SGI Management Center (SMC) Installation and Configuration

Table of Contents

Preface

Audience

This guide is intended for system administrators who will install and configure the SGI Management Center software to manage and control the cluster.

Product Definition

SGI Management Center is actually a suite of products to manage your cluster:

· SGI Management Center for Altix ICE
· SGI Management Center, Standard Edition
· SGI Management Center, Premium Edition

As the name implies, SGI Management Center for Altix ICE is specific to the SGI Altix ICE platform and has a separate manual. Refer to *Related Documentation* on page iv. This manual pertains to all other supported platforms.

For more information on the product definition, see the *SGI Management Center System Administrator's Guide.*

Revision History

Revision	Date	Description
001	April 2010	Describes SGI Management Center 1.0.
002	May 2010	Describes SGI Management Center 1.1.
003	October 2010	Describes SGI Management Center 1.2.
004	January 2011	Describes SGI Management Center 1.3.

Related Documentation

The following documents provide additional information relevant to the SGI Management Center product:

- *SGI Management Center System Administrator's Guide (007-5642-xxx)*
- *SGI Management Center System Quick Start Guide (007-5672-xxx)*
- *SGI Management Center for Altix ICE (007-5718-xxx)*
- *Guide to Administration, Programming Environments, and Tools Available on SGI Altix XE Systems (007-4901-xxx)*

✎ Note

To access the IPMI guide, 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 rlogin command, type the following on a command line:

```
man rlogin
```

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
  46600 Landing Parkway
  Fremont, CA 94538

SGI values your comments and will respond to them promptly.

Chapter 1
Installing SGI Management Center

To set up SGI Management Center in your environment, you must first install the SGI Management Center Server on a Master Host. After your SGI Management Center Server is installed, you can create images to distribute the SGI Management Center 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 SGI Management Center, 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 Platform Managers

· Roamer
· IPMI
· DRAC
· ILO
· Intel Power Node Manager (IPNM)—Powered by Intel Data Center Manager

> ✎ Note
>
> When using Intelligent Platform Management Interface (IPMI), version 2.0 is 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. Roamer provides power control and console access. IPNM provides only power management.

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 SGI Management Center inoperable or impair system functionality. Technical Support is not provided for unapproved system configurations.

SGI Management Center Server

You can run SGI Management Center Server on the following operating systems and architectures:

· SUSE Linux Enterprise Server 11 (also with SP 1)
    · x86_64 hardware
· SUSE Linux Enterprise Server 10 (also with SP 1– 3)
    · x86_64 hardware
· Red Hat Enterprise Linux 6
    · x86_64 hardware
· Red Hat Enterprise Linux 5.0 – 5.5
    · x86_64 hardware
· Community Enterprise Operating System (CentOS) 5.5
    · x86_64 hardware

SGI Management Center Payload Installation

You can run the SGI Management Center Payload Installation on nodes running the following operating systems and architectures:

· SUSE Linux Enterprise Server 11 (also with SP 1)
    · x86_64 hardware
· SUSE Linux Enterprise Server 10 (also with SP 1– 3)
    · x86_64 hardware
· Red Hat Enterprise Linux 6
    · x86_64 hardware
· Red Hat Enterprise Linux 5.0 – 5.5
    · x86_64 hardware
· Community Enterprise Operating System (CentOS) 5.5
    · x86_64 hardware

SGI Management Center Client

You can install and run the SGI Management Center Client on the same operating systems and platforms supported by the SGI Management Center Server as described above.

In addition, you can install the client software on the following Windows platforms:

· Windows 7
· Windows Server 2003
· Windows Server 2008/Windows Server 2008 R2
· Windows Vista
· Windows XP

SGI Management Center 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

SGI Management Center requires the following RPM packages:

· Dynamic Host Configuration Protocol (DHCP)
  Included with your distribution.
· Trivial File Transfer Protocol (TFTP) server
  Included with your distribution. Both tftp and atftp servers are supported.
· Network Time Protocol (NTP) server
  Included with your distribution.
· Jpackage Utilities (jpackage-utils)
  Included with your distribution.
· Telnet client (telnet)
  Included with your distribution.
· IPMItool or Freeipmi
  Required if using IPMI-enabled hosts.

You must enable the DHCP server, TFTP server, NTP server, and IPMI daemon (if using OpenIPMI/ipmitool) to start at system bootup. TFTP, NTP, and IPMI should also be started.

✎ Note

If you do not enable the NTP daemon on the master host, you should set an alternate NTP server when configuring network preferences or bypass the NTP synchronization by entering 127.0.0.1 as the NTP server. See *Configure Network and Email Settings* on page 12. An incorrect NTP configuration can cause the nodes to hang during the SGI boot process.

SGI Foundation Software is required if you want to use the Memory Failure Analysis feature of Management Center. Contact your SGI representative or visit http://www.sgi.com/products/software/sfs.html.

Support for SGI Altix UV Systems

This section pertains to the SGI Altix UV 100 and SGI Altix UV 1000 systems (large-memory systems). This section does not pertain to SGI Altix UV 10 hosts, as they can be treated as general cluster nodes by SGI Management Center.

SGI Management Center support for SGI Altix UV large-memory systems includes the following functionality:

· Partition monitoring
· Provisioning
· Hierarchical tree population
· Event management

In order to support SGI Altix UV large-memory systems, SGI Management Center requires the SGI SMN bundle software to be pre-installed.

✎ Note

For UV systems (with the exception of UV 10 systems), a UV license is required for SGI Management Center. If you install a cluster/server license, some features will not work correctly, such as automatic discovery of blades and chassis management controllers (CMCs).

Configuration (Special Considerations)

The following are SGI Management Center configuration requirements for UV systems:

· A UV license is required for SGI Management Center.
· EFI boot and PXE boot settings should both be enabled in the preferences on the **General** tab. The subnet should also be set to 172.21.0.0 with a subnet mask of 255.255.0.0.
· In the **Platform Management** preferences tab, ensure that **uvcon** is selected as the default device for both power and console support.

With the SMN software running on the SMN, SGI Management Center supports automatic discovery of the CMCs, blades, and the SSI partition. The CMCs and blades are shown in the host tree in the physical view. The SSI partition is shown in the logical view.

Provisioning (Special Considerations)

When a UV license is installed, new images will be set up as UV SSI images by default. In the event that an image is not intended for a UV SSI, right-click on the image and go to **Properties**. Uncheck **UV SSI** and click **Apply**.

A UV SSI image includes a defined EFI boot partition. During the provision, */boot/EFI* will be mounted as VFAT and used for the purpose of an EFI boot. Existing contents of the EFI boot partition will be preserved. Note that SGI Management Center does not create any of the contents of */boot/EFI*.

To provision a UV SSI, right-click on the SSI partition in the logical view of the host tree and mouse over **Provisioning**. Alternatively, you can go to the **Provisioning** tab, select the VCS or working image, and select the SSI partition to provision in the host tree on the left.

Once the provisioning is in progress, you can monitor the status of the blades by right-clicking on a blade in the physical view of the SGI Management Center host tree and clicking **Connect to Console**. To monitor the status of the provisioning of the SSI, right-click on the SSI partition in the logical view and click **Connect to Console**.

Upgrading SGI Management Center

The *SGI Management Center System Administrator's Guide* describes how to upgrade your software from prior versions of the product.

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 SGI Management Center. 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 manual *Guide to Administration, Programming Environments, and Tools Available on SGI Altix XE Systems* (007-4901-xxx) or third-party documentation (in the case of third-party node types).

Configure the Multicast Route

When provisioning nodes, the default multicast configuration may not work properly. You can use the following steps to 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 Management Center is 239.192.0.128.

SLES

1. Enter the following from the command line to temporarily add the route (where *eth1* is the management interface):

    ```
    route add -net 224.0.0.0 netmask 240.0.0.0 dev eth1
    ```

2. Make the change persistent by adding the following to file */etc/sysconfig/network/routes*:

    ```
    224.0.0.0 0.0.0.0 240.0.0.0 eth1 multicast
    ```

RHEL

1. Enter the following from the command line to temporarily add the route (where *eth1* is the management interface):
    route add -net 224.0.0.0 netmask 240.0.0.0 dev eth1

2. Make the change persistent by adding the following to file */etc/sysconfig/network-scripts/route-eth1*:

    ```
    224.0.0.0/4 dev eth1
    ```

Setting the Host Name

By default, SGI Management Center uses the host name *host*. The host name alias needs to resolve to the internal network interface (for example, 10.0.10.1). If it does not resolve to an IP address or if it resolves to a loopback address (such as 127.0.0.1), then startup of the Management Center services will fail. Create an entry in the */etc/hosts* file called *host*. The following is an example:

```
10.0.10.1 admin.default.domain admin host
```

This host name can be changed by setting the *host* and *system.rna.host* values in *$MGR_HOME/@genesis.profile*.

Set Up an SGI Management Center Master Host

After you have installed a Linux distribution and other required software on supported hardware, you are ready to install SGI Management Center Server. (See *Operating System Requirements* on page 2.)

To install SGI Management Center on the master host, you can use any front end for RPM–such as YAST,Yum, the Red Hat Package Management Tool, etc. Add the SGI Management Center CDROM or iso image as an installation source and install the following packages and all dependencies:

· sgimc
· sgimc-server
· sgi-cm-agnostic (Required if you are using the Dynamic Provisioning feature with PBS Professional 10.2 or higher.)
· shout (Only needed if you are installing on an SGI Altix UV System Management Node.)

Other packages such as powerman, conman, and pdsh are provided on the media for convenience and are supported by their software manufacturers. For more information about conman, powerman, and pdsh, see https://computing.llnl.gov/linux/.

Once you have installed the SGI Management Center RPM packages on the master host, you will not be able to start the application GUI until you restart the X session on your host. Alternatively, you can source the */etc/profile.d/mgr.sh* script from the command line:

```
# . /etc/profile.d/mgr.sh
```

Licensing

In order to use SGI Management Center, you will need to obtain a license from SGI. For information about software licensing, refer to the licensing FAQ on the following webpage:

http://www.sgi.com/support/licensing/faq.html

Open the */etc/lk/keys.dat* file in a text editor. Copy and paste the license string, exactly as given, and save the file.

Using the Management Center Interface

The Management Center 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 Management Center

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

1. Open a command line console.

2. Log in as root.

3. On the command line, enter **mgrclient** and press **Enter**.

The Management Center Login is displayed.

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

The Management Center interface is displayed.
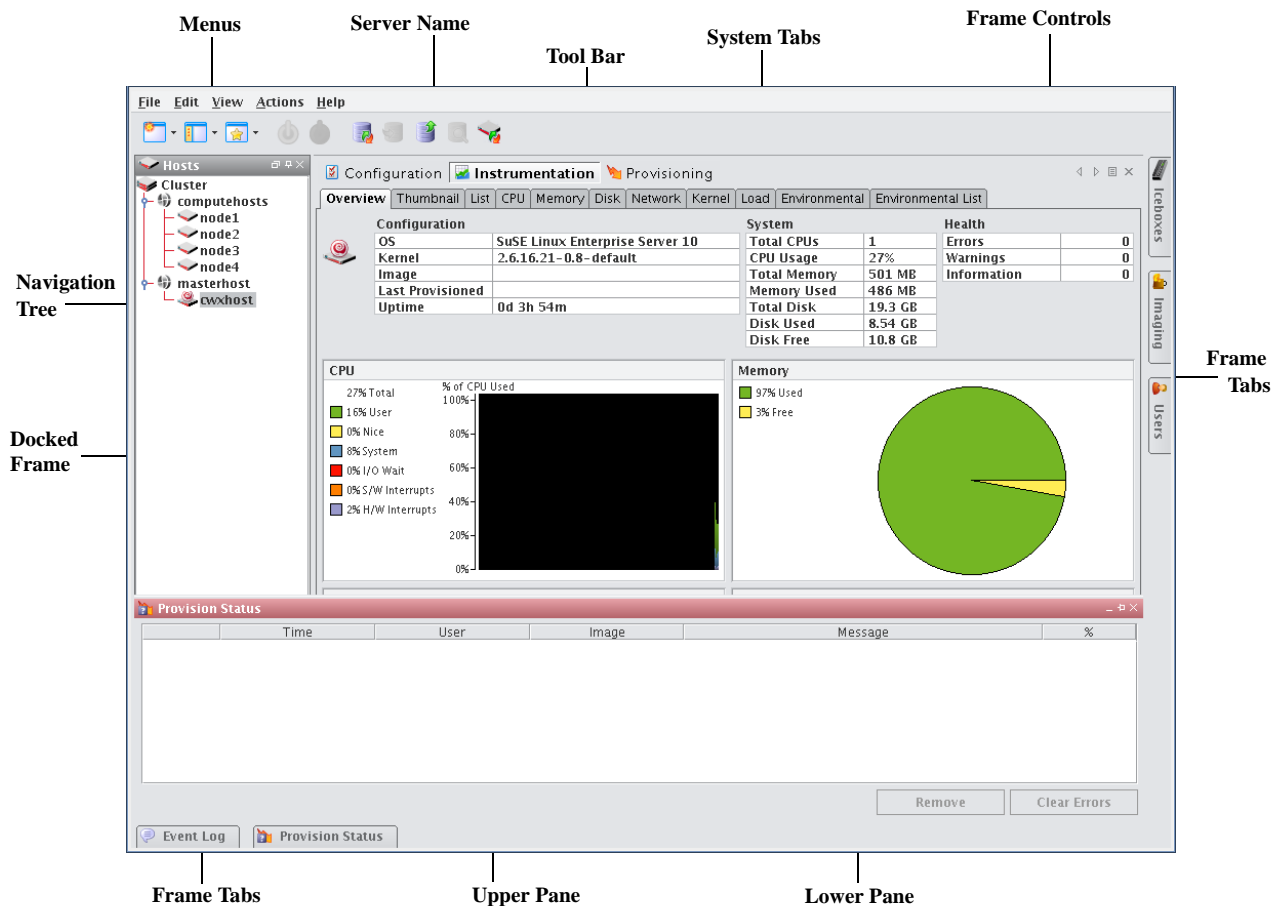
Menus — A collection of pull-down menus that provide access to system features and functionality.

Tool Bar — The tool bar provides quick access to common tasks and features.
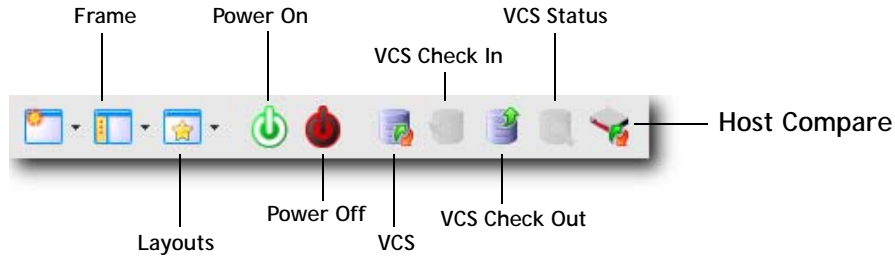
Server Name — The name of the server on which Management Center is running.

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 Management Center Server

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

Example:

```
service mgr status
/etc/init.d/mgr stop
/etc/init.d/mgr start
chkconfig --list mgr
```

Verifying Management Center Services are Running

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

> ✎ **Note**
>
> Management Center includes two distribution services for each provisioning channel pair defined in the preferences.

   · FileService.*<host name>*
   · HostAdministrationService.*<host name>*
   · IceboxAdministrationService
   · ImageAdministrationService
   · InstrumentationService
   · KernelAdministrationService
   · LogMonitoringService
   · NotificationService
   · PayloadAdministrationService
   · PayloadNodeService.*<hostname>*
   · PlatformManagementService
   · PowerMonitoringService
   · ProvisioningService
   · RNA
   · RemoteProcessService.*<hostname>*
   · SynchronizationService
   · TreeMonitoringService
   · VersionService
   · VersionService.*<host_name>*
   · com.sgi.clusterman.server.CommunicationServerFactory

Chapter 2
Configuring SGI Management Center

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 ="eth*x*".

4. Replace "*x*" with the host interface you use.
    Example:

    ```
    DHCPD_INTERFACE="eth1"
    ```

5. Save and close the file.

Configure DHCP Settings

When provisioning occurs, Management Center automatically modifies DHCP settings and restarts the service. If you make manual DHCP modifications and want Management Center 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 subnet 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 Management Center 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 subnet and netmask are mandatory. All other fields are optional.
   - **Base Network Subnet** — The private network used by the cluster (typically a *192.168.x.x* or *10.x.x.x* network). To set the subnet, 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, Management Center is configured to boot using PXE. The default PXE boot configuration utilizes a 3-stage boot process and supports the e1000, e1000e, bnx2, tgz, and r8169 drivers to load the X-SLAM protocol, which uses scalable multicast to provision nodes.

   Management Center also supports booting with zpxe-formatted ROM files using tftp to load the X-SLAM client. Additional zpxe-formatted ROM files which support some common node configurations are installed in the */tftpboot/mgr* directory. Open the preferences dialog from the Edit menu and browse to locate the desired file to use for tftp boot. If the file is located in a different directory, Management Center copies the file to the */tftpboot/mgr* directory.

   To reset Management Center to the default configuration, choose the file *pxelinux.0* in the */tftpboot/mgr* directory. Nodes that are configured with Etherboot client are also supported and will boot without using TFTP.

   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

Management Center 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

Management Center 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 Management Center accounts, and this information cannot be imported into Management Center.

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

Configure TFTP

Management Center places boot files in */tftpboot/mgr*. The tftp or atftp daemon must use */tftpboot* as the home for tftp boot files. If you are using the tftp package that is included with the RHEL 6 distribution, you must change the parameter server_args in */etc/xinetd.d/tftp* to use */tftpboot* as the home for tftp boot files.

Example */etc/xinetd.d/tftp* file:

```
service tftp
{
disable         = no
socket_type     = dgram
protocol        = udp
wait            = yes
user            = root
server          = /usr/sbin/in.tftpd
server_args     = -s /tftpboot
per_source      = 11
cps             = 100 2
flags           = IPv4
}
```

Default User Administration Settings

Management Center creates the following default users, groups, and roles during installation:

After installation, Management Center 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 Management Center 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 Management Center 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, Management Center 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, Management Center 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 Management Center.

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 aNew 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 Management Center 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 Management Center. Changes take effect only after the users close Management Center 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. Management Center 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.

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

Module	Name	Description
Management Center	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 use the following grouping mechanisms:

· *Partitions*
  Partitions include a strict set of hosts that may not be shared with other partitions.
· *Regions*
  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.
· *Racks*
  You can use racks to represent the physical layout of your cluster.

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.

🔍 Tip

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.
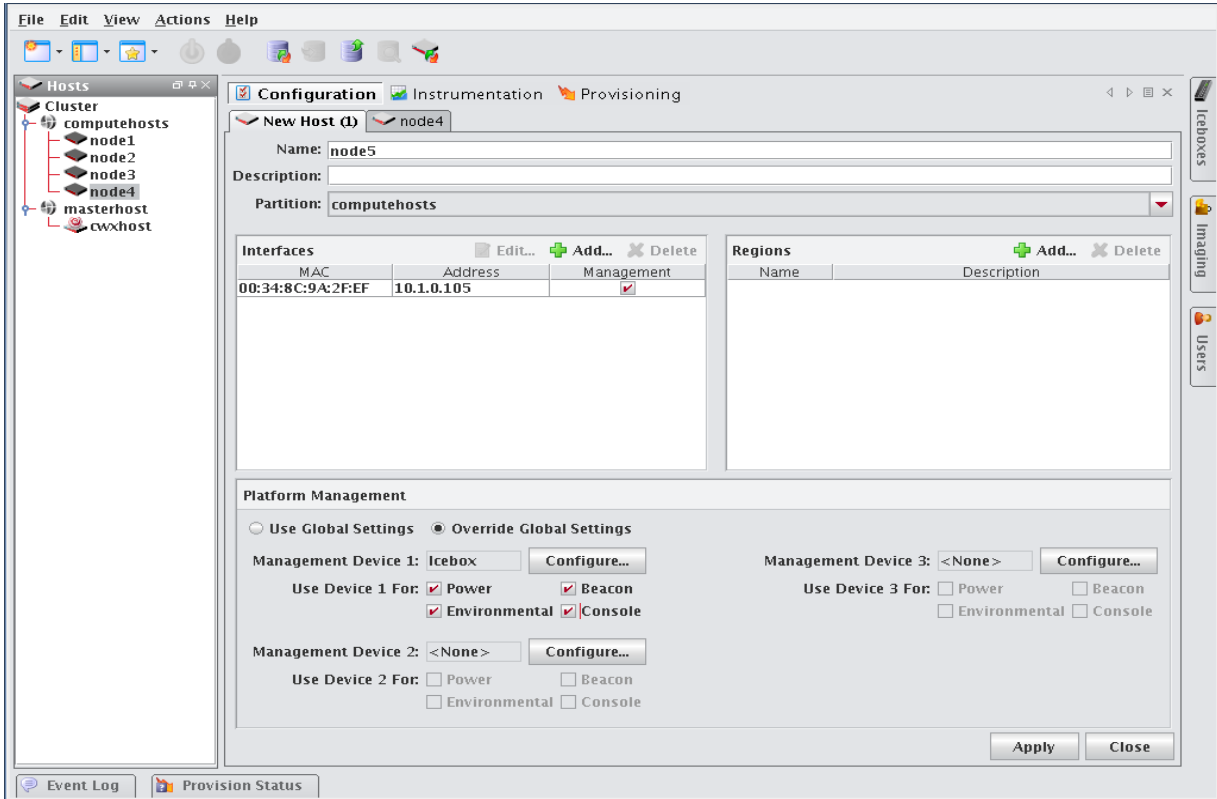
3. Enter the host name.

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 EEPRO100]
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 EEPRO100 ROM address 0x0000
[EEPRO100]Ethernet addr: 00:02:B3:11:03:77

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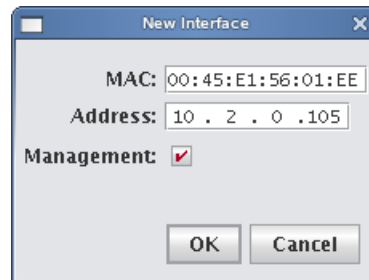
(*If conman is set up and working, this information is also contained in the conman log file for the host—typically located in */var/log/conman/console.n[1-x]*)

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 Management Center interface to manage the host.

    Management Center stores the interface and automatically writes it to *dhcp.conf*.

4. Add any additional interfaces required for this host.

    Management Center 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. Management Center 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.

Intel Power Node Manager (IPNM)

Management Center supports power management through the utilization of the Intel Power Node Manager technology. Enabling this functionality requires special hardware for your system and an appropriate SGI license key to activate this feature. If this feature was not pre-configured for your cluster in Manufacturing, contact SGI for information about enabling this feature.

The Management Center Instrumentation Power Panel uses indirect TCP communications to an external web service provider to accumulate and display the IPNM power monitoring data. This results in a delay in updating the power panel for every tree selection change. Consequently, a several-minute delay is possible for large-scale systems using IPNM.

IPMI, Roamer, ILO, and DRAC

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. Roamer supports power control and console access.

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 *dhcpd.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 OR ROAMER SETTINGS ON YOUR HOST

1. Select **Override Global Settings**.

2. Select **IPMI** or **Roamer** 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.

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

Management Center 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 Management Center imports only host information. Management Center 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 Management Center 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
```
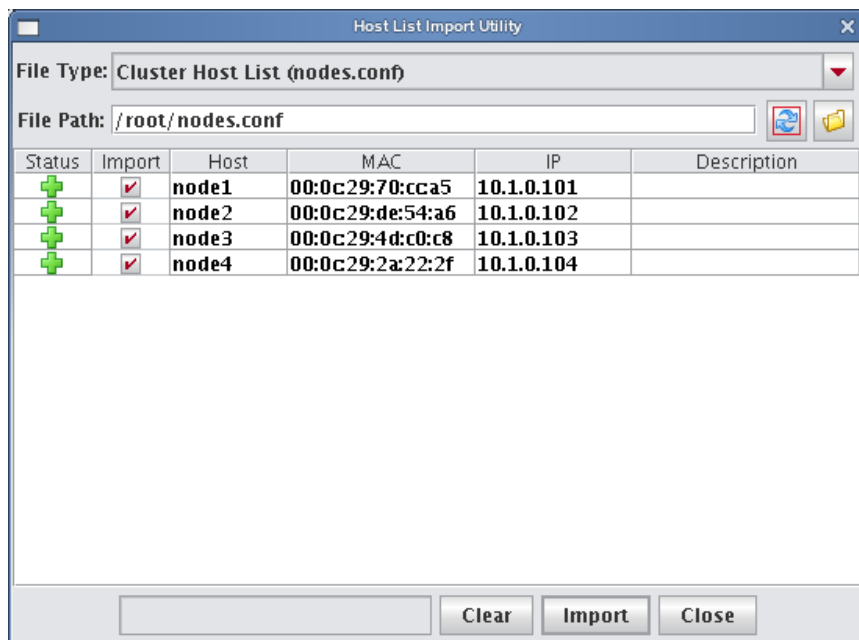
2. Select **Import Host List** from the **File** menu.

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**.

Partitions and Regions

Partitions are used to separate clusters into non-overlapping collections of hosts. Instrumentation, provisoning, power control, and other administration tasks can be performed on this collection of hosts by selecting the partition in the host tree.

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.

10. Click **OK**.

11. Click **Apply**.

Racks

To aid in the management of the cluster, you can use racks to represent the physical layout of the cluster into non-overlapping collections of hosts. If you have hosts which are not assigned a rack, they will appear in a rack labelled **Unassigned**.

Add a Rack

1. Right-click in the **Hosts** navigation tree and select **New Rack** or select **New Rack** from the **File** menu.

2. Enter a rack name.

3. (Optional) Enter a description.

4. In the **Hosts** pane, click **Add** to display the Select Hosts dialog.

5. Select hosts to add to this rack and click **OK**.

6. Click **Apply**.

Edit a Rack

Editing a rack allows you to change previously saved information about a rack. You can edit rack information, alter rack configurations, or remove racks.

1. Select a rack from the host navigation tree.

2. Select **Edit** from the **Edit** menu or right-click on the racks in the host navigation tree and select **Edit**.

3. Use the rack pane to make changes to the rack.

4. Click **Apply** to accept the changes or **Close** to abort this action.

Delete a Rack

✏ **Note**

If you delete a rack, all hosts associated with the rack will be moved to rack **Unassigned**.

1. Select the rack(s) you want to delete from the host navigation tree.

2. Select **Delete** from the **Edit** menu or right-click on the rack(s) in the navigation tree and select **Delete**.

3. Click **OK** to delete the rack(s).

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 32 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 Management Center Version Control System (VCS). Package information and files are stored and may be browsed through Management Center.

> **! 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 Management Center 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**.
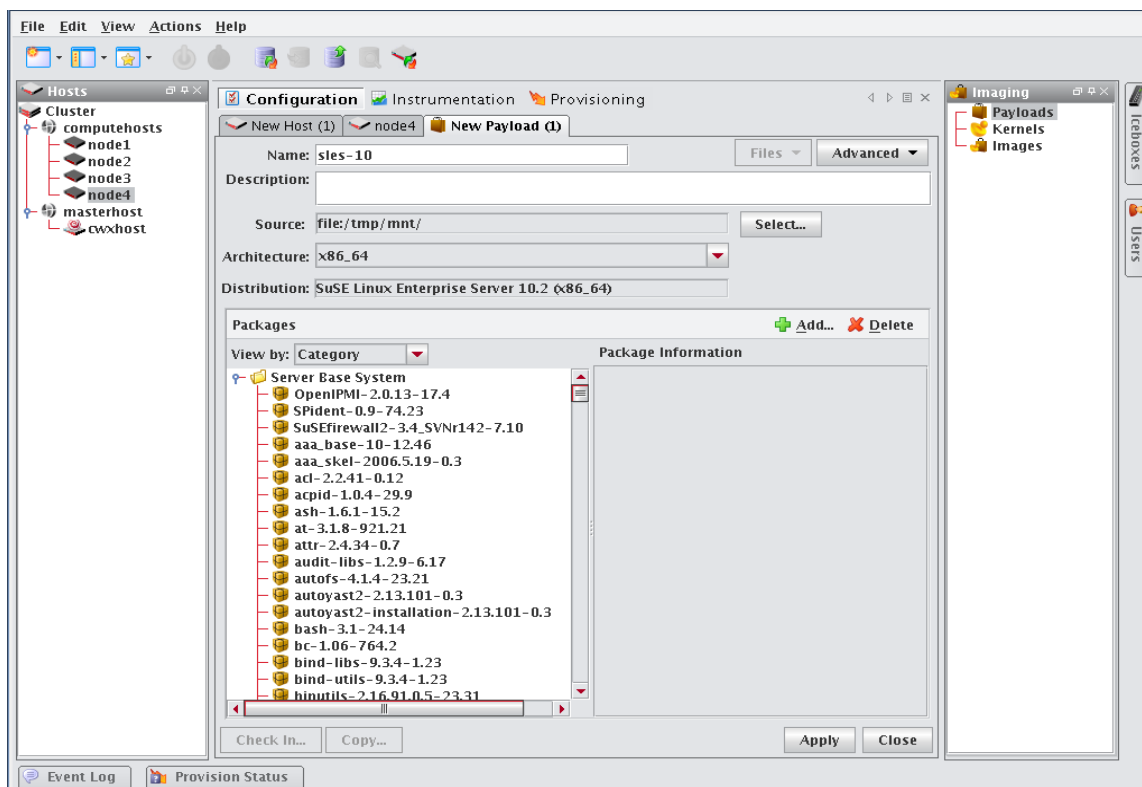

A new payload pane appears.

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 32 and *SUSE Linux Enterprise Server Installations* on page 32 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.
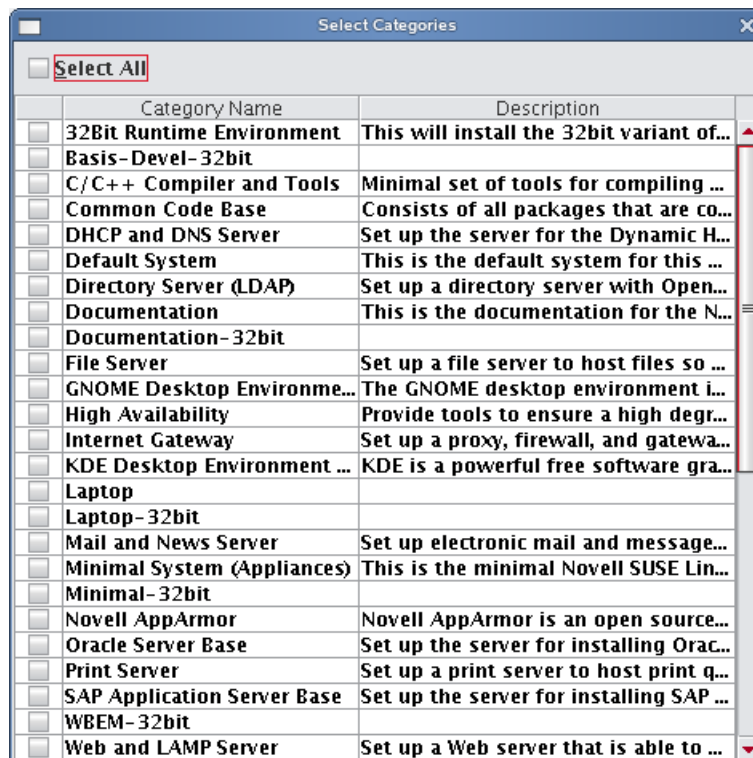

🔍 **Tip**

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

✏️ **Note**

If Management Center 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 Management Center 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.

Select Categories		×
☐ **Select All**		
Category Name	Description	
☐ **32Bit Runtime Environment**	**This will install the 32bit variant of...**	▲
☐ **Basis–Devel–32bit**		
☐ **C/C++ Compiler and Tools**	**Minimal set of tools for compiling ...**	
☐ **Common Code Base**	**Consists of all packages that are co...**	
☐ **DHCP and DNS Server**	**Set up the server for the Dynamic H...**	
☐ **Default System**	**This is the default system for this ...**	
☐ **Directory Server (LDAP)**	**Set up a directory server with Open...**	
☐ **Documentation**	**This is the documentation for the N...**	☰
☐ **Documentation–32bit**		
☐ **File Server**	**Set up a file server to host files so ...**	
☐ **GNOME Desktop Environme...**	**The GNOME desktop environment i...**	
☐ **High Availability**	**Provide tools to ensure a high degr...**	
☐ **Internet Gateway**	**Set up a proxy, firewall, and gatewa...**	
☐ **KDE Desktop Environment ...**	**KDE is a powerful free software gra...**	
☐ **Laptop**		
☐ **Laptop–32bit**		
☐ **Mail and News Server**	**Set up electronic mail and message...**	
☐ **Minimal System (Appliances)**	**This is the minimal Novell SUSE Lin...**	
☐ **Minimal–32bit**		
☐ **Novell AppArmor**	**Novell AppArmor is an open source...**	
☐ **Oracle Server Base**	**Set up the server for installing Orac...**	
☐ **Print Server**	**Set up a print server to host print q...**	
☐ **SAP Application Server Base**	**Set up the server for installing SAP ...**	
☐ **WBEM–32bit**		
☐ **Web and LAMP Server**	**Set up a Web server that is able to ...**	▼

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, Management Center 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, Management Center 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 Management Center into the Payload

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

To install via the SGI Management Center GUI, do the following:

1. Open the **Imaging** frame.

2. Double-click your payload to open it.

3. Click **Add** (to add aditional packages).

4. On the Management Center media, browse to the *sgi/x86_64* directory (if using SLES) or the *RPMS* directory (if using RHEL) and select the following packages:

·   sgimc-payload
·   java-1.6.0-sun

You can also install into the payload from the command line using the RPM "root" parameter. For example:

```
# cd /mnt/cdrom/sgi/x86_64
# rpm -ivh --root=$MGR_HOME/imaging/root/payloads/Compute java-1.6.0-sun-1.6.0.17-
sgi700c1.sles11.x86_64.rpm
# rpm -ivh --root=$MGR_HOME/imaging/root/payloads/Compute sgimc-payload-1.0.0-
sgi700c1.sles11.x86_64.rpm
```

Installation on a Running Altix UV SSI or Cluster Compute Node

The sgimc-payload package can be installed on a running SGI Altix UV SSI or cluster compute node. After installing the sgimc-payload RPM package, you should run the following script:

`/opt/sgi/sgimc/bin/configure_payload.sh`

The following is an example of the configure_payload.sh script with user entries shown in bold:

`configure_payload.sh`

```
This script will configure the SGI Management Center Payload Daemon service to work
correctly with the SGI Management Center server. You will need the following pieces of
information to enable this script to successfully complete the setup.


1. The name of the server where the SGI Managment Center server is running. This name
should match what the server is calling itself in the file /opt/sgi/sgimc/
@genesis.profile. The entry is listed as system.rna.host.
2. The IP address of the server where the SGI Management Center Server is running. This
should be the IP address for the management network.
3. The name that this host is known by in the database of SGI Management Center. This does
not have to match the actual hostname that this system knows itself as.
4. The IP address of this system on the management network. This IP address should also
be present in the SGI Management Center database.


Note: the two host names should be simple host names, not fully qualified domain names.
Please enter the name of the server: host
Please enter the IP address of the server: 172.21.0.1
Please enter the name of this host from the server database: UV00000014-P000
Please enter the IP address of this host: 172.21.1.0
You entered the following:
Server Name = host
Server IP = 172.21.0.1
Host Name = UV00000014-P000
Host IP = 172.21.1.0
The SGI Management Center Payload Daemon service has been configured. You should restart
the Name Service Cache Daemon before attempting to start the payload daemon service.
Execute the following commands to restart the Name Service Cache Daemon and the SGI
Management Center Payload Daemon service:
/etc/init.d/nscd restart
/etc/init.d/mgr restart
```

Note that the script prompts for the host name as defined in the server database–that is, the host name from the host tree in the SGI Management Center GUI. If you are installing on a running UV SSI, use the name of the partition (usually of the format UV*xxxxxxxx*-P*yyy*). If you are installing on a non-UV, generic cluster node, it will be the name of the node in the host tree (for example, n001). If the node is not present in the tree, you must add an entry for it.

Kernel Management

Kernels may be customized for particular applications and used on specific hosts to achieve optimal system performance. Management Center 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 40.

1. Select **New Kernel** from the **File** menu or right-click in the imaging navigation tree and select **New Kernel**.
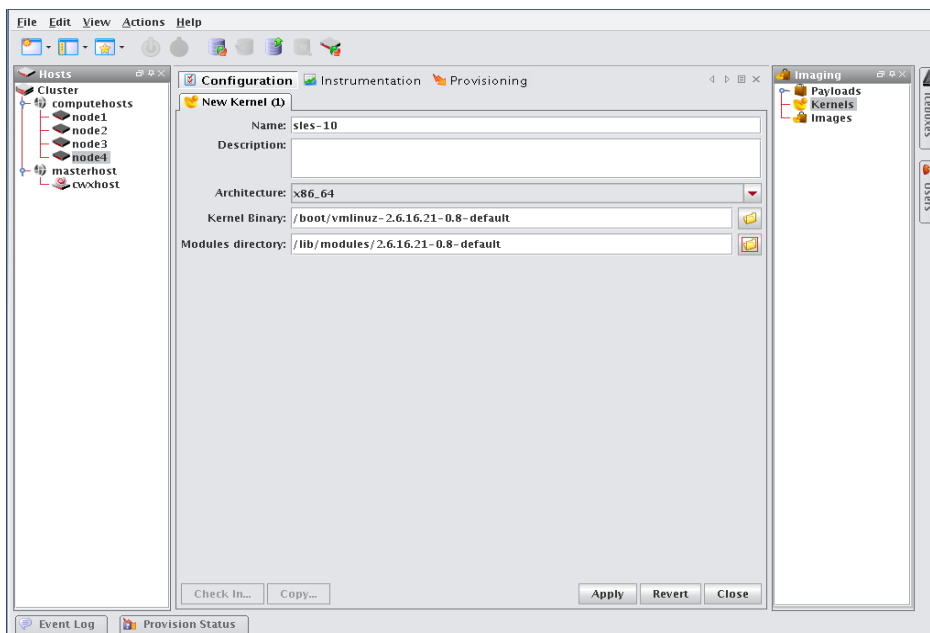
2. Enter the kernel name.
3. (Optional) Enter a kernel description.
4. Select the hardware architecture.
5. Enter the full path to the kernel binary you want to use, or click **Browse** to find the kernel binary you want to use.

! Warning!

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

---
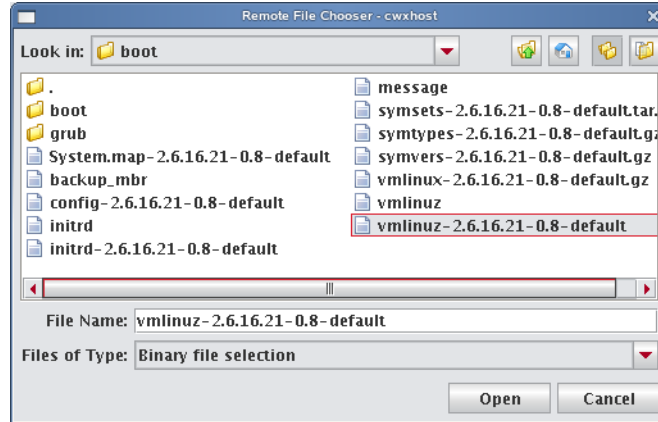
6. Click **Browse** to find the modules directory and click **Open**.

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 38. 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 Management Center 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, Management Center 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 the 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 Management Center 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**.
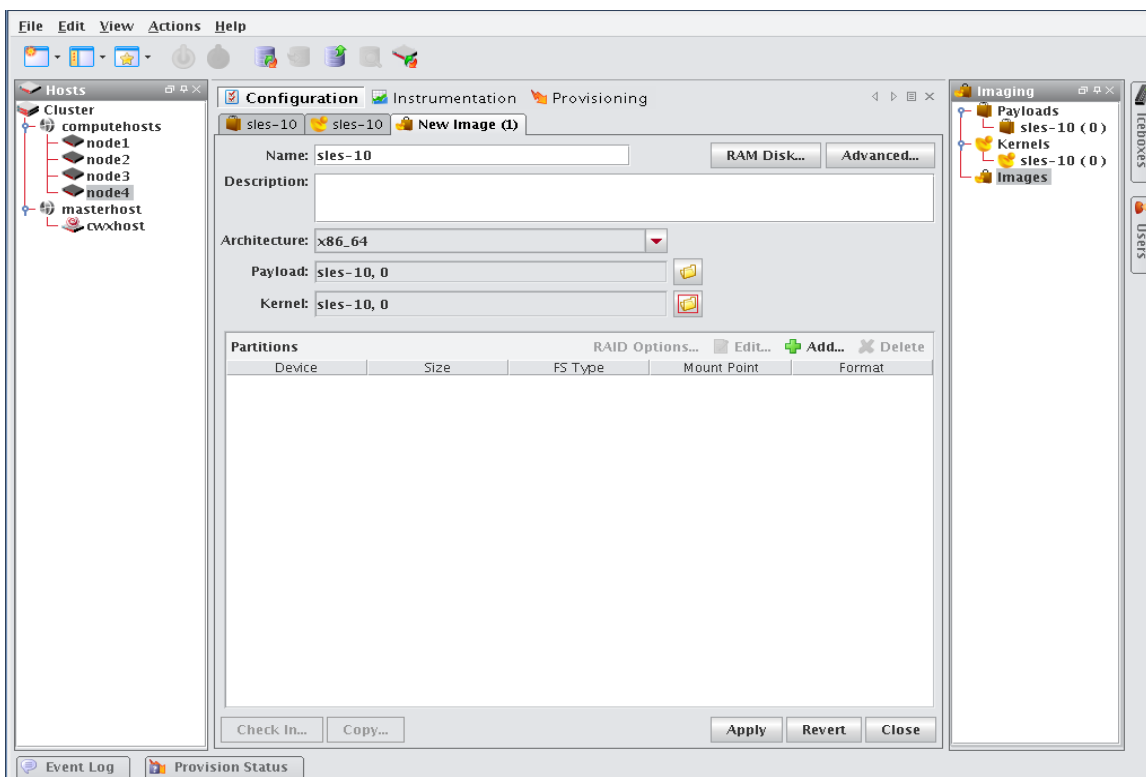
2. Enter a name in the **Name** field.

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 44.

✎ **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 44.

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 44.

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

Management Center 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 *Advanced Imaging Options* on page 42.
image.path	Specifies where to store the downloaded image. By default, */mnt*.

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 46.

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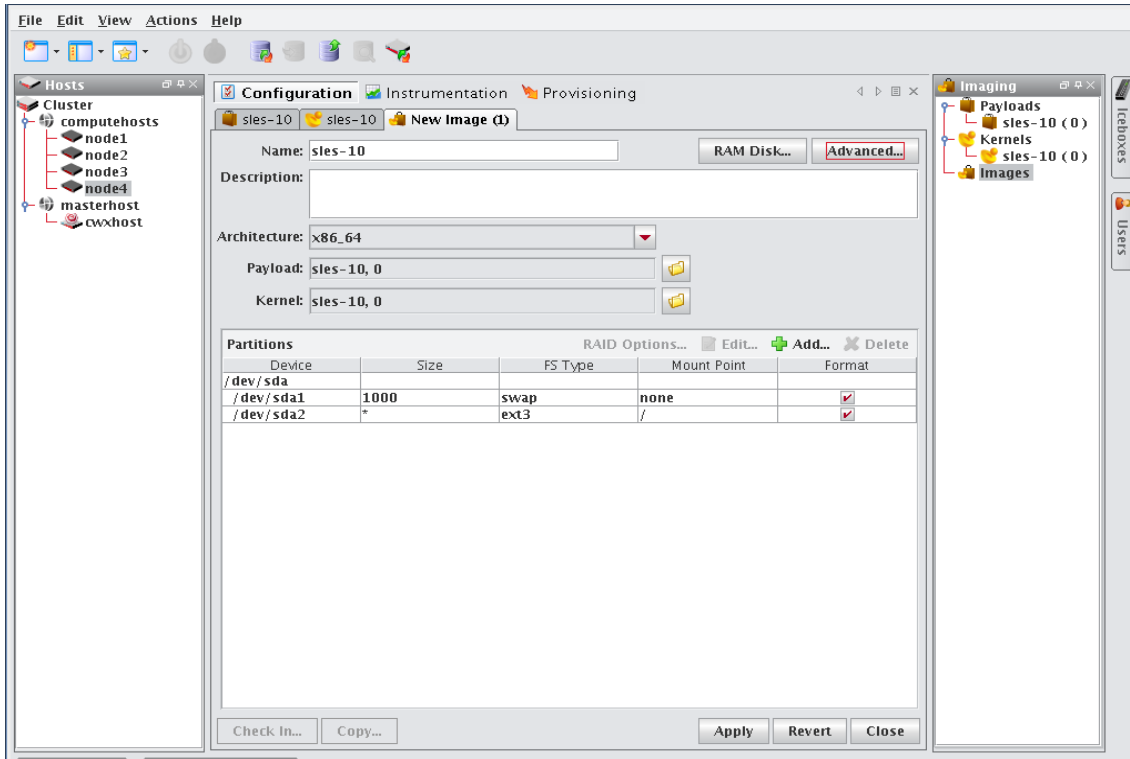
2. In the partitions pane, click **Add**.

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, *$MGR_HOME/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 Management Center.

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 Management Center 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 46 and *boot.profile* on page 43).

✏️ 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

Management Center 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 43.

Provisioning

The Management Center 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**. Management Center 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.

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, *$MGR_HOME/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, *$MGR_HOME/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.

Management Center asks you to confirm your action.

7. Click **Yes** to provision the hosts.

When you click **Yes**, Management Center 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**.

   Management Center 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.

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 Management Center Master Host, compute hosts, Management Center, UPS, high-speed network, storage, and the cabinet.

Management Center Master Host The Management Center 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 Management Center 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 Management Center 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. Management Center 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 Management Center 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., *$MGR_HOME/imaging/<user>/ payloads*). Working copies are not stored in VCS.

Index

F

feedback, documentation v
format partition 45
frames
    controls 8
    dockable 8
Freeipmi 3
fstab 45

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 28
    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 5
    provision 46
    region
        add host to 27
        assign host to 22
    shared 19

I

Icebox
    administration privileges 18
ILO 1, 22, 24
image 43

create 41
management 41
partition 44
privileges, enable imaging 18
provision 46
image.once 43
image.path 43
import
    host list 25
install
    Management Center 6
        into payload 36
Intel Power Node Manager (IPNM) 1, 2, 22
interface
    management 21
IP address 50
    host 21
IPMI 1, 5, 22
IPMItool 3
ipmitool utility 5
IPNM 1, 2, 22

J

Jpackage Utilities 3

K

kernel
    build from source 38
    copy 39
    create 38
    management 38
    upgrade 2

L

licensing 6

M

management
    interface 21
Management Center
    administration
        grant privileges 18
    install into payload 36
    platforms 1
    Premium Edition iii