of the Linux distribution you want to use.

 **Tip**

If you are creating multiple payloads from the same distribution source, it may be faster and easier to copy the distribution onto the hard drive. This also prevents you from having to switch CD-ROMs during the payload creation process. See *Red Hat Installations* on page 30 and *SuSE Linux Enterprise Server Installations* on page 30 for specific details on installing these distributions.

7. (Optional) Enter a Host (only if the selected scheme is **http://** or **ftp://**).
8. (Optional) Enter Username and Password (only if you selected **Use Authentication**).
9. Click **OK** to continue.
10. As the distribution loads, the Task Progress dialog appears. This dialog displays the payload creation status and identifies any errors that occur during the load process.



 **Tip**

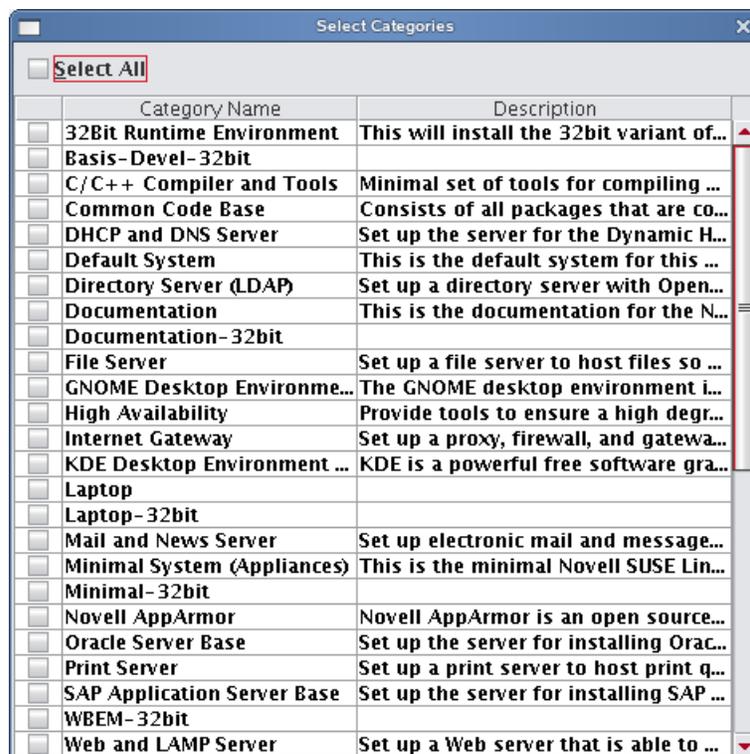
Select Hide on Completion to close the Task Progress dialog if no errors or warnings occur.

 **Note**

If Cluster Manager is unable to detect payload attributes, the Distribution Unknown dialog appears. From this dialog, select the distribution type that most closely resembles your distribution and Cluster Manager will attempt to create your payload.

11. (Optional) In the packages pane, click **Add** to include additional packages in the payload.

The Select Packages dialog appears.



12. Select which payload categories to install or remove by clicking the checkbox next to each package.

Note

When you select a “core” category to include in a payload, Cluster Manager automatically selects packages that are essential in allowing the capability to run. However, you may include additional packages at any time.

13. Click **OK** to continue.
14. (Optional) From the Packages pane, select packages you want to remove from the payload, then click **Delete** in the packages pane.
15. (Optional) Configure any advanced settings you want to apply to the payload.
16. Click **Apply** to save changes and build the payload.
The payload progress dialog appears.

Tip

If an RPM installation error occurs during the payload creation process, Cluster Manager enables the **Details** button and allows you to view which RPM produced the error.

To view error information about a failed command, click the command description field. You may copy the contents of this field and run it from the CLI to view specific details about the error.

17. (Optional) Select any payload files you wish to include with, remove from, or edit from the **File** drop-down menu.
18. (Optional) Click **Check In** to import the new payload into the Version Control System (VCS).

Install Cluster Manager into the Payload

When working with payloads, Cluster Manager requires that each payload contain some basic Cluster Manager services. These services allow Cluster Manager to control various parts of the system, including instrumentation services and the monitoring and event subsystem.

There are two ways to install Cluster Manager into a payload: through the Cluster Manager installer or by adding RPMs included on the Cluster Manager ISO (*\$CMHOME/rpms*) to the payload. The Cluster Manager installer assumes that the payload uses the same OS and architecture as the Master Host. If you are trying to support a different OS or Architecture (32 or 64-bit), use the RPMs on the Cluster Manager ISO. The ISO includes several Cluster Manager and Cluster Manager-Java RPMs, so you must select the RPMs that match your OS and architecture.

Note

By default, RPMs use *host* as the Master Host in the *@genesis.profile* file. You must change this to match the name of your Master Host or make sure *host* can be resolved through the */etc/hosts* file or DNS. If you change the *system.mn.host* in the *@genesis.profile*, you must also change the Transmitter Address in *\$CMHOME/etc/InstrumentationService.profile*.

To Install Cluster Manager into the Payload

1. Access the payload's root directory:

```
cd $CMHOME/imaging/<username>/payloads/<payload>/
```
2. Run the install script from the Cluster Manager Installation:

```
/mnt/cdrom/install.sh
```
3. In the Cluster Manager Installer, select **Payload**.



4. Enter the name of the Cluster Manager Server.

Note

The Cluster Manager Server must be a valid host name that can be resolved through name resolution (using DNS or the */etc/hosts* file).

Tip

If the hosts will communicate with one another, you may prefer to use host names (such as, node1, node2, node3 and so on) rather than IP addresses. Hosts use “name resolution” to convert between numbers and names — most commonly through a local `/etc/hosts` file or with DNS.

For typical Cluster Manager users, the local `/etc/hosts` file already exists on the Master Host computer. To make this file available to all hosts, copy the file into the payload (or simply edit the file in the payload). If you need to create this file, use the `dbx` command to create an `/etc/hosts` formatted list of the hosts in your cluster. To save this list, redirect the output to a file:

```
dbx -f:hosts
```

5. Enter the installation directory (for example, `/opt/sgi/islecm`).
6. Click **Next**.
7. Review the installation settings and click **Install** to continue.
8. After the installation completes, click **Finish**.

Kernel Management

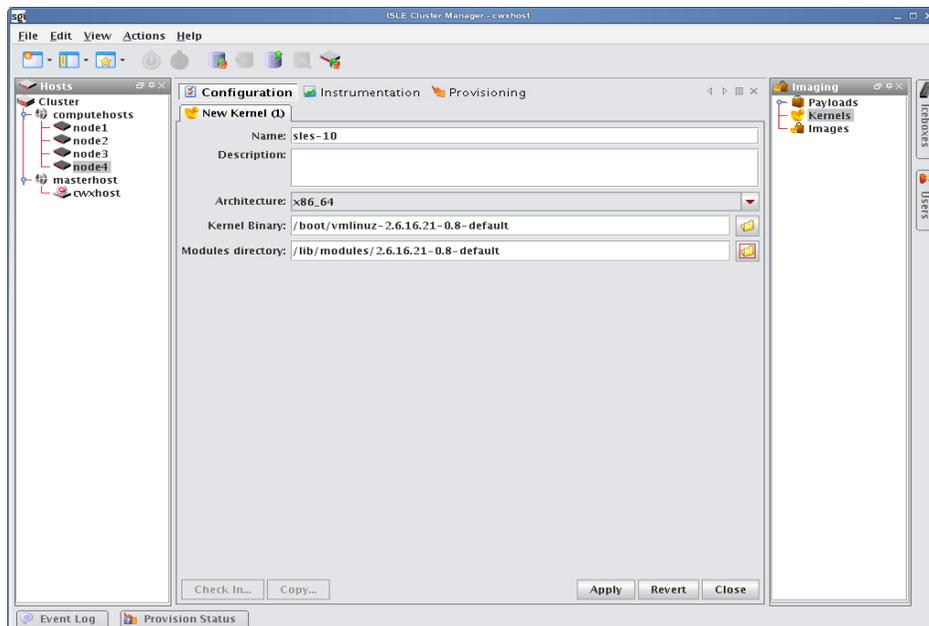
Kernels may be customized for particular applications and used on specific hosts to achieve optimal system performance. Cluster Manager uses VCS to help you manage kernels used on your system.

To Create a Kernel Using an Existing Binary

Note

For information on building a new kernel from source, see *Building a New Kernel from Source* on page 37.

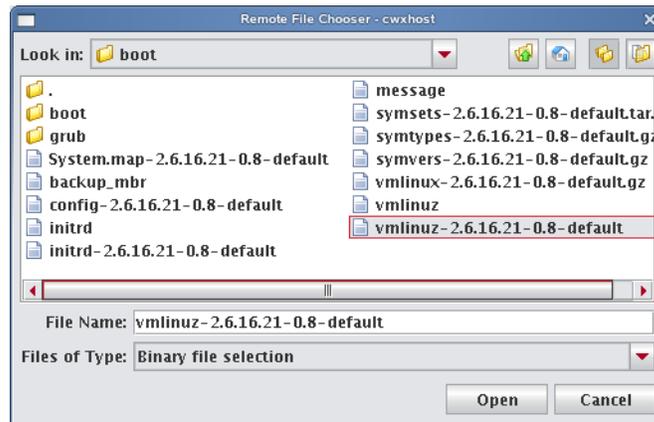
1. Select **New Kernel** from the **File** menu or right-click in the imaging navigation tree and select **New Kernel**.



2. Enter the kernel name.
3. (Optional) Enter a kernel description.
4. Select the hardware architecture.
5. Enter the full path to the kernel binary you want to use, or click **Browse** to find the kernel binary you want to use.

! Warning!

Make sure you select a kernel binary that begins with **vmlinuz** and not **vmlinux**. Selecting **vmlinux** will result in provisioning problems later on.



6. Click **Browse** to find the modules directory and click **Open**.
7. Click **Apply** to create the kernel.
8. (Optional) Click **Check In** to import the kernel into VCS.

To Create a Copy of an Existing Kernel

1. Right-click on an existing kernel in the **Imaging** navigation tree and select **Copy**.

Tip

You may also open a kernel for editing, then click the **Copy** button in the panel.

2. In the Copy Kernel dialog, enter a name for the new kernel.



3. Click **OK**.

Building a New Kernel from Source

If you want to use a stock vendor kernel already loaded on your system, see *To Create a Kernel Using an Existing Binary* on page 35. Otherwise, use the following procedure to build a new kernel from source:

Warning!

Please consult SGI before upgrading your Linux distribution or kernel. Upgrading to a distribution or kernel not approved for use on your system may render Cluster Manager inoperable or otherwise impair system functionality. Technical Support is not provided for unapproved system configurations.

1. Obtain and install the kernel source RPM for your distribution from your distribution CD-ROMs or distribution vendor.

This places the kernel source code under `/usr/src`, typically in a directory named `linux-2.<minor>.<patch>-<revision>` (if building a Red Hat Enterprise Linux kernel, Cluster Manager places the source code into `/usr/src/kernels/2.<minor>.<patch>-<revision>`).

Tip

Because you don't need the kernel source RPM in your payload, install the RPM on the host.

2. If present, review the README file inside the kernel source for instructions on how to build and configure the kernel.

Note

It is highly recommended that you use, or at least base your configuration on one of the vendor's standard kernel configurations.

3. Typically, a standard configuration file is installed in the `/boot` directory, usually as `config-2.<minor>.<patch>-<revision>`.

You may also use a stock configuration file installed as `.config` in the kernel source directory or available in a sub-directory (typically `/configs`) of the kernel source directory.

Tip

To use a stock configuration, copy it to the kernel source directory and run:

```
make oldconfig.
```

4. Build the kernel and its modules using the `make bzImage && make modules` command.

If your distribution uses the Linux 2.4 kernel, use `make dep && make bzImage && make modules` but **DO NOT** install the kernel.

5. From the **File** menu, select **Source Kernel**.

6. In the New Kernel panel, enter the name of the kernel.

7. (Optional) Enter a description of the kernel.

8. Select the hardware architecture.

9. Enter the location of the kernel source (the directory where you unpacked the kernel source) in the **Source Directory** field or click **Browse** to find the source directory and click **Open**.

By default, kernel source files are located in `/usr/src`.

10. (Optional) Enter the binary path of the kernel (for example, *arch/i386/boot/bzImage*) or click **Browse** to find the kernel and click **Open**.
11. Select the modules directory and click **Open**.
12. Click **Apply** to create the kernel.
13. (Optional) Click **Check In** to import the kernel into VCS.

Image Management

Images contain exactly one payload and one kernel, and allow you to implement tailored configurations on various hosts throughout the cluster.

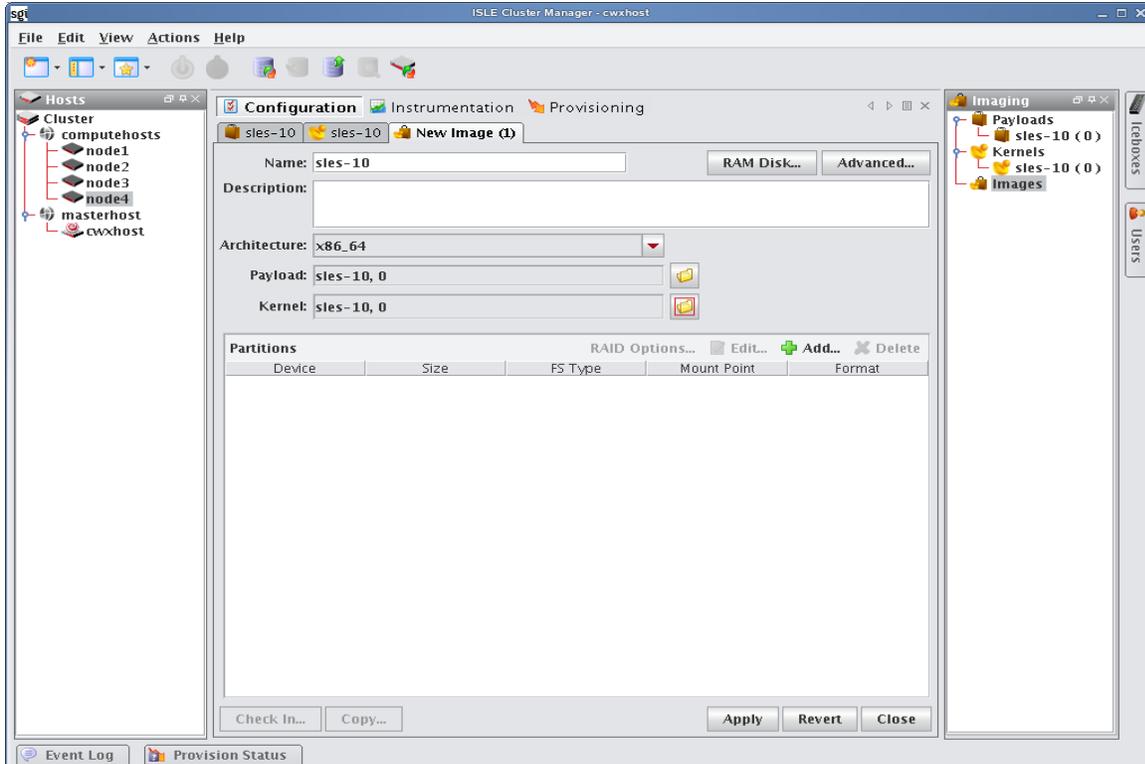
! Warning!

Please consult SGI before upgrading your Linux distribution or kernel. Upgrading to a distribution or kernel not approved for use on your system may render Cluster Manager inoperable or otherwise impair system functionality. Technical Support is not provided for unapproved system configurations.

Create an Image

To Create an Image

1. Select **New Image** from the **File** menu or right-click in the **Imaging** navigation tree and select **New Image**.



2. Enter a name in the **Name** field.
3. (Optional) Enter a description in the **Description** field.

4. Select the architecture supported by the kernel.
5. Click the **Browse** folder icon to select a **Kernel**.
6. Select the **Browse** folder icon to select a **Payload**.
7. Define the partition scheme used for the compute hosts — the partition scheme must include a root (/) partition. See *To Create a Partition for an Image* on page 41.

Note

Kernel support for selected file systems must be included in the selected kernel (or as modules).

8. (Optional) Implement RAID. See *Managing Partitions* on page 41.
9. (Optional) If you need to make modifications to the way hosts boot during the provisioning process, click the **RAM Disk** tab.
10. (Optional) Click the **Advanced** button to display the Advanced Options dialog.
This dialog allows you to configure partitioning behavior and payload download settings (see *Advanced Imaging Options*).
11. Click **Apply**.

Advanced Imaging Options

The Advanced Options dialog allows you to configure partitioning behavior and payload download settings. These settings are persistent, but may be overridden from the Advanced Provisioning Options dialog.



PARTITIONING BEHAVIOR

This option allows you to configure the partition settings used when provisioning a host. You may automatically partition a host if the partitioning scheme changes, always re-create all partitions (including those that are exempt from being overwritten), or choose to never partition the host. See *Managing Partitions* on page 41.

PAYLOAD DOWNLOAD

The payload options allow you to *automatically download* a payload if a newer version is available (or if the current payload is not identical to that contained in the image), *always download* the payload, or choose to *never download* a payload.

boot.profile

Cluster Manager generates the *boot.profile* file each time you save an image (overwriting the previous file in */etc/boot.profile*). The boot profile contains information about the image and is required for the boot process to function properly. You may configure the following temporary parameters:

dmesg.level	The verbosity level (1-8) of the kernel — 1 (the default) is the least verbose and 8 is the most.
partition	Configure the hard drive re-partitioning status (Automatic, Always, Never). By default, Automatic.
partition.once	Override the current drive re-partitioning status (Default, On, Off). By default, Default.
image	Configure the image download behavior (Automatic, Always, Never). By default, Automatic. Always and Never will download the image even if it is up-to-date.
image.once	Override the current image download behavior (Default, On, Off). By default, Default. To view the current download behavior, see <i>Advanced Imaging Options</i> on page 39.
image.path	Specifies where to store the downloaded image. By default, <i>/mnt</i> .

To change the configuration of one of these parameters, add the parameter (for example, *dmesg.level: 7*) to the *boot.profile* and provision using that image. You may also configure most of these values from the GUI. See *Select an Image and Provision* on page 43.

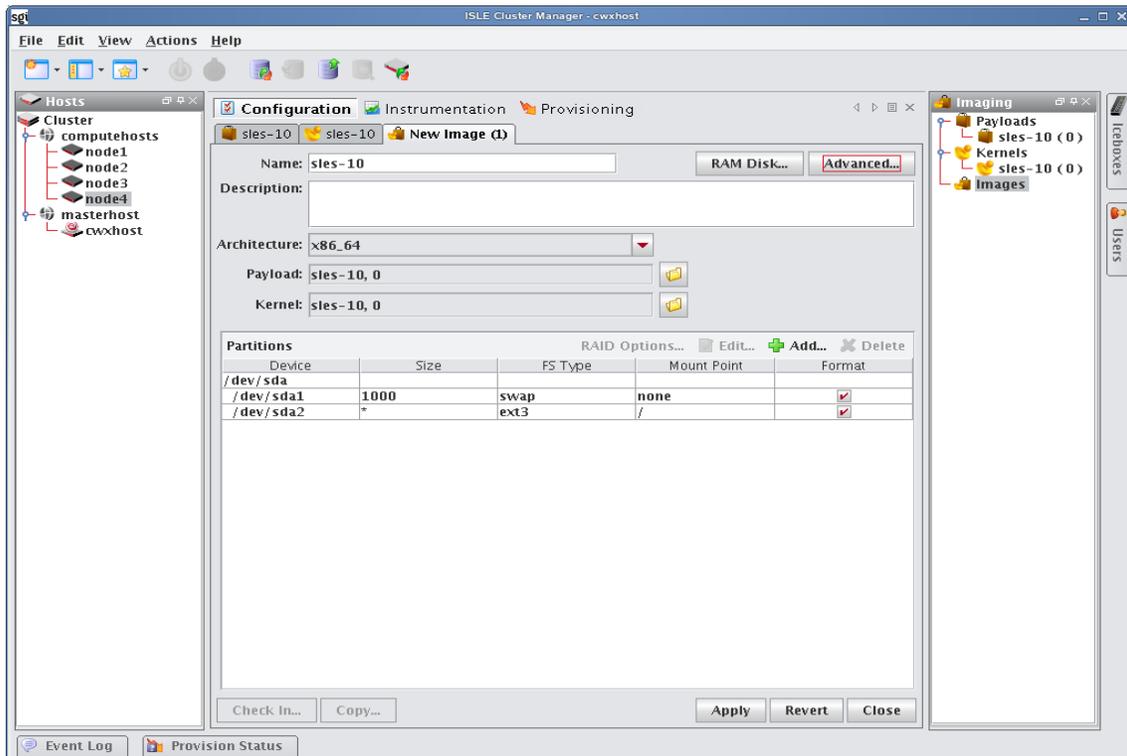
Note

Changes made to image settings remain in effect until the next time you save the image.

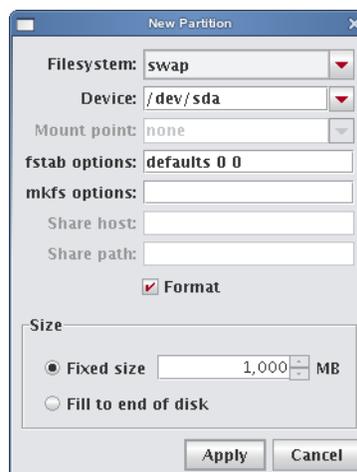
Managing Partitions

To Create a Partition for an Image

1. Right-click an existing image in the **Imaging** navigation tree and select **Edit**.



2. In the partitions pane, click **Add**.



3. In the New Partition dialog, select a file system type from the **Filesystem** drop-down list.
4. Enter the device on which to add the partition or select a device from the drop-down list. Supported devices include the following, but the most common is `/dev/hda` because hosts typically have only one disk and use IDE:

- /dev/hda — Primary IDE Disk
- /dev/hdb — Secondary IDE Disk
- /dev/sda — Primary SCSI Disk
- /dev/sdb — Secondary SCSI Disk

 Tip

If you are using non-standard hosts, you can add additional storage devices to the partitioning drop-down list. The Image Administration Service profile, `$CMHOME/etc/ImageAdministrationService.profile`, allows you to configure non-standard hard drives. This profile contains options that allow you to set the drive name (available when partitioning the disk at the time of creating or modifying an image) and the prefix for a partition on the drive (if one exists). By default these values are commented out, but may be commented in as needed. Once drives are configured, they become available via Cluster Manager.

Profile options are as follows:

`partitioning.devices:cciss/c0d0`

The name of the storage device where the device file is located (e.g., `/dev/cciss/c0d0`).

`partitioning.devices.cciss/c0d0.naming:p`

The partition prefix for the device defined by the previous key (e.g., `cciss/c0d0`).

In this example, the partition will look like `c0d0p1`, `c0d0p2`, and so on.

5. Enter a Mount Point or select one from the drop-down list.
6. (Optional) Enter the `fstab` options. The `/etc/fstab` file controls where directories are mounted and, because Cluster Manager writes and manages the `fstab` on the hosts, any changes made on the hosts are overwritten during provisioning.
7. (Optional) Enter the `mkfs` options to use when creating the file system (such as, file size limits and symlinks). For example, to change the default block size for `ext3` to `4096`, enter `-b 4096` in the `mkfs` options field.
8. (Optional) If creating an NFS mount, enter the NFS host.
9. (Optional) If creating an NFS mount, enter the NFS share.
10. (Optional) Clear the **Format** option to make the partition exempt from being overwritten or formatted when you provision the host. This may be overridden by the **Force Partitioning** option or from the `boot.profile` file (see *Select an Image and Provision* on page 43 and `boot.profile` on page 40).

 Note

After partitioning the hard disks on a host for the first time, you can make a partition on the disk exempt from being overwritten or formatted when you provision the host. However, deciding not to format the partition may have an adverse affect on future payloads — some files may remain from previous payloads. This option is not allowed if the partition sizes change when you provision the host.

11. Select the partition size:

- Fixed size allows you to define the size of the partition (in megabytes).
- Fill to end of disk allows you to create a partition that uses any space that remains after defining partitions with fixed sizes.

Tip

You should allocate slightly more memory than is required on some partitions. To estimate the amount of memory needed by a partition, use the `du -hc` command.

12. Click **Apply**.

13. (Optional) Click **Check In** to import the image into VCS.

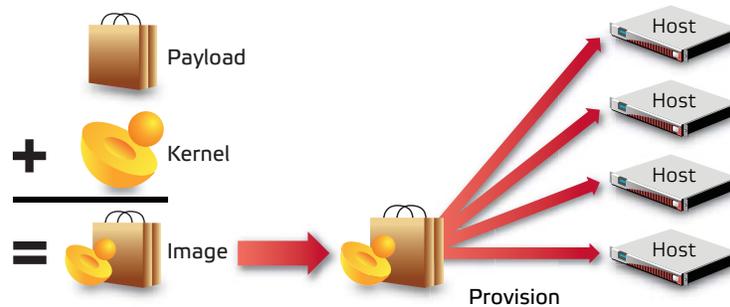
14. Click **Apply**.

Note

Cluster Manager generates the *boot.profile* file each time you save an image. For a description of the information contained in this file, see *boot.profile* on page 40.

Provisioning

The Cluster Manager provisioning service allows you to create an image from a payload and kernel, then apply that image to multiple hosts. When provisioning, you can select a versioned image stored in VCS or use a working copy of an image from your working directory. The following illustration depicts an image that is provisioned to multiple hosts.



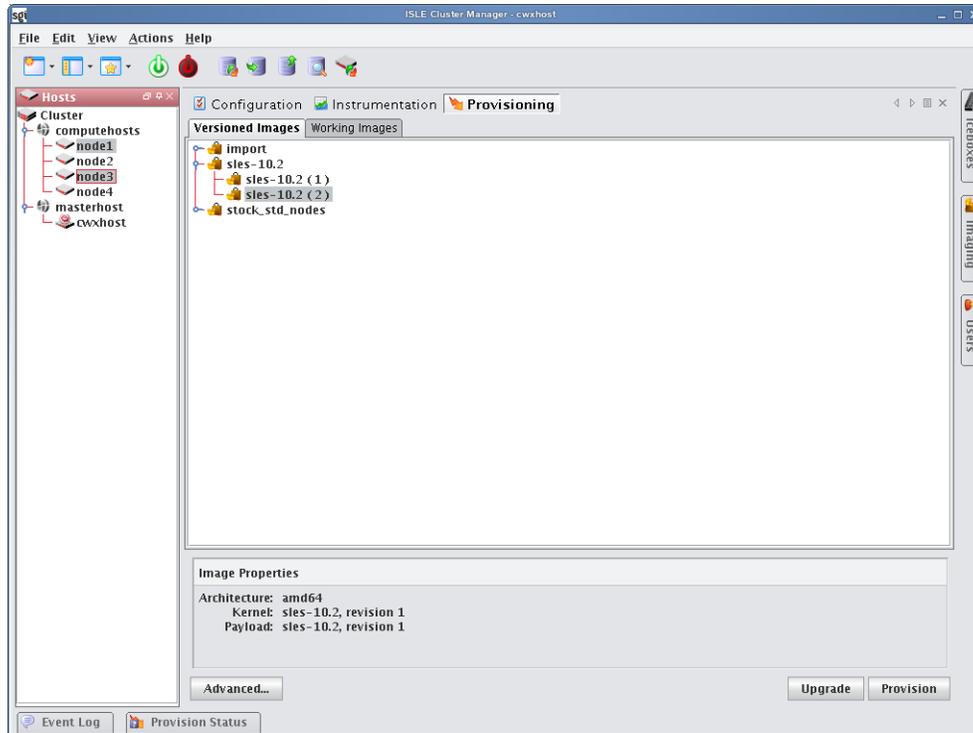
Select an Image and Provision

1. Select the hosts you want to provision from the navigation tree (use the **Shift** or **Ctrl** keys to select multiple hosts).

Tip

If you want to provision a host using the latest revision of an image stored in VCS, you can right-click a host and select **Provision**. Cluster Manager displays a pop-up menu and allows you to select the image you want to use to provision.

2. Select the **Provisioning** tab.



3. Select the **Versioned Images** or **Working Images** tab.

Note

A “Versioned” image is a revision of an image that is checked into VCS. A “Working” image has not been checked into VCS and is currently present in the working area (for example, `$CMHOME/imaging/<user>/images`). This allows you to test changes prior to checking in.

A “Working Copy” of an image is currently present in the working area (for example, `$CMHOME/imaging/<user>/images`). A “Versioned” image is a revision of an image stored in VCS.

4. Select the image you want to use to provision the hosts.
5. (Optional) Click the **Advanced** button to display the Advanced Options dialog.
This dialog allows you to override partitioning, payload, and kernel verbosity settings.
6. Click **Provision** to distribute the image to the selected hosts.
Cluster Manager asks you to confirm your action.
7. Click **Yes** to provision the hosts.

Warning!

When you click **Yes**, Cluster Manager provisions the hosts using the new image. Any pending or running jobs on the selected hosts are lost.

Right-click Provisioning

1. Select the hosts you want to provision from the navigation tree (use the **Shift** or **Ctrl** keys to select multiple hosts).
2. Right-click a host and select **Provision**.
Cluster Manager displays a po-pup menu and allows you to select the image you will use to provision.

 **Note**

Right-click provisioning uses the latest revision of an image stored in VCS.

Index

A

accounts
 disable user 16
 enable 16
add
 group 16
 host 19
 package
 to new payload 33
 user
 to group 16
annotations
 electric shock iv
 note iv
 tip iv
 warning iv

B

block size 42
boot.profile 40

C

cluster 19
 environment 13
Cluster Master
 administration
 grant privileges 18
 install into payload 34
 server, start and stop 9
 services 10

 system requirements 2
compute host (*See* host)
copy
 kernel 36
create
 group 16
 host 19
 image 38
 kernel 35
 multiple payloads from source 32
 partition 41
 password
 user 15
 payload 31
csv 25
Customer service v

D

dbix 25
default user administration settings 14
dhcpd.conf 11
disable
 user account 16
distribution, upgrade 2
dmesg.level 40
dockable frames 9
Documentation
 available via the World Wide Web iii
DRAC 24

E

electric shock iv
 enable
 user account 16
 errors
 RPM 33

F

feedback, documentation iv
 format partition 42
 frames
 controls 9
 dockable 9
 fstab 42

G

GID 13
 group
 add 16
 assign roles to 16
 assign to role 17
 assign user to 16
 GID 13
 grant access to region 16
 power 16
 primary 16
 region, add to 27
 user membership 16

H

hardware
 system requirements 1
 host 19
 add 19
 to partition 27
 administration 19
 grant privileges 18
 import 25
 Master Host 19
 names 35

provision 43
 region
 add host to 27
 assign host to 22
 shared 19

I

Icebox
 administration privileges 18
 ILO 24
 image 40
 create 38
 management 38
 partition 41
 privileges, enable imaging 18
 provision 43
 image.once 40
 image.path 40
 import
 host list 25
 install
 Cluster Master 4
 into payload 34
 interface
 management 21
 IP address 52
 host 21
 IPMI 22

K

kernel
 build from source 35
 copy 36
 create 35
 management 35
 upgrade 2

M

management
 interface 21

Master Host
 definition 19
 system requirements 1
 memory
 estimate partition requirements 43
 mkfs 42

N

name resolution 35
 navigation tree 9
 netmask 52
 NFS 13
 nodes.conf 25
 note iv

O

operating system requirements 2
 override global settings 22

P

package
 add to new payload 33
 partition 19, 26, 40
 add
 host to 27
 region to 27
 create 41
 estimate memory requirements 43
 format 42
 manage 41
 overwrite protection 42
 partitioning behavior 39
 save 42
 size
 fill to end of disk 43
 fixed 43
 partition.once 40
 password
 create new 15
 payload

add
 package to new 33
 attributes, troubleshoot 32
 create 31
 multiple payloads from source 32
 download 39
 install Cluster Master into 34
 management 29
 permissions 17
See role; privileges
 platform management
 DRAC 24
 ILO 24
 IPMI 22
 platform support 1
 power
 group 16
 primary
 group 16
 privileges 17
 Cluster Master 18
 database 18
 host administration 18
 Icebox administration 18
 imaging 18
 instrumentation 18
 logging 18
 power 18
 provisioning 18
 serial 18
 user administration 18
 Product support v
 provision 43
 format partition 42
 right-click 45

R

region 19, 27
 add
 group to 27
 host to 27
 to partition 27

- assign to host 22
- grant group access to 16
- requirements
 - hardware 1
 - operating system 2
 - software 2
- RHEL 53
- right-click menu
 - provisioning 45
- rights
 - See* role; privileges
- role
 - assign group to 17
 - assign to group 16
 - grant privileges and permissions 17
- RPM
 - errors 33

S

- save
 - partition 42
- SLES 53
- software
 - requirements 2
- start Cluster Master server 9
- stop Cluster Master server 9
- system
 - requirements
 - hardware 1
 - operating system 2

T

- task progress dialog 32
- Technical support v
- third-party power controls 16
- tip iv
- toolbar 8
- troubleshooting
 - payload attributes 32
 - RPM errors 33

U

- UID 13, 15
- upgrade
 - distribution 2
 - kernel 2
- user
 - add
 - to group 16
 - administration
 - default settings 14
 - privileges 18
 - assign to group 16
 - group membership 16
 - UID 13, 15

V

- version
 - versioned copy 53

W

- warning iv
- working copy 53

Glossary

Anti-aliasing A technique used to smooth images and text to improve their appearance on screen.

Architecture-independent Allowing hardware or software to function regardless of hardware platform.

Baud rate A unit of measure that describes data transmission rates (in bits per second).

Block size The largest amount of data that the file system will allocate contiguously.

boot.profile A file that contains instructions on how to boot a host.

Boot utilities Utilities added to the RAM Disk that run during the boot process. Boot utilities allow you to create such things as custom, pre-finalized scripts using utilities that are not required for standard Linux versions.

Cluster Clustering is a method of linking multiple computers or compute hosts together to form a unified and more powerful system. These systems can perform complex computations at the same level as a traditional supercomputer by dividing the computations among all of the processors in the cluster, then gathering the data once the computations are completed. A cluster refers to all of the physical elements of your SGI solution, including the Cluster Manager Master Host, compute hosts, Cluster Manager, UPS, high-speed network, storage, and the cabinet.

Cluster Manager Master Host The Cluster Manager Master Host is the host that controls the remaining hosts in a cluster (for large systems, multiple masters may be required). This host is reserved exclusively for managing the cluster and is not typically available to perform tasks assigned to the remaining hosts.

DHCP Dynamic Host Configuration Protocol. Assigns dynamic IP addresses to devices on a network.

Diskless host A host whose operating system and file system are installed into physical memory. This method is generally referred to as RAMfs or TmpFS.

EBI An ELF Binary Image that contains the kernel, kernel options, and a RAM Disk.

Event engine Allows administrators to trigger events based on a change in system status (e.g., when processors rise above a certain temperature or experience a power interruption). Administrators may configure triggers to inform users of a specific event or to take a specific action.

Ext Original extended file system for Linux systems. Provides 255-character filenames and supports files sizes up to 2 Gigabytes.

Ext2 The second extended file system for Linux systems. Offers additional features that make the file system more compatible with other file systems and provides support for file system extensions, larger file sizes (up to 4 Terabytes), symbolic links, and special file types.

Ext3 Provides a journaling extension to the standard ext2 file system on Linux. Journaling reduces time spent recovering a file system, critical in environments where high availability is important.

Group A group refers to an organization with shared or similar needs. A cluster may contain multiple groups with unique or shared rights and privileges. A group may also refer to an administrator-defined collection of hosts within a cluster that perform tasks such as data serving, Web serving, and computational number crunching.

Health monitoring An element of the Instrumentation Service used to track and display the state of all hosts in the system. Health status icons appear next to each host viewed with the instrumentation service or from the navigation tree to provide visual cues about system health. Similar icons appear next to clusters, partitions, and regions to indicate the status of hosts contained therein.

Host An individual server or computer within the cluster that operates in parallel with other hosts in the cluster. Hosts may contain multiple processors.

image.profile A file used to generate *boot.profile*. This file contains information about the image, including the payload, kernel, and partition layout.

IP address A 32-bit number that identifies each sender or receiver of information. In order to transmit or receive information on the network.

Kerberos Kerberos is a network authentication protocol. It is designed to provide strong authentication for client/server applications by using secret-key cryptography.

Kernel The binary kernel, a *.config* file, *System.map*, and modules (if any).

LDAP Lightweight Directory Access Protocol is an Internet protocol that email programs use to look up contact information from a server.

Listener A listener constantly reads and reviews system metrics. Configuring listener thresholds allows you to trigger loggers to address specific issues as they arise.

Logger The action taken when a threshold exceeds its maximum or minimum value. Common logger events include sending messages to the centralized Cluster Manager message log, logging to a file, logging to the serial console, and shutting down the host.

MAC address A hardware address unique to each device installed in the system.

Metrics Used to track logger events and report data to the instrumentation service (where it may be monitored).

MIB Management Information Base. The MIB is a tree-shaped information structure that defines what sort of data can be manipulated via SNMP.

Monitors Monitors run periodically on hosts and provide the metrics that are gathered, processed, and displayed using the Cluster Manager instrumentation service.

Multi-user Allows multiple administrators to simultaneously log into and administer the cluster.

Netmask A string of 0's and 1's that mask or screen out the network part of an IP address so only the host computer portion of the address remains. The binary 1's at the beginning of the mask turn the network ID portion of the IP address into 0's. The binary 0's that follow allow the host ID to remain. A commonly used netmask is 255.255.255.0 (255 is the decimal equivalent of a binary string of eight ones).

NIS Network Information Service makes information available throughout the entire network.

Node See *Host*.

Partition Partitions are used to separate clusters into non-overlapping collections of hosts.

Payload A compressed file system that is downloaded via multicast during the provisioning process.

Plug-ins Programs or utilities added to the boot process that expand system capabilities.

RAID Redundant Array of Independent Disks. Provides a method of accessing multiple, independent disks as if the array were one large disk. Spreading data over multiple disks improves access time and reduces the risk of losing all data if a drive fails.

RAM Disk A small, virtual drive that is created and loaded with the utilities that are required when you provision the host. In order for host provisioning to succeed, the RAM Disk must contain specific boot utilities. Under typical circumstances, you will not need to add boot utilities unless you are creating something such as a custom, pre-finalized script that needs utilities not required by standard Linux versions (e.g., *modprobe*).

RHEL Red Hat Enterprise Linux.

Region A region is a subset of a partition and may share any hosts that belong to the same partition—even if the hosts are currently used by another region.

Role Roles are associated with groups and privileges, and define the functionality assigned to each group.

Secure remote access The ability to monitor and control the cluster from a distant location through an SSL-encrypted connection. Administrators have the benefit of secure remote access to their clusters through any Java-enhanced browser. Cluster Manager can be used remotely, allowing administrators access to the cluster from anywhere in the world.

Secure Shell (SSH) SSH is used to create a secure connection to the CLI. Connections made with SSH are encrypted and safe to use over insecure networks.

SLES SuSE Linux Enterprise Server.

Version branching The ability to modify an existing payload, kernel, or image under version control and check it back into VCS as a new, versioned branch of the original item.

Version Control System (VCS) The Cluster Manager Version Control System allows users with privileges to manage changes to payloads, kernels, or images (similar in nature to managing changes in source code with a version control system such as CVS). The Version Control System supports common Check-Out and Check-In operations.

Versioned copy A versioned copy of a payload, kernel, or image is stored in VCS.

Working copy A working copy of a payload, kernel, or image is currently present in the working area only (e.g., `$CMHOME/imaging/<user>/payloads`). Working copies are not stored in VCS.

