SGI Management Center (SMC) System Administrator’s Guide

007-5642-006

COPYRIGHT

© 2010, 2011 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 Silicon Graphics, 46600 Landing Parkway, Fremont, CA 94538.

TRADEMARKS AND ATTRIBUTIONS

Silicon Graphics, SGI, the SGI logo, SGI Prism, and Altix are trademarks or registered trademarks of Silicon Graphics International Corp. or its subsidiaries 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. NVIDIA is a registered trademark of NVIDIA Corporation in the United States and/or other countries. PBS Professional is a trademark of Altair Grid Technologies, a subsidiary of Altair Engineering, Inc. 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, Windows Server, and Windows Vista 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 .... **ix**
Product Definition ix
Audience ix
Revision History ix
Related Documentation x
Annotations xi
Product Support xii
Reader Comments xii

Chapter 1
Getting Started **1**
System Requirements 1
Minimum Hardware Requirements 1
Operating System Requirements 2
Software Requirements 3
Setting the Host Name 4
Set Up an SGI Management Center Master Host 4
Server Installation 4
Client Installation 8
Advanced Scale-Out Configuration 10
Licensing 12
Importing Existing Hosts 12
Starting and Stopping the SGI Management Center Server 12
Verifying SGI Management Center Services Are Running 13

Chapter 2
Introduction to SGI Management Center **15**
Overview 15
Product Definition 15

Comprehensive System Monitoring 16
Version Controlled Image Management 17
Fast Multicast Provisioning 17
Memory Failure Analysis 17
Support for SGI Altix UV Systems 17
Support of SGI Prism XL Systems 19
Auto Node Discovery 19
Using the Management Center Interface 20
Starting Management Center 20
Customizing the Interface 23
Customizing System Tabs 23
Dockable Frames 24
Layouts 24
Chapter 3
Preferences and Settings 27
Preferences 27
General 27
Configure Network and Email Settings 28
Platform Management 29
Applications 38
Provisioning Settings 39
Configuring IPMI 41
Configure the IPMI BMC 41
Configure the ipmitool_options.profile 42
Configure the Payload and Kernel 42
Configure the Master Host and Management Center 43
Configuring DHCP 44
Configure DHCP Settings 44
Configure Multicast Routes 45
Configure TFTP 45
Chapter 4
Cluster Configuration 47
Clustered Environments 47
Setting Up Your Cluster 48
Adding Hosts 49
Configure Platform Management 51
Edit a Host 54
Find a Host 55
Delete a Host 55
Import Hosts 55
Host Power Controls 58
Console 60
Partitions 61
Adding Partitions 61
Editing Partitions 62
Deleting Partitions 62

Regions 63
  Creating Regions 63
  Editing Regions 64
  Deleting Regions 65
Racks 66
  Adding Racks 67
  Editing Racks 67
  Deleting Racks 67

Chapter 5

User Administration 69
Default User Administration Settings 70
  Adding a User 70
  Editing User Accounts 72
  Disabling a User Account 72
  Deleting a User Account 72
Groups 73
  Adding a Group 73
  Editing a Group 75
  Deleting a Group 75
Roles 76
  Adding a Role 76
  Editing a Role 78
  Deleting Roles 78
  Privileges 79

Chapter 6

Imaging, Version Control, and Provisioning 81
Overview 81
Payload Management 82
  Configuring a Payload Source 82
  Creating a Payload 84
  Importing Kernel Parameters from a Running Host 89
  Adding a Package to an Existing Payload 90
  Remove a Payload Package 93
  Payload File Configuration 95
  Payload Authentication Management 96
  Payload Local User and Group Account Management 98
  Add and Update Payload Files or Directories 102
  Edit a Payload File with the Text Editor 103
  Delete Payload Files 104
  Delete a Payload 104
  Install Management Center into the Payload 105
  Installation on a Running Altix UV SSI or Cluster Compute Node 106
Kernel Management 107
  Create a Kernel 107
  Edit a Kernel 113
  Delete a Kernel 115

Image Management 116
  Create an Image 116
  Delete an Image 119
  Managing Partitions 120
  RAID Partitions 123
  Edit a Partition 125
  Delete a Partition 127
  User-Defined File Systems.... 128
  Diskless Hosts.... 131
  RAM Disk.... 134
  Plug-ins for the Boot Process 136
Version Control System (VCS) 140
  Version Control.... 140
  Version Branching 141
  Version Control Check-in 142
  Version Control Check-out.... 143
  VCS Management 143
  VCS Host Compare.... 145
Provisioning 147
  Select an Image and Provision.... 147
  VCS Upgrade 150
  Advanced Provisioning Options 151

Chapter 7 ....
Instrumentation and Events 153
Instrumentation 153
  States 154
  Event Log.... 154
  Menu Controls 155
  Overview Tab.... 156
  Thumbnail Tab.... 157
  List Tab.... 158
  CPU Tab 159
  Memory Tab.... 160
  Disk Tab 161
  Network Tab 162
  Kernel Tab 163
  Load Tab.... 164
  Environmental Tab.... 165
  Environmental List Tab.... 166
  GPU Tab 167
  Power Tab.... 168
Failure Analysis 171
Management Center Monitoring and Event Subsystem 172
  Monitors.... 173
  Custom Monitors.... 181
  Metrics.... 185
  Event Listeners 189
  Loggers 198

Chapter 8

Upgrading SGI Management Center 201
General Tasks 201
Upgrading from a Previous Version of SGI Management Center 202
Upgrading from SGI ISLE Cluster Manager 2.x 203
Additional Upgrade Notes 205
Backing Up Provisioned Nodes........ 205
Changing the Default RNA Host Name........ 205
Deleting Old Administration Node from Host Tree........ 205

Chapter 9

Using the Discover Interface 207
Software Requirements 207
The Graphical Interface 208
The Command-Line Interface 212

Chapter 10

Troubleshooting 213
Debug Logs 214
Support Information Tool 214
Startup Daemon Fails on the Master Host 214
Nodes in Provisioning or Unknown State after Provisioning 215
Temperatures and Fan Speeds Not Registering 215
Inordinately High CPU Usage on Head Node 215
Insufficient Number of Provisioning Channels 215
Kernel Modules Not Loading on Compute Nodes 216
Command-line Boot Parameters Not Honored 216
Payload Check-in Error 216
Invalid or Expired License Message 217
Resource Usage Too High on Head Node 217
Altix UV Provisioning Stops While Loading Kernel 218

Chapter 11

Command-Line Interface 219
Command-Line Syntax and Conventions 219
CLI Commands 220
conman 226
cwhost 229
cwpower 238
cwprovision 240
cwuser 243

dbix 249
dbx 250
imgr 251
kmgr 252
pdcp 253
pdsh 255
pmgr 258
powerman 259
vcs 261

Appendix 265

Pre-configured Metrics 265
CPU 265
Disk 266
Kernel 266
Load 267
Memory 267
Network 268

Glossary 269

Index 273

Preface

The *SGI Management Center System Administrator's Guide* is written in modular style where each section builds upon another to deliver progressively advanced scenarios and configurations. Depending on your system configuration and implementation, certain sections of this guide may be optional, but warrant your attention as the needs of your system evolve. This guide assumes that you, the reader, have a working knowledge of Linux.

Product Definition

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

- SGI Management Center, Standard Edition
- SGI Management Center, Premium Cluster Edition
- SGI Management Center, Premium ICE Edition
- SGI Management Center, Premium UV Edition

As the name implies, SGI Management Center Premium ICE Edition is specific to the SGI Altix ICE platform and has a separate manual. Refer to *Related Documentation* on page x. This manual pertains to all other supported platforms. For more information about the product editions and features, see *Product Definition* on page 15.

Audience

This guide's intended audience is the system administrator who will be working with the SGI Management Center software to manage and control the cluster.

Revision History

Revision	Date	Description
001	April 2010	Supports SGI Management Center 1.0.
002	May 2010	Supports SGI Management Center 1.1.
003	October 2010	Supports SGI Management Center 1.2.
004	January 2011	Supports SGI Management Center 1.3.
005	July 2011	Supports SGI Management Center 1.4.
006	October 2011	Supports SGI Management Center 1.5.

Related Documentation

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

- *SGI Management Center (SMC) Installation and Configuration Guide (007-5643-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
Getting Started

To set up SGI Management Center in your environment, you must first install 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 (DCM)

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/DCM 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– 4)
  - x86_64 hardware
- Red Hat Enterprise Linux 6.0 and 6.1
  - x86_64 hardware
- Red Hat Enterprise Linux 5.0 – 5.7
  - x86_64 hardware
- Community Enterprise Operating System (CentOS) 6.0
  - x86_64 hardware
- Community Enterprise Operating System (CentOS) 5.0 – 5.7
  - 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– 4)
  - x86_64 hardware
- Red Hat Enterprise Linux 6.0 and 6.1
  - x86_64 hardware
- Red Hat Enterprise Linux 5.0 – 5.7
  - x86_64 hardware
- Community Enterprise Operating System (CentOS) 6.0
  - x86_64 hardware
- Community Enterprise Operating System (CentOS) 5.0 – 5.7
  - 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 earlier.

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 28. 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.

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

Setting the Host Name

By default, SGI Management Center uses the host name `admin`. 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 SGI Management Center services will fail. Create an entry in the */etc/hosts* file called `admin`. The following is an example:

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

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.) Ensure that your host name is set properly. (See the preceding section *Setting the Host Name* on page 4.)

Server Installation

Using the following topics, this section describes server installation instructions for various Linux distributions.

- *RPM Packages* on page 5
- *RPM Front Ends* on page 5
- *Pre-Installation Steps* on page 5
- *Installing the Server* on page 8

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
```

Note

By default, the SGI Management Center password is *root*. For information on how to change this password, see *Editing User Accounts* on page 72. When you provision a host, SGI Management Center sets up a root account for your hosts.

! Warning!

If the management network is something other than 10.0.0.0 following an installation or upgrade, you need to log in as root and update it in SGI Management Center preferences. See *Preferences* on page 27.

RPM Packages

To install SGI Management Center on the master host, you need to install the `sgimc` package. This package and notable dependent packages follow:

- `sgimc`
- `sgimc-server`
- `sgi-cm-agnostic` (Required if you are using the Dynamic Provisioning feature with PBS Professional 10.2 or higher.)
- `shout` (Required for UV systems excluding UV10s. Install with YUM or Zypper.)
- Discover packages (Required for the Discover interface on non-UV systems.)

The Discover packages include the following:

- `discover`
- `discover-common`
- `discover-server`
- `sgi-common-python`
- `sgi-management-device`
- `cattr`

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

RPM Front Ends

To install SGI Management Center, you can use any front end for RPM. The suggested instructions in this section use YUM and Zypper.

Pre-Installation Steps

Do the following to prepare for installation:

1. Set up package repositories for the Linux distribution you are using.
   Consult the documentation for your Linux distribution for details on how to set up the package repositories for your distribution.
2. Mount the DVD or ISO for SGI Management Center.

   Example for DVD:

   ```
   # mount -o loop /dev/sdc /media/CDROM
   ```

   Example for ISO:

   ```
   # mount -o loop sgimc-1.5-cd1-media-rhel6-x86_64.iso /mnt/iso3
   ```

   Substitute the correct ISO name where necessary—for example, `sgimc-1.5-cd1-media-sles11-x86_64.iso`.

Note

Installation on Red Hat Linux will only succeed if you have the Red Hat distribution packages set up for installation via YUM. This should be the case already if you have registered for updates with the Red Hat Network (RHN). Generally, CentOS Linux also will be set up for installing distribution packages via YUM. Otherwise, please consult the YUM documentation to see how to set up YUM to install distribution packages.

PREPARING TO INSTALL ON RED HAT ENTERPRISE LINUX 6.X AND CENTOS 6.X

The following instructions use the YUM front end for RPM.

1. Change directories to the location where the ISO is mounted.

   The following example requires the ISO to be mounted to */mnt/iso3*:

   ```
   # cd /mnt/iso3
   ```

2. Run the provided script to create the YUM configuration file:

   ```
   # ./create-yum-config-file
   ```

   A temporary YUM repository file will be created in */tmp/yum-sgi.conf*.

3. If you want the new YUM repository to be permanent, copy it to */etc/yum.repos.d* and install `sgimc`.

   Example:

   ```
   # cp /tmp/yum-sgi.conf /etc/yum.repos.d/yum-sgi.repo
   ```

Note

The `create-yum-config-file` script supports configuring repository data for multiple SGI products in one repository file. If you plan to run the script additional times for other SGI products such as SGI Foundation Software or SGI Performance Suite, do this before copying the file.

4. If you do not want the YUM repository to be permanent, you can install `sgimc` from a temporary file:

   ```
   # yum -c /tmp/yum-sgi.conf install sgimc
   ```

   YUM will resolve all dependencies and ask for confirmation that you wish to install the packages.

PREPARING TO INSTALL ON RED HAT ENTERPRISE LINUX 5.X AND CENTOS 5.X

Create a YUM repository in */etc/yum.repos.d*.

Example:

```
# /etc/yum.repos.d/yum-sgi.repo:

[SGI-Management-Center-1.5]
name=SGI-Management-Center-1.5
baseurl=file:///mnt/iso3
enabled=1
gpgkey=file:///mnt/iso3/RPM-GPG-KEY-sgi
gpgcheck=1
```

PREPARING TO INSTALL ON SUSE LINUX ENTERPRISE SERVER 11.X

Do the following:

1. Ensure that a Zypper repository exists for Linux distribution packages.

   Management Center installation will fail if you do not have distribution packages set up to install via Zypper. To show which repositories are available, use the following command:

   ```
   # zypper lr
   ```

   If you do not already have distribution packages set up to install via Zypper, create such a repository, as shown in the following example. The following example assumes the SLES 11 SP1 ISO is mounted at */mnt/iso1*:

   ```
   # zypper ar dir:///mnt/iso1 'SUSE-Linux-Enterprise-Server-11-SP1'
   ```

2. Add a repository for SGI Management Center.

   The following example requires the Management Center ISO to be mounted at */mnt/iso3*):

   ```
   # zypper ar dir:///mnt/iso3 'SGI Management Center 1.5'
   ```

PREPARING TO INSTALL ON SUSE LINUX ENTERPRISE SERVER 10.X

Do the following:

1. Ensure that a Zypper repository exists for Linux distribution packages.

   Management Center installation will fail if you do not have distribution packages set up to install via Zypper. To show which repositories are available, use the following command:

   ```
   # zypper sl
   ```

   If you do not already have distribution packages set up to install via Zypper, create such a repository, as shown in the following example. The following example assumes the SLES 11 SP1 ISO is mounted at */mnt/iso1*:

   ```
   # zypper sa dir:///mnt/iso3 'SUSE-Linux-Enterprise-Server-11-SP1'
   ```

2. Add a repository for SGI Management Center.

   The following example requires the Management Center ISO to be mounted at */mnt/iso3*):

   ```
   # zypper sa dir:///mnt/iso3 'SGI Management Center 1.5'
   ```

Installing the Server

To install the SGI Management Center Server, install the required packages, as shown in the following examples.

Note

If you are installing on an SGI Altix UV system, going to use the PBS Dynamic Provisoning feature, or going to use the Discover feature, refer to the required packages described in *RPM Packages* on page 5.

Example (SLES):

```
# zypper install sgimc
```

Zypper will resolve all dependencies and ask for confirmation that you wish to install the packages.

Note

You may see dependency failures regarding `sgimc-tftp` and `sgimc-tftpboot` followed by Zypper trying those packages again and succeeding. This is normal behavior for Zypper on SLES 10.x with interdependent packages.

Example (RHEL):

```
# yum install sgimc
```

YUM will resolve all dependencies and ask for confirmation that you wish to install the packages.

Client Installation

The client allows you to remotely manage your cluster from a computer that is not part of the cluster. The client installation also gives you superior performance because it significantly reduces network traffic. You can install the client on a computer running Linux or Windows.

Linux Client Installation

To install the Linux client, set up the desired software repositories as described in section *Pre-Installation Steps* on page 5 and install the following package from the SGI Management Center media:

`sgimc-client`

Example (SLES):

```
# zypper install sgimc-client
```

Example (RHEL):

```
# yum install sgimc-client
```

Windows Client Using the Management Center Installer

1. Insert the SGI Management Center CD in your CD/DVD-ROM drive and allow the SGI Management Center installer to launch.

   If the installer does not start automatically, launch the installer manually (assuming the CDROM drive is `d:`):

   ```
   d:\windows\launch_installer.vbs
   ```
2. Specify the Installation Directory and Host Name, then click **Next**.

Note

The SGI Management Center Server or Master Host must use a valid host name that can be resolved through name resolution (for example, DNS, */etc/hosts*). For information on changing the name of the Master Host, see *Renaming the Management Center Master Host* on page 54.

3. Review the installation settings and click **Install** to continue.

Note

Java 6 is required for the Windows client. If the installer detects no Java 6, it will start an installation of Sun Java 6 JRE for you. The Java installer will ask you to accept the applicable license agreement.

4. After the installation completes, click **Finish**.
5. When you finish installing SGI Management Center, use Explorer to navigate to the installation directory.
6. Copy the SGI Management Center shortcut to your desktop.
7. Use the desktop shortcut to launch SGI Management Center.

Tip

You can also start SGI Management Center from the command-line interface. For example, if you installed to the default location *`c:\program files\sgi`* , enter the following:

```
c:\program files\sgi\sgimc\bin\mgrclient.vbs
```

To run SGI Management Center from a remote share, map the network drive where you installed SGI Management Center and create a copy of the shortcut on your local machine.

Windows Clients and Connect to Console Feature

In Windows 7, Windows Server 2008 R2, and Windows Vista, you may need to enable the Telnet client before you can use the Connect to Console feature of SGI Management Center. You can enable the Telnet client by doing the following:

1. Open the **Control Panel**.
2. Select **Programs**.
3. Select **Turn Window Programs On or Off**.
4. Click the appropriate checkbox to enable the Telnet client.

Advanced Scale-Out Configuration

To configure an SGI Management Center system for scale-out functionality past the default 4096 compute nodes, multiple instances of the SGI Management Center must be present. With this scale-out methodology, the SGI Management Center can support numerous groups of 4096 compute nodes to scale up to tens of thousands of nodes.

Prerequisites

This advanced configuration requires the following prerequisites:

- A shared filesystem on the host node (SAN, NAS, NFS, etc.)
- More than one host or service node running an instance of the SGI Management Center
- Support for IGMP multicast routing in the cluster environment
- Proper configuration of the SGI Management Center

Configuration

The following steps describe how to configure the SGI Management Center for scale-out:

1. Designate one system to be the *primary host* for the SGI Management Center.

This system will manage the first 4096 compute nodes and will be utilized for image, kernel, and payload management.

2. Add multiple service nodes to accommodate the desired node count.

Each host can manage 4096 compute nodes. For example, 32,768 compute nodes require 1 primary host and 7 service nodes.

3. Install the SGI Management Center on all of the participating host and service nodes.
4. Populate the various SGI Management Center databases with their respective 4096 compute nodes.
5. Export the *$MGR_HOME/vcs* directory on the primary host across the shared filesytem for the service nodes.

For NFS:

```
# /opt/sgi/sgimc/vcs 10.0.10.*(rw,sync,no_root_squash)
```

Note

The primary host will be the only system managing the VCS mechanism. The other subordinate service node directories will not be populated or managed.

6. From the participating service nodes, mount the shared filesystem.

For NFS:

```
# mount master:/opt/sgi/sgimc/vcs /opt/sgi/sgimc/vcs
```

7. Modify the IGMP multicast base addresses on the participating services nodes from their default settings.

This is accomplished through the SGI Management Center GUI:

Edit —> **Preferences** —> **Provisioning**

In this scenario, the following is an example of the base address layout for IGMP multicasting:

`master   239.192.0.128` (No change required from default configuration.)

`service1 239.192.1.128`

`service2 239.192.2.128`

`service1 239.192.3.128`

`service2 239.192.4.128`

`service1 239.192.5.128`

`service2 239.192.6.128`

`service1 239.192.7.128`

Remember to modify your IGMP multicast routing tables as well on these nodes.

For example:

`239.192.0.0/24, 239.192.1.0/24, 239.192.2.0/24,` etc.

8. Configure your images, kernels, and payloads on the primary host for your cluster.

Tip

The primary host can be utilized for validation of images, kernels, and payloads for your system using the *working* and *versioned* check-out mechanism. This can be useful in provisioning the primary group of 4096 compute nodes initially and ensuring desired functionality.

Provisioning

Do the following to provision the cluster:

1. Provision the primary host.

The primary host will manage the VCS and working copies of the images, kernels, and payloads for the subordinate service nodes.

Note

The separate provisioning of each block of 4096 compute nodes does not imply that you cannot boot all nodes simultaneously. It only means that the SGI Management Center instances are sharing the same VCS imaging database. This is to avoid complications within the VCS system.

2. Log in to each service node and start the SGI Management Center GUI.

The VCS entries from the primary host will be populated on these nodes.

3. From each service node, provision the corresponding compute nodes.

Note

If the SGI Management Center GUIs are open and you make changes to the primary host VCS entries, you will need to refresh the service node GUIs to see the modifications. You can do this by toggling between the **Working Images** and **Versioned Images** tabs.

Instrumentation

Each instance of the SGI Management Center will monitor the environmental, thermal, and other metric data from their assigned compute node groups. In order for each compute node to know where (which instance of SGI Management Center) to send its instrumentation data, you must modify the image on each compute node after installation.

To make this modification to the image, do the following:

1. Examine the script *scaleout_prefinalize.sh* carefully to determine whether or not you need to modify the script for your particular installation.

The script is in directory *$MGR_HOME/payload/utilities*.

2. Add the script as a *prefinalize* script for the image that you will be provisioning to your hosts.

Follow the instructions in section *RAM Disk* on page 134.

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.

For information on the various product editions of SGI Management Center, see *Product Definition* on page ix.

Importing Existing Hosts

After your SGI Management Center installation is complete, you can import existing hosts with the Import Host List option in the **File** menu. See *Import Hosts* on page 55.

Starting and Stopping the SGI Management Center Server

SGI Management Center services are started and stopped from scripts that exist in */etc/init.d*. SGI Management Center, typically installed in */opt/sgi/sgimc*, is controlled by one of these services—this allows you to manage SGI Management Center services using standard Linux tools such as `chkconfig` and `service`. Standard functions for services include start, stop, restart, and status. For example:

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

Verifying SGI 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

SGI 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

Introduction to SGI Management Center

Overview

SGI Management Center reduces the total cost of cluster ownership by streamlining and simplifying all aspects of cluster management. Through a single point of control, you can automate repetitive installation and configuration tasks. SGI Management Center automates problem determination and system recovery and monitors and reports health information and resource utilization.

SGI Management Center provides administrators with increased power and flexibility in controlling cluster system resources. Improved scalability and performance allows SGI Management Center to manage cluster systems of any size. Version-controlled provisioning allows administrators to easily install the operating system (OS) and applications on all hosts in the cluster and facilitates changes to an individual host or group of hosts.

Product Definition

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

- SGI Management Center, Standard Edition
- SGI Management Center, Premium Cluster Edition
- SGI Management Center, Premium ICE Edition
- SGI Management Center, Premium UV Edition

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

The following table charts the default features of the various product editions for Management Center. Note that feature Power Management is not a default feature of any of the product editions. It must be purchased as an add-on.

Feature	Standard Cluster	Premium Cluster	Premium ICE	Premium UV
Provisioning & BIOS Update	X	X	X	X
Image Management	X	X	X	X
Power On/Off Control	X	X	X	X
Console Connection	X	X	X	X
Instrumentation	X	X	X	X
Host Management	X	X	X	X
User Management	X	X	X	X
Dashboard View	X	X	X	X
Failure Analysis		X	X	X
High Availability		X	X	X
GPU Monitoring		X	X	X
Power Monitoring		X	X	X
UV SSI Management				X
Power Management (Optional)				

Comprehensive System Monitoring

SGI Management Center uses multiple monitoring features to improve system efficiency. These monitors allow you to examine system functionality from individual host components to the application level and help track system health, trends, and bottlenecks. With the information collected through these monitors, you can more easily plan for future computing needs—the more efficiently your cluster system operates, the more jobs it can run. Over the life of your system, you can accelerate research and time-to-market.

SGI Management Center provides results in near real-time and uses only a minute amount of the CPU. All data is displayed in a portable and easy-to-deploy Java-based GUI that runs on both Linux and Windows. Monitoring values include CPU usage, disk I/O, file system usage, kernel and operating system information, CPU load, memory usage, network information and bandwidth, and swap usage. Administrators may also write plug-ins to add functionality or monitor a specific device or application.

Noteworthy is the power monitoring and power management provided by SGI Managment Center together with Intel Data Center Manager (DCM) technology. This integration includes power state control, power utilization monitoring, power policy definition and enforcement, and power cost tracking and trending.

Version Controlled Image Management

Version control greatly simplifies the task of cluster administration by allowing system administrators to track upgrades and changes to the system image. If a problem arises with a system image, system administrators can even revert to a previous, more robust version of the image. By allowing system administrators to update the operating system and other applications both quickly and efficiently, version control ensures that organizations receive the highest return on their cluster system investment.

Tip

In cases where only minor changes are made to VCS-controlled images, SGI Management Center allows you to apply updates without re-provisioning. See *VCS Upgrade* on page 150.

Fast Multicast Provisioning

Thanks to fast multicast provisioning, SGI Management Center can add or update new images in a matter of minutes—no matter how many hosts your system contains. This saves time by allowing system administrators to quickly provision and incrementally update the cluster system as needed; and since updates take only a few minutes, this means less down-time and fewer system administration headaches.

Memory Failure Analysis

SGI Management Center supports failure analysis for memory errors via memlog, a software component of SGI Foundation Software. When the memlog software is installed and functional, the **Failure Analysis** panel within the SGI Management Center will be populated with pertinent memory data.

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.

Features

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

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

- A UV license is required for SGI Management Center.
- **EFI Boot** and **Direct 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 (shown in the following figure). The SSI partition is shown in the logical view.

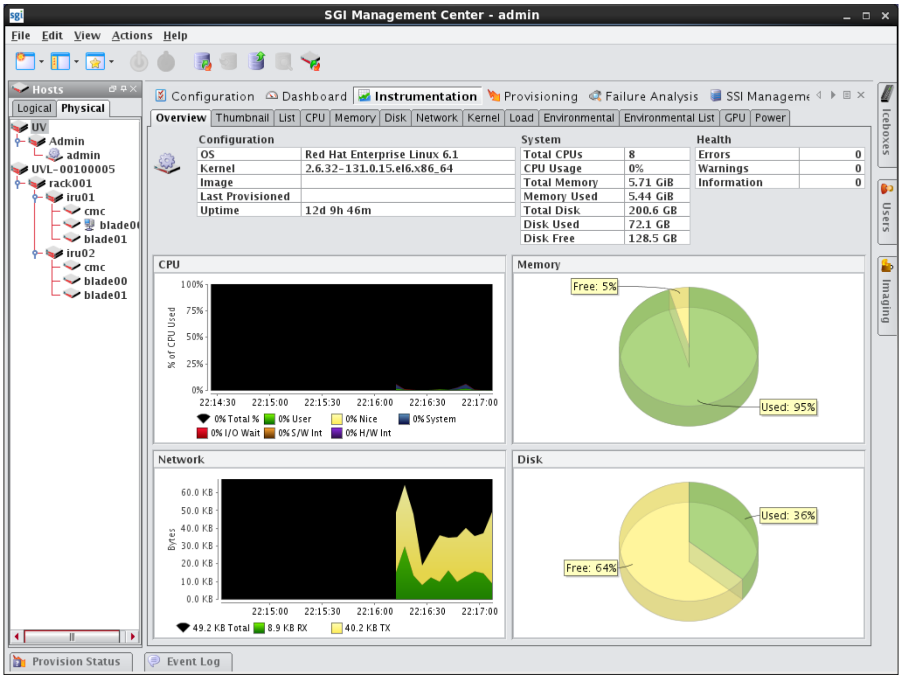
SGI Management Center - admin
File Edit View Actions Help
Hosts
Logical Physical
UV
Admin
admin
UVL-00100005
rack001
iru01
cmc
blade00
blade01
iru02
cmc
blade00
blade01
Configuration Dashboard Instrumentation Provisioning Failure Analysis SSI Manageme
Overview Thumbnail List CPU Memory Disk Network Kernel Load Environmental Environmental List GPU Power
Iceboxes
Users
Imaging
CPU
100%
75%
50%
25%
0%
% of CPU Used
22:14:30 22:15:00 22:15:30 22:16:00 22:16:30 22:17:00
0% Total % 0% User 0% Nice 0% System
0% I/O Wait 0% S/W Int 0% H/W Int
Memory
Free: 5%
Used: 95%
Network
60.0 KB
50.0 KB
40.0 KB
30.0 KB
20.0 KB
10.0 KB
0.0 KB
Bytes
22:15:00 22:15:30 22:16:00 22:16:30 22:17:00
49.2 KB Total 8.9 KB RX 40.2 KB TX
Disk
Used: 36%
Free: 64%
Provision Status
Event Log

Configuration

OS	Red Hat Enterprise Linux 6.1
Kernel	2.6.32-131.0.15.el6.x86_64
Image	
Last Provisioned	
Uptime	12d 9h 46m

System

Total CPUs	8
CPU Usage	0%
Total Memory	5.71 GiB
Memory Used	5.44 GiB
Total Disk	200.6 GB
Disk Used	72.1 GB
Disk Free	128.5 GB

Health

Errors	0
Warnings	0
Information	0

Provisioning

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**.

Kernel Parameters

The UV SSI automatically determines the best kernel boot parameters for your UV hardware. Consult the UV documentation for details. If these kernel parameters change (for example, due to a hardware change), SGI Management Center will send a warning alert in the **Event Log** panel. The warning will indicate that the kernel parameters have changed and will instruct you to import the new kernel parameters for the appropriate kernel.

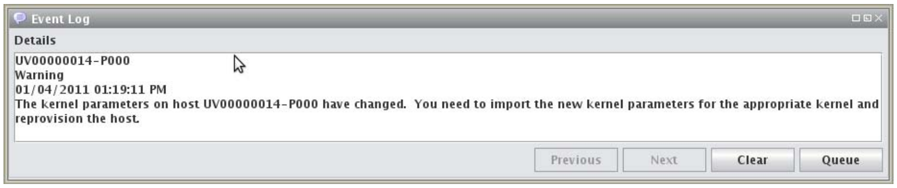

```
UV00000014-P000
Warning
01/04/2011 01:19:11 PM
The kernel parameters on host UV00000014-P000 have changed.  You need to import the new kernel parameters for the appropriate kernel and
reprovision the host.
```

Section *Importing Kernel Parameters from a Running Host* on page 89 describes how you import the kernel parameters.

Support of SGI Prism XL Systems

SGI Management Center provides limited support for the SGI Prism XL platforms. Currently, this extends to the monitoring of NVIDIA and AMD GPUs (items like temperature, fan speed, memory usage, and ECC). For a listing of GPU solutions supported by SGI, see http://www.sgi.com/pdfs/4235.pdf.

SGI Management Center requires special licensing for the GPU Monitoring feature.

Auto Node Discovery

SGI Management Center provides the Discover interface to greatly assist you in adding new compute nodes to your host tree. Discover will determine the pertinent MAC addresses for the new nodes. For more details and usage information, see *Using the Discover Interface* on page 207.

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.

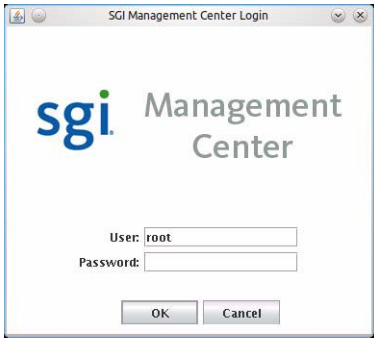

User: root

Password:

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

The Management Center interface is displayed.

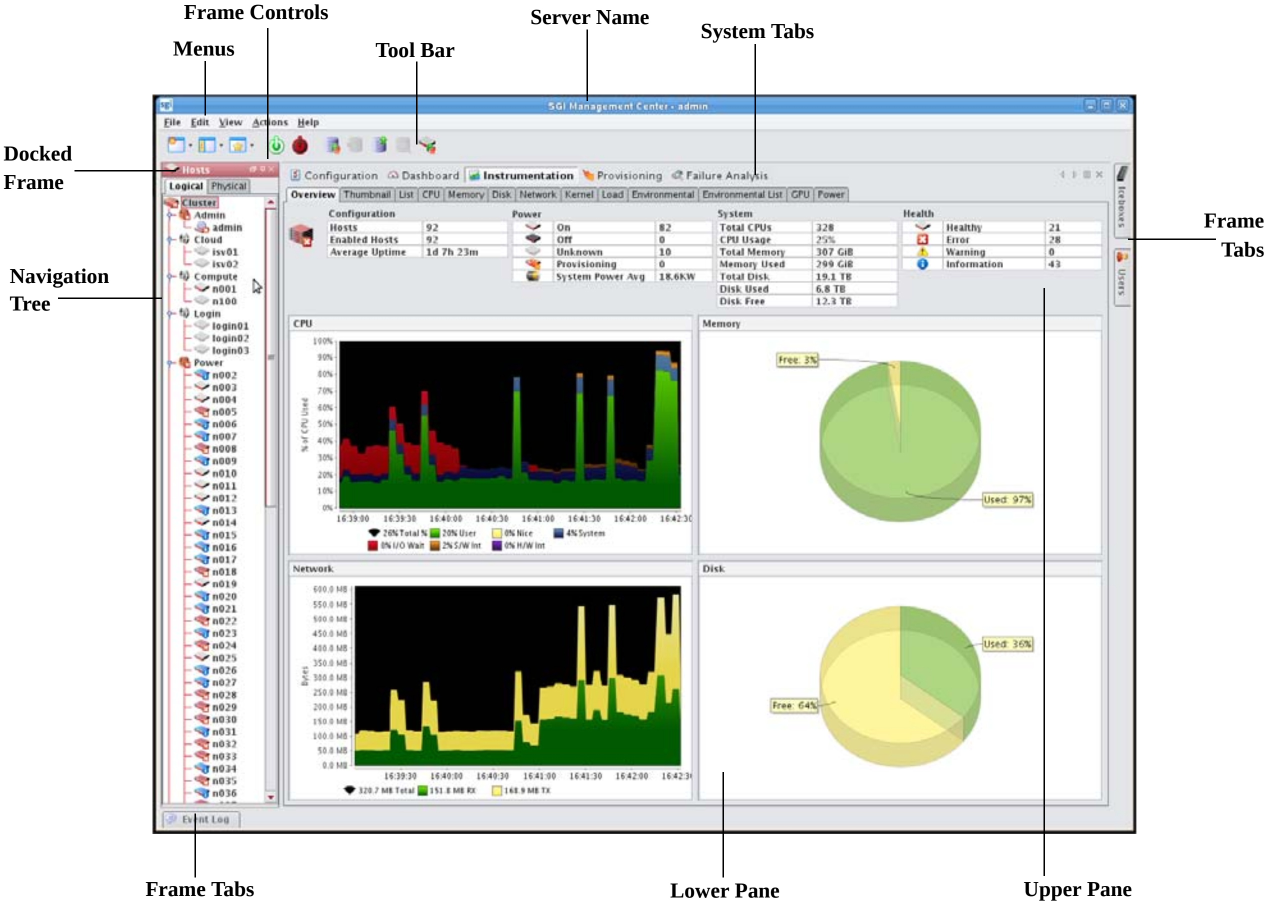
Frame Controls
Menus
Tool Bar
Server Name
System Tabs
Docked Frame
Navigation Tree
Frame Tabs
Frame Tabs
Lower Pane
Upper Pane
SGI Management Center - admin
File Edit View Actions Help
Hosts
Logical Physical
Configuration Dashboard Instrumentation Provisioning Failure Analysis
Overview Thumbnail List CPU Memory Disk Network Kernel Load Environmental Environmental List GPU Power
CPU
Memory
Free: 3%
Used: 97%
Network
Disk
Used: 36%
Free: 64%
Event Log
Iceboxes
Users

Configuration	
Hosts	92
Enabled Hosts	92
Average Uptime	1d 7h 23m

Power	
On	82
Off	0
Unknown	10
Provisioning	0
System Power Avg	18.6KW

System	
Total CPUs	328
CPU Usage	25%
Total Memory	307 GiB
Memory Used	299 GiB
Total Disk	19.1 TB
Disk Used	6.8 TB
Disk Free	12.3 TB

Health	
Healthy	21
Error	28
Warning	0
Information	43

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.

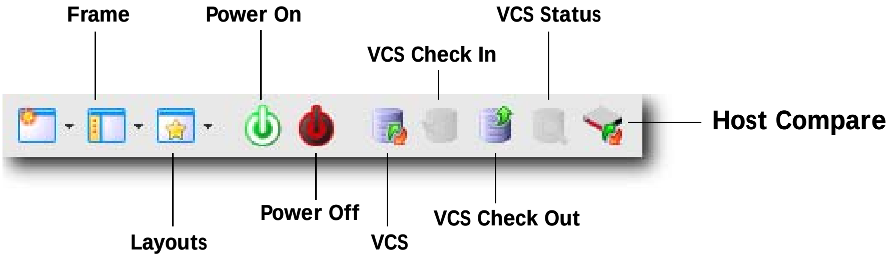
Frame
Power On
VCS Status
VCS Check In
Host Compare
Layouts
Power Off
VCS
VCS Check Out

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.

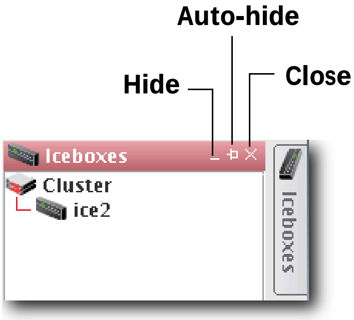

Frames — Provide you with specific control over common aspects of cluster systems (for example, imaging and user accounts). Each frame tab opens a frame containing a navigation tree that allows you to manage system components easily. The navigation tree is found in most frames and is used to help organize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

Customizing the Interface

Management Center is flexible and can be modified to meet your specific needs. For example, you can arrange the interface to make it easier to view multiple frames or configure it to display only those items related to a particular task (such as provisioning). You can save each view as a custom layout and easily toggle between saved views—which is helpful if you have multiple users administering your clusters.

Customizing System Tabs

The system tabs include **Configuration**, **Dashboard**, **Instrumentation**, and **Provisioning**. You can open, close, or rearrange these system tabs as needed.

Closing Tabs

1. Click and display the tab you want to close.
2. Click the **Close** icon at the right of the tab's pane.

Opening Tabs

- From the **View** menu, select **Tabs >** *<name of tab>*.

  or
- Select the **Tab** you want to display.

  or
- Click the **Tab** icon in the tool bar and select the name of the tab you want to display.

Arranging Tabs

REORDERING TABS

- Click a tab and drag it to a new position.

CREATING SPLIT PANE VIEWS

1. Right-click a tab and select **New Horizontal Group** or **New Vertical Group**.
2. The tab appears in a new horizontal or vertical pane.

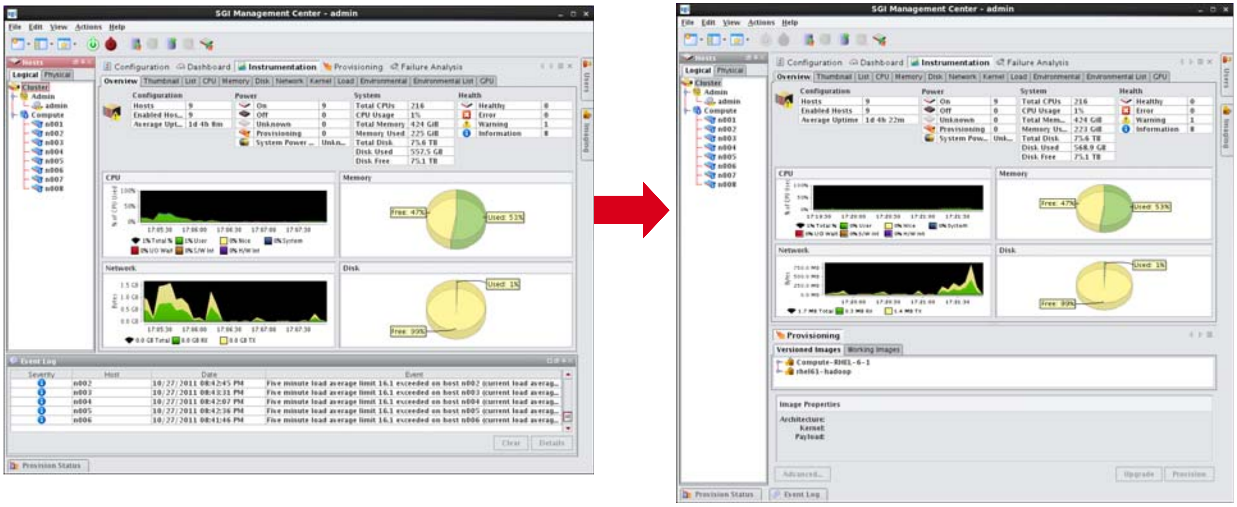
SGI Management Center - admin
Configuration
Dashboard
Instrumentation
Provisioning
Failure Analysis
CPU
Memory
Network
Disk
Event Log
Provision Status
SGI Management Center - admin
Configuration
Dashboard
Instrumentation
Failure Analysis
CPU
Memory
Network
Disk
Provisioning
Versioned Images
Working Images
Image Properties
Provision Status
Event Log

Tip

To move tabs between groups, right-click the tab and select **Move to Next Tab Group**. You can also drag and drop tabs between groups.

Dockable Frames

Management Center dockable frames can be opened, closed and repositioned to meet your needs.

Note

Before you can reposition a frame, you must click the **Auto-Hide** button to make the frame always visible. See *Frame Controls — Lets you dock, un-dock, hide, minimize, and close frames.* on page 22.

To Move a Dockable Frame

1. Open the frame and toggle the **Auto-Hide** button to **Off**.
2. Click the frame's title bar and drag it to a new position in the interface.

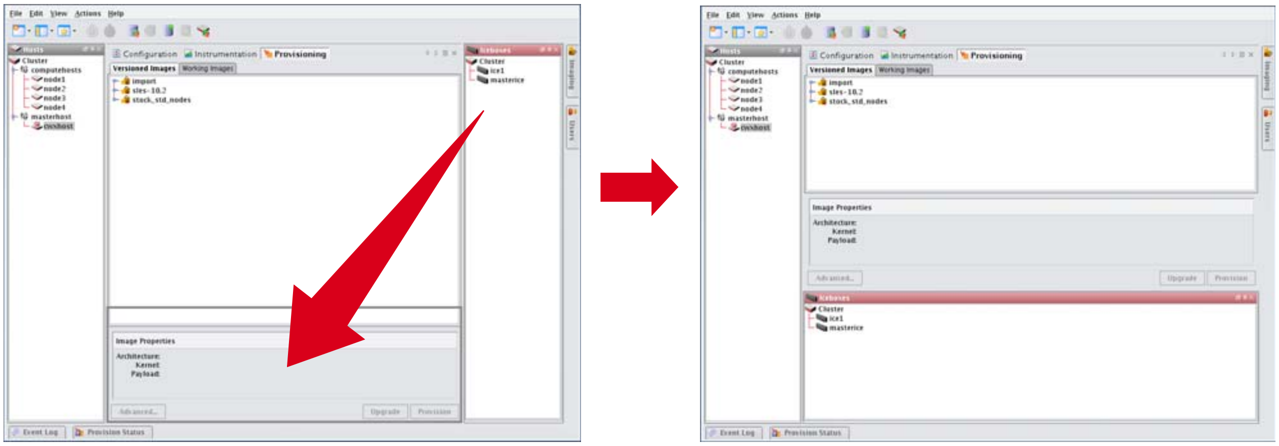
File Edit View Actions Help
Configuration
Instrumentation
Provisioning
Versioned Images
Working Images
Image Properties
Architecture:
Kernel:
Payload:
Event Log
Provision Status

Layouts

Customized views of the Management Center interface are easily saved and accessed from the **View** menu or opened with the **Layouts** button on the toolbar.

Saving the Current Layout

1. From the **View** menu, select **Manage Layouts**.
2. Select **Save Current Layout**.

Tip

To overwrite an existing layout with the current view, move the mouse over the layout and select **Overwrite** from the popup menu.

3. Enter the name of the new layout and click **OK**.

Opening a Saved Layout

1. On the tool bar, click the **Layouts** button, or select **Layouts** from the **View** menu.
2. Select the layout you want to open from the popup menu.

Renaming a Layout

1. From the **View** menu, select **Manage Layouts**.
2. Move the mouse over the layout you want to rename and select **Rename** from the popup menu.
3. Enter the new name of the layout and click **OK**.

Adding a Layout Description

1. From the **View** menu, select **Manage Layouts**.
2. Move the mouse over the layout you want to add a description to and select **Describe** from the popup menu.
3. In the Input dialog, enter a brief description and click **OK**.

Deleting a Layout

1. From the **View** menu, select **Manage Layouts**.
2. Move the mouse over the layout you want to delete and select **Delete** from the popup menu.

Setting a Default Layout

The default layout launches every time you start Management Center (by default, the System Default layout). To change the default layout:

1. From the **View** menu, select **Manage Layouts**.
2. Move the mouse over the layout you want to set as the default and select **Set as Default** from the popup menu.

Chapter 3

Preferences and Settings

Preferences

Management Center preferences allow you to configure the global settings and default behavior for your cluster. Preferences include general settings, platform management configurations, applications, and provisioning. Although these settings apply to the entire cluster, you may override certain preferences as needed (such as, provisioning).

You can access preferences by selecting **Preferences** from the **Edit** menu.

General

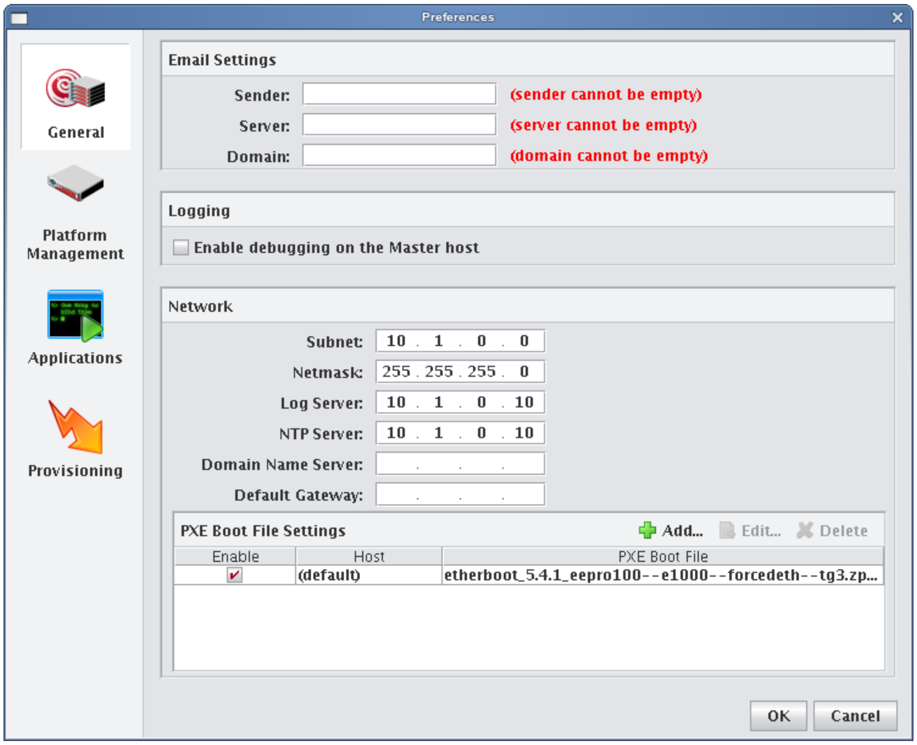
Preferences
General
Platform Management
Applications
Provisioning
Email Settings
Logging
Network
PXE Boot File Settings
Add...
Edit...
Delete
OK
Cancel

Sender: (sender cannot be empty)

Server: (server cannot be empty)

Domain: (domain cannot be empty)

☐ Enable debugging on the Master host

Subnet: 10 . 1 . 0 . 0

Netmask: 255 . 255 . 255 . 0

Log Server: 10 . 1 . 0 . 10

NTP Server: 10 . 1 . 0 . 10

Domain Name Server: . . .

Default Gateway: . . .

Enable	Host	PXE Boot File
☑	(default)	etherboot_5.4.1_eepro100--e1000--forcedeth--tg3.zp...

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.
- **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.

Platform Management

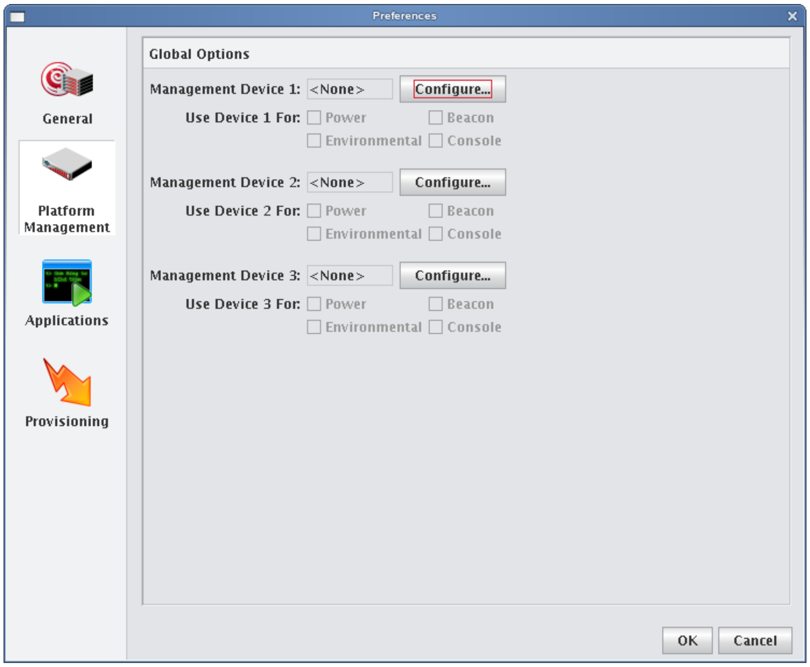

Management Device 1: <None>
Use Device 1 For: ☐ Power ☐ Beacon ☐ Environmental ☐ Console

Management Device 2: <None>
Use Device 2 For: ☐ Power ☐ Beacon ☐ Environmental ☐ Console

Management Device 3: <None>
Use Device 3 For: ☐ Power ☐ Beacon ☐ Environmental ☐ Console

Global Options

Management Center supports multiple platform management interfaces. This is useful if you are using multiple platforms for system management (for instance, one interface for power management and another for environmental monitoring). The global options section of the preferences dialog allows you to set the default options used for the majority of hosts in the system, although some hosts may still need additional configuration.

Set the most common options by configuring Device 1, 2, and 3. From the configuration dialog, select the Platform Management Device Type you want to use: Icebox (not currently supported), Roamer, IPMI, FreeIPMI, DRAC, ILO, Powerman, or Conman). The check boxes below each device indicate which features are available to be managed by the device. If you configure multiple devices, you can select or clear these check boxes to indicate which device will manage this feature. See also *Intel Data Center Manager (DCM)* on page 32.

Note

Not all management controllers have the same feature set. DRAC and ILO support only power control and Conman supports only serial console control. Roamer supports serial console control and power control. IPMI supports all features, but you may want to use other interfaces for power and serial console control and IPMI for controlling the beacon and environmental monitoring.

Roamer, IPMI, DRAC, ILO

Roamer, IPMI, ILO, and DRAC are typically configured using only global options. Prior to configuring the Management Device, you must configure the IP address of the controller. This is typically set dynamically via DHCP but can also be set statically. If you choose to assign the IP address dynamically, add the device's MAC address and IP address as a secondary interface under each Management Center host. This causes Management Center to automatically add an entry for the interface in the */etc/dhcpd.conf* file and attempt to configure it via DHCP.

Management Center provides three Management Device IP Address Types: Dynamic, Relative, and Static. These address types are described below, but you may also want to refer to *IPMI* on page 51 and *DRAC and ILO* on page 53 for additional information and examples.

Note

To configure platform management to use a remote power control device such as IPMI, ILO, or DRAC, you must first create the power control user. See *Configure the Master Host and Management Center* on page 43.

In order for DRAC to successfully control power on DRAC-enabled hosts, you must install the racadm utility on the Master Host. You may obtain the racadm RPM, *mgmtst-racadm-4.5.0-335.1386.rpm*, from the */misc* directory of the Management Center CD or from SGI technical support.

Dynamic If you are setting up the Management Device dynamically and the device's interface MAC address is an offset of the management interface, set the Management Device IP Address Type to Dynamic and enter the MAC Address Offset. This is typical for IPMI implementations with on-board BMC controllers. For example, a host whose management interface MAC address is 00:11:22:33:44:55 might have a Management Device with a MAC address of 00:11:22:33:44:58. In this case, the MAC offset would be 00:00:00:00:00:03 (Greater Than).

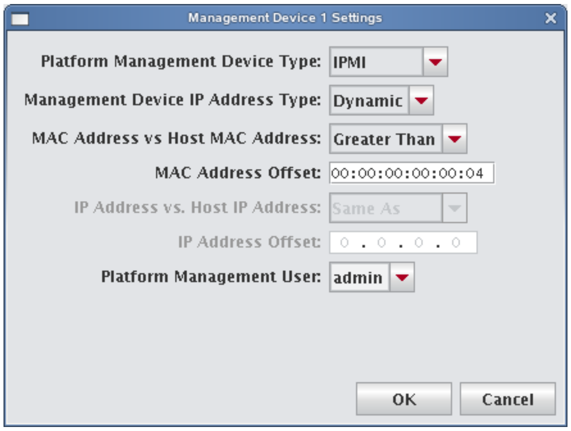

Management Device 1 Settings

Platform Management Device Type: IPMI
Management Device IP Address Type: Dynamic
MAC Address vs Host MAC Address: Greater Than
MAC Address Offset: 00:00:00:00:00:04
IP Address vs. Host IP Address: Same As
IP Address Offset: 0 . 0 . 0 . 0
Platform Management User: admin

Relative If you are setting up the Management Device dynamically or statically and the device's interface IP address is an offset of the management interface, set the Management Device IP address type to Relative and use the IP Address Offset. This is typical when using ILO or an IPMI controller with an add-on BMC daughter card. For example, a host with an IP address of 10.0.0.1 might have a Management Device with an IP address of 10.0.2.1. In this case, the IP offset would be 0.0.2.0 (Greater).

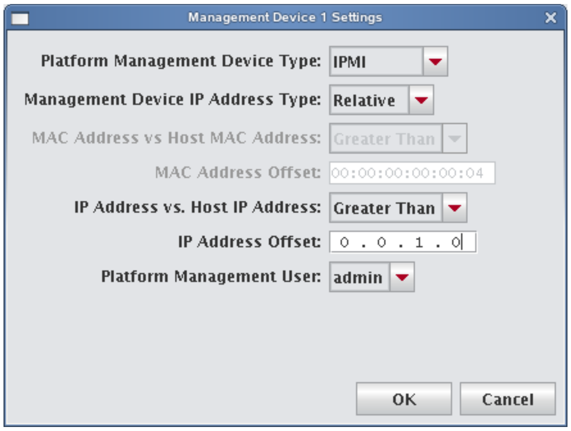

Platform Management Device Type: IPMI
Management Device IP Address Type: Relative
MAC Address vs Host MAC Address: Greater Than
MAC Address Offset: 00:00:00:00:00:04
IP Address vs. Host IP Address: Greater Than
IP Address Offset: 0 . 0 . 1 . 0
Platform Management User: admin

Static If you are setting up the Management Device dynamically or statically and the device's interface MAC address or IP address does not correlate with either the MAC or IP address of the management interface, set the Management Device IP address type to Static—this is not typical. If you select Static, you must configure the IP address manually on a per-host basis.

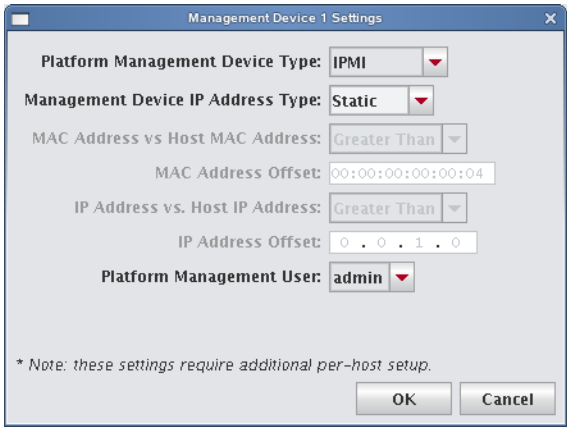

Platform Management Device Type: IPMI
Management Device IP Address Type: Static
MAC Address vs Host MAC Address: Greater Than
MAC Address Offset: 00:00:00:00:00:04
IP Address vs. Host IP Address: Greater Than
IP Address Offset: 0 . 0 . 1 . 0
Platform Management User: admin

* Note: these settings require additional per-host setup.

Conman

Conman is a serial console management program designed to support a large number of devices simultaneously. Conman supports multiple serial controllers (including IPMI) and provides continuous serial logging and multiplexing that allows you to share a serial connection for logging and access, or between multiple consoles.

Tip

Conman is available under the GPL and is installed by default on SGI systems. Conman can be obtained from SGI as RPM packages or from *http://home.gna.org/conman/*.

Prior to selecting **Conman** for serial access, you must install the conman RPM on the Master Host, then configure conman by defining the serial devices and consoles in */etc/conman.conf*. Additional information on conman is available from the man pages by entering man conman.conf.

Note

Before you can begin using conman, you must start its daemon, `conmand` (installed as `/etc/init.d/conmand`). For information on using conman, see *conman* on page 226.

Intel Data Center Manager (DCM)

Management Center supports power monitoring and scaling through the utilization of the Intel Data Center Manager (DCM) technology. DCM provides Management Center with a robust power management API. This integration includes power state control, power utilization monitoring, power policy definition and enforcement, and power cost tracking and trending.

SPECIAL HARDWARE AND LICENSING

Enabling DCM requires special hardware for your system and an appropriate SGI license key for feature activation. If this feature was not pre-configured for your cluster in Manufacturing, contact SGI for information about enabling this feature.

SGI Management Center provides power monitoring for each node configured for DCM-based management. DCM-based management may be enabled for any node with a DCM-compatible design. Such units include Intel-based chipsets with IPMI and Intel Power Node Manager (IPNM) as well as other vendor technologies providing standards-based power management. Typically, PMBUS-compliant hardware is required.

ENABLING DCM

To enable DCM-based management for compliant hardware, do the following:

1. Configure the appropriate IPMI endpoint parameters.
   Use the Platform Management pane of the host-configure panel (**Configuration** tab) or the global preferences dialogue:

   Edit —> Preferences —> Platform Management

   See section *Roamer, IPMI, DRAC, ILO* on page 30 for more information on configuring the platform management preferences for IPMI.

2. Add a DCM management device.
   You can use the same interfaces as in step 1. It is not necessary to target the DCM management device toward any particular management subsystem (power, beacon, environmental or console) but you must enter the DCM access information in the management device settings dialogue. When you select **Configure** in the **Platform Management** section to enable a new management device, you must specify values for the fields in the following table:

Field	Description
Management Device Type	Select DCM.
DCM Primary Host Name	Provides the DCM server hostname or IP address used to access the DCM service.
Nameplate Power	Provides the maximum PSU output power for the monitored equipment. (Using the global options pane will create a per-node default nameplate power.)
Platform Management User	Provides authenticated access to the node management infrastructure.

Once configured, SGI Management Center will synchronize with DCM and achieve model consistency. Management Center will begin to receive event updates from the DCM service. You can access power monitoring data through **Instrumentation** —> **Power** . See *Power Tab* on page 168.

POWER POLICY MANAGEMENT

Management Center provides power policy management capability for licensed customers whose systems comply with the aforementioned requirements for Management Center DCM integration. You can access the managment functions in two ways:

- **Action** —> **Power**
- Right-clicking on a cluster-tree target in the **Hosts** side bar.

Selecting **Policy** from the **Power** sub-menu displays the **Policy Management** dialog box.

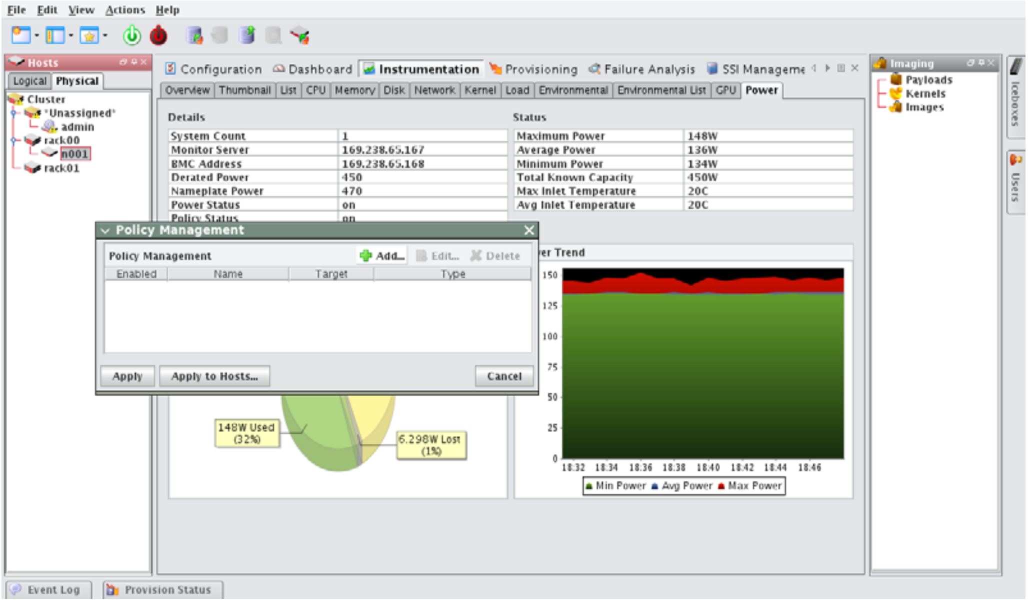
File
Edit
View
Actions
Help
Hosts
Logical
Physical
Cluster
'Unassigned'
admin
rack00
n001
rack01
Configuration
Dashboard
Instrumentation
Provisioning
Failure Analysis
SSI Manageme
Overview
Thumbnail
List
CPU
Memory
Disk
Network
Kernel
Load
Environmental
Environmental List
GPU
Power
Policy Management
Policy Management
Add...
Edit...
Delete
Enabled
Name
Target
Type
Apply
Apply to Hosts...
Cancel
148W Used (32%)
6.298W Lost (1%)
er Trend
150
125
100
75
50
25
0
18:32
18:34
18:36
18:38
18:40
18:42
18:44
18:46
Min Power
Avg Power
Max Power
Imaging
Payloads
Kernels
Images
Iceboxes
Users
Event Log
Provision Status

Details

System Count	1
Monitor Server	169.238.65.167
BMC Address	169.238.65.168
Derated Power	450
Nameplate Power	470
Power Status	on
Policy Status	on

Status

Maximum Power	148W
Average Power	136W
Minimum Power	134W
Total Known Capacity	450W
Max Inlet Temperature	20C
Avg Inlet Temperature	20C

Any existing policy definitions will be displayed.

Adding New Policies

New policies may be defined for the selected physical target or for the cluster as a whole (if selected) by clicking **Add** from the **Policy Management** dialog box. The GUI displays the **Add Custom Policy** dialog box for policy definition.

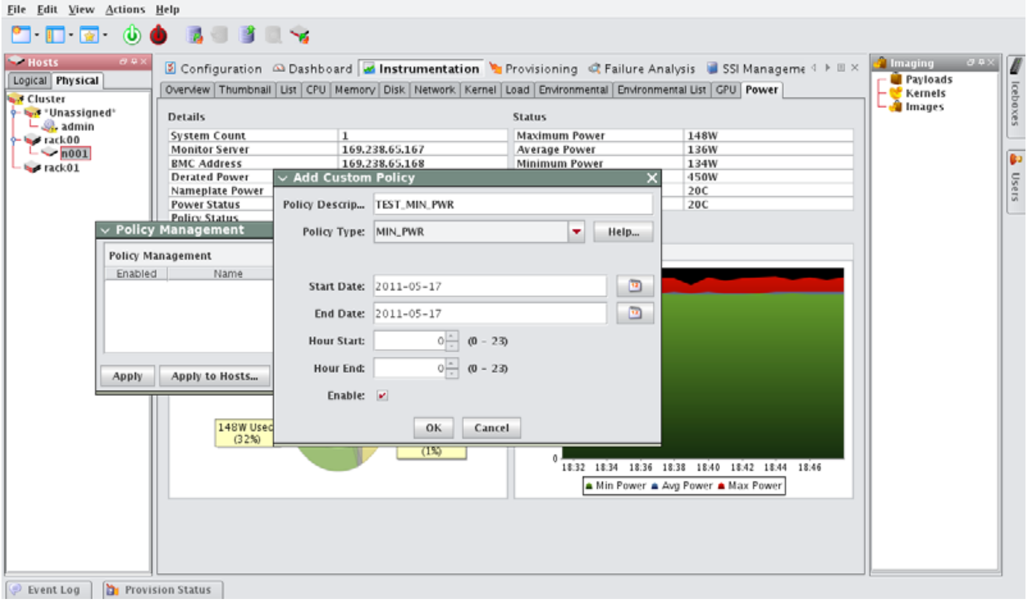
File Edit View Actions Help
Hosts
Logical Physical
Cluster
Unassigned
admin
rack00
n001
rack01
Configuration Dashboard Instrumentation Provisioning Failure Analysis SSI Manageme
Overview Thumbnail List CPU Memory Disk Network Kernel Load Environmental Environmental List GPU Power
Imaging
Payloads
Kernels
Images
Iceboxes
Users
Policy Management
Enabled Name
Apply Apply to Hosts...
148W Used (32%)
(1%)
0
18:32 18:34 18:36 18:38 18:40 18:42 18:44 18:46
Min Power Avg Power Max Power
Event Log Provision Status

Details

System Count	1
Monitor Server	169.238.65.167
BMC Address	169.238.65.168
Derated Power	
Nameplate Power	
Power Status	
Policy Status	

Status

Maximum Power	148W
Average Power	136W
Minimum Power	134W
	450W
	20C
	20C

Add Custom Policy

Policy Descrip...: TEST_MIN_PWR
Policy Type: MIN_PWR
Start Date: 2011-05-17
End Date: 2011-05-17
Hour Start: 0 (0 - 23)
Hour End: 0 (0 - 23)
Enable: ☑

Help...
OK
Cancel

You can select one of the policy types described in the table below:

Policy Type	Description
CUSTOM_PWR_LIMIT	Defines a maximum power threshold for the endpoint. If you select **CUSTOM_PWR_LIMIT**, you will be prompted to enter Power Limit, (an integer representing the allowable power threshold in Watts).
MIN_PWR	Defines a policy which enforces the minimum operational power for the selected target (an entity or a group of entities).
MIN_PWR_ON_INLET_ TEMP_TRIGGER	Establishes a **MIN_PWR** policy that is enforced upon detection of a specific temperature threshold within the target.
STATIC_PWR_LIMIT	Constructs immutable power budgets for entities that have special characteristics. Policies that are defined on groups of nodes (a rack or cluster, for example) may overlap with those defined for individual nodes and, typically, the policies attempt to achieve the highest possible power savings. **STATIC_PWR_LIMIT** may be used to preserve a specific power allocation budget for targeted endpoints within the cluster. If further **CUSTOM_PWR_LIMIT** restrictions are placed upon the endpoint directly or by the application of a group policy, the **STATIC_PWR_LIMIT** budget will take precedent.

The updated **Policy Management** dialog box reflects the new addition.

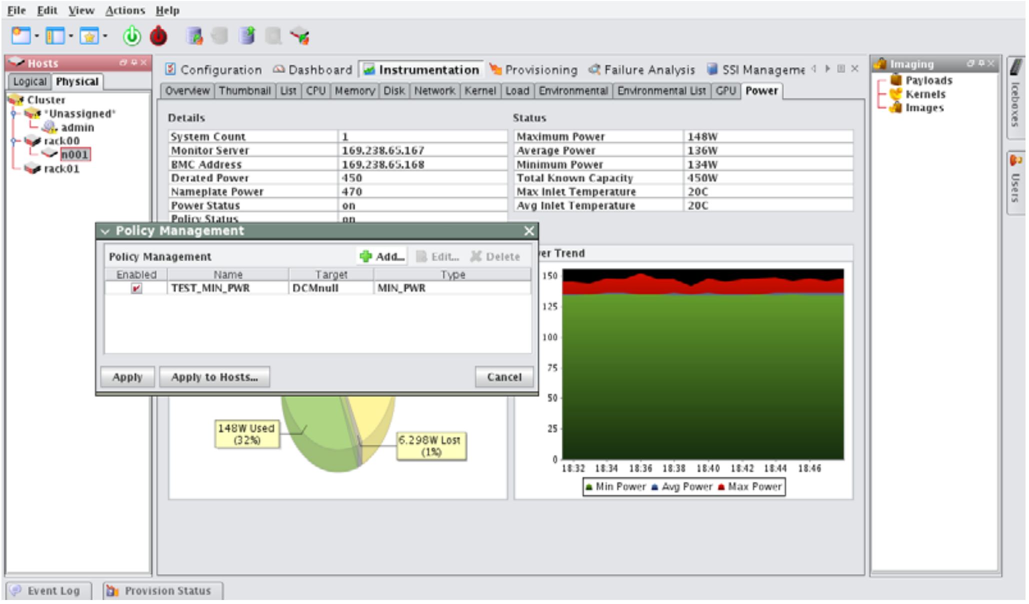
File Edit View Actions Help
Hosts
Logical Physical
Cluster
Unassigned
admin
rack00
n001
rack01
Configuration Dashboard Instrumentation Provisioning Failure Analysis SSI Manageme
Overview Thumbnail List CPU Memory Disk Network Kernel Load Environmental Environmental List GPU Power
Imaging
Payloads
Kernels
Images
Iceboxes
Users
148W Used (32%)
6.298W Lost (1%)
er Trend
150
125
100
75
50
25
0
18:32 18:34 18:36 18:38 18:40 18:42 18:44 18:46
Min Power Avg Power Max Power
Add... Edit... Delete
Apply Apply to Hosts... Cancel
Event Log Provision Status

Details

System Count	1
Monitor Server	169.238.65.167
BMC Address	169.238.65.168
Derated Power	450
Nameplate Power	470
Power Status	on
Policy Status	on

Status

Maximum Power	148W
Average Power	136W
Minimum Power	134W
Total Known Capacity	450W
Max Inlet Temperature	20C
Avg Inlet Temperature	20C

Policy Management

Policy Management

Enabled	Name	Target	Type
☑	TEST_MIN_PWR	DCMnull	MIN_PWR

Activating Policies

You can activate a power policy in an on-demand fashion or non-interactively via schedule definition. To do so interactively (on demand), use the **Edit** feature of the **Policy Management** dialog box. Simply select the **Enable** check box in the resulting display.

You can create and enable a schedule definition upon policy creation by entering valid values in the following fields:

- **Start Date**
- **End Date**
- **Hour Start**
- **Hour End**

You can later access these fields using the **Edit** feature from the **Policy Management** dialog box. The **Hour Start** to **Hour End** interval defines how the policy will be activated on a day-to-day basis.

Disabling Policies

A policy may be disabled at any time either by deleting the policy from the **Policy Management** dialog box (using the **Delete** button) or by using the **Edit** feature and de-selecting the **Enable** check box.

POWER COST TRACKING AND TRENDING

There are two aspects to the power cost monitoring provided by SGI Management Center:

- Daily cost tracking and hourly cost breakdown
- Real-time power cost trending

Daily Cost Tracking and Hourly Cost Breakdown

Management Center can provide the various equipment energy costs contributing to the operating costs of the data center and can provide a daily cost summary for the last 24-hour period. You can access power cost data through **Instrumentation —> Power** . See *Power Tab* on page 168.

Real-time Power Cost Trending

Cost trending corresponds to the billing rate (hourly trending) for servers, racks, and clusters. This capability tracks hourly costs for each endpoint against their power utilization characteristics over time. This feature aids in capacity planning for the data center. Management Center can calculate and display cost trends for power usage for systems. When combined with tactical power policy activation, this information enables business-critical decisions for managing operational costs which may be planned or executed on the fly.

Enabling Power Cost Tracking

As with establishing a power policy, you must enable DCM. See *Enabling DCM* on page 32. As shown in the following figure, the DCM configuration panel includes the following two entries:

- **Cooling Power Multiplier**
- **Power Cost per KWh**

You must make entries for these fields.

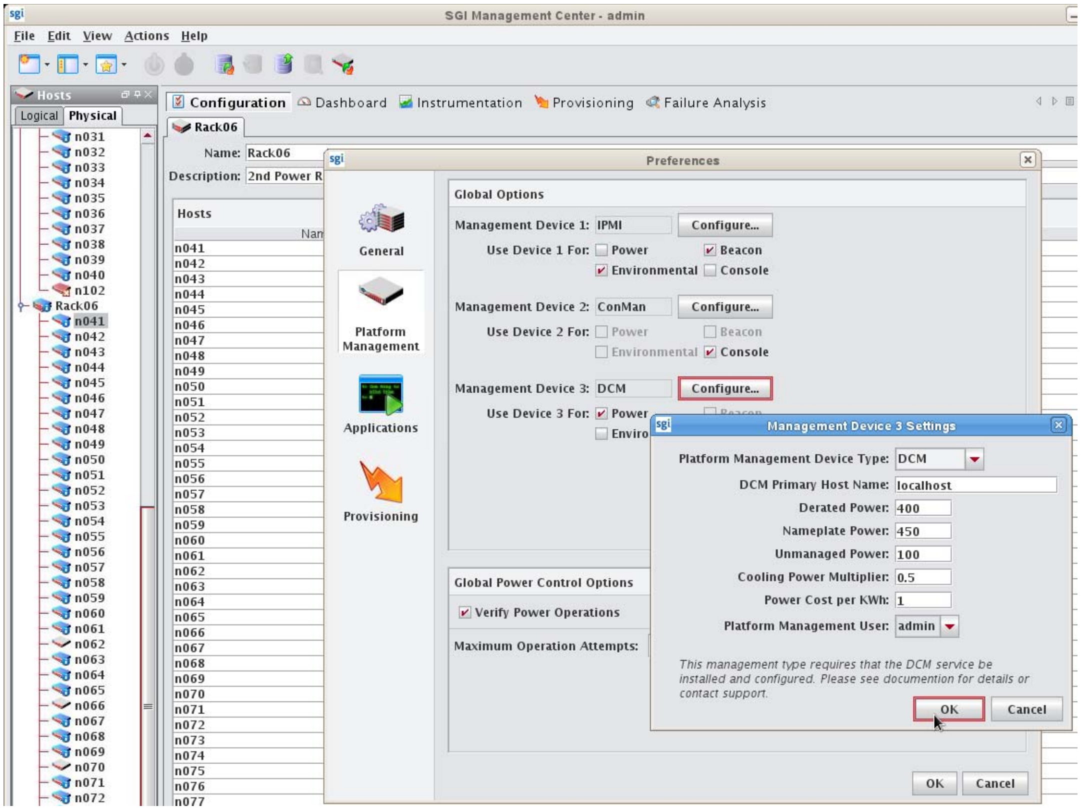
SGI Management Center - admin
File Edit View Actions Help
Hosts
Logical Physical
Configuration Dashboard Instrumentation Provisioning Failure Analysis
Rack06
General
Platform Management
Applications
Provisioning
Configure...
OK
Cancel

Name: Rack06
Description: 2nd Power R

Preferences

Global Options

Management Device 1: IPMI
Use Device 1 For: ☐ Power ☑ Beacon ☑ Environmental ☐ Console

Management Device 2: ConMan
Use Device 2 For: ☐ Power ☐ Beacon ☐ Environmental ☑ Console

Management Device 3: DCM
Use Device 3 For: ☑ Power

Global Power Control Options
☑ Verify Power Operations
Maximum Operation Attempts:

Management Device 3 Settings

Platform Management Device Type: DCM
DCM Primary Host Name: localhost
Derated Power: 400
Nameplate Power: 450
Unmanaged Power: 100
Cooling Power Multiplier: 0.5
Power Cost per KWh: 1
Platform Management User: admin

This management type requires that the DCM service be installed and configured. Please see documentation for details or contact support.

The **Cooling Power Multiplier** field represents the rate (currency-agnostic) at which energy is billed per kilowatt-hour.

The **Power Cost per KWh** field represents the estimated ratio of cooling power over operating load for the system. It is used to calculate the rate of cooling energy expenditure for each endpoint and the system as a whole. The following equations summarize the use of the cooling power multiplier:

```
IT_PWR = Unmanaged_Equip_Power + Managed_Equip_Power + Static_Power
COOLING_PWR = IT_PWR * Cooling_Power_Multiplier
COOLING_ENERGY = COOLING_PWR / 3600 * Aggregation_Period
COOLING_COST = COOLING_ENERGY / 1000 * Power_Cost_per_KWh
```

Applications

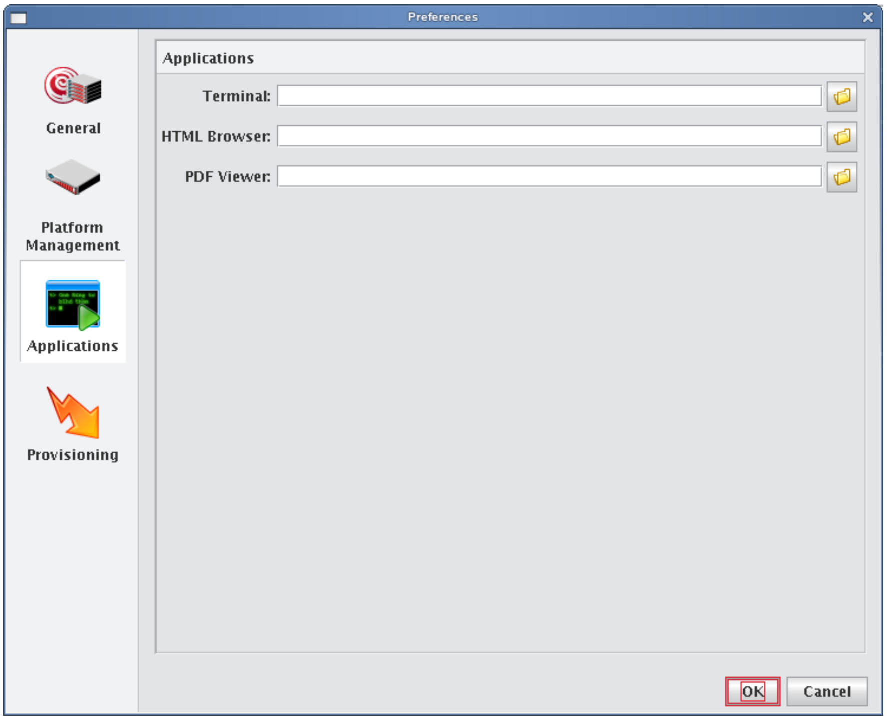

APPLICATIONS

The applications option allows you to select the default applications used for specific actions and file types.

Terminal Enter the executable path of the application you want to use for your terminal window. The terminal application is used when opening a serial console to the host. By default, Management Center uses an xterm with the following options:

```
xterm -geom 80x25 -T “Console of {host}” -sb -bg black -fg gray -sl 1000 -e /usr/bin/
telnet {system.rna.host} {port}
```

Note

The Management Center terminal field supports the use of the following variables:

- `{host}` The host name used to set the console name (optional).
- `{system.rna.host}` The host name of the Master Host (required).
- `{port}` The dynamic port set by the Master Host (required).

Management Center uses any terminal that supports spawning an external command (usually the '`-e`' flag). The full path to the terminal and the '`-e /usr/bin/telnet {system.rna.host}{port}`' statement are the only requirements. All other items are optional. Consider the following examples:

Cygwin terminal on Windows:

```
C:\cygwin\bin\rxvt.exe -sr -sl 10000 -fg white -bg black -fn fixedsys -fb fixedsys -T
“Console of {host}” -tn cygwin -e /usr/bin/telnet {system.rna.host} {port}
```

Simple white xterm on Linux:

```
/usr/bin/xterm -e /usr/bin/telnet {system.rna.host} {port}
```

Gnome-terminal on Linux:

```
/usr/bin/gnome-terminal -t “Console of {host}” -e /usr/bin/telnet {system.rna.host}
{port}
```

Note

If you use Konsole or Gnome-terminal, you can use the default settings used by your desktop.

HTML Browser Enter the executable path of the application you want to use as your HTML browser. On Linux, the default browser is Firefox. On Windows, Management Center uses your default browser.

PDF Viewer Enter the executable path of the application you want to use to view PDFs such as the SGI Management Center System Administrator’s Guide or Release Notes. On Linux, the defaults are Acrobat Reader then xpdf. On Windows, Management Center uses your default PDF reader.

Provisioning Settings

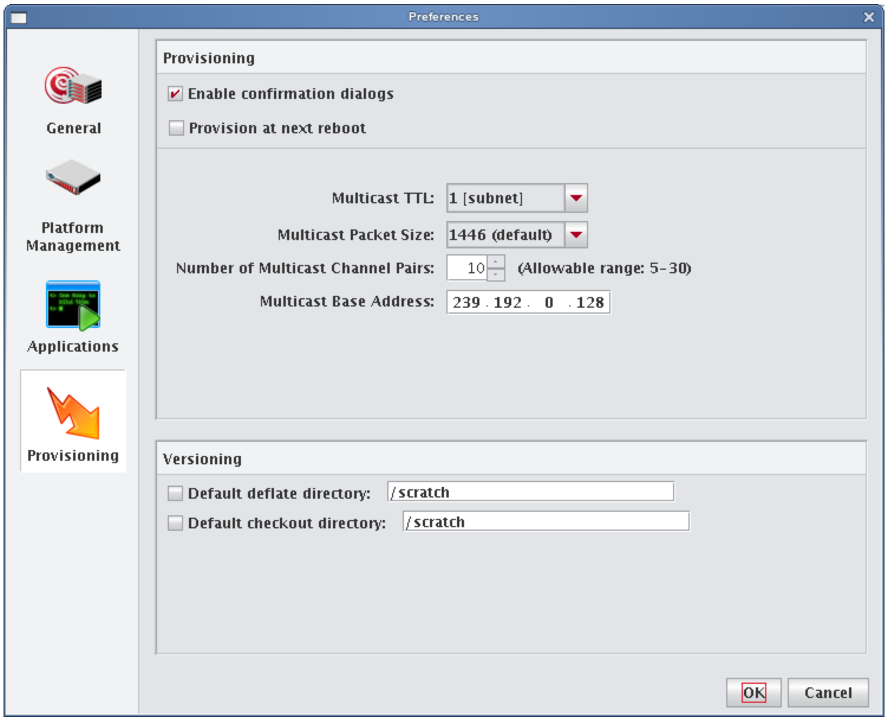

Provisioning

☑ Enable confirmation dialogs

☐ Provision at next reboot

Multicast TTL: 1 [subnet]

Multicast Packet Size: 1446 (default)

Number of Multicast Channel Pairs: 10 (Allowable range: 5-30)

Multicast Base Address: 239 . 192 . 0 . 128

Versioning

☐ Default deflate directory: /scratch

☐ Default checkout directory: /scratch

Provisioning

These settings let you control the default provisioning behavior.

Note

You can overwrite these settings from the Advanced Provisioning dialog. See *Advanced Provisioning Options* on page 151.

Enable Confirmation Dialogs Select this option if you want to display a confirmation dialog when you provision hosts.

Provision at Next Reboot When checked, hosts are not provisioned until you reboot them manually or with a script. When unchecked, Management Center automatically restarts hosts or powers them down to begin the provisioning process.

Multicast TTL Sets the Multicast TTL or *Time-To-Live* on a multicast packet. The default, 1, restricts multicast packets to the subnet (the cluster's internal network). If you are using multicast across networks and multiple switches across a private network, select 32. If you plan to use multicast across a company WAN, use 64 (the maximum TTL that multicast supports).

Multicast Packet Size Sets the maximum size of multicast packets (by default, 1446).

Number of Multicast Channel Pairs Management Center uses one channel for downloading the kernel and RAMdisk, and another channel for downloading the payload. Typically, you will need only one channel per image used; however, depending on the number of images in use on the system, you may require additional multicast channels. If you run out of channels, a "No Available Channels" error occurs when you attempt to provision. By default, 10 channel pairs are configured on your system.

Multicast Base Address The multicast base address specifies what multicast subnet you will use, starting at the last octet and increasing by 1. By default, Management Center sets the base multicast address to 239.192.0.128 with 10 channels, which uses addresses from 239.192.0.128-137. If you have multiple Management Center Master Hosts on the same network, they should use a different subnet or different ranges within that subnet. For example, Master 1 might use 239.192.0.128-137 and Master 2 might use 239.192.1.128-137. Other multicast ranges such as 224.0.0.x may also be suitable for your network.

! Warning!

If you change your multicast base address, you must verify that the multicast default route includes the new base address. See *Configure Multicast Routes* on page 45 for information on configuring multicast routes.

Note

After changing multicast settings, you must restart your server.

Specify the Download Path of the Payload During the provisioning process, Management Center downloads the payload to the host's root directory. Depending on the size of the payload, this may require a very large root partition. To use a smaller root partition, you may download the payload to a different partition by specifying the *image.path* in *$MGR_HOME/etc/ProvisioningService.profile*:

```
# Uncomment this variable to manually set the download path
# of the provisioning image file (__image__). The default
# path (/mnt) is shown in the example below
# image.path=/mnt
```

Note

You should be aware of the following:

1. The directory you select must belong to its own partition (for example, if you are downloading to */scratch*, it must belong to its own partition).

2. During this point of the provisioning process, the file system is still mounted by the RAMdisk. Because of this, you must include */mnt* in the *image.path*. For example, to mount */scratch*, the *image.path* would be */mnt/scratch*.

In either case, Management Center reverts to the root partition if the partition doesn't exist or if the path is wrong.

Versioning

These options allow you to configure default directories used to check items in and out of VCS and to open large files created when importing a payload.

Default Checkout Directory When enabled, Management Center uses this directory as a scratch directory for checking items in and out of VCS. Use this if you have limited space on the partition containing $MGR_HOME.

Default Deflate Directory When enabled, this option allows you to specify an alternate path in which to open large files created when importing a payload. Use this if you have limited space on the partition containing $MGR_HOME.

Configuring IPMI

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 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 ipmitool_options.profile

By default, the Management Center uses the Lanplus interface to send remote commands to IPMI BMC(s). Use the file *$MGR_HOME/etc/ipmitool_options.profile* if the BMC needs specific options to be passed on to the ipmitool command line (for example, oem type or encryption settings). After making any changes, enter **/etc/init.d/mgr restart** to restart the services.

Example:

```
# Use standard options globally
ipmitool.power._default_=-I lanplus
ipmitool.status._default_=-I lanplus
ipmitool.sol._default_=-I lanplus


# Use Intel OEM Options for Node n015
ipmitool.power.n015=-I lanplus -o intelplus
ipmitool.status.n015=-I lanplus -o intelplus
ipmitool.sol.n015=-I lanplus -o intelplus
```

Configure the Payload and Kernel

Before you can begin using IPMI, you must configure the kernel and payload to support IPMI as follows:

1. Add the following modules (available in *drivers/char/ipmi* under the kernel modules tree) to the kernel with which you will be provisioning:

   ```
   ipmi_devintf
   ipmi_si
   ipmi_msghandler
   ```

2. In the kernel parameters, set the serial console and baud rate.

   For SGI clusters, the defaults are **ttyS1** and **115200**, respectively.

3. Install either OpenIPMI or Freeipmi into the payload if neither is already installed.You can obtain OpenIPMI from the SLES or RHEL distribution CD/DVD. Freeipmi can be obtained from *http://www.gnu.org/software/freeipmi*.

4. If you are using OpenIPMI, run the following command to enable the ipmi daemon on the master host:

   ```
   chroot $MGR_HOME “chkconfig ipmi on”
   ```

Note

OpenIPMI requires the kernel binary RPM installed in the payload in order for the ipmi daemon to run properly.

Configure the Master Host and Management Center

The following section describes how to configure your Master Host to use IPMI and how to prepare your system (through Management Center) for IPMI control:

1. Install Ipmitool on the Master Host to allow you to perform IPMI-related tasks such as powering off hosts, executing beacon operations, and activating SOL.

Note

The SDR cache is created in *$MGR_HOME/ipmi/sdrcache.dat* on each host. If the *$MGR_HOME/ipmi* directory or the *sdrcache.dat* file cannot be created, monitoring will fail.

2. Start Management Center.
3. Create a new user (see *Adding a User* on page 70).
4. Assign the new user the name and password configured for BMC controllers (for SGI systems, *admin* and *ipmi*).
   This gives you full access to IPMI controls on the hosts.
5. Assign the user to the power group and make power the primary group for the user.

Note

This user is not required for monitoring temperature and fans but is required for power control and beaconing. This user cannot log into Management Center.

6. In the Platform Management pane, select Override Global Settings.
7. Select IPMI as the Platform Management Device Type.
8. 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.
9. (Optional) Select the MAC Address vs. Host MAC Address type:
   - **A.** Not Related
   - **B.** Greater Than
   - **C.** Less Than
10. (Optional) Enter the MAC Address Offset.
11. Select the MAC Address to use to manage this host.
12. (Optional) Select the IP Address vs. Host IP Address type:
    - **A.** Greater Than
    - **B.** Less Than
    - **C.** Same As

13. (Optional) Enter the IP address offset from the management interface.

14. (Optional) Enter the IP address for the host.

15. Select a Platform Management User.

✎ **Note**

Users must belong to Power as their primary group to appear in this list. See *Groups* on page 73.

---

16. Click **OK**.

Configuring 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 Multicast Routes

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
   ```

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
}
```

Chapter 4

Cluster Configuration

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). This host, commonly referred to as the *Management Center 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.

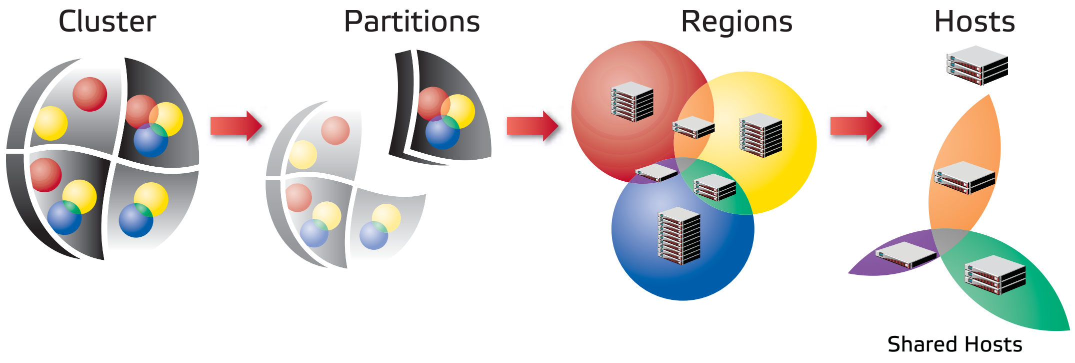

Setting Up Your Cluster

Management Center divides system configuration into several components:

- *Adding Hosts* on page 49
- *Partitions* on page 61
- *Regions* on page 63
- *Racks* on page 66
- *User Administration* on page 69
- *Adding a User* on page 70
- *Groups* on page 73
- *Roles* on page 76

Adding Hosts

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 55.

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.

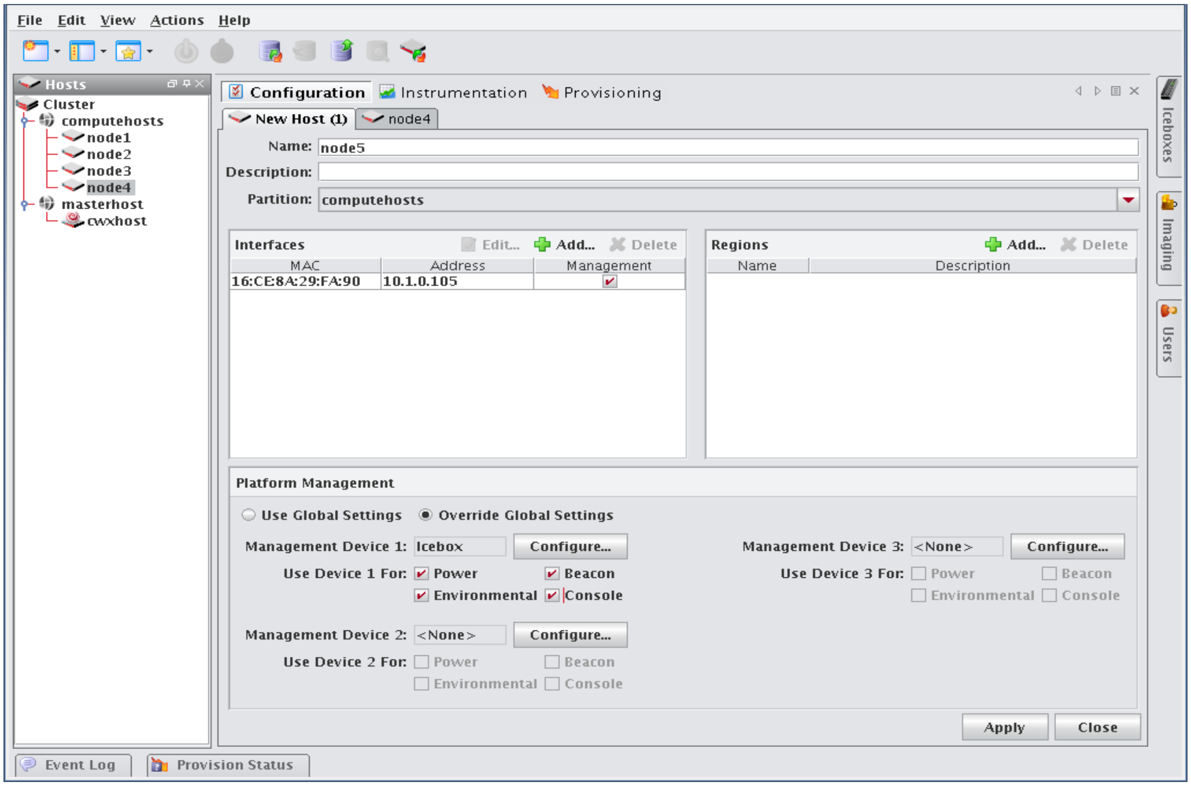
File
Edit
View
Actions
Help
Hosts
Cluster
computehosts
node1
node2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
New Host (1)
node4
Edit...
Add...
Delete
Configure...
Apply
Close
Iceboxes
Imaging
Users
Event Log
Provision Status

Name: node5
Description:
Partition: computehosts

Interfaces

MAC	Address	Management
16:CE:8A:29:FA:90	10.1.0.105	☑

Regions

Name	Description

Platform Management

○ Use Global Settings ● Override Global Settings

Management Device 1: Icebox
Use Device 1 For: ☑ Power ☑ Beacon ☑ Environmental ☑ Console
Management Device 2: <None>
Use Device 2 For: ☐ Power ☐ Beacon ☐ Environmental ☐ Console
Management Device 3: <None>
Use Device 3 For: ☐ Power ☐ Beacon ☐ Environmental ☐ Console

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 or 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.

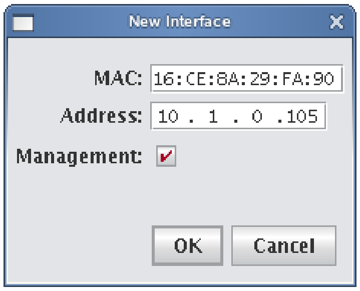

MAC: 16:CE:8A:29:FA:90
Address: 10 . 1 . 0 .105
Management: ☑

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.

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 50). 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.

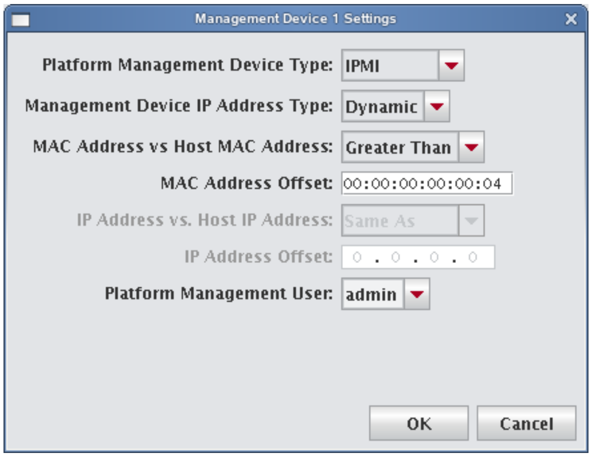

Platform Management Device Type: IPMI
Management Device IP Address Type: Dynamic
MAC Address vs Host MAC Address: Greater Than
MAC Address Offset: 00:00:00:00:00:04
IP Address vs. Host IP Address: Same As
IP Address Offset: 0 . 0 . 0 . 0
Platform Management User: admin

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.

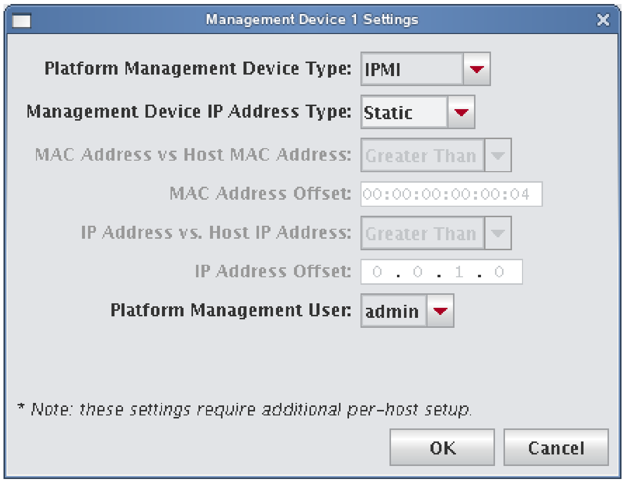

Management Device 1 Settings

Platform Management Device Type: IPMI
Management Device IP Address Type: Static
MAC Address vs Host MAC Address: Greater Than
MAC Address Offset: 00:00:00:00:00:04
IP Address vs. Host IP Address: Greater Than
IP Address Offset: 0 . 0 . 1 . 0
Platform Management User: admin

* Note: these settings require additional per-host setup.

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.

Edit a Host

Editing hosts allows you to change information previously saved about a host, edit host configurations, or move hosts in and out of partitions and regions.

To Edit a Host

1. Select a host from the host navigation tree. (To select multiple hosts, use the **Shift** or **Ctrl** keys.)
2. Select **Edit** from the **Edit** menu or right-click the hosts in the navigation tree and select **Edit**.

   Management Center displays the host pane for each selected host. From this view, you can make changes to the hosts.

! Warning!

Changing the name of the Master Host may prevent the cluster from functioning correctly. For information on changing the name of the Master Host, see *Renaming the Management Center Master Host* on page 54.

---

3. Click **Apply**.

Renaming the Management Center Master Host

Before changing the name of the Master Host, consider applications that require the use of this name (for example, job schedulers, mpi "machines" files, and other third-party software). In some cases, you may need to consult with application vendors regarding special instructions for changing the name of the Master Host.

When you change the Master Host name, all Management Center services, hosts, and clients must be able to resolve the new name. To ensure that your system functions properly after renaming the Master Host, you must update the host name in several files. To rename your Master Host:

1. Select the Master Host in the host navigation tree.
2. Select **Edit** from the **Edit** menu or right-click on the Master Host and select **Edit**. Management Center displays the host pane.
3. In the host pane, enter a new name and click **Apply**.
4. Exit Management Center.
5. In a command line, enter `/etc/init.d/mgr stop` to shut down Management Center services on the system.
6. On the Master Host, edit the *$MGR_HOME/@genesis.profile* to use the new name (*system.rna.host*).
7. On the Master Host, edit the *$MGR_HOME/etc/Activator.profile* and change all instances of the host name to use the new name.
8. Add the new Master Host name to the alias list in */etc/hosts*. For example:

   ```
   10.168.18.3 host.sgi.com  host  <new_name>
   ```
9. Restart Management Center.

Find a Host

To Find a Host in the Host Navigation Tree

1. Select **Find** from the **Edit** menu.

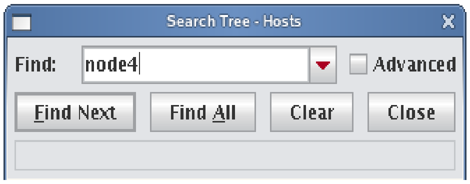
Search Tree - Hosts
Find:
node4
☐ Advanced
Find Next
Find All
Clear
Close

2. Enter the name of a host and click **Find Next**.
3. (Optional) Click **Advanced** to enable more extensive search options.

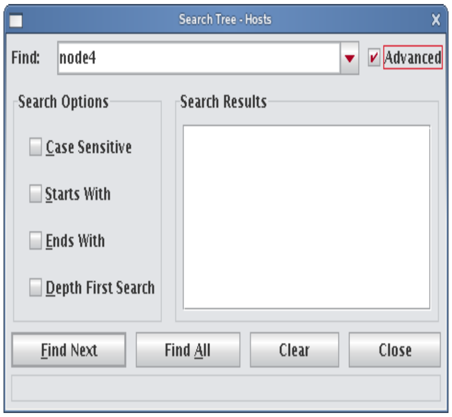
Search Tree - Hosts
Find:
node4
☑ Advanced
Search Options
☐ Case Sensitive
☐ Starts With
☐ Ends With
☐ Depth First Search
Search Results
Find Next
Find All
Clear
Close

Delete a Host

Deleting a host removes it from the cluster.

To Delete a Host

1. Select the host you want to delete from the host navigation tree. (To select multiple hosts, use the Shift or Ctrl keys).
2. Select **Delete** from the **Edit** menu or right-click the selected hosts in the navigation tree and select **Delete**.
   Management Center asks you to confirm your action.
3. Click **OK** to delete the hosts.

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.

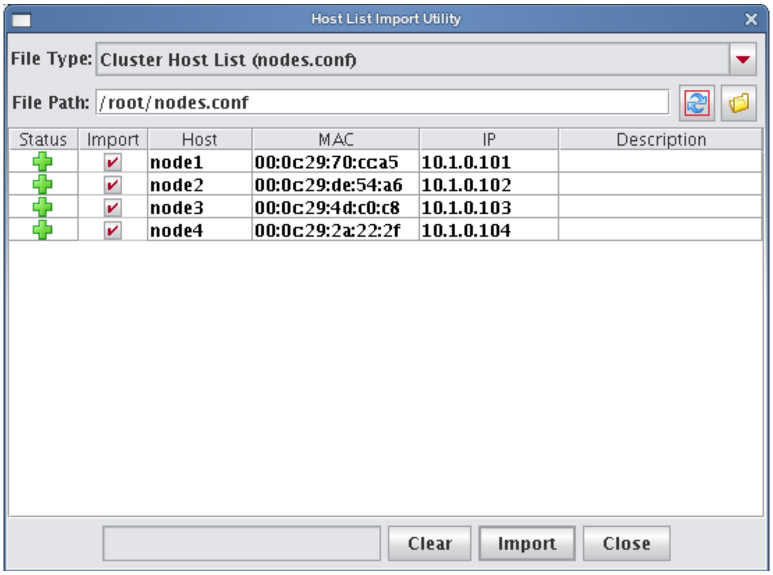

File Type: Cluster Host List (nodes.conf)

File Path: /root/nodes.conf

Status	Import	Host	MAC	IP	Description
✚	☑	node1	00:0c29:70:cca5	10.1.0.101	
✚	☑	node2	00:0c29:de:54:a6	10.1.0.102	
✚	☑	node3	00:0c29:4d:c0:c8	10.1.0.103	
✚	☑	node4	00:0c29:2a:22:2f	10.1.0.104	

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**.

Host Power Controls

The Power Management feature provides you with the ability to remotely reset, power up, power down, and cycle power to hosts installed in your system. Power status information for each host is available through the instrumentation tab. See *Overview Tab* on page 156 and *Thumbnail Tab* on page 157.

System

The System options in the right-click menu execute power-related events on the hosts.

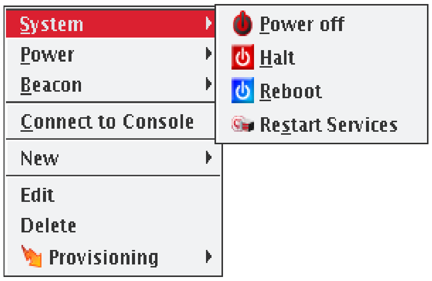

POWER OFF

Issues the Linux `/sbin/poweroff` command to stop all applications and services running on the host and, if the hardware allows, to power off the host. If you have used the `/sbin/shutdown` command to successfully shut down and reboot hosts at the next power cycle, you should be safe to enable this option. To enable shutdown, set the **shutdown.button.enable** option in *HostAdministrationService.profile* to true.

! Warning!

Using the shutdown option requires that the BIOS is enabled to support boot at power up — the default behavior for LinuxBIOS. This setting, also referred to as Power State Control or Power On Boot, is typically enabled for most server-type motherboards.

If you do not enable this BIOS setting, hosts that are shut down may become unusable until you press the power button on each host. For the location of your host power switch, please consult your host installation documentation.

Note

The power connection to the host remains active unless you click Off. To return the host to normal operational status, cycle the power.

HALT

Issues the Linux `/sbin/halt` command to stop all applications and services running on the host and, if the hardware allows, power off the host.

REBOOT

Shuts down and restarts all applications and services on the host.

RESTART SERVICES

Restarts the Management Center services on the selected hosts.

Note

You cannot restart Management Center services on the Master Host from the GUI. You must perform this action from the CLI.

Power

The Power options in the right-click menu execute power-related events from your power management device.

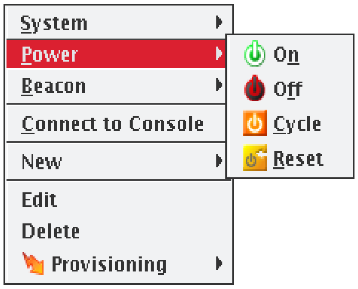
System
Power
Beacon
Connect to Console
New
Edit
Delete
Provisioning
On
Off
Cycle
Reset

ON

Turn on power to the host.

Note

If you are unable to power a host on or off, the port may be locked.

OFF

Immediately turn off power to the host.

CYCLE

Turn off the power, then back on. This is useful for multiple hosts.

RESET

Send a signal to the motherboard to perform a soft boot of the host.

Beacon

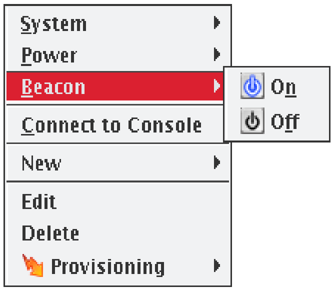
System
Power
Beacon
Connect to Console
New
Edit
Delete
Provisioning
On
Off

BEACON ON

To identify a specific host in a cluster for troubleshooting purposes, click Beacon On to flash a light from the host. Use the Shift and Ctrl keys to select multiple hosts.By default, the beacon icon appears next to the selected host(s) for 180 seconds. You can change this default time by changing the `timeout.beacon.seconds` parameter in file $MGR_HOME`/etc/PlatformManagementService.profile`.

Note

The beacon function works only if the hardware installed in your cluster supports beacons (i.e., the hosts support IPMI or ILO).

BEACON OFF

Turn off the beacon.

Console

Connecting to the console allows you to monitor activity on a host-by-host basis. When you connect to the console, Management Center opens a terminal window for each host and allows you to view host activity or execute bash and other general command-line operations necessary for troubleshooting. You can also use the console to apply specific configurations or enhancements to a payload that you can import and use at a later time.

To Connect to the Console

Note

Before you can connect to a console, you must configure the platform management settings for your hosts to direct them to the serial device they will use (such as IPMI or Conman) and enable the **Console** option. You must install and configure conman it before you can use it. See *Conman* on page 32 and *conman* on page 226 for additional information about configuration and CLI controls.

1. Select the host on which you want to open a console from the host navigation tree.

   To select multiple hosts, use the *Shift* or *Ctrl* keys.
2. Right-click on a selected host and select **Connect to Console**.
   A console opens for each host.
3. Enter bash or other general CLI commands as needed to configure the host.
4. When finished, close the console.

Roamer KVM

When using Roamer-enabled nodes, you can connect to the Roamer KVM from the Management Center GUI. When you connect to the KVM console, Management Center opens a console window using Java. This allows you to control the host as you would with a keyboard, monitor, and mouse.

Note

Before you can connect to a console, you must configure the platform management settings for your host to use Roamer and enable the **Console** option. This configures the host to use Roamer for the serial console.

To connect to the Roamer KVM, use the following steps:

1. Select the host on which you want to open a console from the host navigation tree.
2. Right-click on a selected host and select **Connect to KVM**.
   A Java KVM console opens for the host.
3. When finished, close the console.

Partitions

You can use partitions to group clusters into non-overlapping collections of hosts. Instrumentation, provisioning, power control, and administrative tasks can be performed on this collection of hosts by selecting the partition in the host tree.

Adding Partitions

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

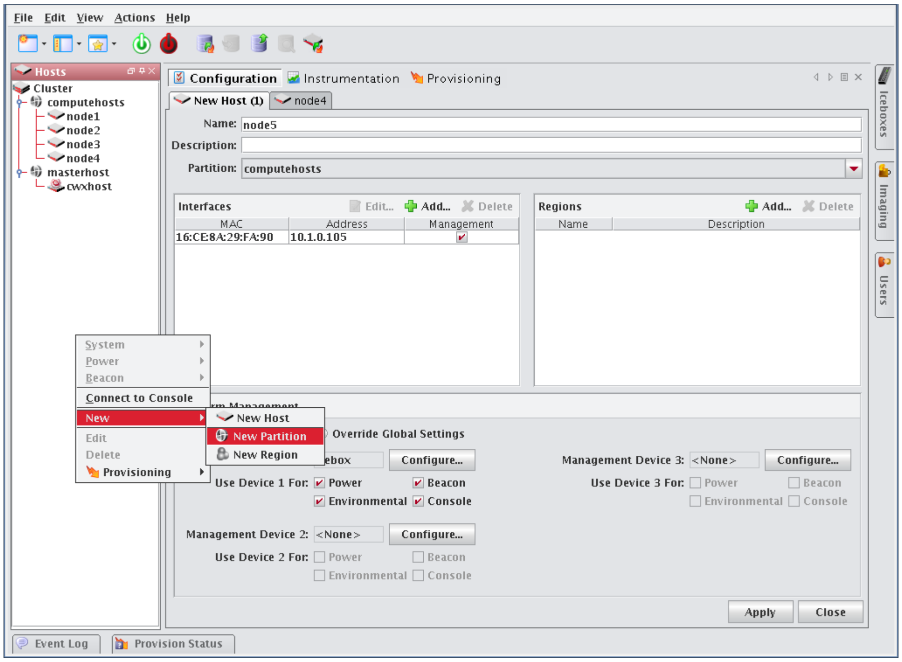
File Edit View Actions Help
Hosts
Cluster
computehosts
node1
node2
node3
node4
masterhost
cwxhost
System
Power
Beacon
Connect to Console
New
Edit
Delete
Provisioning
New Host
New Partition
New Region
Configuration
Instrumentation
Provisioning
New Host (1)
node4
Interfaces
Edit...
Add...
Delete
Regions
Name
Description
Override Global Settings
Configure...
Apply
Close
Event Log
Provision Status
Iceboxes
Imaging
Users

Name: node5
Description:
Partition: computehosts

MAC	Address	Management
16:CE:8A:29:FA:90	10.1.0.105	☑

Use Device 1 For: ☑ Power ☑ Beacon ☑ Environmental ☑ Console
Management Device 2: <None>
Use Device 2 For: ☐ Power ☐ Beacon ☐ Environmental ☐ Console
Management Device 3: <None>
Use Device 3 For: ☐ Power ☐ Beacon ☐ Environmental ☐ Console

2. Enter a partition name.
3. (Optional) Enter a description.
4. In the **Hosts** pane, click **Add** to display the Select Hosts dialog.

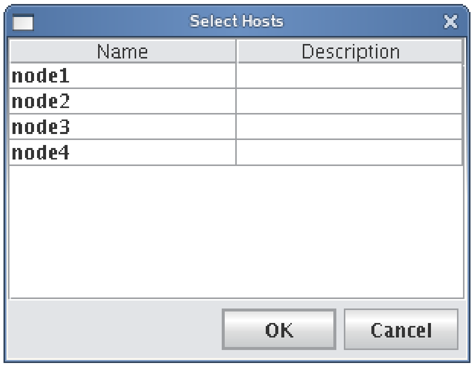
Select Hosts
OK
Cancel

Name	Description
node1	
node2	
node3	
node4	

5. Select hosts to add to this partition and click **OK**.
6. Click **Apply**.

Editing Partitions

Editing a partition allows you to change previously saved information about a partition. You can edit or remove regions, alter partition configurations, disable partitions, or remove partitions from the host.

1. Select a partition from the host navigation tree.
2. Select **Edit** from the **Edit** menu or right-click on the partitions in the host navigation tree and select **Edit**.
3. Use the **Partition** pane to make changes to the partition).

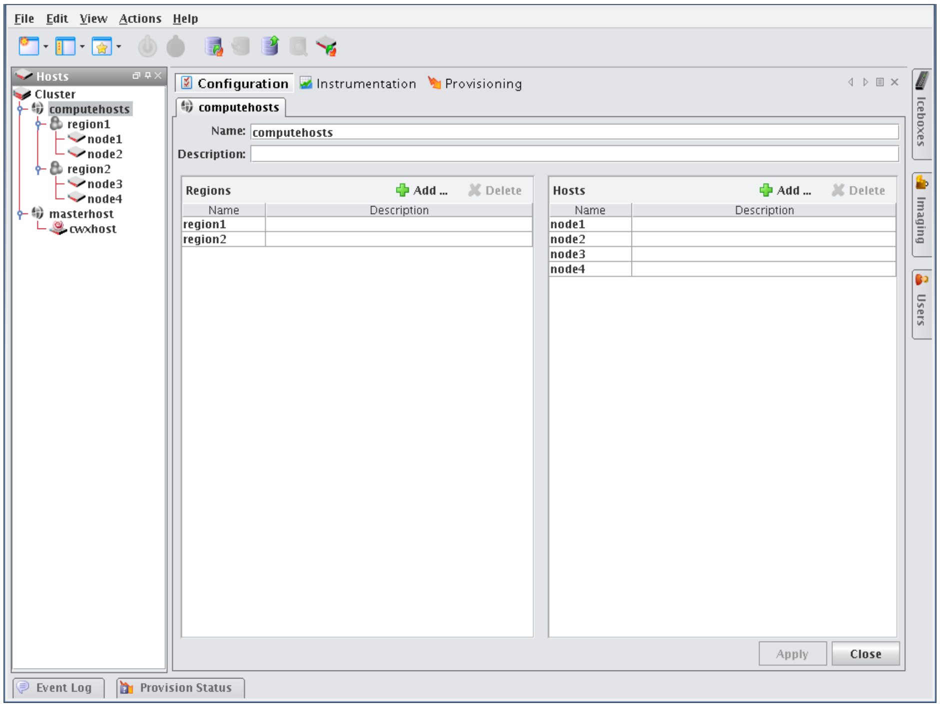
File
Edit
View
Actions
Help
Hosts
Cluster
computehosts
region1
node1
node2
region2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
computehosts
Regions
Add ...
Delete
Hosts
Add ...
Delete
Apply
Close
Event Log
Provision Status
Iceboxes
Imaging
Users

Name: computehosts
Description:

Name	Description
region1	
region2	

Name	Description
node1	
node2	
node3	
node4	

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

Deleting Partitions

Deleting a partition allows you to remove unused partitions from the system.

Note

If you delete a partition, all regions and hosts associated with the partition will move to the default partition. To delete regions and hosts, refer to *Regions* on page 63 and *Adding Hosts* on page 49.

1. Select the partitions you want to delete from the host navigation tree.
2. Select **Delete** from the **Edit** menu or right-click on the partitions in the navigation tree and select **Delete**.
3. Click **OK** to delet the partitions.

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.

Creating Regions

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

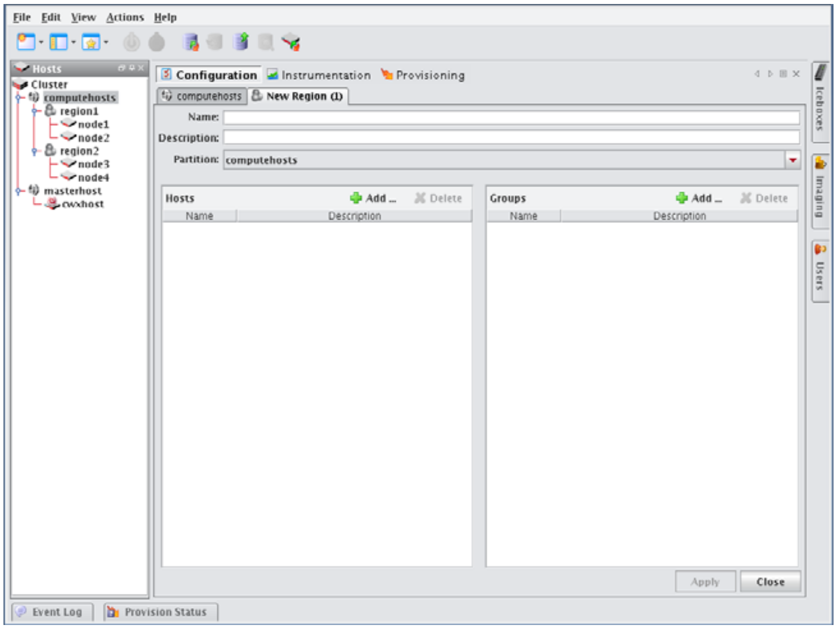
File
Edit
View
Actions
Help
Hosts
Cluster
computehosts
region1
node1
node2
region2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
computehosts
New Region (1)
Name:
Description:
Hosts
Add ...
Delete
Name
Description
Groups
Add ...
Delete
Name
Description
Apply
Close
Event Log
Provision Status
Iceboxes
Imaging
Users

Partition: computehosts

2. Enter the name.
3. (Optional) Enter a description.
4. (Optional) Select the name of the partition you want to assign the region to from the drop-down list.

Note

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

5. In the **Hosts** pane, click **Add** to assign hosts to the region.
6. In the Select Hosts dialog, select the hosts you want to add to the region.

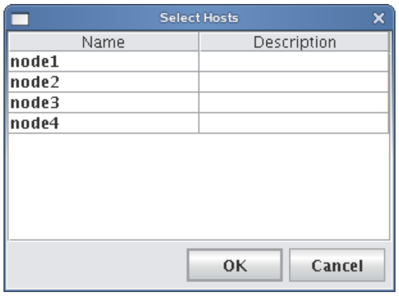
Select Hosts
OK
Cancel

Name	Description
node1	
node2	
node3	
node4	

7. Click **OK** to add the hosts.

8. In the **Groups** pane, click **Add**.

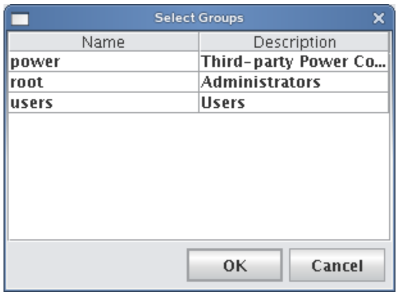
Select Groups
OK
Cancel

Name	Description
power	Third-party Power Co...
root	Administrators
users	Users

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** to add the groups.

11. Click **Apply**.

Editing Regions

Editing regions allows you to change previously saved information about a region or to modify region memberships by adding or removing groups or hosts.

1. Select a region from the host navigation tree.

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

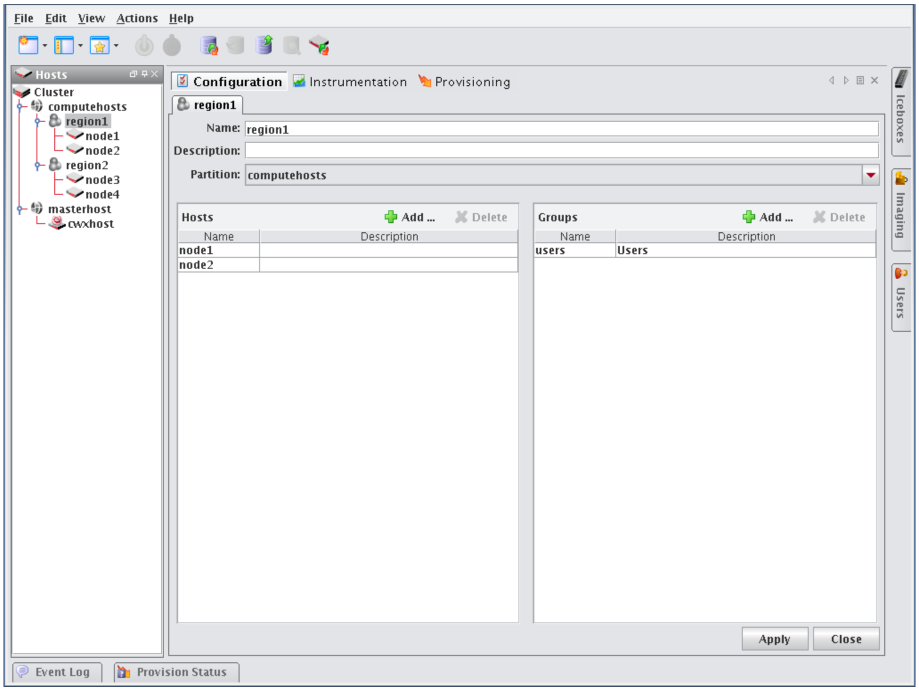
File Edit View Actions Help
Hosts
Cluster
computehosts
region1
node1
node2
region2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
region1
Add ...
Delete
Add ...
Delete
Apply
Close
Iceboxes
Imaging
Users
Event Log
Provision Status

Name: region1
Description:
Partition: computehosts

Hosts

Name	Description
node1	
node2	

Groups

Name	Description
users	Users

3. Make changes to the regions (such as adding or deleting hosts).

4. Click **Apply**.

Deleting Regions

Deleting a region allows you to remove unused regions from the system.

1. Select the region you want to delete from the host navigation tree.
2. Select **Delete** from the **Edit** menu or right-click on the regions in the navigation tree and select **Delete**.
3. Click **OK** to remove the regions.

Note

If you delete a region, all hosts associated with the region return to the partition to which the region belonged. If the region was not part of a partition, the hosts move to the default partition.

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**.

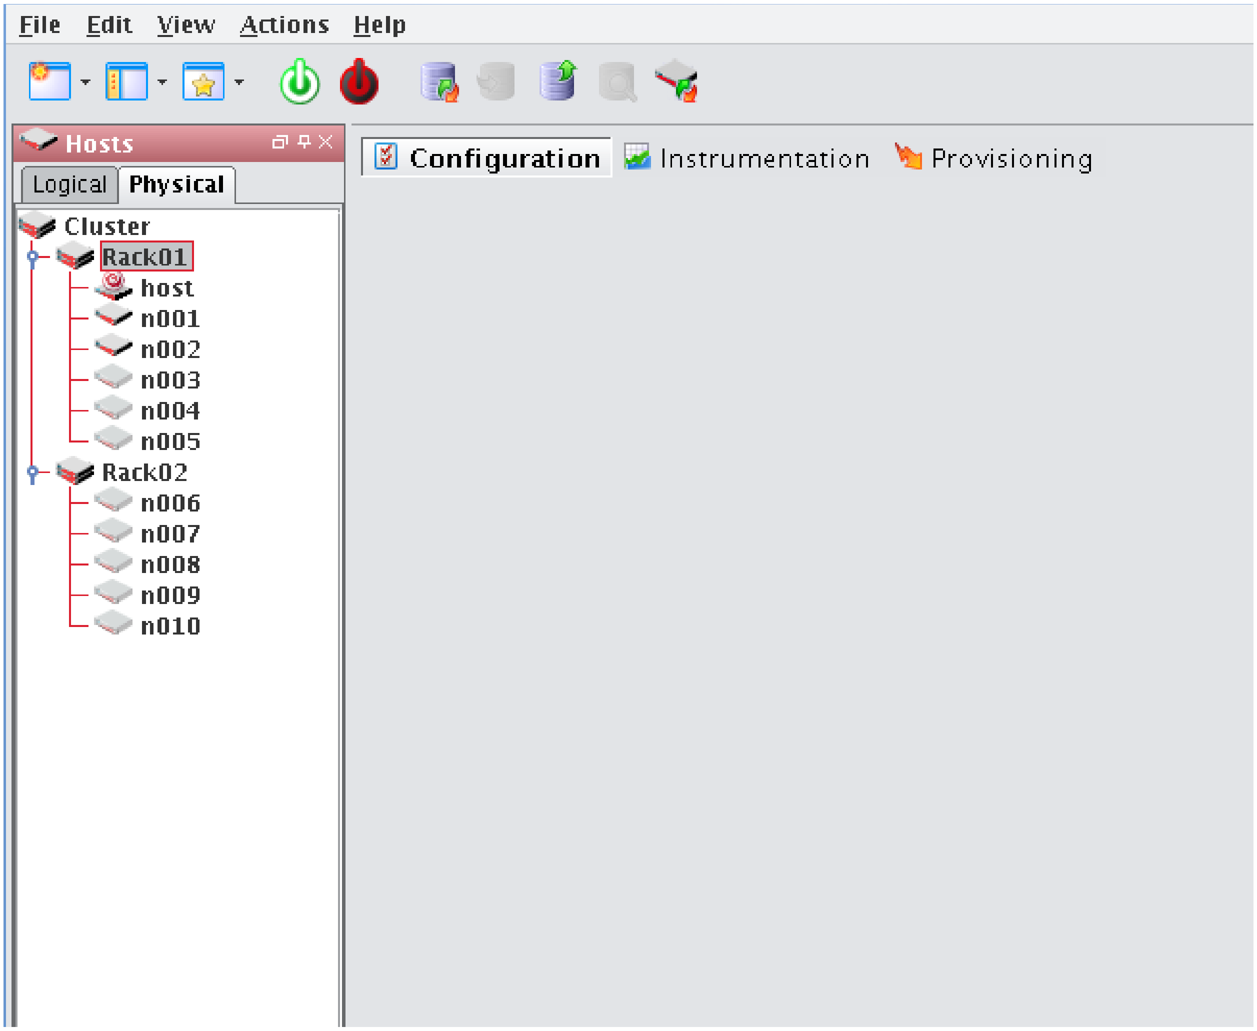
File
Edit
View
Actions
Help
Hosts
Logical
Physical
Cluster
Rack01
host
n001
n002
n003
n004
n005
Rack02
n006
n007
n008
n009
n010
Configuration
Instrumentation
Provisioning

Adding Racks

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**.

Editing Racks

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.

Deleting Racks

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 5

User Administration

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 group(s) to which they belong.

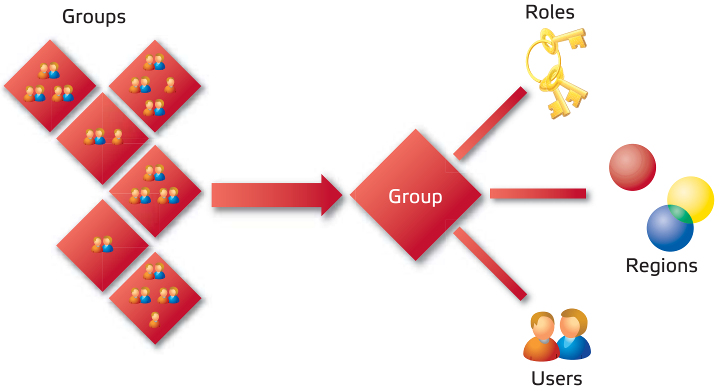

For example, a user assigned to multiple groups (as indicated by the following diagram) has different rights and privileges within each group. This flexibility allows you to establish several types of user roles: full administration, group administration, user, or guest.

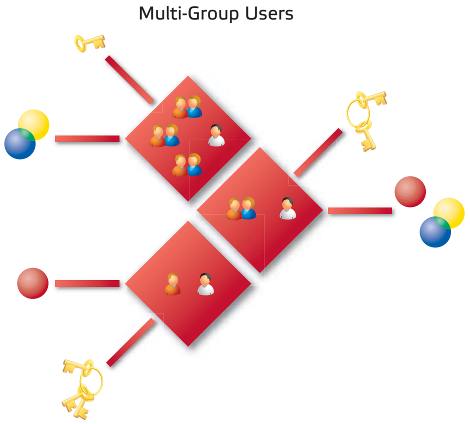

Note

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, ensure that the user and group IDs (UIDs and GIDs, respectively) match up between the accounts on the Master Host and Management Center. Otherwise, NFS may not work properly.

Default User Administration Settings

Management Center implements the following structure during the installation process:

- The root and guest user accounts are created.
- The root, power, and users groups are created.
- The root and user roles are created.
- All privileges allowed by the installed license are created.

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.

Adding 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**.

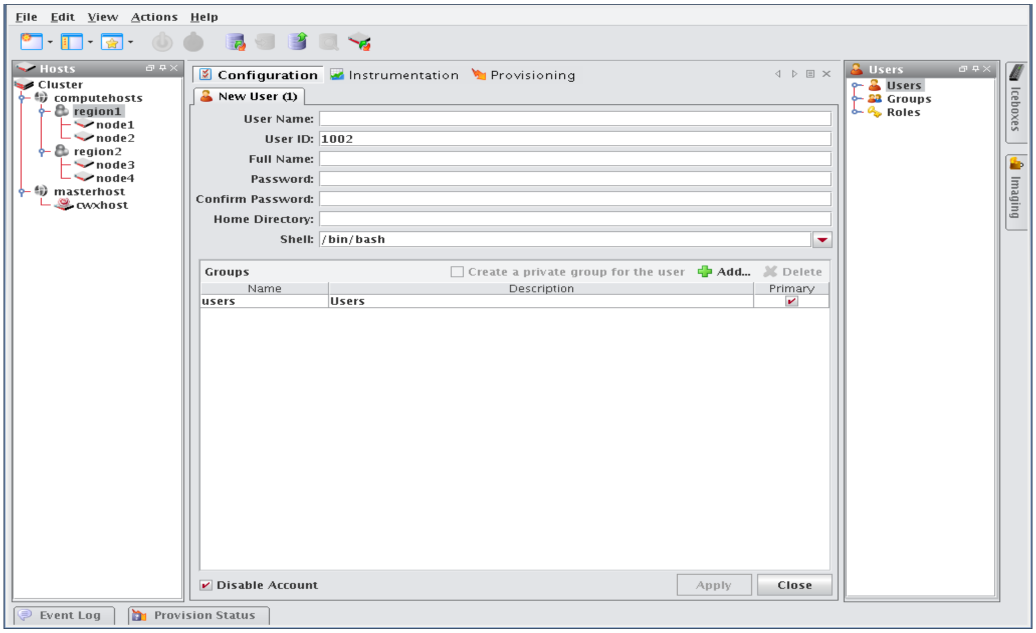
File
Edit
View
Actions
Help
Hosts
Cluster
computehosts
region1
node1
node2
region2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
New User (1)
User Name:
Full Name:
Password:
Confirm Password:
Home Directory:
Groups
Create a private group for the user
Add...
Delete
Name
Description
Primary
Apply
Close
Users
Users
Groups
Roles
Iceboxes
Imaging
Event Log
Provision Status

User ID: 1002

Shell: /bin/bash

Name	Description	Primary
users	Users	☑

☑ Disable Account

2. Enter the user’s login name.
3. (Optional) Management Center assigns a system-generated user ID. Enter any changes to the ID in the User ID field.
4. Enter the user’s first and last name in the **Full Name** field.

Note

The Management Center UID must match the system UID.

Tip

If a user already has an account and you would like to apply the account to the Master Host and compute hosts, add the user to your payload during payload creation. When you provision, Management Center creates the account on the hosts. See *Payload Local User and Group Account Management* on page 98.

5. Enter and confirm a user password.
6. (Optional) Enter a home directory (for example, */home/username*).
7. (Optional) Enter a shell for this user or select an existing one from the drop-down list. (By default, Management Center uses */bin/bash*.)
8. Click **Apply**.

Defining User Groups

The groups pane allows you to identify the group(s) 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. See *Roles* on page 76.

1. To add the user to a group, click **Add**.

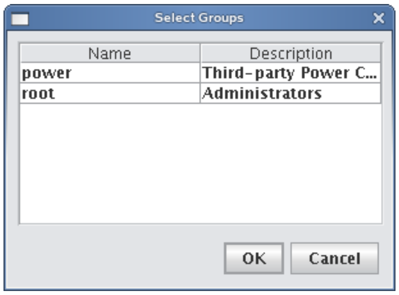

Name	Description
power	Third-party Power C...
root	Administrators

2. Select the groups you want associated with the user.

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. See *Power* on page 73.

3. Click **OK**.
4. (Optional) Select **Create a private group for the user** to create a new group with the same name as the user.
5. (Optional) Check **Disable Account** to prevent users from logging into this account and to exclude this account from future payloads without deleting the account.

Editing User Accounts

Editing a user account allows you to change information previously saved about a user.

1. Select a user from the user navigation tree.
2. Select **Edit** from the **Edit** menu or right-click a user in the navigation tree and select **Edit**.

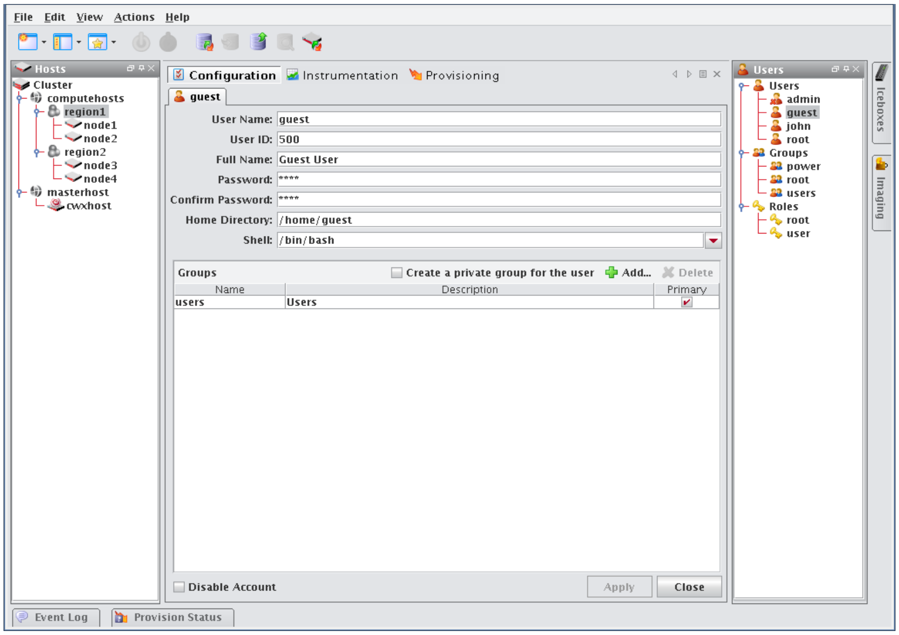
File
Edit
View
Actions
Help
Hosts
Cluster
computehosts
region1
node1
node2
region2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
guest
Users
Users
admin
guest
john
root
Groups
power
root
users
Roles
root
user
Iceboxes
Imaging
Event Log
Provision Status
Apply
Close

User Name: guest
User ID: 500
Full Name: Guest User
Password: ****
Confirm Password: ****
Home Directory: /home/guest
Shell: /bin/bash

Groups ☐ Create a private group for the user Add... Delete

Name	Description	Primary
users	Users	☑

☐ Disable Account

3. Click **Apply**.

Disabling a User Account

Disabling a user account allows you to render the account temporarily inoperative without removing it.

1. Select a user from the user navigation tree.
2. Select **Edit** from the **Edit** menu or right-click a user in the navigation tree and select Edit.
3. Select **Disable Account**.
4. Click **Apply**.

Deleting a User Account

Deleting a user allows you to remove unused user accounts from the system. To temporarily disable a user account, see *Disabling a User Account* on page 72.

Note

You cannot remove the root user.

To Delete a User

1. Select the users you want to delete from the user navigation tree.
2. Select **Delete** from the **Edit** menu or right-click the user names and select **Delete**.
3. Click **OK** to remove the users.

Groups

The following sections outline the fundamentals of adding, editing, and deleting groups. By default, Management Center enables the following groups, but you can create new groups as needed:

Power The power group contains the user names and passwords that will be used to manage IPMI and other 3rd-party power controllers. By default, this group has no role associated with it, so users assigned to this group cannot typically log into Management Center. Although temperature and fan monitoring do not require that a user is assigned to this group, you must assign a user to the power group in order to use power control and beaconing for IPMI-enabled devices.

Note

When using third-party power controls such as IPMI, the power group must be the primary group for *all* users who will access these controls (see *Defining User Groups* on page 71). Users who belong to the power group cannot log into Management Center.

Root The root group typically contains users with full administrative privileges.

Users The users group typically includes all users with access to the cluster. By default, the Users group is associated with the Users role. Management Center automatically assigns all users to the “users” group.

Adding a Group

Adding groups creates a collection of users with shared or similar needs (for example, 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**.

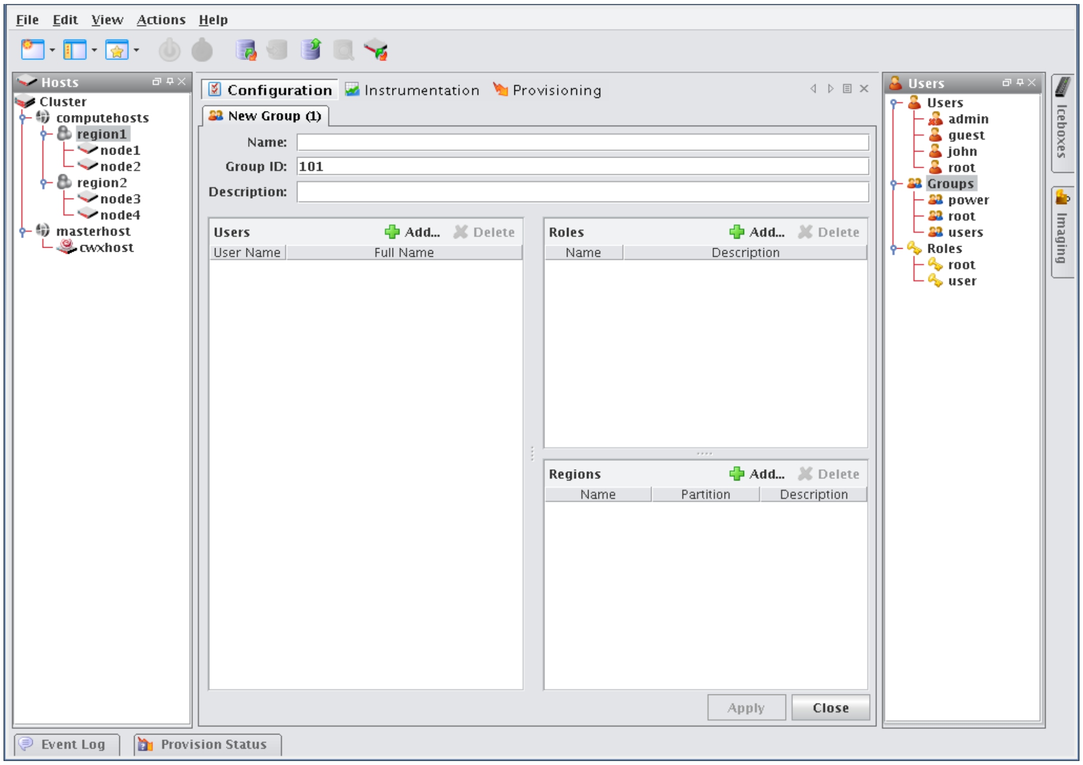
File
Edit
View
Actions
Help
Hosts
Cluster
computehosts
region1
node1
node2
region2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
New Group (1)
Name:
Description:
Users
Add...
Delete
User Name
Full Name
Roles
Add...
Delete
Name
Description
Regions
Add...
Delete
Name
Partition
Description
Apply
Close
Users
Users
admin
guest
john
root
Groups
power
root
users
Roles
root
user
Iceboxes
Imaging
Event Log
Provision Status

Group ID: 101

2. Enter the group name.

3. (Optional) Management Center assigns a system-generated Group ID. Enter any changes to the ID in the Group ID field.
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. 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. See *Roles* on page 76.

1. To add a user to the group, click **Add**.

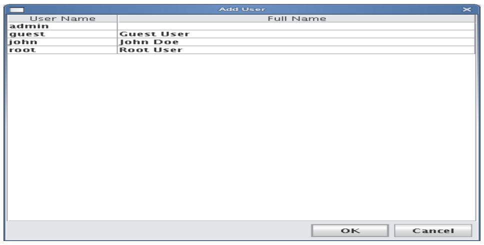

User Name	Full Name
admin	
guest	Guest User
john	John Doe
root	Root User

2. Select the users to add to the group (use the **Shift** or **Ctrl** keys to select multiple users.
3. Click **OK**.

ASSIGN ROLES

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

1. Click **Add** in the **Roles** field.

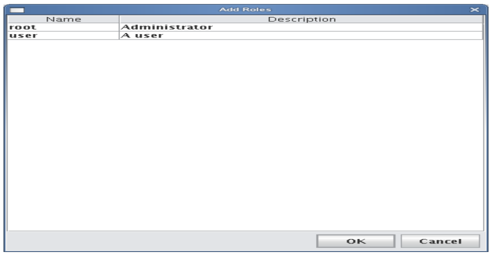

Name	Description
root	Administrator
user	A user

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. See *User Administration* on page 69.

1. Click **Add** in the **Regions** field.

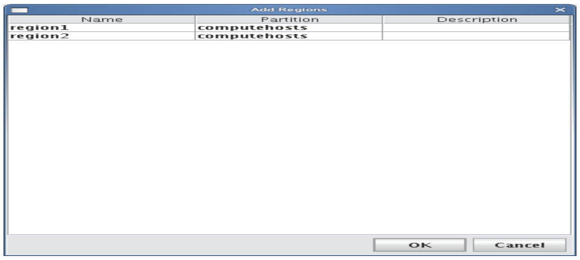

Name	Partition	Description
region1	computehosts	
region2	computehosts	

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

Editing a Group

Editing a group allows you to change previously saved information about a group or modify group memberships by adding or removing users.

1. Select a group from the user navigation tree.
2. Select **Edit** from the **Edit** menu or right-click a group name in the navigation tree and select **Edit**.

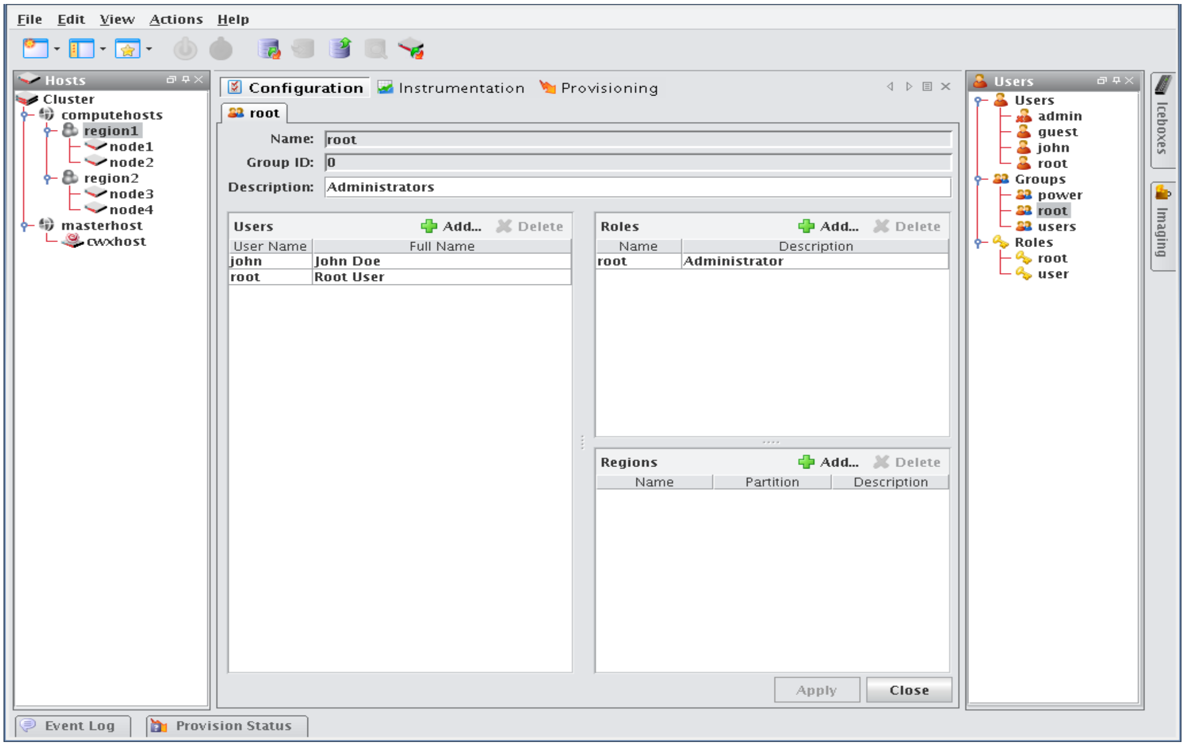
File
Edit
View
Actions
Help
Hosts
Cluster
computehosts
region1
node1
node2
region2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
root
Add...
Delete
Apply
Close
Users
Users
admin
guest
john
root
Groups
power
root
users
Roles
root
user
Iceboxes
Imaging
Event Log
Provision Status

Name: root
Group ID: 0
Description: Administrators

Users

User Name	Full Name
john	John Doe
root	Root User

Roles

Name	Description
root	Administrator

Regions

Name	Partition	Description

3. Make changes by adding or deleting users, roles, and regions as needed.
4. Click **Apply**.

Deleting a Group

Deleting a group allows you to remove unused groups from the system.

1. Select the groups you want to delete from the user navigation tree.
2. Select **Delete** from the **Edit** menu or right-click group names in the navigation tree and select **Delete**.
3. Click **OK**.

Note

You cannot remove the root, power, or users groups.

Roles

The following sections outline the fundamentals of adding, editing, and deleting roles. Roles are associated with groups and privileges, and define the functionality assigned to each group. Several groups can use the same role.

Adding 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**.

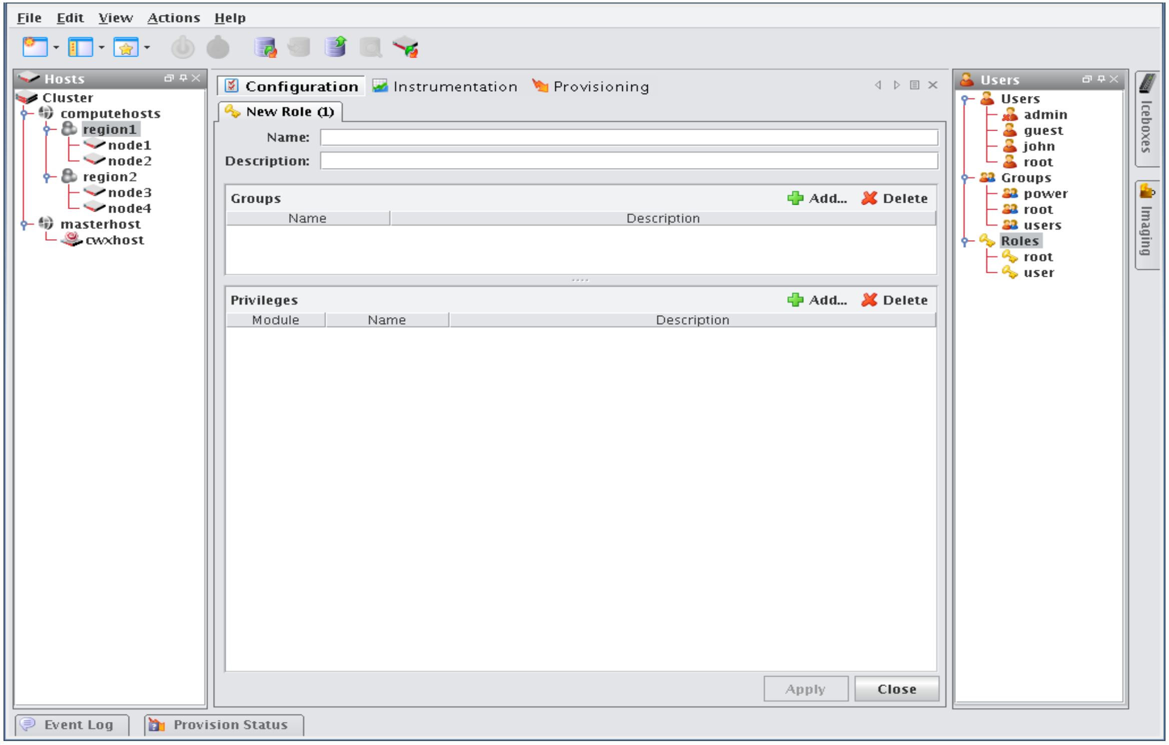
File Edit View Actions Help
Hosts
Cluster
computehosts
region1
node1
node2
region2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
New Role (1)
Name:
Description:
Groups
Add...
Delete
Name
Description
Privileges
Add...
Delete
Module
Name
Description
Apply
Close
Users
Users
admin
guest
john
root
Groups
power
root
users
Roles
root
user
Iceboxes
Imaging
Event Log
Provision Status

2. Enter the 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.

ASSIGNING GROUPS TO ROLES

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

1. Click **Add** in the **Groups** pane.

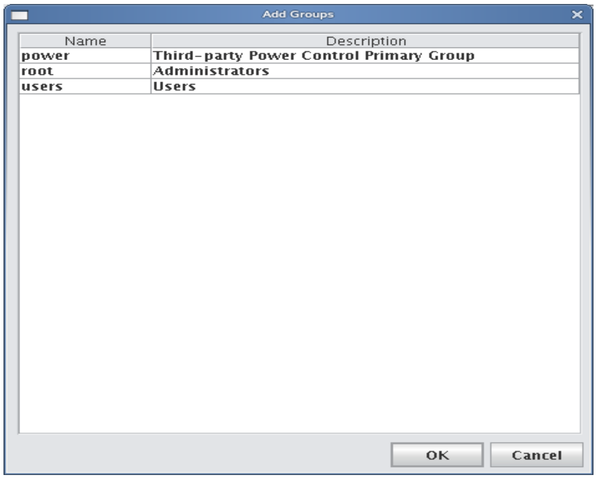

Name	Description
power	Third-party Power Control Primary Group
root	Administrators
users	Users

2. Select the groups you want to assign to the role.
3. Click **OK**.

GRANTING PRIVILEGES

The Privileges pane allows you to assign permissions to a role. Any user with the role will have these permissions in the system. See *Privileges* on page 79.

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

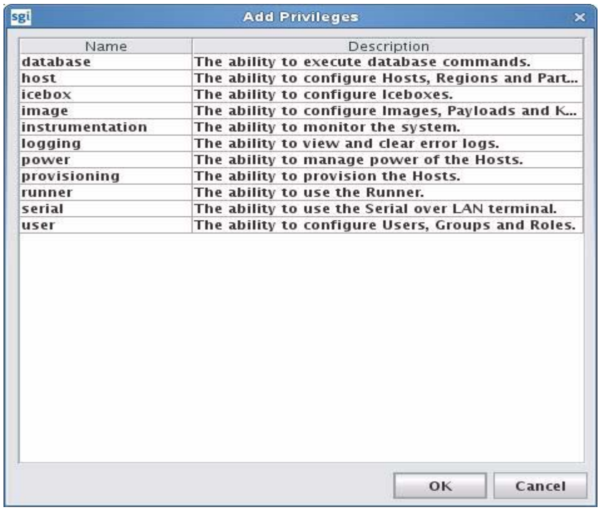

Name	Description
database	The ability to execute database commands.
host	The ability to configure Hosts, Regions and Part...
icebox	The ability to configure Iceboxes.
image	The ability to configure Images, Payloads and K...
instrumentation	The ability to monitor the system.
logging	The ability to view and clear error logs.
power	The ability to manage power of the Hosts.
provisioning	The ability to provision the Hosts.
runner	The ability to use the Runner.
serial	The ability to use the Serial over LAN terminal.
user	The ability to configure Users, Groups and Roles.

3. Click **OK**.

Editing a Role

Editing roles allows you to modify privileges defined for a group.

1. Select a role from the user navigation tree.
2. Select **Edit** from the **Edit** menu or right-click role names in the navigation tree and select **Edit**..

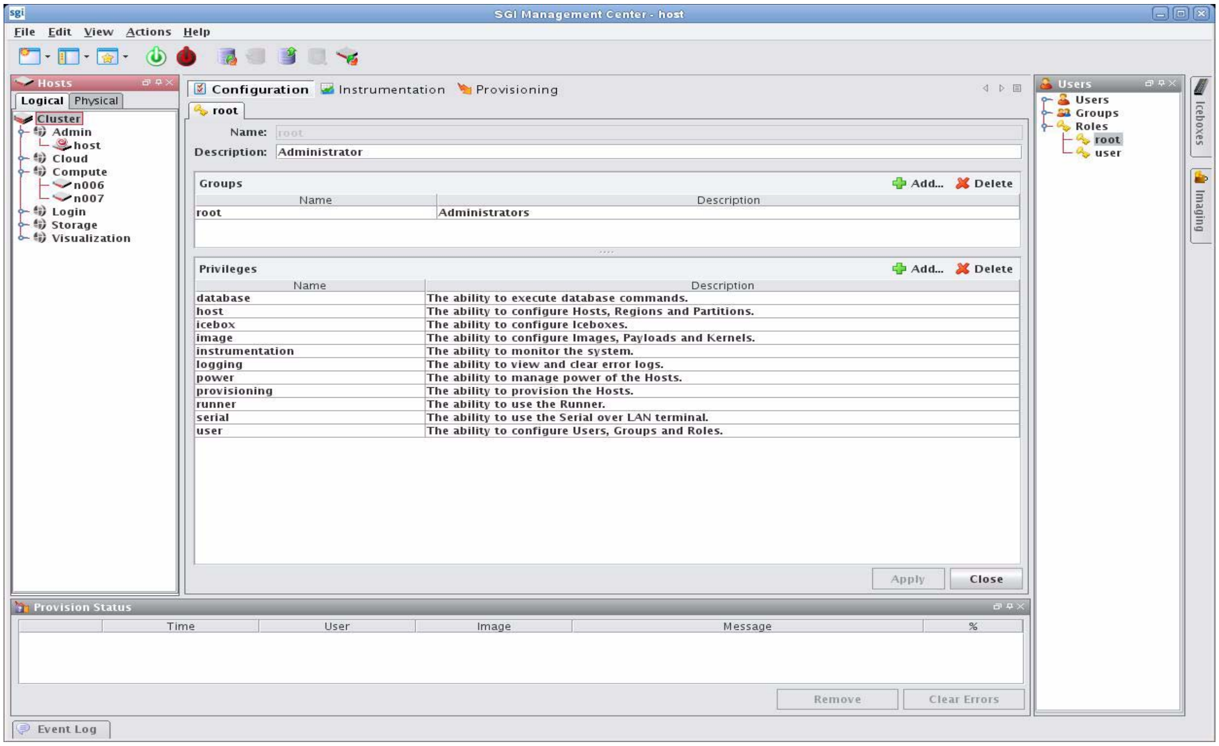
SGI Management Center - host
File Edit View Actions Help
Hosts
Logical Physical
Cluster
Admin
host
Cloud
Compute
n006
n007
Login
Storage
Visualization
Configuration
Instrumentation
Provisioning
root
Add...
Delete
Apply
Close
Provision Status
Time
User
Image
Message
%
Remove
Clear Errors
Event Log
Users
Users
Groups
Roles
root
user
Iceboxes
Imaging

Name: root

Description: Administrator

Groups

Name	Description
root	Administrators

Privileges

Name	Description
database	The ability to execute database commands.
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 of the Hosts.
provisioning	The ability to provision the Hosts.
runner	The ability to use the Runner.
serial	The ability to use the Serial over LAN terminal.
user	The ability to configure Users, Groups and Roles.

3. Make changes as needed and click **Apply**.

Note

Deleting a role will not affect the privileges of a user that is currently logged into Management Center. Changes will take effect only after you restart the Management Center client.

Deleting Roles

Deleting a role removes any user privileges assigned to the role.

1. Select the role you want to delete from the user navigation tree.
2. Select **Delete** from the **Edit** menu or right-click role names in the navigation tree and select **Delete**.
3. Click **OK**.

Note

Deleting a role does not affect the privileges of a user that is currently logged into Management Center. Changes take effect only after you restart the Management Center client. Also note that you cannot delete the root role.

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 will have different roles and access within each group. This flexibility allows you to establish several types of roles you can assign to users: full administration, group administration, user, or guest. See *User Administration* on page 69. 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 6

Imaging, Version Control, and Provisioning

Overview

Management Center version-controlled image management allows you to create and store images that can be used to install and configure hosts in your system. An image may contain file system information, utilities used for provisioning, one payload, and one kernel—although you may create and store many payloads and kernels. The payload contains the operating system, applications, libraries, configuration files, locale and time zone settings, file system structure, selected local user and group accounts (managed by Management Center), and any centralized user authentication settings to install on each host (e.g., NIS, LDAP, and Kerberos). The kernel is the Linux kernel.

Note

For a list of Management Center-supported operating systems, see *Operating System Requirements* on page 2.

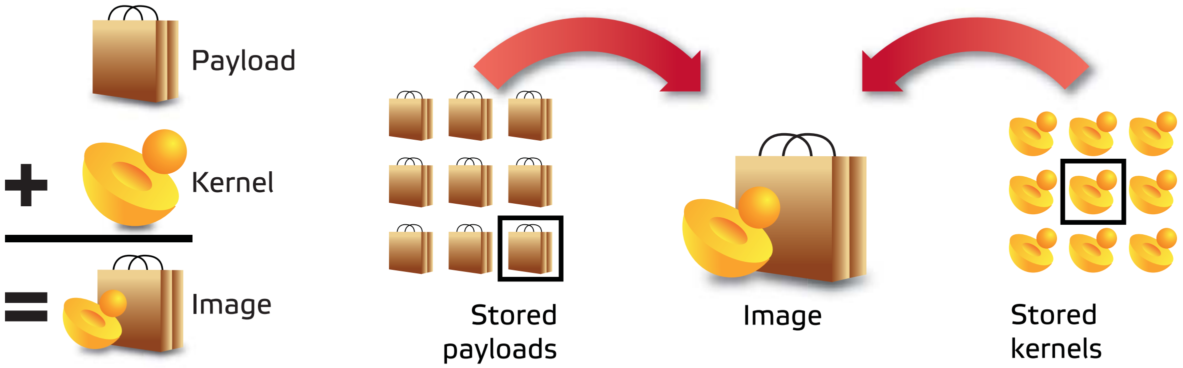

This chapter provides both GUI and command-line interface directions to assist you in configuring and maintaining images, and in using them to provision hosts. The image configuration process allows you to select a kernel and payload, and also configures the boot utilities and partition layout. Once the new image is complete, you can check it into the Version Control System and provision hosts with the new image. See *Version Control System (VCS)* on page 140 and *Provisioning* on page 147.

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 83 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 (e.g., 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.

Creating 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 from a Linux distribution:

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

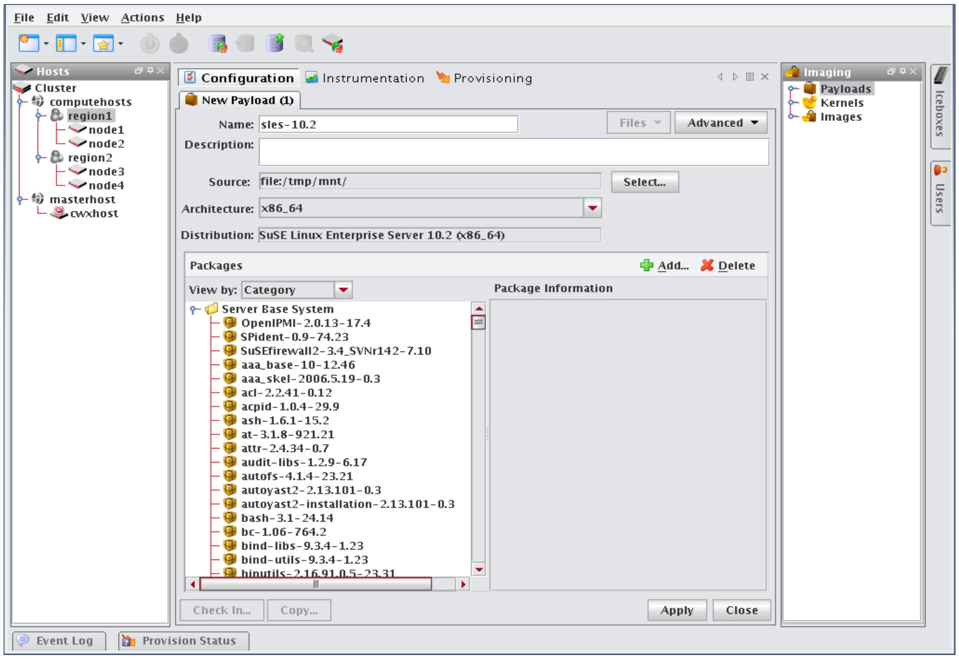
File
Edit
View
Actions
Help
Hosts
Cluster
computehosts
region1
node1
node2
region2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
New Payload (1)
Files
Advanced
Select...
Packages
Add...
Delete
Package Information
Check In...
Copy...
Apply
Close
Imaging
Payloads
Kernels
Images
Iceboxes
Users
Event Log
Provision Status

Name: sles-10.2

Description:

Source: file:/tmp/mnt/

Architecture: x86_64

Distribution: SuSE Linux Enterprise Server 10.2 (x86_64)

View by: Category

Server Base System
- OpenIPMI-2.0.13-17.4
- SPident-0.9-74.23
- SuSEfirewall2-3.4_SVNr142-7.10
- aaa_base-10-12.46
- aaa_skel-2006.5.19-0.3
- acl-2.2.41-0.12
- acpid-1.0.4-29.9
- ash-1.6.1-15.2
- at-3.1.8-921.21
- attr-2.4.34-0.7
- audit-libs-1.2.9-6.17
- autofs-4.1.4-23.21
- autoyast2-2.13.101-0.3
- autoyast2-installation-2.13.101-0.3
- bash-3.1-24.14
- bc-1.06-764.2
- bind-libs-9.3.4-1.23
- bind-utils-9.3.4-1.23
- binutils-2.16.91.0.5-23.31

Tip

To create a new payload using a payload from a host you have already configured, see *Importing a Payload from an Existing Host* on page 87.

2. Enter a payload name.
3. (Optional) Enter a description.

4. Click **Select** next to the **Source** field.

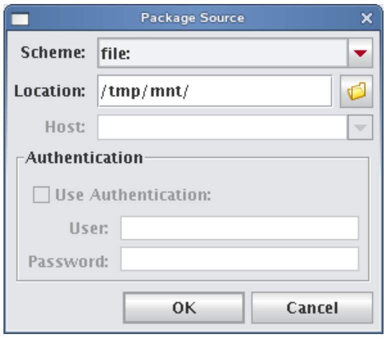

Scheme: file:
Location: /tmp/mnt/

5. Select the **Scheme** (file, http://, or ftp://) from the drop-down list.
6. Enter the location of the top level directory for the Linux distribution or, click the **Browse** icon if you selected the File scheme to locate the directory.

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 82 and *SUSE Linux Enterprise Server Installations* on page 83 for specific details on installing these distributions.

7. (Optional) If you select **http://** or **ftp://**, enter a host.
8. (Optional) If you select **Use Authentication**, enter a username and password.
9. Click **OK**.

   As the distribution loads, the progress of the payload creation is displayed along with the operation status messages.

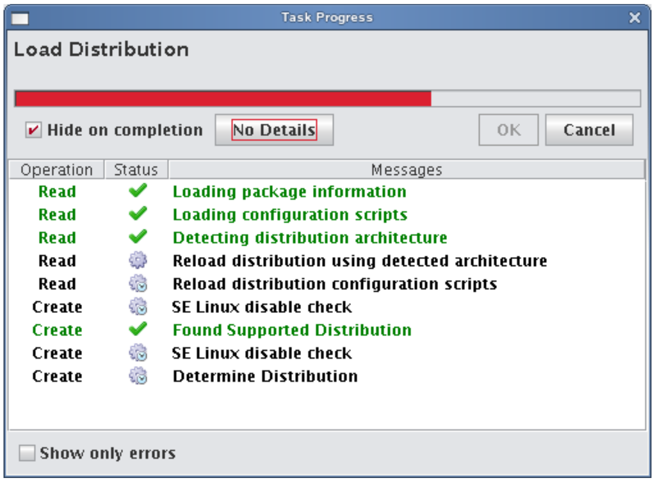

☑ Hide on completion

Operation	Status	Messages
Read	✓	Loading package information
Read	✓	Loading configuration scripts
Read	✓	Detecting distribution architecture
Read	⚙	Reload distribution using detected architecture
Read	⚙	Reload distribution configuration scripts
Create	⚙	SE Linux disable check
Create	✓	Found Supported Distribution
Create	⚙	SE Linux disable check
Create	⚙	Determine Distribution

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.

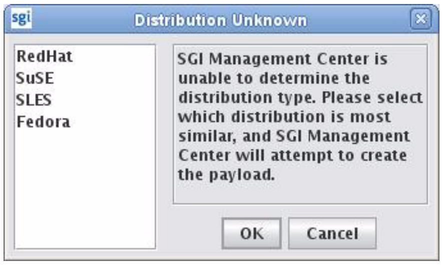

SGI Management Center is unable to determine the distribution type. Please select which distribution is most similar, and SGI Management Center will attempt to create the payload.

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

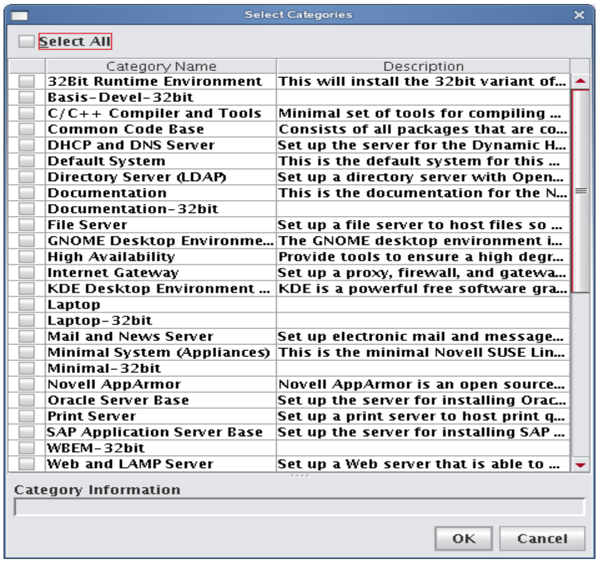

	Category Name	Description
☐	32Bit Runtime Environment	This will install the 32bit variant of...
☐	Basis-Devel-32bit	
☐	C/C++ Compiler and Tools	Minimal set of tools for compiling ...
☐	Common Code Base	Consists of all packages that are co...
☐	DHCP and DNS Server	Set up the server for the Dynamic H...
☐	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 ...

11. 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. See *Adding a Package to an Existing Payload* on page 90.

12. Click **OK**.
13. (Optional) From the Packages pane, select packages you want to remove from the payload, then click **Delete** in the packages pane.
14. (Optional) Configure advanced settings you want to apply to the payload. See *Payload File Configuration* on page 95, *Payload Authentication Management* on page 96, and *Payload Local User and Group Account Management* on page 98.

15. Click **Apply**.

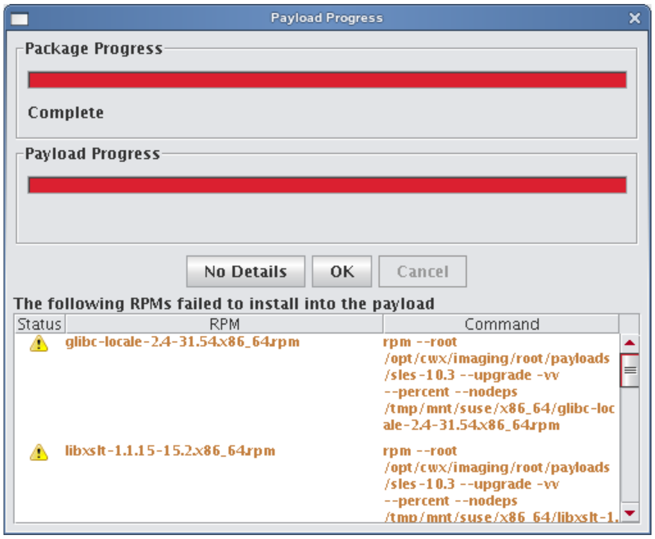

The following RPMs failed to install into the payload

Status	RPM	Command
⚠	glibc-locale-2.4-31.54.x86_64.rpm	rpm --root /opt/cwx/imaging/root/payloads /sles-10.3 --upgrade -vv --percent --nodeps /tmp/mnt/suse/x86_64/glibc-locale-2.4-31.54.x86_64.rpm
⚠	libxslt-1.1.15-15.2.x86_64.rpm	rpm --root /opt/cwx/imaging/root/payloads /sles-10.3 --upgrade -vv --percent --nodeps /tmp/mnt/suse/x86_64/libxslt-1.

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.

16. (Optional) Select any payload files you wish to include with, remove from, or edit from the **File** drop-down list. See *Add and Update Payload Files or Directories* on page 102.

17. (Optional) Click **Check In** to import the new payload into VCS. See also *Version Control System (VCS)* on page 140.

Creating a Copy of an Existing Payload

1. Right-click on a payload in the imaging navigation tree and select **Copy**.

Tip

If a payload is open in the GUI, click **Copy** in the lower left of the panel to create a copy of the payload.

Note

When you copy of a payload, Management Center creates a *working* copy of the payload — in other words, the payload that is checked out into the *$MGR_HOME/imaging/<username>/payloads directory.* To create a copy of a versioned payload, use *VCS Management* on page 143.

2. In the **Copy Payload** dialog, enter the name of the new payload and click **OK**.

Importing a Payload from an Existing Host

Creating a payload from an existing host is helpful in situations where a specific host is already configured the way you want it. This feature allows you to create new payloads that use the configuration and distribute the image to other hosts.

Note

On RHEL, temporarily disable SE Linux while importing the payload. If you do not require SE Linux, you may want to leave it disabled.

To disable SE Linux:

1. Navigate to the **Imaging** tab.
2. Select the kernel you are using and edit the kernel parameters.
3. Add `selinux=0` as a parameter.
4. Reboot the host and import the payload.

3. Select **Import Payload** from the **File** menu.

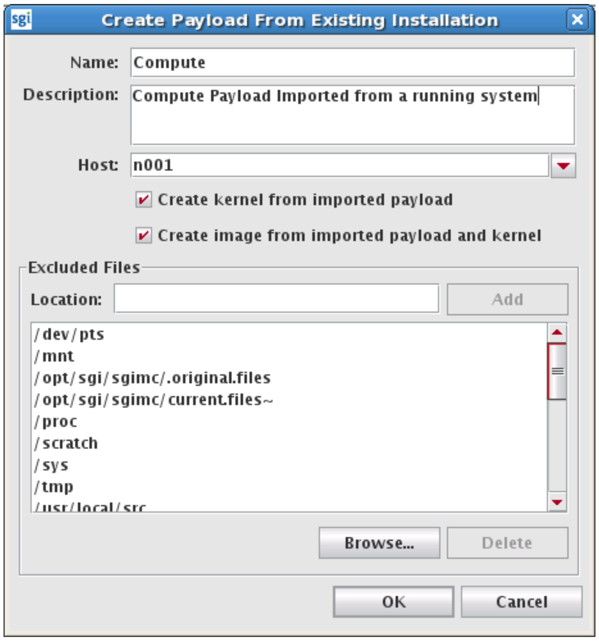

Name: Compute
Description: Compute Payload Imported from a running system
Host: n001
☑ Create kernel from imported payload
☑ Create image from imported payload and kernel

/dev/pts
/mnt
/opt/sgi/sgimc/.original.files
/opt/sgi/sgimc/current.files~
/proc
/scratch
/sys
/tmp
/usr/local/src

Tip

You can also import a payload using `pmgr` from the command line. See *pmgr* on page 258.

4. Enter a payload name.

5. (Optional) Enter a description.

6. Enter the host name you are creating the payload from or select a host from the drop-down list.

7. Use the following check box selections to indicate whether or not an image and kernel should be created:

* **Create kernel from imported payload** creates a kernel with the same name as the payload and populates the list of modules in the kernel to match that of the running host**.**

* **Create image from imported payload and kernel** creates an image and attempts to re-create all local filesystems from the list of partitions on the running host.

Note

- In the imported kernel, check the list of kernel modules that is generated and the list of kernel boot parameters that are generated. You may need to customize them according to your needs.

- In the imported image, check the partition scheme and add/remove partitions as necessary.

- LVM is not supported. If you have LVM partitions on the running host, you will need to create traditional partitions in the image manually.

- If you have any remote filesystem mounts on the running host, such as static NFS mounts, they will not be defined in the new image. They will need to be defined manually.

8. (Optional) Review the **Excluded Files** list and remove any files you want to exclude from the payload.

! Warning!

If you include a symlink when creating a payload, excluding the target produces a dangling symbolic link. This link may cause an exception and abort payload creation when Management Center attempts to repair missing directories.

9. (Optional) Enter the location of any file you want to exclude from the payload and click **Add**. Click Browse to locate a file on your system.
10. Click **OK**.

Importing Kernel Parameters from a Running Host

Management Center allows you to import the list of kernel parameters from a running host. This is especially useful when you have imported a payload (and kernel/image) from a running host and you want the kernel parameters to match those of the running host.

To import the kernel parameters, do the following:

1. Open the imaging pane, find the desired kernel, and double-click it.
2. On the resulting kernel configuration panel (shown in the following figure), click the **Import...** button.

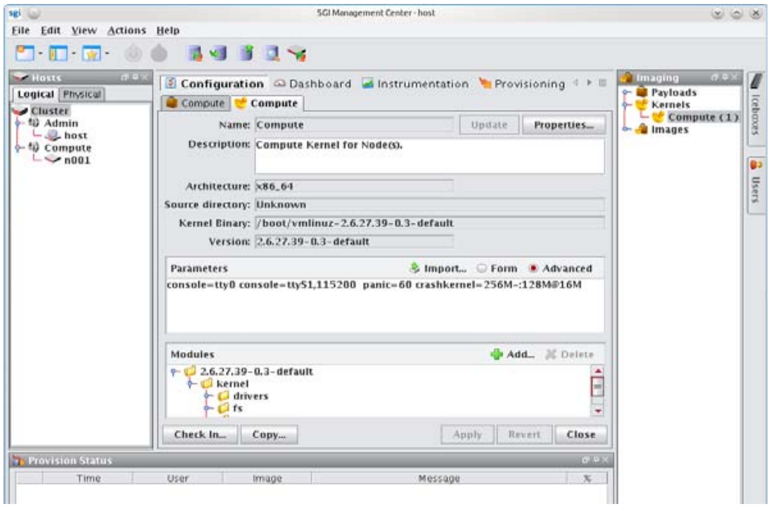

SGI Management Center - host
File
Edit
View
Actions
Help
Hosts
Logical
Physical
Cluster
Admin
host
Compute
n001
Configuration
Dashboard
Instrumentation
Provisioning
Compute
Compute
Update
Properties...
Import...
○ Form
● Advanced
Add...
Delete
Check In...
Copy...
Apply
Revert
Close
Imaging
Payloads
Kernels
Compute (1)
Images
Iceboxes
Users
Provision Status
Time
User
Image
Message

Name: Compute
Description: Compute Kernel for Node(s).
Architecture: x86_64
Source directory: Unknown
Kernel Binary: /boot/vmlinuz-2.6.27.39-0.3-default
Version: 2.6.27.39-0.3-default

Parameters

```
console=tty0 console=ttyS1,115200  panic=60 crashkernel=256M-:128M@16M
```

Modules

```
2.6.27.39-0.3-default
  kernel
    drivers
    fs
```

3. In the **Import Kernel Parameters** window that appears, select the desired host and click **Import**.

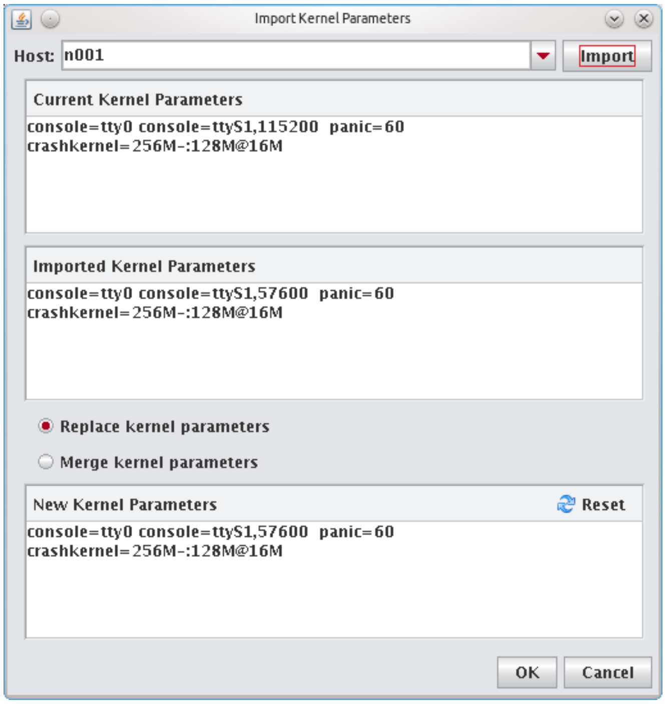

Host: n001

Current Kernel Parameters

```
console=tty0 console=ttyS1,115200  panic=60
crashkernel=256M-:128M@16M
```

Imported Kernel Parameters

```
console=tty0 console=ttyS1,57600  panic=60
crashkernel=256M-:128M@16M
```

● Replace kernel parameters
○ Merge kernel parameters

New Kernel Parameters

```
console=tty0 console=ttyS1,57600  panic=60
crashkernel=256M-:128M@16M
```

4. Select either **Replace kernel parameters** or **Merge kernel parameters** and click the **OK** button.
5. Examine the list of kernel parameters that was generated and make any needed changes.
6. Update the kernel and image and check in any changes.

Adding a Package to an Existing Payload

Adding a package to a payload allows you to make additions or changes to the default Linux installation. For a list of supported distributions, see *Operating System Requirements* on page 2. If you add packages to the payload that contain new or updated kernel modules and complications occur (or if the modules are needed to boot the system), then you should create a new kernel. See *To Create a Kernel from a Payload* on page 109.

To add a package, do the following:

1. Right-click a payload name in the imaging navigation tree and select **Edit**.

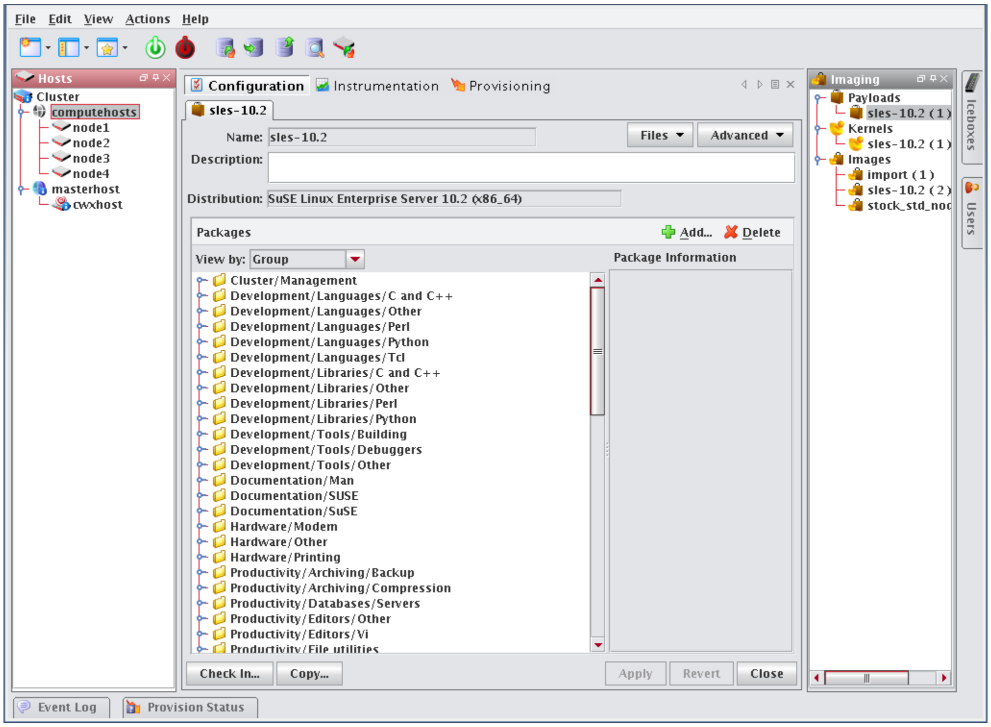
File Edit View Actions Help
Hosts
Cluster
computehosts
node1
node2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
sles-10.2
Files
Advanced
Packages
Add...
Delete
Package Information
Cluster/Management
Development/Languages/C and C++
Development/Languages/Other
Development/Languages/Perl
Development/Languages/Python
Development/Languages/Tcl
Development/Libraries/C and C++
Development/Libraries/Other
Development/Libraries/Perl
Development/Libraries/Python
Development/Tools/Building
Development/Tools/Debuggers
Development/Tools/Other
Documentation/Man
Documentation/SUSE
Documentation/SuSE
Hardware/Modem
Hardware/Other
Hardware/Printing
Productivity/Archiving/Backup
Productivity/Archiving/Compression
Productivity/Databases/Servers
Productivity/Editors/Other
Productivity/Editors/Vi
Productivity/File utilities
Check In...
Copy...
Apply
Revert
Close
Imaging
Payloads
sles-10.2 (1)
Kernels
sles-10.2 (1)
Images
import (1)
sles-10.2 (2)
stock_std_noc
Iceboxes
Users
Event Log
Provision Status

Name: sles-10.2
Description:
Distribution: SuSE Linux Enterprise Server 10.2 (x86_64)
View by: Group

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

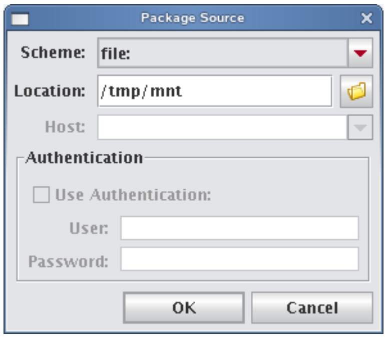
Package Source
Host:
Authentication
☐ Use Authentication:
User:
Password:
OK
Cancel

Scheme: file:
Location: /tmp/mnt

3. Select a scheme (file, http://, or ftp://).
4. Enter the Location of the top level directory for the Linux distribution, a directory containing RPM packages, or the location of an individual package. If you selected the **File** scheme, click the **Browse** icon to locate the package.

Note

If the browse button does not launch a dialog, a DNS name resolution error may exist. The DNS server name must be specified in the client — not the IP address.

If you have several packages in a directory, select the directory. Management Center displays all packages in the directory — you can choose which packages you want to install. Management Center resolves package dependencies (see *Payload Package Dependency Checks* on page 93).

5. (Optional) If you selected **http://** or **ftp://**, enter a host.
6. (Optional) If you selected **Use Authentication**, enter a username and password.
7. Click **OK**.

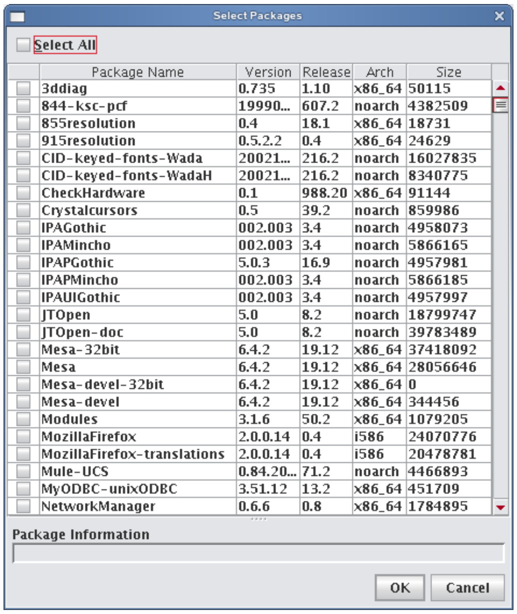

	Package Name	Version	Release	Arch	Size
☐	3ddiag	0.735	1.10	x86_64	50115
☐	844-ksc-pcf	19990...	607.2	noarch	4382509
☐	855resolution	0.4	18.1	x86_64	18731
☐	915resolution	0.5.2.2	0.4	x86_64	24629
☐	CID-keyed-fonts-Wada	20021...	216.2	noarch	16027835
☐	CID-keyed-fonts-WadaH	20021...	216.2	noarch	8340775
☐	CheckHardware	0.1	988.20	x86_64	91144
☐	Crystalcursors	0.5	39.2	noarch	859986
☐	IPAGothic	002.003	3.4	noarch	4958073
☐	IPAMincho	002.003	3.4	noarch	5866165
☐	IPAPGothic	5.0.3	16.9	noarch	4957981
☐	IPAPMincho	002.003	3.4	noarch	5866185
☐	IPAUIGothic	002.003	3.4	noarch	4957997
☐	JTOpen	5.0	8.2	noarch	18799747
☐	JTOpen-doc	5.0	8.2	noarch	39783489
☐	Mesa-32bit	6.4.2	19.12	x86_64	37418092
☐	Mesa	6.4.2	19.12	x86_64	28056646
☐	Mesa-devel-32bit	6.4.2	19.12	x86_64	0
☐	Mesa-devel	6.4.2	19.12	x86_64	344456
☐	Modules	3.1.6	50.2	x86_64	1079205
☐	MozillaFirefox	2.0.0.14	0.4	i586	24070776
☐	MozillaFirefox-translations	2.0.0.14	0.4	i586	20478781
☐	Mule-UCS	0.84.20...	71.2	noarch	4466893
☐	MyODBC-unixODBC	3.51.12	13.2	x86_64	451709
☐	NetworkManager	0.6.6	0.8	x86_64	1784895

8. Select the packages you want to install.
9. Click **OK**.
10. Click **Apply** to save changes.

Note

Before adding the package, Management Center performs a package dependency check. See *Payload Package Dependency Checks* on page 93 for information about dependency errors.

11. Click **Check In** to check the payload into VCS.
12. Update the image to use the new payload.
13. Re-provision the hosts with the new image or update the payload on the hosts using *VCS Upgrade* on page 150.

Remove a Payload Package

The Packages pane of the payload panel provides a view into the current packages installed in the payload. See also *Payload Package Dependency Checks* on page 93.

To Remove a Payload Package

1. Right-click on a payload in the imaging navigation tree and select **Edit**. The payload panel appears.
2. From the package list in the Packages pane, select a package group or expand the group to view individual packages.

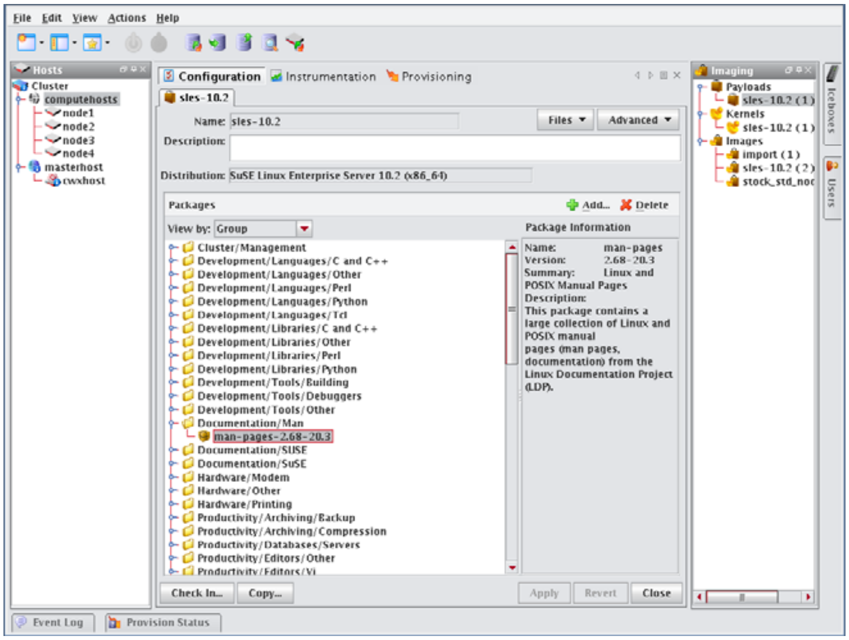
File
Edit
View
Actions
Help
Hosts
Cluster
computehosts
node1
node2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
sles-10.2
Files
Advanced
Packages
Add...
Delete
Package Information
Cluster/Management
Development/Languages/C and C++
Development/Languages/Other
Development/Languages/Perl
Development/Languages/Python
Development/Languages/Tcl
Development/Libraries/C and C++
Development/Libraries/Other
Development/Libraries/Perl
Development/Libraries/Python
Development/Tools/Building
Development/Tools/Debuggers
Development/Tools/Other
Documentation/Man
man-pages-2.68-20.3
Documentation/SUSE
Documentation/SuSE
Hardware/Modem
Hardware/Other
Hardware/Printing
Productivity/Archiving/Backup
Productivity/Archiving/Compression
Productivity/Databases/Servers
Productivity/Editors/Other
Productivity/Editors/Vi
Check In...
Copy...
Apply
Revert
Close
Imaging
Payloads
sles-10.2 (1)
Kernels
sles-10.2 (1)
Images
import (1)
sles-10.2 (2)
stock_std_noc
Iceboxes
Users
Event Log
Provision Status

Name: sles-10.2

Description:

Distribution: SuSE Linux Enterprise Server 10.2 (x86_64)

View by: Group

Name: man-pages
Version: 2.68-20.3
Summary: Linux and POSIX Manual Pages
Description:
This package contains a large collection of Linux and POSIX manual pages (man pages, documentation) from the Linux Documentation Project (LDP).

Tip

To view individual packages instead of package groups, change the View Packages By option.

3. Click **Delete**.
4. Click **OK** to remove the packages.
5. Click **Apply** to save changes.

Note

Before adding the package, Management Center performs a package dependency check. See *Payload Package Dependency Checks* on page 93 for information about dependency errors.

Payload Package Dependency Checks

Before performing package addition, update, or removal, Management Center performs a package dependency check. Any failures identified through the dependency check are displayed in the Resolve Dependency Failures dialog. From this dialog, you can choose a course of action to address the failure(s).

ADDING A PACKAGE

When adding a package, you may correct dependency failures by selecting one of the following options:

- Add packages needed to resolve dependency failures.
- Ignore packages that have dependency failures.
- Force package installation, ignoring dependency failures.

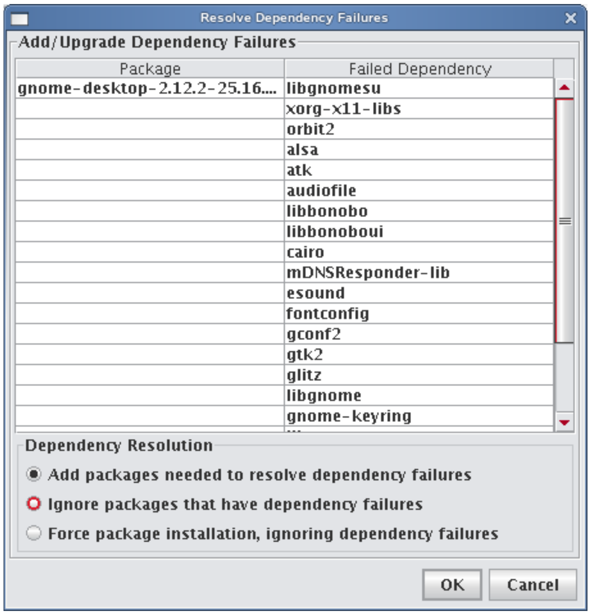

Add/Upgrade Dependency Failures

Package	Failed Dependency
gnome-desktop-2.12.2-25.16....	libgnomesu
	xorg-x11-libs
	orbit2
	alsa
	atk
	audiofile
	libbonobo
	libbonoboui
	cairo
	mDNSResponder-lib
	esound
	fontconfig
	gconf2
	gtk2
	glitz
	libgnome
	gnome-keyring

Dependency Resolution

● Add packages needed to resolve dependency failures
○ Ignore packages that have dependency failures
○ Force package installation, ignoring dependency failures

REMOVING A PACKAGE

When removing a package, you may correct dependency failures by selecting one of the following options:

- Ignore packages that have dependency failures.
- Force package deletion, ignoring dependency failures.

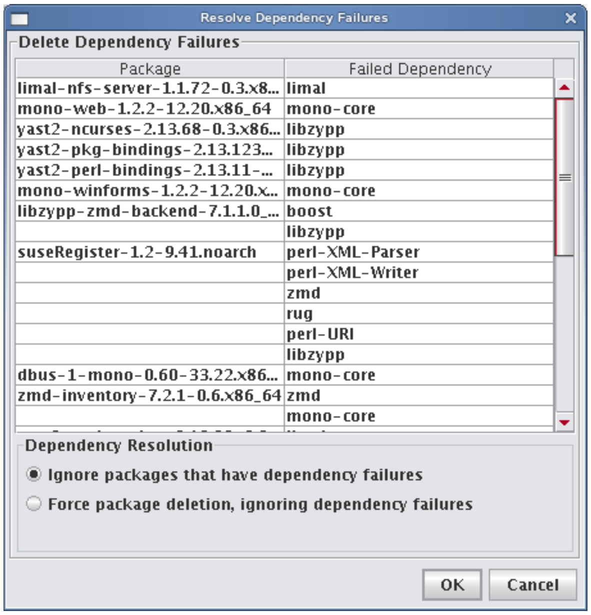

Delete Dependency Failures

Package	Failed Dependency
limal-nfs-server-1.1.72-0.3.x8...	limal
mono-web-1.2.2-12.20.x86_64	mono-core
yast2-ncurses-2.13.68-0.3.x86...	libzypp
yast2-pkg-bindings-2.13.123...	libzypp
yast2-perl-bindings-2.13.11-...	libzypp
mono-winforms-1.2.2-12.20.x...	mono-core
libzypp-zmd-backend-7.1.1.0_...	boost
	libzypp
suseRegister-1.2-9.41.noarch	perl-XML-Parser
	perl-XML-Writer
	zmd
	rug
	perl-URI
	libzypp
dbus-1-mono-0.60-33.22.x86...	mono-core
zmd-inventory-7.2.1-0.6.x86_64	zmd
	mono-core

Dependency Resolution

● Ignore packages that have dependency failures
○ Force package deletion, ignoring dependency failures

Payload File Configuration

Payload file configuration allows you to set up configuration options when creating or editing a payload including: DHCP Network, Network, Serial Console, Virtual Console, and more. When you click Apply, the scripts that correspond to the selected item(s) run on the payload. It is important to note that the selected script(s) run at the time you click Apply—this list is *not* an indication of scripts that have run at some point on the system.

Note

The list of options available is based on the distribution selected. The options displayed in the example below are SUSE-based distributions (SUSE Linux Enterprise Server 10).

To Configure a Payload

1. Right-click on a payload in the imaging navigation tree and select **Edit**.

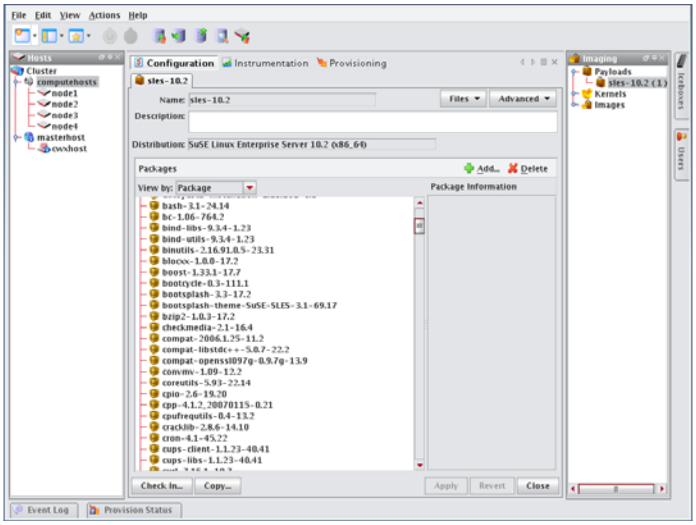
File
Edit
View
Actions
Help
Hosts
Cluster
computehosts
node1
node2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
sles-10.2
Files
Advanced
Packages
Add...
Delete
Package Information
Check In...
Copy...
Apply
Revert
Close
Imaging
Payloads
sles-10.2 (1)
Kernels
Images
Iceboxes
Users
Event Log
Provision Status

Name: sles-10.2

Description:

Distribution: SuSE Linux Enterprise Server 10.2 (x86_64)

View by: Package

```
bash-3.1-24.14
bc-1.06-764.2
bind-libs-9.3.4-1.23
bind-utils-9.3.4-1.23
binutils-2.16.91.0.5-23.31
blocxx-1.0.0-17.2
boost-1.33.1-17.7
bootcycle-0.3-111.1
bootsplash-3.3-17.2
bootsplash-theme-SuSE-SLES-3.1-69.17
bzip2-1.0.3-17.2
checkmedia-2.1-16.4
compat-2006.1.25-11.2
compat-libstdc++-5.0.7-22.2
compat-openssl097g-0.9.7g-13.9
convmv-1.09-12.2
coreutils-5.93-22.14
cpio-2.6-19.20
cpp-4.1.2_20070115-0.21
cpufrequtils-0.4-13.2
cracklib-2.8.6-14.10
cron-4.1-45.22
cups-client-1.1.23-40.41
cups-libs-1.1.23-40.41
```

2. Select **Configuration** from the **Advanced** drop-down list and click the check box by each script you want to enable.

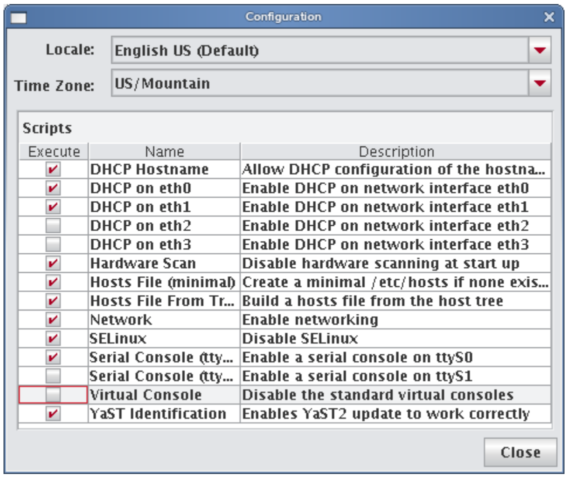
Configuration
Close

Locale: English US (Default)

Time Zone: US/Mountain

Scripts

Execute	Name	Description
☑	DHCP Hostname	Allow DHCP configuration of the hostna...
☑	DHCP on eth0	Enable DHCP on network interface eth0
☑	DHCP on eth1	Enable DHCP on network interface eth1
☐	DHCP on eth2	Enable DHCP on network interface eth2
☐	DHCP on eth3	Enable DHCP on network interface eth3
☑	Hardware Scan	Disable hardware scanning at start up
☑	Hosts File (minimal)	Create a minimal /etc/hosts if none exis...
☑	Hosts File From Tr...	Build a hosts file from the host tree
☑	Network	Enable networking
☑	SELinux	Disable SELinux
☑	Serial Console (tty...	Enable a serial console on ttyS0
☐	Serial Console (tty...	Enable a serial console on ttyS1
☐	Virtual Console	Disable the standard virtual consoles
☑	YaST Identification	Enables YaST2 update to work correctly

3. Click **Apply**.

Payload Authentication Management

Payload Authentication manages the authentication settings for the payload. This option allows you to enable, disable, or modify the settings for supported remote authentication schemes. Management Center supports the following remote authentication schemes:

- Network Information Service (NIS)
- Lightweight Directory Access Protocol (LDAP)
- Kerberos (a network authentication protocol)

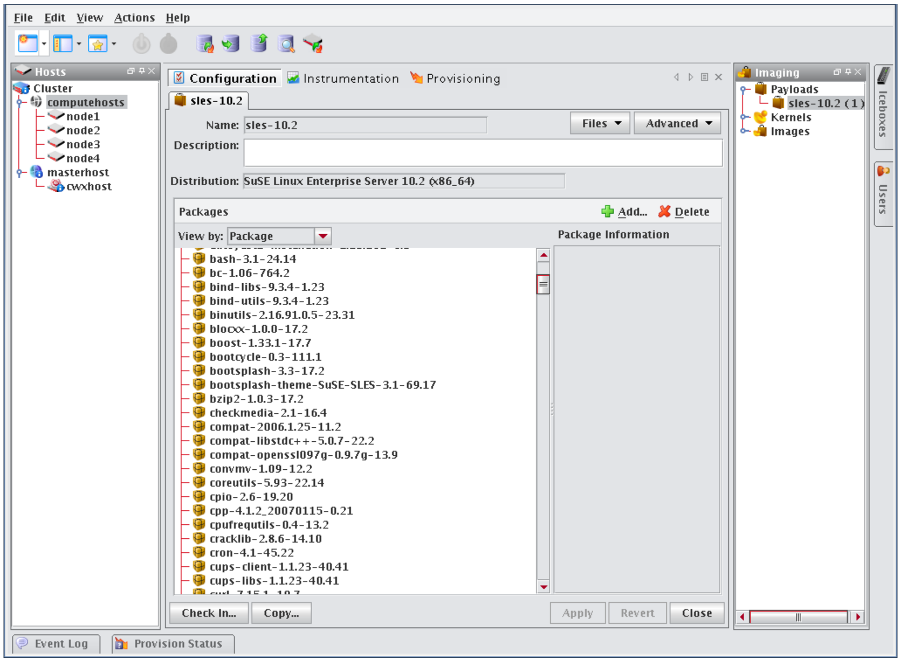
File Edit View Actions Help
Hosts
Cluster
computehosts
node1
node2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
sles-10.2
Files
Advanced
Packages
Add...
Delete
View by:
Package Information
bash-3.1-24.14
bc-1.06-764.2
bind-libs-9.3.4-1.23
bind-utils-9.3.4-1.23
binutils-2.16.91.0.5-23.31
blocxx-1.0.0-17.2
boost-1.33.1-17.7
bootcycle-0.3-111.1
bootsplash-3.3-17.2
bootsplash-theme-SuSE-SLES-3.1-69.17
bzip2-1.0.3-17.2
checkmedia-2.1-16.4
compat-2006.1.25-11.2
compat-libstdc++-5.0.7-22.2
compat-openssl097g-0.9.7g-13.9
convmv-1.09-12.2
coreutils-5.93-22.14
cpio-2.6-19.20
cpp-4.1.2_20070115-0.21
cpufrequtils-0.4-13.2
cracklib-2.8.6-14.10
cron-4.1-45.22
cups-client-1.1.23-40.41
cups-libs-1.1.23-40.41
Check In...
Copy...
Apply
Revert
Close
Imaging
Payloads
sles-10.2 (1)
Kernels
Images
Iceboxes
Users
Event Log
Provision Status

Name: sles-10.2

Description:

Distribution: SuSE Linux Enterprise Server 10.2 (x86_64)

View by: Package

To Configure NIS Authentication

1. Right-click on a payload in the imaging navigation tree and select Edit. The payload panel appears.
2. Select Authentication from the Advanced pull-down menu. The Authentication dialog appears.
3. Select the NIS tab.
    - **A.** Click the Use NIS option.

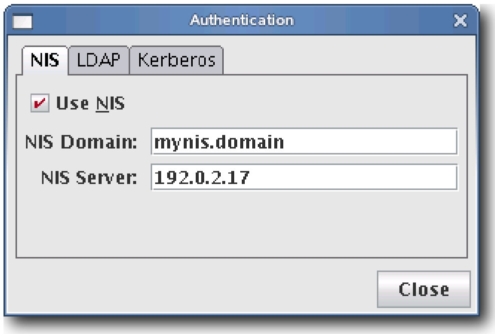
Authentication
NIS
LDAP
Kerberos
Close

☑ Use NIS

NIS Domain: mynis.domain

NIS Server: 192.0.2.17

- **B.** Enter the NIS domain.
- **C.** (Optional) Enter the NIS Server.

4. Click Close.
5. Click Apply to save changes. Click Revert or Close to abort this action.

To Configure LDAP Authentication

1. Right-click on a payload in the imaging navigation tree and select Edit. The payload panel appears.
2. Select Authentication from the Advanced pull-down menu. The Authentication dialog appears.
3. Select the LDAP tab.
   A. Click the Use LDAP option.

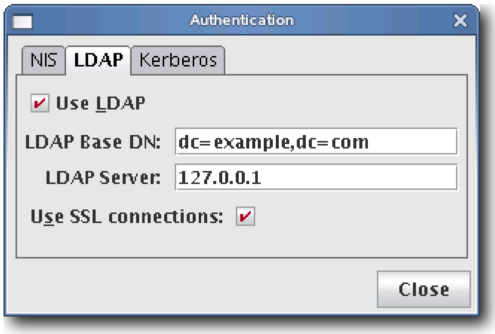

☑ Use LDAP
LDAP Base DN: dc=example,dc=com
LDAP Server: 127.0.0.1
Use SSL connections: ☑

   B. Enter the LDAP Base DN (Distinguished Name).
   C. Enter the LDAP Server.
   D. (Optional) Click Use SSL connections if you want to connect to the LDAP server via SSL.
4. Click Close.
5. Click Apply to save changes. Click Revert or Close to abort this action.

To Configure Kerberos Authentication

1. Right-click on a payload in the imaging navigation tree and select Edit. The payload panel appears.
2. Select Authentication from the Advanced pull-down menu. The Authentication dialog appears.
3. Select the Kerberos tab.
   A. Click the Use Kerberos option.

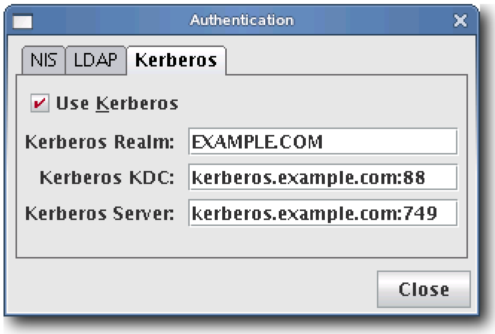

☑ Use Kerberos
Kerberos Realm: EXAMPLE.COM
Kerberos KDC: kerberos.example.com:88
Kerberos Server: kerberos.example.com:749

   B. Enter the Kerberos Realm.
   C. Enter the Kerberos KDC (Key Distribution Center).
   D. Enter the Kerberos Server.
4. Click Close.
5. Click Apply to save changes. Click Revert or Close to abort this action.

Payload Local User and Group Account Management

The Local Accounts payload management option provides a means for managing local accounts in payloads. This option allows you to:

- Add a local user or group account known to Management Center to the payload (see *User Administration* on page 69).
- Delete a local user or group account from the payload.

Note

Local account management does not support moving local accounts from the host.

Local user and group accounts that are reserved for system use do not display and cannot be added or deleted. The root account is added automatically. Management Center handles group dependencies.

Tip

Software that requires you to add groups (e.g., Myrinet Group) can be managed through user accounts.

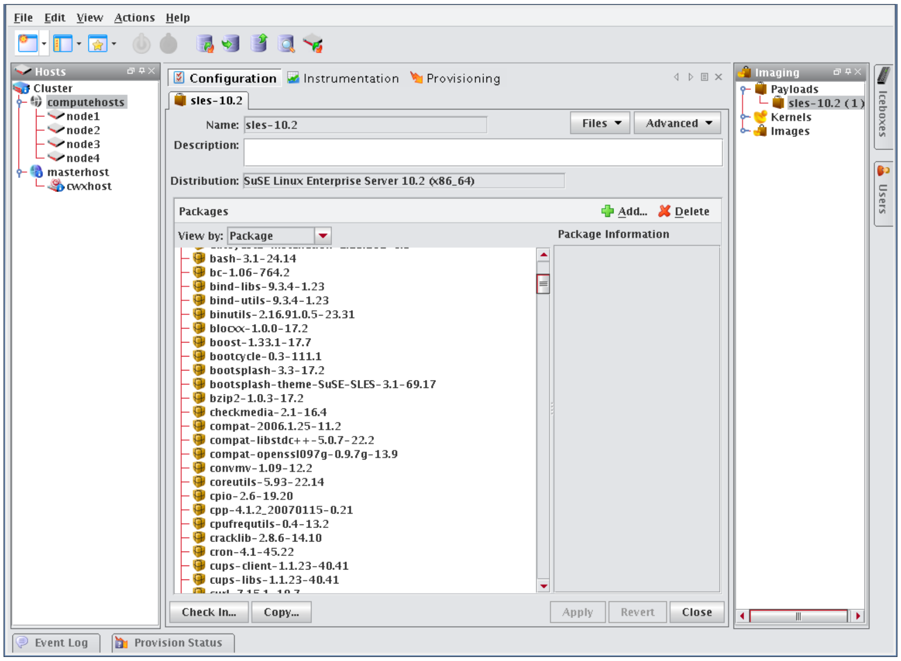
File
Edit
View
Actions
Help
Hosts
Cluster
computehosts
node1
node2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
sles-10.2
Files
Advanced
Packages
Add...
Delete
View by: Package
Package Information
Check In...
Copy...
Apply
Revert
Close
Imaging
Payloads
sles-10.2 (1)
Kernels
Images
Iceboxes
Users
Event Log
Provision Status

Name: sles-10.2
Description:
Distribution: SuSE Linux Enterprise Server 10.2 (x86_64)

bash-3.1-24.14
bc-1.06-764.2
bind-libs-9.3.4-1.23
bind-utils-9.3.4-1.23
binutils-2.16.91.0.5-23.31
blocxx-1.0.0-17.2
boost-1.33.1-17.7
bootcycle-0.3-111.1
bootsplash-3.3-17.2
bootsplash-theme-SuSE-SLES-3.1-69.17
bzip2-1.0.3-17.2
checkmedia-2.1-16.4
compat-2006.1.25-11.2
compat-libstdc++-5.0.7-22.2
compat-openssl097g-0.9.7g-13.9
convmv-1.09-12.2
coreutils-5.93-22.14
cpio-2.6-19.20
cpp-4.1.2_20070115-0.21
cpufrequtils-0.4-13.2
cracklib-2.8.6-14.10
cron-4.1-45.22
cups-client-1.1.23-40.41
cups-libs-1.1.23-40.41

Local User Accounts

TO ADD A LOCAL USER ACCOUNT TO A PAYLOAD

1. Right-click on a payload in the imaging navigation tree and select Edit. The payload panel appears.
2. Select Local Accounts from the Advanced pull-down menu. The Local Accounts dialog appears.

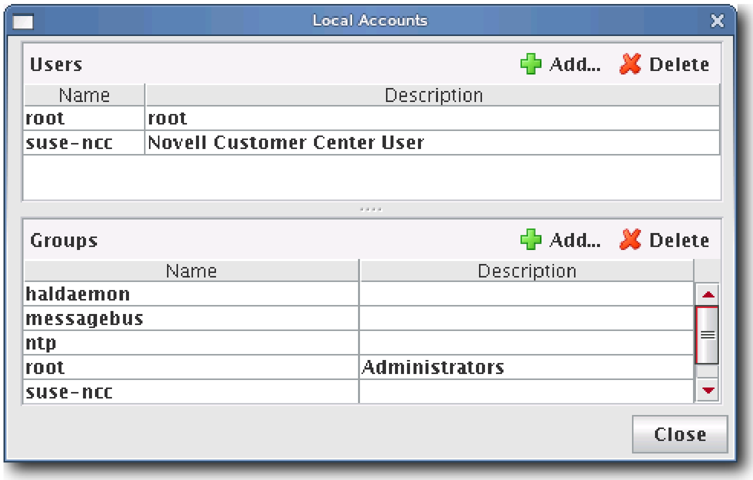

Name	Description
root	root
suse-ncc	Novell Customer Center User

Name	Description
haldaemon	
messagebus	
ntp	
root	Administrators
suse-ncc	

3. In the Users pane, click Add. The Add User dialog appears.

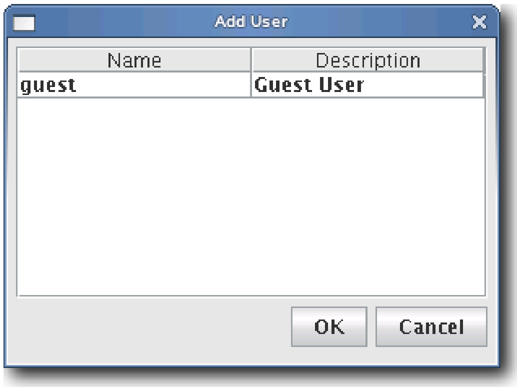

Name	Description
guest	Guest User

4. Select the user(s) to add to the payload (use the Shift or Ctrl keys to select multiple users).
5. Click OK to add the user(s) or click Cancel to abort this action.
6. Click Apply to save changes. Click Revert or Close to abort this action.

DELETE A LOCAL USER ACCOUNT FROM A PAYLOAD

1. Right-click on a payload in the imaging navigation tree and select Edit. The payload panel appears.
2. Select Local Accounts from the Advanced pull-down menu. The Local Accounts dialog appears.

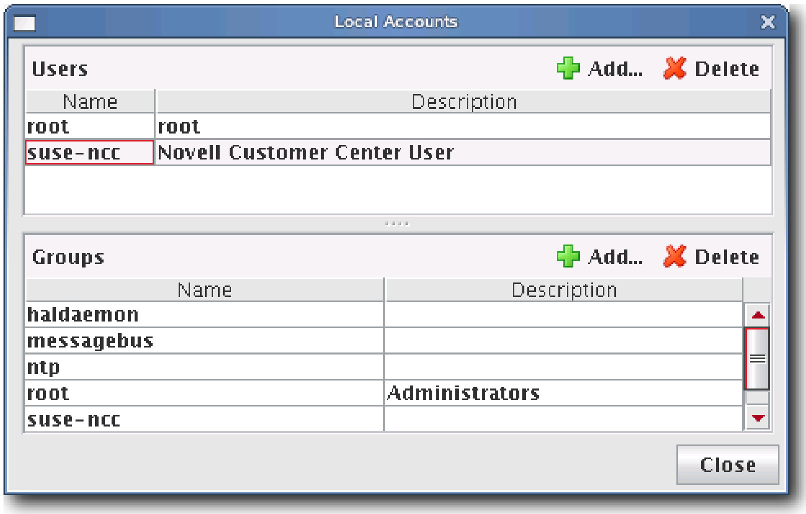

Name	Description
root	root
suse-ncc	Novell Customer Center User

Name	Description
haldaemon	
messagebus	
ntp	
root	Administrators
suse-ncc	

3. Select the user(s) to remove from the payload (use the Shift or Ctrl keys to select multiple users).
4. Click Delete to remove the user(s).
5. Click Close.
6. Click Apply to complete the process. Click Revert or Close to abort this action.

Group User Accounts

ADD A GROUP USER ACCOUNT TO A PAYLOAD

1. Right-click on a payload in the imaging navigation tree and select Edit. The payload panel appears.
2. Select Local Accounts from the Advanced pull-down menu. The Local Accounts dialog appears.

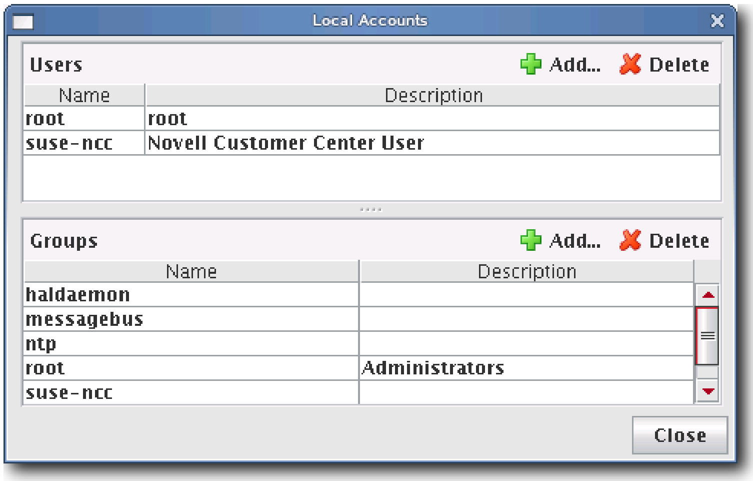

Name	Description
root	root
suse-ncc	Novell Customer Center User

Name	Description
haldaemon	
messagebus	
ntp	
root	Administrators
suse-ncc	

3. In the Groups pane, click Add. The Add Group dialog appears.

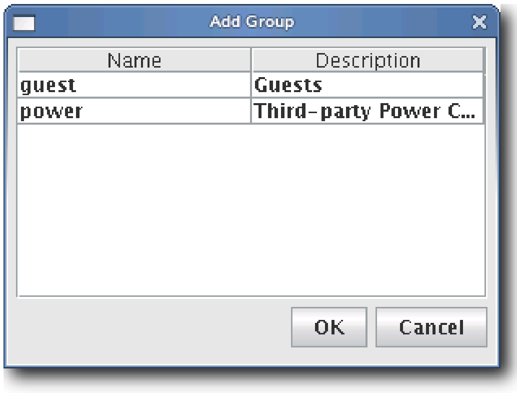

Name	Description
guest	Guests
power	Third-party Power C...

4. Select the group(s) to add to the payload (use the Shift or Ctrl keys to select multiple users).
5. Click OK to add the group(s) or click Cancel to abort this action.
6. Click Apply to complete the process. Click Revert or Close to abort this action.

Delete a Group User Account from a Payload

1. Right-click on a payload in the imaging navigation tree and select Edit. The payload panel appears.
2. Select Local Accounts from the Advanced pull-down menu. The Local Accounts dialog appears.

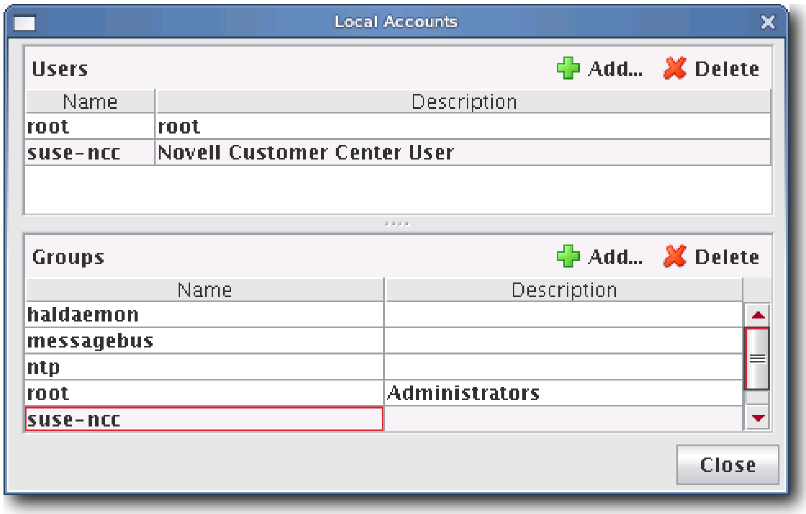

Name	Description
root	root
suse-ncc	Novell Customer Center User

Name	Description
haldaemon	
messagebus	
ntp	
root	Administrators
suse-ncc	

3. Select the group(s) to remove from the payload (use the Shift or Ctrl keys to select multiple groups).
4. Click Delete to remove the group(s).
5. Click Close.
6. Click Apply to complete the process. Click Revert or Close to abort this action.

Add and Update Payload Files or Directories

Adding and updating payload files allows you to select a file or directory from the Master Host's file system and copy it into the payload.

To Add or Update a Payload File or Directory

1. Right-click on a payload in the imaging navigation tree and select Edit. The payload panel appears.

2. Select Add File from the Files pull-down menu. The Add File or Directory dialog appears.

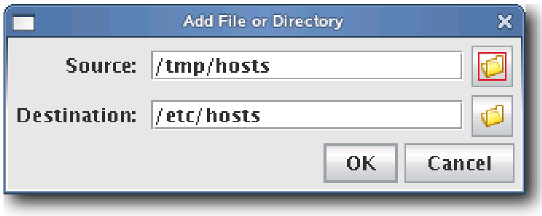

Source: /tmp/hosts
Destination: /etc/hosts

3. Enter the source for the new file in the Source field or click Browse to locate the source.

4. Enter the destination for the new file in the Destination field or click Browse to select the destination.

Note

The destination specified is relative to the payload root.

5. Click OK to save changes or click Cancel to abort this action.

6. Click Apply to complete the process. Click Revert or Close to abort this action.

Tip

If a working copy of a payload is available, you can enter the payload directory and make changes to the payload manually from the CLI. Working copies of payloads are stored at:

$MGR_HOME/imaging/<username>/payloads/<payload_name>

From this directory, enter chroot to change the directory to your root (/) directory. After making changes, check the payload into VCS.

Edit a Payload File with the Text Editor

Management Center allows you to edit payload files with a text editor. Files edited in this manner are treated as plain text and only basic editing tools such as insert, cut, and paste are available.

To Edit a Payload File with the Text Editor

1. Right-click on a payload in the imaging navigation tree and select Edit. The payload panel appears.
2. Select Edit File from the Files pull-down menu. The Remote File Chooser appears.

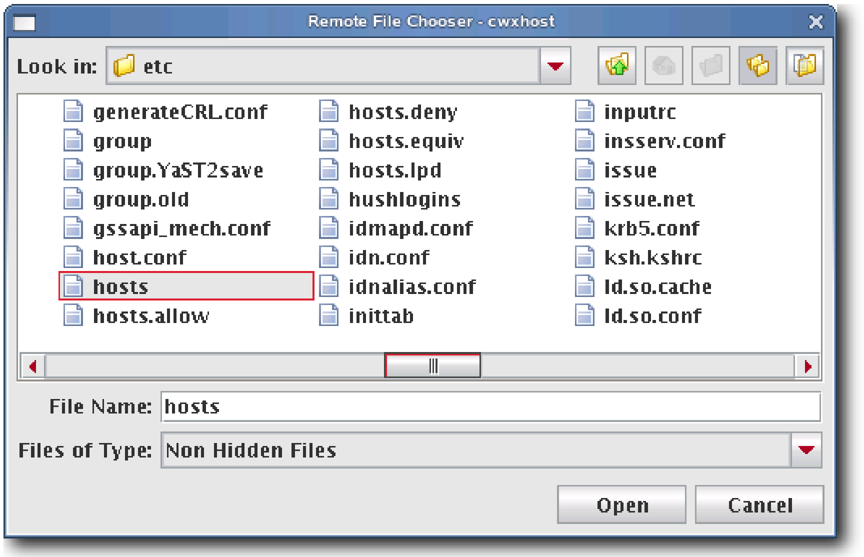
Remote File Chooser - cwxhost
generateCRL.conf
group
group.YaST2save
group.old
gssapi_mech.conf
host.conf
hosts
hosts.allow
hosts.deny
hosts.equiv
hosts.lpd
hushlogins
idmapd.conf
idn.conf
idnalias.conf
inittab
inputrc
insserv.conf
issue
issue.net
krb5.conf
ksh.kshrc
ld.so.cache
ld.so.conf
Open
Cancel

Look in: etc

File Name: hosts

Files of Type: Non Hidden Files

3. Select the file to edit and click Open. The text editor window appears.

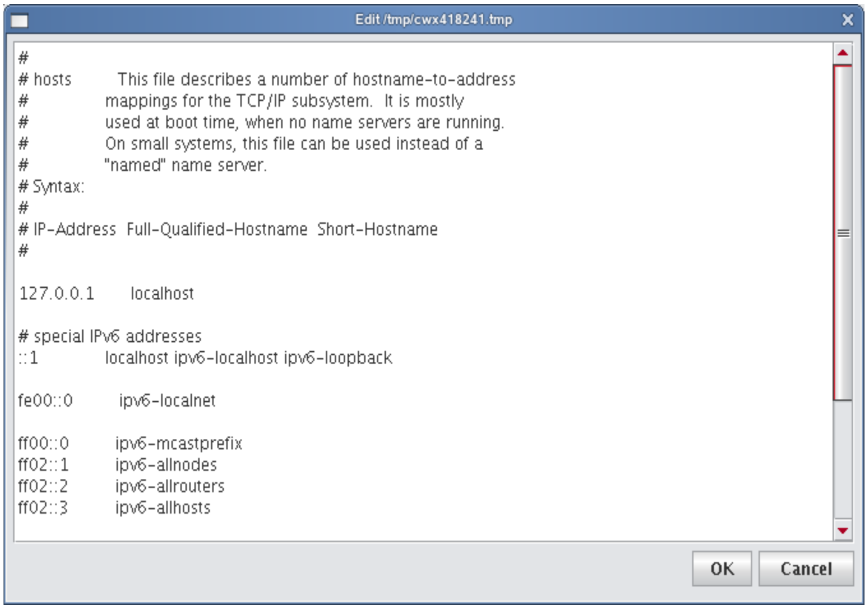
Edit /tmp/cwx418241.tmp
OK
Cancel

```
#
# hosts         This file describes a number of hostname-to-address
#               mappings for the TCP/IP subsystem.  It is mostly
#               used at boot time, when no name servers are running.
#               On small systems, this file can be used instead of a
#               "named" name server.
# Syntax:
#
# IP-Address  Full-Qualified-Hostname  Short-Hostname
#

127.0.0.1       localhost

# special IPv6 addresses
::1             localhost ipv6-localhost ipv6-loopback

fe00::0         ipv6-localnet

ff00::0         ipv6-mcastprefix
ff02::1         ipv6-allnodes
ff02::2         ipv6-allrouters
ff02::3         ipv6-allhosts
```

4. Edit the file as necessary, then click OK to save changes or click Cancel to abort this action.
5. Click Apply to complete the configuration. Click Revert or Close to abort this action.

Delete Payload Files

Deleting a payload file allows you to exclude a specific file(s) from a payload.

To Delete a File from a Payload

1. Right-click on a payload in the imaging navigation tree and select Edit. The payload panel appears.
2. Select Delete File from the Files pull-down menu. The Remote File Chooser appears.

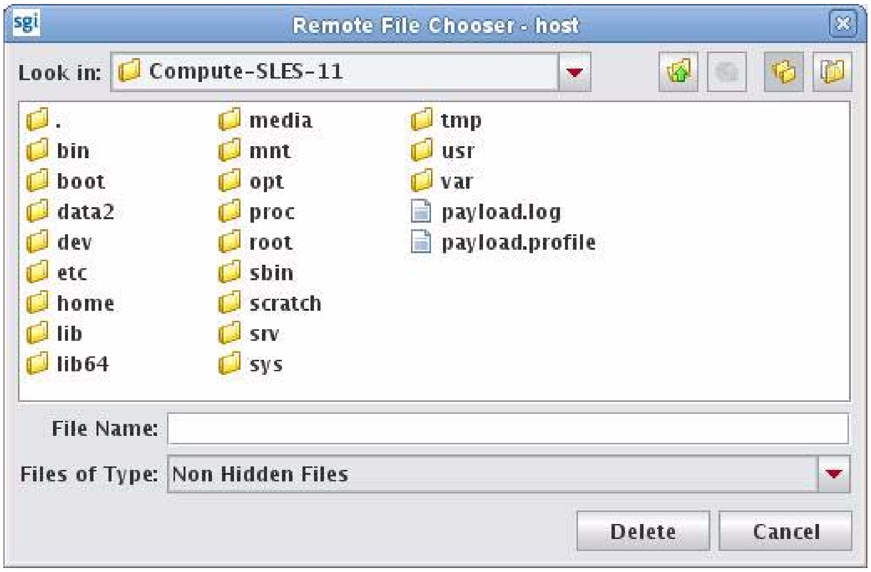

Look in: Compute-SLES-11

Files of Type: Non Hidden Files

sgi
Remote File Chooser - host
.
bin
boot
data2
dev
etc
home
lib
lib64
media
mnt
opt
proc
root
sbin
scratch
srv
sys
tmp
usr
var
payload.log
payload.profile
File Name:
Delete
Cancel

3. Select the file(s) you want to remove, then click Delete to remove the files or Cancel to abort this action.
4. Click Apply to complete the process. Click Revert or Close to abort this action.

Delete a Payload

To Delete a Working Copy of a Payload

Note

Before you delete the working copy of your payload, use the VCS status option to verify that the payload is checked in. See *Version Control System (VCS)* on page 140 for details on using version control.

Tip

Once you check the payload into VCS, you may remove the directory from within your working user directory (e.g., to save space):

$MGR_HOME/imaging/<username>/payloads/<name>

To verify that your changes were checked in, use the VCS status option. See *Version Control System (VCS)* on page 140 for details on using the version control system.

1. Right-click on a payload in the imaging navigation tree and select Delete.
2. Management Center asks you to confirm your action.

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.

Using the GUI

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`

From the Command Line

To install into the payload from the command line, set up the desired software repositories (See *Pre-Installation Steps* on page 5.) and then install the `sgimc-payload` package.

Example (RHEL)

For RHEL installations, you can use the Yum option `--installroot` to specify the payload `chroot` environment. However, you must first copy any applicable repositories from */etc/yum.repos.d/* into the payload `chroot`.

```
# cp /etc/yum.repos.d/yum-sgi.repo /opt/sgi/sgimc/imaging/root/payloads/Compute/etc/
yum.repos.d/yum-sgi.repo
```

```
# yum install sgimc-payload --installroot=/opt/sgi/sgimc/imaging/root/payloads/Compute
```

Example (SLES)

For SLES installations, you can use Zypper option `--root` to specify the payload `chroot` environment. However, you must first set up the same Zypper repositories in the payload.

This example assumes you have a SLES Zypper repository for your SLES distribution ISO mounted in */mnt/iso1* and that the installation media for SGI Managment Center is mounted in */mnt/iso3*.

1. Check which repositories are already configured:

```
# zypper --root=/opt/sgi/sgimc/imaging/root/payloads/Compute lr
```

2. Add repositories to the payload:

```
# zypper --root=/opt/sgi/sgimc/imaging/root/payloads/Compute ar dir:///mnt/iso1 'SuSE
Linux Enterprise Server 11 SP1'
# zypper --root=/opt/sgi/sgimc/imaging/root/payloads/Compute ar dir:///mnt/iso3 'SGI
Management Center 1.5'
```

3. Install the `sgimc-payload` package into the payload.

```
# zypper --root=/opt/sgi/sgimc/imaging/root/payloads/Compute install sgimc-payload
```

The `sgimc-payload` package will be installed into the payload along with all dependencies.

Note

If you are using SLES 10.x, use `sa` instead of `ar` to add the repository with Zypper. Likewise, to check for which repositories are already configured, use `sl` instead of `lr`.

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.

Create a Kernel

The following sections review the steps necessary to create a kernel for use in provisioning your cluster.

To Create a Kernel Using an Existing Binary

Note

For information on building a new kernel from source, see *To Build a New Kernel from Source* on page 110.

1. Select New Kernel from the File menu or right-click in the imaging navigation tree and select New Kernel. A new kernel pane appears.

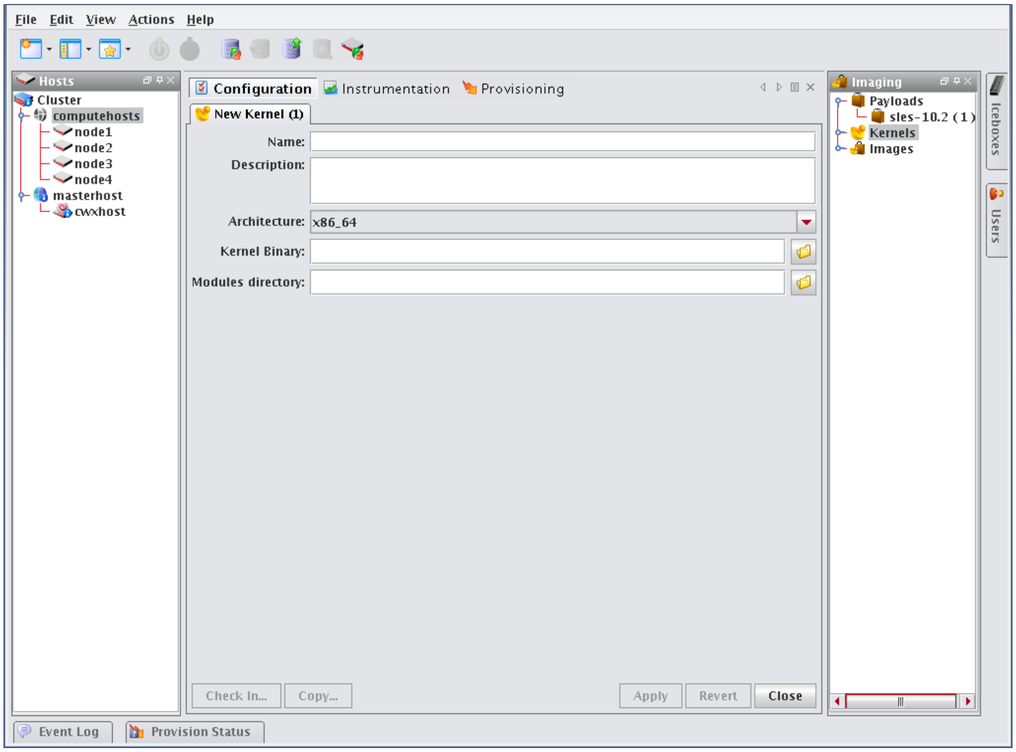
File
Edit
View
Actions
Help
Hosts
Cluster
computehosts
node1
node2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
New Kernel (1)
Name:
Description:
Kernel Binary:
Modules directory:
Check In...
Copy...
Apply
Revert
Close
Imaging
Payloads
sles-10.2 (1)
Kernels
Images
Iceboxes
Users
Event Log
Provision Status

Architecture: x86_64

2. Enter the name of the Kernel.
3. (Optional) Enter a description of the kernel.
4. Select the hardware architecture.

5. Specify the full path to the kernel binary or click Browse to open the Remote File Chooser and select the kernel binary.

! Warning!

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

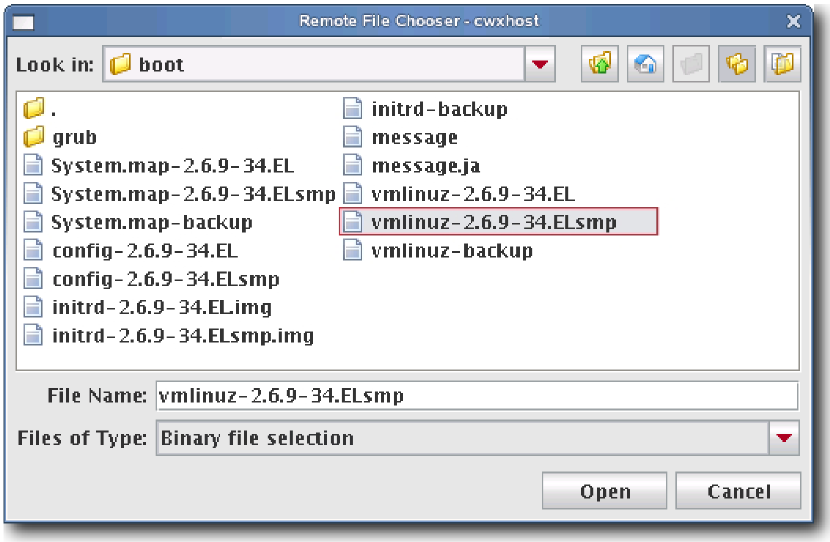

Remote File Chooser - cwxhost

Look in: boot

```
.                                initrd-backup
grub                             message
System.map-2.6.9-34.EL           message.ja
System.map-2.6.9-34.ELsmp        vmlinuz-2.6.9-34.EL
System.map-backup                vmlinuz-2.6.9-34.ELsmp
config-2.6.9-34.EL               vmlinuz-backup
config-2.6.9-34.ELsmp
initrd-2.6.9-34.EL.img
initrd-2.6.9-34.ELsmp.img
```

File Name: vmlinuz-2.6.9-34.ELsmp

Files of Type: Binary file selection

6. Specify the location of the modules directory (e.g., */lib/modules*) or click Browse to open the Remote File Chooser.

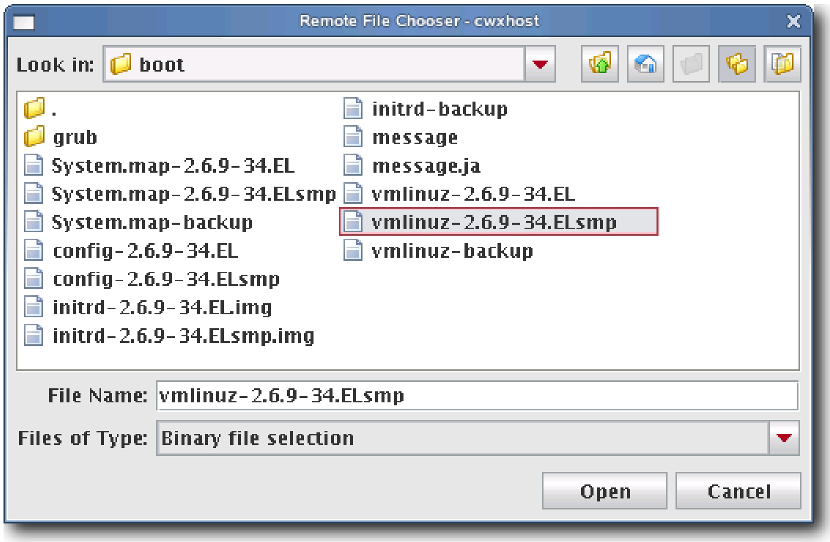

Remote File Chooser - cwxhost

Look in: boot

```
.                                initrd-backup
grub                             message
System.map-2.6.9-34.EL           message.ja
System.map-2.6.9-34.ELsmp        vmlinuz-2.6.9-34.EL
System.map-backup                vmlinuz-2.6.9-34.ELsmp
config-2.6.9-34.EL               vmlinuz-backup
config-2.6.9-34.ELsmp
initrd-2.6.9-34.EL.img
initrd-2.6.9-34.ELsmp.img
```

File Name: vmlinuz-2.6.9-34.ELsmp

Files of Type: Binary file selection

7. Select the modules directory and click Open.
8. Click Apply to create the kernel. Click Revert or Close to abort this action.

9. (Optional) Click Check In to import the kernel into VCS.

Note

To make configuration changes to the kernel, see *Edit a Kernel* on page 113.

To Create a Kernel from a Payload

If you have specific packages in your payload that contain specific kernel binaries or modules, you may need to create your kernel from the payload. This ensures that the modules and kernel binary in the kernel match exactly with what is contained in the payload.

To create the kernel from the payload, follow steps 1 through 9 from the preceding section *To Create a Kernel Using an Existing Binary* on page 107 but with the following modifications/clarifications:

- When selecting the binary (step 5), browse to *$MGR_HOME/imaging/root/payloads/<payload name>/boot/* (for example, */opt/sgi/sgimc/imaging/root/payloads/Compute/boot/*).
- After you select the correct kernel binary and browse for the modules directory (step 6), Management Center will default to a path inside your payload (for example, */opt/sgi/sgimc/imaging/root/payloads/Compute/lib/modules*). Select the appropriate modules directory and click **Open**.
- If you are creating a kernel to replace an existing kernel, you can see the list of modules that was used in the old kernel by examining */opt/sgi/sgimc/imaging/root/kernels/<kernel name>/kernel.profile*.

To Create a Copy of an Existing Kernel

1. Right-click on a kernel in the imaging navigation tree and select Copy.
2. Select a kernel from the navigation tree, then right-click on the payload and select Copy.

Tip

You may also open a kernel for editing, then click the Copy button at the lower left of the panel.

3. Management Center prompts you for the name of the new kernel.

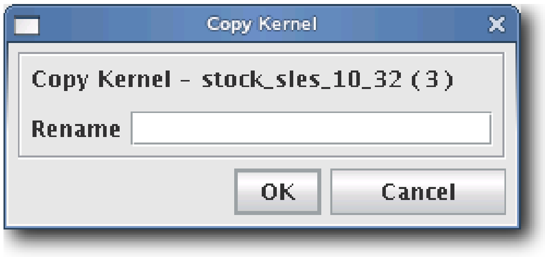

4. Enter the name of the new kernel and click OK. Click Cancel to abort this action.

To Build 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 107. 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. Select Source Kernel from the File menu. A new kernel pane appears.

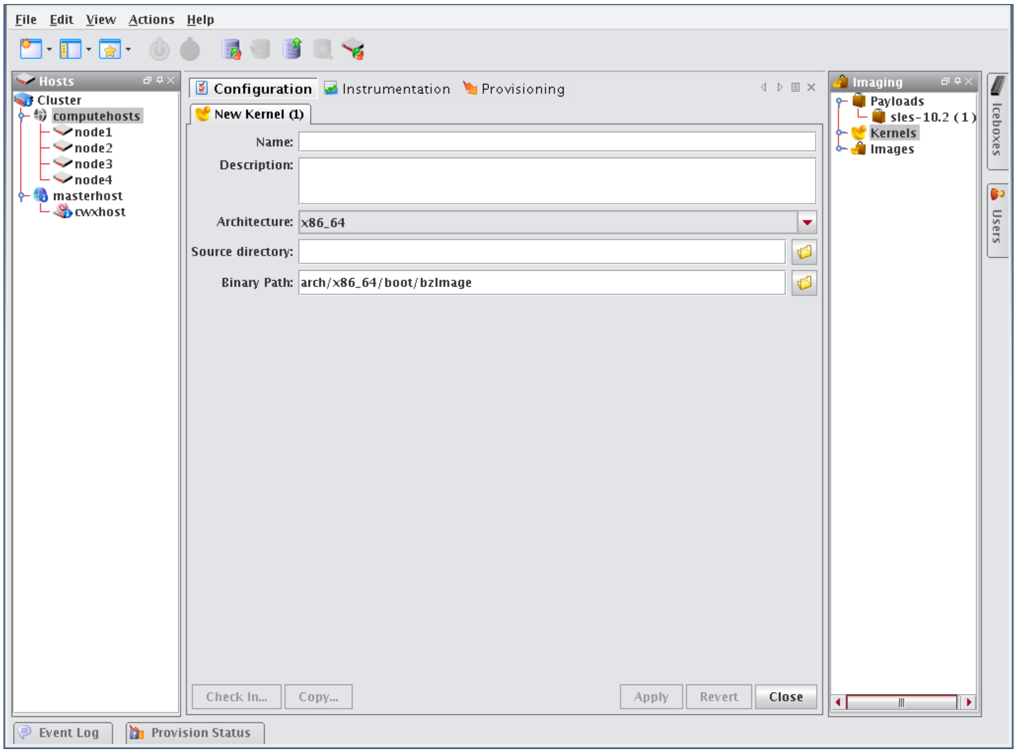
File Edit View Actions Help
Hosts
Cluster
computehosts
node1
node2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
New Kernel (1)
Imaging
Payloads
sles-10.2 (1)
Kernels
Images
Iceboxes
Users
Check In...
Copy...
Apply
Revert
Close
Event Log
Provision Status

Name:
Description:
Architecture: x86_64
Source directory:
Binary Path: arch/x86_64/boot/bzImage

6. 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 (i.e., where you unpacked the kernel source) in the Source Directory field or click Browse to open the Remote File Chooser. By default, kernel source files are located in */usr/src*.

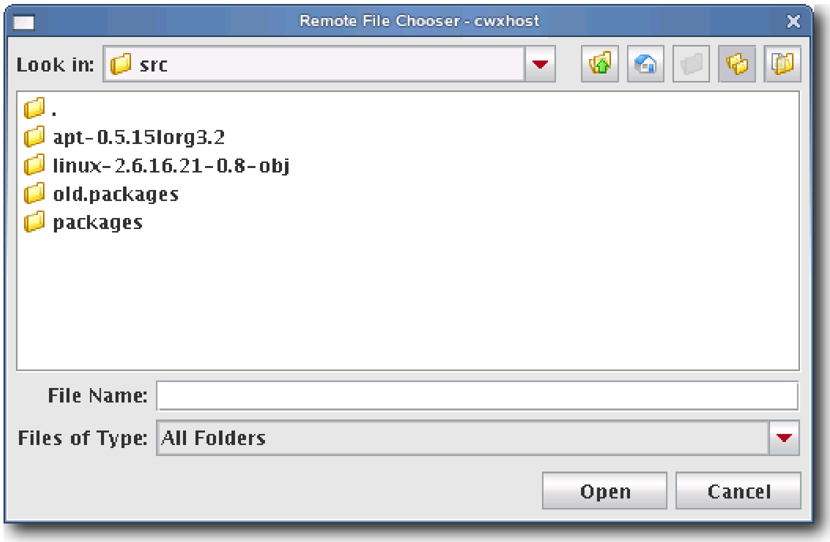
Remote File Chooser - cwxhost
Open
Cancel

Look in: src

.
apt-0.5.15lorg3.2
linux-2.6.16.21-0.8-obj
old.packages
packages

File Name:
Files of Type: All Folders

10. Select the source directory and click Open.
11. (Optional) Enter the binary path of the kernel (e.g., arch/i386/boot/bzImage) or click Browse to open the Remote File Chooser.

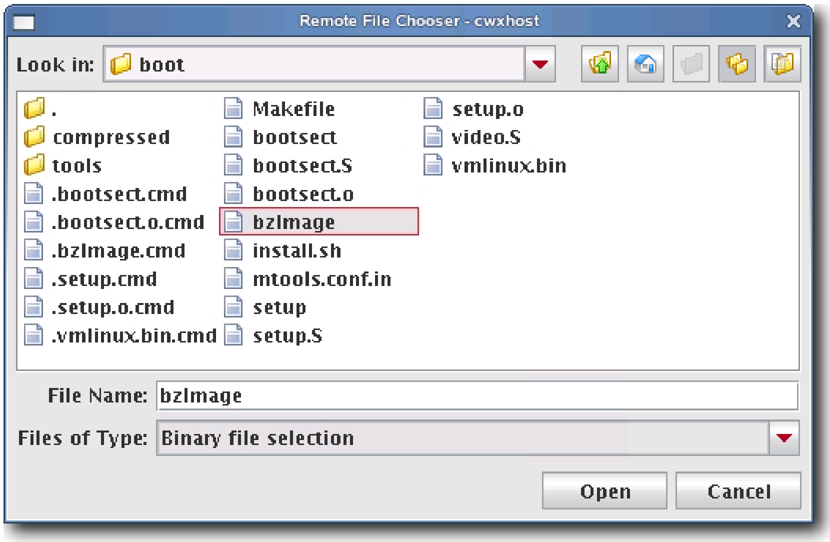
Remote File Chooser - cwxhost
.
compressed
tools
.bootsect.cmd
.bootsect.o.cmd
.bzImage.cmd
.setup.cmd
.setup.o.cmd
.vmlinux.bin.cmd
Makefile
bootsect
bootsect.S
bootsect.o
bzImage
install.sh
mtools.conf.in
setup
setup.S
setup.o
video.S
vmlinux.bin
Open
Cancel

Look in: boot

File Name: bzImage

Files of Type: Binary file selection

12. Select the modules directory and click Open.
13. Click Apply to create the kernel. Click Revert or Close to abort this action.
14. (Optional) Click Check In to import the kernel into VCS.

Note

To make configuration changes to the kernel, see *Edit a Kernel* on page 113.

Edit a Kernel

To Edit a Kernel

1. Right-click a kernel in the imaging navigation tree and select Edit.

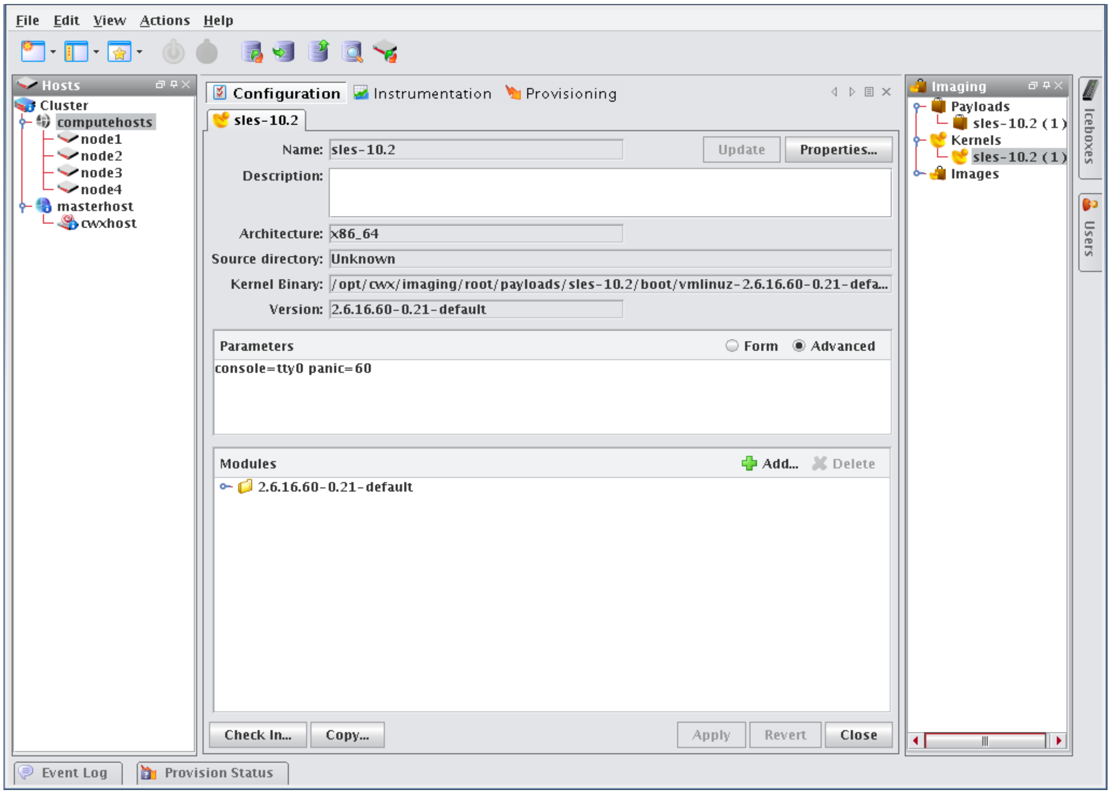
File Edit View Actions Help
Hosts
Cluster
computehosts
node1
node2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
sles-10.2
Update
Properties...
Parameters
Modules
Add...
Delete
Check In...
Copy...
Apply
Revert
Close
Imaging
Payloads
sles-10.2 (1)
Kernels
sles-10.2 (1)
Images
Iceboxes
Users
Event Log
Provision Status

Name: sles-10.2
Description:
Architecture: x86_64
Source directory: Unknown
Kernel Binary: /opt/cwx/imaging/root/payloads/sles-10.2/boot/vmlinuz-2.6.16.60-0.21-defa...
Version: 2.6.16.60-0.21-default

Parameters ○ Form ● Advanced

```
console=tty0 panic=60
```

Modules: 2.6.16.60-0.21-default

2. (Optional) Edit the kernel's description in the Description field.
3. (Optional) Click Update to update a kernel that has been recompiled for some reason (e.g., a change in kernel configuration). Management Center updates the kernel based on the Source Directory and Binary Path used when you created the kernel. See *To Create a Kernel Using an Existing Binary* on page 107.
4. (Optional) Click Properties to view the **.config* and *System.map* files for the kernel (if they existed when you imported the kernel).
5. (Optional) Edit the Parameters pane using the Form or Advanced view. The form view organizes and displays the basic required options and provides the default values required for IPMI. The Advanced view allows you to view all configurations in an editable text field and allows you to configure the kernel's command-line parameters string.
   - **A.** Select Serial Console to specify which console (tty0 or tty1) you will use to communicate with hosts.
   - **B.** Select Baud Rate to change the baud rate used on your system.
   - **C.** Select RAMdisk Size to change the size of the RAMdisk configured on your system.
6. (Optional) In the modules pane, click Add to include new modules in this kernel. You may select modules individually (files ending in **.ko*) or you can add a directory and allow Management Center to automatically select all modules and directories recursively. See *Modules* on page 114.
7. (Optional) In the modules pane, select any module(s) you want to remove from the kernel and click Delete.
8. Click Apply to complete the process. Click Revert or Close to abort this action.

9. (Optional) Click Check In to commit changes to the kernel into VCS.

10. (Optional) Click Copy to create a copy of this kernel. See *To Create a Copy of an Existing Kernel* on page 109.

MODULES

Many provisioning systems use a basic kernel to boot and provision the host, then reboot with an optimized kernel that will run on the host. Management Center requires only a single kernel to boot and run; however, you must compile any additional functionality into the kernel (i.e., *monolithic*) or add loadable kernel modules to the kernel (i.e., *modular*). Management Center loads the modules during the provisioning process.

Note

If you encounter problems when provisioning hosts on your cluster, check to see that you compiled your kernel correctly. If you compiled a modular kernel, you must include ethernet or file system modules before the host can provision properly. Use the serial console to watch the host boot.

Tip

In some cases, it may be necessary to install kernel modules on a host during the provisioning process, but not load them at boot time. Because an image ties a kernel and payload together, modules can be copied to the host by adding them to an image rather than adding them to a payload.

To add modules to an image, run mkdir -p ramdisk/lib/modules from the images directory. For example, if you were running as root and your image name were *ComputeHost*:

```
cd $MGR_HOME/imaging/root/images/ComputeHost
mkdir -p ramdisk/lib/modules/<linux name & version>/kernel/
mkdir -p ramdisk/lib/modules/<kernel name with version>/kernel/net/e1000
```

Then copy the modules you want to an appropriate subdirectory of the modules directory:

```
cp /usr/src/linux/drivers/net/e1000/e1000.ko
ramdisk/lib/modules/<linux name & version>/kernel/net
ramdisk/lib/modules/<linux name & version>/kernel/net/e1000/
```

You may wish to look at your local */lib/modules* directory if you have questions about the directory structure. During the boot process, the kernel automatically loads the modules that were selected in the kernel configuration screen. The additional modules will be copied to the host during the finalize stage. This method keeps the payload independent from the kernel and allows you to load the modules after the host boots.

Delete a Kernel

To Delete a Working Copy of a Kernel

1. Select the **Imaging** tab.

2. Right-click on the kernel in the imaging navigation tree and select Delete.

3. Management Center asks you to confirm your action.

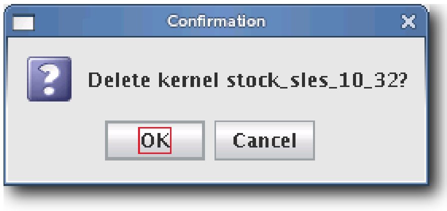

Delete kernel stock_sles_10_32?

Note

Before you delete the working copy of your kernel, check VCS to verify that the kernel is checked in. See *Version Control System (VCS)* on page 140 for details on using version control.

Tip

Once you check the kernel into VCS, you may delete the working copy of the kernel from your working directory (e.g., to save space).

$MGR_HOME/imaging/<*username*>/<*kernel*>/<*name*>

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. A New Image pane appears.

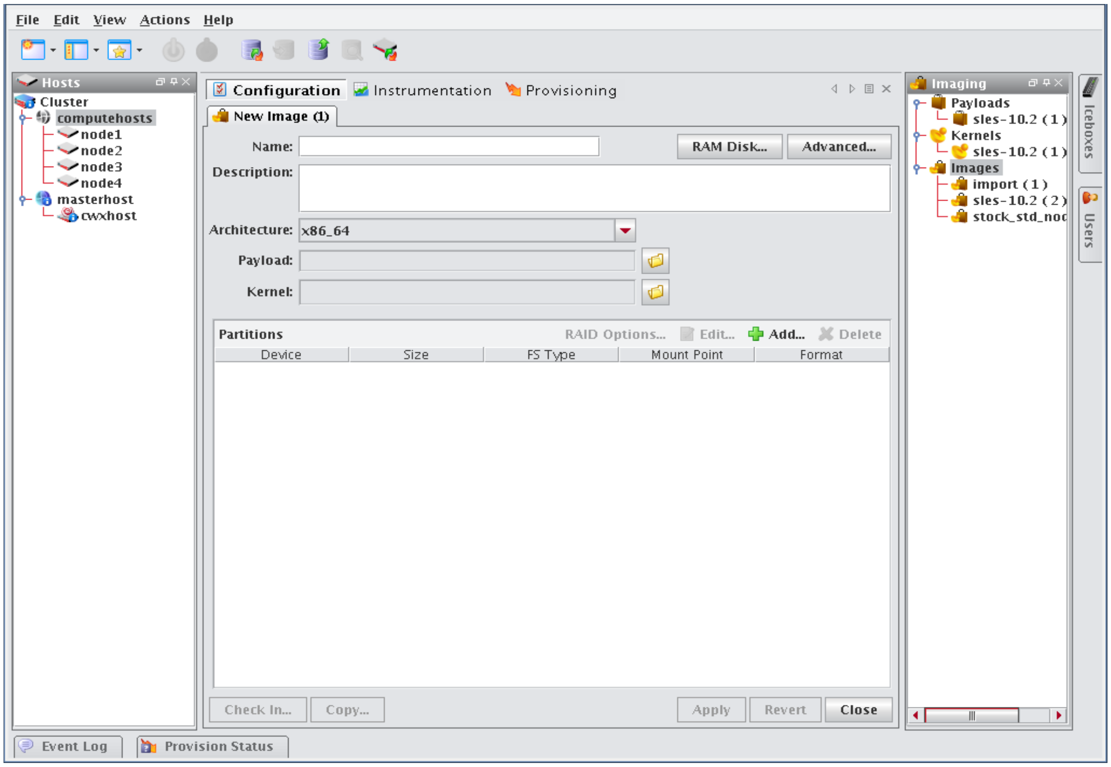
File
Edit
View
Actions
Help
Hosts
Cluster
computehosts
node1
node2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
New Image (1)
Name:
RAM Disk...
Advanced...
Description:
Payload:
Kernel:
Partitions
RAID Options...
Edit...
Add...
Delete
Device
Size
FS Type
Mount Point
Format
Check In...
Copy...
Apply
Revert
Close
Imaging
Payloads
sles-10.2 (1)
Kernels
sles-10.2 (1)
Images
import (1)
sles-10.2 (2)
stock_std_noc
Iceboxes
Users
Event Log
Provision Status

Architecture: x86_64

2. Enter the name of the new image in the Name field.
3. (Optional) Enter a description of the new image in the Description field.
4. Select the architecture supported by the kernel.
5. Select a Kernel by clicking Browse. To install additional kernel modules that do not load at boot time, see *Modules* on page 114.
6. Select a Payload by clicking Browse.

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 120.

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 120.
9. (Optional) If you need to make modifications to the way hosts boot during the provisioning process, select the RAM Disk tab. See *RAM Disk* on page 134.
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 to complete the process. Click Revert or Close to abort this action.

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. See *Advanced Provisioning Options* on page 151.

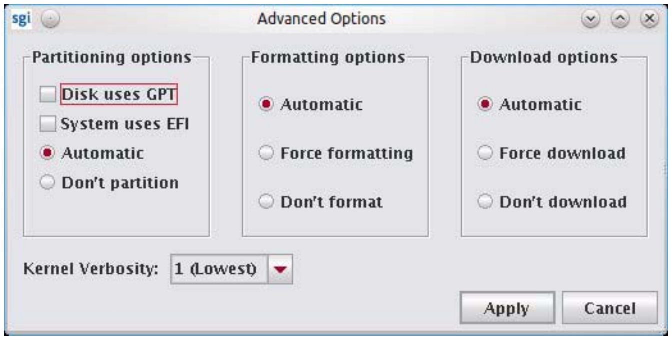

Partitioning options
☐ Disk uses GPT
☐ System uses EFI
● Automatic
○ Don't partition

Formatting options
● Automatic
○ Force formatting
○ Don't format

Download options
● Automatic
○ Force download
○ Don't download

Kernel Verbosity: 1 (Lowest)

PARTITIONING OPTIONS

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 or choose to never partition the host. You may also specify if the image should use GPT partition tables or EFI. See *Managing Partitions* on page 120.

FORMATTING OPTIONS

These options allow you to configure the partition formatting settings used when provisioning a host. You may automatically format when drives need to be formatted (for example, if the payload or the partitioning scheme changes), always re-create all partitions (including those that are exempt from being overwritten), or choose to never format.

DOWNLOAD OPTIONS

These 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.

KERNEL VERBOSITIY

The kernel verbosity level (1–8) allows you to control debug messages displayed by the kernel during provisioning. The default value 1 is the least verbose and 8 is the most.

boot.profile

Management Center generates the file, *boot.profile*, each time you save an image (overwriting the previous file in */etc/boot.profil*e). 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 117.
image.path	Specifies where to store the downloaded image. By default, */mnt*.

To change the configuration of one of these parameters, add the parameter (e.g., 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 147.

Note

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

To Create a Copy of an Existing Image

1. Select the **Imaging** tab.
2. Select an image from the navigation tree, then right-click on the image and select Copy.

Tip

You may also open an image for editing, then click the Copy button.

3. Management Center prompts you for the name of the new image.

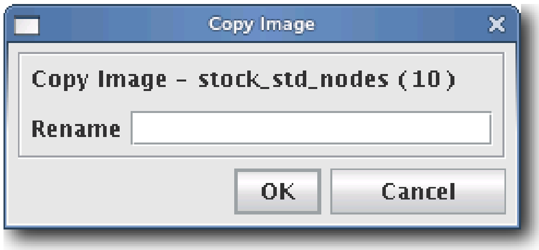

Copy Image - stock_std_nodes (10)

4. Enter the name of the new image and click OK. Click Cancel to abort this action.

Delete an Image

To Delete a Working Copy of an Image

1. Right-click an image in the imaging navigation tree and select Delete.
2. Management Center asks you to confirm your action.

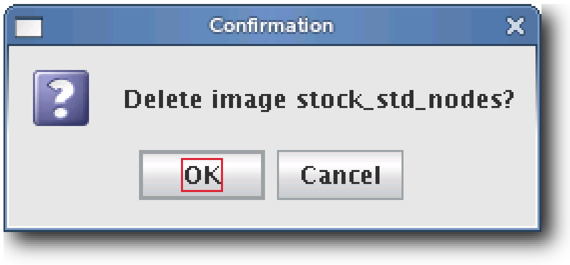

Delete image stock_std_nodes?

Tip

Once you check the image into VCS, you may remove the directory from within your working user directory (e.g., to save space).
$MGR_HOME/imaging/<*username*>/*images*/<*name*>

To verify that your changes were checked in, use the VCS status option. See *Version Control System (VCS)* on page 140 for details on using version control.

Managing Partitions

To Create a Partition for an Image

1. Right-click on an image in the imaging navigation tree and select Edit. The Image panel appears.

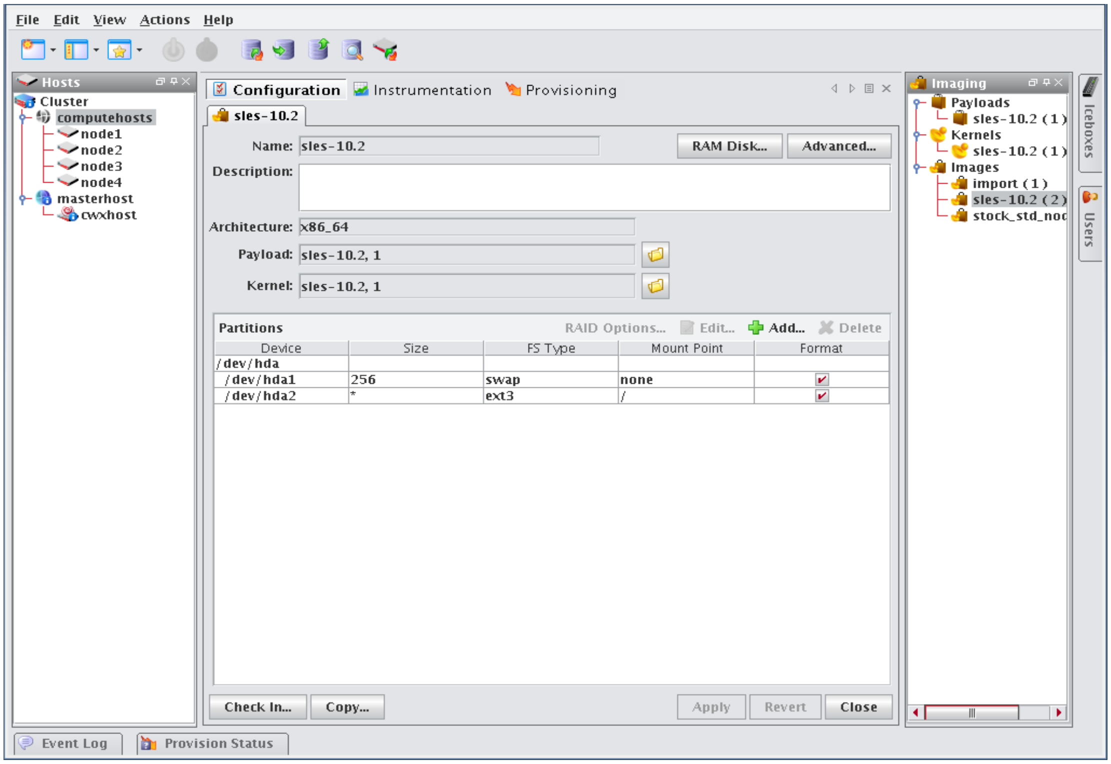
File
Edit
View
Actions
Help
Hosts
Cluster
computehosts
node1
node2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
sles-10.2
RAM Disk...
Advanced...
RAID Options...
Edit...
Add...
Delete
Check In...
Copy...
Apply
Revert
Close
Imaging
Payloads
sles-10.2 (1)
Kernels
sles-10.2 (1)
Images
import (1)
sles-10.2 (2)
stock_std_noc
Iceboxes
Users
Event Log
Provision Status

Name: sles-10.2

Description:

Architecture: x86_64

Payload: sles-10.2, 1

Kernel: sles-10.2, 1

Partitions

Device	Size	FS Type	Mount Point	Format
/dev/hda				
/dev/hda1	256	swap	none	☑
/dev/hda2	*	ext3	/	☑

2. In the partitions pane, click Add to create a new partition. The New Partition dialog appears.

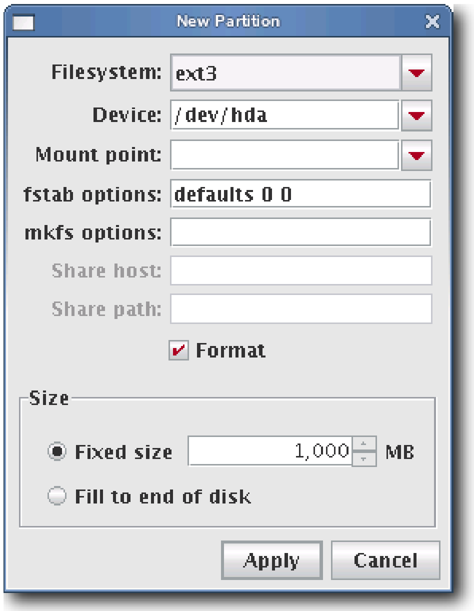

Filesystem: ext3
Device: /dev/hda
Mount point:
fstab options: defaults 0 0
mkfs options:
Share host:
Share path:
☑ Format
Size
● Fixed size 1,000 MB
○ Fill to end of disk

3. Select a file system type from the Filesystem pull-down menu. To create a diskless host, see *Diskless Hosts* on page 131.
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 pull-down menu.

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 (i.e., file size limits, symlinks, journalling). 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) Un-check 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 formatting option or from the *boot.profile* (see *Select an Image and Provision* on page 147 and *boot.profile* on page 118).

Note

After partitioning the hard disk(s) 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.

For nodes with external storage, detach the storage when provisioning. The discovery order may present the external storage first and, consequently, Management Center will use the storage for the filesystems it manages.

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

Tip

It is wise to 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 to save changes or click Cancel to abort this action.
13. (Optional) Click Check In to import the image into VCS.
14. Click Apply to complete the process. Click Revert or Close to abort this action.

Note

Management Center generates the file, *boot.profile*, each time you save an image. For a description of the information contained in this file, see *boot.profile* on page 118.

RAID Partitions

To Create a RAID Partition

Note

When adding a RAID partition, the host typically requires two disks and at least two previously created software RAID partitions (one per disk).

1. Right-click on an image in the imaging navigation tree and select Edit. The image pane appears.
2. In the partitions pane, click Add to create the appropriate number of software RAID partitions for the RAID you are creating. See *To Create a Partition for an Image* on page 120.

Note

The RAID button is disabled until you create at least two RAID partitions.

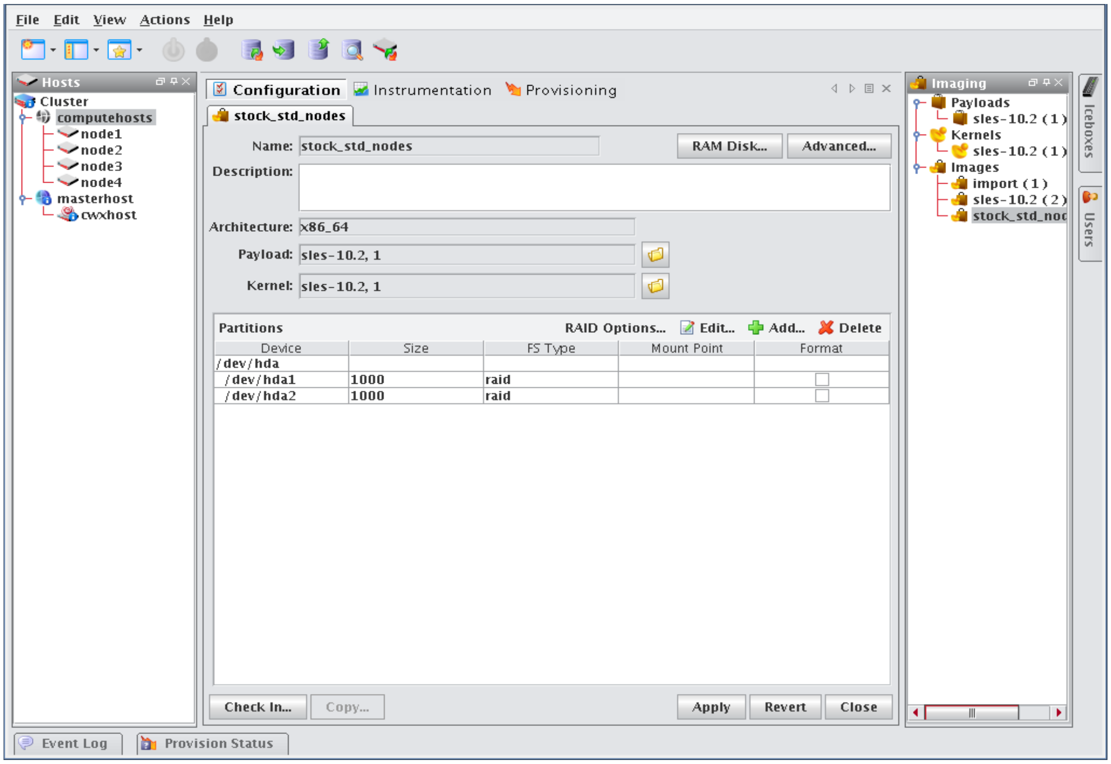
File
Edit
View
Actions
Help
Hosts
Cluster
computehosts
node1
node2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
stock_std_nodes
RAM Disk...
Advanced...
Description:
Partitions
RAID Options...
Edit...
Add...
Delete
Check In...
Copy...
Apply
Revert
Close
Imaging
Payloads
sles-10.2 (1)
Kernels
sles-10.2 (1)
Images
import (1)
sles-10.2 (2)
stock_std_noc
Iceboxes
Users
Event Log
Provision Status

Name: stock_std_nodes

Architecture: x86_64

Payload: sles-10.2, 1

Kernel: sles-10.2, 1

Device	Size	FS Type	Mount Point	Format
/dev/hda				
/dev/hda1	1000	raid		☐
/dev/hda2	1000	raid		☐

3. Click the RAID button to assign the partitions a file system, mount point, and RAID level. The Add RAID dialogue appears.

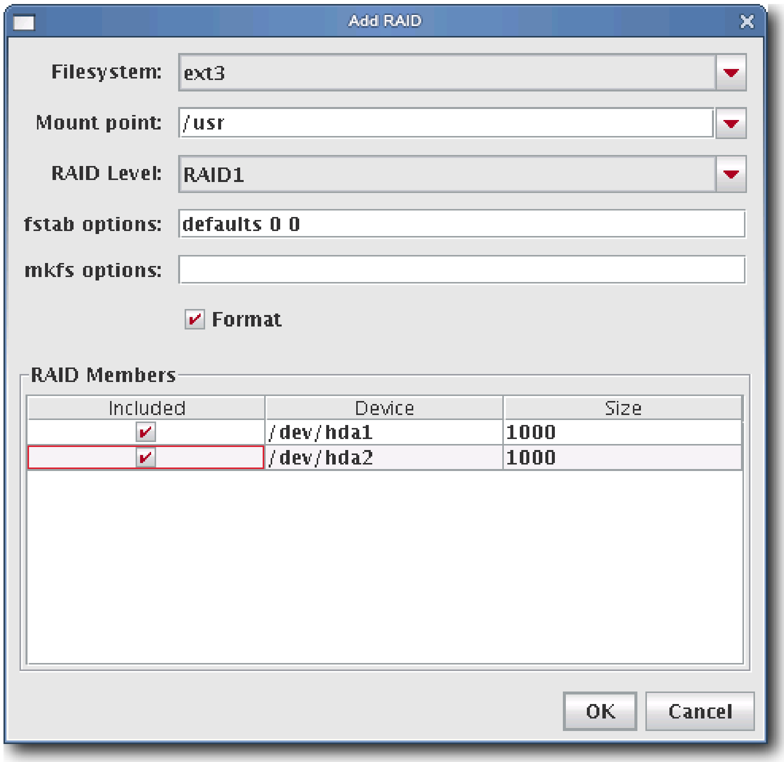

Filesystem: ext3

Mount point: /usr

RAID Level: RAID1

fstab options: defaults 0 0

mkfs options:

☑ Format

RAID Members

Included	Device	Size
☑	/dev/hda1	1000
☑	/dev/hda2	1000

4. Select a file system type from the Filesystem pull-down menu.
5. Enter a Mount point or select one from the pull-down menu.
6. Select a RAID level from the RAID Level pull-down menu. This level affects the size of the resulting RAID and the number of RAID partitions required to create it (e.g., RAID0 and RAID1 require 2 RAID partitions, RAID5 requires 3 RAID partitions).
7. (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.
8. (Optional) Enter the mkfs options to use when creating the file system (i.e., file size limits, symlinks, journalling). For example, to change the default block size for ext3 to 4096, enter -b 4096 in the mkfs field.
9. From the RAID Members list, select the currently unused RAID partitions to include in this RAID.
10. Click OK to save changes or click Cancel to abort this action.
11. Click Apply to complete the process. Click Revert or Close to abort this action.

Edit a Partition

To Edit a Partition on an Image

1. Right-click an image in the imaging navigation tree and select Edit. The image panel appears.
2. In the partitions pane, select the partition you want to edit from the list of partitions.

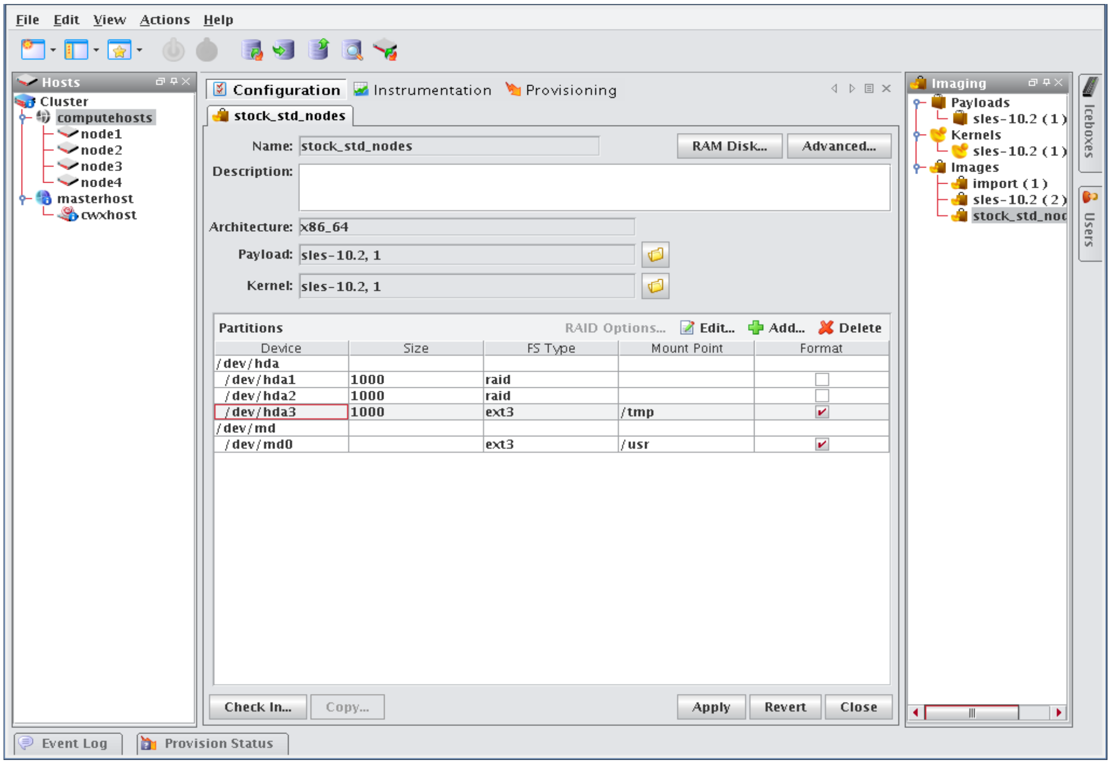
File
Edit
View
Actions
Help
Hosts
Cluster
computehosts
node1
node2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
stock_std_nodes
RAM Disk...
Advanced...
Partitions
RAID Options...
Edit...
Add...
Delete
Check In...
Copy...
Apply
Revert
Close
Imaging
Payloads
sles-10.2 (1)
Kernels
sles-10.2 (1)
Images
import (1)
sles-10.2 (2)
stock_std_noc
Iceboxes
Users
Event Log
Provision Status

Name: stock_std_nodes

Description:

Architecture: x86_64

Payload: sles-10.2, 1

Kernel: sles-10.2, 1

Device	Size	FS Type	Mount Point	Format
/dev/hda				
/dev/hda1	1000	raid		☐
/dev/hda2	1000	raid		☐
/dev/hda3	1000	ext3	/tmp	☑
/dev/md				
/dev/md0		ext3	/usr	☑

3. Click Edit in the partitions pane. The Edit Partition dialog appears.

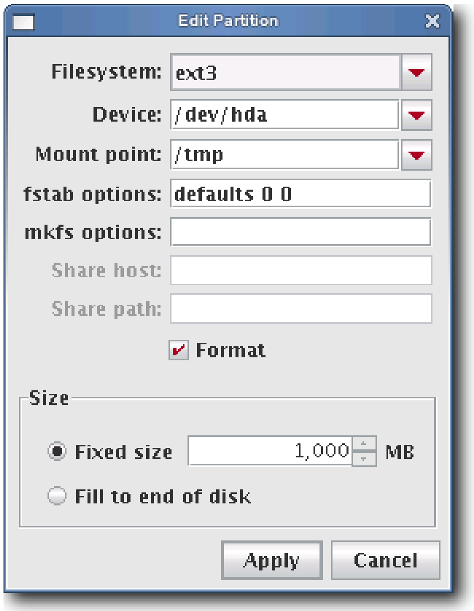

Edit Partition

Filesystem: ext3

Device: /dev/hda

Mount point: /tmp

fstab options: defaults 0 0

mkfs options:

Share host:

Share path:

☑ Format

Size

● Fixed size 1,000 MB

○ Fill to end of disk

4. Make any necessary changes to the partition, then click Apply to accept the changes. Click Cancel to abort this action.

Delete a Partition

To Delete a Partition from an Image

1. Right-click an image in the imaging navigation tree and select Edit. The image panel appears.

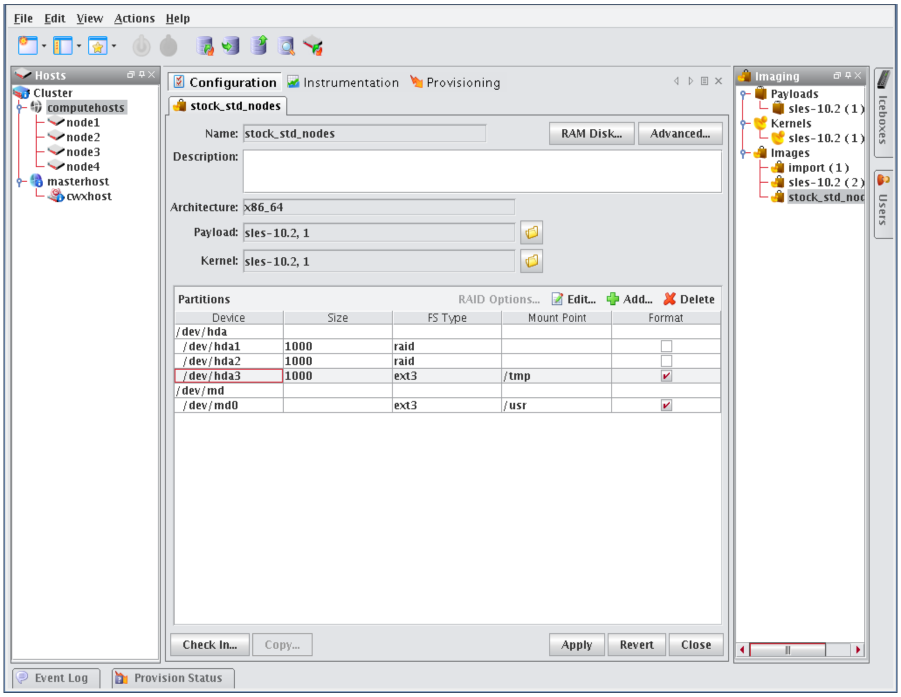

File
Edit
View
Actions
Help
Hosts
Cluster
computehosts
node1
node2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
stock_std_nodes
RAM Disk...
Advanced...
Description:
Partitions
RAID Options...
Edit...
Add...
Delete
Check In...
Copy...
Apply
Revert
Close
Imaging
Payloads
sles-10.2 (1)
Kernels
sles-10.2 (1)
Images
import (1)
sles-10.2 (2)
stock_std_noc
Iceboxes
Users
Event Log
Provision Status

Name: stock_std_nodes

Architecture: x86_64

Payload: sles-10.2, 1

Kernel: sles-10.2, 1

Device	Size	FS Type	Mount Point	Format
/dev/hda				
/dev/hda1	1000	raid		☐
/dev/hda2	1000	raid		☐
/dev/hda3	1000	ext3	/tmp	☑
/dev/md				
/dev/md0		ext3	/usr	☑

2. From the partitions pane, select the partition you want to delete from the list of partitions. To select multiple partitions, use the Shift or Ctrl keys.
3. Click Delete.

User-Defined File Systems

Establishing a user-defined file system allows you to create a raw partition that you may format with a file system not supported by Management Center.

To Create a Partition with a User-defined File System

1. Right-click on an image in the imaging navigation tree and select Edit. The image panel appears.

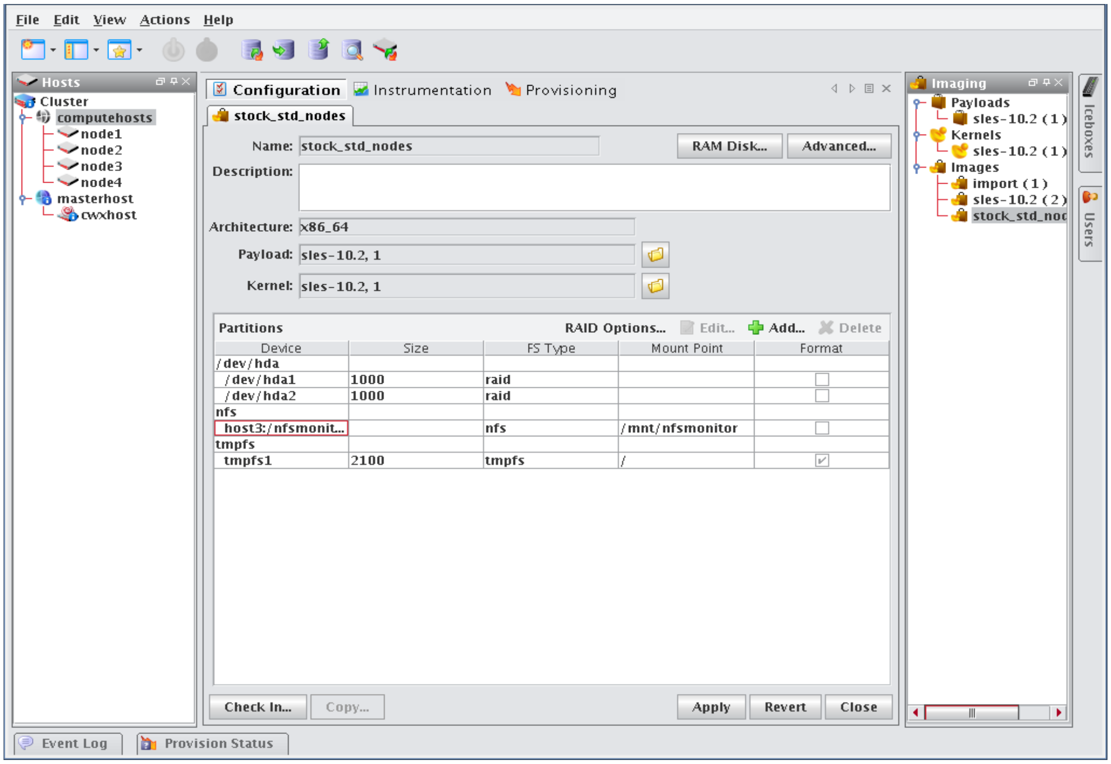
File Edit View Actions Help
Hosts
Cluster
computehosts
node1
node2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
stock_std_nodes
RAM Disk...
Advanced...
Partitions
RAID Options...
Edit...
Add...
Delete
Check In...
Copy...
Apply
Revert
Close
Imaging
Payloads
sles-10.2 (1)
Kernels
sles-10.2 (1)
Images
import (1)
sles-10.2 (2)
stock_std_noc
Iceboxes
Users
Event Log
Provision Status

Name: stock_std_nodes

Description:

Architecture: x86_64

Payload: sles-10.2, 1

Kernel: sles-10.2, 1

Device	Size	FS Type	Mount Point	Format
/dev/hda				
/dev/hda1	1000	raid		☐
/dev/hda2	1000	raid		☐
nfs				
host3:/nfsmonit...		nfs	/mnt/nfsmonitor	☐
tmpfs				
tmpfs1	2100	tmpfs	/	☑

2. From the partitions pane, click Add. The New Partition dialog appears.

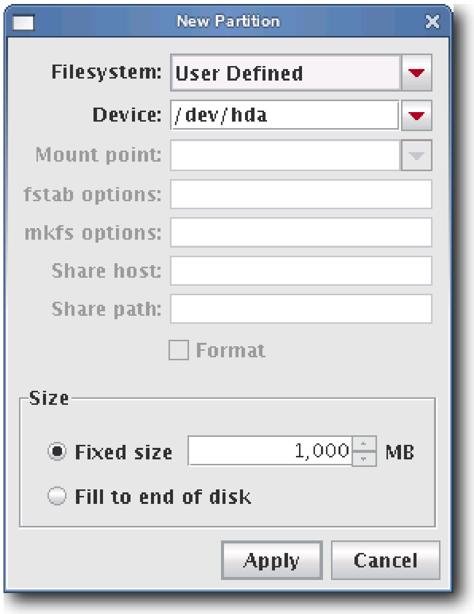

3. Select User Defined from the Filesystem pull-down menu.

4. Enter the device on which to add the partition or select a device from the pull-down menu. 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. Create a plug-in to create the user-defined file system. Everything required to build and mount the file system will need to be included in the RAMdisk. Kernel modules needed to support the file system must be added to the kernel you selected. See *Plug-ins for the Boot Process* on page 136.

6. Select the partition size:
   - Fixed partition size allows you to define the size of the partition (in MBs).
   - Fill to end of disk allows you to create a partition that uses any space that remains after defining partitions with fixed sizes.
7. Click Apply to save changes or click Cancel to abort this action.
8. Click Check In to import the image into VCS.
9. Click Apply to complete the process. Click Revert or Close to abort this action.

Note

Management Center generates the file, *boot.profile*, each time you save an image. See *boot.profile* on page 118 for a description of the information contained in this file.

Diskless Hosts

Management Center provides support for diskless hosts. For optimal performance, Management Center implements diskless hosts by installing the operating system into the host's physical memory, generally referred to as RAMfs or TmpFS. Because the OS is stored in memory, it is recommended that you use a minimal Linux installation to avoid consuming excess memory. An optimized Linux installation is typically around 100-150MB, but may be as small as 30MB depending on which libraries are installed. Management Center also supports local scratch or swap space on the hosts.

Note

Potentially large directories like */home* should never be stored in RAM. Rather, they should be shared through a global storage solution.

! Warning!

When using diskless hosts, the file system is stored in memory. Changes made to the host's file system will be lost when the host reboots. If changes are required, make them in the payload first.

SGI offers secure diskless systems for classified environments. These include integration of micro installation with a globally mounted file system and scripts that optimize and simplify diskless management. Additional options for diskless systems are available through SGI Professional Services. Please contact SGI or speak with your SGI representative for more information.

To Configure a Diskless Host

1. Right-click on an image in the imaging navigation tree and select Edit. The image panel appears.

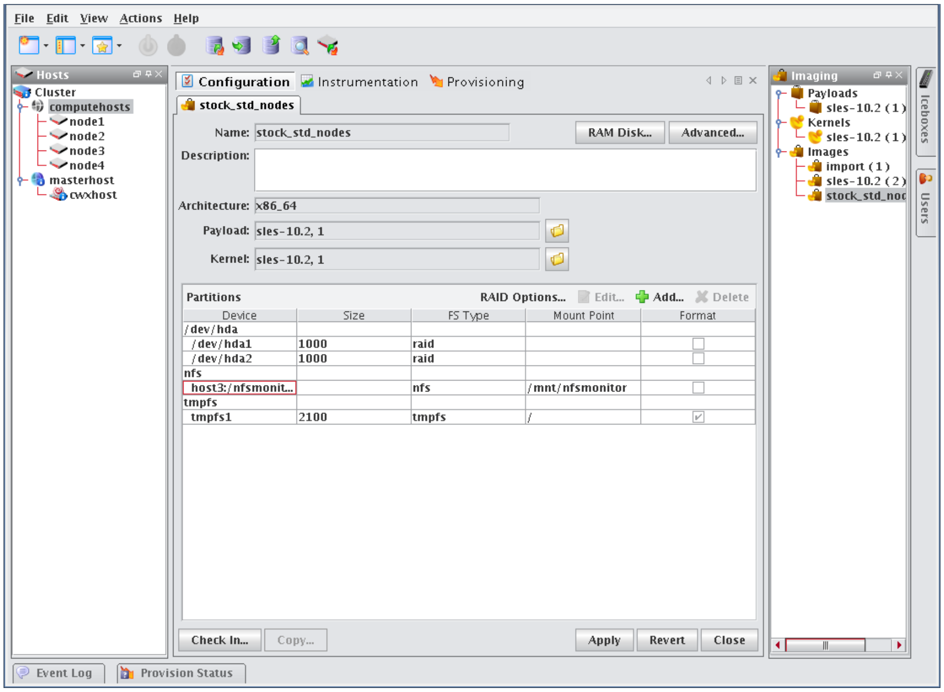
File
Edit
View
Actions
Help
Hosts
Cluster
computehosts
node1
node2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
stock_std_nodes
RAM Disk...
Advanced...
Partitions
RAID Options...
Edit...
Add...
Delete
Check In...
Copy...
Apply
Revert
Close
Imaging
Payloads
sles-10.2 (1)
Kernels
sles-10.2 (1)
Images
import (1)
sles-10.2 (2)
stock_std_noc
Iceboxes
Users
Event Log
Provision Status

Name: stock_std_nodes

Description:

Architecture: x86_64

Payload: sles-10.2, 1

Kernel: sles-10.2, 1

Device	Size	FS Type	Mount Point	Format
/dev/hda				
/dev/hda1	1000	raid		☐
/dev/hda2	1000	raid		☐
nfs				
host3:/nfsmonit...		nfs	/mnt/nfsmonitor	☐
tmpfs				
tmpfs1	2100	tmpfs	/	☑

2. From the partitions pane, click Add. The New Partition dialog appears.

3. Select the tmpfs or nfs file system type from the Filesystem pull-down menu.

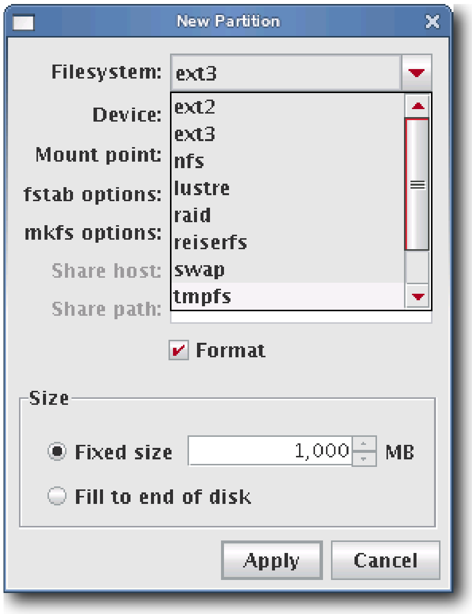
New Partition
Filesystem: ext3
ext2
ext3
nfs
lustre
raid
reiserfs
swap
tmpfs
Device:
Mount point:
fstab options:
mkfs options:
Share host:
Share path:
☑ Format
Size
● Fixed size 1,000 MB
○ Fill to end of disk
Apply
Cancel

Note

Although diskless hosts may use either tmpfs or nfs partitions, they must use only one type. If you are converting or editing a diskless host, change all partitions to the same type.

4. Enter the Mount Point or select one from the pull-down menu (diskless hosts use root "/" as the mount point).

Tip

In most Linux installations, the majority of the OS is stored in the */usr* directory. To help conserve memory, you may elect to share the */usr* directory via NFS or another global file system.

5. (Optional) Enter the fstab options. The */etc/fstab* file controls where directories are mounted.

Note

Because Management Center writes and manages the fstab on the hosts, any changes made on the hosts are overwritten during provisioning.

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

Tip

It is wise to 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.

It is important to note that memory allocated to a partition is not permanently consumed. For example, consider programs that need to write temporary files in a */tmp* partition. Although you may configure the partition to use a maximum of 50 MB of memory, the actual amount used depends on the contents of the partition. If the */tmp* partition is empty, the amount of memory used is 0 MB.

---

7. Click Apply to save changes or click Cancel to abort this action.
8. Click Check In to import the image into VCS.

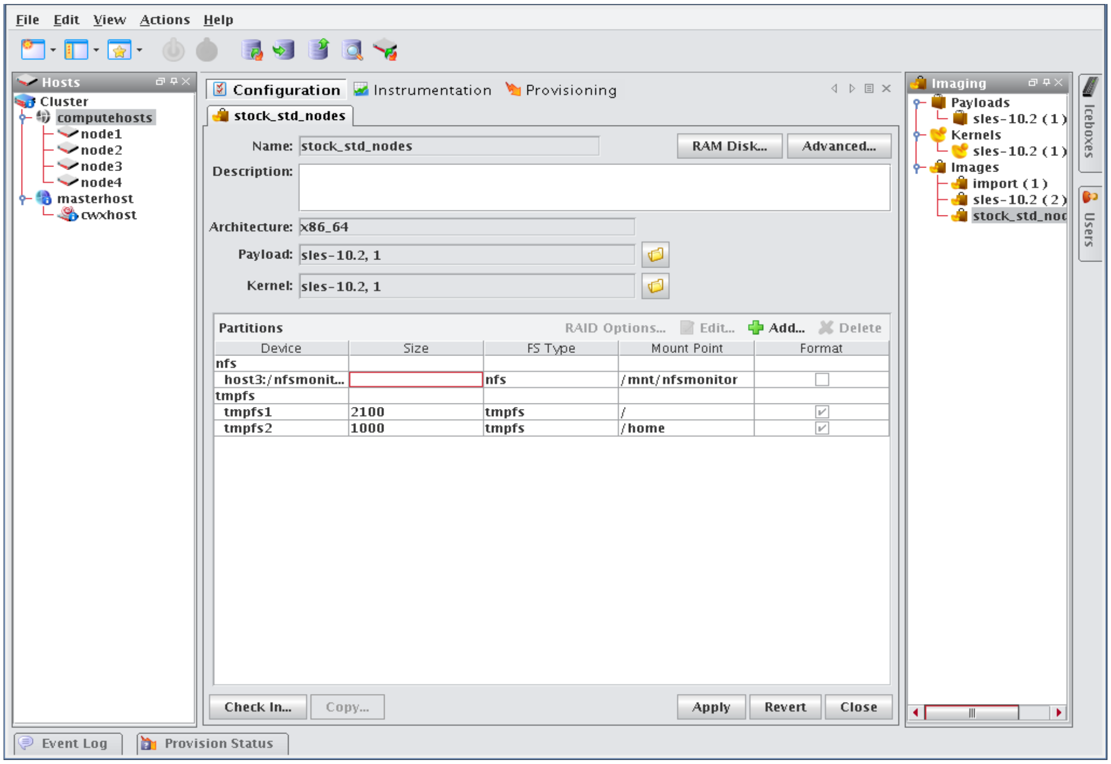
File
Edit
View
Actions
Help
Hosts
Cluster
computehosts
node1
node2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
stock_std_nodes
RAM Disk...
Advanced...
Partitions
RAID Options...
Edit...
Add...
Delete
Check In...
Copy...
Apply
Revert
Close
Imaging
Payloads
sles-10.2 (1)
Kernels
sles-10.2 (1)
Images
import (1)
sles-10.2 (2)
stock_std_noc
Iceboxes
Users
Event Log
Provision Status

Name: stock_std_nodes

Description:

Architecture: x86_64

Payload: sles-10.2, 1

Kernel: sles-10.2, 1

Device	Size	FS Type	Mount Point	Format
nfs				
host3:/nfsmonit...		nfs	/mnt/nfsmonitor	☐
tmpfs				
tmpfs1	2100	tmpfs	/	☑
tmpfs2	1000	tmpfs	/home	☑

9. Click Apply to complete the process. Click Revert or Close to abort this action.

Note

Management Center generates the file, *boot.profile*, each time you save an image. See *boot.profile* on page 118 for a description of the information contained in this file.

RAM Disk

The RAM Disk is a small disk image that is created and loaded with the utilities required to provision the host. When the host first powers on, it loads the kernel and mounts the RAM Disk as the root file system. 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).

Note

Management Center uses two *skeletal* RAM Disks—one for ia32 and another for both AMD-64 and EM64T. These skeletal disks are located in *$MGR_HOME/ramdisks* and should never be modified manually. All changes must be performed through Management Center or in *$MGR_HOME/imaging/<username>/images/<image_name>/ramdisk*.

! Warning!

Modifications made to the skeletal RAM Disks are permanent for ALL images.

To Add Boot Utilities

Adding boot utilities to the RAM Disk allows you to create such things as custom, pre-finalized scripts using utilities that are not required for standard Linux versions.

1. Right-click on an image n the imaging navigation tree and select Edit. The image panel appears.

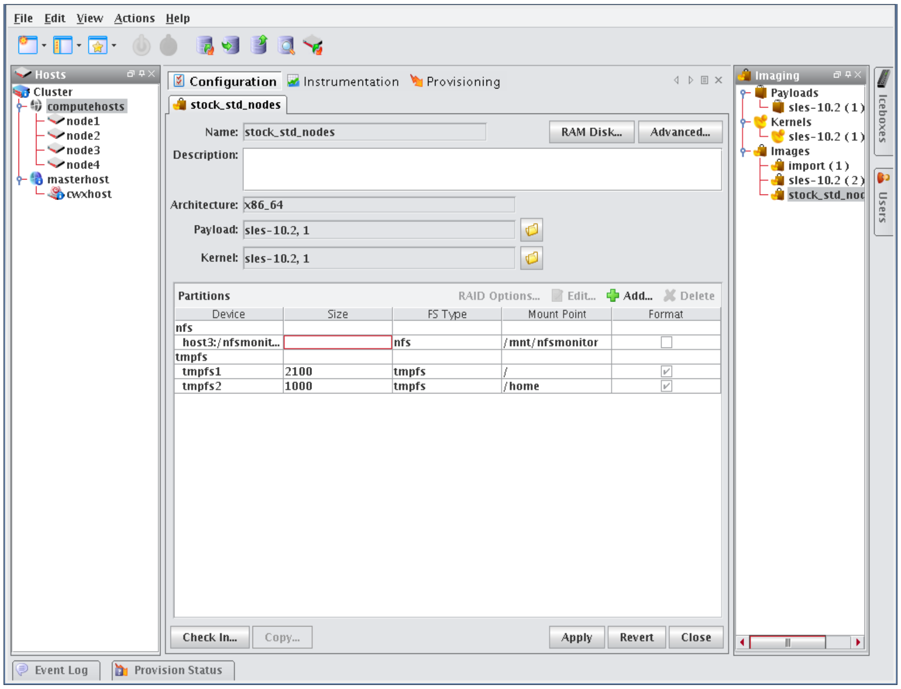
File
Edit
View
Actions
Help
Hosts
Cluster
computehosts
node1
node2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
stock_std_nodes
RAM Disk...
Advanced...
RAID Options...
Edit...
Add...
Delete
Check In...
Copy...
Apply
Revert
Close
Imaging
Payloads
sles-10.2 (1)
Kernels
sles-10.2 (1)
Images
import (1)
sles-10.2 (2)
stock_std_noc
Iceboxes
Users
Event Log
Provision Status

Name: stock_std_nodes

Description:

Architecture: x86_64

Payload: sles-10.2, 1

Kernel: sles-10.2, 1

Partitions

Device	Size	FS Type	Mount Point	Format
nfs				
host3:/nfsmonit...		nfs	/mnt/nfsmonitor	☐
tmpfs				
tmpfs1	2100	tmpfs	/	☑
tmpfs2	1000	tmpfs	/home	☑

2. Click the RAM Disk button. The RAM Disk dialog appears. Default files from the skeleton RAM Disk are grayed out—any changes or updates appear in black.

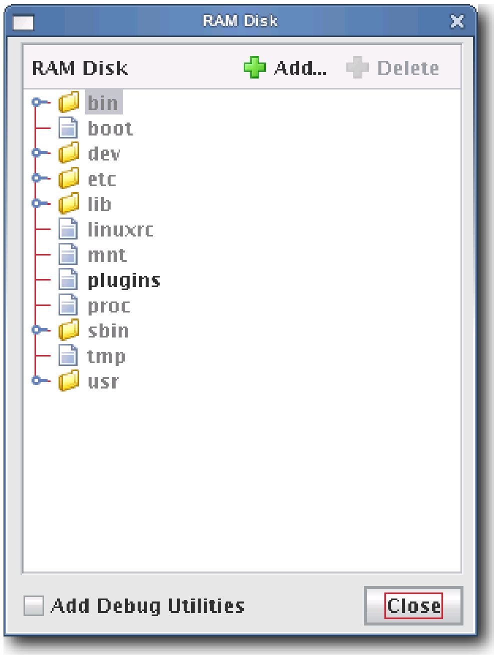
RAM Disk
RAM Disk
Add...
Delete
bin
boot
dev
etc
lib
linuxrc
mnt
plugins
proc
sbin
tmp
usr
Add Debug Utilities
Close

3. Click Add. The Add File to RAM Disk dialog appears.

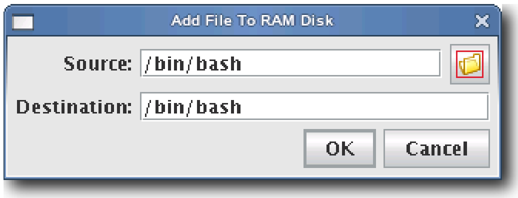
Add File To RAM Disk
OK
Cancel

Source: /bin/bash
Destination: /bin/bash

4. Enter the boot utility path in the Source field or click Browse to locate a utility.
5. Specify the Destination location in which to install the boot utility in the RAM Disk file system.
6. Click OK to install the boot utility or click Cancel to abort this action.
7. (Optional) Select Add Debug Utilities to apply additional debugging utilities to the RAM Disk.
8. Click Apply to complete the process. Click Revert or Close to abort this action.

Note

Management Center generates the file, *boot.profile*, each time you save an image. See *boot.profile* on page 118 for a description of the information contained in this file.

Plug-ins for the Boot Process

A host requires a boot process to initialize hardware, load drivers, and complete the necessary tasks to initiate a login prompt. The boot process is composed of five main stages and allows you to include additional plug-ins at each stage to expand system capabilities. During the boot process, the system moves from stage to stage installing any plug-ins specified. If you do not specify any plug-ins, the host will boot using the built-in boot process. The boot process is as follows:

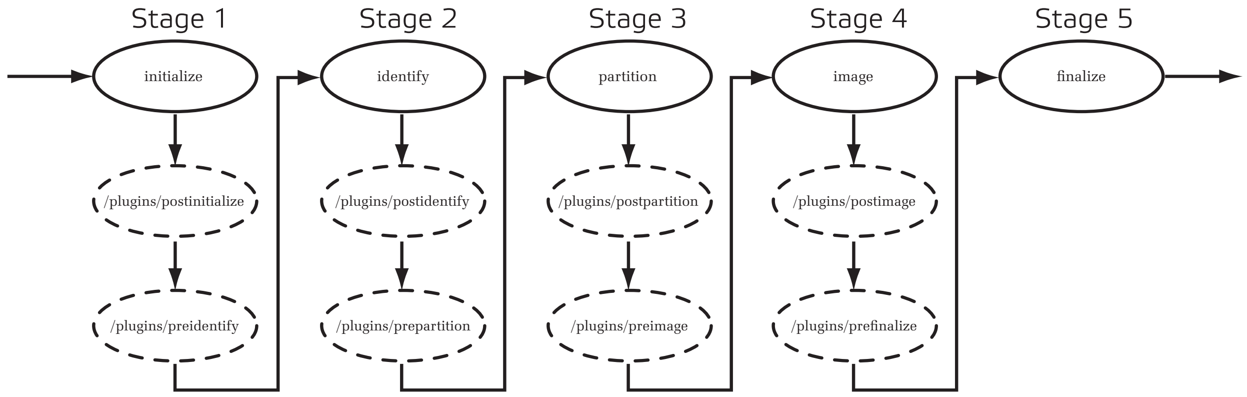
Stage 1
Stage 2
Stage 3
Stage 4
Stage 5
initialize
identify
partition
image
finalize
/plugins/postinitialize
/plugins/postidentify
/plugins/postpartition
/plugins/postimage
/plugins/preidentify
/plugins/prepartition
/plugins/preimage
/plugins/prefinalize

· initialize	Stage one creates writable directories and loads any kernel modules.
· identify	Stage two uses DHCP to get the IP address and host name.
· partition	Stage three creates partitions and file systems.
· image	Stage four downloads and extracts the payload.
· finalize	Stage five configures Management Center services to run with the host name retrieved from DHCP.

Note

All plug-ins must be added inside the RAM Disk under */plugins/<filename>*.

The provisioning plug-in scripts run each time the node is booted with an image that contained that plug-in at the time it was provisioned. This is not dependent on whether or not a new payload is being downloaded or similar situations. Plug-in scripts should be written in such a way that running it multiple times against the same installed payload will not cause problems.

To Add a Plug-in

The following example depicts how to run a script during the boot process.

1. Write a shell or Perl script to run during the boot process. For example, to run a script immediately after partitioning a drive, name the script *postpartition* and add it to the plugins directory in the RAMdisk (i.e., */plugins/<filename>*).

Note

You must add all necessary utilities for your plug-in script to the RAM Disk. For example, if you use a Perl script as a plug-in, you must add the Perl binary and all necessary shared libraries and modules to the RAM Disk. The shared libraries for a utility may be determined using the ldd(1) command. Please note that adding these items significantly increases the size of the RAM Disk. See *To Add Boot Utilities* on page 134.

2. Right-click on an image in the imaging navigation tree and select Edit. The image panel appears.

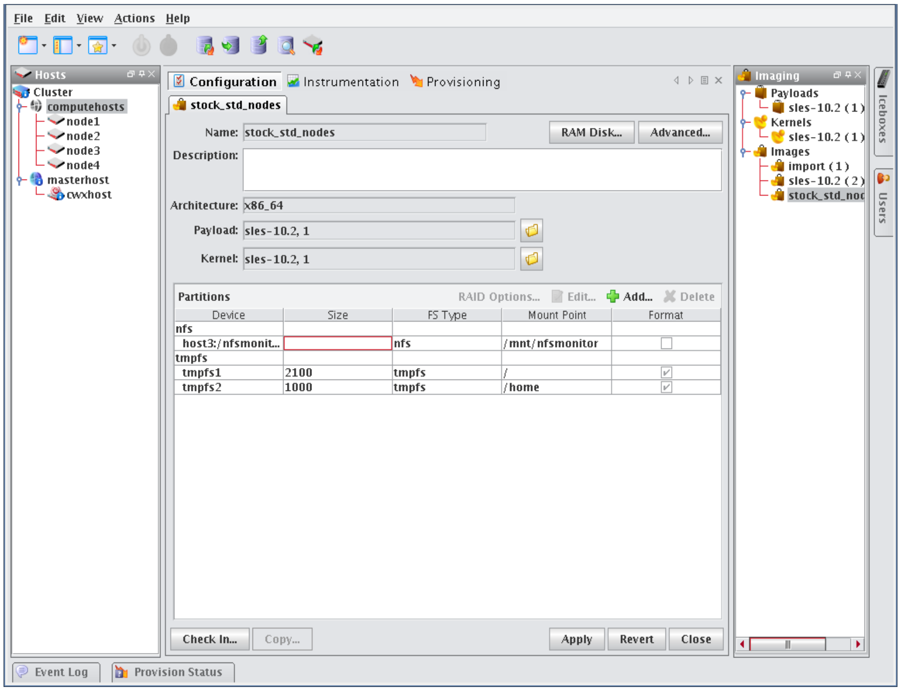
File
Edit
View
Actions
Help
Hosts
Cluster
computehosts
node1
node2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
stock_std_nodes
RAM Disk...
Advanced...
Partitions
RAID Options...
Edit...
Add...
Delete
Check In...
Copy...
Apply
Revert
Close
Imaging
Payloads
sles-10.2 (1)
Kernels
sles-10.2 (1)
Images
import (1)
sles-10.2 (2)
stock_std_noc
Iceboxes
Users
Event Log
Provision Status

Name: stock_std_nodes
Description:
Architecture: x86_64
Payload: sles-10.2, 1
Kernel: sles-10.2, 1

Device	Size	FS Type	Mount Point	Format
nfs				
host3:/nfsmonit...		nfs	/mnt/nfsmonitor	☐
tmpfs				
tmpfs1	2100	tmpfs	/	☑
tmpfs2	1000	tmpfs	/home	☑

3. Click the RAM Disk button. The RAM Disk dialog appears.

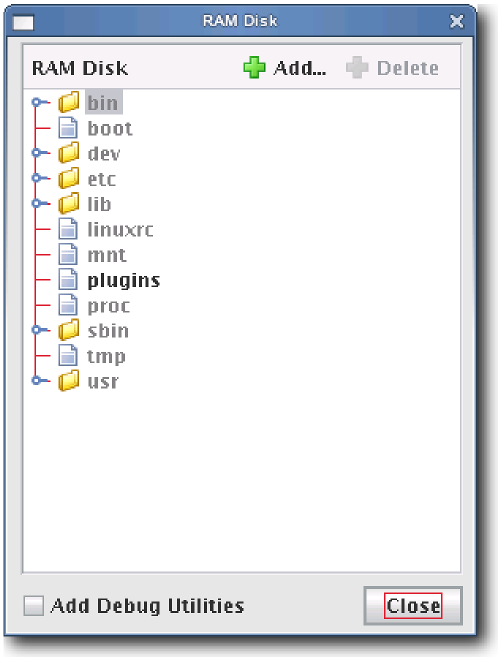
RAM Disk
RAM Disk
Add...
Delete
bin
boot
dev
etc
lib
linuxrc
mnt
plugins
proc
sbin
tmp
usr
Add Debug Utilities
Close

4. Click Add. The Add File To RAM Disk dialog appears.

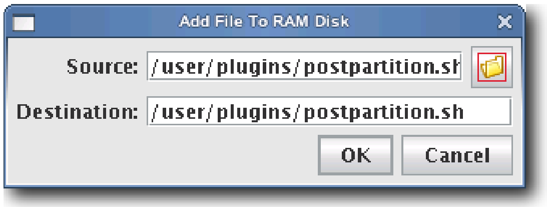
Add File To RAM Disk
OK
Cancel

Source: /user/plugins/postpartition.sh
Destination: /user/plugins/postpartition.sh

5. Enter the boot utility path in the Source field or click Browse to locate a plug-in.
6. Specify the installation location in the Destination field.

Note

All scripts must be installed in the */plugins/* directory. However, you can overwrite other utilities.

7. Click OK to install the utility or click Cancel to abort this action. The new plugins appear in the RAM Disk dialog.

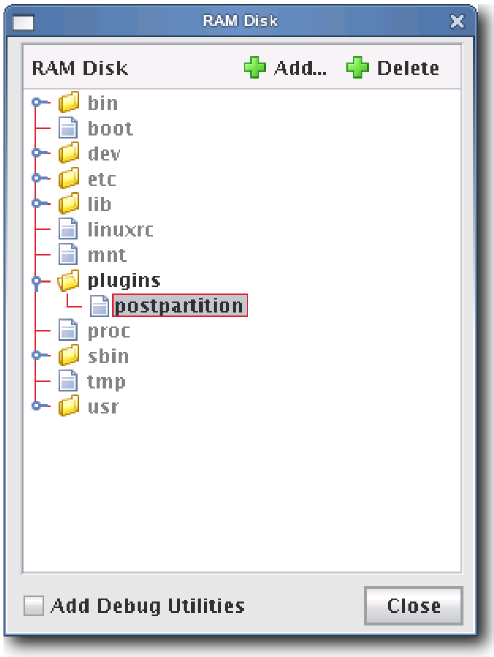
RAM Disk
RAM Disk
Add...
Delete
bin
boot
dev
etc
lib
linuxrc
mnt
plugins
postpartition
proc
sbin
tmp
usr
Add Debug Utilities
Close

8. (Optional) Select Add Debug Utilities to apply additional debugging utilities to the RAM Disk.
9. Click Close.
10. Click Apply to complete the process. Click Revert or Close to abort this action.

Note

Management Center generates the file, *boot.profile*, each time you save an image. See *boot.profile* on page 118 for a description of the information contained in this file.

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). The Version Control System is accessed via the VCS menu and supports common Check-Out and Check-In operations. Items are version controlled by the user—when an item is checked out, it can be modified locally and checked back in. For information on initially placing a payload, kernel, or image under version control, see *Payload Management* on page 82, *Kernel Management* on page 107, or *Image Management* on page 116.

Tip

You can also use VCS Management to copy a payload, kernel, or image and create a new version. See *VCS Management* on page 143.

Version Control

The following diagram illustrates version control for a kernel. The process begins with a working copy of a kernel that is checked into VCS as a versioned kernel. The kernel is then checked out of VCS, modified (as a working copy of the kernel), and checked back into VCS as a new version of the original kernel.

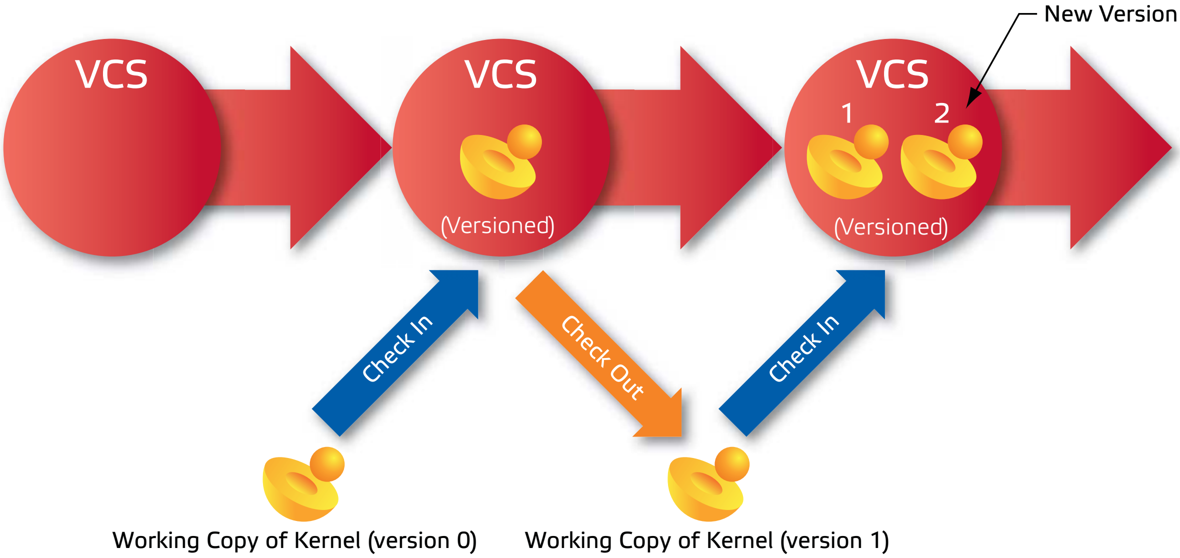

Tip

If another user checks out a copy of the same item you are working with and checks it back into VCS before you do, you must either discard your changes and check out the latest version of the item or create a new branch that does not contain the items checked in by the other user.

A *Working Copy* of a payload, kernel, or image is currently present in the working area (e.g. *$MGR_HOME/imaging/<user>/payloads*). A *Versioned* payload, kernel, or image is a revision of a payload, kernel, or image stored in VCS. Management Center displays payloads, kernels, or images that are currently checked out of VCS in the imaging tree. These items may be edited only while they are checked out, but you may check them into VCS to store your changes. If you are not using a working copy of an item (e.g., it is checked into VCS), you can delete it to conserve space.

Version Branching

Image management works with VCS to allow you to branch any payload, kernel, or image under version control arbitrarily from any version. Suppose, for example, that a payload under version control was gradually optimized to suit specific hardware contained in a cluster. If the optimization were performed in stages (where each stage was a different VCS revision), VCS would contain multiple versions of the payload.

Now suppose that you added some new hosts with slightly different hardware specifications to the cluster, but the last few revisions of the payload use optimizations that are incompatible with the new hardware. Using the version branching feature, you could create a new branch of the payload based on an older version that does not contain the offending optimizations. The new branch could be used with the new hosts, while the remaining hosts could use the original payload.

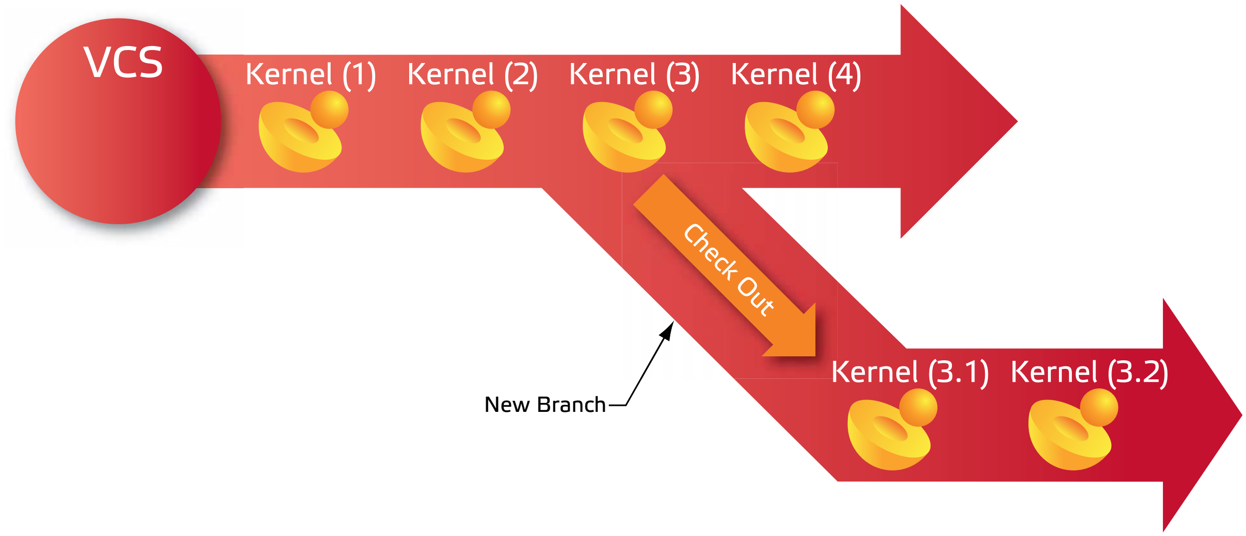

Version Control Check-in

To Check In a Payload, Kernel, or Image

1. After making changes to a payload, kernel, or image, click Check In or select Check In from the VCS menu. The VCS Import dialog appears.

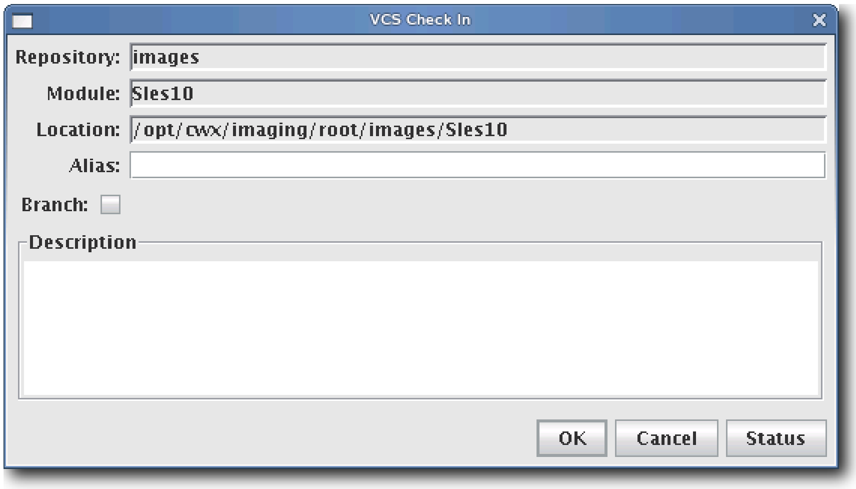

Repository: images
Module: Sles10
Location: /opt/cwx/imaging/root/images/Sles10
Alias:
Branch: ☐

2. (Optional) Enter an alias to use when referring to this version. The alias is the name displayed in the VCS Log between the parentheses:

   1(<*Alias*>)
   February 26, 2004 9:14:17 AM MST, root
   Description of changes...

3. (Optional) Select Branch to create a new branch of this item. Do not select this option if you want Management Center to create a new revision on the current branch.

Note

If another user checks out a copy of the same item you are working with and checks it back into VCS before you do, you must either discard your changes and check out the latest version of the item or create a new branch that does not contain the items checked in by the other user.

4. Click OK to continue or click Cancel to abort this action.

Tip

VCS Check In may fail if you have insufficient disk space. To monitor the amount of available disk space, configure the disk space monitor to log this information, e-mail the administrator, or run a script when disk space is low. See *Management Center Monitoring and Event Subsystem* on page 172 for details.

Version Control Check-out

To Check Out a Payload, Kernel, or Image

1. Select Check Out from the VCS option in the Actions menu. The VCS Check Out dialog appears.

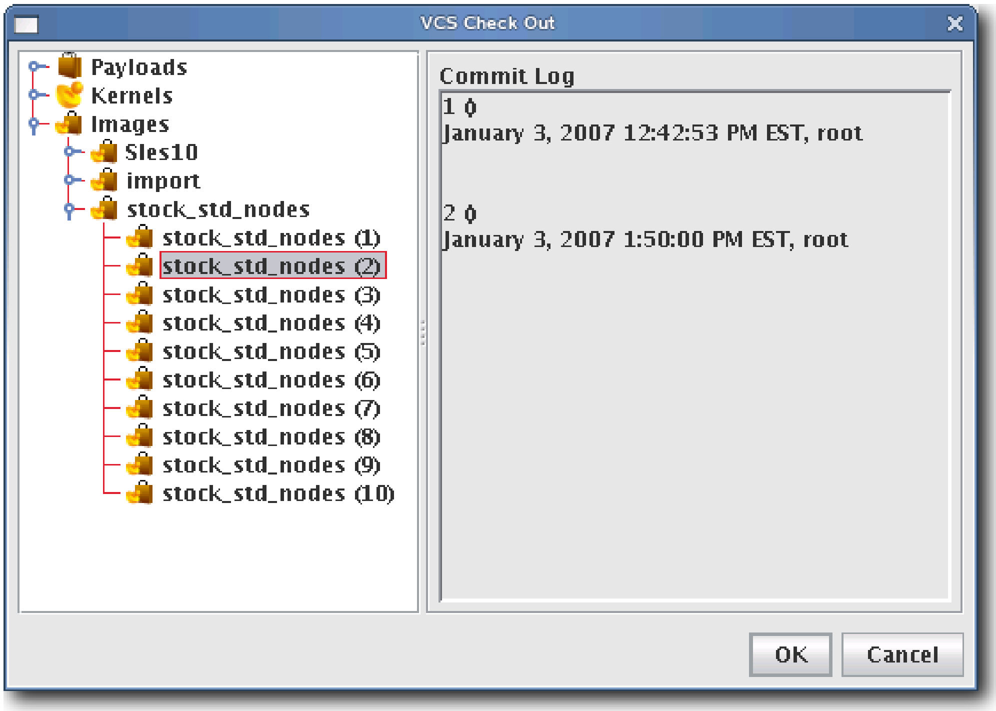
VCS Check Out
Payloads
Kernels
Images
Sles10
import
stock_std_nodes
stock_std_nodes (1)
stock_std_nodes (2)
stock_std_nodes (3)
stock_std_nodes (4)
stock_std_nodes (5)
stock_std_nodes (6)
stock_std_nodes (7)
stock_std_nodes (8)
stock_std_nodes (9)
stock_std_nodes (10)
Commit Log
OK
Cancel

```
1 0
January 3, 2007 12:42:53 PM EST, root


2 0
January 3, 2007 1:50:00 PM EST, root
```

2. Select the payload, kernel, or image you want to check out of VCS (use the Shift or Ctrl keys to select multiple items).

Note

When you check out a payload, kernel, or image, Management Center creates a working copy of the item. If you check out the root of a payload, kernel, or image, Management Center selects the tip revision.

! Warning!

Every time a user creates a payload (or checks a payload out of VCS), Management Center stores a working copy of the payload in the user's *$MGR_HOME/imaging* directory. To accommodate this process, Management Center requires a minimum of 10 GB of disk space. Once the payload is checked into VCS, the user may safely remove the contents of the imaging directory.

3. Click OK. Management Center places the item(s) into a working directory where you may make changes. Click Cancel to abort this action.

VCS Management

The VCS management console allows you to copy, delete, or view the change history for a particular package, kernel, or image.

To Launch the VCS Management Console

1. Select Manage from the VCS option in the Actions menu. The VCS Management dialog appears.
2. Select a payload, kernel, or image for which to display a change history.

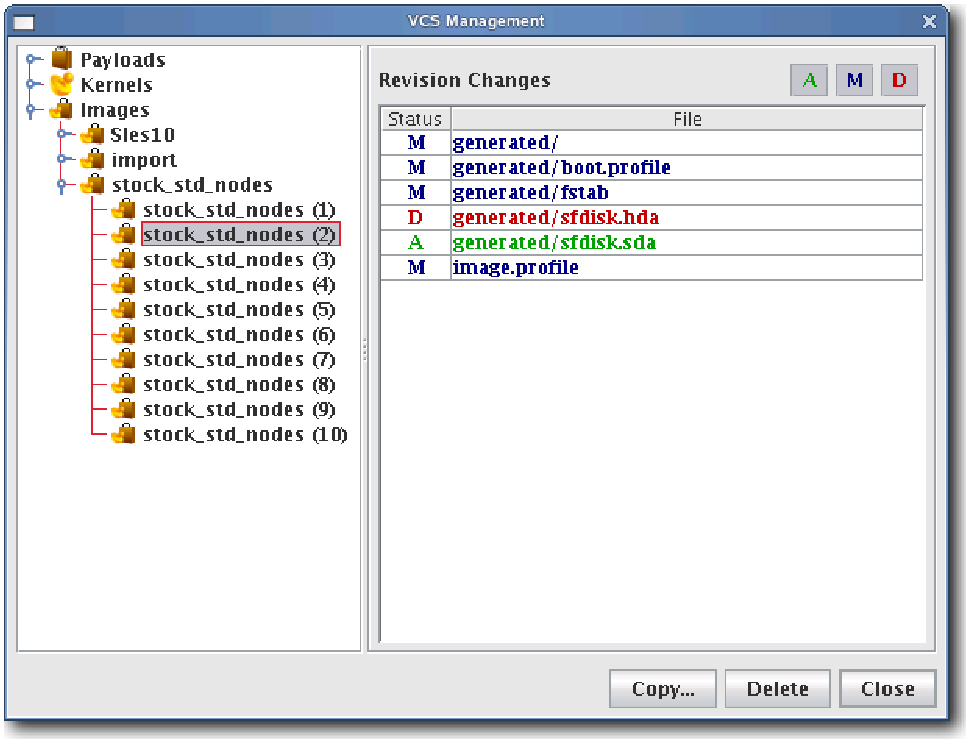
VCS Management
Payloads
Kernels
Images
Sles10
import
stock_std_nodes
stock_std_nodes (1)
stock_std_nodes (2)
stock_std_nodes (3)
stock_std_nodes (4)
stock_std_nodes (5)
stock_std_nodes (6)
stock_std_nodes (7)
stock_std_nodes (8)
stock_std_nodes (9)
stock_std_nodes (10)
Revision Changes
A
M
D
Copy...
Delete
Close

Status	File
M	generated/
M	generated/boot.profile
M	generated/fstab
D	generated/sfdisk.hda
A	generated/sfdisk.sda
M	image.profile

Tip

Click the Add (A), Modify (M), or Delete (D) options to include or exclude specific information.

3. To remove a versioned payload, kernel, or image from VCS, select the item from the navigation tree and click Delete. When deleting a version of any item, all subsequent versions are also deleted (i.e., deleting version 4 also removes versions 5, 6, and so on).

! Warning!

If you select Payloads, Kernels, or Images from the navigation tree, clicking Delete will remove *ALL* payloads, kernels, or images from the system.

4. To copy a payload, kernel, or image, right-click on the item in the navigation tree and select Copy. Management Center prompts you for a new name, then creates a new copy of the item in VCS.

VCS Host Compare

The Host Compare feature allows you to compare the payload currently installed on a host with the latest version of the payload stored in VCS. This is useful when determining whether or not to re-provision a host with a new payload. Similar to the VCS Management Console, this option displays all additions, modifications, and deletions made to the payload since you last used it to provision the host.

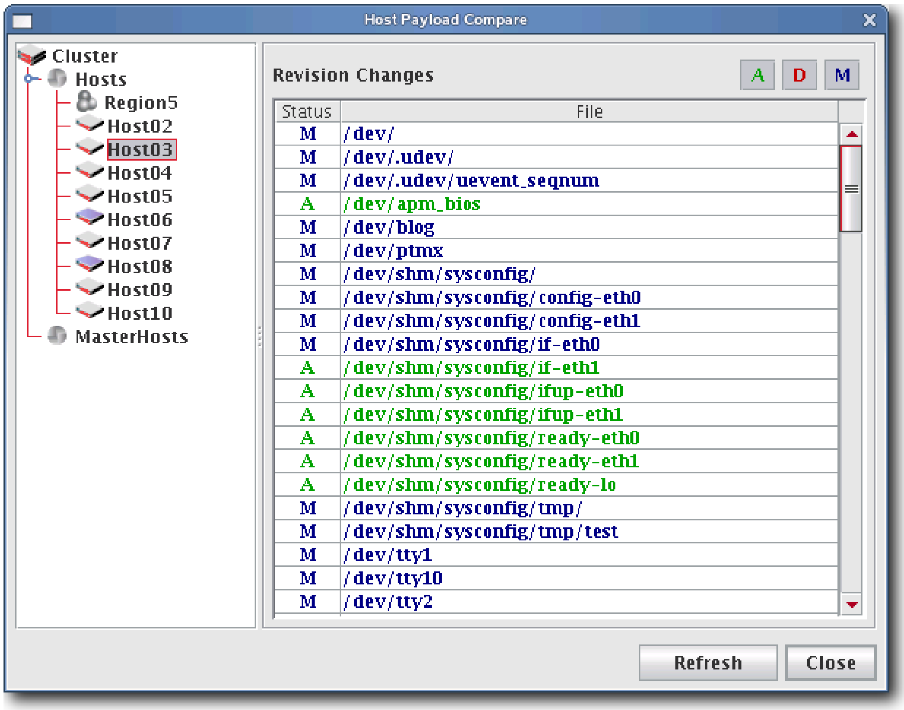
Host Payload Compare
Cluster
Hosts
Region5
Host02
Host03
Host04
Host05
Host06
Host07
Host08
Host09
Host10
MasterHosts
Revision Changes
A
D
M
Refresh
Close

Status	File
M	/dev/
M	/dev/.udev/
M	/dev/.udev/uevent_seqnum
A	/dev/apm_bios
M	/dev/blog
M	/dev/ptmx
M	/dev/shm/sysconfig/
M	/dev/shm/sysconfig/config-eth0
M	/dev/shm/sysconfig/config-eth1
M	/dev/shm/sysconfig/if-eth0
A	/dev/shm/sysconfig/if-eth1
A	/dev/shm/sysconfig/ifup-eth0
A	/dev/shm/sysconfig/ifup-eth1
A	/dev/shm/sysconfig/ready-eth0
A	/dev/shm/sysconfig/ready-eth1
A	/dev/shm/sysconfig/ready-lo
M	/dev/shm/sysconfig/tmp/
M	/dev/shm/sysconfig/tmp/test
M	/dev/tty1
M	/dev/tty10
M	/dev/tty2

To Exclude Files from the Comparison List

1. Open the file, *$MGR_HOME/etc/exclude.files* (a copy of this file should exist on all hosts):

```
proc
dev/pts
etc/ssh/ssh_host_dsa_key
etc/ssh/ssh_host_dsa_key.pub
etc/ssh/ssh_host_key
etc/ssh/ssh_host_key.pub
etc/ssh/ssh_host_rsa_key
etc/ssh/ssh_host_rsa_key.pub
media
mnt
root/.ssh
scratch
sys
tmp
usr/local/src
usr/share/doc
usr/src
var/cache/
var/lock
```

```
var/log
var/run
var/spool/anacron
var/spool/at
var/spool/atjobs
var/spool/atspool
var/spool/clientmqueue
var/spool/cron
var/spool/mail
var/spool/mqueue
var/tmp
```

2. Edit the file as needed, then save your changes.

Tip

It is best to edit this file while it is in the payload so it can be copied to all hosts.

VersionControlService.profile

Management Center uses *VersionControlService.profile*, a global default exclude list that is not distribution-specific. You may add files or directories to this list to prevent Management Center from checking them into VCS—particularly helpful when importing payloads from the working directory. To remove items from the exclusion list, comment them out of the profile.

Also contained in the *VersionControlService.profile*, the deflate.temp:/<*dir*> parameter allows you to specify an alternate path for large files created while importing a payload.

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.

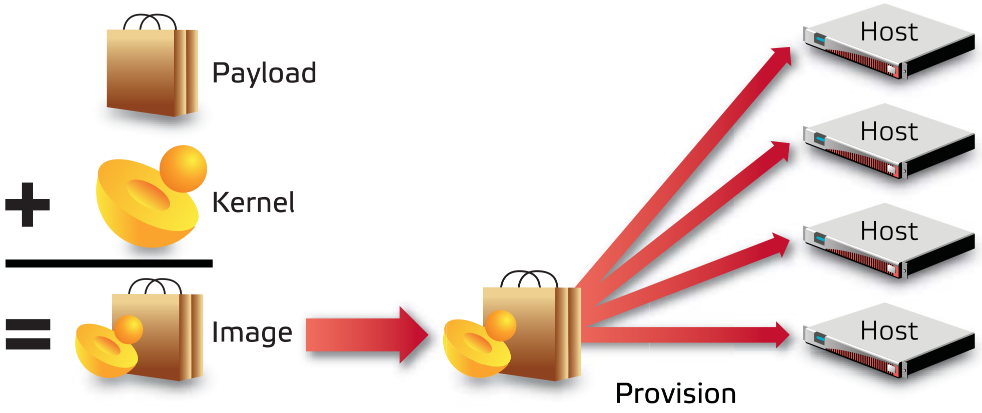

Select an Image and Provision

To Select an Image and Provision

1. Select the host(s) 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 popup menu and allows you to select the image you want to use to provision. If you have made only minor changes to an image and want to upgrade your hosts to use the new image, see *VCS Upgrade* on page 150.

2. Select the Provisioning tab.

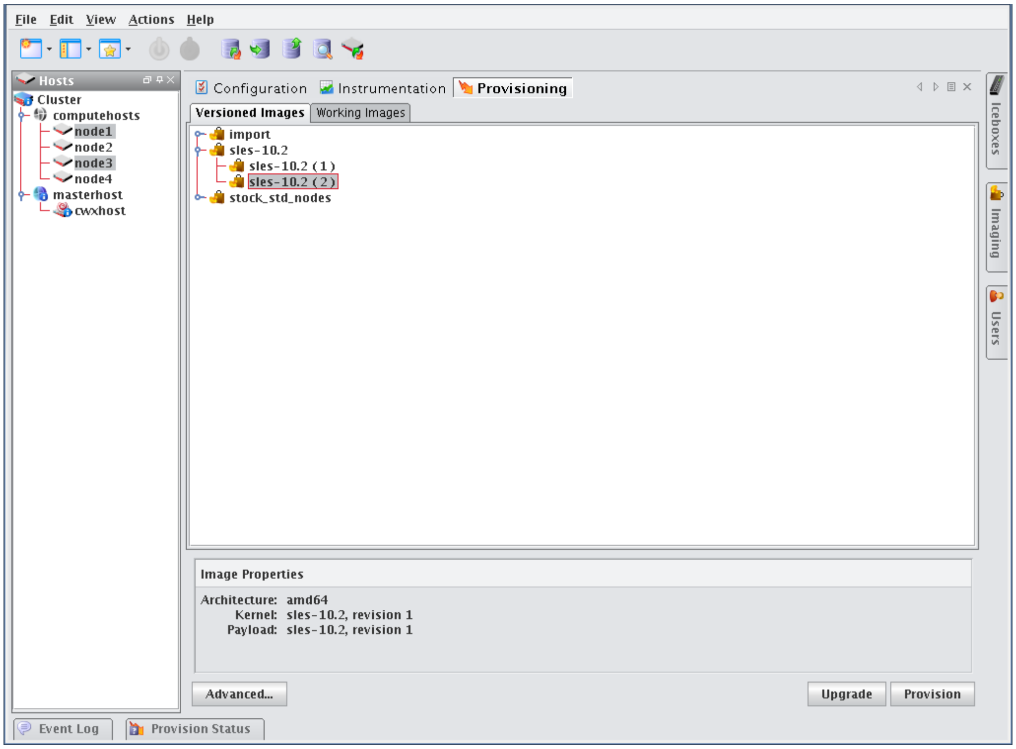
File Edit View Actions Help
Hosts
Cluster
computehosts
node1
node2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
Versioned Images
Working Images
import
sles-10.2
sles-10.2 (1)
sles-10.2 (2)
stock_std_nodes
Iceboxes
Imaging
Users
Advanced...
Upgrade
Provision
Event Log
Provision Status

Image Properties

Architecture: amd64
Kernel: sles-10.2, revision 1
Payload: sles-10.2, revision 1

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 (e.g., *$MGR_HOME/imaging/<user>/images*). This allows you to test changes prior to checking in. See *Version Control System (VCS)* on page 140 for details on using the version control system.

A *Working Copy* of an image is currently present in the working area (e.g., *$MGR_HOME/imaging/<user>/images*). A *Versioned* image is a revision of an image stored in VCS. See *Version Control System (VCS)* on page 140 for details on using the version control system.

4. Select the image you want to use to provision the host(s).
5. (Optional) Click the Advanced button to display the Advanced Options dialog (see *Advanced Provisioning Options* on page 151). 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 host(s) or click No to abort this action.

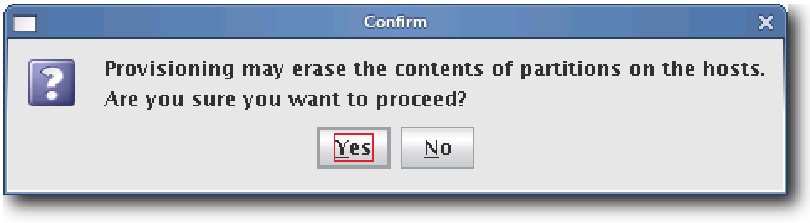

Provisioning may erase the contents of partitions on the hosts.
Are you sure you want to proceed?

! Warning!

When you click Yes, Management Center re-provisions the hosts using the new image. Any pending or running jobs on the selected host(s) are lost.

Tip

To disable the provisioning confirmation dialog, see *Provisioning* on page 40.

Right-click Provisioning

1. Select the host(s) 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 popup 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.

VCS Upgrade

VCS Upgrade is a quick, easy way to make small changes to hosts. Unlike provisioning (which requires rebooting the host and reformatting its hard drive), the VCS Upgrade feature copies the VCS revision to the host and inflates it while the host is running. Using the upgrade feature requires that you check all changes into the payload, that the payload revision is updated in the image, and that you check in the image.

The update feature will update only those hosts with files managed by the payload and will not affect the running kernel or file system information. If there are changes to the kernel or image, they will not take place until the host is re-provisioned with that image. You cannot “downgrade” a host by using an older version of a payload.

! Warning!

Major changes made to hosts should be done using provisioning. This ensures that all hosts are homogenous and takes full advantage of multicast. Also, VCS Upgrade leaves the image and payload on the host out of sync from what is available in the VCS repository—for this reason, SGI recommends that you use *Advanced Provisioning Options* on page 151 to schedule the hosts to be re-provisioned with the selected image the next time they reboot.

To Upgrade a Host(s)

1. Select the Provisioning tab.
2. Open the Versioned Images tab and select the image you want to use to upgrade the host(s)

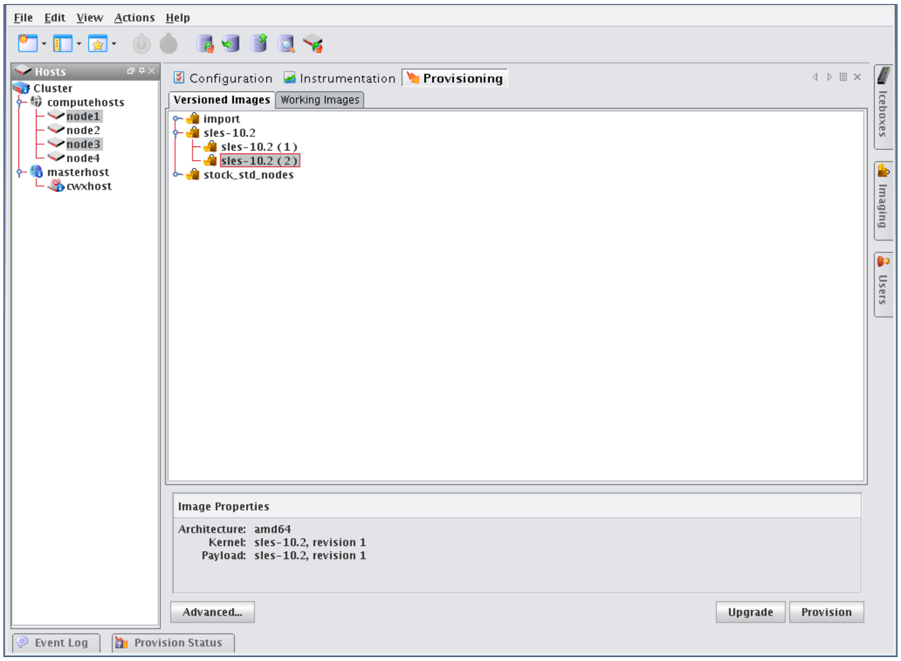
File
Edit
View
Actions
Help
Hosts
Cluster
computehosts
node1
node2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
Versioned Images
Working Images
import
sles-10.2
sles-10.2 (1)
sles-10.2 (2)
stock_std_nodes
Iceboxes
Imaging
Users
Advanced...
Upgrade
Provision
Event Log
Provision Status

Image Properties

Architecture: amd64
Kernel: sles-10.2, revision 1
Payload: sles-10.2, revision 1

3. Select the host(s) you want to upgrade from the navigation tree (use the Shift or Ctrl keys to select multiple hosts).
4. Click Update to update the image to the selected hosts. As the operation begins, a status dialog appears.

Advanced Provisioning Options

The Advanced Options dialog allows you to temporarily modify partitioning behavior, payload download settings, and Kernel verbosity. These settings are not persistent, they simply override those configurations made using the Advanced Image Options dialog. See *Advanced Imaging Options* on page 117.

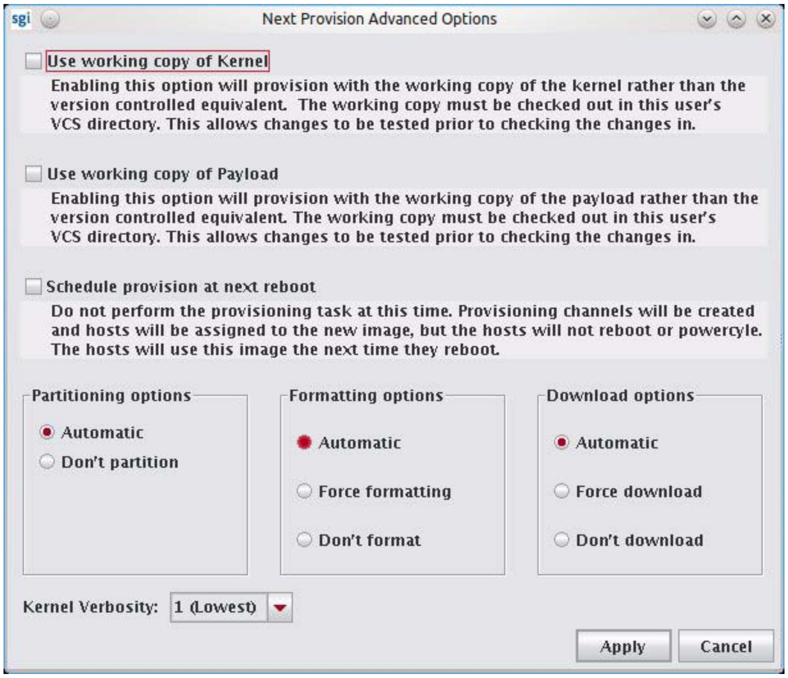

☐ Use working copy of Kernel

Enabling this option will provision with the working copy of the kernel rather than the version controlled equivalent. The working copy must be checked out in this user's VCS directory. This allows changes to be tested prior to checking the changes in.

☐ Use working copy of Payload

Enabling this option will provision with the working copy of the payload rather than the version controlled equivalent. The working copy must be checked out in this user's VCS directory. This allows changes to be tested prior to checking the changes in.

☐ Schedule provision at next reboot

Do not perform the provisioning task at this time. Provisioning channels will be created and hosts will be assigned to the new image, but the hosts will not reboot or powercyle. The hosts will use this image the next time they reboot.

Partitioning options
● Automatic
○ Don't partition

Formatting options
● Automatic
○ Force formatting
○ Don't format

Download options
● Automatic
○ Force download
○ Don't download

Kernel Verbosity: 1 (Lowest)

USE WORKING COPY OF KERNEL

Enable this option to use the working copy of the kernel in place of its version-controlled equivalent. This allows you to test your changes prior to checking them in.

USE WORKING COPY OF PAYLOAD

Enable this option to use the working copy of the payload in place of its version-controlled equivalent. Because working copies of payloads are often shared, hosts associated with the working copy are updated to use the latest version when they reboot—but only if the payload was modified or used to provision other hosts.

SCHEDULE PROVISION AT NEXT REBOOT

Enable this option to postpone provisioning until the next time you reboot the hosts. Provisioning channels are created and hosts are assigned to the new image, but the hosts cannot reboot or cycle power without being provisioned.

Tip

To change the default scheduled provisioning setting, see *Provisioning* on page 40.

Scheduling a provision at next reboot can be especially useful when used with PBS. For example, you may make updates to a payload, then schedule provisioning to occur only after the current tasks are complete. To do this, the root user (who must be allowed to submit jobs) can submit a job to each host instructing it to reboot.

The root user can submit jobs to PBS only if *acl_roots* is configured. To configure acl_roots, run *qmgr* and enter the following from the qmgr prompt:

```
qmgr: set server acl_roots += root
```

If you already set up additional ACLs, you will also need to add root to those ACLs. For example, suppose you have an acl_users list that allows access to a queue, *workq*. The command to add root to the ACL would be:

```
# set queue workq acl_users += root
```

The following is a sample PBS script you might use to reboot hosts:

```
################################################
#!/bin/bash
for i in `seq 1 64`
do
        echo \#PBS -N Reboot_n$i > Reboot_n$i.pbs
        echo \#PBS -joe >> Reboot_n$i.pbs
        echo \#PBS -V >> Reboot_n$i.pbs
        echo \#PBS -l nodes=n$i >> Reboot_n$i.pbs
        echo \#PBS -q workq >> Reboot_n$i.pbs
        echo \#PBS -o /dev/null >> Reboot_n$i.pbs
        echo \/sbin\/reboot >> Reboot_n$i.pbs
        echo done >> Reboot_n$i.pbs

        qsub < Reboot_n$i.pbs

        rm Reboot_n$i.pbs
done
################################################
```

PARTITIONING OPTIONS

This option allows you override the current partition settings. You can automatically partition an iamge if the partition changes or choose not to re-partition drives.

FORMATTING OPTIONS

You can automatically format partitions if the payload or partitioning scheme changes, force formatting of all partitions—including those that are exempt from being overwritten (see *Partitions* on page 61), or choose not to format.

PAYLOAD DOWNLOAD OPTIONS

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), force Management Center to download a new copy of the image—regardless of the image status, or choose not to download a payload.

KERNEL VERBOSITY

The kernel verbosity level (1-8) allows you to control debug messages displayed by the kernel during provisioning. The default value, 1, is the least verbose and 8 is the most.

Chapter 7

Instrumentation and Events

Instrumentation

The Management Center instrumentation service provides the ability to monitor system health and activity for every host in the cluster. Hosts may be monitored collectively to provide a general system overview, or individually to allow you to view the configuration of a particular host (useful when diagnosing problems with a particular host or configuration). From the Instrumentation tab, you can view statistical data for the following areas:

- Overview
- Thumbnail
- List
- CPU
- Memory
- Disk
- Network
- Kernel
- Load
- Environmental
- Environmental List
- GPU
- Power

Note

When monitoring the Management Center Master Host, the name of the Master Host must match the name assigned in *$MGR_HOME/@genesis.profile*.

Tip

When using the Management Center client by exporting an X session over an SSH connection, enabling the gradient fill and anti-aliasing options for instrumentation may adversely affect the performance of the GUI. This is common on slower systems. To improve system performance, disable the Gradient Fill and Anti-Aliasing options under the View menu. For best performance, install a Management Center Client.

States

Management Center uses the following icons to provide visual cues about system status. These icons appear next to each host viewed with the instrumentation service or from the navigation tree. Similar icons appear next to clusters, partitions, and regions to indicate the status of hosts contained therein.

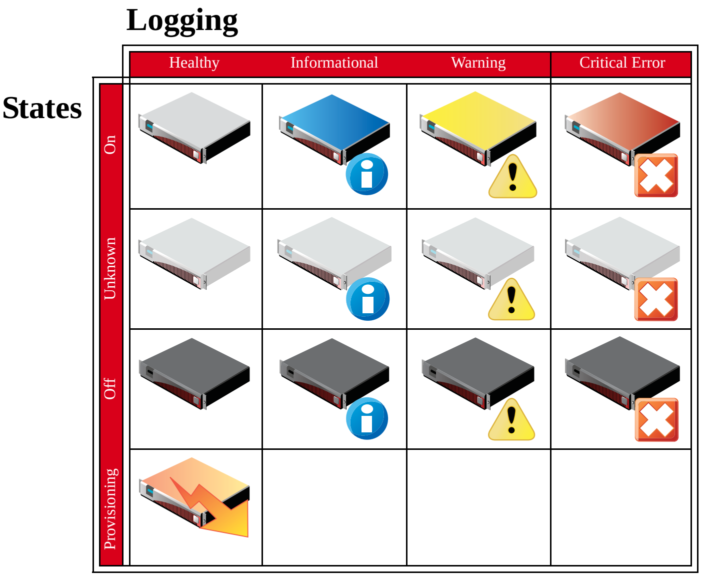

Event Log

Management Center also tracks events logged for each host in the cluster. The Management Center event log is located on the instrumentation overview screen. If you select multiple hosts (or a container such as a cluster, partition, or region), the log shows messages for any host in the selection. If you select a single host, the event log shows messages for this host only. Events have three severity levels: error, warning, and information. For additional details on instrumentation event monitoring, see *Management Center Monitoring and Event Subsystem* on page 172.

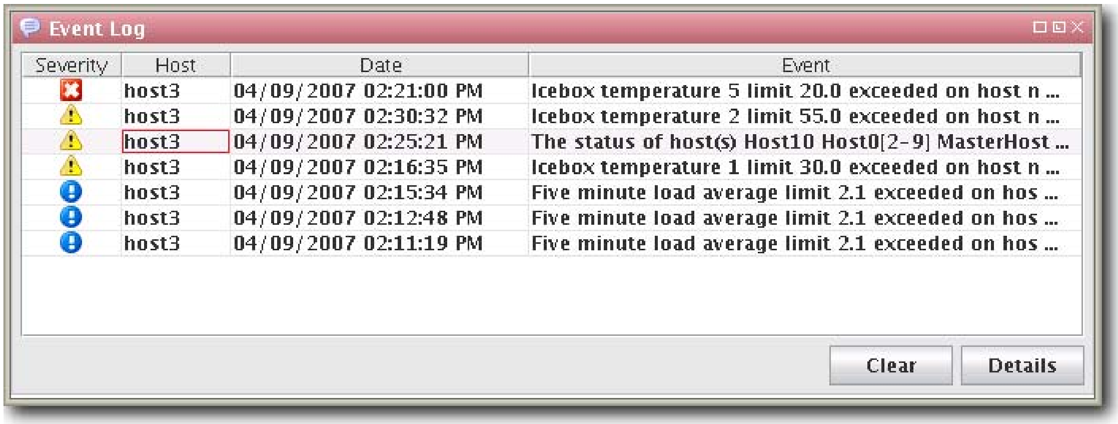

Severity	Host	Date	Event
☒	host3	04/09/2007 02:21:00 PM	Icebox temperature 5 limit 20.0 exceeded on host n ...
⚠	host3	04/09/2007 02:30:32 PM	Icebox temperature 2 limit 55.0 exceeded on host n ...
⚠	host3	04/09/2007 02:25:21 PM	The status of host(s) Host10 Host0[2-9] MasterHost ...
⚠	host3	04/09/2007 02:16:35 PM	Icebox temperature 1 limit 30.0 exceeded on host n ...
ℹ	host3	04/09/2007 02:15:34 PM	Five minute load average limit 2.1 exceeded on hos ...
ℹ	host3	04/09/2007 02:12:48 PM	Five minute load average limit 2.1 exceeded on hos ...
ℹ	host3	04/09/2007 02:11:19 PM	Five minute load average limit 2.1 exceeded on hos ...

Menu Controls

The output for the instrumentation service is easily configured and displayed using menu controls located in the View menu.

View Menu

Metrics Select and display custom metrics defined for your system—this option is not available to all tab views. See *Metrics* on page 185 for information on defining metrics.

Interval Set the frequency (in seconds) with which to gather and display data—10, 5, or 1.

Layout Arrange how the instrumentation panel displays information.

Filter List hosts that are in specific states (Thumbnail tab only).

Size Change the display size of thumbnails (Small, Medium, Large).

Sort Organize and display statistical data according to the name or state of the host(s).

Temperatures Select the format in which to display temperatures (Celsius, Fahrenheit).

Anti-Aliasing Apply smoothing to line graphs.

Gradient Fill Apply fill colors to line graphs.

Overview Tab

The Overview tab provides details about the configuration, power status, resource utilization, and health status of the host(s) selected in the host navigation tree. Selecting a Cluster, Partition, or Region in the tree displays all hosts contained in it. See *States* on page 154 for a list of system health indicators and *Event Log* on page 154 for information regarding messages generated by the host(s).

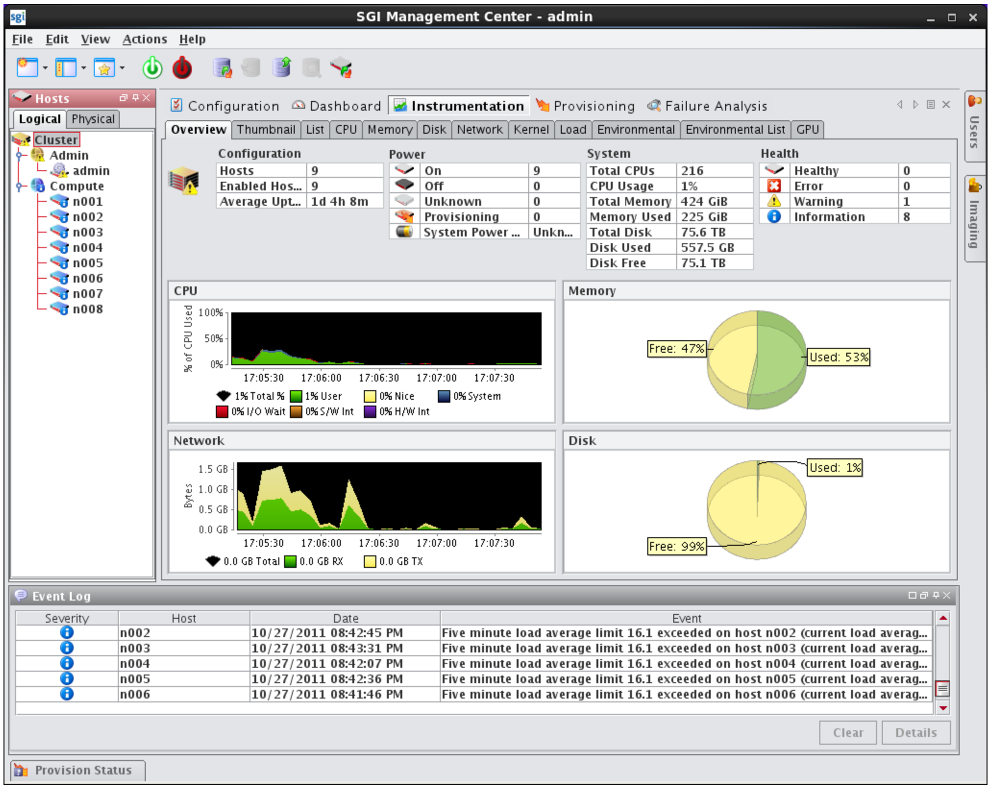
SGI Management Center - admin
File Edit View Actions Help
Hosts
Logical Physical
Cluster
Admin
admin
Compute
n001
n002
n003
n004
n005
n006
n007
n008
Configuration Dashboard Instrumentation Provisioning Failure Analysis
Overview Thumbnail List CPU Memory Disk Network Kernel Load Environmental Environmental List GPU
Users
Imaging
CPU
100%
50%
0%
% of CPU Used
17:05:30 17:06:00 17:06:30 17:07:00 17:07:30
1% Total % 1% User 0% Nice 0% System
0% I/O Wait 0% S/W Int 0% H/W Int
Memory
Free: 47%
Used: 53%
Network
1.5 GB
1.0 GB
0.5 GB
0.0 GB
Bytes
17:05:30 17:06:00 17:06:30 17:07:00 17:07:30
0.0 GB Total 0.0 GB RX 0.0 GB TX
Disk
Used: 1%
Free: 99%
Event Log
Clear
Details
Provision Status

Configuration

Hosts	9
Enabled Hos...	9
Average Upt...	1d 4h 8m

Power

On	9
Off	0
Unknown	0
Provisioning	0
System Power ...	Unkn...

System

Total CPUs	216
CPU Usage	1%
Total Memory	424 GiB
Memory Used	225 GiB
Total Disk	75.6 TB
Disk Used	557.5 GB
Disk Free	75.1 TB

Health

Healthy	0
Error	0
Warning	1
Information	8

Severity	Host	Date	Event
ⓘ	n002	10/27/2011 08:42:45 PM	Five minute load average limit 16.1 exceeded on host n002 (current load averag...
ⓘ	n003	10/27/2011 08:43:31 PM	Five minute load average limit 16.1 exceeded on host n003 (current load averag...
ⓘ	n004	10/27/2011 08:42:07 PM	Five minute load average limit 16.1 exceeded on host n004 (current load averag...
ⓘ	n005	10/27/2011 08:42:36 PM	Five minute load average limit 16.1 exceeded on host n005 (current load averag...
ⓘ	n006	10/27/2011 08:41:46 PM	Five minute load average limit 16.1 exceeded on host n006 (current load averag...

Thumbnail Tab

The Thumbnail tab displays a graphical representation of the system health, event log status, CPU usage, memory availability, and disk space. From the View menu, you may filter hosts to display only those in a specific state, resize the thumbnails, or sort the hosts by name or state. See *States* on page 154 for a list of system health indicators.

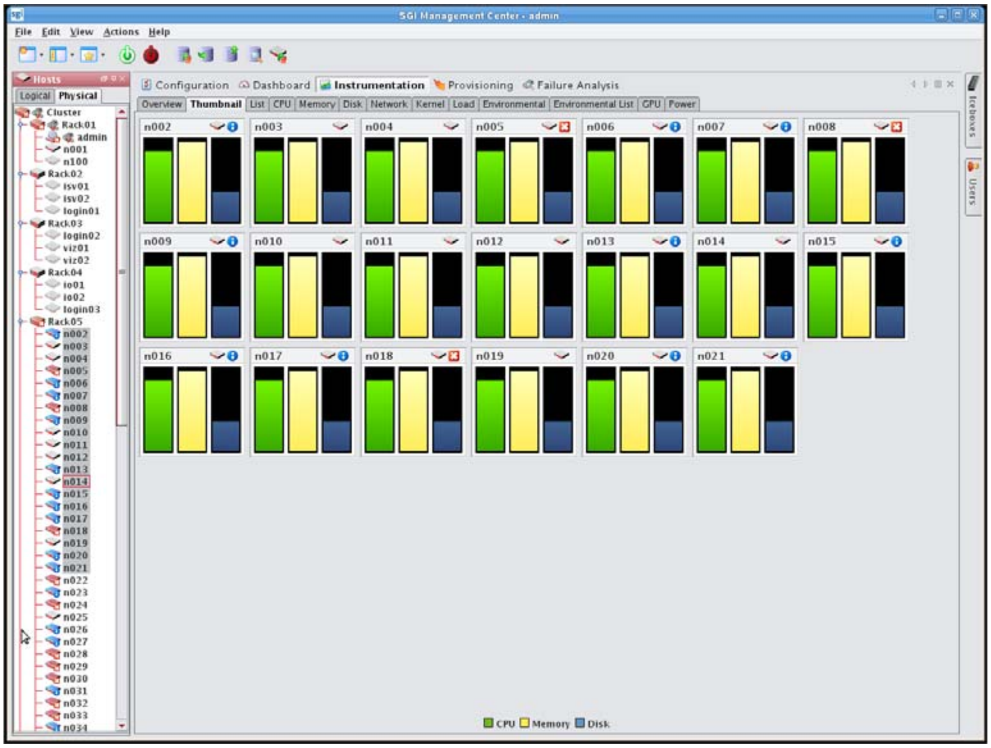
SGI Management Center - admin
File Edit View Actions Help
Hosts
Logical Physical
Cluster
Rack01
admin
n001
n100
Rack02
isv01
isv02
login01
Rack03
login02
viz01
viz02
Rack04
io01
io02
login03
Rack05
n002
n003
n004
n005
n006
n007
n008
n009
n010
n011
n012
n013
n014
n015
n016
n017
n018
n019
n020
n021
n022
n023
n024
n025
n026
n027
n028
n029
n030
n031
n032
n033
n034
Configuration
Dashboard
Instrumentation
Provisioning
Failure Analysis
Overview
Thumbnail
List
CPU
Memory
Disk
Network
Kernel
Load
Environmental
Environmental List
GPU
Power
n002
n003
n004
n005
n006
n007
n008
n009
n010
n011
n012
n013
n014
n015
n016
n017
n018
n019
n020
n021
CPU
Memory
Disk
Iceboxes
Users

List Tab

The List tab displays all pre-configured and custom metrics being observed by the instrumentation service. To add metrics to this list, select Metrics from the View menu. To create new metrics, see *Instrumentation* on page 153.

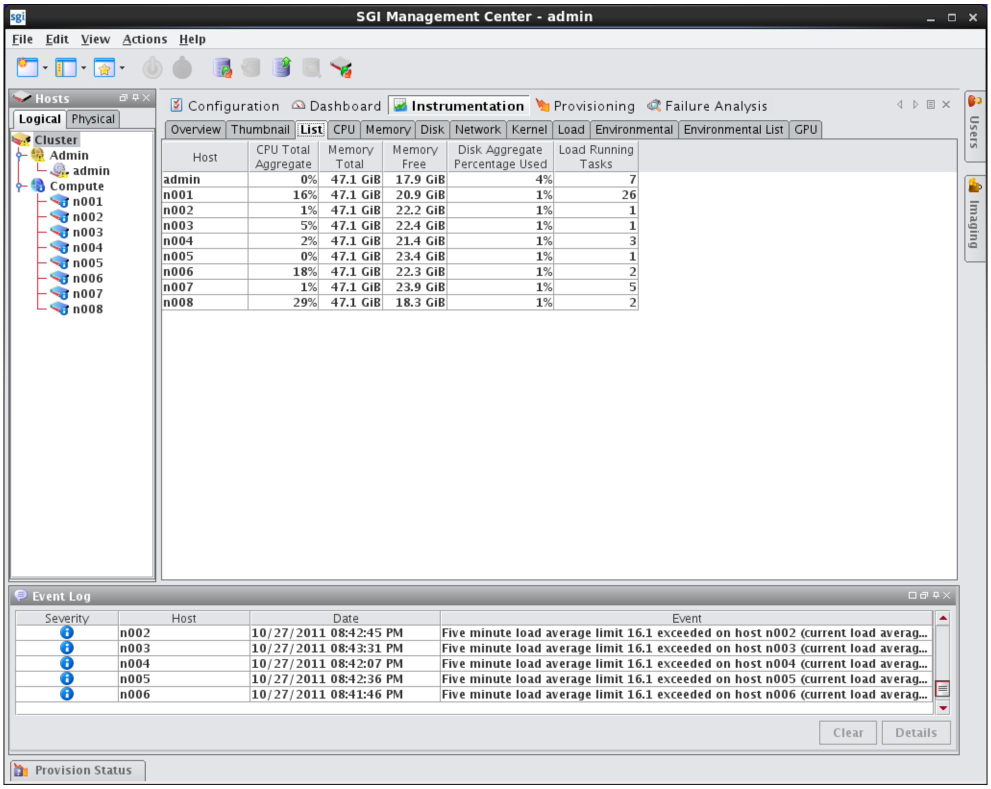
SGI Management Center - admin
File
Edit
View
Actions
Help
Hosts
Logical
Physical
Cluster
Admin
admin
Compute
n001
n002
n003
n004
n005
n006
n007
n008
Configuration
Dashboard
Instrumentation
Provisioning
Failure Analysis
Overview
Thumbnail
List
CPU
Memory
Disk
Network
Kernel
Load
Environmental
Environmental List
GPU
Users
Imaging
Event Log
Clear
Details
Provision Status

Host	CPU Total Aggregate	Memory Total	Memory Free	Disk Aggregate Percentage Used	Load Running Tasks
admin	0%	47.1 GiB	17.9 GiB	4%	7
n001	16%	47.1 GiB	20.9 GiB	1%	26
n002	1%	47.1 GiB	22.2 GiB	1%	1
n003	5%	47.1 GiB	22.4 GiB	1%	1
n004	2%	47.1 GiB	21.4 GiB	1%	3
n005	0%	47.1 GiB	23.4 GiB	1%	1
n006	18%	47.1 GiB	22.3 GiB	1%	2
n007	1%	47.1 GiB	23.9 GiB	1%	5
n008	29%	47.1 GiB	18.3 GiB	1%	2

Severity	Host	Date	Event
ℹ	n002	10/27/2011 08:42:45 PM	Five minute load average limit 16.1 exceeded on host n002 (current load averag...
ℹ	n003	10/27/2011 08:43:31 PM	Five minute load average limit 16.1 exceeded on host n003 (current load averag...
ℹ	n004	10/27/2011 08:42:07 PM	Five minute load average limit 16.1 exceeded on host n004 (current load averag...
ℹ	n005	10/27/2011 08:42:36 PM	Five minute load average limit 16.1 exceeded on host n005 (current load averag...
ℹ	n006	10/27/2011 08:41:46 PM	Five minute load average limit 16.1 exceeded on host n006 (current load averag...

Tip

You may copy and paste the contents of list view tables for use in other applications.

CPU Tab

Select the CPU tab to monitor the CPU utilization for the selected host(s).

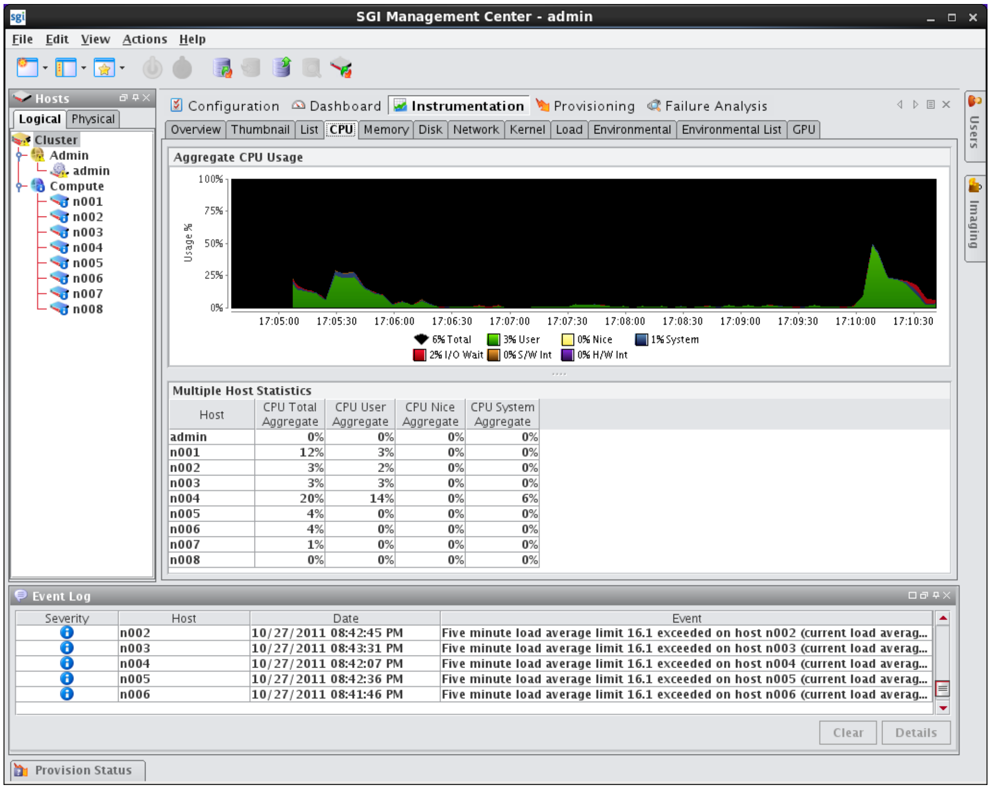
SGI Management Center - admin
File Edit View Actions Help
Hosts
Logical Physical
Cluster
Admin
admin
Compute
n001
n002
n003
n004
n005
n006
n007
n008
Configuration Dashboard Instrumentation Provisioning Failure Analysis
Overview Thumbnail List CPU Memory Disk Network Kernel Load Environmental Environmental List GPU
Users
Imaging
Aggregate CPU Usage
100%
75%
50%
25%
0%
Usage %
17:05:00 17:05:30 17:06:00 17:06:30 17:07:00 17:07:30 17:08:00 17:08:30 17:09:00 17:09:30 17:10:00 17:10:30
6% Total
3% User
0% Nice
1% System
2% I/O Wait
0% S/W Int
0% H/W Int
Event Log
Clear
Details
Provision Status

Multiple Host Statistics

Host	CPU Total Aggregate	CPU User Aggregate	CPU Nice Aggregate	CPU System Aggregate
admin	0%	0%	0%	0%
n001	12%	3%	0%	0%
n002	3%	2%	0%	0%
n003	3%	3%	0%	0%
n004	20%	14%	0%	6%
n005	4%	0%	0%	0%
n006	4%	0%	0%	0%
n007	1%	0%	0%	0%
n008	0%	0%	0%	0%

Event Log

Severity	Host	Date	Event
ℹ	n002	10/27/2011 08:42:45 PM	Five minute load average limit 16.1 exceeded on host n002 (current load averag...
ℹ	n003	10/27/2011 08:43:31 PM	Five minute load average limit 16.1 exceeded on host n003 (current load averag...
ℹ	n004	10/27/2011 08:42:07 PM	Five minute load average limit 16.1 exceeded on host n004 (current load averag...
ℹ	n005	10/27/2011 08:42:36 PM	Five minute load average limit 16.1 exceeded on host n005 (current load averag...
ℹ	n006	10/27/2011 08:41:46 PM	Five minute load average limit 16.1 exceeded on host n006 (current load averag...

Memory Tab

Select the Memory tab to monitor the physical and virtual memory utilization for the selected host(s).

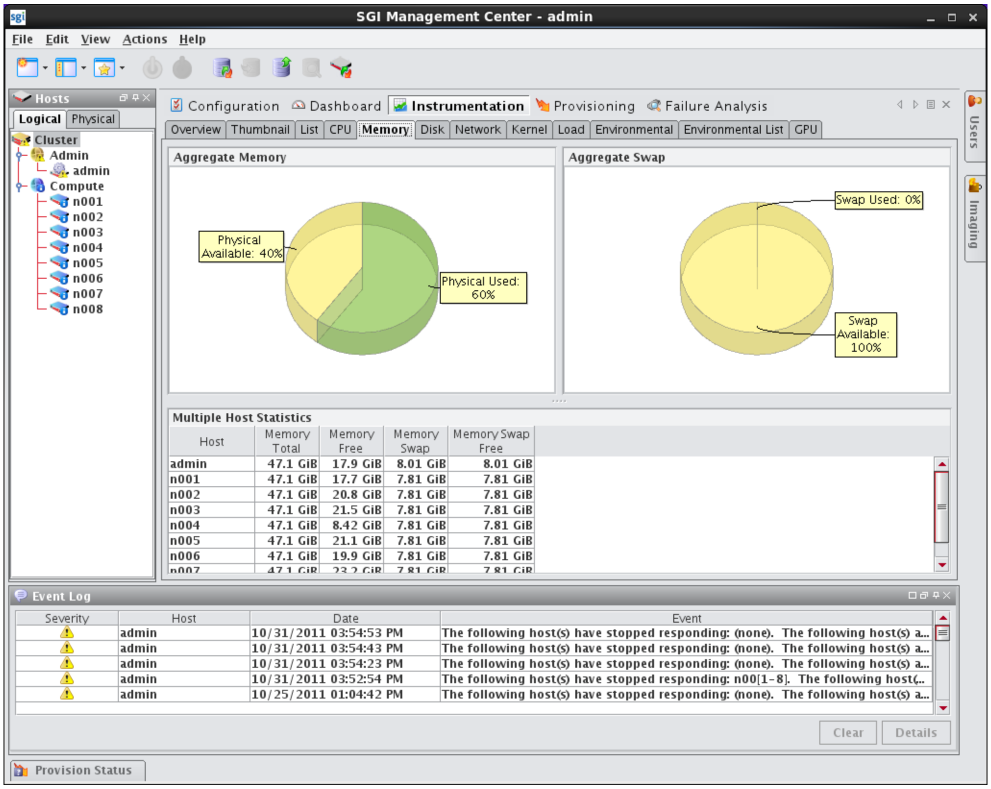
SGI Management Center - admin
File Edit View Actions Help
Hosts
Logical Physical
Cluster
Admin
admin
Compute
n001
n002
n003
n004
n005
n006
n007
n008
Configuration Dashboard Instrumentation Provisioning Failure Analysis
Overview Thumbnail List CPU Memory Disk Network Kernel Load Environmental Environmental List GPU
Aggregate Memory
Physical Available: 40%
Physical Used: 60%
Aggregate Swap
Swap Used: 0%
Swap Available: 100%
Users
Imaging
Clear
Details
Provision Status

Multiple Host Statistics

Host	Memory Total	Memory Free	Memory Swap	Memory Swap Free
admin	47.1 GiB	17.9 GiB	8.01 GiB	8.01 GiB
n001	47.1 GiB	17.7 GiB	7.81 GiB	7.81 GiB
n002	47.1 GiB	20.8 GiB	7.81 GiB	7.81 GiB
n003	47.1 GiB	21.5 GiB	7.81 GiB	7.81 GiB
n004	47.1 GiB	8.42 GiB	7.81 GiB	7.81 GiB
n005	47.1 GiB	21.1 GiB	7.81 GiB	7.81 GiB
n006	47.1 GiB	19.9 GiB	7.81 GiB	7.81 GiB
n007	47.1 GiB	23.2 GiB	7.81 GiB	7.81 GiB

Event Log

Severity	Host	Date	Event
⚠	admin	10/31/2011 03:54:53 PM	The following host(s) have stopped responding: (none). The following host(s) a...
⚠	admin	10/31/2011 03:54:43 PM	The following host(s) have stopped responding: (none). The following host(s) a...
⚠	admin	10/31/2011 03:54:23 PM	The following host(s) have stopped responding: (none). The following host(s) a...
⚠	admin	10/31/2011 03:52:54 PM	The following host(s) have stopped responding: n00[1-8]. The following host(...
⚠	admin	10/25/2011 01:04:42 PM	The following host(s) have stopped responding: (none). The following host(s) a...

Disk Tab

Select the Disk tab to monitor the disk I/O and usage for the selected host(s).

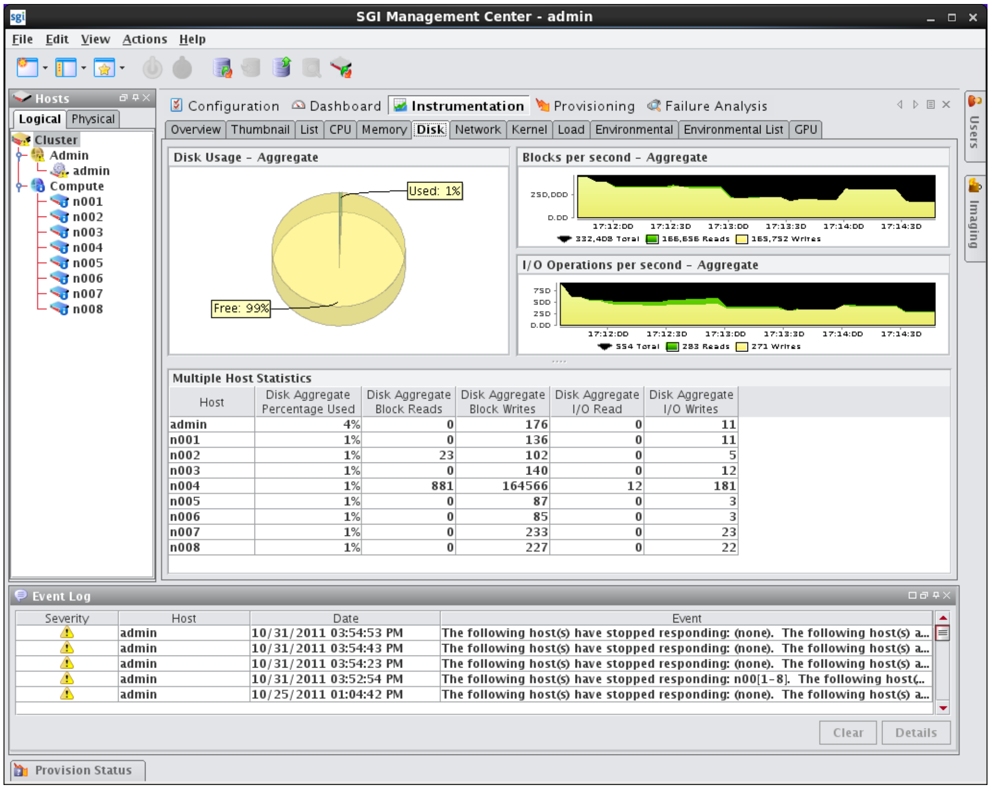
SGI Management Center - admin
File Edit View Actions Help
Hosts
Logical Physical
Cluster
Admin
admin
Compute
n001
n002
n003
n004
n005
n006
n007
n008
Configuration Dashboard Instrumentation Provisioning Failure Analysis
Overview Thumbnail List CPU Memory Disk Network Kernel Load Environmental Environmental List GPU
Disk Usage - Aggregate
Used: 1%
Free: 99%
Blocks per second - Aggregate
250,000
0.00
17:12:00 17:12:30 17:13:00 17:13:30 17:14:00 17:14:30
332,408 Total 166,656 Reads 165,752 Writes
I/O Operations per second - Aggregate
750
500
250
0.00
17:12:00 17:12:30 17:13:00 17:13:30 17:14:00 17:14:30
554 Total 283 Reads 271 Writes
Users
Imaging
Event Log
Clear
Details
Provision Status

Multiple Host Statistics

Host	Disk Aggregate Percentage Used	Disk Aggregate Block Reads	Disk Aggregate Block Writes	Disk Aggregate I/O Read	Disk Aggregate I/O Writes
admin	4%	0	176	0	11
n001	1%	0	136	0	11
n002	1%	23	102	0	5
n003	1%	0	140	0	12
n004	1%	881	164566	12	181
n005	1%	0	87	0	3
n006	1%	0	85	0	3
n007	1%	0	233	0	23
n008	1%	0	227	0	22

Severity	Host	Date	Event
⚠	admin	10/31/2011 03:54:53 PM	The following host(s) have stopped responding: (none). The following host(s) a...
⚠	admin	10/31/2011 03:54:43 PM	The following host(s) have stopped responding: (none). The following host(s) a...
⚠	admin	10/31/2011 03:54:23 PM	The following host(s) have stopped responding: (none). The following host(s) a...
⚠	admin	10/31/2011 03:52:54 PM	The following host(s) have stopped responding: n00[1-8]. The following host(...
⚠	admin	10/25/2011 01:04:42 PM	The following host(s) have stopped responding: (none). The following host(s) a...

Network Tab

Select the Network tab to monitor packet transmissions and errors for the selected host(s).

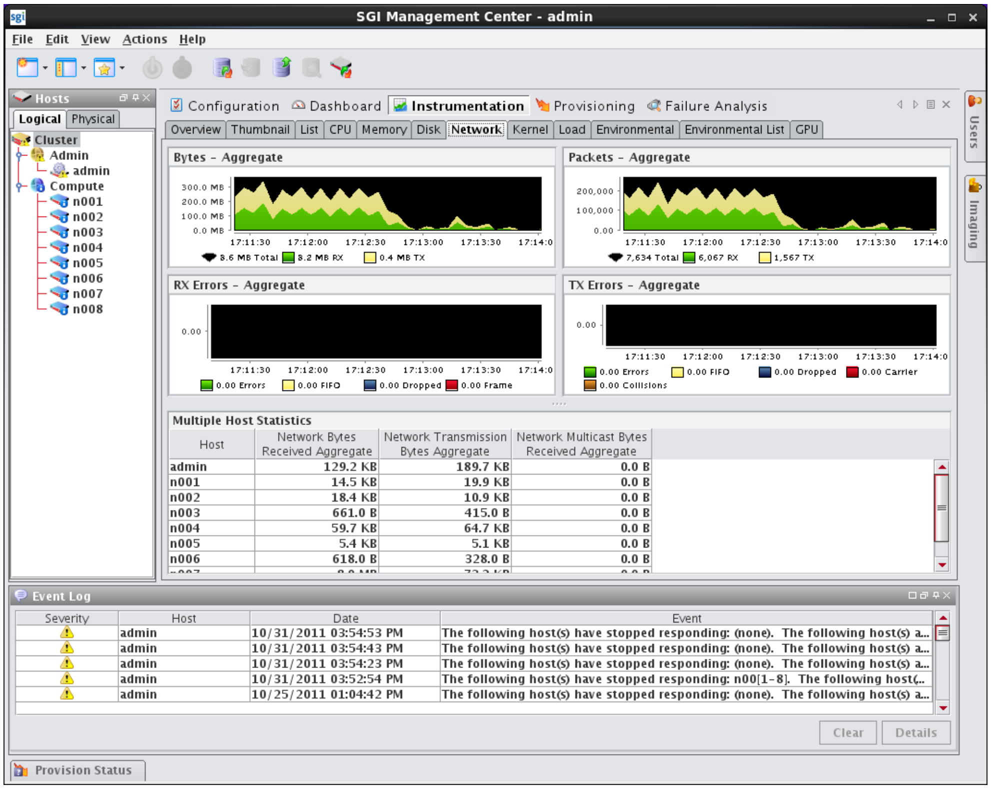
SGI Management Center - admin
File Edit View Actions Help
Hosts
Logical Physical
Cluster
Admin
admin
Compute
n001
n002
n003
n004
n005
n006
n007
n008
Configuration Dashboard Instrumentation Provisioning Failure Analysis
Overview Thumbnail List CPU Memory Disk Network Kernel Load Environmental Environmental List GPU
Users
Imaging
Bytes - Aggregate
300.0 MB
200.0 MB
100.0 MB
0.0 MB
17:11:30 17:12:00 17:12:30 17:13:00 17:13:30 17:14:0
8.6 MB Total 8.2 MB RX 0.4 MB TX
Packets - Aggregate
200,000
100,000
0.00
17:11:30 17:12:00 17:12:30 17:13:00 17:13:30 17:14:0
7,634 Total 6,067 RX 1,567 TX
RX Errors - Aggregate
0.00
17:11:30 17:12:00 17:12:30 17:13:00 17:13:30 17:14:0
0.00 Errors 0.00 FIFO 0.00 Dropped 0.00 Frame
TX Errors - Aggregate
0.00
17:11:30 17:12:00 17:12:30 17:13:00 17:13:30 17:14:0
0.00 Errors 0.00 FIFO 0.00 Dropped 0.00 Carrier
0.00 Collisions
Event Log
Clear
Details
Provision Status

Multiple Host Statistics

Host	Network Bytes Received Aggregate	Network Transmission Bytes Aggregate	Network Multicast Bytes Received Aggregate
admin	129.2 KB	189.7 KB	0.0 B
n001	14.5 KB	19.9 KB	0.0 B
n002	18.4 KB	10.9 KB	0.0 B
n003	661.0 B	415.0 B	0.0 B
n004	59.7 KB	64.7 KB	0.0 B
n005	5.4 KB	5.1 KB	0.0 B
n006	618.0 B	328.0 B	0.0 B

Event Log

Severity	Host	Date	Event
⚠	admin	10/31/2011 03:54:53 PM	The following host(s) have stopped responding: (none). The following host(s) a...
⚠	admin	10/31/2011 03:54:43 PM	The following host(s) have stopped responding: (none). The following host(s) a...
⚠	admin	10/31/2011 03:54:23 PM	The following host(s) have stopped responding: (none). The following host(s) a...
⚠	admin	10/31/2011 03:52:54 PM	The following host(s) have stopped responding: n00[1-8]. The following host(...
⚠	admin	10/25/2011 01:04:42 PM	The following host(s) have stopped responding: (none). The following host(s) a...

Kernel Tab

Select the Kernel tab to monitor the kernel information for the selected host(s).

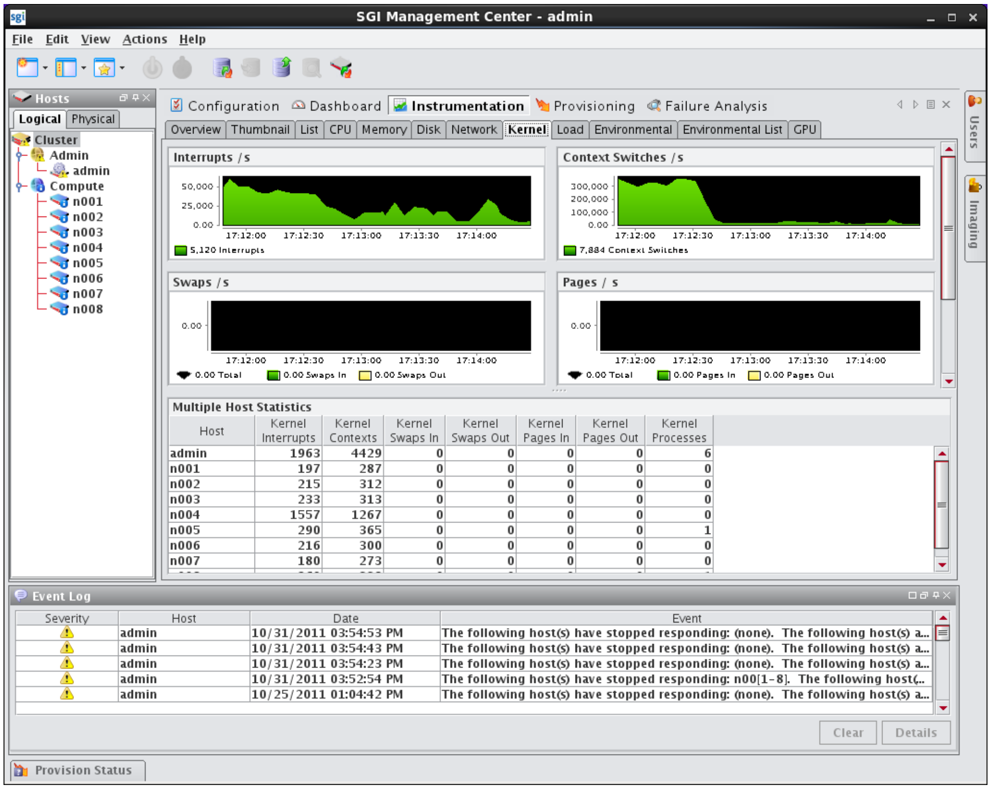
SGI Management Center - admin
File Edit View Actions Help
Hosts
Logical Physical
Cluster
Admin
admin
Compute
n001
n002
n003
n004
n005
n006
n007
n008
Configuration Dashboard Instrumentation Provisioning Failure Analysis
Overview Thumbnail List CPU Memory Disk Network Kernel Load Environmental Environmental List GPU
Interrupts / s
50,000
25,000
0.00
17:12:00 17:12:30 17:13:00 17:13:30 17:14:00
5,120 Interrupts
Context Switches / s
300,000
200,000
100,000
0.00
17:12:00 17:12:30 17:13:00 17:13:30 17:14:00
7,884 Context Switches
Swaps / s
0.00
17:12:00 17:12:30 17:13:00 17:13:30 17:14:00
0.00 Total
0.00 Swaps In
0.00 Swaps Out
Pages / s
0.00
17:12:00 17:12:30 17:13:00 17:13:30 17:14:00
0.00 Total
0.00 Pages In
0.00 Pages Out
Users
Imaging
Clear
Details
Provision Status

Multiple Host Statistics

Host	Kernel Interrupts	Kernel Contexts	Kernel Swaps In	Kernel Swaps Out	Kernel Pages In	Kernel Pages Out	Kernel Processes
admin	1963	4429	0	0	0	0	6
n001	197	287	0	0	0	0	0
n002	215	312	0	0	0	0	0
n003	233	313	0	0	0	0	0
n004	1557	1267	0	0	0	0	0
n005	290	365	0	0	0	0	1
n006	216	300	0	0	0	0	0
n007	180	273	0	0	0	0	0

Event Log

Severity	Host	Date	Event
⚠	admin	10/31/2011 03:54:53 PM	The following host(s) have stopped responding: (none). The following host(s) a...
⚠	admin	10/31/2011 03:54:43 PM	The following host(s) have stopped responding: (none). The following host(s) a...
⚠	admin	10/31/2011 03:54:23 PM	The following host(s) have stopped responding: (none). The following host(s) a...
⚠	admin	10/31/2011 03:52:54 PM	The following host(s) have stopped responding: n00[1-8]. The following host(...
⚠	admin	10/25/2011 01:04:42 PM	The following host(s) have stopped responding: (none). The following host(s) a...

Load Tab

Select the Load tab to monitor the load placed on the selected host(s).

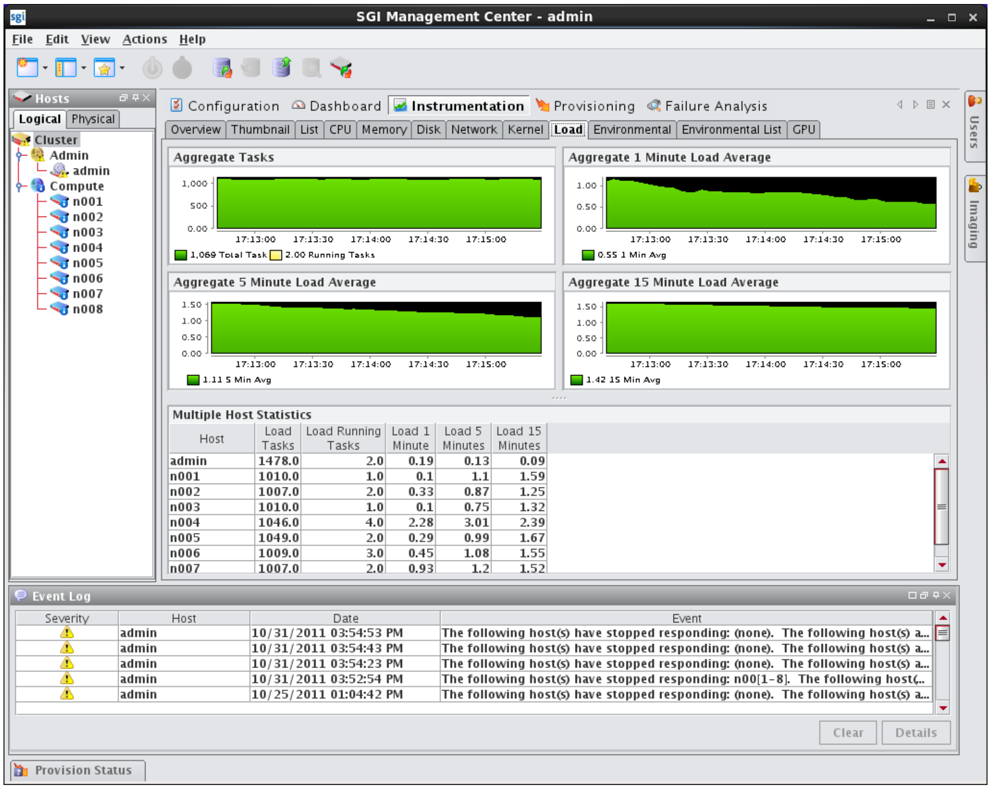
SGI Management Center - admin
File Edit View Actions Help
Hosts
Logical Physical
Cluster
Admin
admin
Compute
n001
n002
n003
n004
n005
n006
n007
n008
Configuration Dashboard Instrumentation Provisioning Failure Analysis
Overview Thumbnail List CPU Memory Disk Network Kernel Load Environmental Environmental List GPU
Aggregate Tasks
1,000
500
0.00
17:13:00 17:13:30 17:14:00 17:14:30 17:15:00
1,069 Total Task
2.00 Running Tasks
Aggregate 1 Minute Load Average
1.00
0.50
0.00
0.55 1 Min Avg
Aggregate 5 Minute Load Average
1.50
1.00
0.50
0.00
1.11 5 Min Avg
Aggregate 15 Minute Load Average
1.42 15 Min Avg
Users
Imaging
Clear
Details
Provision Status

Multiple Host Statistics

Host	Load Tasks	Load Running Tasks	Load 1 Minute	Load 5 Minutes	Load 15 Minutes
admin	1478.0	2.0	0.19	0.13	0.09
n001	1010.0	1.0	0.1	1.1	1.59
n002	1007.0	2.0	0.33	0.87	1.25
n003	1010.0	1.0	0.1	0.75	1.32
n004	1046.0	4.0	2.28	3.01	2.39
n005	1049.0	2.0	0.29	0.99	1.67
n006	1009.0	3.0	0.45	1.08	1.55
n007	1007.0	2.0	0.93	1.2	1.52

Event Log

Severity	Host	Date	Event
⚠	admin	10/31/2011 03:54:53 PM	The following host(s) have stopped responding: (none). The following host(s) a...
⚠	admin	10/31/2011 03:54:43 PM	The following host(s) have stopped responding: (none). The following host(s) a...
⚠	admin	10/31/2011 03:54:23 PM	The following host(s) have stopped responding: (none). The following host(s) a...
⚠	admin	10/31/2011 03:52:54 PM	The following host(s) have stopped responding: n00[1-8]. The following host(...
⚠	admin	10/25/2011 01:04:42 PM	The following host(s) have stopped responding: (none). The following host(s) a...

Environmental Tab

Select the Environmental tab to view the temperature summary readings for the selected host(s). Each summary contains up to five temperature readings—four processor temperatures followed by the ambient host temperature (which requires an Icecard). On hosts that support IPMI, these temperature readings differ slightly—two processor temperatures, two power supply temperatures, and the ambient host temperature.

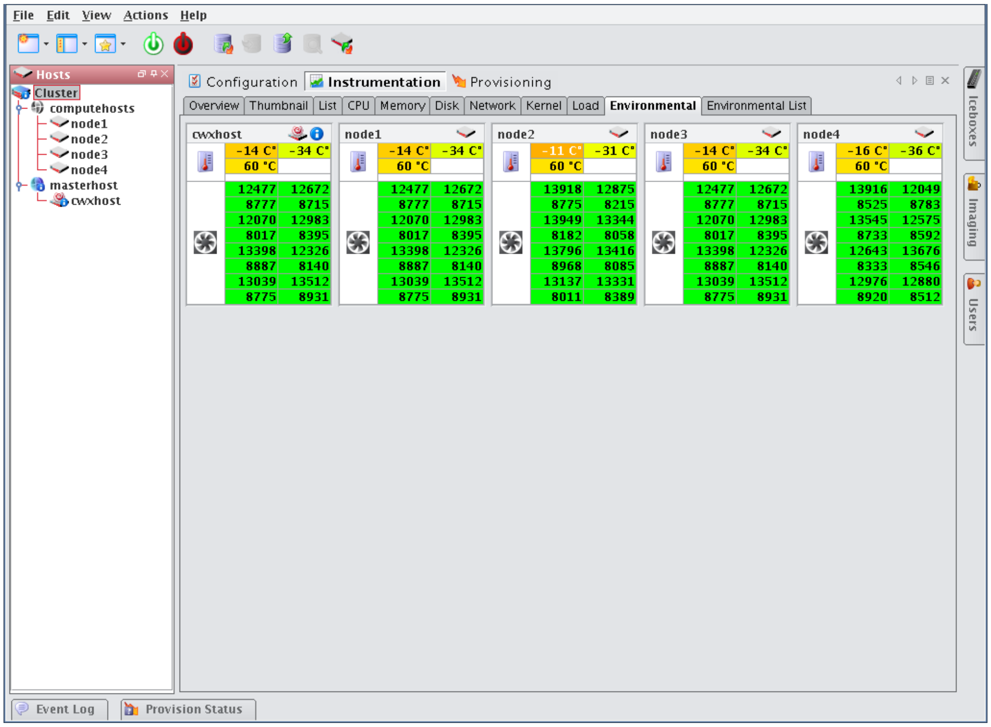
File
Edit
View
Actions
Help
Hosts
Cluster
computehosts
node1
node2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
Overview
Thumbnail
List
CPU
Memory
Disk
Network
Kernel
Load
Environmental
Environmental List
Iceboxes
Imaging
Users
Event Log
Provision Status

cwxhost		node1		node2		node3		node4	
-14 C°	-34 C°	-14 C°	-34 C°	-11 C°	-31 C°	-14 C°	-34 C°	-16 C°	-36 C°
60 °C		60 °C		60 °C		60 °C		60 °C	
12477	12672	12477	12672	13918	12875	12477	12672	13916	12049
8777	8715	8777	8715	8775	8215	8777	8715	8525	8783
12070	12983	12070	12983	13949	13344	12070	12983	13545	12575
8017	8395	8017	8395	8182	8058	8017	8395	8733	8592
13398	12326	13398	12326	13796	13416	13398	12326	12643	13676
8887	8140	8887	8140	8968	8085	8887	8140	8333	8546
13039	13512	13039	13512	13137	13331	13039	13512	12976	12880
8775	8931	8775	8931	8011	8389	8775	8931	8920	8512

Note

The processor temperature readings for IPMI-based hosts indicate the amount of temperature change that must occur before the CPU’s thermal control circuitry activates to prevent damage to the CPU. These are not actual CPU temperatures.

From the Environmental tab, you can access the following options from the View menu:

Filter Filter and display hosts based on error status

Size Change the size of the thumbnail view (small, medium, or large). Small thumbnails support a mouse-over function to display a host summary.

Temperatures Set temperature options to display values as Celsius or Fahrenheit. Temperatures range from Green (Cool) to Yellow (warm) to Red (Hot). Fan speeds follow the same convention—slow or stopped fans appear in red.

Environmental List Tab

Select the Environmental tab to open the list view of the temperature summary readings for the selected host(s).

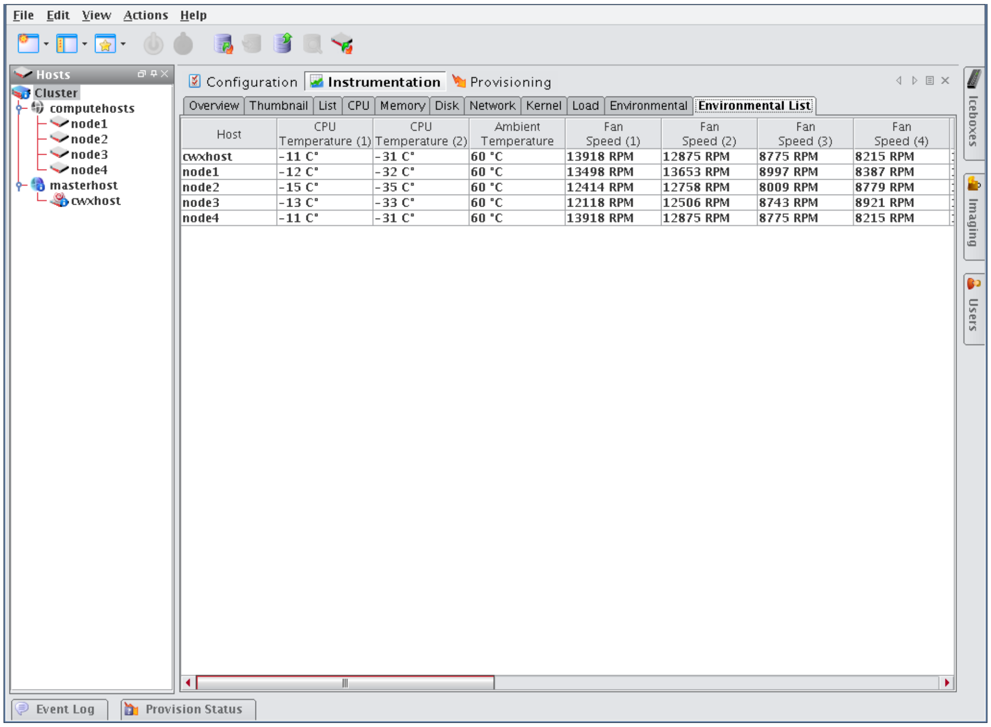
File Edit View Actions Help
Hosts
Cluster
computehosts
node1
node2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
Overview
Thumbnail
List
CPU
Memory
Disk
Network
Kernel
Load
Environmental
Environmental List
Iceboxes
Imaging
Users
Event Log
Provision Status

Host	CPU Temperature (1)	CPU Temperature (2)	Ambient Temperature	Fan Speed (1)	Fan Speed (2)	Fan Speed (3)	Fan Speed (4)
cwxhost	-11 C°	-31 C°	60 °C	13918 RPM	12875 RPM	8775 RPM	8215 RPM
node1	-12 C°	-32 C°	60 °C	13498 RPM	13653 RPM	8997 RPM	8387 RPM
node2	-15 C°	-35 C°	60 °C	12414 RPM	12758 RPM	8009 RPM	8779 RPM
node3	-13 C°	-33 C°	60 °C	12118 RPM	12506 RPM	8743 RPM	8921 RPM
node4	-11 C°	-31 C°	60 °C	13918 RPM	12875 RPM	8775 RPM	8215 RPM

GPU Tab

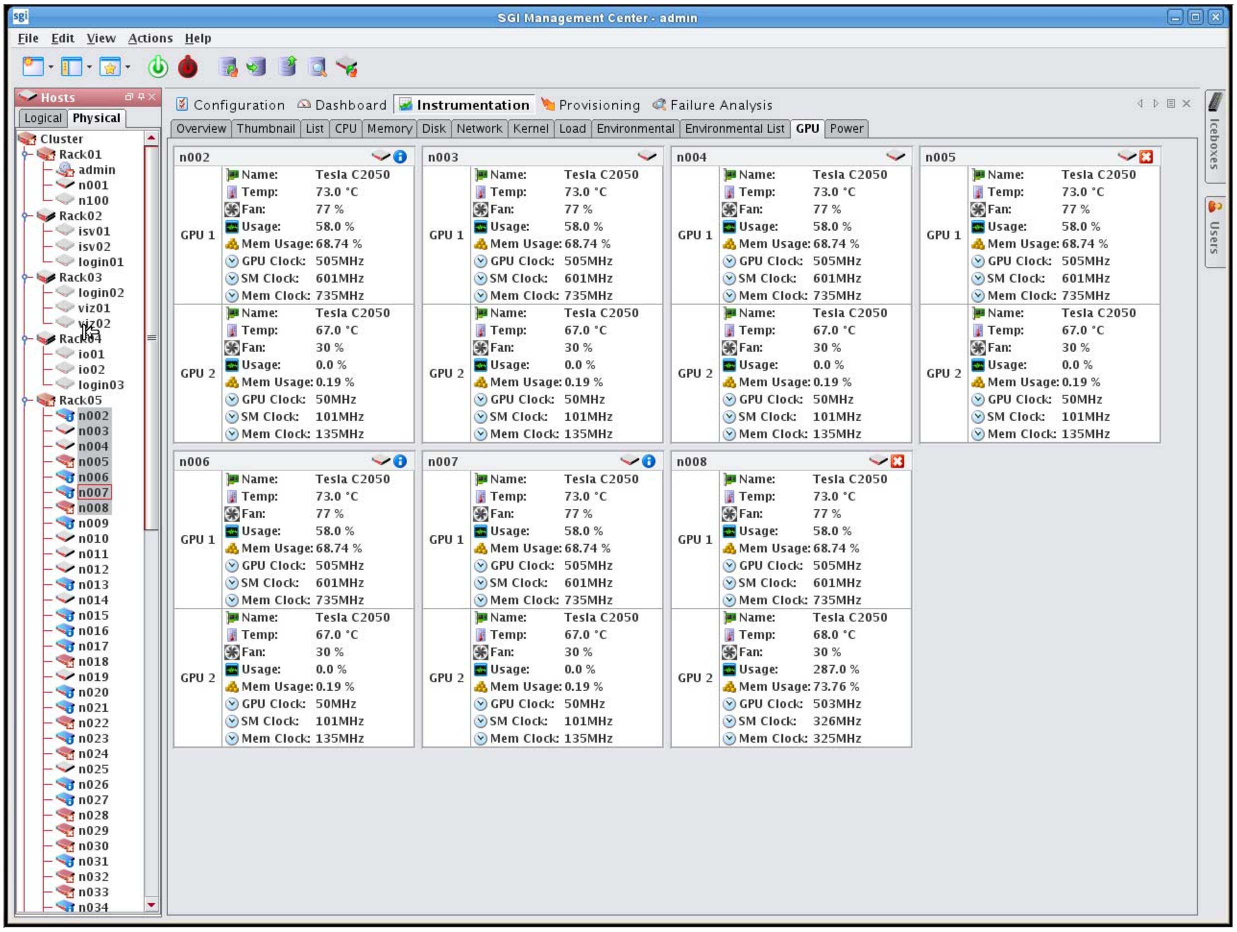
SGI Management Center - admin
File
Edit
View
Actions
Help
Hosts
Logical
Physical
Cluster
Rack01
admin
n001
n100
Rack02
isv01
isv02
login01
Rack03
login02
viz01
viz02
Rack04
io01
io02
login03
Rack05
n002
n003
n004
n005
n006
n007
n008
n009
n010
n011
n012
n013
n014
n015
n016
n017
n018
n019
n020
n021
n022
n023
n024
n025
n026
n027
n028
n029
n030
n031
n032
n033
n034
Configuration
Dashboard
Instrumentation
Provisioning
Failure Analysis
Overview
Thumbnail
List
CPU
Memory
Disk
Network
Kernel
Load
Environmental
Environmental List
GPU
Power
Iceboxes
Users

n002

GPU 1
Name: Tesla C2050
Temp: 73.0 °C
Fan: 77 %
Usage: 58.0 %
Mem Usage: 68.74 %
GPU Clock: 505MHz
SM Clock: 601MHz
Mem Clock: 735MHz

GPU 2
Name: Tesla C2050
Temp: 67.0 °C
Fan: 30 %
Usage: 0.0 %
Mem Usage: 0.19 %
GPU Clock: 50MHz
SM Clock: 101MHz
Mem Clock: 135MHz

n003

GPU 1
Name: Tesla C2050
Temp: 73.0 °C
Fan: 77 %
Usage: 58.0 %
Mem Usage: 68.74 %
GPU Clock: 505MHz
SM Clock: 601MHz
Mem Clock: 735MHz

GPU 2
Name: Tesla C2050
Temp: 67.0 °C
Fan: 30 %
Usage: 0.0 %
Mem Usage: 0.19 %
GPU Clock: 50MHz
SM Clock: 101MHz
Mem Clock: 135MHz

n004

GPU 1
Name: Tesla C2050
Temp: 73.0 °C
Fan: 77 %
Usage: 58.0 %
Mem Usage: 68.74 %
GPU Clock: 505MHz
SM Clock: 601MHz
Mem Clock: 735MHz

GPU 2
Name: Tesla C2050
Temp: 67.0 °C
Fan: 30 %
Usage: 0.0 %
Mem Usage: 0.19 %
GPU Clock: 50MHz
SM Clock: 101MHz
Mem Clock: 135MHz

n005

GPU 1
Name: Tesla C2050
Temp: 73.0 °C
Fan: 77 %
Usage: 58.0 %
Mem Usage: 68.74 %
GPU Clock: 505MHz
SM Clock: 601MHz
Mem Clock: 735MHz

GPU 2
Name: Tesla C2050
Temp: 67.0 °C
Fan: 30 %
Usage: 0.0 %
Mem Usage: 0.19 %
GPU Clock: 50MHz
SM Clock: 101MHz
Mem Clock: 135MHz

n006

GPU 1
Name: Tesla C2050
Temp: 73.0 °C
Fan: 77 %
Usage: 58.0 %
Mem Usage: 68.74 %
GPU Clock: 505MHz
SM Clock: 601MHz
Mem Clock: 735MHz

GPU 2
Name: Tesla C2050
Temp: 67.0 °C
Fan: 30 %
Usage: 0.0 %
Mem Usage: 0.19 %
GPU Clock: 50MHz
SM Clock: 101MHz
Mem Clock: 135MHz

n007

GPU 1
Name: Tesla C2050
Temp: 73.0 °C
Fan: 77 %
Usage: 58.0 %
Mem Usage: 68.74 %
GPU Clock: 505MHz
SM Clock: 601MHz
Mem Clock: 735MHz

GPU 2
Name: Tesla C2050
Temp: 67.0 °C
Fan: 30 %
Usage: 0.0 %
Mem Usage: 0.19 %
GPU Clock: 50MHz
SM Clock: 101MHz
Mem Clock: 135MHz

n008

GPU 1
Name: Tesla C2050
Temp: 73.0 °C
Fan: 77 %
Usage: 58.0 %
Mem Usage: 68.74 %
GPU Clock: 505MHz
SM Clock: 601MHz
Mem Clock: 735MHz

GPU 2
Name: Tesla C2050
Temp: 68.0 °C
Fan: 30 %
Usage: 287.0 %
Mem Usage: 73.76 %
GPU Clock: 503MHz
SM Clock: 326MHz
Mem Clock: 325MHz

As shown in the panel above, SGI Management Center provides monitoring of supported GPUs (items like temperature, fan speed, memory usage, and ECC). For a listing of GPU solutions supported by SGI see http://www.sgi.com/pdfs/4235.pdf.

Power Tab

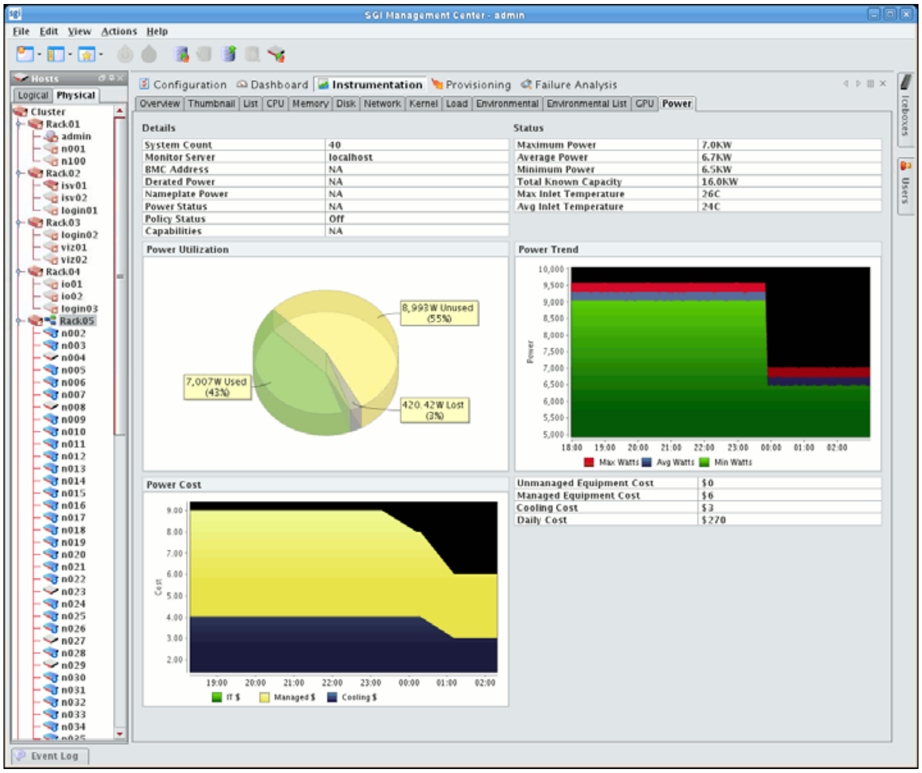
SGI Management Center - admin
File Edit View Actions Help
Hosts
Logical Physical
Cluster
Rack01
admin
n001
n100
Rack02
isv01
isv02
login01
Rack03
login02
viz01
viz02
Rack04
io01
io02
login03
Rack05
Event Log
Configuration Dashboard Instrumentation Provisioning Failure Analysis
Overview Thumbnail List CPU Memory Disk Network Kernel Load Environmental Environmental List GPU Power
Iceboxes
Users
Power Utilization
8,993W Unused (55%)
7,007W Used (43%)
420.42W Lost (3%)
Power Trend
10,000 9,500 9,000 8,500 8,000 7,500 7,000 6,500 6,000 5,500 5,000
Power
18:00 19:00 20:00 21:00 22:00 23:00 00:00 01:00 02:00
Max Watts Avg Watts Min Watts
Power Cost
9.00 8.00 7.00 6.00 5.00 4.00 3.00 2.00
Cost
19:00 20:00 21:00 22:00 23:00 00:00 01:00 02:00
IT $ Managed $ Cooling $

Details

System Count	40
Monitor Server	localhost
BMC Address	NA
Derated Power	NA
Nameplate Power	NA
Power Status	NA
Policy Status	Off
Capabilities	NA

Status

Maximum Power	7.0KW
Average Power	6.7KW
Minimum Power	6.5KW
Total Known Capacity	16.0KW
Max Inlet Temperature	26C
Avg Inlet Temperature	24C

Unmanaged Equipment Cost	$0
Managed Equipment Cost	$6
Cooling Cost	$3
Daily Cost	$270

Note

The SGI Management Center DCM integration uses indirect TCP communications through an external web service provider to accumulate and display power monitoring data. This results in a delay in updating instrumentation for every tree selection change. Consequently, the waiting time for initial **Power** panel updates for large-scale systems using DCM may be several minutes in duration.

The following are the primary components of the **Power** panel:

- **Details** table
- **Status** table
- **Power Utilization** pie chart
- **Power Trend** chart
- **Power Cost** chart
- Power Cost Details table

This section describes these components. For more details, see *Power Policy Management* on page 33 and *Power Cost Tracking and Trending* on page 36.

Details Table

The Details table contains configuration details for various power-related entities.

Attribute	Description
System Count	Displays the number of systems associated with the selected entity, if available.
Monitor Server	Displays the power service provider location currently in use.
BMC Address	Shows the endpoint BMC address currently in use (if several or none are offered, this entry is unavailable.)
Derated Power	Displays the calculated derated power for the endpoint.
Nameplate Power	Displays the calculated or configured nameplate power for the endpoint.
Power Status	Indicates the mechanical power state of the endpoint.
Policy Status	Indicates the policy activation status for the endpoint.
Capabilities	Describes the discovered power monitoring/MGT capabilities for the endpoint.

Status Table

The **Status** table contains power sampling and measurement data.

Data Item	Description
Maximum Power	Total maximum power measurement recorded in any monitoring cycle for all sampling intervals within the aggregation period applied for the selected entity (which may be a group). max{ sum_T1(max_N1{ P1, P2,...,Pn}, ..., max_Nn{...}), ..., sum_Tn(...)}
Average Power	Sum of the entity/group mean power measurements as given by the sum of the arithmetic mean of power measurements for all sub-nodes within the specified entity/group for all sampling intervals within the aggregation period. avg{ sum_T1(avg_N1{ P1, P2, ..., Pn}, ..., avg_Nn{...}), ..., sum_Tn(...)}
Minimum Power	Result calculated much like **Maximum Power**, but with floor. Pn is the last monitoring cycle in a sampling interval. Tn is the last sampling interval in an aggregation period.
Total Known Capacity	A measured, calculated and/or configured sum of all of the derated power components for the selected endpoint.
Max Inlet Temperature **Avg Inlet Temperature**	Taken from a prescribed IPMI/SMBUS-accessible inlet-air sensor for the endpoint (if available).

Power Utilization Pie Chart

The **Power Utilization** pie chart displays the three metrics in the following chart:

Metric	Description
Used	Indicates the last instantaneously sampled power measurement for the endpoint.
Unused	Calculated as the total-known-capacity less the used figure.
Lost	Calculated using the configured or estimated power factor for the endpoint, and is generally an estimate on the efficiency of the power distribution for a node or rack.

Power Trend Chart

The **Power Trend** chart displays the same miniumum, average, and maximum metrics that are shown in the **Status** table but does so in an interactive, historical chart. This chart features user-defined zooming and scaling and you can save the chart for reporting purposes. Simply right-click on the chart or click and drag the mouse to explore further.

Power Cost Chart

The **Power Cost** chart displays energy costs over time for each endpoint in the cluster. The trend incorporates tracking against kilowatt hours and leverages the globally configured parameters **Power Cost per KWh** and **Cooling Power Multiplier** (See *Enabling Power Cost Tracking* on page 37). Therefore, each chart datapoint expresses energy cost for the last hour of utilization. The chart incorporates the three metrics described in the following table:

Metric	Description
IT $	Information Technology equipment energy costs, expressed as the per-hour cost for actively-monitored equipment plus statically metered unmanaged equipment.
Managed $	Managed equipment energy costs, expressed as the per-hour cost for actively-monitored equipment.
Cooling $	Cooling energy costs, expressed as the per-hour cooling energy cost as outlined in section *Enabling Power Cost Tracking* on page 37.

Power Cost Details Table

The Power Cost Details table consists of the following metrics:

Metric	Description
Unmanaged Equipment Cost	The per-hour cost for non-monitored equipment whose power utilization characteristics are statically configured.
Managed Equipment Cost	The per-hour cost for actively-monitored equipment.
Cooling Cost	The per-hour cooling energy cost as outlined above.
Daily Cost	The cost for all unmanaged, managed and cooling power utilization for the last 24 hours.

Failure Analysis

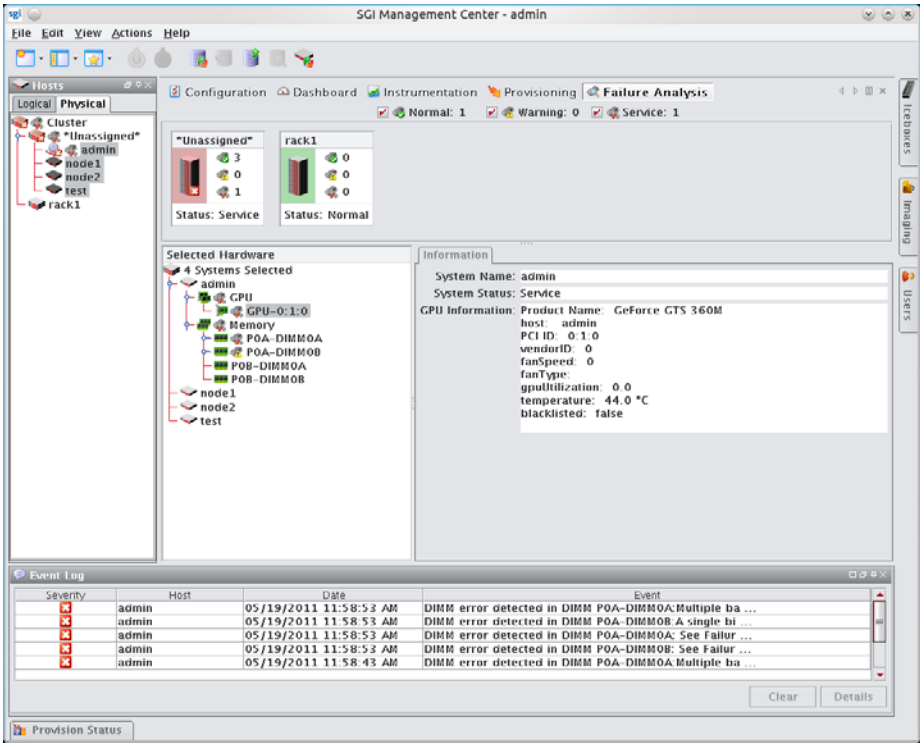
SGI Management Center - admin
File Edit View Actions Help
Hosts
Logical Physical
Cluster
Unassigned
admin
node1
node2
test
rack1
Configuration Dashboard Instrumentation Provisioning Failure Analysis
Normal: 1 Warning: 0 Service: 1
Unassigned
3
0
1
Status: Service
rack1
0
0
0
Status: Normal
Iceboxes
Imaging
Users
Selected Hardware
4 Systems Selected
admin
GPU
GPU-0:1:0
Memory
P0A-DIMM0A
P0A-DIMM0B
P0B-DIMM0A
P0B-DIMM0B
node1
node2
test
Information
Event Log
Clear
Details
Provision Status

System Name: admin
System Status: Service
GPU Information: Product Name: GeForce GTS 360M
host: admin
PCI ID: 0:1:0
vendorID: 0
fanSpeed: 0
fanType:
gpuUtilization: 0.0
temperature: 44.0 °C
blacklisted: false

Severity	Host	Date	Event
☒	admin	05/19/2011 11:58:53 AM	DIMM error detected in DIMM P0A-DIMM0A:Multiple ba ...
☒	admin	05/19/2011 11:58:53 AM	DIMM error detected in DIMM P0A-DIMM0B:A single bi ...
☒	admin	05/19/2011 11:58:53 AM	DIMM error detected in DIMM P0A-DIMM0A: See Failur ...
☒	admin	05/19/2011 11:58:53 AM	DIMM error detected in DIMM P0A-DIMM0B: See Failur ...
☒	admin	05/19/2011 11:58:43 AM	DIMM error detected in DIMM P0A-DIMM0A:Multiple ba ...

SGI Management Center supports failure analysis for memory errors via memlog, a software component of SGI Foundation Software.

Management Center Monitoring and Event Subsystem

Management Center uses a monitoring and event system to track system values. This system includes monitors, metrics, listeners, and loggers that collect values from the cluster, then display this information using the Management Center instrumentation GUI (see *Instrumentation* on page 153). You can extend the standard monitoring and event system to include custom values and set thresholds for user-defined events. For example:

- Monitoring custom values using scripts.
- Displaying custom values in the Management Center list view.
- Setting thresholds on values and taking an action if these thresholds are exceeded.
- Logging custom error conditions in the Management Center log.
- Running custom scripts as event actions.

Monitors run at a set interval and collect information from each host. Listeners receive information about metrics from the instrumentation service, then determine if the values are reasonable. If a listener determines that a metric is above or below a set threshold, the listener triggers a logger to take a specific action.

Typically, configuration files are host-specific and are located in the *$MGR_HOME/etc* directory. If you modify the configuration files, you can copy them into the payload to make them available on each host after you provision.

Note

By default, Management Center creates a backup of the *$MGR_HOME/etc* directory during installation and copies it to *$MGR_HOME/etc.bak.<date>.<timestamp>*

Monitors

Management Center Monitors run periodically on the cluster and provide metrics that are gathered, processed, and displayed using the Management Center instrumentation GUI. Using monitors allows you to “tune” Management Center to meet your exact system needs by enabling or disabling specific monitors or by setting the rate at which monitors run. In cases where pre-defined monitors simply do not meet your specific needs, Management Center also allows you to create custom monitors (see *Custom Monitors* on page 181). The following table lists the Management Center default monitors.

Monitor Name	Interval
NFS Client	5
NFS Server	5
BlueSmoke	500
Disk	5
Disk Space	60
Identity	5
Kernel	5
LinuxBIOS	86400
Load	15
Memory	5
Network	5
Uptime	60
Environmental	5

All standard Management Center monitors are configured in the *InstrumentationMonitors.profile* in the *$MGR_HOME/etc* directory. The format of the monitor configuration in the file is generally as follows (where <*time*> is in milliseconds):

```
<name>: com.lnxi.instrumentation.server.<monitor_name>
<name>.interval: <time>
```

When working with standard monitors, it is strongly recommended that you leave all monitors enabled—however, you can increase how often these monitors run. Raising the interval can reduce CPU time and network use for monitoring. Because Management Center uses very little CPU processing time on the compute hosts, values as high as 1 second (1000 milliseconds) are nearly undetectable. By default, some monitors are set to run at 5 seconds (5000 milliseconds) or longer.

Note

When monitoring the Management Center Master Host, the name of the Master Host must match the name assigned in *$MGR_HOME/@genesis.profile*.

To Enable or Disable a Monitor

1. Select Event Administration from the Edit menu. The Event Administration dialog appears.

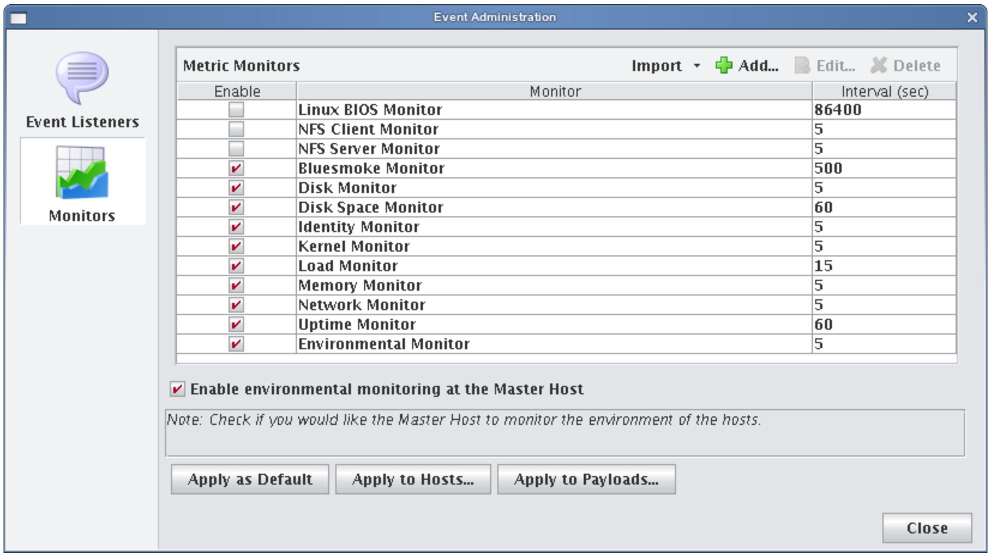
Event Administration
Event Listeners
Monitors
Metric Monitors
Import
Add...
Edit...
Delete
Apply as Default
Apply to Hosts...
Apply to Payloads...
Close

Enable	Monitor	Interval (sec)
☐	Linux BIOS Monitor	86400
☐	NFS Client Monitor	5
☐	NFS Server Monitor	5
☑	Bluesmoke Monitor	500
☑	Disk Monitor	5
☑	Disk Space Monitor	60
☑	Identity Monitor	5
☑	Kernel Monitor	5
☑	Load Monitor	15
☑	Memory Monitor	5
☑	Network Monitor	5
☑	Uptime Monitor	60
☑	Environmental Monitor	5

☑ Enable environmental monitoring at the Master Host

Note: Check if you would like the Master Host to monitor the environment of the hosts.

2. Select Monitors.
3. Check or un-check the box next to each monitor you want to enable or disable.
4. (Optional) Click Apply as Default to apply the listener configuration as the default on the Master Host and payload. Management Center saves the listeners in *InstrumentationMonitors.profile.default*.
5. (Optional) Click Apply to Hosts to apply the monitor to a specific host(s). The Export to Hosts dialog appears.

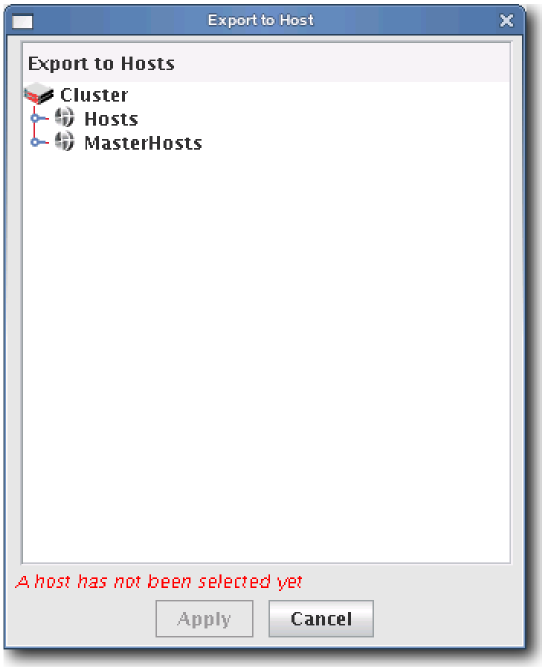
Export to Host
Export to Hosts
Cluster
Hosts
MasterHosts
Apply
Cancel

A host has not been selected yet

   A. Select the host(s) to which to export the monitors from the navigation tree.

   B. Click Apply to save changes or click Close to abort this action.

6. (Optional) Click Apply to Payloads to include these monitors as part of a payload. The Export to Payloads appears.

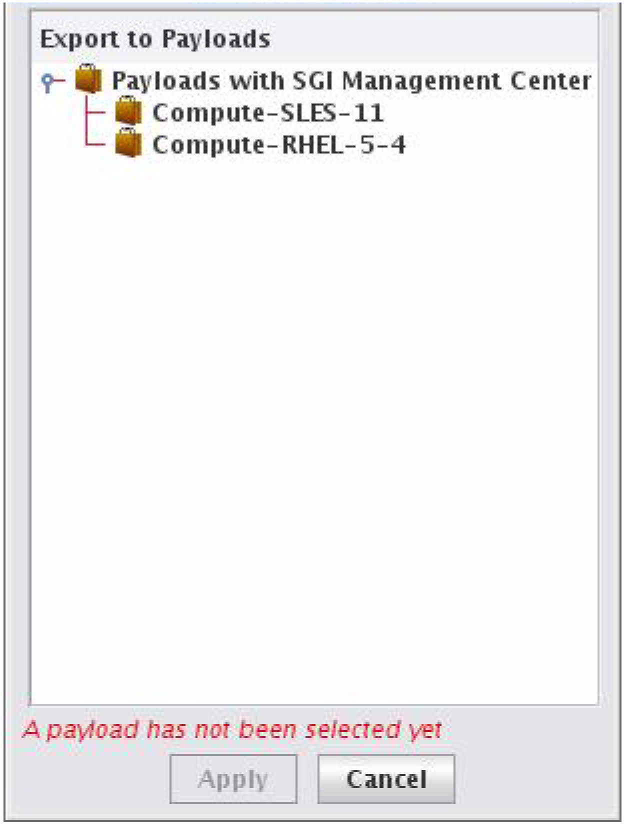
Export to Payloads
Payloads with SGI Management Center
Compute-SLES-11
Compute-RHEL-5-4
Apply
Cancel

A payload has not been selected yet

   A. Select the payload(s) to which to apply the monitors.
   B. Click Apply to save changes or click Cancel to abort this action.

7. Click Close to complete this action and close the Event Administration dialog.

Note

If you click close without applying your changes, all modifications will be lost.

To Add a Monitor

1. Select Event Administration from the Edit menu. The Event Administration dialog appears.
2. Select Monitors.
3. Click Add. The Add Custom Monitor dialog appears.

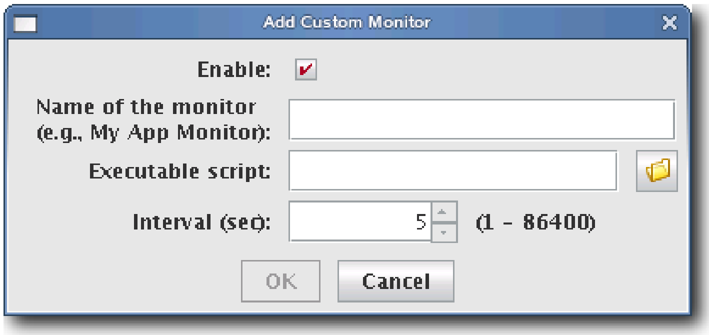
Add Custom Monitor
Enable: ☑
Name of the monitor (e.g., My App Monitor):
Executable script:
Interval (sec): 5
(1 - 86400)
OK
Cancel

Tip

For information on creating a custom monitor, see *Custom Monitors* on page 181.

4. Enter the name of the monitor.
5. Enter the path of the executable script used for this monitor or click browse to locate the script.

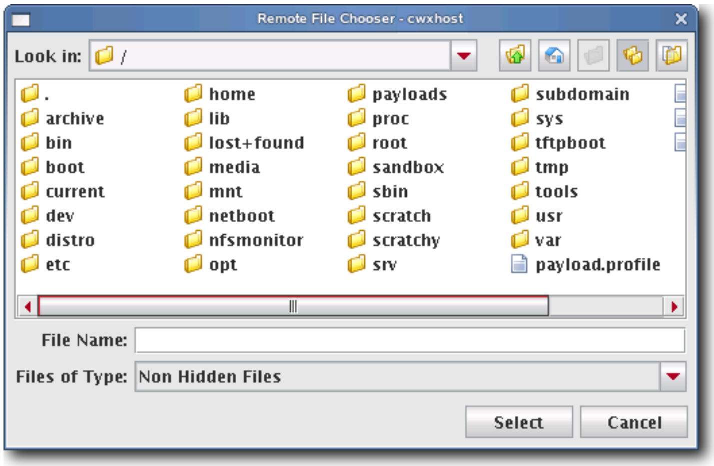
Remote File Chooser - cwxhost
Look in: /
.
archive
bin
boot
current
dev
distro
etc
home
lib
lost+found
media
mnt
netboot
nfsmonitor
opt
payloads
proc
root
sandbox
sbin
scratch
scratchy
srv
subdomain
sys
tftpboot
tmp
tools
usr
var
payload.profile
File Name:
Files of Type: Non Hidden Files
Select
Cancel

6. Enter the monitoring interval (in seconds).
7. Check the Enable option to activate the listener.
8. Click OK to continue or click Cancel to abort this action.
9. (Optional) Apply the monitor to hosts or payloads.

10. Click Apply as Default to save the monitor.

Note

When you add a monitor and click Apply as Default, Management Center saves the monitor as one of the default monitors—all future payloads will contain the new monitor. Furthermore, the new monitor will be included any time you install Management Center into a payload.

11. Click Close.

To Import Monitors

IMPORT FROM HOST

1. Select Event Administration from the Edit menu. The Event Administration dialog appears.

2. Select Monitors.

3. Click Import and select Import from Host. The Import from Hosts dialog appears.

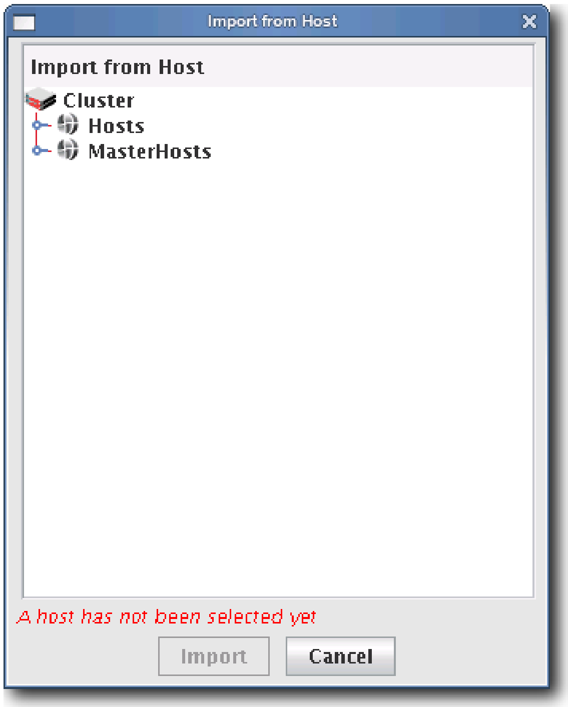
Import from Host
Import from Host
Cluster
Hosts
MasterHosts
Import
Cancel

A host has not been selected yet

4. Select the host from which to import listeners and click Import. Click Cancel to abort this action.

IMPORT FROM PAYLOAD

1. Select Event Administration from the Edit menu. The Event Administration dialog appears.
2. Select Monitors.
3. Click Import and select Import from Payload. The Import from Payloads dialog appears.

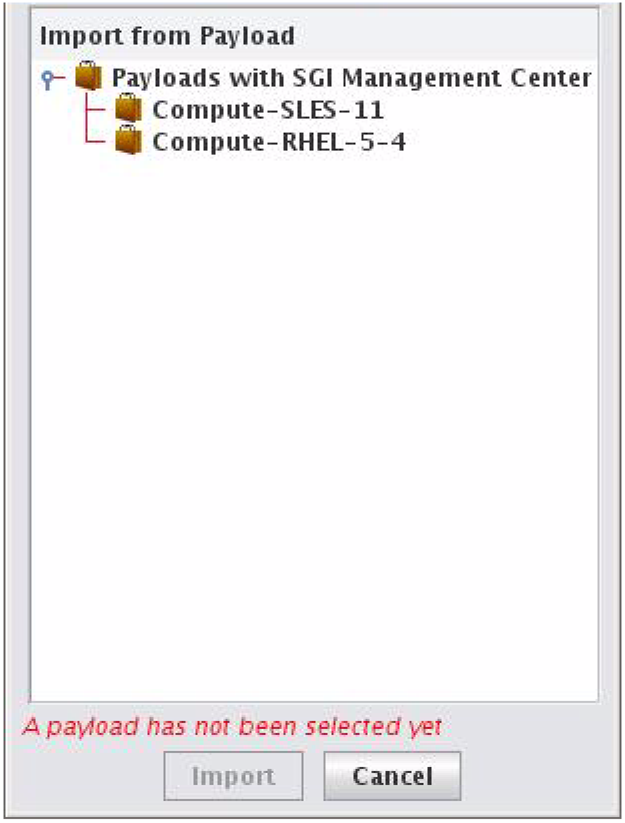

A payload has not been selected yet

4. Select the payload from which to import listeners and click Import. Click Cancel to abort this action.

IMPORT DEFAULT

1. Select Event Administration from the Edit menu. The Event Administration dialog appears.
2. Select Monitors.
3. Click Import and select Import Default. Management Center restores all monitors stored as default monitors in *InstrumentationMonitors.profile.default*. See *To Enable or Disable a Listener* on page 190 for information on adding default listeners.

RESTORE FACTORY SETTINGS

1. Select Event Administration from the Edit menu. The Event Administration dialog appears.
2. Select Monitors.
3. Click Import and select Restore Factory Settings. Management Center reverts the default monitors that shipped with Management Center.

To Edit a Monitor

1. Select Event Administration from the Edit menu. The Event Administration dialog appears.
2. Select Monitors.
3. Double-click a monitor in the list or select the monitor and click Edit. The edit dialog appears.

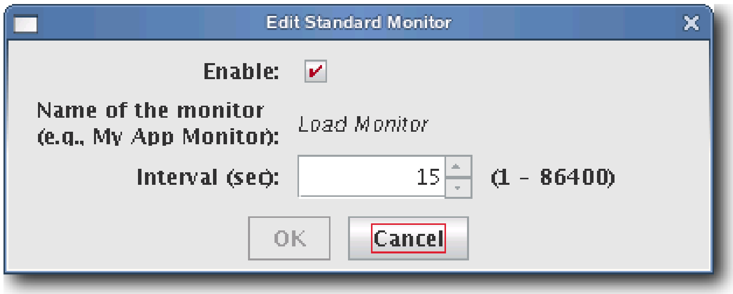

Enable: ☑

Name of the monitor (e.g., My App Monitor): Load Monitor

Interval (sec): 15 (1 - 86400)

4. Make any necessary modifications, then click OK to apply your changes. Click Cancel to abort this action.
5. (Optional) Apply the monitor to hosts or payloads.
6. Click Apply as Default to save the monitor.

Note

When you change a monitor and click Apply as Default, Management Center saves the monitor as one of the default monitors—all future payloads will contain the new monitor. Furthermore, the new monitor will be included any time you install Management Center into a payload.

7. Click Close.

To Delete a Monitor

1. Select Event Administration from the Edit menu. The Event Administration dialog appears.
2. Select Monitors.
3. Select a listener from the list and click Delete.

Note

You cannot delete Management Center default monitors—these monitors can be disabled only.

4. Management Center asks you to confirm your action.

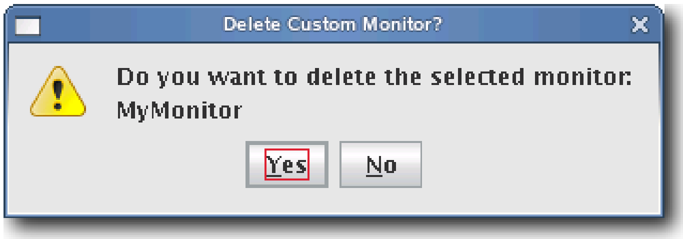

Do you want to delete the selected monitor: MyMonitor

5. Click Yes to delete the listener or click No to abort this action.

Custom Monitors

Custom monitors are added by creating a new monitor with the Management Center GUI and including a user-defined program or script that returns information in a format Management Center can process.

Note

The name must be unique for each monitor.

! Warning!

Test scripts carefully! Running an invalid script may cause undesired results with Management Center.

Because monitors typically invoke a script (e.g., bash, perl), using values of less than 5 seconds is not recommended (but is supported). To use a custom monitor, the program or script called by the monitor must return values to STDOUT in *key:value* pairs that use the following format:

```
hosts.<hostname>.<name>.<key1>:<value1>\n
hosts.<hostname>.<name>.<key2>:<value2>\n
```

The <*hostname*> refers to the name of the host from which you are running the script.

Note

When monitoring the Management Center Master Host, the name of the Master Host must match the name assigned in *$MGR_HOME/@genesis.profile*.

The <*name*> is the same name used in the *InstrumentationMonitors.profile.default*.

The <*key*> parameter refers to what is being monitored.

The <*value*> is the return value for that key. The script can return one or more items as long as they all have a key and value. The value can be any string or number, but the script is responsible for the formatting. The \n at the end is a newline character (required).

To Add a Custom Monitor

Note

You must configure new metrics as part of this process. See *Custom Metrics Example* on page 187 for a continuation of this example.

1. Open *InstrumentationMonitors.profile.default* from *$MGR_HOME/etc*.
2. Add the new monitor to the custom monitors profile. The following example uses perl to monitor how many users are logged into a host. The script returns two values: how many people are logged in and who the people are. The script name is *$MGR_HOME/bin/who.pl* and returns *who.who* and *who.count*.

```
#!/usr/bin/perl -w
# Basic modules are allowed
use IO::File;
use Sys::Hostname;

    $host = hostname;

    my @users;

    # This opens the program and runs it. Don't forget the '|' on the end
    my $fh = new IO::File('/usr/bin/who |');
    # If the program was started
    if (defined $fh) {
        # Then loop through its output until you get an eof.
        while (defined($line = <$fh>)) {
            if ($line =~ m/^\w+.*/) {
                $line =~ m/^(\w+).*$/;
                push(@users,$1);
            }
        }
        # Close the file.
        $fh->close();
    }

    # Remove duplicate entries of who.
    %seen = ();
    foreach $item (@users) {
        push(@uniq, $item) unless $seen{$item}++;
        }
    # Count how many items are in the array for our count
        $count = scalar(@uniq);

    # Rather than an array of values, just return a single text string;
    foreach $users(@uniq) {
        $who .= “$users,”;
        }
        chop($who);

    print “hosts.”. $host . “.who.count:” . $count .“\n”;
    print “hosts.”. $host . “.who.who:” . join(“,”, $who).“\n”;
```

When you run the script on host “n2” (assuming that perl and the perl modules above are installed correctly), the following prints to STDOUT:

```
[root@n2 root]# ./who.pl
hosts.n2.who.count:1
hosts.n2.who.who:root
```

Note

The script MUST exist on the hosts that will run this monitor. Therefore, you must either copy this script to each host (*$MGR_HOME/bin*) or configure the payload to include the script and provision the hosts with the new payload.

3. Restart Management Center services.
4. Select Event Administration from the Edit menu. The Event Administration dialog appears.
5. Select Monitors—Management Center displays the new monitor.

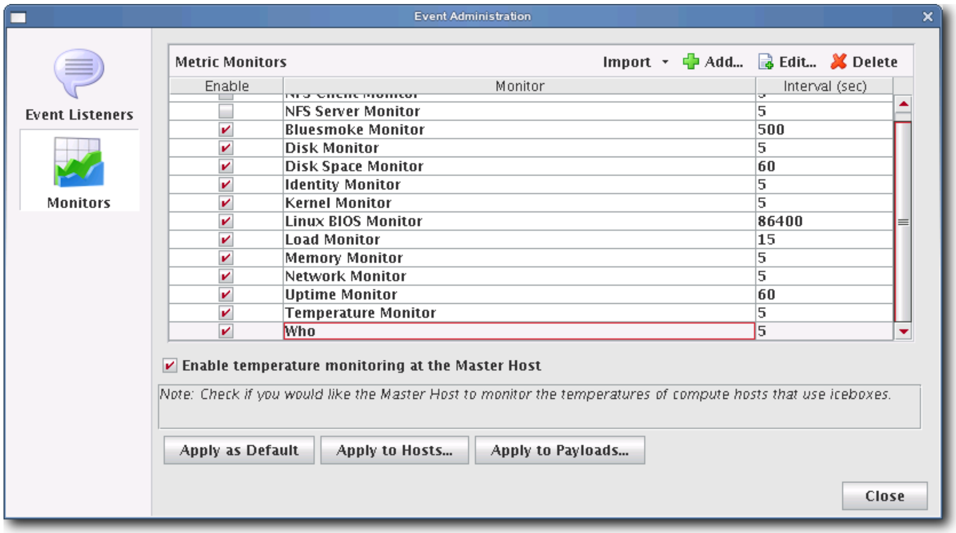
Event Administration
Event Listeners
Monitors
Metric Monitors
Import
Add...
Edit...
Delete
Apply as Default
Apply to Hosts...
Apply to Payloads...
Close

Enable	Monitor	Interval (sec)
☐	NFS Server Monitor	5
☑	Bluesmoke Monitor	500
☑	Disk Monitor	5
☑	Disk Space Monitor	60
☑	Identity Monitor	5
☑	Kernel Monitor	5
☑	Linux BIOS Monitor	86400
☑	Load Monitor	15
☑	Memory Monitor	5
☑	Network Monitor	5
☑	Uptime Monitor	60
☑	Temperature Monitor	5
☑	Who	5

☑ Enable temperature monitoring at the Master Host

Note: Check if you would like the Master Host to monitor the temperatures of compute hosts that use iceboxes.

6. (Optional) Open and edit the listener as needed.

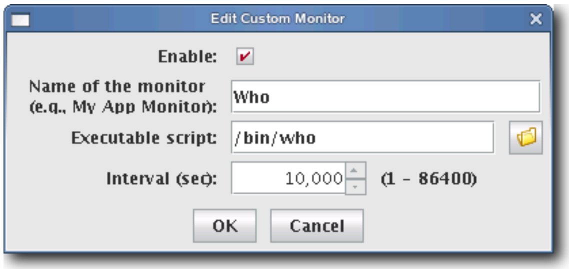
Edit Custom Monitor
OK
Cancel

Enable: ☑
Name of the monitor (e.g., My App Monitor): Who
Executable script: /bin/who
Interval (sec): 10,000 (1 - 86400)

7. Apply the monitor to the host(s) you want to monitor.

Note

When applying listeners to a host, the image used to provision the host must use a payload that contains Management Center. See *Install Management Center into the Payload* on page 105.

8. (Optional) Apply the monitor to payloads.

9. Click Apply as Default to save the monitor.

Note

When you add a monitor and click Apply as Default, Management Center saves the monitor as one of the default monitors—all future payloads will contain the new monitor. Furthermore, the new monitor will be included any time you install Management Center into a payload.

10. Click Close.

Metrics

Metrics refer to data collected by monitors that is processed and displayed by the Management Center instrumentation service. The types of metrics collected are feature-specific and Management Center allows you to view metrics for an individual host or group of hosts. For a list of available metrics, see *Pre-configured Metrics* on page 265.

Note

Before you can display a custom metric, you must define a custom monitor to collect the data. See *Custom Monitors* on page 181.

To Display Custom Metrics

1. Select the Instrumentation tab.
2. Select the host(s) for which you want to display metrics in the host navigation tree.
3. Select the List tab.

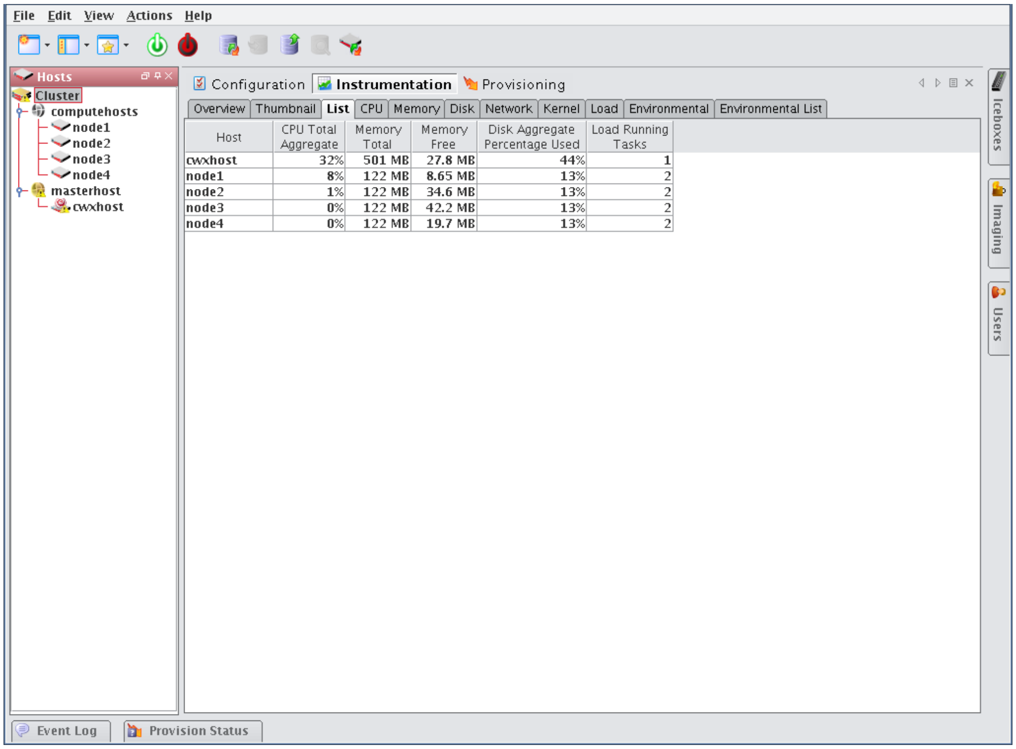
File
Edit
View
Actions
Help
Hosts
Cluster
computehosts
node1
node2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
Overview
Thumbnail
List
CPU
Memory
Disk
Network
Kernel
Load
Environmental
Environmental List
Iceboxes
Imaging
Users
Event Log
Provision Status

Host	CPU Total Aggregate	Memory Total	Memory Free	Disk Aggregate Percentage Used	Load Running Tasks
cwxhost	32%	501 MB	27.8 MB	44%	1
node1	8%	122 MB	8.65 MB	13%	2
node2	1%	122 MB	34.6 MB	13%	2
node3	0%	122 MB	42.2 MB	13%	2
node4	0%	122 MB	19.7 MB	13%	2

4. Select Metrics from the Edit menu. The Metric Selector appears.

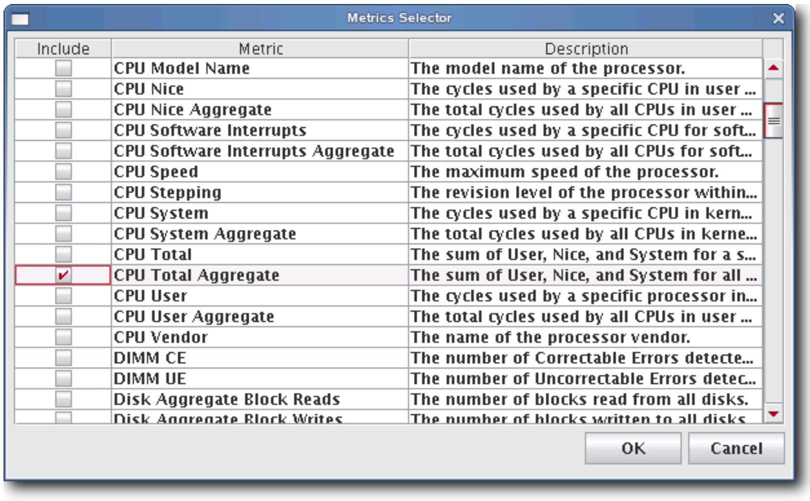

Include	Metric	Description
☐	CPU Model Name	The model name of the processor.
☐	CPU Nice	The cycles used by a specific CPU in user ...
☐	CPU Nice Aggregate	The total cycles used by all CPUs in user ...
☐	CPU Software Interrupts	The cycles used by a specific CPU for soft...
☐	CPU Software Interrupts Aggregate	The total cycles used by all CPUs for soft...
☐	CPU Speed	The maximum speed of the processor.
☐	CPU Stepping	The revision level of the processor within...
☐	CPU System	The cycles used by a specific CPU in kern...
☐	CPU System Aggregate	The total cycles used by all CPUs in kerne...
☐	CPU Total	The sum of User, Nice, and System for a s...
☑	CPU Total Aggregate	The sum of User, Nice, and System for all ...
☐	CPU User	The cycles used by a specific processor in...
☐	CPU User Aggregate	The total cycles used by all CPUs in user ...
☐	CPU Vendor	The name of the processor vendor.
☐	DIMM CE	The number of Correctable Errors detecte...
☐	DIMM UE	The number of Uncorrectable Errors detec...
☐	Disk Aggregate Block Reads	The number of blocks read from all disks.
☐	Disk Aggregate Block Writes	The number of blocks written to all disks.

5. Select the metrics you want to include, then click OK. The metrics appear in the List tab.

Metrics Selector

The Metrics Selector reads from *Metrics.profile* in the *$MGR_HOME/etc* directory on each Management Center client. You may add custom metrics to this profile by making additions in the proper file format:

```
hosts.<name>.<key>.label:<metric_title>
hosts.<name>.<key>.description:<description>
hosts.<name>.<key>.type:java.lang.<type>
hosts.<name>.<key>.pattern:<pattern>
```

- The <*name*> is the host name.
- The <*metric_title*> is the title displayed in the Management Center list monitoring view and in the metric selector dialog.
- The <*description*> indicates what the monitor does and appears in the metric selector dialog.
- The <*type*> is either “Number” or “String.” Numbers are right justified and Strings are left-justified in the Management Center list view.
- The <*pattern*> helps set the column width for the Management Center list monitoring view. The column width should reflect the number of characters typically returned by the value. If the returned value has 10-12 characters, the pattern would be 12 zeros (000000000000). For example, if the returned value is a percent, the pattern should be “100%” or 4 zeros (0000).

CUSTOM METRICS EXAMPLE

Continuing with the example introduced in *To Add a Custom Monitor* on page 182, add the following to the *Metrics.profile* on the Management Center client—then restart the client:

```
hosts.who.count.label=Who Count
hosts.who.count.description=Number of users logged in.
hosts.who.count.type=java.lang.Number
hosts.who.count.pattern=00
hosts.who.who.label=Who's On
hosts.who.who.description=Who's logged in.
hosts.who.who.type=java.lang.String
hosts.who.who.pattern=0000000
```

The new metrics appear in the Metrics Selector dialog.

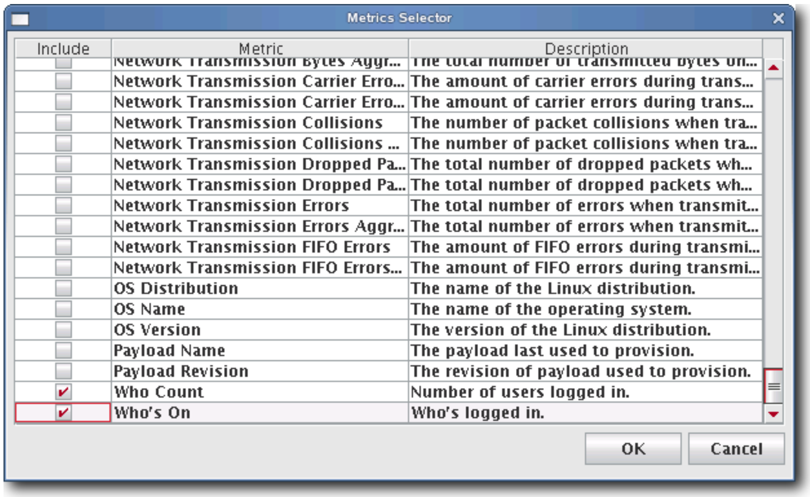

Include	Metric	Description
☐	Network Transmission Bytes Aggr...	The total number of transmitted bytes on...
☐	Network Transmission Carrier Erro...	The amount of carrier errors during trans...
☐	Network Transmission Carrier Erro...	The amount of carrier errors during trans...
☐	Network Transmission Collisions	The number of packet collisions when tra...
☐	Network Transmission Collisions ...	The number of packet collisions when tra...
☐	Network Transmission Dropped Pa...	The total number of dropped packets wh...
☐	Network Transmission Dropped Pa...	The total number of dropped packets wh...
☐	Network Transmission Errors	The total number of errors when transmit...
☐	Network Transmission Errors Aggr...	The total number of errors when transmit...
☐	Network Transmission FIFO Errors	The amount of FIFO errors during transmi...
☐	Network Transmission FIFO Errors...	The amount of FIFO errors during transmi...
☐	OS Distribution	The name of the Linux distribution.
☐	OS Name	The name of the operating system.
☐	OS Version	The version of the Linux distribution.
☐	Payload Name	The payload last used to provision.
☐	Payload Revision	The revision of payload used to provision.
☑	Who Count	Number of users logged in.
☑	Who's On	Who's logged in.

The “who” additions also appear in the Instrumentation List view:

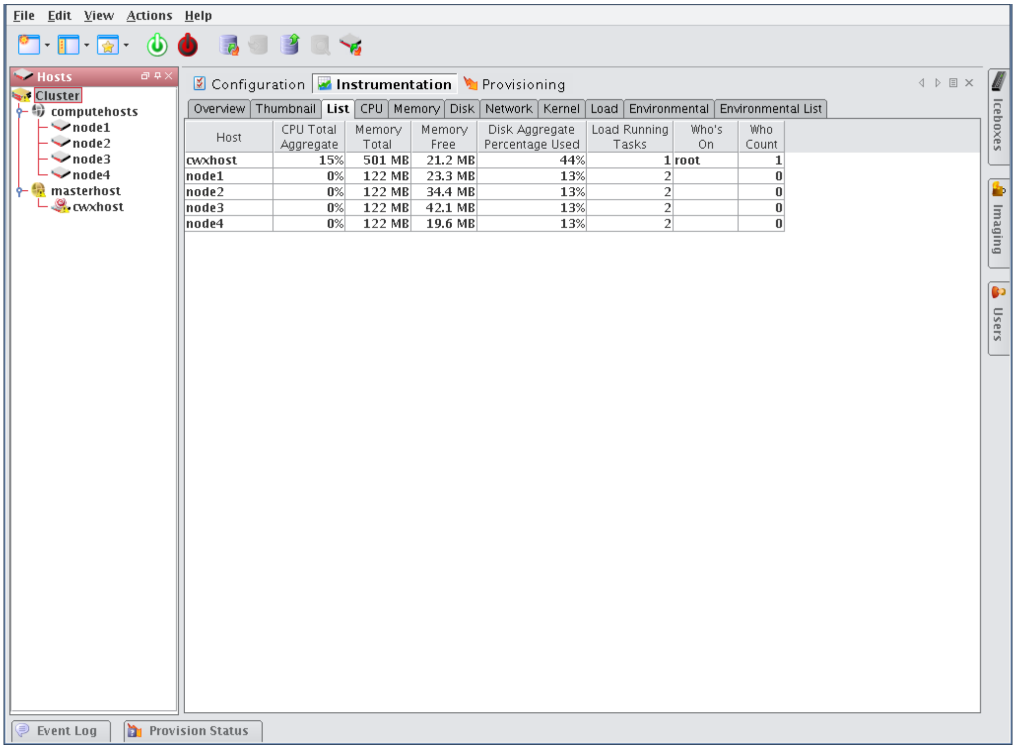
File Edit View Actions Help
Hosts
Cluster
computehosts
node1
node2
node3
node4
masterhost
cwxhost
Configuration
Instrumentation
Provisioning
Overview
Thumbnail
List
CPU
Memory
Disk
Network
Kernel
Load
Environmental
Environmental List
Iceboxes
Imaging
Users
Event Log
Provision Status

Host	CPU Total Aggregate	Memory Total	Memory Free	Disk Aggregate Percentage Used	Load Running Tasks	Who's On	Who Count
cwxhost	15%	501 MB	21.2 MB	44%	1	root	1
node1	0%	122 MB	23.3 MB	13%	2		0
node2	0%	122 MB	34.4 MB	13%	2		0
node3	0%	122 MB	42.1 MB	13%	2		0
node4	0%	122 MB	19.6 MB	13%	2		0

Event Listeners

Event Listeners allow you to easily monitor your cluster and trigger events (loggers) when you exceed specific thresholds. Event listeners may be configured on specific hosts (including the Master Host) and included on payloads that contain Management Center (see *Install Management Center into the Payload* on page 105). By default, Management Center includes a basic collection of listeners, but allows you to add custom listeners as needed. You may also import listeners from an existing host or payload, import the default listeners, or restore the factory settings. The following table lists the default listeners:

Listener Name	Threshold	Message
System Swap Space	512000000	Master Host is using swap space.
Memory (EDAC) (Correctable Errors)	500	Memory Error Detection and Correction (EDAC—AKA BlueSmoke) detected {2} correctable memory error(s) on host {3}.
Memory (EDAC) (Uncorrectable Errors)	1	Memory Error Detection and Correction (EDAC—AKA BlueSmoke) detected {2} correctable memory error(s) on host {3}.
LinuxBIOS Bootmode	0	Management Center has detected that LinuxBIOS is running in Fallbackmode. This may indicate an error with BIOS settings. As a result, this host may not be running at full performance.
System Load	2.1	Five minute load average limit {0} exceeded on host {3} (current load average {2})
Host Power Status		The following host(s) have stopped responding: {0}. The following host(s) are still not responding - {1}. This may be due to the host(s) failing, network congestion, or Management Center services being stopped.
Ambient Temperature (Warning)	55	Ambient Temperature limit {0} exceeded on host {3} (current temperature {2}).
Ambient Temperature (Error)		Ambient Temperature limit {0} exceeded on host {3} (current temperature {2}). Shutting down.
CPU Temperature (Warning)		CPU Temperature 1 limit {0} exceeded on host {3} (current temperature {2}).
CPU Temperature (Error)		CPU Temperature 1 limit {0} exceeded on host {3} (current temperature {2}).

Tip

The temperature listener is divided into a CPU temperature listener and an ambient temperature listener. The CPU temperature listener is triggered by any CPU and the CPU that trips it is specified in the message. By separating the ambient temperature, Management Center supports a negative threshold for PEKI temperatures and a positive threshold for ambient temperatures.

To Enable or Disable a Listener

1. Select Event Administration from the Edit menu. The Event Administration dialog appears.

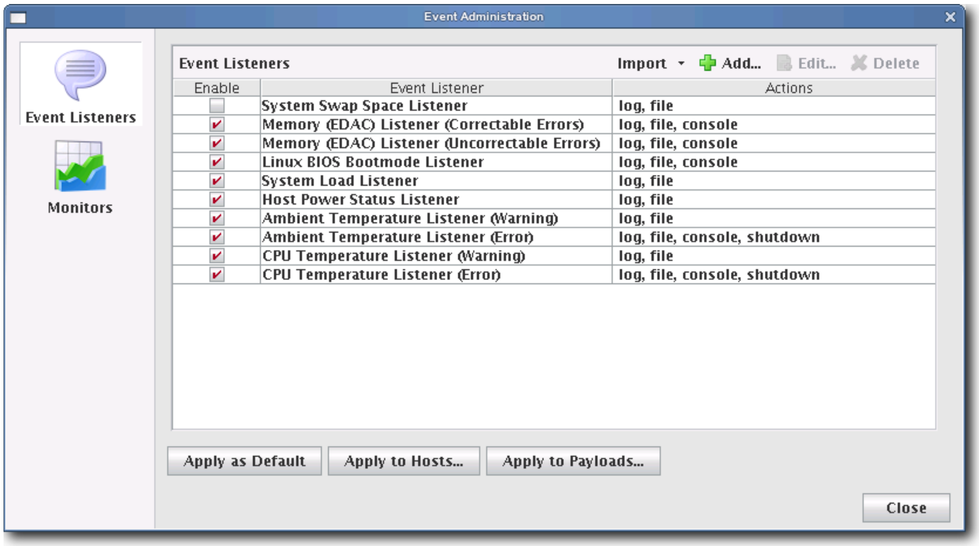
Event Administration
Event Listeners
Monitors
Event Listeners
Import
Add...
Edit...
Delete
Apply as Default
Apply to Hosts...
Apply to Payloads...
Close

Enable	Event Listener	Actions
☐	System Swap Space Listener	log, file
☑	Memory (EDAC) Listener (Correctable Errors)	log, file, console
☑	Memory (EDAC) Listener (Uncorrectable Errors)	log, file, console
☑	Linux BIOS Bootmode Listener	log, file, console
☑	System Load Listener	log, file
☑	Host Power Status Listener	log, file
☑	Ambient Temperature Listener (Warning)	log, file
☑	Ambient Temperature Listener (Error)	log, file, console, shutdown
☑	CPU Temperature Listener (Warning)	log, file
☑	CPU Temperature Listener (Error)	log, file, console, shutdown

2. Select Event Listeners.
3. Check or un-check the box next to each listener you want to enable or disable.
4. (Optional) Click Apply as Default to apply the listener configuration as the default on the Master Host and payload. Management Center saves the listeners in *InstrumentationListeners.profile.default*.
5. (Optional) Click Apply to Hosts to apply the listener to a specific host(s). The Export to Hosts dialog appears.

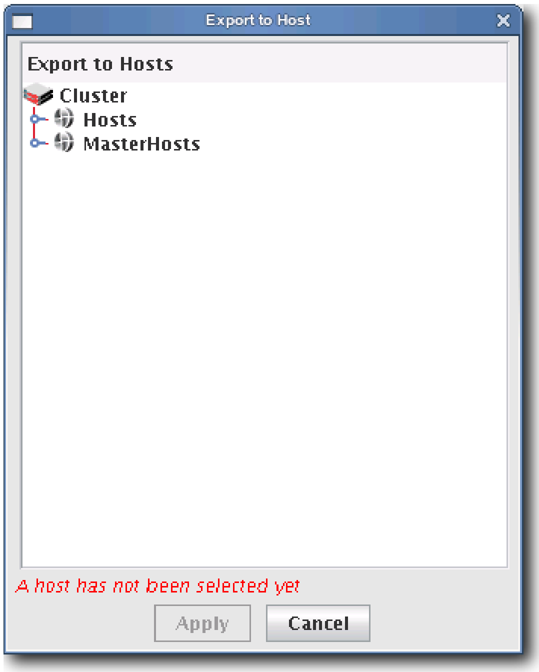
Export to Host
Export to Hosts
Cluster
Hosts
MasterHosts
Apply
Cancel

A host has not been selected yet

   A. Select the host(s) to which to export the listeners from the navigation tree.

  B. Click Apply to save changes or click Close to abort this action.

6. (Optional) Click Apply to Payloads to include these listeners as part of a payload. The Export to Payloads appears.

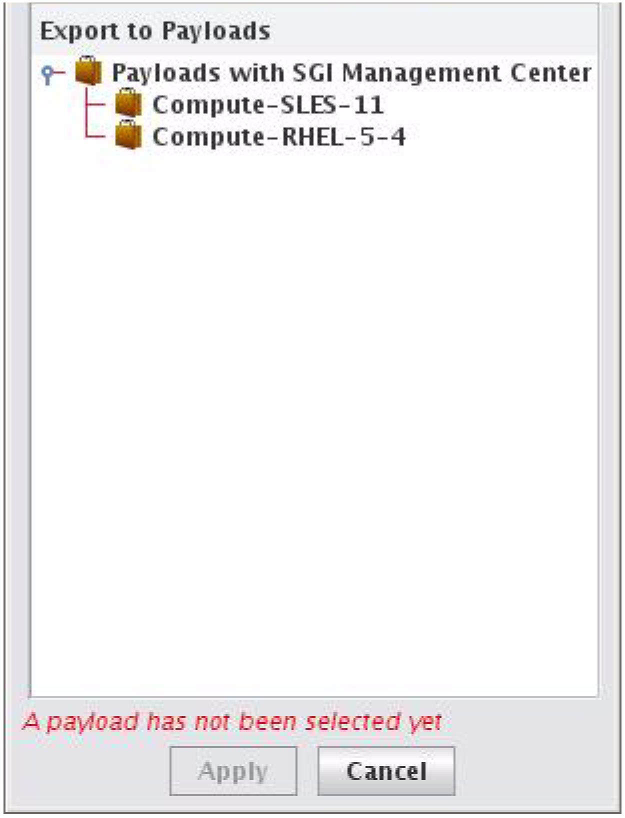
Export to Payloads
Payloads with SGI Management Center
Compute-SLES-11
Compute-RHEL-5-4
Apply
Cancel

A payload has not been selected yet

  A. Select the payload(s) to which to apply the listeners.
  B. Click Apply to save changes or click Cancel to abort this action.

7. Click Close to complete this action and close the Event Administration dialog.

Note

If you click close without applying your changes, all modifications will be lost.

To Add a Listener

1. Select Event Administration from the Edit menu. The Event Administration dialog appears.
2. Select Event Listeners.
3. Click Add. The Add Listener Metric dialog appears.

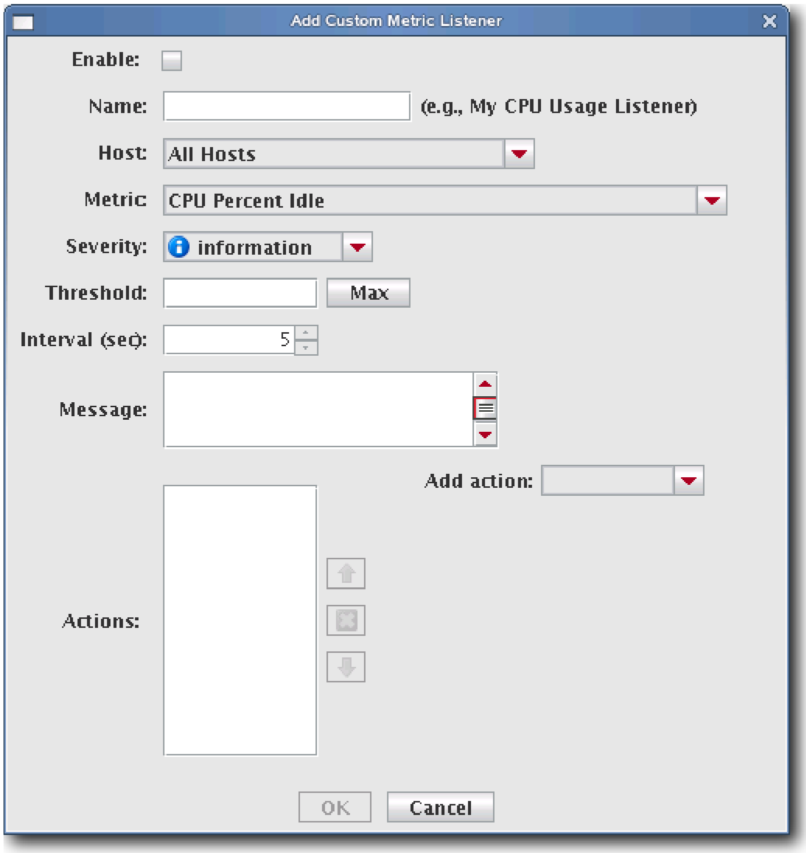

Add Custom Metric Listener

Enable: ☐

Name: (e.g., My CPU Usage Listener)

Host: All Hosts

Metric: CPU Percent Idle

Severity: information

Threshold:

Interval (sec): 5

Message:

Add action:

Actions:

4. Enter the name of the listener.
5. Select the host(s) on which to enable the listener.
6. Select the metric to monitor. For a list of available metrics, see *Pre-configured Metrics* on page 265.

Note

If you write a custom monitor and want to use one or more of the metrics from that monitor, you must edit the *CustomMetrics.profile* to include the metrics, then restart Management Center—otherwise, no custom listeners will be defined. *CustomMetrics.profile* uses the same format as *Metrics.profile*, discussed in *Metrics Selector* on page 187.

7. Specify the severity level of the event (Information, Warning, Error).
8. Enter the threshold for the metric and click the Max/Min button to specify whether this value is the maximum of minimum threshold.
9. Enter the monitoring interval (in seconds).
10. Enter a message to display with this listener.

Note

The *message* is user-configurable and contains the content of the log message or e-mail message. Several variables are available in the message:
{0} = Threshold
{1} = Metric Name
{2} = Metric Value at the time the listener was triggered
{3} = Hostname

11. Add actions to perform if this event is triggered. Available actions are listed in the following table:

Action	Description
email	Send an event notification e-mail to a comma-delimited list of recipients.
script	Executes a user-selected script when triggered.
snmp	Sends SNMP messages to a user-specified trap host.
beacon	Turns the beacon on for the host.
console	Sends event information to the console.
file	Sends event information to *$MGR_HOME/log/event.log*
halt	Halts the host on which HostAdministrationService is running (user-specified).
log	Displays event information in the Event Log GUI.
pbsoff	Automatically set the host status to offline. The pbsoff action requires some additional configuration. See *PBS Configuration* on page 194.
powercycle	Cycles power to the host.
poweron	Powers the host on.
poweroff	Powers the host off.
reboot	Soft reboots the host.
shutdown	Shut down the host.
syslog	Sends an event message to the syslog.

Tip

The Actions list allows you to configure the order in which actions should occur. You may also click Delete to remove an action from the list.

12. Check the Enable option to activate the listener.

13. Click OK to continue or click Cancel to abort this action.

14. Click Apply as Default to save the listener.

Note

When you add a listener and click Apply as Default, Management Center saves the listener as one of the default listeners—all future payloads will contain the new listener. Furthermore, the new listener will be included any time you install Management Center into a payload.

15. Click Close.

PBS CONFIGURATION

The pbsoff action uses the pbsnodes command. This command is installed on the hosts as part of the PBS package—however, the PBS server is not typically configured to authenticate from other hosts in the system. In order for the pbsoff action to be successful, you must allow pbsnodes to run from the hosts. To do this, set the pbs manager via qmgr:

```
qmgr -c “set server managers = root@*.<cluster>.<domain>.<base>”
```

For example:

```
qmgr -c “set server managers = root@*.engr.mycompany.com”
```

Tip

You can test this configuration by running the following command on one of the hosts:

```
pbsnodes -o <hostname>
```

To Edit a Listener

1. Select Event Administration from the Edit menu. The Event Administration dialog appears.
2. Select Event Listeners.
3. Double-click a listener in the list or select the listener and click Edit. The edit dialog appears.

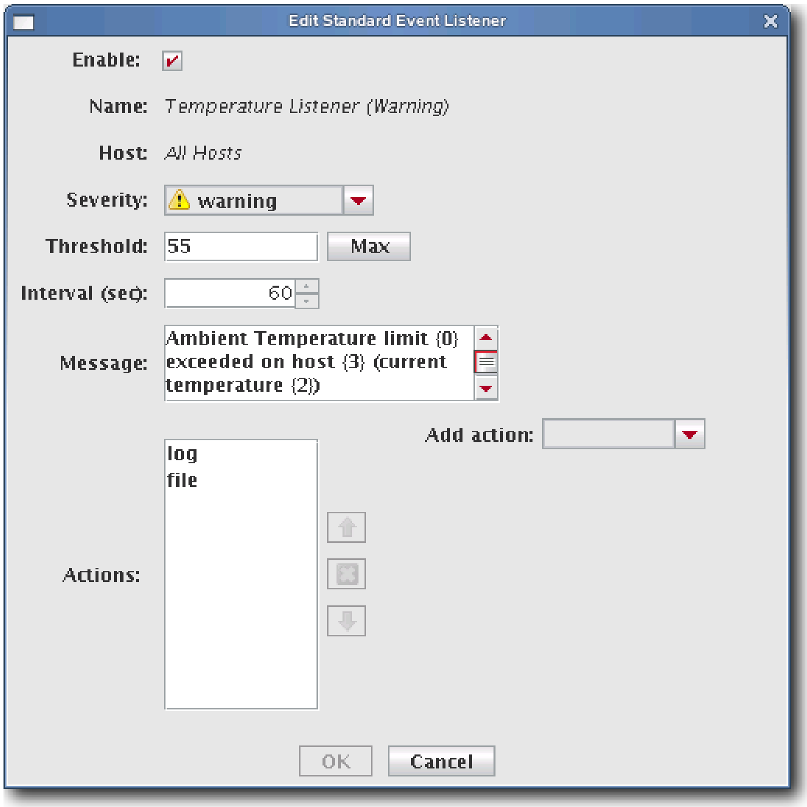

Enable: ☑

Name: Temperature Listener (Warning)

Host: All Hosts

Severity: warning

Threshold: 55

Interval (sec): 60

Message: Ambient Temperature limit {0} exceeded on host {3} (current temperature {2})

Actions:
log
file

4. Make any necessary modifications, then click OK to apply your changes. Click Cancel to abort this action.
5. (Optional) Apply the listener to hosts or payloads.
6. Click Apply as Default to save the listener.

Note

When you change a listener and click Apply as Default, Management Center saves the listener as one of the default listeners—all future payloads will contain the new listener. Furthermore, the new listener will be included any time you install Management Center into a payload.

7. Click Close.

To Delete a Listener

1. Select Event Administration from the Edit menu. The Event Administration dialog appears.
2. Select a listener from the list and click Delete.
3. Management Center asks you to confirm your action.

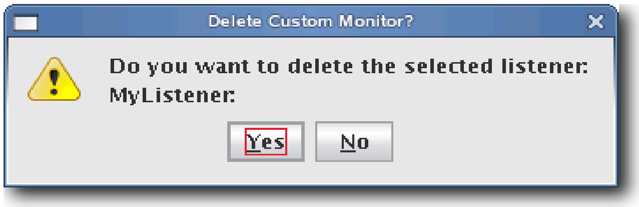
Delete Custom Monitor?
Yes
No

Do you want to delete the selected listener: MyListener:

4. Click Yes to delete the listener or click No to abort this action.

To Import Listeners

IMPORT FROM HOST

1. Select Event Administration from the Edit menu. The Event Administration dialog appears.
2. Select Event Listeners.
3. Click Import and select Import from Host. The Import from Hosts dialog appears.

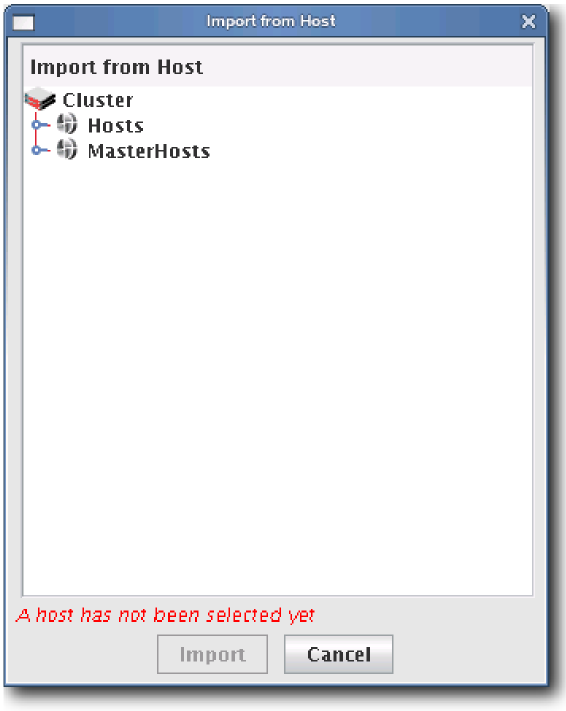
Import from Host
Import from Host
Cluster
Hosts
MasterHosts
Import
Cancel

A host has not been selected yet

4. Select the host from which to import listeners and click Import. Click Cancel to abort this action.

IMPORT FROM PAYLOAD

1. Select Event Administration from the Edit menu. The Event Administration dialog appears.
2. Select Event Listeners.
3. Click Import and select Import from Payload. The Import from Payloads dialog appears.

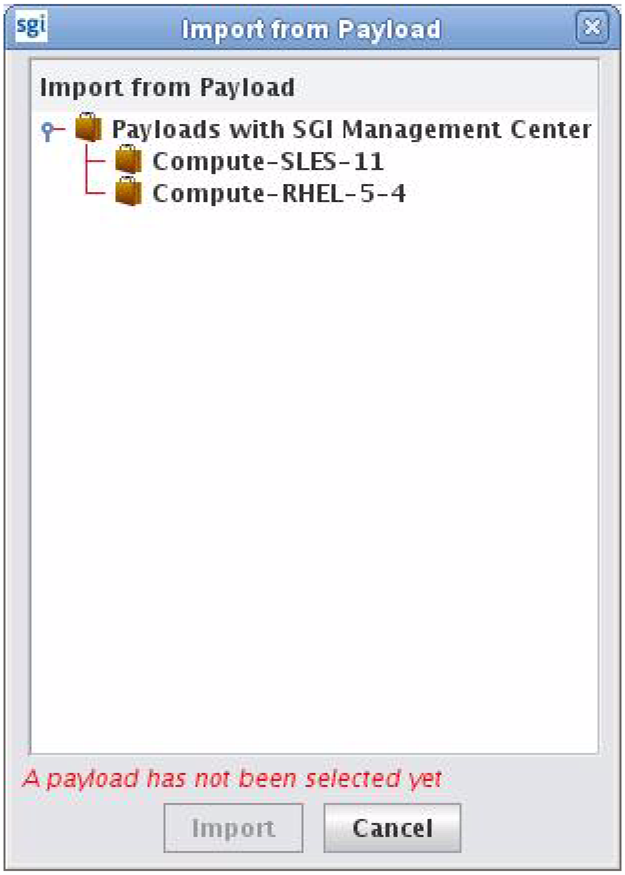
sgi
Import from Payload
Import from Payload
Payloads with SGI Management Center
Compute-SLES-11
Compute-RHEL-5-4

A payload has not been selected yet

Import
Cancel

4. Select the payload from which to import listeners and click Import. Click Cancel to abort this action.

IMPORT DEFAULT

1. Select Event Administration from the Edit menu. The Event Administration dialog appears.
2. Select Event Listeners.
3. Click Import and select Import Default. Management Center restores all listeners stored as default listeners in *InstrumentationListeners.profile.default*. See *To Enable or Disable a Listener* on page 190 for information on adding default listeners.

RESTORE FACTORY SETTINGS

1. Select Event Administration from the Edit menu. The Event Administration dialog appears.
2. Select Event Listeners.
3. Click Import and select Restore Factory Settings. Management Center reverts the default listeners that shipped with Management Center.

Loggers

Loggers refer to actions taken when a threshold exceeds its maximum or minimum value. Common logger events include sending messages to the centralized Management Center event log, logging to a file, logging to the serial console, and shutting down the host.

MANAGEMENT CENTER EVENT LOG

The event log is located on the instrumentation overview screen. If you select multiple hosts (or a container such as a cluster, partition, or region), the log shows messages for any host in the selection. If you select a single host, the message log shows messages for this host only. Messages have three severity levels: error, warning, and informational.

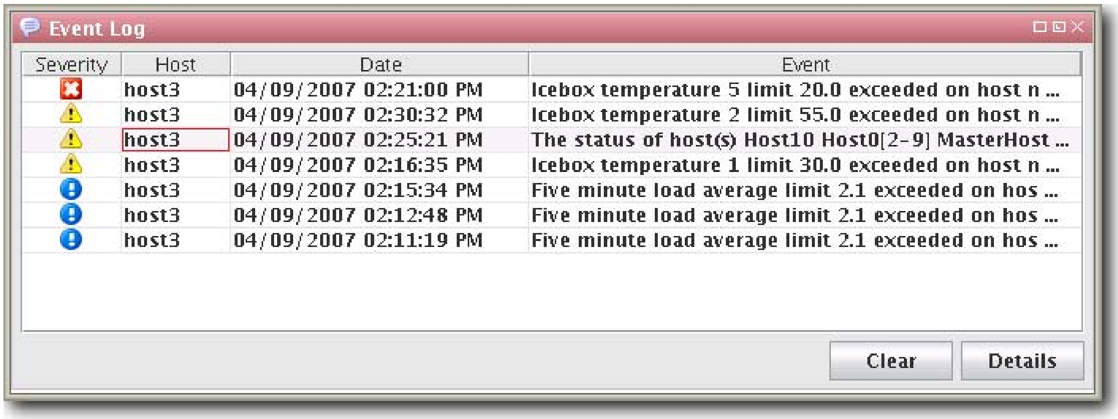

Severity	Host	Date	Event
✖	host3	04/09/2007 02:21:00 PM	Icebox temperature 5 limit 20.0 exceeded on host n ...
⚠	host3	04/09/2007 02:30:32 PM	Icebox temperature 2 limit 55.0 exceeded on host n ...
⚠	host3	04/09/2007 02:25:21 PM	The status of host(s) Host10 Host0[2-9] MasterHost ...
⚠	host3	04/09/2007 02:16:35 PM	Icebox temperature 1 limit 30.0 exceeded on host n ...
ℹ	host3	04/09/2007 02:15:34 PM	Five minute load average limit 2.1 exceeded on hos ...
ℹ	host3	04/09/2007 02:12:48 PM	Five minute load average limit 2.1 exceeded on hos ...
ℹ	host3	04/09/2007 02:11:19 PM	Five minute load average limit 2.1 exceeded on hos ...

TEMPLATEFORMATTER

You may extend the abilities of pre-configured and custom loggers (located in *$MGR_HOME/etc/Logging.profile*) using the template field of the TemplateFormatter. The template field allows you to configure the types of messages displayed by loggers. For example, the message template type used in the following example is %m:

```
formatters.com.lnxi.instrumentation.event: \
    com.xeroone.logging.TemplateFormatter
formatters.com.lnxi.instrumentation.event.template: %m
```

The following table contains a list of supported message templates:

Template	Description
%N	Sequential record number. This number resets each time the virtual machine restarts.
%T	Creation time.
%C	Channel.
%S	Severity.
%M	Message.
%E	Event.
%EN	Event name.
%ET	Event trace.
%AN	Application name.
%AM	Application moniker.
%AST	Application start time.

Template	Description
%AV	Application version.
%HN	Host name.
%HM	Host moniker.
%MS	Memory size.
%MF	Memory free.
%OSN	Operating system name.
%OSV	Operating system version.
%%	Literal % character.
"	Literal ' (single quote) character.
'	Escape character for quoted text.

Chapter 8

Upgrading SGI Management Center

Using the following topics, this chapter describes how you upgrade SGI Management Center:

- *General Tasks* on page 201
- *Upgrading from a Previous Version of SGI Management Center* on page 202
- *Upgrading from SGI ISLE Cluster Manager 2.x* on page 203
- *Additional Upgrade Notes* on page 205

General Tasks

Regardless of the base version, you will need to do the following:

- Back up the product database.
- Upgrade the `sgimc-payload` and `java-1.6.0-sun` packages in all existing payloads. To install the packages into the payload, see *Install Management Center into the Payload* on page 105.
- Update any existing images to use the new payload revision and check in the image.
  Updating the image also refreshes the list of RAMdisk files that need to be copied the next time you provision. If you do not perform this step, you may get `error running cp` messages in the provisioning status panel.
- If you are upgrading from a previous version of SGI Management Center and you are using the Discover feature, upgrade the packages listed in section *Software Requirements* on page 207. Use YUM or Zypper.

The following sections describe upgrading items specific to the base version from which you are upgrading.

! Warning!

If you upgrade from a cluster running SGI Management Center 1.3 or older, the contents of all partitions on a node will be lost the first time you provision.

Upgrading from a Previous Version of SGI Management Center

The following steps outline the best practices for upgrading from a previous version of SGI Management Center.

1. Back up the database as a safeguard.

   To do this, issue the following command as `root`:
   ```
   # dbix -x > <filename>
   ```
   For example:
   ```
   # dbix -x > /root/SMC1.0-backup.dbix
   ```
2. Set up the desired software repositories (See *Pre-Installation Steps* on page 5.) and then install the updated packages.

   Example (RHEL):

   ```
   # yum install sgimc db48 libdb_java-4_8 ice ice-java ice-libs mkelfImage java-1.6.0-sun java-1.6.0-sun-fonts
   ```

   Example (SLES):

   ```
   # zypper install sgimc db48 libdb_java-4_8 ice ice-java ice-libs mkelfImage java-1.6.0-sun java-1.6.0-sun-fonts
   ```

Note

Ensure that you have installed `shout` if on an SGI Altix UV system and `sgi-cm-agnostic` if you are using Altair PBS Professional Dynamic Provisioning.

After installing the new `shout` package, you should configure the shout daemon to be enabled on boot and restart the `shout` daemon:

```
# chkconfig shout on
# /etc/init.d/shout restart
```

3. Copy over custom settings, if applicable.

   Upon upgrade, SGI Management Center backs up all the previous files in */opt/sgi/sgimc/etc/<datestamp>* with a symbolic link called */opt/sgi/sgimc/etc.bak* that points to the latest backup. You may copy custom settings that were present in the previous SGI Management Center installation. For example, if you had custom settings applied for Instrumentation and Event listeners, you can copy to */opt/sgi/sgimc/etc* the settings from the following files:

   /opt/sgi/sgimc/etc.bak/InstrumentationMonitors.profile
   /opt/sgi/sgimc/etc.bak/InstrumentationListeners.profile

   You can perform a `diff` to see what has changed. For example:

   ```
   # diff -r /opt/sgi/sgimc/etc /opt/sgi/sgimc/etc.bak
   ```
4. Restart the services:

   ```
   # /etc/init.d/mgr restart
   ```

Upgrading from SGI ISLE Cluster Manager 2.x

If you are upgrading from SGI ISLE Cluster Manager 2.x to SGI Management Center, you should use the following process:

1. Back up the dbix database and the settings directory.

   For example:

   ```
   # dbix -x > /data1/islecm_backup.dbix
   # cp -a /opt/sgi/islecm/etc /data1/islecm_etc_backup
   ```

   You should check that there is text in the dbix backup file in order to confirm that the dbix backup succeeded.

2. Set up the desired software repositories (See *Pre-Installation Steps* on page 5.) and then install the updated packages.

   Example (RHEL):

   ```
   # yum install sgimc db48 libdb_java-4_8 ice ice-java ice-libs mkelfImage java-1.6.0-
   sun java-1.6.0-sun-fonts
   ```

   Example (SLES):

   ```
   # zypper install sgimc db48 libdb_java-4_8 ice ice-java ice-libs mkelfImage java-
   1.6.0-sun java-1.6.0-sun-fonts
   ```

   The appropriate packages will be replaced.

Note

Ensure that you have installed `shout` if on an SGI Altix UV system and `sgi-cm-agnostic` if you are using Altair PBS Professional Dynamic Provisioning.

After installing the new `shout` package, you should configure the `shout` daemon to be enabled on boot and restart the `shout` daemon:

```
# chkconfig shout on
# /etc/init.d/shout restart
```

3. Log out and back in to source the environment.

4. Restart the services in the new environment:

   ```
   # /etc/init.d/mgr restart
   ```

5. Move over the imaging and vcs contents to their new home.

   Note:The following steps assume the destination directories are empty or non-existent. Make sure there are no files or directories in the destination with the same name as the files being moved.

   ```
   # mkdir -p /opt/sgi/sgimc/vcs
   # mkdir -p /opt/sgi/sgimc/imaging/root/payloads
   # mkdir -p /opt/sgi/sgimc/imaging/root/kernels
   # mkdir -p /opt/sgi/sgimc/imaging/root/images
   # mv /opt/sgi/islecm/vcs/* /opt/sgi/sgimc/vcs/

   # mv /opt/sgi/islecm/imaging/root/payloads/* /opt/sgi/sgimc/imaging/root/payloads/
   # mv /opt/sgi/islecm/imaging/root/kernels/* /opt/sgi/sgimc/imaging/root/kernels/
   # mv /opt/sgi/islecm/imaging/root/images/* /opt/sgi/sgimc/imaging/root/images/
   ```
6. Copy over custom settings if applicable.

   You may copy custom settings that were present in the ISLE Cluster Manager installation. The files are kept in directory */opt/sgi/islecm/etc*. For example, if you had custom settings applied for Instrumentation and Event listeners, you can copy to */opt/sgi/sgimc/etc* the settings from the following files*:*

   /data1/islecm_etc_backup/InstrumentationMonitors.profile
   /data1/islecm_etc_backup/InstrumentationListeners.profile

   You can perform a `diff` to see what has changed. For example:

   ```
   # diff -r /data1/islecm_etc_backup /opt/sgi/sgimc/etc
   ```

   NOTE: Most of the files end in *.profile* and do not contain static path information. If you copy a custom version of *exclude.files*, change all references from */opt/sgi/islecm* to */opt/sgi/sgimc*. Example:

   ```
   # cd /opt/sgi/sgimc/etc
   # sed -i "s/opt\/sgi\/islecm/opt\/sgi\/sgimc/g" exclude.files
   ```
7. Restore the dbix database:

   ```
   # dbix -i < /data1/islecm_backup/islecm_backup.dbix
   ```
8. Restart the services:

   ```
   # /etc/init.d/mgr restart
   ```

Note

You can safely remove the `jdk` package (for example, `jdk-1.5.0_17-fcs`) if it was installed for SGI ISLE Cluster Manager and you do not need it otherwise. Any leftover `islecm-java` packages from previous versions of SGI ISLE Cluster Manager may also be removed. Similarly, you can remove the `db46` package if it was installed with a previous version of SGI Management Center and is no longer being used otherwise.

Additional Upgrade Notes

In addition to the tasks described in section *General Tasks* on page 201, you must perform special tasks when you upgrade from some earlier base versions. This section describes these tasks.

Backing Up Provisioned Nodes

Note

Applicable if upgrading from base versions prior to SGI Management Center 1.4.

Back up any critical data that may be on any of the provisioned nodes. The first time you check in an existing image and use it to provision, all partitions will be re-formatted due to a re-alignment of the partition table(s). Thereafter, node boot-up and provisioning will examine the disks on each boot in order to determine if a re-provision is necessary.

Changing the Default RNA Host Name

Note

Applicable if upgrading from base versions prior to SGI Management Center 1.4.

Older versions of SGI Management Center used the default host name `host`. Starting with SGI Management Center 1.4, the default RNA host name is `admin`. Ensure that an alias exists for `admin` in file */etc/hosts* on the master host and in all payloads.

Example:

```
10.0.10.1 admin.default.domain admin loghost
```

Deleting Old Administration Node from Host Tree

Note

Applicable if upgrading from base versions prior to SGI Management Center 1.4.

Delete the old administration node entry from the Management Center host tree if that node is present from the earlier release. For example, you should delete any old administration node entries from the host tree named `host` or `cwxhost`. To do this, right-click on the appropriate host entity and click **Delete**.

After deleting the old entry, run the Configurator to add the `admin` node entry to the host tree:

1. In the **File** pull-down menu, click **Configure SGI Management Center**.
2. Select the appropriate Ethernet interface (for example, `10.0.10.1` / `eth1`)
3. Check the box **Add SGI Management Center Server to the Host tree**.

Chapter 9

Using the Discover Interface

When you add new nodes to your cluster, you must provide profile information about the new nodes to SGI Management Center—notably, information like the the MAC addresses of the compute nodes and of the BMCs in the new compute nodes. SGI Management Center provides the Discover interface to assist you in determining (*discovering*) the pertinent MAC addresses and adding the new nodes to your cluster.

There is both a graphical and command-line interface for Discover. This chapter describes how you can use each to discover compute nodes.

Note

This does not pertain to SGI Altix UV large-memory platforms. For those platforms, SGI Management Center uses the system management node (SMN) and its associated SMN software bundle to discover its chassis management controllers (CMCs) and blades.

Software Requirements

In order to use Discover, a premium-licensed feature, you need to install the following packages:

- discover
- discover-common
- discover-server
- sgi-common-python
- sgi-management-device
- cattr

After installing these packages, you must also start the Discover daemon:

```
# service discoverd start
```

The Graphical Interface

Use the following procedure to add new nodes:

1. Select **Discover** from the **File** menu.
2. In the initial window, press **Add** to add the new nodes.
3. After the addition of the new nodes, press **Start**.

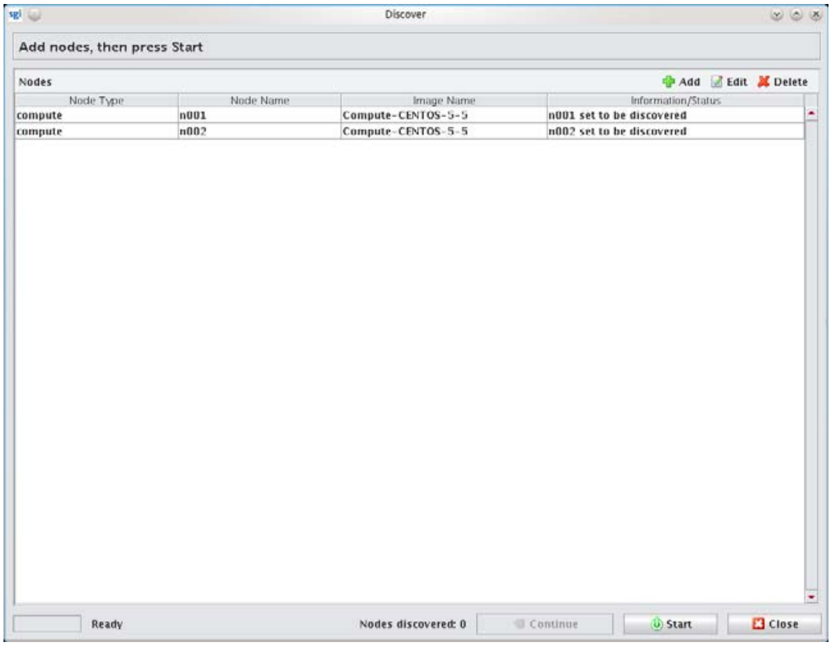
Discover
Add nodes, then press Start
Nodes
Add
Edit
Delete
Ready
Continue
Start
Close

Node Type	Node Name	Image Name	Information/Status
compute	n001	Compute-CENTOS-5-5	n001 set to be discovered
compute	n002	Compute-CENTOS-5-5	n002 set to be discovered

Nodes discovered: 0

4. Apply power to the nodes but do not turn the nodes on.

The system searches for MAC addresses and determines if they belong to a BMC, Ethernet switch, or unknown device. The **Continue** button is enabled as soon as at least one BMC is found.

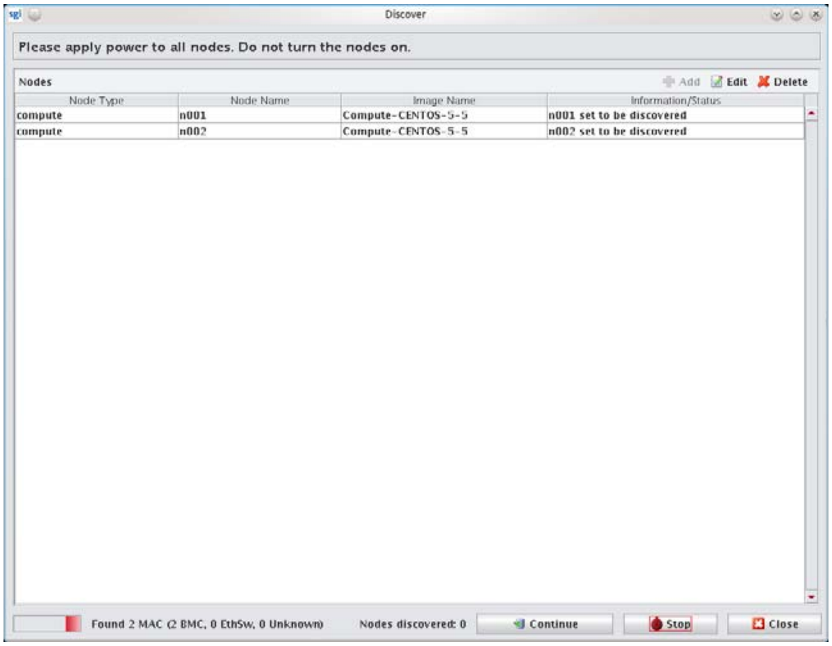
Discover
Nodes
Add
Edit
Delete
Continue
Stop
Close

Please apply power to all nodes. Do not turn the nodes on.

Node Type	Node Name	Image Name	Information/Status
compute	n001	Compute-CENTOS-5-5	n001 set to be discovered
compute	n002	Compute-CENTOS-5-5	n002 set to be discovered

Found 2 MAC (2 BMC, 0 EthSw, 0 Unknown)

Nodes discovered: 0

5. Once Discover finds the MACs for the BMCs, press **Continue**.

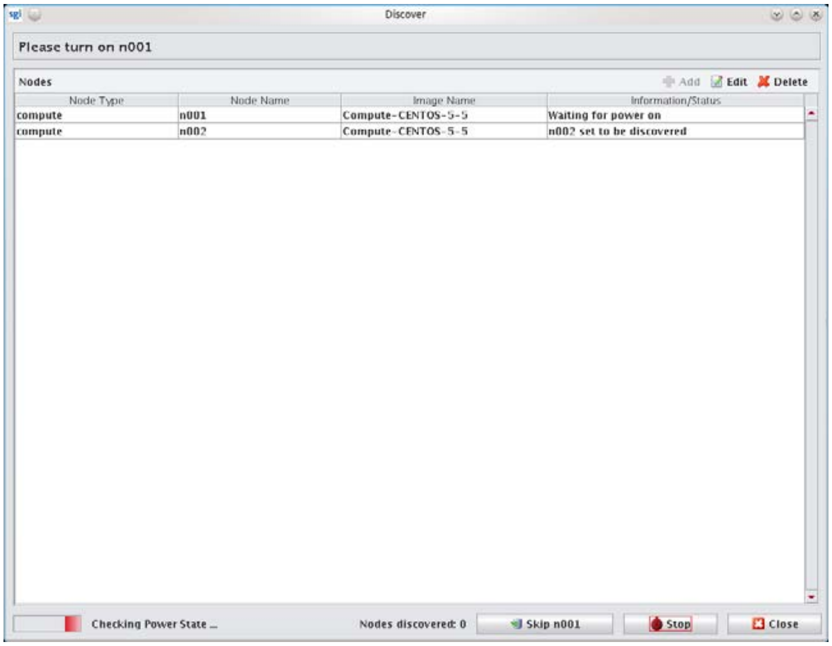
Discover
Nodes
Add
Edit
Delete
Skip n001
Stop
Close

Please turn on n001

Node Type	Node Name	Image Name	Information/Status
compute	n001	Compute-CENTOS-5-5	Waiting for power on
compute	n002	Compute-CENTOS-5-5	n002 set to be discovered

Checking Power State ...

Nodes discovered: 0

6. Power on the designated nodes as requested by Discover.

As Discover detects the BMCs, it associates the MACs with the BMCs in order.

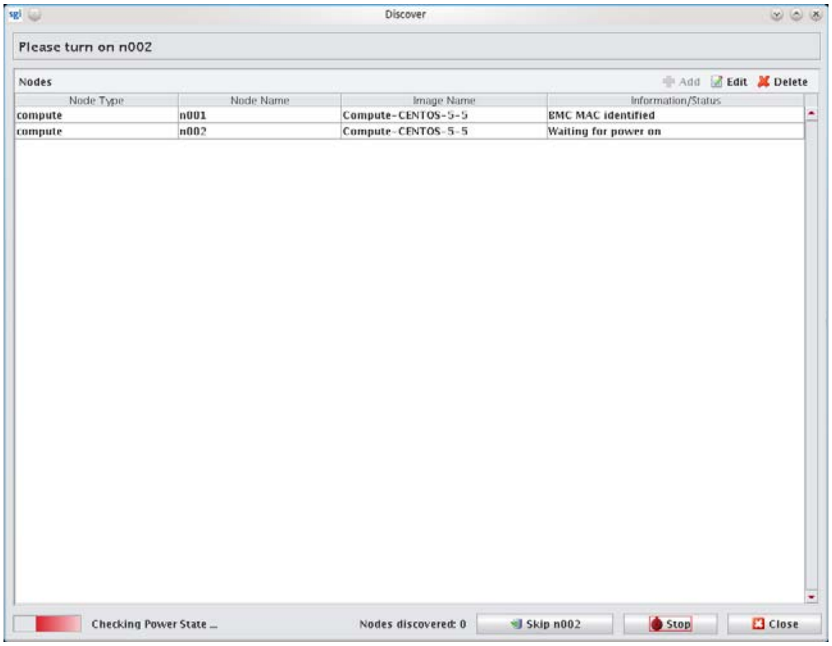
Discover
Nodes
Add
Edit
Delete
Skip n002
Stop
Close

Please turn on n002

Node Type	Node Name	Image Name	Information/Status
compute	n001	Compute-CENTOS-5-5	BMC MAC identified
compute	n002	Compute-CENTOS-5-5	Waiting for power on

Checking Power State ...

Nodes discovered: 0

Once that the BMCs have been ordered, the Discover will determine the system MAC for each node.

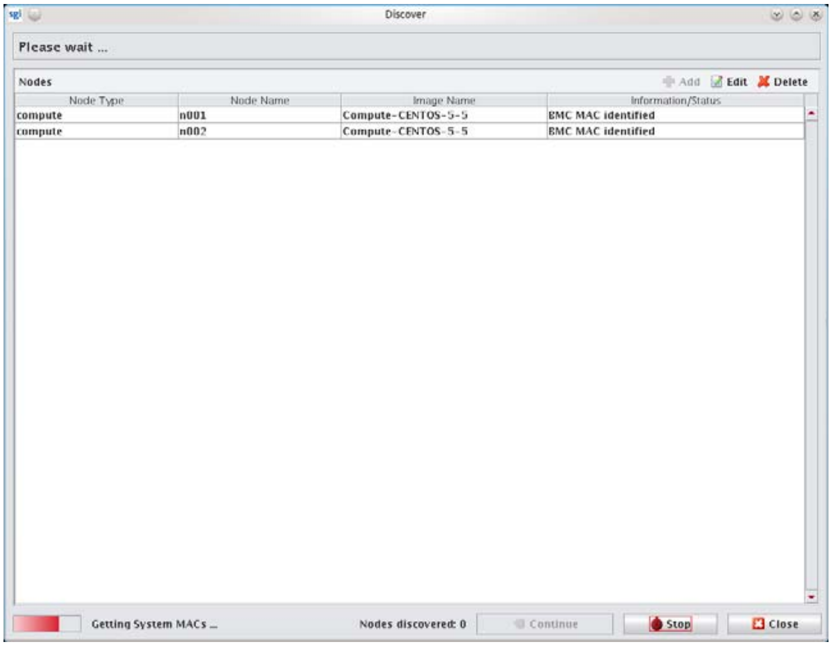
Discover
Nodes
Add
Edit
Delete
Continue
Stop
Close

Please wait ...

Node Type	Node Name	Image Name	Information/Status
compute	n001	Compute-CENTOS-5-5	BMC MAC identified
compute	n002	Compute-CENTOS-5-5	BMC MAC identified

Getting System MACs ...

Nodes discovered: 0

After discovering the system MACs, the nodes will be added to the hosts tree.

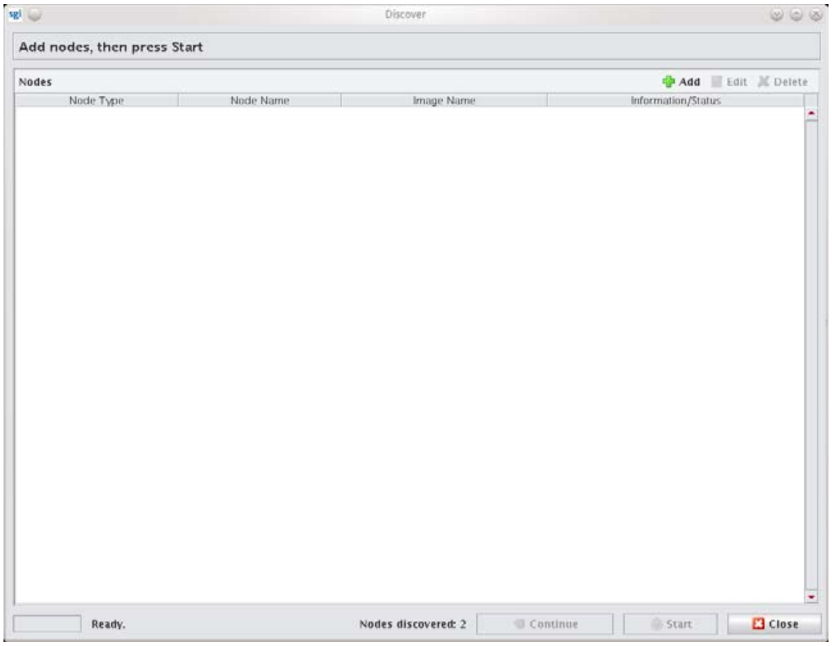

Discover
Nodes
Add
Edit
Delete
Node Type
Node Name
Image Name
Information/Status
Continue
Start
Close

Add nodes, then press Start

Ready.

Nodes discovered: 2

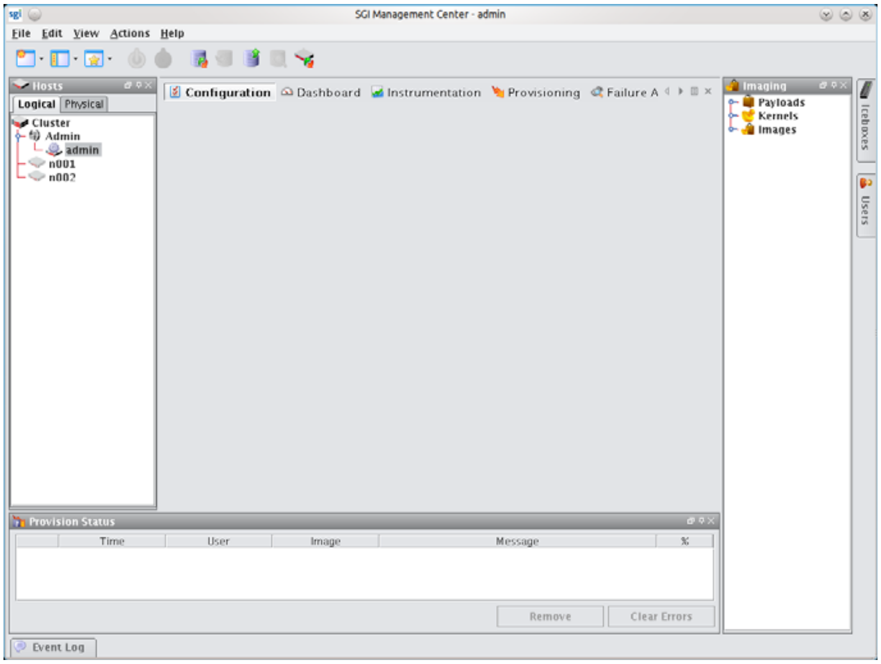

SGI Management Center - admin
File
Edit
View
Actions
Help
Hosts
Logical
Physical
Cluster
Admin
admin
n001
n002
Configuration
Dashboard
Instrumentation
Provisioning
Failure A
Imaging
Payloads
Kernels
Images
Iceboxes
Users
Provision Status
Time
User
Image
Message
%
Remove
Clear Errors
Event Log

The Command-Line Interface

The Discover command-line interface is somewhat parallel to that of the graphical interface. The following example illustrates the process:

```
# discover --compute 0 --compute 1
```

First, you will be asked to apply power to all nodes:

```
*** Please apply power to all nodes. Do not turn the nodes on.
```

Discover will then show you the BMCs as they are detected:

```
Found BMC: 00:30:48:32:1b:89 (1 remaining)
Found BMC: 00:30:48:94:1a:a3 (0 remaining)
```

Once all BMCs are found, Discover verifies that they are all powered off. If nodes are on, you will be asked to power them off. Once all nodes are in a powered off state you will see the following:

```
Verification that all nodes are currently powered off was successful.
```

Next, Discover determines the order of the BMCs by having you power on each node in order:

```
 *** Please turn ON: compute0

 compute0’s BMC is 00:30:48:94:1a:a3

 *** Please turn ON: compute1

 compute1’s BMC is 00:30:48:32:1b:89
```

After ordering is done, Discover will finish without the need for any further interaction. During this process, discovery determines the MAC addresses for the systems themselves:

```
Starting fully automated portion of discovery.

compute0 bmc:00:30:48:94:1a:a3 system:00:30:48:34:2a:b0 (1 remaining)
compute1 bmc:00:30:48:32:1b:89 system:00:30:48:32:91:9a (0 remaining)
```

After all MACs are collected, the nodes are added and discovery is done:

```
Adding nodes...
Discovery complete.
```

You may now provision the nodes normally.

Chapter 10

Troubleshooting

This chapter describes some troubleshooting steps for various problems that may arise. If you encounter a problem not listed here or the suggested solution does not work, contact SGI Customer Support. See *Product Support* on page xii.

The following topics appear in this chapter:

- *Debug Logs* on page 214
- *Support Information Tool* on page 214
- *Startup Daemon Fails on the Master Host* on page 214
- *Nodes in Provisioning or Unknown State after Provisioning* on page 215
- *Temperatures and Fan Speeds Not Registering* on page 215
- *Inordinately High CPU Usage on Head Node* on page 215
- *Insufficient Number of Provisioning Channels* on page 215
- *Kernel Modules Not Loading on Compute Nodes* on page 216
- *Command-line Boot Parameters Not Honored* on page 216
- *Payload Check-in Error* on page 216
- *Invalid or Expired License Message* on page 217
- *Resource Usage Too High on Head Node* on page 217
- *Altix UV Provisioning Stops While Loading Kernel* on page 218

Debug Logs

When you are encountering problems with SGI Management Center, it is often helpful to turn on debugging. In the Management Center GUI, go to the **Preferences** screen:

Edit —> Preferences

Check the box next to **Enable debugging on the master host**.

Alternatively, you can turn on debugging by modifying file */opt/sgi/sgimc/etc/system-clustermanager.profile* on any master host, client install node, or payload install node. Add the following text to the file:

```
system.logging:com.lnxi.debug
logging.level: DEBUG
```

The following logs are generated:

/opt/sgi/sgimc/log/debug.log
/opt/sgi/sgimc/log/SGIMC-server.log
/tmp/SGIMC-<username>.log

Once you reproduce the problem that is occurring, examine these logs for information about the cause of the problem.

Note

There are often log entries (like warnings and/or exceptions) that are not an indication of an actual problem.

Support Information Tool

To generate a set of support information to send to SGI Product Support, run the following command:

```
# /opt/sgi/sgimc/bin/generate_support_info.sh
```

This script generates a compressed tarball, *sgimc_info_<datestamp>.tgz*, which contains the troubleshooting information. Send this file to SGI Product Support.

Startup Daemon Fails on the Master Host

Symptom

Upon trying to start the Management Center GUI or any one of the command-line utilities, you may see a Java Runtime Exception (`Cannot assign requested address`).

Resolution

Ensure that the RNA host name (default: `admin`) of the master host is present in the in the /etc/hosts file.

Example:

```
10.0.10.1          admin.default.domain          admin loghost
```

Note

This should not be a loopback address such as 127.0.0.1. You can find the RNA host name by checking */opt/sgi/sgimc/@genesis.profile*.

Nodes in Provisioning or Unknown State after Provisioning

Symptom

After provisioning, the nodes remain in the provisioning or unknown state.

Resolution

Make sure that the host name of the node and the host name of the master host appear in the */etc/hosts* file on the node.

Example:

```
10.0.10.1      admin.default.domain      admin loghost
10.0.1.1       n001.default.domain       n001
```

Temperatures and Fan Speeds Not Registering

Symptom

Temperatures and fan speeds are not appearing in the **Environmental** tab under **Instrumentation** for a particular host.

Resolution

This can occur if the IPMI cache file is corrupted or the BMC has been replaced by a different hardware type. You can clear the IPMI cache using the following command:

```
# rm /opt/sgi/sgimc/ipmi/*
```

Inordinately High CPU Usage on Head Node

Symptom

During provisioning, if the payload takes vast amounts of time to download or fails to download completely, you may see the following on the screen with no progress:

```
Downloading      0% [                ]
```

Resolution

The most common cause of this behavior is a misconfiguration of IGMP in the management network switches. Please verify multicast routing is enabled on the switch. In some cases, you may need to enable IGMP Snooping. For some switches, it is required to disable spanning tree protocol or enable RTSP / Edge Routing. Consult your switch documentation for information about how to configure your switch.

Insufficient Number of Provisioning Channels

Symptom

When you attempt to provision, you see a message that indicates there are no more provisioning channels.

Resolution

In the Management Center GUI, do the following:

1. Go to the **Preferences** screen:

   Edit —> Preferences

2. Click **Provisioning**.
3. Raise the number of multicast channels. (The maximum is 30.)

If the preceding steps fail to resolve the issue, you may have too many combinations of payloads/kernels assigned to various nodes in the system. This can be resolved by re-provisioning all nodes with one image, which will free the rest of the available payload/kernel combinations.

Kernel Modules Not Loading on Compute Nodes

Symptom

The expected kernel modules do not load upon booting the compute nodes.

Resolution

You should open the kernel in Management Center GUI and select the desired kernel modules to load. Adding them to */etc/sysconfig/kernel* in the payload and running mkinitrd will not work, since nodes are booted from the network and the list of modules depends on the modules defined in the Management Center kernel. All module dependencies must be added as well.

Command-line Boot Parameters Not Honored

Symptom

Command-line parameters listed in */boot/grub/menu.lst* or */boot/efi/efi/SuSE/elilo.conf* do not take effect.

Resolution

Nodes that are provisioned by Management Center do not boot from grub or EFI. In order to specify command-line parameters, do the following:

1. Open the specified kernel in Management Center GUI.
2. In the **Parameters** frame, select **Advanced**.
3. Enter the desired command-line parameters (for example, `cgroup_disable=memory`).

Payload Check-in Error

Symptom

When attempting to check in a payload, you receive the error message `Error checking in <payload name>`.

Resolution

Try the following:

1. Ensure you have not run out of disk space in the filesystem that contains /opt/sgi/sgimc.
2. Ensure that you do not have any extra mounts in the payload. For example, unmount `sys` or `dev` in */opt/sgi/sgimc/imaging/root/payloads/<name>* .
3. With debugging enabled, re-attempt the check-in and then examine */opt/sgi/sgimc/log/debug.log* for additional information about the error.

Invalid or Expired License Message

Symptom

Even though you have installed a license, you receive the following message:

```
The license is invalid or has expired.
```

Resolution

One or more the following might solve the problem:

- Restart the Management Center daemon:

  ```
  # service mgr restart
  ```

- Inspect file */etc/lk/keys.dat* for corruption or missing characters (for example—missing lines or an unmatched, single quotation mark).
- Ensure that the system time/date is correct.
- Run the following command and examine the output for information about the current state of the license:

  ```
  # lk_verify
  ```

Resource Usage Too High on Head Node

Symptom

The head node is experiencing high load levels from the Management Center application. The DNA process may be showing abnormally high CPU usage.

Resolution

Management Center is actively managing power and instrumentation on the cluster. This can take a significant amount of resources especially on a large cluster. For this reason, do not use the master host as a compute resource. However, if Management Center is persistently experiencing an inordinately high level of resource usage and you have already tried restarting the Management Center daemon, it may help to disable IPMI environmental gathering from the master host. Each host will then monitor its own temperature and fan speeds and send the data to the master host.

To disable environmental monitoring from the master host, do the following:

1. In the Management Center GUI, go to the **Event Administration** screen.

   Edit —> Event Administration

2. Click **Monitoring**.
3. Un-check the box for **Enable environmental gathering from the master host**.
4. Click **Apply As Default**.
5. Click **Apply to Hosts** and proceed to select the desired hosts.
6. To make the change permanent, click **Apply to Payloads**.
7. Select all desired payloads and click **Apply**.

Note

You must restart the Management Center daemon on the master host and on all nodes.

Altix UV Provisioning Stops While Loading Kernel

Symptom

When provisioning an Altix UV system, the kernel console output stops after the following messages are displayed on the screen:

```
.ELILO v3.12 for EFI/x86_64
Loading kernel kexec-boot/bzImage... done
Loading file kexec-boot/initramfs.gz...done
```

Resolution

In the Management Center GUI, open the general network preferences and put a check in the box for **Direct PXE Boot**. Re-provision the Altix UV system.

Chapter 11

Command-Line Interface

Command-Line Syntax and Conventions

CLI commands documented in this guide adhere to the following rules—commands entered incorrectly may produce the “Command not recognized” error message.

Convention	Description
xyz	Items in bold indicate mandatory parameters or keywords (e.g., all).
<*variable*>	<> Angle brackets and *italics* indicate a user-defined variable (e.g., an IP address or host name)
[x]	[ ] Square brackets indicate optional items.
[x|y|z]	[|] Square brackets with a vertical bar indicate a choice of an *optional* value.
{x|y|z}	{ | } Braces with a vertical bar indicate a choice of a *required* value.
[x{y|z}]	[{ | }] A combination of square brackets and braces with vertical bars indicates a required choice of an optional parameter.

Tip

Help for all CLI commands is available through man pages. To access the man pages, enter man page from the CLI. The cwx man page describes all command-line utilities available in Management Center.

Note

All CLI command arguments documented in this chapter are shown using colon notation only ({--partition:|-p:}). You may also use a space or an equal sign (i.e., --description , -M=) with these arguments.

CLI Commands

Note

Most of the CLI commands outlined in this chapter are exclusive to the Management Center Master Host.

CLI Commands

```
conman {
[[-b <host>[ <host> ...<host_n>]]|
[-d <destination>[:<port>]]|
[-e <character>]|
[-f]|
[-F <file_name>]|
[-h]|
[-j]|
[-l <file_name>]|
[-L]|
[-m]|
[-q]|
[-Q]|
[-r]|
[-v]|
[-V]]
<host_console>
}
```

CLI Commands

```
cwhost {
[partadd [{--description:|-d:} <partition_description>] [--enable:] [--disable:]
[{--regions:|-R} <region1>[,<region2>...]] [{--hosts:|-h} <host1>[,<host2>...]]
<partition>|
[partmod {[{--name:|-n:} <partition_name>] [{--description:|-d:}
<partition_description>]
[--enable:] [--disable:] [{--regions:|-R} <region1>[,<region2>...]]
[{--hosts:|-h} <host1>[,<host2>...]]} <partition>]|
[partdel <partition_name>]|
[partshow [<partition_1>[ <partition_2> ...<partition_n>]]]|
[regionadd [{--description:|-d:} <region_description>] [{--partition:|-p:}
<partition_description>]
[--enable:] [--disable:] [{--hosts:|-h} <host1>[,<host2>...]]
[{--groups:|-g} <group1>[,<group2>...]] <region>]|
[regionmod {--name:|-n:} <region> [{--description:|-d:} <region_description>]
[{--partition:|-p:} <partition_description>] [--enable:] [--disable:]
[{--hosts:|-h} <host1>[,<host2>...]]
[{--groups:|-g} <group1>[,<group2>...]] <region>]<region>]|
[regiondel <region>]|
[regionshow [<region_1>[ <region_2> ...<region_n>]]]|
[hostadd <host1> <mac1> <ip1>[ <host2> <mac2> <ip2>] [{--description:|-d:}
<host_description>]
[--enable:] [--disable:] [{--partition:|-p:} <partition_description>]
[{--regions:|-R:} <region_1>[,<region_2>,...<region_n>]]
[{--iceboxes:|-i:} <icebox_1>:<port>[,<icebox_2>:<port>,...<icebox_3>:<port>]]]|
[hostmod <host> [{--name:|-n:} <host>] [{--interfaces:|-I}
<mac1>|<ip1>[,<mac2>|<ip2>]]
[{--description:|-d:} <host_description>] [--enable:] [--disable:]
[{--partition:|-p:} <partition_description>]
[{--regions:|-R:} <region_1>[,<region_2>,...<region_n>]]
[{--iceboxes:|-i:} <icebox_1>:<port>[,<icebox_2>:<port>,...<icebox_3>:<port>]]]|
[hostdel <host>]|
[hostshow [<host_1>[ <host_1> ...<host_n>]]]|
[ifaceadd <host> <mac> <ip> [{--management:|-M:}]]|
[ifacemod <mac>|<ip> [{--management:|-M:}] [--mac:|-m:} <mac>] [{--ip:|-i:} <ip>]
[{--hostname:|-h:} <host>]]|
[ifacedel <mac>|<ip>]|
[ifaceshow [<mac_1>|<ip_1>[ <mac_2>|<ip_2> ...<mac_n>|<ip_n>]]]|
[iceboxadd <icebox> <mac> <ip> [{--description:|-d:} <icebox_description>]
[{--password:|-p:} <password>] [{--hosts:|-h:} <host1>:<port1>[,<host2>:<port2>...]]]|
[iceboxmod <icebox> [{--name:|-n:} <icebox>] [{--mac:|-m:} <mac>] [{--ip:|-i:} <ip>]
[{--description:|-d:} <icebox_description>] [{--password:|-p:} <password>]
[{--hosts:|-h:} <host1>:<port1>[,<host2>:<port2>...]]]|
[iceboxdel <icebox>]|
[iceboxshow [<icebox_1>[ <icebox_2> ...<icebox_n>]]]|
[inflate <host-range1>[ <host-range2> ...]]|
[deflate <host1>[ <host2> ...]]|
[{--verbose|-v}]|
[-signature]|
[{-usage|-help|-?}]
}
```

CLI Commands

```
cwpower {
{
[--on:|-1:]|
[--off:|-0:]|
[--cycle:|-C:]|
[--reset:|-R:]|
[--powerstatus:|-S:]|
[--reboot:|-r:]|
[--halt:|-h:]|
[--down:|-d:]|
[--hoststatus:|-s:]|
[--flash|-f]|
[--unflash|-u]|
[--beacon|-b]|
[--severity|-e]|
[{--verbose:|-v:} [--progressive:|-p:]]
}
<host_1>[ <host_1> ...<host_n>]|
[-signature]|
[{-usage|-help|-?}]
}
```

```
cwprovision {
[{--download-path:|-d:}<path>
{--image:|-i:}<image>
{--image.revision:|-I:}<revision>
{--kernel:|-k:}[<kernel>]
[{--kernel-log-level:|-l:}[<level>]]
{--payload:|-p:}[<payload>]
[{--payload-download:|-D:}yes|no|default]
[--update --payload.revision:<revision>]
[{--repartition:|-R:}yes|no|default]
[{--working-image:|-w:}<name>]|
[{--next-reboot:|-n:}]]|
[{--query-last-image:|-q} [--uncompressed-hostnames:|-u]]
<host_1>[ <host_1> ...<host_n>]}|
[-signature]|
[{-usage|-help|-?}]
}
```

CLI Commands

```
cwuser {
[useradd [{--description:|-c:}“<description>”] [{--home:|-d:}<home_directory>]
[{--group:|-g:}<primary_group>]
[{--groups:|-G:}<secondary_group_1>[,<secondary_group_2>,...<secondary_group_n>]]
[{--password:|-p:}<encrypted_password>] [{--shell:|-s:}<shell>] [{--uid:|-u:}<uid>]
[{--enable:|-U}] [{--disable:|-L:}] [{--normal:|-n:}] <user>]|
[usermod [{--description:|-c:}“<description>”] [{--home:|-d:}<home_directory>]
[{--group:|-g:}<primary_group>]
[{--groups:|-G:}<secondary_group_1>[,<secondary_group_2>,...<secondary_group_n>]]
[{--password:|-p:}<encrypted_password>] [{--shell:|-s:}<shell>] [{--uid:|-u:}<uid>]
[{--enable:|-U}] [{--disable:|-L:}] [{--name:|-l:}<user>] <user>]|
[userdel <user>]|
[usershow [<user_1>[ <user_2> ...<user_n>]]]|
[passwd <user>]|
[encryptpasswd]|
[groupadd [{--description:|-d:}“<description>”] [{--gid:|-g:}<gid>]
[[{--roles:|-r:}<role_1>] [,<role_2>...<role_n>]]
[{--regions:|-R:}<region_1>[,<region_2>...<region_3>]] <group>]|
[groupmod [{--description:|-d:}“<description>”] [{--gid:|-g:}<gid>]
[[{--roles:|-r:}<role_1>] [,<role_2>,...<role_n>]]
[{--regions:|-R:}<region_1>[,<region_2>,...<region_3>]]
[{--name:|-n:}<group>] <group>]|
[groupdel <group>]|
[groupshow [<group_1>[ <group_2> ...<group_n>]]]|
[roleadd [{--description:|-d:}“<description>”]
[{--privileges:|-p:}<privilege_1>[,<privilege_2>,...<privilege_n>]] <role>]|
[rolemod [{--description:|-d:}“<description>”]
[{--privileges:|-p:}<privilege_1>[,<privilege_2>,...<privilege_n>]]
[{--name:|-n:}<role>] <role>]|
[roledel <role>]|
[roleshow [<role_1>[ <role_2> ...<role_n>]]]|
[privshow [<privilege_1>[ <privilege_2> ...<privilege_n>]]]|
[{--verbose|-v}]|
[-signature]|
[{-usage|-help|-?}]
}
```

```
dbix {
[{-d|--delete} <context_1>[ <context_2> ...<context_n>]]|
[{-i|--import} <context>] |
[{-x|--export} <context_1>[ <context_2> ...<context_n>]]|
[{-usage|-help|-?}]
}
```

```
dbx {
[{--domain:|-d} <domain>] [{--format:|-f:} <format>] [{-usage|-help|-?}] [-
runtime[:verbose]]
[-signature] [-splash]
}
```

CLI Commands

```
imgr {
{--image:|-i:}<image> [{--kernel:|-k:}<kernel>] [{--kernel-
revision:|-K:}<kernel_revision>]
[{--payload:|-p:}<payload>] [{--payload.revision:|-P:}<payload_revision>]
[{--force:|-f:}] [{--list:|-l:}]|
[{-usage|-help|-?}]
}
```

```
kmgr {
{--name:|-n:}<name> [{--description:|-d:}“<description>”]
{--path:|-p:}<path_to_Linux_kernel_source> [{--kernel:|-k:}<name_of_binary>]
[{--architecture:|-a:}<architecture>] [{--modules:|-m:}] [{--binary:|-b:}] [{--list:|-
l:}]|
[{-usage|-help|-?}]
}
```

```
pdcp {[
[-w <host>[,<host>...,<host_n>]]|
[-x <host>[,<host>...,<host_n>]]|
[-a]|
[-i]|
[-r]|
[-p]|
[-q]|
[-f <number>]|
[-l <user>]|
[-t <seconds>]|
[-d]]
<source>[ <source>... <source_n>]
<destination>
}
```

```
pdsh {
[[-w <host>[,<host>...,<host_n>]]|
[-x <host>[,<host>...,<host_n>]]|
[-a]|
[-i]|
[-q]|
[-f <number>]|
[-s]|
[-l <user>]|
[-t <seconds>]|
[-u <seconds>]|
[-n <tasks_per_host>]|
[-d]|
[-S]|
<host>[,<host>...,<host_n>]]
<command>
}
```

CLI Commands

```
pmgr {
[[{--description:|-d:}“<description>”] [{--include:|-i:}<include_file_or_directory>]
[{--include-from:|-I:}<file_containing_list>] [{--location:|-l:}<location_dir>]
[{--silent:|-s:}<silent>]
[{--exclude:|-x:}<exclude_file_or_dir>]] [{--exclude-from:|-X:}<file_containing_list>]
<payload_name>| [{-usage|-help|-?}]
}
```

```
powerman {
[[{--on|-1}]|
[{--off|-0}]|
[{--cycle|-c}]|
[{--reset|-r}]|
[{--flash|-f}]|
[{--unflash|-u}]|
[{--list|-l}]|
[{--query|-q}]|
[{--node|-n}]|
[{--beacon|-b}]|
[{--temp|-t}]|
[{--help|-h}]|
[{--license|-L}]|
[{--destination|-d} host[:port]]|
[{--version|-V}]|
[{--device|-D}]|
[{--telemetry|-T}]|
[{--exprange|-x}]]
<host>[ <host> ...<host_n>]
}
```

```
vcs {
[{identify| id}]|
[status]|
[include <files>]|
[exclude <files>]|
[archive <filename>]|
[import -R:<repository> -M:<module> [-n:<name>] [-d:“<description>”] [<files>]]|
[commit [-n:<name>] [-d:“<description>”] [<files>]]|
[branch [-n:<name>] [-d:“<description>”] [<files>]]|
[{checkout | co} -R:<repository> -M:<module> [-r:<revision>|<branch>|<name>]]|
[{update | up} [-r:<revision>|<branch>|<name>] [<files>]]|
[name [-R:<repository>] [-M:<module>] [-r:<revision>|<branch>|<name>] <text>]|
[describe [-R:<repository>] [-M:<module>] [-r:<revision>|<branch>|<name>] <text>]|
[{narrate | log} [-R:<repository> -M:<module>] [-r:<revision>|<branch>|<name>]]|
[iterate [-R:<repository> [-M:<module> [-r:<revision>|<branch>|<name>]]]]|
[list]|
[{-usage|-help|-?}]
}
```

conman

```
conman {
[[-b <host>[ <host> ...<host_n>]]|
[-d <destination>[:<port>]]|
[-e <character>]|
[-f]|
[-F <file_name>]|
[-h]|
[-j]|
[-l <file_name>]|
[-L]|
[-m]|
[-q]|
[-Q]|
[-r]|
[-v]|
[-V]]
<host_console>
}
```

Description

The Conman client allows you to connect to remote consoles managed by conmand. Console names are separated by spaces or commas and matched to the configuration via globbing. Regular expression matching can be enabled with the -r option.

Conman supports three console access modes: *monitor* (read-only), *interactive* (read-write), and *broadcast* (write-only). Unless otherwise specified, conman opens the console session in interactive mode (the default).

Note

To use Conman for serial access (that is, as your platform management device), Conman must be installed on the Master Host and the console(s) must be configured in */etc/conman.conf*. The Conman daemon (installed as */etc/init.d/conmand*) must also be started.

You can obtain Conman from *http://home.gna.org/conman/*. Additional information on Conman is available from the man pages by entering man conman.conf.

Parameters

[-b *<host>***[** *<host>* ...*<host_n>***]]**	(Optional) Broadcast to multiple host consoles (write-only). You may enter a range of hosts or a space-delimited list of hosts (e.g., **host[1-4 7 9]**). Data sent by the client is copied to all specified consoles in parallel, but console output is not sent back to the client. You can use this option in conjunction with **-f** or **-j**.
[-d *<destination>***[:***<port>***]]**	(Optional) Specify the location of the conmand daemon, overriding the default [127.0.0.1:7890]. This location may contain a host name or IP address and be followed by an optional colon and port number.
[-e *<character>***]**	(Optional) Specify the client escape character, overriding the default (&).

[-f]	(Optional) Specify that write-access to the console should be forced, thereby stealing the console away from existing clients with write privileges. As connections are terminated, conmand informs the original clients of who perpetrated the theft.
[-F <*file_name*>**]**	(Optional) Read console names or patterns from a file with the specified name. Only one console name may be specified per line. Leading and trailing white space, blank lines, and comments (i.e., lines beginning with a **#**) are ignored.
[-h]	(Optional) Display a summary of the command-line options.
[-j]	(Optional) Specify that write-access to the console should be joined, thereby sharing the console with existing clients that have write privileges. As privileges are granted, conmand informs the original clients that privileges have been granted to new clients.
[-l <*file_name*>**]**	(Optional) Log console session output to a file with the specified name.
[-L]	(Optional) Display license information.
[-m]	(Optional) Monitor a console (read-only).
[-q]	(Optional) Query conmand for consoles matching the specified names or patterns. Output from this query can be saved to file for use with the **-F** option.
[-Q]	(Optional) Enable quiet-mode, suppressing informational messages. This mode can be toggled on and off from within a console session via the **&Q** escape.
[-r]	(Optional) Match console names via regular expressions instead of globbing.
[-v]	(Optional) Enable verbose mode.
[-V]	(Optional) Display version information.
<*host_console*>	The name of the host to which to connect.

ESCAPE CHARACTERS

Conman supports the following escapes and assumes the default escape character (&):

&?	Display a list of all escapes currently available.
&.	Terminate the connection.
&&	Send a single escape character.
&B	Send a serial-break to the remote console.
&F	Switch from read-only to read-write via a force.
&I	Display information about the connection.
&J	Switch from read-only to read-write via a join.
&L	Replay the last 4KB of console output. This escape requires that logging is enabled for the console in the conmand configuration.
&M	Switch from read-write to read-only.
&Q	Toggle quiet-mode to display or suppress informational messages.
&R	Reset the host associated with this console. This escape requires that **resetcmd** is specified in the conmand configuration.
&Z	Suspend the client.

ENVIRONMENT

The following environment variables may be used to override default settings.

CONMAN_HOST	Specifies the host name or IP address at which to contact conmand, but may be overridden with the **-d** command-line option. Although a port number separated by a colon may follow the host name (i.e., host:port), the CONMAN_PORT environment variable takes precedence. If you do not specify a host, the default host IP address (127.0.0.1) is used.
CONMAN_PORT	Specifies the port on which to contact conmand, but may be overridden by the **-d** command-line option. If not set, the default port (7890) is used.
CONMAN_ESCAPE	The first character of this variable specifies the escape character, but may be overridden by the **-e** command-line option. If not set, the default escape character (&) is used.

! Warning!

Client and server communications are not yet encrypted.

Example 1

To connect to host console n1, enter:

```
conman n1
```

Note

Once in conman, enter &. to exit or &? to display a list of conman commands.

Example 2

To broadcast (write-only) to multiple hosts, enter:

```
conman -b n[1-10]
```

Tip

To view the output of broadcast commands on a group of hosts, use the conmen command before you begin entering commands from conman. Conmen opens a new window for each host and displays the host output.

For example, the following command opens new consoles for hosts n2-n4:

```
conmen n[2-4]
```

cwhost

```
cwhost {
[partadd [{--description:|-d:} <partition_description>] [--enable:] [--disable:]
[{--regions:|-R} <region1>[,<region2>...]] [{--hosts:|-h} <host1>[,<host2>...]]
<partition>|
[partmod {[{--name:|-n:} <partition_name>] [{--description:|-d:}
<partition_description>]
[--enable:] [--disable:] [{--regions:|-R} <region1>[,<region2>...]]
[{--hosts:|-h} <host1>[,<host2>...]]} <partition>]|
[partdel <partition_name>]|
[partshow [<partition_1>[ <partition_2> ...<partition_n>]]]|
[regionadd [{--description:|-d:} <region_description>] [{--partition:|-p:}
<partition_description>]
[--enable:] [--disable:] [{--hosts:|-h} <host1>[,<host2>...]]
[{--groups:|-g} <group1>[,<group2>...]] <region>]|
[regionmod {--name:|-n:} <region> [{--description:|-d:} <region_description>]
[{--partition:|-p:} <partition_description>] [--enable:] [--disable:]
[{--hosts:|-h} <host1>[,<host2>...]]
[{--groups:|-g} <group1>[,<group2>...]] <region>]<region>]|
[regiondel <region>]|
[regionshow [<region_1>[ <region_2> ...<region_n>]]]|
[hostadd <host1> <mac1> <ip1>[ <host2> <mac2> <ip2>] [{--description:|-d:}
<host_description>]
[--enable:] [--disable:] [{--partition:|-p:} <partition_description>]
[{--regions:|-R:} <region_1>[,<region_2>,...<region_n>]]
[{--iceboxes:|-i:} <icebox_1>:<port>[,<icebox_2>:<port>,...<icebox_3>:<port>]]]|
[hostmod <host> [{--name:|-n:} <host>] [{--interfaces:|-I} <mac1>|<ip1>[,<mac2>|<ip2>]]
[{--description:|-d:} <host_description>] [--enable:] [--disable:]
[{--partition:|-p:} <partition_description>]
[{--regions:|-R:} <region_1>[,<region_2>,...<region_n>]]
[{--iceboxes:|-i:} <icebox_1>:<port>[,<icebox_2>:<port>,...<icebox_3>:<port>]]]|
[hostdel <host>]|
[hostshow [<host_1>[ <host_1> ...<host_n>]]]|
[ifaceadd <host> <mac> <ip> [{--management:|-M:}]]|
[ifacemod <mac>|<ip> [{--management:|-M:}] [--mac:|-m:} <mac>] [{--ip:|-i:} <ip>]
[{--hostname:|-h:} <host>]]|
[ifacedel <mac>|<ip>]|
[ifaceshow [<mac_1>|<ip_1>[ <mac_2>|<ip_2> ...<mac_n>|<ip_n>]]]|
[iceboxadd <icebox> <mac> <ip> [{--description:|-d:} <icebox_description>]
[{--password:|-p:} <password>] [{--hosts:|-h:} <host1>:<port1>[,<host2>:<port2>...]]]|
[iceboxmod <icebox> [{--name:|-n:} <icebox>] [{--mac:|-m:} <mac>] [{--ip:|-i:} <ip>]
[{--description:|-d:} <icebox_description>] [{--password:|-p:} <password>]
[{--hosts:|-h:} <host1>:<port1>[,<host2>:<port2>...]]]|
[iceboxdel <icebox>]|
[iceboxshow [<icebox_1>[ <icebox_2> ...<icebox_n>]]]|
[inflate <host-range1>[ <host-range2> ...]]|
[deflate <host1>[ <host2> ...]]|
[{--verbose|-v}]|
[-signature]|
[{-usage|-help|-?}]
}
```

Description

The Host Administration (cwhost) utility allows you to add, modify, view the current state of, or delete any partition, region, host, interface, or Icebox in your cluster.

Subcommands

partadd

Add a partition to the cluster.

[{--description:|-d:} *<partition_description>***]**

(Optional) A brief description of the partition. If you do not specify a description, this field remains blank.

[--enable:] [--disable:]

(Optional) Indicates whether or not the partition is enabled. If you do not specify this option, Management Center will enable the partition.

[{--regions:|-R} *<region1>***[**,*<region2>...***]]**

(Optional) The list of regions that are members of this partition. If you do not specify any regions, none are included in the partition.

[{--hosts:|-h} *<host1>***[**,*<host2>...***]]**

(Optional) The list of hosts that are members of this partition. If you do not specify any hosts, none are included in the partition.

<partition>

The name of the partition to add.

partmod

Modify a partition on the cluster. Unchanged entries remain the same.

[{--name:|-n:} *<partition_name>***]**

(Optional) Change the partition name. If you do not specify a name, Management Center uses the current partition name.

[{--description:|-d:} *<partition_description>***]**

(Optional) A brief description of the partition. If you do not specify a description, Management Center uses the current partition description.

[--enable:] [--disable:]

(Optional) Indicates whether or not the partition is enabled. If you do not specify this option, the partition remains in its original state.

[{--regions:|-R} *<region1>***[**,*<region2>...***]]**

(Optional) The list of regions that are members of this partition. If you do not specify any regions, the partition remains in its original state.

[{--hosts:|-h} *<host1>***[**,*<host2>...***]]**

(Optional) The list of hosts that are members of this partition. If you do not specify any hosts, the partition remains in its original state.

<partition>

The name of the partition to add.

partdel

Delete a partition from the cluster.

<*partition_name>*

The name of the partition to delete.

partshow

Display the current settings for a partition(s).

[*<partition_1>***[** *<partition_2> ...<partition_n>***]]**

(Optional) The name(s) of the partition(s) for which to display the current settings. Multiple entries are delimited by spaces. Leave this option blank to display all partitions.

regionadd

Add a region to a partition.

[{--description:|-d:} *<region_description>***]**	(Optional) A brief description of the region. If you do not specify a description, this field remains blank.
[{--partition:|-p:} *<partition_description>***]**	(Optional) The partition to which this region belongs. If you do not specify a partition, Management Center assigns the region to the default or unassigned partition.
[--enable:] [--disable:]	(Optional) Indicates whether or not the region is enabled. If you do not specify this option, Management Center will enable the region.
[{--hosts:|-h} *<host1>***[**,*<host2>...***]]**	(Optional) The list of hosts that are members of this region. If you do not specify this option, the region will not contain any member hosts.
[{--groups:|-g} *<group1>***[**,*<group2>...***]]**	(Optional) The list of groups that may access this region. If you do not specify this option, the region will not be available to any groups.
<region>	The name of the new region.

regionmod

Modify a region on the cluster. Unchanged entries remain the same.

{--name:|-n:} *<region>*	(Optional) Change the region name. If you do not specify a name, Management Center uses the current region name.
[{--description:|-d:} *<region_description>***]**	(Optional) A brief description of the region. If you do not specify a description, Management Center uses the current region description.
[{--partition:|-p:} *<partition_description>***]**	(Optional) The partition to which this region belongs. If you do not specify a partition, Management Center assigns the region to the original partition specified.
[--enable:] [--disable:]	(Optional) Indicates whether or not the region is enabled. If you do not specify this option, the region remains in its original state.
[{--hosts:|-h} *<host1>***[**,*<host2>...***]]**	(Optional) The list of hosts that are members of this region. If you do not specify any hosts, the region remains in its original state.
[{--groups:|-g} *<group1>***[**,*<group2>...***]]**	(Optional) The list of groups that may access this region. If you do not specify any groups, the region remains in its original state.
<region>	The name of the region to modify.

regiondel

Delete a region from the cluster.

<region>	The name of the region to delete.

regionshow

Display the current settings for a region(s).

[*<region_1>*[*<region_2> ...<region_n>*]]

(Optional) The name of the region(s) for which to display the current settings. Multiple entries are delimited by spaces. Leave this option blank to display all regions.

`hostadd`

Add a host to the cluster.

<host1> <mac1> <ip1>[*<host2> <mac2> <ip2>*]

The name of each new host, its MAC address, and its IP address. The first host specified is the management interface. Multiple entries are space-delimited.

[{--description:|-d:} *<host_description>***]**

(Optional) A brief description of the host. If you do not specify a description, this field remains blank.

[--enable:] [--disable:]

(Optional) Indicates whether or not the host is enabled. If you do not specify this option, Management Center enables the host.

[{--partition:|-p:} *<partition_description>***]**

(Optional) The partition to which this host belongs. If you do not specify a partition, Management Center assigns the host to the default or unassigned partition.

[{--regions:|-r:} *<region_1>*[,*<region_2>,...<region_n>*]]

(Optional) The region(s) to which this host belongs. If you do not specify a region, Management Center does not assign the host to any region. Multiple entries are comma-delimited.

[{--iceboxes:|-i:} *<icebox_1>:<port>*[,*<icebox_2>:<port>,...<icebox_3>:<port>*]]

(Optional) The Icebox(es) and port(s) to which this host is connected. If you do not specify an Icebox and port, Management Center assumes that the host is not connected to an Icebox. Multiple entries are comma-delimited.

`hostmod`

Modify a host on the cluster—unchanged entries remain the same.

<host>

The name of the host to modify.

{--name:|-n:} *<host>*

The host’s new name.

[{--interfaces:|-I} *<mac1>|<ip1>*[,*<mac2>|<ip2>*]]

(Optional) A list of interfaces with which this host is associated. If none of the specified interfaces are management interfaces, Management Center marks the first interface as the management interface.

[{--description:|-d:} *<host_description>***]**

(Optional) A brief description of the host. If you do not specify a description, Management Center uses the current host description.

[--enable: {yes|no}]

(Optional) Indicates whether or not the host is enabled. If you do not specify this option, the host remains in its original state.

[{--partition:|-p:} *<partition_description>***]**

(Optional) The partition to which this host belongs. If you do not specify a partition, the host remains associated with the original partition specified.

[{--regions:|-r:} *<region_1>***[,***<region_2>,...<region_n>***]]**	(Optional) The region(s) to which this host belongs. If you do not specify a partition, the host will not belong to any region. Multiple entries are comma-delimited.
[{--iceboxes:|-i:} *<icebox_1>***:***<port>***[,***<icebox_2>***:***<port>,...<icebox_3>***:***<port>***]]**	(Optional) The Iceboxes and ports to which this host is connected. If you do not specify an Icebox and port, Management Center assumes that the host is not connected to an Icebox. Multiple entries are comma-delimited.

hostdel

Delete a host.

<host>	The name of the host to delete.

hostshow

Display the current settings for a host(s).

[*<host_1>***[** *<host_1> ...<host_n>***]]**	(Optional) The name of the host(s) for which to display the current settings. Multiple entries are delimited by spaces. Leave this option blank to display all hosts.

ifaceadd

Add an interface to the cluster.

<host>	The name of the host on which you added the interface.
<mac>	The MAC address of the interface.
<ip>	The IP address of the interface.
[{--management:|-M:}]	(Optional) Specify whether or not this interface is a management interface. If you do not specify this option, Management Center assumes that this interface is not a management interface.

ifacemod

Modify an interface on the cluster—unchanged entries remain the same.

<mac>	The MAC address of the interface.
<ip>	The IP address of the interface.
[{--management:|-M:}]	(Optional) Specify whether or not this interface is a management interface. If you do not specify this option, the interface remains in its original state.
[--mac:|-m:} *<mac>***]**	(Optional) Change the interface's hardware or MAC address.
[{--ip:|-i:} *<ip>***]**	(Optional) Change the interface's IP address.
[{--hostname:|-h:} *<host>***]**	(Optional) Change the host to which this interface belongs.

ifacedel

Delete an interface from the cluster.

<mac>	The MAC address of the interface to delete.
<ip>	The IP address of the interface to delete.

ifaceshow

Display the current settings for an interface(s).

[<*mac_1*>**|**<*ip_1*>**[** <*mac_2*>**|**<*ip_2*> ...<*mac_n*>**|**<*ip_n*>**]]**
(Optional) The MAC or IP address(es) of the interface(s) for which to display the current settings. Multiple entries are delimited by spaces. Leave this option blank to display all interfaces.

iceboxadd

Add an Icebox to the cluster.

<*host*>	The name of the new Icebox.
<*mac*>	The MAC address of the new Icebox.
<*ip*>	The IP address of the new Icebox.

[{--description:|-d:} <*icebox_description*>**]**
(Optional) A brief description of the Icebox. If you do not specify a description, this field remains blank.

[{--password:|-p:} <*password*>**]**
(Optional) The Icebox's administrative password. If you do not specify a password, Management Center uses the default password "*icebox*".

[{--hosts:|-h:} <*host1*>**:**<*port1*>**[,**<*host2*>**:**<*port2*>...**]]**

(Optional) A list of hosts connected to the Icebox and the ports to which they are connected. If you do not specify this option, Management Center assumes that the hosts are not connected to an Icebox.

iceboxmod

Modify an Icebox on the cluster—unchanged entries remain the same.

<*icebox*> The name of the Icebox to modify.

[{--name:|-n:} <*icebox*>**]**
(Optional) The Icebox's new name.

[{--mac:|-m:} <*mac*>**]**	(Optional) Change the Icebox's hardware or MAC address.
[{--ip:|-i:} <*ip*>**]**	(Optional) Change the Icebox's IP address.

[{--description:|-d:} <*icebox_description*>**]**
(Optional) A brief description of the Icebox. If you do not specify a description, Management Center uses the current Icebox description.

[{--password:|-p:} <*password*>**]**
(Optional) The Icebox's administrative password. If you do not specify a password, Management Center uses the original password.

[{--hosts:|-h:} <*host1*>**:**<*port1*>**[,**<*host2*>**:**<*port2*>...**]]**

(Optional) A list of hosts connected to the Icebox and the ports to which they are connected. If you do not specify this option, Management Center assumes that the hosts remain in their original state.

iceboxdel

Delete a Management Center Icebox.

<*icebox*> The name of the Icebox to delete.

iceboxshow

Display the current settings for an Icebox(es).

[<*icebox_1*>**[** <*icebox_2*> **...**<*icebox_n*>**]]**	(Optional) The Icebox(es) for which to display the current setting(s). Multiple entries are delimited by spaces. Leave this option blank to display all Iceboxes.
inflate <*host-range1*>**[** <*host-range2*> **...]**	(Optional) Allows you to change between full and compressed host list format. Inflate the specified host range(s) to display a full list of hosts.
deflate <*host1*>**[** <*host2*> **...]**	(Optional) Allows you to change between full and compressed host list format. Deflate the specified host range(s) to display a compressed host list.
[{--verbose|-v}]	(Optional) Display verbose output when performing operations. This option is common to all subcommands.
[-signature]	(Optional) Displays the application signature. The application signature contains the name, description, version, and build information of this application.
[{-usage|-help|-?}]	(Optional) Display help information for the command and exit. All other options are ignored.

Examples

EXAMPLE 1

View the layout of the system:

```
cwhost hostshow
```

EXAMPLE 2

Get details of the system:

```
cwhost hostshow -v
```

EXAMPLE 3

Create a region called *group1*:

```
cwhost regionadd group1
```

EXAMPLE 4

Add a host to region *group1* with the host name *n1*, the mac *0005b342afe1*, and the IP address *10.0.0.1*:

```
cwhost hostadd -r:group1 n1 0005b342afe1 10.0.0.1
```

EXAMPLE 5

Add host *n2* to the *group1* region:

```
cwhost hostmod -r:group1 n2
```

EXAMPLE 6

Deflate the host list *n1*, *n2*, *n3*, and *n4*:

```
cwhost deflate n1 n2 n3 n4
```

```
n[1-4]
```

EXAMPLE 7

Inflate the host list *n[1-4]*:

```
cwhost inflate n[1-4]
```

```
n1
n2
n3
n4
```

cwpower

```
cwpower {
{
[--on:|-1:]|
[--off:|-0:]|
[--cycle:|-C:]|
[--reset:|-R:]|
[--powerstatus:|-S:]|
[--reboot:|-r:]|
[--halt:|-h:]|
[--down:|-d:]|
[--hoststatus:|-s:]|
[--flash|-f]|
[--unflash|-u]|
[--beacon|-b]|
[--severity|-e]|
[{--verbose:|-v:} [--progressive:|-p:]]
}
<host_1>[ <host_1> ...<host_n>]|
[-signature]|
[{-usage|-help|-?}]
}
```

Description

The Power Administration (cwpower) utility allows you to perform power administration operations on a host(s) within the cluster. Operations include power on, power off, power cycle, reset, reboot, halt, and power down (a soft power off). You may also query the current power status of a particular host(s).

Note

You may specify only one power administration operation option each time you use the cwpower command.

Parameters

[--on|-1]	(Optional) Turn on power to the specified host(s).
[--off|-0]	(Optional) Turn off power to the specified host(s).
[--cycle|-C]	(Optional) Cycle power to the specified host(s).
[--reset|-R]	(Optional) Perform a hardware reset for the specified host(s).
[--powerstatus|-S]	(Optional) Query the hard power status for the specified host(s).
[--reboot|-r]	(Optional) Reboot the specified host(s).
[--halt|-h]	(Optional) Halt the specified host(s).
[--down|-d]	(Optional) Execute a soft power down on the specified host(s).
[--hoststatus|-s]	(Optional) Query the host administration power status for the specified host(s).
[--flash|-f]	(Optional) Turn the beacon on for the specified host(s).
[--unflash|-u]	(Optional) Turn the beacon off for the specified host(s).
[--beacon|-b]	(Optional) Report the beacon status for the specified host(s).
[--severity|-e]|	(Optional) Report the error status for the specified host(s).

[{--verbose|-v} [--progressive|-p]]	(Optional) Change the standard output to verbose. Output displays the power status of each host, one per line. To display output as information becomes available, select the progressive option—progressive output is not guaranteed to be sorted and is not summarized.
*<host_1>***[** *<host_1>* ...*<host_n>***]**	The name of the host(s) for which to execute the specified operation. You may enter a range of hosts or a space-delimited list of hosts (e.g., **host[1-4 7 9]**).
[-signature]	(Optional) Displays the application signature. The application signature contains the name, description, version, and build information of this application.
[{-usage|-help|-?}]	(Optional) Display help information for the command and exit. All other options are ignored.

Examples

EXAMPLE 1

To Power on hosts 1–10:

```
cwpower -1 n[1-10]
```

EXAMPLE 2

Power off host 1:

```
cwpower -0 n1
```

EXAMPLE 3

Power cycle hosts 2–5:

```
cwpower -C n[2-5]
```

EXAMPLE 4

Check the status (On, Off, Unknown, Provisioning) of hosts 1–10:

```
cwpower -s n[1-10]
```

cwprovision

```
cwprovision {
[{--download-path:|-d:}<path>
{--image:|-i:}<image>
{--image.revision:|-I:}<revision>
{--kernel:|-k:}[<kernel>]
[{--kernel-log-level:|-l:}[<level>]]
{--payload:|-p:}[<payload>]
[{--payload-download:|-D:}yes|no|default]
[--update --payload.revision:<revision>]
[{--repartition:|-R:}yes|no|default]
[{--working-image:|-w:}<name>]|
[{--next-reboot:|-n:}]]|
[{--query-last-image:|-q} [--uncompressed-hostnames:|-u]]
<host_1>[ <host_1> ...<host_n>]}|
[-signature]|
[{-usage|-help|-?}]
}
```

Description

The Provisioning (cwprovision) utility allows you to provision or update a host(s) on the cluster and use working copies to override the kernel and payload associated with the image. See *Provisioning* on page 147 and *Version Control System (VCS)* on page 140.

Parameters

{--download-path:|-d:}*<path>*	The path to which to download the image during the boot process (by default, */mnt*).
{--image:|-i:}*<image>*	The image to use to provision the host(s). Unless you specify the working image option, Management Center assumes that the image is a version-controlled image.
{--image.revision:|-I:}*<revision>*	The revision of the image to use to provision the host(s). If you specify a branch revision, Management Center uses the tip revision of the branch. If you do not specify a revision or a working image, Management Center uses the tip revision of the image. Revisions may be specified either numerically or by alias.

Note

The image.revision option is not available in conjunction with the working-image option.

{--kernel:|-k:}[*<kernel>***]**	The working copy of the kernel associated with the image used to provision the host(s). The name is required only if two or more working copies of the kernel exist.
[{--kernel-log-level:|-l:}[*<level>***]]**	Select the kernel verbosity level used to control debug messages. This level may range from 1 (the least verbose) to 8 (the most verbose). By default, the verbosity level is 1.

Power Management{--payload:|-p:}[<*payload*>**]**

The working copy of the payload associated with the image used to provision the host(s). The name is required only if two or more working copies of the payload exist.

[{--payload-download:|-D:}yes|no|default]

(Optional) Specify whether or not to force a download of the payload to the host during this provisioning operation. The default option automatically detects whether or not to download the payload. See *Advanced Provisioning Options* on page 151.

[--update --payload.revision:<*revision*>**]**

Update the host(s) with the version of the payload specified.

[{--repartition:|-R:}yes|no|default]

(Optional) Specify whether or not to force a repartition of the host during this provisioning operation. The default option automatically detects whether or not to repartition the host. See *Advanced Provisioning Options* on page 151.

[{--working-image:|-w:}<*name*>]

(Optional) Use the working copy of the specified image to provision the host(s).

Note

The working-image option is not available in conjunction with the image.revision option.

[{--next-reboot:|-n:}] (Optional) Provision the selected host(s) after the next reboot.

[{--query-last-image:|-q}] (Optional) Display the name and revision of the last image used to provision the host(s). By default, this option displays a list of compressed host names and their corresponding images. To change this format, use the uncompressed-hostnames option. The uncompressed format displays hosts and images in a colon-separated list that is easily parsed by command-line tools. Each line follows the format:

<*host_name*>**:[VCS| Working] Image:**<*image_name*>:
{<*VCS_revision*>|<*user_name*>}**:** <*kernel*>**:**<*payload*>

Tip

The kernel and payload specify zero (0) if you use the VCS version and one (1) if you use the working version to override the kernel or payload using the advanced provisioning options.

Note

The query-last-image option can display image and host information even if the host is down.

[{--uncompressed-hostnames:|-u}]

(Optional) Select this option to change the output format for query-last-image to list one host name and corresponding image per line. This option can be used only with query-last-image.

<*host_1*>**[** <*host_1*> ...<*host_n*>**]**

The name of the host(s) to provision. You may enter a range of hosts or a space-delimited list of hosts (e.g., host[1-4 7 9]).

[-signature] (Optional) Displays the application signature. The application signature contains the name, description, version, and build information of this application.

[{-usage|-help|-?}] (Optional) Display help information for the command and exit. All other options are ignored.

Examples

Tip

Use vcs iterate -R:images to see what images are available for provisioning. For a list of working images, use imgr --list.

EXAMPLE 1

To provision hosts 2–4 with image *Compute_Host*:

```
cwprovision -i:Compute_Host n[2-4]
```

EXAMPLE 2

To provision hosts 2–4 with an older version (version 3) of the image *Compute_Host*:

```
cwprovision -i:Compute_Host -I:3 n[2-4]
```

EXAMPLE 3

To set advanced options to force re-partitioning and download the payload for hosts 2–4:

```
cwprovision -i:Compute_Host -I:3 -R:yes -D:yes n[2-4]
```

EXAMPLE 4

To provision hosts 2–10 after the next reboot:

```
cwprovision -i:rhel4_img --next-reboot n[2-10]
```

EXAMPLE 5

To update hosts 6-8 with revision 9 of the payload:

```
cwprovision --update --payload.revision:9 n[6-8]
```

cwuser

```
cwuser {
[useradd [{--description:|-c:}“<description>”] [{--home:|-d:}<home_directory>]
[{--group:|-g:}<primary_group>]
[{--groups:|-G:}<secondary_group_1>[,<secondary_group_2>,...<secondary_group_n>]]
[{--password:|-p:}<encrypted_password>] [{--shell:|-s:}<shell>] [{--uid:|-u:}<uid>]
[{--enable:|-U}] [{--disable:|-L:}] [{--normal:|-n:}] <user>]|
[usermod [{--description:|-c:}“<description>”] [{--home:|-d:}<home_directory>]
[{--group:|-g:}<primary_group>]
[{--groups:|-G:}<secondary_group_1>[,<secondary_group_2>,...<secondary_group_n>]]
[{--password:|-p:}<encrypted_password>] [{--shell:|-s:}<shell>] [{--uid:|-u:}<uid>]
[{--enable:|-U}] [{--disable:|-L:}] [{--name:|-l:}<user>] <user>]|
[userdel <user>]|
[usershow [<user_1>[ <user_2> ...<user_n>]]]|
[passwd <user>]|
[encryptpasswd]|
[groupadd [{--description:|-d:}“<description>”] [{--gid:|-g:}<gid>]
[[{--roles:|-r:}<role_1>] [,<role_2>...<role_n>]]
[{--regions:|-R:}<region_1>[,<region_2>...<region_3>]] <group>]|
[groupmod [{--description:|-d:}“<description>”] [{--gid:|-g:}<gid>]
[[{--roles:|-r:}<role_1>] [,<role_2>,...<role_n>]]
[{--regions:|-R:}<region_1>[,<region_2>,...<region_3>]]
[{--name:|-n:}<group>] <group>]|
[groupdel <group>]|
[groupshow [<group_1>[ <group_2> ...<group_n>]]]|
[roleadd [{--description:|-d:}“<description>”]
[{--privileges:|-p:}<privilege_1>[,<privilege_2>,...<privilege_n>]] <role>]|
[rolemod [{--description:|-d:}“<description>”]
[{--privileges:|-p:}<privilege_1>[,<privilege_2>,...<privilege_n>]]
[{--name:|-n:}<role>] <role>]|
[roledel <role>]|
[roleshow [<role_1>[ <role_2> ...<role_n>]]]|
[privshow [<privilege_1>[ <privilege_2> ...<privilege_n>]]]|
[{--verbose|-v}]|
[-signature]|
[{-usage|-help|-?}]
}
```

Description

The User Administration (cwuser) utility allows you to perform user, group, and role administration operations on the cluster. Operations include adding, modifying, deleting, and displaying the current state of users, groups, and roles.

Subcommands

useradd

Add a Management Center user account.

[{--description:|-c:}“<description>”]

The user’s description (e.g., the user’s full name). If you do not specify a description, this field remains blank.

[{--home:|-d:}<home_directory>]

The user’s home directory (by default, */home/<user>*).

[{--group:|-g:}<*primary_group*>**]**

The user's primary group. You may enter the group name or its numerical gid. If you do not enter a primary group, Management Center will do one of the following:

Red Hat Linux

Create a group with the same name as the user and assign the primary group to that group (unless you specify the **[--normal:|-n:]** option).

SuSE Linux

The primary group for the user is the default group specified for users, usually *users*.

[{--groups:|-G:}<*secondary_group_1*>**[,**<*secondary_group_2*>*,...*<*secondary_group_n*>**]]**

The secondary group(s) to which the user belongs. If you do not specify this option, the user belongs to no secondary groups. Multiple entries are delimited by commas.

[{--password:|-p:}<*encrypted_password*>**]**

The user's encrypted password. If you do not specify a password, Management Center disables the account.

[{--shell:|-s:}<*shell*>**]**

The user's login shell. If you do not specify this option, Management Center assigns */bin/bash* as the user's login shell.

[{--uid:|-u:}<*uid*>**]**

The user's uid. If you do not specify a uid, Management Center assigns the first available uid greater than 499.

[{--enable:|-U}] [{--disable:|-L:}]

These options allow you to enable or disable the user's account. The -U (unlock) and -L (lock) options are provided for compatibility with the useradd utility and allow you to enable and disable the user's account respectively. If you do not specify either of these options, the user's account is enabled by default (unless no password is supplied).

[{--normal:|-n:}]

If you do not specify a group for the user on Red Hat Linux, Management Center will behave as it does with most other versions of Linux. The user's primary group uses the default user group, *users*.

<*user*>

The user's login name.

usermod

Modify an existing Management Center user account.

[{--description:|-c:}"<*description*>**"]**

The user's description (e.g., the user's full name). If you do not specify a description, Management Center uses the current description.

[{--home:|-d:}<*home_directory*>**]**

The user's home directory. If left blank, the current home directory.

[{--group:|-g:}<*primary_group*>**]**

The user's primary group. You may enter the group name or its numerical gid. If you do not enter a primary group, Management Center uses the current group assignment.

[{--groups:|-G:}<*secondary_group_1*>**[,**<*secondary_group_2*>*,...*<*secondary_group_n*>**]]**

The secondary group(s) to which the user belongs. If you do not specify this option, Management Center assigns the user to any secondary groups previously assigned. Multiple entries are delimited by commas.

[{--password:|-p:}*<encrypted_password>***]**	Change the user’s encrypted password. If you do not specify a password, Management Center uses the current password.
[{--shell:|-s:}*<shell>***]**	The user’s login shell. If you do not specify this option, Management Center uses the login shell previously assigned to the user.
[{--uid:|-u:}*<uid>***]**	The user’s uid. If you do not specify a uid, Management Center uses the current uid.
[{--enable:|-U}] [{--disable:|-L:}]	These options allow you to enable or disable the user’s account. The -U (unlock) and -L (lock) options are provided for compatibility with the useradd utility and allow you to enable and disable the user’s account respectively. If you do not specify either of these options, the user’s account is enabled by default (unless no password is supplied).
[{--name:|-l:}*<user>***]**	Change the login name for the user’s account. If you do not specify this option, Management Center uses the previous login name.
<user>	The user’s login name.

userdel

Delete a Management Center user account.

<user>	The user’s login name.

usershow

Display the current settings for Management Center user(s).

[*<user_1>***[** *<user_2> ...<user_n>***]]**	(Optional) The user’s(s’) login name(s). Multiple entries are delimited by spaces. Leave this option blank to display all users.

passwd

Alter the password for a Management Center user. After making the change, Management Center prompts you to re-enter the password.

<user>	The user’s login name.

encryptpasswd

This option allows you to encrypt a clear text password into the Management Center encrypted format and display it on screen. You may then copy and paste the encrypted password when creating a new user account. See example on page 248.

Note

Encrypted password strings often contain characters with which the Linux shell has problems. To overcome this, encrypted text must be escaped using single quotes:

```
cwuser usermod '-p:$1$Jx^VLEZy$/7SmJmEbmbVMQW13kxaIg.' john
```

---

groupadd

Add a group to Management Center.

[{--description:|-d:}“*<description>***”]**	The group’s description. If you do not specify a description, this field remains blank.
[{--gid:|-g:}*<gid>***]**	The group’s gid. If you do not specify a gid, Management Center assigns the first available gid greater than 499.

[{--roles:|-r:}_<role_1>_**[,**_<role_2>,...<role_n>_**]]**	The roles associated with the group. If you do not specify a role(s), the group is not associated with any roles. Multiple entries are delimited by commas.
[{--regions:|-R:}_<region_1>_**[,**_<region_2>,...<region_3>_**]]**	The region(s) associated with the group. If you do not specify a region(s), Management Center does not associate the group with any regions. Multiple entries are delimited by commas.
<group>	Group name.

groupmod

Modify an existing Management Center group.

[{--description:|-d:}"_<description>_**"]**	The group's description. If you do not specify a description, Management Center uses the current group description.
[{--gid:|-g:}_<gid>_**]**	The group's gid. If you do not specify a gid, Management Center uses the gid previously assigned.
[{--roles:|-r:}_<role_1>_**[,**_<role_2>,...<role_n>_**]]**	The roles associated with the group. If you do not specify a role(s), the group maintains its previous role associations. Multiple entries are delimited by commas.
[{--regions:|-R:}_<region_1>_**[,**_<region_2>,...<region_3>_**]]**	The regions associated with the group. If you do not specify a region(s), Management Center maintains the current region associations. Multiple entries are delimited by commas.
[{--name:|-n:}_<group>_**]**	Use this option to change the group name. If you do not specify a name, the group name remains unchanged.
<group>	Current group name.

groupdel

Delete a Management Center group.

<group>	Group name.

groupshow

Display the current settings for Management Center group(s).

[_<group_1>_**[** _<group_2> ...<group_n>_**]]**	(Optional) Group name(s) for which to display the current settings. Multiple entries are delimited by spaces. Leave this option blank to display all groups.

roleadd

Add a role to the Management Center database.

[{--description:|-d:}"_<description>_**"]**	The role's description. If you do not specify a role description, this field remains blank.
[{--privileges:|-p:}_<privilege_1>_**[,**_<privilege_2>,...<privilege_n>_**]]**	The privileges associated with the role. If you do not specify a privilege(s), Management Center does not assign any privileges to the role. Multiple entries are delimited by commas.
<role>	The name of the role.

rolemod

Modify an existing Management Center role.

[{--description:|-d:}"*<description>***"]**
The role's description. If you do not specify a description for the role, Management Center uses the current description.

[{--privileges:|-p:}*<privilege_1>***[,***<privilege_2>,...<privilege_n>***]]**
The privileges associated with the role. If you do not specify a privilege(s), Management Center uses current privilege associations. Multiple entries are delimited by commas.

[{--name:|-n:}*<role>***]** Use this option to change the name of the role. If you do not specify a name, the role name remains unchanged.

<role> The name of the current role.

roledel

Delete a Management Center role.

<role> The name of the role to delete.

roleshow

Display the current settings for Management Center role(s).

[*<role_1>***[** *<role_2>* **...***<role_n>***]]**
(Optional) The name of the role(s) for which to display the current settings. Multiple entries are delimited by spaces. Leave this option blank to display all roles.

privshow

Display the current settings for Management Center privilege(s).

[*<privilege_1>***[** *<privilege_2> ...<privilege_n>***]]**
(Optional) The privilege(s) for which to display the current settings. Multiple entries are delimited by spaces. Leave this option blank to display all privileges.

[{--verbose|-v}] (Optional) Display verbose output when performing operations. This option is common to all subcommands.

[-signature] (Optional) Displays the application signature. The application signature contains the name, description, version, and build information of this application.

[{-usage|-help|-?}] (Optional) Display help information for the command and exit. All other options are ignored.

Examples

EXAMPLE 1

Display the current users in the system:

```
cwuser usershow -v
```

EXAMPLE 2

Add the user *john* to the *users* group:

```
cwuser useradd -g:users john
```

Note

John’s account will be disabled until you add a password.

EXAMPLE 3

Add an encrypted password to a new user account:

```
cwuser encryptpasswd
<Enter, then verify password>
```

The command outputs an encrypted string to use when creating the new account.

```
$1$Jx^VLEZy$/7SmJmEbmbVMQW13kxaIg
```

Note

Because encrypted password strings often contain characters with which the Linux shell has problems, encrypted text and user names containing spaces (e.g., John Johnson) must be escaped using single quotes.

Create the new user account using the encrypted password.

```
cwuser useradd '-p:$1$Jx^VLEZy$/7SmJmEbmbVMQW13kxaIg.' -d:/home/john -s:/bin/bash -
uid:510 -g:users -c:”John Johnson” john
```

dbix

```
dbix {
[{-d|--delete} <context_1>[ <context_2> ...<context_n>]]|
[{-i|--import} <context>] |
[{-x|--export} <context_1>[ <context_2> ...<context_n>]]|
[{-usage|-help|-?}]
}
```

Description

The dbix application provides support for importing, exporting, and deleting Management Center database entries. The application uses the standard input and output streams for reading and writing data, and the delete and export options accept an optional space-delimited list of contexts (a *context* refers to the path to the database attributes on which to perform the operation).

Parameters

[{-d|--delete} *<context_1>***[** *<context_2> ...<context_n>***]]**
Delete entries under the specified context(s).

[{-i|--import} *<context>***]**
Import entries from stdin.

[{-x|--export} *<context_1>***[** *<context_2> ...<context_n>***]]**
Export entries for the specified context(s) to stdout.

[{-usage|-help|-?}] (Optional) Display help information for the command and exit. All other options are ignored.

Examples

EXAMPLE 1

Export the entire database to a file:

```
dbix -x > cwx.4.0-May.20.2007.db
```

EXAMPLE 2

Export the hosts section of the database to a file:

```
dbix -x hosts > cwx.4.0-hosts.db
```

EXAMPLE 3

Delete the entire database:

```
dbix -d
(confirm action)
```

EXAMPLE 4

Import a new database (or additions):

```
dbix -i < cwx.4.0-new_hosts.db
```

dbx

```
dbx {
[{--domain:|-d} <domain>] [{--format:|-f:} <format>] [{-usage|-help|-?}] [-
runtime[:verbose]]
[-signature] [-splash]
}
```

Description

This utility exports specific file formats from the database. Supported formats include a simple host name list typically used for mpich, pdsh, etc., an IP address to host name map (*/etc/hosts*), and configuration files for powerman and conman.

Parameters

Note

Arguments and option values are case sensitive. Option names are not.

[{--domain:|-d} *<domain>***]**	(Optional) Domain name.
[{--format:|-f:} *<format>***]**	(Optional) Output file format. Supported formats are defined as follows: **names** Simple host name list. **hosts** IP address to host name map. **powerman** Powerman configuration file. **conman** Conman configuration file.
[{-usage|-help|-?}]	(Optional) Display help information for the command and exit. All other options are ignored.
[-runtime[:verbose]]	(Optional) Provides specific information about the current Java runtime environment.
[-signature]	(Optional) Displays the application signature. The application signature contains the name, description, version, and build information of this application.
[-splash]	(Optional) Enables the presentation of the application caption or splash screen. By default, on.

Examples

EXAMPLE 1

Use dbx to configure a powerman.conf file:

```
dbx -f:conman > /etc/conman.conf
```

EXAMPLE 2

Use dbx to configure a hosts file:

```
dbx -f:hosts -d:sgi.com > /etc/hosts
```

imgr

```
imgr {
{--image:|-i:}<image> [{--kernel:|-k:}<kernel>] [{--kernel-
revision:|-K:}<kernel_revision>]
[{--payload:|-p:}<payload>] [{--payload.revision:|-P:}<payload_revision>]
[{--force:|-f:}] [{--list:|-l:}]|
[{-usage|-help|-?}]
}
```

Description

The imgr command is used to modify the kernel or payload of an existing image. To create a new image, please refer to *Image Management* on page 116. The Imaging CLI allows you to perform the following operations:

- Specify a kernel for an image
- Specify a payload for an image

Note

If you change a kernel or payload, Management Center rebuilds the image but still requires that you commit the image to VCS. See *vcs* on page 261.

Parameters

{--image:|-i:}<*image*>	The name of the image to modify. By default, Management Center selects the version of the image that was most recently checked in.
[{--kernel:|-k:}<*kernel*>**]**	(Optional) The name of the kernel to modify.
[{--kernel-revision:|-K:}	(Optional) Specify which kernel revision to use. If you do not specify a revision, you will be asked whether or not to use the latest revision.
[{--payload:|-p:}<*payload*>**]**	(Optional) The name of the payload.
[{--payload.revision:|-P:}<*payload_revision*>**]**	(Optional) Specify which payload revision to use. If you do not specify a revision, you will be asked whether or not to use the latest revision.
[{--force:|-f:}]	(Optional) Select the force option to automatically select the latest revision of a payload or kernel. Selecting this option suppresses the prompt that asks you whether or not to use the latest revision.
[{--list:|-l:}]	(Optional) Display a list of working images.
[{-usage|-help|-?}]	(Optional) Display help information for the command and exit. All other options are ignored.

Examples

Update image *Compute* to use revision *4* of *kernel-2.4*:

```
imgr -i:Compute -k:linux-2.4 -K:4
```

To use the latest revision of a payload in an image:

```
imgr -i:MyImage -p:MyPayload
You have not specified the payload revision (latest is 1)
Using latest revisions, continue (yes/no)?
yes
```

kmgr

```
kmgr {
{--name:|-n:}<name> [{--description:|-d:}“<description>”]
{--path:|-p:}<path_to_Linux_kernel_source> [{--kernel:|-k:}<name_of_binary>]
[{--architecture:|-a:}<architecture>] [{--modules:|-m:}] [{--binary:|-b:}] [{--list:|-
l:}]|
[{-usage|-help|-?}]
}
```

Description

The kmgr command is used to create a kernel package from a binary kernel or from a kernel source directory. The utility copies the binary kernel, *.config*, *System.map*, and modules to the kernel directory.

Parameters

{--name:|-n:}*<name>*	The kernel name.
[{--description:|-d:}:“*<description>***”]**	(Optional) A brief description of the kernel.
{--path:|-p:}*<path_to_Linux_kernel_source>*	The path to the kernel source.
[{--kernel:|-k:}*<name_of_binary>***]**	(Optional) The binary name of the kernel. By default, **arch/***<architecture_selected>***/boot/bzImage**
[{--architecture:|-a:}*<architecture>***]**	(Optional) The kernel architecture: amd64 or ia32 (by default, ia32).
[{--modules:|-m:}]	(Optional) The absolute path to *lib/modules/<kernel_version>*.
[{--binary:|-b:}]	(Optional) Enable support for binary kernels.
[{--list:|-l:}]	(Optional) Display a list of working kernels.
[{-usage|-help|-?}]	(Optional) Display help information for the command and exit. All other options are ignored.

Example 1

Create a new kernel named *linux-2.4*:

```
kmgr -n:linux-2.4 -p:/usr/src/linux-2.4.20-8 -a:i386
```

Example 2

Create a new kernel, *linux-2.6*, from a binary kernel:

```
kmgr -b -n:linux-2.6 -k:/boot/vmlinuz-2.6.16-smp -a:x86_64 -d:”Linux 2.6.16 SMP kernel”
```

pdcp

```
pdcp {[
[-w <host>[,<host>...,<host_n>]]|
[-x <host>[,<host>...,<host_n>]]|
[-a]|
[-i]|
[-r]|
[-p]|
[-q]|
[-f <number>]|
[-l <user>]|
[-t <seconds>]|
[-d]]
<source>[ <source>... <source_n>]
<destination>
}
```

Description

Pdcp is a parallel copy command used to copy files from a Master Host to all or selected hosts in the cluster. Unlike rcp which copies files only to an individual host, pdcp can copy files to multiple remote hosts in parallel. When pdcp receives SIGINT (Ctrl+C), it lists the status of current threads. A second SIGINT within one second terminates the program.

Parameters

TARGET HOST LIST OPTIONS

Note

If you do not specify any of the following options, the WCOLL environment variable must point to a file that contains a list of hosts, one per line.

[-w *<host>***[,***<host>***...,***<host_n>***]]**

(Optional) Execute this operation on the specified host(s). You may enter a range of hosts or a comma-delimited list of hosts (e.g., **host[1-4,7,9]**). Any list that consists of a single “-” character causes pdsh to read the target hosts from stdin, one per line.

Note

No spaces are allowed in comma-delimited lists.

[-x *<host>***[,***<host>***...,***<host_n>***]]**	(Optional) Exclude the specified hosts from this operation. You may enter a range of hosts or a comma-delimited list of hosts (e.g., **host[1-4,7,9]**). You may use this option in conjunction with other target host list options such as **-a**.
[-a]	(Optional) Perform this operation on all hosts in the cluster.
[-i]	(Optional) Use this option in conjunction with **-a** or **-g** to request canonical host names. By default, pdsh uses reliable host names.

Note

Gender or -g classifications are not currently supported in this version of pdsh.

[-r]	(Optional) Copy recursively.
[-p]	(Optional) Preserve modification time and modes.
[-q]	(Optional) List option values and target hosts.
[-f <*number*>**]**	(Optional) Set the maximum number of simultaneous remote copies (by default, 32).
[-l <*user*>**]**	(Optional) This option allows you to copy files as another user, subject to authorization. For BSD rcmd, the invoking user and system must be listed in the user's **.rhosts* file (even for root).
[-t <*seconds*>**]**	(Optional) Set the connect time-out (by default, 10 seconds)—this is concurrent with the normal socket level time-out.
[-d]	(Optional) Include more complete thread status when receiving SIGINT and, when finished, display connect and command time statistics on stderr.
<*source*>**[** <*source*>... <*source_n*>**]**	List the source file(s) you want to copy from the Master Host. To copy multiple files, enter a space-delimited list of files (e.g., **pdcp -a /source1 /source2 /source3 / destination**).

Note

The destination is always the last file in the list.

<*destination*>	The location to which to copy the file. The destination is set off from the source by a space.

Example 1

Copy */etc/hosts* to foo01–foo05:

```
pdcp -w foo[01-05] /etc/hosts /etc
```

Example 2

Copy */etc/hosts* to foo0 and foo2–foo5:

```
pdcp -w foo[0-5] -x foo1 /etc/hosts /etc
```

Example 3

To copy a file to all hosts in the cluster:

```
pdcp -a /etc/hosts /etc/
```

Example 4

To copy a directory recursively:

```
pdcp -a -r /scratch/dir /scratch
```

Example 5

To copy multiple files to a directory

```
pdcp -a /etc/passwd /etc/shadow /etc/group /etc
```

pdsh

```
pdsh {
[[-w <host>[,<host>...,<host_n>]]|
[-x <host>[,<host>...,<host_n>]]|
[-a]|
[-i]|
[-q]|
[-f <number>]|
[-s]|
[-l <user>]|
[-t <seconds>]|
[-u <seconds>]|
[-n <tasks_per_host>]|
[-d]|
[-S]|
<host>[,<host>...,<host_n>]]
<command>
}
```

Description

Note

To use pdsh, it must be installed and configured. You can obtain pdsh from *http://sourceforge.net/projects/pdsh/*.

Pdsh is a variant of the rsh command. However, unlike rsh which runs commands only on an individual host, pdsh allows you to issue parallel commands on groups of hosts. When pdsh receives SIGINT (Ctrl+C), it lists the status of current threads. A second SIGINT within one second terminates the program. If set, the DSHPATH environment variable is the PATH for the remote shell.

If a command is not specified on the command line, pdsh runs interactively, prompting for commands, then executing them when terminated with a carriage return. In interactive mode, target hosts that time-out on the first command are not contacted for subsequent commands. Commands prefaced with an exclamation point are executed on the local system.

Parameters

TARGET HOST LIST OPTIONS

[-w *<host>***[,***<host>***...,***<host_n>***]]**

(Optional) Execute this operation on the specified host(s). You may enter a range of hosts or a comma-delimited list of hosts (e.g., **host[1-4,7,9]**). Any list that consists of a single “-” character causes pdsh to read the target hosts from stdin, one per line.

Note

No spaces are allowed in comma-delimited lists.

[-x *<host>***[,***<host>***...,***<host_n>***]]**

(Optional) Exclude the specified hosts from this operation. You may enter a range of hosts or a comma-delimited list of hosts (e.g., **host[1-4,7,9]**). You may use this option in conjunction with other target host list options such as **-a**.

[-a]

(Optional) Perform this operation on all hosts in the cluster. By default, a list of all hosts installed in the cluster is available under */etc/pdsh/machines*.

[-i]	(Optional) Use this option in conjunction with **-a** or **-g** to request canonical host names. By default, pdsh uses reliable host names.

Note

Gender or -g classifications are not currently supported in this version of pdsh.

[-q]	(Optional) List option values and target hosts.
[-f <*number*>**]**	(Optional) Set the maximum number of simultaneous remote commands (by default, 32).
[-s]	(Optional) Combine the remote command stderr with stdout. Combining these commands saves one socket per connection but breaks remote cleanup when pdsh is interrupted with a Ctrl+C.
[-l <*user*>**]**	(Optional) This option allows you to run remote commands as another user, subject to authorization. For BSD rcmd, the invoking user and system must be listed in the user's **.rhosts* file (even for root).
[-t <*seconds*>**]**	(Optional) Set the connect time-out (by default, 10 seconds)—this is concurrent with the normal socket level time-out.
[-u <*seconds*>**]**	(Optional) Limit the amount of time a remote command is allowed to execute (by default, no limit is defined).
[-n <*tasks_per_host*>**]**	(Optional) Set the number of tasks spawned per host. In order for this to be effective, the underlying remote shell service must support spawning multiple tasks.
[-d]	(Optional) Include more complete thread status when receiving SIGINT and, when finished, display connect and command time statistics on stderr.
[-S]	(Optional) Return the largest of the remote command return values.
<*host*>**[,**<*host*>*...*,<*host_n*>**]**	The name of the host(s) on which to execute the specified operation. You may enter a range of hosts or a comma-delimited list of hosts (e.g., **host[1-4,7,9]**).

Note

No spaces are allowed in comma-delimited lists.

<*command*>	The command you want to execute on the host(s).

Example 1

Run a command on foo7 and foo9–foo15:

```
pdsh -w foo[7,9-15] <command>
```

Example 2

Run a command on foo0 and foo2–foo5:

```
pdsh -w foo[0-5] -x foo1 <command>
```

Example 3

In some instances, it is preferable to run pdsh commands using a pdsh shell. To open the shell for a specific group of hosts, enter the following:

```
pdsh -w foo[0-5]
```

From the shell, you may enter commands without specifying the host names:

```
pdsh> date
```

To exit the pdsh shell, type exit.

pmgr

```
pmgr {
[[{--description:|-d:}“<description>”] [{--include:|-i:}<include_file_or_directory>]
[{--include-from:|-I:}<file_containing_list>] [{--location:|-l:}<location_dir>]
[{--silent:|-s:}<silent>]
[{--exclude:|-x:}<exclude_file_or_dir>]] [{--exclude-from:|-X:}<file_containing_list>]
<payload_name>| [{-usage|-help|-?}]
}
```

Description

The pmgr utility generates a Management Center payload from an existing Linux installation to use on a specified host—however, Management Center services must be running on the remote host. An exclude list (or file) allows you to manage which files and directories you want to exclude from the payload (e.g., remote NFS mounted directories or */proc*).

Parameters

[-d:“*<description>***”]**	(Optional) The description of the payload.
[-i:*<include_file_or_directory>***]**	(Optional) Enter the name of the file or directory to include in the payload. When you specify a directory, the payload will include all files and subdirectories contained in the directory.

Tip

To include a previously excluded item (i.e., a file or directory contained in an excluded directory), enter the name of the file or subdirectory.

[{--include-from:|-I:}*<file_containing_list>***]**	(Optional) Enter the name of the file that contains a list of all files to include in the payload.
[-l:*<location_dir>***]**	(Optional) The directory in which to create the payload. By default, the user's payload working directory with the payload name appended.
[-s:*<silent>***]**	(Optional) Omit all output other than errors, including the payload creation progress meter and final summary. This is useful when scripting pmgr.
[-x:*<exclude_file_or_dir>***]**	(Optional) Exclude the named file or directory from the payload. Excluding a directory excludes all files and subdirectories.
[{--exclude-from:|-X:}*<file_containing_list>***]**	(Optional) Enter the name of the file that contains a list of all files to exclude from the payload.
<payload_name>	The name of the payload.
[{-usage|-help|-?}]	(Optional) Display help information for the command and exit. All other options are ignored.

Example

The following example demonstrates how to create a new payload from an existing host installation, *n2*, and exclude some unwanted directories from the payload:

```
pmgr -x:/proc:/home:/var/log:/dev/pts:/mnt -h=n2 n2_payload
```

powerman

```
powerman {
[[{--on|-1}]|
[{--off|-0}]|
[{--cycle|-c}]|
[{--reset|-r}]|
[{--flash|-f}]|
[{--unflash|-u}]|
[{--list|-l}]|
[{--query|-q}]|
[{--node|-n}]|
[{--beacon|-b}]|
[{--temp|-t}]|
[{--help|-h}]|
[{--license|-L}]|
[{--destination|-d} host[:port]]|
[{--version|-V}]|
[{--device|-D}]|
[{--telemetry|-T}]|
[{--exprange|-x}]]
<host>[ <host> ...<host_n>]
}
```

Description

Note

To use Powerman for power control (that is, as your platform management device), Powerman must be installed and configured. You can obtain Powerman from *http://sourceforge.net/projects/powerman/*.

Powerman offers power management controls for hosts in clustered environments. Controls include power on, power off, and power cycle via remote power control (RPC) devices. Target host names are mapped to plugs on RPC devices in *powerman.conf*.

Parameters

[{--on|-1}]	(Optional) Power hosts On.
[{--off|-0}]	(Optional) Power hosts Off.
[{--cycle|-c}]	(Optional) Cycle power to hosts.
[{--reset|-r}]	(Optional) Assert hardware reset for hosts (if implemented by RPC).
[{--flash|-f}]	(Optional) Turn beacon On for hosts (if implemented by RPC).
[{--unflash|-u}]	(Optional) Turn beacon Off for hosts (if implemented by RPC).
[{--list|-l}]	(Optional) List available hosts. If possible, output is compressed into host ranges.
[{--query|-q}]	(Optional) Query plug status of a host(s). If you do not specify a host(s), powerman queries the plug status of all hosts. Status is not cached—powermand queries the appropriate RPC's each time you use this option. Hosts connected to RPC's that cannot be contacted (e.g., due to network failure) are reported as status unknown. If possible, output is compressed into host ranges.
[{--node|-n}]	(Optional) Query host power status (if implemented by RPC). If you do not specify a host(s), powerman queries the power status of all hosts. Please note that this option

	returns the host's power status only, not its operational status. A host in the Off state could be On at the plug and operating in standby power mode.
[{--beacon|-b}]	(Optional) Query beacon status (if implemented by RPC). If you do not specify a host(s), powerman queries the beacon status of all hosts.
[{--temp|-t}]	(Optional) Query host temperature (if implemented by RPC). If you do not specify a host(s), powerman queries the temperature of all hosts. Temperature information is not interpreted by powerman and is reported as received from the RPC on one line per host, prefixed by the host name.
[{--help|-h}]	(Optional) Display option summary.
[{--license|-L}]	(Optional) Show powerman license information.
[{--destination|-d} host[:port]]	(Optional) Connect to a powerman daemon on a non-default host and optional port.
[{--version|-V}]	(Optional) Display the powerman version number.
[{--device|-D}]	(Optional) Display RPC status information. If you specify a host(s), powerman displays only RPC's that match the host list.
[{--telemetry|-T}]	(Optional) Displays RPC telemetry information as commands are processed. This is useful for debugging device scripts.
[{--exprange|-x}]	(Optional) Expand host ranges in query responses.
*<host>***[** *<host> ...<host_n>***]**	The name of the host(s) on which to execute the specified operation. You may enter a range of hosts or a space- or comma-delimited list of hosts (e.g., **host[1-4 7 9]** or **host[1-4 7,9]**).

FILES

```
/usr/sbin/powermand
/usr/bin/powerman
/usr/bin/pm
/etc/powerman/powerman.conf
/etc/powerman/*.dev
```

Example 1

To power on hosts bar, baz, and n01–n05:

```
powerman --on bar baz n[01-05]
```

Example 2

To turn off hosts n4 and n7–n9:

```
powerman -0 n4,n[7-9]
```

vcs

```
vcs {
[{identify| id}]|
[status]|
[include <files>]|
[exclude <files>]|
[archive <filename>]|
[import -R:<repository> -M:<module> [-n:<name>] [-d:“<description>”] [<files>]]|
[commit [-n:<name>] [-d:“<description>”] [<files>]]|
[branch [-n:<name>] [-d:“<description>”] [<files>]]|
[{checkout | co} -R:<repository> -M:<module> [-r:<revision>|<branch>|<name>]]|
[{update | up} [-r:<revision>|<branch>|<name>] [<files>]]|
[name [-R:<repository>] [-M:<module>] [-r:<revision>|<branch>|<name>] <text>]|
[describe [-R:<repository>] [-M:<module>] [-r:<revision>|<branch>|<name>] <text>]|
[{narrate | log} [-R:<repository> -M:<module>] [-r:<revision>|<branch>|<name>]]|
[iterate [-R:<repository> [-M:<module> [-r:<revision>|<branch>|<name>]]]]|
[list]|
[{-usage|-help|-?}]
}
```

Description

Manage version controlled directories within Management Center.

Parameters

[{identify| id}]	(Optional) Display information about the module contained in the current working directory.
[status]	(Optional) Display the status of the files within the current working directory including whether they have been added (A), modified (M) or deleted (D).
[include *<files>***]**	(Optional) Add provided list of files to the include list. You may also use this option to override a specific file exclusion.
[exclude *<files>***]**	(Optional) Add provided list of files to the exclude list. Excluding files allows you to remove files that may cause problems (e.g., when trying to archive files).
[archive *<filename>***]**	(Optional) Create an archive of the current working directory in the given file. This option may be used to archive a host and include it in VCS as a payload.
[import -R:*<repository>* **-M:***<module>* **[-n:***<name>***] [-d:“***<description>***”] [***<files>***]]**	(Optional) Create a new module with the provided list of files or all of the current working directory.
[commit [-n:*<name>***] [-d:“***<description>***”] [***<files>***]]**	(Optional) Insert a new revision in the module using the provided list of files or any working copy modifications.
[branch [-n:*<name>***] [-d:“***<description>***”] [***<files>***]]**	(Optional) Insert a new revision that is not on tip using the provided list of files or any working copy modifications.
[{checkout| co} -R:*<repository>* **-M:***<module>* **[-r:***<revision>*|*<branch>*|*<name>***]]**	(Optional) Retrieve an existing revision from a module. The contents of the module will be stored in a new directory named after the module.

[{update| up} [-r:<*revision*>|<*branch*>|<*name*>**] [**<*files*>**]]**	(Optional) Update the current directory to use the latest tip revision of a branch (*3.4*), the main trunk of a specific branch (*4*), or a branch with a specific name (*Golden*). The files option allows you to update a specific file contained in a payload.
[name [-R:<*repository*>**] [-M:**<*module*>**] [-r:**<*revision*>|<*branch*>|<*name*>**]** <*text*>**]**	(Optional) Add, modify or delete the optional name or alias of a revision. Names are unique revision identifiers for the entire module. A blank for the name will delete the previous value.
[describe [-R:<*repository*>**] [-M:**<*module*>**] [-r:**<*revision*>|<*branch*>|<*name*>**]** <*text*>**]**	(Optional) Add, modify or delete the optional description of a revision. A blank for the description will delete the previous value.
[{narrate| log} [-R:<*repository*> **-M:**<*module*>**] [-r:**<*revision*>|<*branch*>|<*name*>**]]**	(Optional) Display the history of a module revision.
[iterate [-R:<*repository*> **[-M:**<*module*> **[-r:**<*revision*>|<*branch*>|<*name*>**]]]]**	(Optional) Display the organizational information of the version service.
[list]	(Optional) Display a list of all category types (payloads, kernels, and images) that have been checked into VCS.
[{-usage|-help|-?}]	(Optional) Display help information for the command and exit. All other options are ignored.

Examples

EXAMPLE 1

Display a list of images contained in the Version Control System:

```
vcs iterate -R:images
```

EXAMPLE 2

Display a list of files that have changed since the last time the Compute payload was checked out:

```
cd $MGR_HOME/imaging/root/payloads/Compute
vcs status
```

EXAMPLE 3

List current versions of all category types (payloads, kernels, and images) checked into VCS:

```
vcs list
```

```
Images
---------------------------------------------------------
    MyImage (1) - Kernel: MyKernel (3) Payload: MyPayload (6.1.4)
    TestImage (1) - Kernel: Compute (2) Payload: SLES10 (23)

Kernels
---------------------------------------------------------
    MyKernel (5)
    Compute (2)

Payloads
---------------------------------------------------------
    MyPayload (6.1.7)
    SLES9 (34)
    SLES10 (23)
```

EXAMPLE 4

Check out a specific revision, *8*, of a version controlled payload named *Compute*:

```
vcs checkout -R:payloads -M:Compute -r:8
```

EXAMPLE 5

Use VCS to make sure you have the latest revision of what was originally checked out in the previous example:

```
cd $MGR_HOME/imaging/<username>/payloads/Compute
vcs update
```

Appendix

Pre-configured Metrics

The *CustomMetrics.profile* is the file used to define which metrics are available in the Add a Custom Metric Listener dialog. The *Metrics.profile* is the file used to define which metrics are available from the Metrics Selector dialog to view in the instrumentation service.

Both the *Metrics.profile* and *CustomMetrics.profile* use the same format and need to be edited only if you have written a custom monitoring script and configured it as a custom monitor. Then, if you want to:

- Display the custom metrics in the List View, add the new metrics to the *Metrics.profile*.
- Set thresholds on the custom metrics, add the new metrics to the *CustomMetrics.profile*.

CPU

Metric Name	Format and Description
CPU Percent Idle Aggregate	`hosts.{host.moniker}cpu.idle.pattern=100%`
	Percentage of time the CPU is idle.
CPU Percent I/O Wait Aggregate	`hosts.{host.moniker}.cpu.iowait.pattern=100%`
	The total cycles used by all CPUs waiting for I/O.
CPU Percent Nice Aggregate	`hosts.{host.moniker}.cpu.nice.pattern=100%`
	The total cycles used by all CPUs in user mode with low priority.
CPU Percent System Aggregate	`hosts.{host.moniker}.cpu.system.pattern=100%`
	The total cycles used by all CPUs in kernel mode.
CPU Percent User Aggregate	`hosts.{host.moniker}.cpu.user.pattern=100%`
	The total cycles used by all CPUs in user mode.

Disk

Metric Name	Format and Description
Disk Reads (blocks per second)	`hosts.{host.moniker}.disks.hda.block.reads.pattern=000000`
	The number of blocks read from a disk.
Disk Writes (blocks per second)	`hosts.{host.moniker}.disks.hda.block.writes.pattern=000000`
	The number of blocks written to a disk.
Disk I/O Read (bytes per second)	`hosts.{host.moniker}.disks.hda.io.reads.pattern=000000`
	The number of I/O reads from a disk.
Disk I/O Writes (bytes per second)	`hosts.{host.moniker}.disks.hda.io.writes.pattern=000000`
	The number of I/O writes to a disk.
Disk (hda[1-4]) Capacity Used (bytes)	`hosts.{host.moniker}.disks.hda1.capacity.used.pattern=0,000 MB`
	The disk capacity used for disk hda1, hda2, hda3, or hda4.
Disk (hda[1-4]) Capacity Free (bytes)	`hosts.{host.moniker}.disks.hda1.capacity.free.pattern=0,000 MB`
	The disk capacity free for disk hda1, hda2, hda3, or hda4.
Disk (hda[1-4]) Percentage Used	`hosts.{host.moniker}.disks.hda1.percentage.used.pattern=100%`
	The disk percentage used for disk hda1, hda2, hda3, or hda4.
Disk (sda[1-4]) Capacity Used (bytes)	`hosts.{host.moniker}.disks.sda1.capacity.used.pattern=0,000 MB`
	The disk capacity used for disk sda1, sda2, sda3, or sda4.
Disk (sda[1-4]) Capacity Free (bytes)	`hosts.{host.moniker}.disks.sda1.capacity.free.pattern=0,000 MB`
	The disk capacity free for disk sda1, sda2, sda3, or sda4.
Disk (sda[1-4]) Percentage Used	`hosts.{host.moniker}.disks.sda1.percentage.used.pattern=100%`
	The disk percentage used for disk sda1, sda2, sda3, or sda4.

Kernel

Metric Name	Format and Description
Kernel Context Switches (per second)	`hosts.{host.moniker}.kernel.contexts.pattern=000000`
	The number of context switches the system has undergone.
Kernel Interrupts (per second)	`hosts.{host.moniker}.kernel.interrupts.pattern=000000`
	The number of interrupts received from the system since boot.
Kernel Running Processes	`hosts.{host.moniker}.kernel.processes.pattern=00000`
	The number of forks since boot.

Metric Name	Format and Description
Kernel Swaps In	`hosts.{host.moniker}.kernel.swaps.in.pattern=000000` The number of swap pages that have been brought in.
Kernel Swaps Out	`hosts.{host.moniker}.kernel.swaps.out.pattern=000000` The number of swap pages that have been sent out.

Load

Metric Name	Format and Description
Load - 15 Minute	`hosts.{host.moniker}.load.15m.pattern=0.00` The number of tasks in the run state averaged over 15 minutes.
Load - 1 Minute	`hosts.{host.moniker}.load.1m.pattern=0.00` The number of tasks in the run state averaged over 1 minute.
Load - 5 Minute	`hosts.{host.moniker}.load.5m.pattern=0.00` The number of tasks in the run state averaged over 5 minutes.

Memory

Metric Name	Format and Description
Memory Active (bytes)	`hosts.{host.moniker}.memory.active.pattern=0,000 MB` The amount of active memory.
Memory Cached (bytes)	`hosts.{host.moniker}.memory.cached.pattern=0,000 MB` The amount of cached memory.
Memory Used (bytes)	`hosts.{host.moniker}.memory.committed.pattern=0,000 MB` The amount of used memory.
Memory Free (bytes)	`hosts.{host.moniker}.memory.free.pattern=0,000 MB` The total amount of free memory.
Memory Swap Cached (bytes)	`hosts.{host.moniker}.memory.swap.cached.pattern=0,000 MB` The amount of cached swap.
Memory Swap Free (bytes)	`hosts.{host.moniker}.memory.swap.free.pattern=0,000 MB` The amount of free swap space.
Memory Total (bytes)	`hosts.{host.moniker}.memory.total.pattern=0,000 MB` The total amount of memory.

Network

Metric Name	Format and Description
Network (eth0) Bytes Received (per second)	`hosts.{host.moniker}.network.eth0.rx.bytes.pattern=0,000 MB` The total number of bytes received on all interfaces.
Network (eth0) Packets Received (per second)	`hosts.{host.moniker}.network.eth0.rx.packets.pattern=0,000 MB` The total number of received packets on all interfaces.
Network (eth0) Bytes Transmitted (per second)	`hosts.{host.moniker}.network.eth0.tx.bytes.pattern=0,000 MB` The total number of transmitted bytes on all interfaces.
Network (eth0) Packets Transmitted (per second)	`hosts.{host.moniker}.network.eth0.tx.packets.pattern=0,000 MB` The total number of packets transmitted on all interfaces.
Network (eth1) Bytes Received (per second)	`hosts.{host.moniker}.network.eth1.rx.bytes.pattern=0,000 MB` The total number of bytes received on all interfaces.
Network (eth1) Packets Received (per second)	`hosts.{host.moniker}.network.eth1.rx.packets.pattern=0,000 MB` The total number of received packets on all interfaces.
Network (eth1) Bytes Transmitted (per second)	`hosts.{host.moniker}.network.eth1.tx.bytes.pattern=0,000 MB` The total number of transmitted bytes on all interfaces.
Network (eth1) Packets Transmitted (per second)	`hosts.{host.moniker}.network.eth1.tx.packets.pattern=0,000 MB` The total number of packets transmitted on all interfaces.

Glossary

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

Architecture-independent Allows 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.

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

A

accounts
  disable user 71
  enable 71
  manage group 98
  manage local 98
acl_roots 152
add
  boot utilities 134
  custom monitors 182
  directory to payload 102
  file to payload 102
  group 73
    user account to payload 100
  host 19, 49, 207
  kernel modules without loading 114
  listener 192
  local user account to payload 99
  monitor 176
  package
    to existing payload 90
  partition 61
  plug-in 137
  RAID partition 123
  role 76
  user 70
    to group 74
administration levels 69
Altix ICE ix, 15
Altix UV systems 3, 17, 106, 218
AMD GPUs 19
annotations
  electric shock xi
  note xi
  tip xi
  warning xi
anti-aliasing 153, 155
appearance
  interface 23
applications preferences 38
apply listeners
  as default 174, 190
  to hosts 174, 190
  to payloads 175, 191
authentication
  management, payload 96
auto node discovery 19, 207

B

beacon
  turn off 60
  turn on 59
block size 113, 122
boot
  process, plug-ins for 136
  troubleshooting 216
  utilities, add 134
boot.profile 118, 135
branch, version 141

C

chassis management controllers (CMCs) 17
check into VCS
  image 142
  kernel 142
  payload 142
check out of VCS
  image 143
  kernel 143
  payload 143
CLI 269
client platforms 3
cluster 47, 69

- environment 69
- host administration 229
- power administration 238
- provisioning 240
- system monitoring 153
- user administration 243

CMCs 17
command-line interface 219, 269
- conman 226
- cwhost 229
- cwpower 238
- cwprovision 240
- cwuser 243
- dbix 249
- dbx 250
- imgr 251
- kmgr 252
- pdcp 253
- pdsh 255
- pmgr 258
- powerman 259
- vcs 261

compute host (*See* host)
configure
- NIS 96

conman 29, 226
connect to console 60
console 60
copy
- from VCS 144
- image 118
- kernel 109
- payload 87

CPU
- metrics 265
- tab 159
- utilization 159

create
- group 73
- host 49
- image 116
- kernel 107
- kernel from binary 252
- multiple payloads from source 85
- partition 61, 120
- password
  - Icebox 235
  - user 71, 244
- payload 84
- region 63
- role 76

csv 55
Customer Service xii, 213, 214
customize the interface 23
cwhost 229
cwpower 238
cwprovision 240
cwuser 243

D

Data Center Manager (DCM) 2, 32
dbix 55, 249
dbx 250
DCM 2, 32
debugging 214
default user administration settings 70
delete
- all payloads, kernels, and images 144
- file(s) from payload 104
- group 75
  - account from payload 101
- host 55
- image partition 127
- listener 196
- local user account from payload 100
- monitor 180
- package from payload 90
- partition 62
- payload 104
- region 65
- role 78
- user account 72
- working copy of image 119
- working copy of kernel 115
- working copy of payload 104

dependency checks, package 93
DHCP 3, 44
dhcpd.conf 44
disable
- anti-aliasing 153
- gradient fill 153
- Kerberos 97
- LDAP 97
- listener 190
- monitor 174
- NIS 96
- user account 71, 72

Discover interface 19, 207
disk

- aggregate usage 161
- fill to end of 132
- I/O 161
- metrics 266
- tab 161

diskless hosts 131
- configure 131
- mount point 132

display
- custom metrics 185

distribution, upgrade 2
dmesg.level 118
DNS name resolution 92
dockable frames 22, 24
Documentation
- available via the World Wide Web x

DRAC 1, 29, 53
Dynamic Host Configuration Protocol (DHCP) 3

E

edit
- group 75
- host 54
- Icebox password 235
- image partition 125
- kernel 113
- listener 195
- monitor 180
- partition 62
- password 72, 245
- payload 93
  - using text editor 103
- region 64
- role 78
- user account 72

electric shock xi
enable
- anti-aliasing 153
- gradient fill 153
- Kerberos 97
- LDAP 97
- listeners 190
- monitor 174
- NIS 96
- user account 71

environment monitoring 166, 172
environmental tab 165
errors
- messages 198
- RPM 87
- troubleshooting 213

event
- listeners 189
- log 154, 198
- monitoring 172

exclude
- files and directories from VCS 146

exclude file(s) from payload 89

F

features, Management Center 16
feedback, documentation xi
file
- exclude file(s) from payload 89
- system, user-defined 128

fill to end of disk 132
filter 155, 157
find
- host 55

format partition 122
frames
- controls 22
- dockable 22, 24

FreeIPMI 29
Freeipmi 3
fstab 122, 132

G

general preferences 27
general tab 156
GID 70, 74
GPU monitoring 19, 167
gradient fill 153, 155
group 69, 73
- add 73
  - account to payload 100
- assign roles to 74
- assign to role 77
- assign user to 71
- delete 75
  - account from payload 101
- edit 75
- GID 70, 74
- grant access to region 74
- power 71, 73
- primary 71
- region, add to 64

    - root 73
    - user membership 71
    - users 73

H

- halt host 58
- hardware
    - system requirements 1
- health
    - monitoring 153
        - event log 154
        - system status icons 154
    - status 156
- host 47, 69
    - add 19, 49, 207
        - to partition 61
    - administration 47, 69
        - grant privileges 79
    - beacon
        - turn off 60
        - turn on 59
    - CLI administration 230
    - configure
        - diskless host 131
    - cycle power to 59
    - delete 55
    - diskless 131
    - edit 54
    - event log 154
    - find 55
    - halt 58
    - import 55
    - load monitoring 164
    - Management Center Master 47
        - rename 54
    - names 4
    - power
        - turn off 58, 59
        - turn on 59
    - power management 58
    - provision 147
        - using CLI 240
    - reboot 58
    - region
        - add host to 63
        - assign host to 51
    - reset 59
    - shared 47, 69
    - shut down 58
    - states 154
    - upgrade 150

I

- Icebox
    - administration privileges 79
    - create password for 235
    - modify password 235
- icons, system status 154
- ILO 1, 29, 53
- image 81, 118
    - add modules without loading 114
    - check into VCS 142
    - check out from VCS 143
    - CLI controls 251
    - copy 118
    - create 116
    - delete all 144
    - delete partition 127
    - delete working copy of image 119
    - edit image partition 125
    - management 116
    - partition 120
    - privileges, enable imaging 79
    - provision 147
        - select image 147
    - versioned 140
    - working copy 140
- image.once 118
- image.path 118
- imgr 251
- import
    - binary kernel 252
    - default listeners 197
    - host list 55
    - listener 196
    - listener from payload 197
    - listeners from payload 197
    - monitors
        - from payload 179
    - monitors from host 178
- informational messages 198
- install
    - Management Center 4
        - client 8
        - into payload 105
- instrumentation 153
    - CPU utilization 159
    - custom monitors 181

- disk
  - aggregate usage 161
  - I/O 161
- enhance performance 153
- event log 154
- health status 156
- host load 164
- kernel information 163
- list view 158
- memory utilization 160
- menu controls 155
- metrics, define 185
- metrics, pre-configured 265
- monitoring and event subsystem 172
- packet transmissions 162
- power status 156
- resource utilization 156
- system configuration 156
- system status 154
  - overview 156
- temperature readings 165, 166
- thumbnail view 157

Intel Data Center Manager (DCM) 32
- DCM
  - Data Center Manager (DCM) 1

Intel Power Node Manager (IPNM) 1, 2, 32
interface
- customized appearance 23
- management 50
- map 20
- split-pane view 23

interval 155
IP address 270
- host 50

IPMI 1, 29, 41, 51
IPMItool 3
ipmitool utility 41
IPNM 1, 2, 32
ISLE Cluster Manager, upgrading from 203

J

Jpackage Utilities 3

K

Kerberos 97
kernel 81
- build from source 107
- check into VCS 142
- check out from VCS 143
- CLI controls 252
- copy 109
- create 107
- create from binary 252
- delete all 144
- delete working copy of kernel 115
- edit 113
- install modules without loading 114
- loadable modules 114
- management 107
- metrics 266
- modular 114
- monolithic 114
- troubleshooting 216
- upgrade 2
- verbosity level 117, 152, 240
- versioned 140
- working copy 140

kernel tab 163
kmgr 252

L

layouts 24
- open saved 25
- save 24
- set default 25

LDAP 97
licensing 12, 217
links, dangling symbolic 89
list view 155, 158
listeners 172, 197
- add 192
- apply as default 174, 190
- apply to hosts 174, 190
- apply to payloads 175, 191
- delete 196
- disable 190
- edit 195
- enable 190
- event 189
- import 196

load
- metrics 267

load tab 164
loadable kernel modules 114
loggers 172, 198
- TemplateFormatter 198

M

MAC addresses 207
maintenance operations 12
management
    interface 50
    management network 4
    VCS 144
Management Center
    administration
        grant privileges 79
    features 16
    install into payload 105
    install on client 8
    interface
        customize 23
        map 20
        split-pane view 23
    introduction 15
    Master Host
        rename 54
    preferences 27
        applications 38
        general 27
        platform management 29
        provisioning settings 39
    Premium Cluster Edition ix, 15
    Premium ICE Edition ix, 15
    Premium UV Edition ix, 15
    product definition ix, 15
    server, start and stop 12
    services 13
    Standard Cluster Edition ix, 15
    upgrading 201
Managment Center
    platforms 1
    system requirements 2
Master Host
    definition 47
    rename 54
    system requirements 1
memlog 3, 17, 171
memory
    estimate partition requirements 122, 132
    metrics 267
    utilization 160
Memory Failure Analysis 3
memory failure analysis 17, 171
memory tab 160
metrics 185, 265
    alignment 187
    CPU 265
    custom 187
    disk 266
    display custom 185
    instrumentation service 155
    kernel 266
    load 267
    memory 267
    metrics selector 187
    network 268
mkfs 122
modules
    install without loading 114
    loadable kernel 114
modules subtab 114
monitoring
    event 172
    system health 153
monitors 172, 173
    add 176
    add custom 182
    custom 181
    delete 180
    disable 174
    edit 180
    enable 174
    import from host 178
    import from payload 179
multicast
    route configuration 45

N

navigation tree 22
netmask 270
network metrics 268
network tab 162
Network Time Protocol (NTP) 3
NFS 70
NIS 96
nodes.conf 55
note xi
NTP 3
NVIDIA GPUs 19

O

open
    layouts 25

operating system requirements 2
override global settings 51
overview, system status 156

P

package
  add to existing payload 90
  dependency checks 93
  remove from payload 93
packet transmissions 162
partition 47, 61, 69, 118
  add 61
    host to 61
    RAID 123
  create 120
    user-defined file system 128
  delete 62
  delete from image 127
  edit 62
  edit image partition 125
  estimate memory requirements 122, 132
  format 122
  manage 120
  overwrite protection 122
  partition this time 152
  partitioning behavior 117
  save 122
  size
    fill to end of disk 122, 130, 132
    fixed 122, 130, 132
partition.once 118
password
  create Icebox 235
  create new 71, 244
  encrypt 245
  modify 72, 245
  modify Icebox 235
payload 81
  account management, local user 98
  add
    directory to 102
    file to 102
    group user account to 100
    local user account to 99
    package to existing 90
  attributes, troubleshoot 86
  authentication management 96
  check into VCS 142
  check out from VCS 143
  check-in error 216
  CLI controls 258
  configure 95
  copy 87
  create 84
    multiple payloads from source 85
  dangling symbolic links 89
  delete 104
    file(s) from payload 104
    group account from payload 101
    local user account from payload 100
    working copy of payload 104
  delete all 144
  download this time 152
  edit
    using CLI 102
    with text editor 103
  exclude file(s) 89
  file configuration 95
  group account management 98
  install Management Center into 105
  management 82
  package dependency checks 93
  pmgr 258
  remove package from 93
  script, enable 95
  update directory 102
  update file 102
  versioned 140
  working copy 140
PBS 151
PBS Professional 5
pdcp 253
pdsh 255
PEKI temperatures 189
permissions 79
  See role; privileges
physical memory utilization 160
plaforms, Management Center 1
platform management 1
  DRAC 53
  ILO 53
  IPMI 51
platform management preferences 29
plug-ins
  add 137
  for boot process 136
power 59
  CLI administration 238
  control 58

    cost tracking and trending 16, 32, 36, 168
    cycle
        to host 59
    group 71, 73
    management 16, 32, 168
    management, host 58
    monitoring 16, 32, 168
    policy 16, 32
    powerman 259
    status 156
    turn off
        to host 59
    turn off host 58
    turn on
        to host 59
powerman 29, 259
pre-configured metrics 265
preferences
    applications 38
    general 27
    Management Center 27
    platform management 29
    provisioning settings 39
Premium Cluster Edition, Management Center ix, 15
Premium ICE Edition, Management Center ix, 15
Premium UV Edition, Management Center ix, 15
primary
    group 71
Prism XL platforms 19
privileges 79
    change user 78
    database 79
    host administration 79
    Icebox administration 79
    imaging 79
    instrumentation 79
    logging 79
    Management Center 79
    power 79
    provisioning 79
    serial 79
    user administration 79
problems 213
product definition, Management Center ix, 15
Product Support xii, 213, 214
provision 147
    CLI controls 240
    disable confirmation dialog 149
    enable confirmation dialog 149
    format partition 122
    provisioning settings preferences 39
    right-click 149
    schedule at next reboot 151
    select an image 147
    troubleshooting 215, 218

Q

qmgr 152

R

racks 47, 66
RAID 123
RAM Disk 134
    block size 113
RAMfs 131
reboot host 58
region 47, 63, 69
    add
        group to 64
        host to 63
    assign to host 51
    create 63
    delete 65
    edit 64
    grant group access to 74
remove
    file(s) from payload 104
    group 75
    group account from payload 101
    host 55
    local user account from payload 100
    package from payload 93
    partition 62
    region 65
    role 78
    user account 72
rename
    host 54
    Management Center Master Host 54
requirements
    hardware 1
    operating system 2
    software 3
reset
    host 59
resource
    utilization 156
restore factory settings 197

RHEL 270
right-click menu 58
    connect to console 60
    provisioning 149
rights
    See role; privileges
Roamer 1, 29
Roamer KVM 60
role 69, 76
    add 76
    assign group to 77
    assign to group 74
    delete 78
    edit 78
    grant privileges and permissions 77
root group 73
routes, multicast 45
RPM
    errors 87
RPM front ends 5

S

save
    layouts 24
    partition 122
scalability 10
schedule provision at next reboot 151
script, enable in payload 95
search, tree 55
server platforms 2
SGI Altix UV systems 3, 17, 106, 218
SGI Foundation Software 17, 171
SGI Prism XL platforms 19
shut down a host 58
size
    thumbnail 157
SLES 271
SMN 17, 18
SMN bundle software 3, 17
software
    requirements 3
sort 155
split-pane view 23
SSL 97
Standard Cluster Edition, Management Center ix, 15
start Management Center server 12
state, host 154
stop Management Center server 12
symbolic links, dangling 89
symlink 89
system
    configuration 156
    health 153
    requirements
        hardware 1
        operating system 2
    status
        event log 154
        icons 154
        overview 156
system management node (SMN) 3, 17, 18

T

task progress dialog 85
Technical Support xii, 213, 214
Telnet client 3
temperature
    changing thresholds 172
    monitoring 155, 166
    PEKI 189
    readings 165
    troubleshooting 215
TemplateFormatter 198
TFTP 3, 45
third-party power controls 71, 73
thumbnail
    size 155
    view 157
thumbnail view 155
tip xi
TmpFS 131
toolbar 21
transmissions, packet 162
Trivial File Transfer Protocol (TFTP) 3, 45
troubleshooting
    general 213
    payload attributes 86
    RPM errors 87

U

UID 70, 71
upgrade
    distribution 2
    kernel 2
    VCS upgrade 150
upgrading Management Center 201
user 69, 70

    - add 70
        - local user account to payload 99
        - to group 74
    - administration 69
        - default settings 70
        - privileges 79
    - assign to group 71
    - CLI administration 243
    - delete
        - local user account from payload 100
    - delete account 72
    - disable account 72
    - edit account 72
    - group membership 71
    - multi-group 69
    - UID 70, 71
- user-defined file system 128
- users group 73
- UV systems 3, 17, 106, 218

V

- VCS 140
    - branch 142
    - CLI controls 261
    - command-line controls 261
    - copy 144
    - exclude files and directories 146
    - management console 144
    - upgrade 150
- verbosity level, kernel 117, 152, 240
- version
    - branching 141
    - control system 140
        - check into 142
        - check out 143
        - vcs command 261
    - versioned copy 140, 271
- VersionControlService.profile 146
- virtual memory utilization 160

W

- warning xi
- warning messages 198
- Windows clients 3
- working copy 140, 271

Y

- YUM 5

Z

- Zypper 5