



**SGI® Management Center™ (SMC)
Installation and Configuration Guide
for Clusters**

007-6359-004

COPYRIGHT

© 2014–2015 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.

The SMC software stack depends on several open source packages which require attribution. They are as follows:

c3:

C3 version 3.1.2: Cluster Command & Control Suite Oak Ridge National Laboratory, Oak Ridge, TN, Authors: M.Brim, R.Flanery, G.A.Geist, B.Luethke, S.L.Scott (C) 2001 All Rights Reserved NOTICE Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted provided that the above copyright notice appear in all copies and that both the copyright notice and this permission notice appear in supporting documentation. Neither the Oak Ridge National Laboratory nor the Authors make any representations about the suitability of this software for any purpose. This software is provided "as is" without express or implied warranty. The C3 tools were funded by the U.S. Department of Energy.

conserver:

Copyright (c) 2000, conserver.com All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. - Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. - Neither the name of conserver.com nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission. THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Copyright (c) 1998, GNAC, Inc. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: - Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. - Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. - Neither the name of GNAC, Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission. THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Copyright 1992 Purdue Research Foundation, West Lafayette, Indiana 47907. All rights reserved. This software is not subject to any license of the American Telephone and Telegraph Company or the Regents of the University of California. Permission is granted to anyone to use this software for any purpose on any computer system, and to alter it and redistribute it freely, subject to the following restrictions: 1. Neither the authors nor Purdue University are responsible for any consequences of the use of this software. 2. The

origin of this software must not be misrepresented, either by explicit claim or by omission. Credit to the authors and Purdue University must appear in documentation and sources. 3. Altered versions must be plainly marked as such, and must not be misrepresented as being the original software. 4. This notice may not be removed or altered.

Copyright (c) 1990 The Ohio State University. All rights reserved. Redistribution and use in source and binary forms are permitted provided that: (1) source distributions retain this entire copyright notice and comment, and (2) distributions including binaries display the following acknowledgment: "This product includes software developed by The Ohio State University and its contributors" in the documentation or other materials provided with the distribution and in all advertising materials mentioning features or use of this software. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission. THIS SOFTWARE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation. This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

pysqlite:

Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

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.

TRADEMARKS AND ATTRIBUTIONS

Altix, ICE, Performance Co-Pilot, Rackable, SGI, the SGI logo, and Supportfolio are trademarks or registered trademarks of Silicon Graphics International Corp. or its subsidiaries in the United States and other countries.

Altair is a registered trademark and PBS Professional is a trademark of Altair Engineering, Inc. Intel, Itanium, Phi, and Xeon are trademarks or registered trademarks of Intel Corporation. InfiniBand is a trademark of the InfiniBand Trade Association. Connect-IB, InfiniScale, and Mellanox are registered trademarks of Mellanox Technologies. Linux is a registered trademark of Linus Torvalds. LSI Logic and MegaRAID are registered trademarks of the LSI Logic Corporation. CentOS, Red Hat, and Red Hat Enterprise Linux are registered trademarks of Red Hat, Inc., in the United States and other countries. SLES, SUSE, and YAST are registered trademarks of SUSE LLC in the United States and other countries.

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

New Features

This revision includes the following new features:

- Information about SMC support for Red Hat Enterprise Linux (RHEL) 7.1.
- Miscellaneous editorial and technical corrections and additions.

Documentation that pertains specifically to SLES 12 can be ignored in this revision because SMC support for SLES 12 is deferred.

Record of Revision

Version	Description
001	November 2014 Original publication. This revision supports the SGI Management Center 3.0 release and the SGI Foundation Software 2.11 release.
002	February 2015 Includes miscellaneous updates and corrections to the previous revision.
003	May 2015 This manual supports the the SGI Management Center 3.1 release.
004	July 2015 This manual supports the the SGI Management Center 3.1.2 release.

Contents

About This Guide	xvii
Related Publications	xvii
Obtaining Publications	xix
Conventions	xix
Reader Comments	xx
1. System Software Overview	1
About SGI Cluster Computer Systems	1
About SGI Rackable Clusters	2
About SGI ICE Clusters	4
About High Availability Nodes in SGI Clusters	4
About Managed Nodes and Unmanaged Nodes	5
SGI Cluster System Node Images	5
Operating System Support	8
SGI Cluster Networks	8
SGI Rackable Networks	8
SGI ICE Networks	10
2. Customizing a Factory-installed SGI Cluster (RHEL 6 and SLES 11)	17
About Customizing a Factory-installed Cluster	17
Obtaining Information	18
Changing the Password and Specifying Network Information	19
Completing the Customization	23
(Conditional) Pushing Changes to SGI ICE Compute Nodes	26
Verifying the System and Configuring Additional Features	28
007-6359-004	ix

3. Customizing a Factory-installed SGI Cluster (RHEL 7 and SLES 12)	31
About Customizing a Factory-installed Cluster	31
Obtaining Information	32
Changing the Password and Specifying Network Information	33
Completing the Customization	38
(Conditional) Pushing Changes to SGI ICE Compute Nodes	41
Verifying the System and Configuring Additional Features	43
4. Installing and Configuring an SGI Cluster System	45
About Performing a New Installation and Configuring the Software on an SGI Cluster	46
(Conditional) Migrating From SMC 1.7.x to SMC 3.x	50
Preparing the Payloads for Migration	50
Backing Up the Payloads and Running the Migration Script	52
Planning the Image Installation Method	53
Preparing to Install Software on a Cluster	56
(Conditional) Setting a Static IP Address for the Baseboard Management Controller (BMC) in the Admin Node	57
(Optional) Configuring a High Availability Admin Node or a High Availability Rack Leader Controller (RLC)	59
Booting the System	59
Configuring the Operating System on the Admin Node	65
Configuring RHEL 7.x on the Admin Node	65
Configuring RHEL 6.x on the Admin Node	69
Configuring SLES 11 on the Admin Node	74
Configuring the Cluster	79
(Conditional) Configuring External Domain Name Service (DNS) Servers	93
Synchronizing the Software Repository, Installing Software Updates, Downloading Debugging Packages, and Cloning the Images	94

(Conditional) Downloading the Intel Manycore Platform Software Stack (MPSS) Software and Creating Images	97
Downloading the MPSS Software From the Intel Corporation	97
Creating Images for the SGI ICE Compute Nodes That Include MIC Devices	99
Creating Images for the Flat Compute Nodes That Include MIC Devices	103
(Conditional) Running the Migration Script and Editing the Cluster Definition File	108
Configuring the Switches	109
About the Cluster Definition File	110
Verifying the Switch Cabling	115
Configuring Management Switches With a Cluster Definition File	118
Configuring Management Switches Without a Cluster Definition File	122
(Conditional) Configuring the Cooling Racks and Cooling Distribution Units (CDUs) on the MCell Network's Switch Ports	126
Configuring the Cluster With the <code>discover</code> Command	129
<code>discover</code> Command Example 1	133
<code>discover</code> Command Example 2	134
<code>discover</code> Command Example 3	136
<code>discover</code> Command Example 4	136
<code>discover</code> Command Example 5	136
<code>discover</code> Command Example 6	137
Verifying Power Operations and Configuring Power Management	137
(Conditional) Verifying That the <code>smcrbd</code> daemon and the <code>icerackd</code> Daemon are Running	143
(Optional) Configuring a Backup Domain Name Service (DNS) Server	146
(Conditional) Configuring the InfiniBand Subnetworks	147
Configuring the InfiniBand Subnetworks	148
Configuring Additional Features	152
5. Configuring Additional Features	155

About Additional Features	155
Configuring Array Services for MPI Programs	155
Planning the Configuration	156
Preparing the Images	159
Configuring the Authentication Files in the New System Images on the Admin Node	161
Permitting Remote Access to the Compute Services Node	162
Preventing Remote Access to the Service Node	163
Distributing Images to all the Nodes in the Array	164
Power Cycling the Nodes and Pushing Out the New Images	166
(Conditional) Enabling the Mellanox OpenFabrics Enterprise Distribution (OFED) for Linux Software	167
Troubleshooting Configuration Changes	174
6. Troubleshooting	175
About Troubleshooting	176
Using the <code>switchconfig</code> Command	176
SGI ICE Compute Nodes Are Taking Too Long To Boot (SGI ICE Clusters Only)	178
Verify the Bonding Mode on the Rack Leader Controller (RLC) (SGI ICE Clusters Only)	179
<code>cimage --push-rack</code> Pushes Too Many (or Too Few) Expansions (SGI ICE Clusters Only)	182
Cannot ping the CMCs from the Rack Leader Controller (RLC) (SGI ICE Clusters Only)	182
Restarting the <code>blademond</code> Daemon (SGI ICE Clusters Only)	185
Log Files	185
CMC <code>slot_map</code> / <code>blademond</code> Debugging Hints	186
Resolving CMC Slot Map Ordering Problems	187
In <code>tmpfs</code> Mode, File Has Date in the Future Warnings	189
Ensuring Hardware Clock Has the Correct Time	189
Troubleshooting a Rack Leader Controller (RLC) With Misconfigured Switch Information (SGI ICE Clusters Only)	190

Switch Wiring Rules	192
Admin Node eth2 Link in the Bond is Down	193
Installing SGI Tempo Versions Older than 2.9.0	193
Booting Nodes With iPXE After an Upgrade	194
Miniroot Operations	194
About the SMC Miniroot	195
Entering Rescue Mode	196
Booting an SGI Flat Compute Node on an Installed Cluster	196
Phase 1 — Initiating the Boot	197
Phase 2 — Loading the Node’s Kernel	198
Phase 3 — Loading the Miniroot	200
Phase 4 — Starting the Operating System on the Node	200
Appendix A. YAST Navigation	201
Appendix B. Subnetwork Information	203
About the Cluster Subnetworks	203
SGI Rackable Subnetworks	204
SGI Rackable Component Naming Conventions	204
SGI Rackable Address Ranges and VLANs for Management Software and Application Software	205
SGI Rackable System Management Network Topology Overview	207
SGI ICE VLANs and Subnetworks	208
SGI ICE Component Naming Conventions	208
SGI ICE Address Ranges and VLANs for Management Software and Application Software	210
SGI ICE System Control Configuration	213
Appendix C. SGI ICE MCell Network IP Addresses	217

Appendix D. Partition Layout Information	223
About the Partition Layout on SGI Clusters	223
About the Current Release's Partition Layout	225
Partition Layout for a One-slot Cluster	225
Partition Layout for a Two-slot Cluster (Default)	226
Partition Layout for a Five-slot Cluster	227
About the Legacy Partition Layout	228
Legacy Partition Layout for a One-slot SGI ICE Cluster	229
Legacy Partition Layout for a Two-slot SGI ICE Cluster	229
Legacy Partition Layout for a Five-slot SGI ICE Cluster	230
Appendix E. Specifying Configuration Attributes	233
About Configuration Attributes	233
UDPCast Options	234
edns_udp_size	234
udpcast_max_bitrate	234
udpcast_max_wait	235
udpcast_mcast_rdv_addr	235
udpcast_min_receivers	236
udpcast_min_wait	236
udpcast_rexmit_hello_interval	237
udpcast_ttl	238
VLAN and General Network Options	239
head_vlan	239
mcell_network	240
mcell_vlan	240
mgmt_vlan_end	240
mgmt_vlan_start	241

rack_vlan_end	241
rack_vlan_start	242
redundant_mgmt_network	242
switch_mgmt_network	242
Console Server Options	243
conserver_logging	243
conserver_ondemand	244
Miscellaneous Options	244
blademond_scan_interval	245
cluster_domain	245
dhcp_bootfile	246
discover_skip_switchconfig	247
max_rack_irus	247
mic	248
my_sql_replication	248
Index	249

About This Guide

This guide is a reference document for people who install and configure SGI® ICE™ cluster computer systems and SGI Rackable® cluster computer systems. SGI Rackable clusters are sometimes called *flat cluster* computer systems because of their nonhierarchical structure. This manual describes how to perform general system discovery, installation, configuration, and operations.

Note: The SMC 3.1.2 release supports the following operating systems: RHEL 7.1, RHEL 6.6, SLES 11 SP3, and CentOS 6.6. This documentation includes information about SMC support for SLES 12, but SMC support for SLES 12 is deferred.

Related Publications

The SGI Foundation Software release notes and the SGI Performance Suite release notes contain information about the specific software packages provided in those products. The release notes also list SGI publications that provide information about the products. The release notes are available in the following locations:

- Online at Supportfolio. After you log into Supportfolio, you can access the release notes. The SGI Foundation Software release notes are posted to the following website:

https://support.sgi.com/content_request/194480/index.html

The SGI Performance Suite release notes are posted to the following website:

https://support.sgi.com/content_request/786853/index.html

Note: You must sign into Supportfolio, at <https://support.sgi.com/login>, in order for the preceding links to work.

- On the product media. The release notes reside in a text file in the `/docs` directory on the product media. For example, `/docs/SGI-MPI-1.x-readme.txt`.
- On the system. After installation, the release notes and other product documentation reside in the `/usr/share/doc/packages/product` directory.

All SGI publications are available on the Technical Publications Library at the following website:

<http://docs.sgi.com>

The following software documentation might be useful to you:

- *SGI Management Center (SMC) Administration Guide for Clusters*, publication 007-6358-xxx

This manual explains how to use manage an SGI ICE cluster or an SGI Rackable cluster.

- *SGI Foundation Software (SFS) User Guide*, which explains how to use SFS to optimize system performance.
- *SGI MPI and SGI SHMEM User Guide*, which describes the industry-standard message passing protocol optimized for SGI computers. This manual describes how to tune the run-time environment to improve the performance of an MPI message passing application on SGI computers. None of these ways involve application code changes.
- *MPInside Reference Guide*, which documents the SGI MPInside profiling tool.

SGI's hardware documentation might also interest you. SGI creates hardware manuals that are specific to each product line. The hardware documentation typically includes a system architecture overview and describes the major components. It also provides the standard procedures for powering on and powering off the system, basic troubleshooting information, and important safety and regulatory specifications.

The following procedure explains how to retrieve a list of hardware manuals for your system.

Procedure 0-1 To retrieve hardware documentation

1. Type the following URL into the address bar of your browser:

`docs.sgi.com`

2. In the search box on the Techpubs Library, narrow your search as follows:

- In the **search** field, type the model of your SGI system.

For example, type one of the following: "UV 2000", "ICE X", Rackable.

Remember to enclose hardware model names in quotation marks (" ") if the hardware model name includes a space character.

- Check **Search only titles**.
- Check **Show only 1 hit/book**.
- Click **search**.

In addition to SGI documentation, the following documentation from other sources might interest you:

- SUSE documentation for SLES 12 and for SLES 11 SP3
- Red Hat documentation for Red Hat Enterprise Linux (RHEL) 7.1, RHEL 6.6, and CentOS 6.6
- Intel compiler documentation
- Intel documentation about Xeon architecture

Obtaining Publications

You can obtain SGI documentation in the following ways:

- See the SGI Technical Publications Library at the following website:

<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 view man pages by typing `man title` on a command line.

Conventions

The following conventions are used throughout this document:

Convention	Meaning
<code>command</code>	This fixed-space font denotes literal items such as commands, files, routines, path names, signals, messages, and programming language structures.

<i>variable</i>	Italic typeface denotes variable entries and words or concepts being defined.
user input	This bold, fixed-space font denotes literal items that the user enters in interactive sessions. (Output is shown in nonbold, fixed-space font.)
[]	Brackets enclose optional portions of a command or directive line.
...	Ellipses indicate that a preceding element can be repeated.

Reader Comments

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

You can contact SGI in either 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:
<http://www.sgi.com/support/supportcenters.html>

SGI values your comments and will respond to them promptly.

System Software Overview

This chapter includes the following topics:

- "About SGI Cluster Computer Systems" on page 1
- "SGI Cluster System Node Images" on page 5
- "Operating System Support" on page 8
- "SGI Cluster Networks" on page 8

About SGI Cluster Computer Systems

Figure 1-1 on page 2 shows an SGI Rackable cluster and an SGI ICE cluster. The SGI ICE family includes the following computers systems:

- SGI[®] ICE™ XA
- SGI[®] ICE™ X
- SGI[®] Altix[®] ICE 8400
- SGI[®] Altix[®] ICE 8200. SGI does not support RHEL 7.x or SLES 12 on SGI Altix ICE 8200 platforms.

Each type of cluster includes an admin node and flat compute nodes. As Figure 1-1 on page 2 shows, the admin node and the flat compute nodes attach directly to the management network. Admin nodes are sometimes referred to as *system admin controller* (SAC) nodes.

In an SGI ICE cluster, the SGI ICE compute nodes are configured in a hierarchical way, under a rack leader controller (RLC). In the SGI ICE cluster, it is the RLC that attaches to the management network, not the compute nodes.

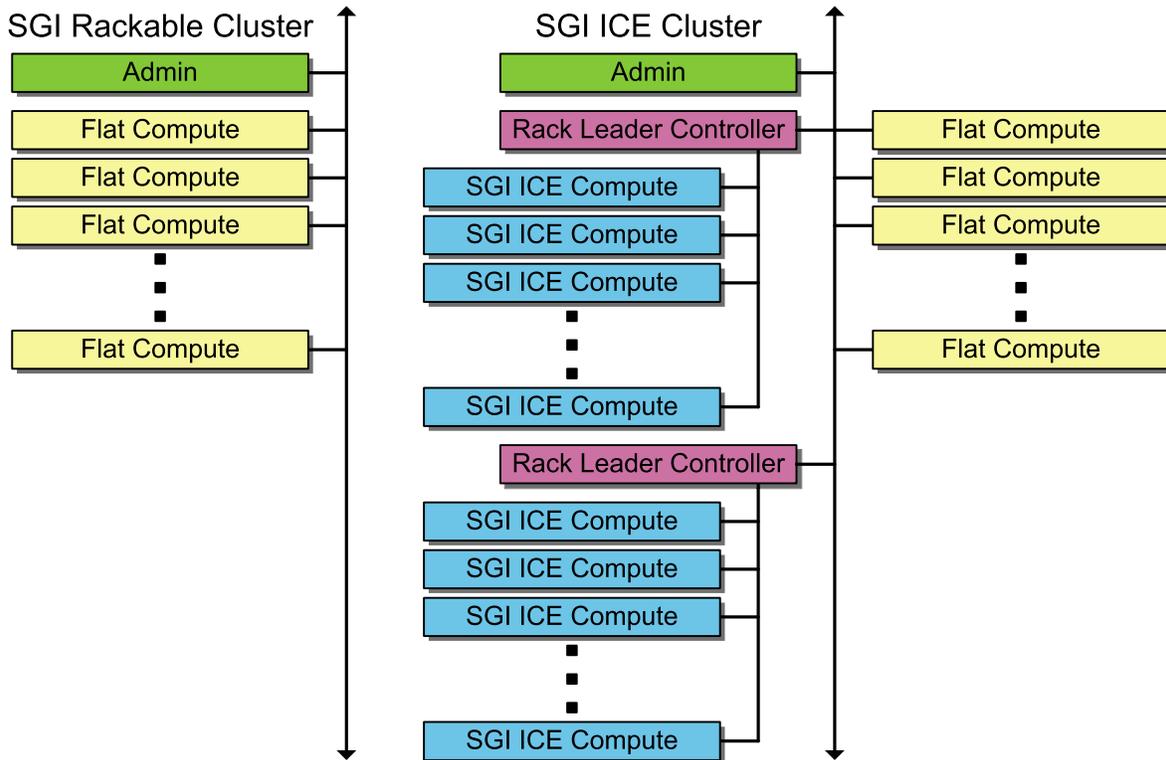


Figure 1-1 SGI Rackable Clusters and SGI ICE Clusters

The following topics describe the SGI cluster systems:

- "About SGI Rackable Clusters" on page 2
- "About SGI ICE Clusters" on page 4
- "About High Availability Nodes in SGI Clusters" on page 4
- "About Managed Nodes and Unmanaged Nodes" on page 5

About SGI Rackable Clusters

The nodes in an SGI Rackable cluster have the following roles:

- The admin node is the cluster's administrative node. This is the node from which you install software and manage the cluster. SGI Management Center (SMC) software resides on the admin node. SMC enables you to install, provision, configure, and manage the SGI Rackable cluster computing system.

Each cluster has an admin node. The admin node hosts the original, factory-installed copies of the software images for each component. System administrators log into the admin node to run system management commands, to modify component images, and to perform system-wide operations.

The SMC software distribution includes the master system image for the admin node. During the installation and configuration process, the installation software creates the master system images for the other components in the cluster. As you customize the system for your site, you modify the component-specific system images on the admin node and push the updated images to the other nodes.

- The compute nodes in an SGI Rackable cluster are called *flat compute nodes* because they are not configured in a hierarchical structure. An admin node can manage thousands of compute nodes, depending on the cluster's workload. The compute nodes all receive a hostname and an IP address during the configuration process. You can configure flat compute nodes with one or more of the following types of user services:
 - Login services. These services allow an end user to log in and then, for example, run or monitor MPI jobs.
 - Batch scheduling services. You can install workload schedulers such as Altair's PBS Professional, Adaptive Computing's Moab, SLURM, or TORQUE.
 - I/O gateway. On a small system, you can combine the I/O gateway, login services, and batch scheduling on the same compute node.

The I/O gateway services connect the cluster to your site network. You can configure one or more of the following protocols on the node: network file system (NFS), network address translation (NAT), or network information service (NIS).

- Storage. A compute node with storage is a network attached (NAS) appliance bundle that provides InfiniBand attached storage for the system.
- Object storage server. This service is used in Lustre File Storage configurations.
- Metadata server. This service is used in Lustre File Storage configurations.

SGI recommends the following login practices:

- Only the system administrator should be able to log into the admin node. SGI recommends that sites prohibit end-user access to the admin node.
- User services can be configured on the compute nodes, and end users can have access to these compute nodes.

About SGI ICE Clusters

On an SGI ICE cluster, the admin node and the flat compute nodes have the same characteristics that they have in an SGI Rackable cluster. That is, an admin node can support flat compute nodes configured for services such as logging in, batch computing, I/O, gateway, or storage. SGI recommends that in an SGI ICE cluster, you install user services on the flat compute nodes. For more information about the flat compute nodes, see the following:

"About SGI Rackable Clusters" on page 2

The hierarchical structure of the RLCs and the SGI ICE compute nodes is unique to the SGI ICE cluster. The hierarchical design enables these computing systems to be provisioned quickly. Master software images for each type of node in the cluster reside on the admin node. When an SGI ICE system is configured, the admin node pushes the software images to the flat compute nodes and to the RLCs. Each RLC pushes the compute node images to the SGI ICE compute nodes that reside in its rack.

An SGI ICE admin node can support many RLCs, each of which can manage hundreds of SGI ICE compute nodes. The characteristics of the RLCs and SGI ICE compute nodes are as follows:

- The RLC's role is to manage a set of SGI ICE compute nodes in a rack. SGI ICE clusters have at least one RLC.
- The SGI ICE compute nodes are simplified, typically diskless compute nodes that reside in a rack with an RLC. These compute nodes require an RLC for services, infrastructure, and support. An RLC can manage up to 288 SGI ICE compute nodes, each running its own Linux operating system, in a rack. The exact number of SGI ICE compute nodes that an RLC can manage depends on the specific model number of your SGI ICE cluster.

About High Availability Nodes in SGI Clusters

If you have an SGI ICE cluster, you can configure the admin node for high availability (HA) operations. An HA admin node consists of two physical nodes

dedicated to the admin node role and one virtual machine (VM). The VM hosts the functioning admin node, and the VM resides on one host at a time. The host upon which the VM resides is the *active node*. The other node is the *passive node*.

If you have an SGI ICE cluster, you can also configure the RLC for HA operations. In this case, two physical nodes and one virtual machine are dedicated to the RLC role, and the RLC VM resides on one of the RLC nodes at a time.

About Managed Nodes and Unmanaged Nodes

On an SGI cluster system, the admin node manages the vast majority of flat compute nodes, RLCs, and SGI ICE compute nodes.

Some clusters also include nodes that are unmanaged. These unmanaged nodes receive an IP address and a hostname from the admin node when you include them on the `discover` command, but all other operations occur outside the SMC environment.

The SGI SMC documentation addresses only the nodes that SMC actively manages. The documentation does not pertain to the unmanaged nodes that are attached to the cluster but are managed separately.

SGI Cluster System Node Images

Figure 1-2 on page 6 shows two example cluster systems. The upper figure is an SGI ICE cluster that also includes some flat compute nodes, and the lower figure is an SGI Rackable cluster. These depictions show example clusters with different operating systems running on different node types.

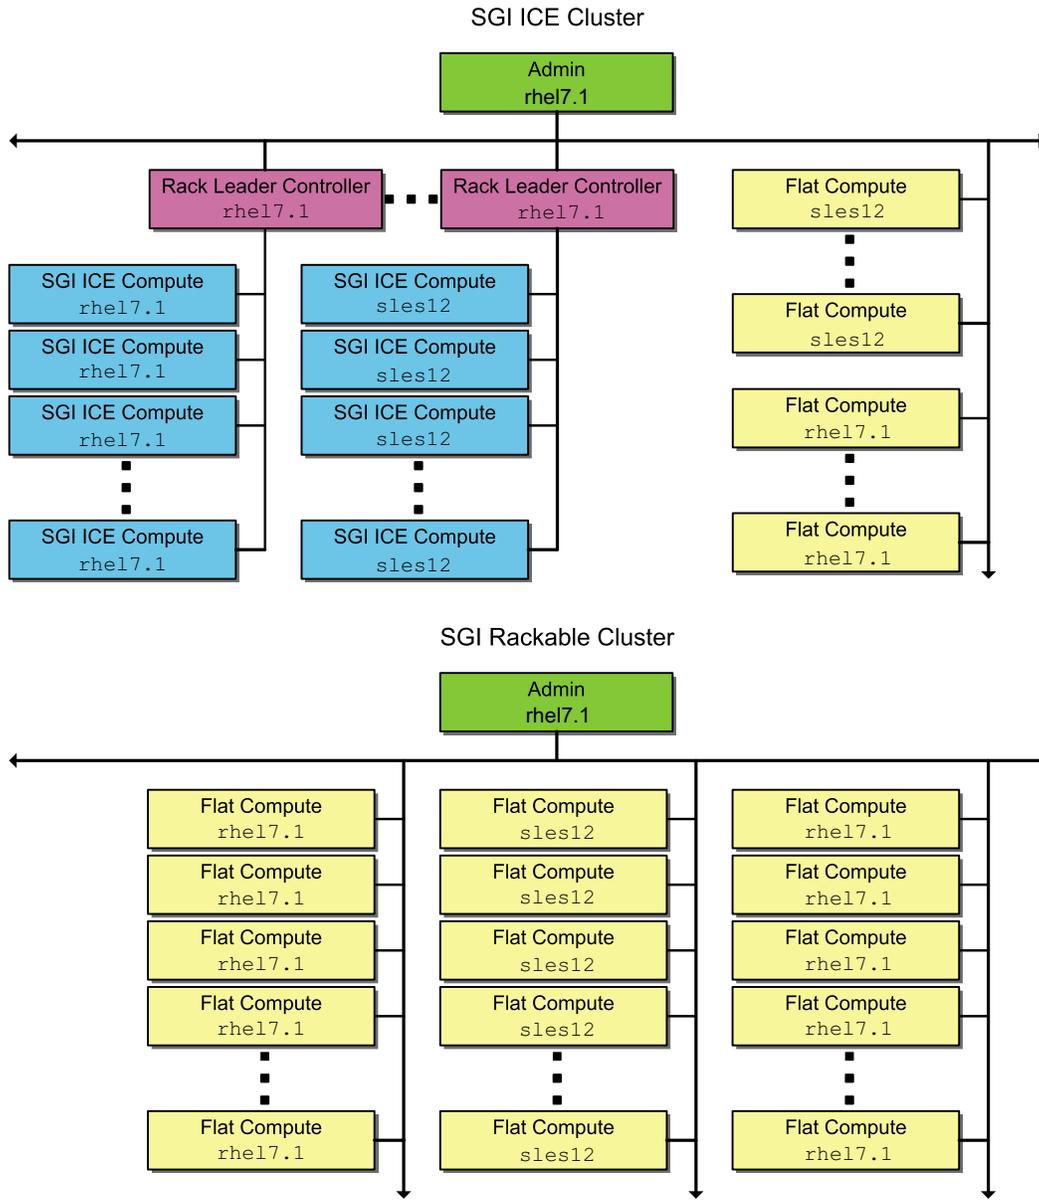


Figure 1-2 Cluster Nodes and Their Operating Systems

Figure 1-2 on page 6 shows the image names for the software images on each node. As the example shows, you can configure nodes with different operating system images within the same cluster. On an SGI ICE cluster, the admin node and RLC nodes must be installed with the same operating system. The SGI ICE compute nodes can host a different operating system. The image names shown in the figure are the default names for the factory-installed system images. The system images for the cluster nodes are unique to each type of node.

Table 1-1 on page 7 lists the nodes in each type of SGI cluster and shows the images that reside on each node. The table shows the default, factory-given name for each image.

Table 1-1 Cluster Nodes and Software Image Names

SGI Rackable Nodes	SGI ICE Nodes	Node System Image Name
Admin node	Admin node	—Not applicable— Installation of the admin node is facilitated by the admin install DVD (sgi-mgmtsuite-install).
Flat compute nodes	Flat compute nodes	<i>os_name</i> For example, rhel7.1
—Not applicable—	Rack leader controller (RLC)	lead- <i>os_name</i> For example, lead-sles12
—Not applicable—	SGI ICE compute nodes	ice- <i>os_name</i> For example, ice-sles11sp3

If you modify the image to include site-specific software, it is typical to copy the image, update the copy, and give the modified image a new name. Copying is also referred to as *cloning* in the cluster documentation. When you use a new image name for your changes, you preserve the ability to refer back to the original image.

Whenever you add or modify the software on a node, you can use the SGI Management Center software's version control software to manage multiple versions of each node's software. The version control system facilitates the following:

- Storage. You can have many versions of each individual software image, and each version is easily retrieved.

- Experimentation. Each software image is tagged with a version number, so you can easily enable and disable specific versions of the software images.

Operating System Support

The SGI cluster computer systems support the following operating systems:

- Red Hat Enterprise Linux (RHEL) 6.6 and RHEL 7.1
- SLES 11 SP3
- CentOS 6.6

Note: In SGI documentation, you can assume that feature descriptions for RHEL platforms also pertain to CentOS platforms unless otherwise noted.

SGI Cluster Networks

The following topics explain the SGI cluster networks:

- "SGI Rackable Networks" on page 8
- "SGI ICE Networks" on page 10

SGI Rackable Networks

Figure 1-3 on page 9 shows the logical networks of an SGI Rackable cluster.

SGI Rackable Cluster Management Network (Logical)

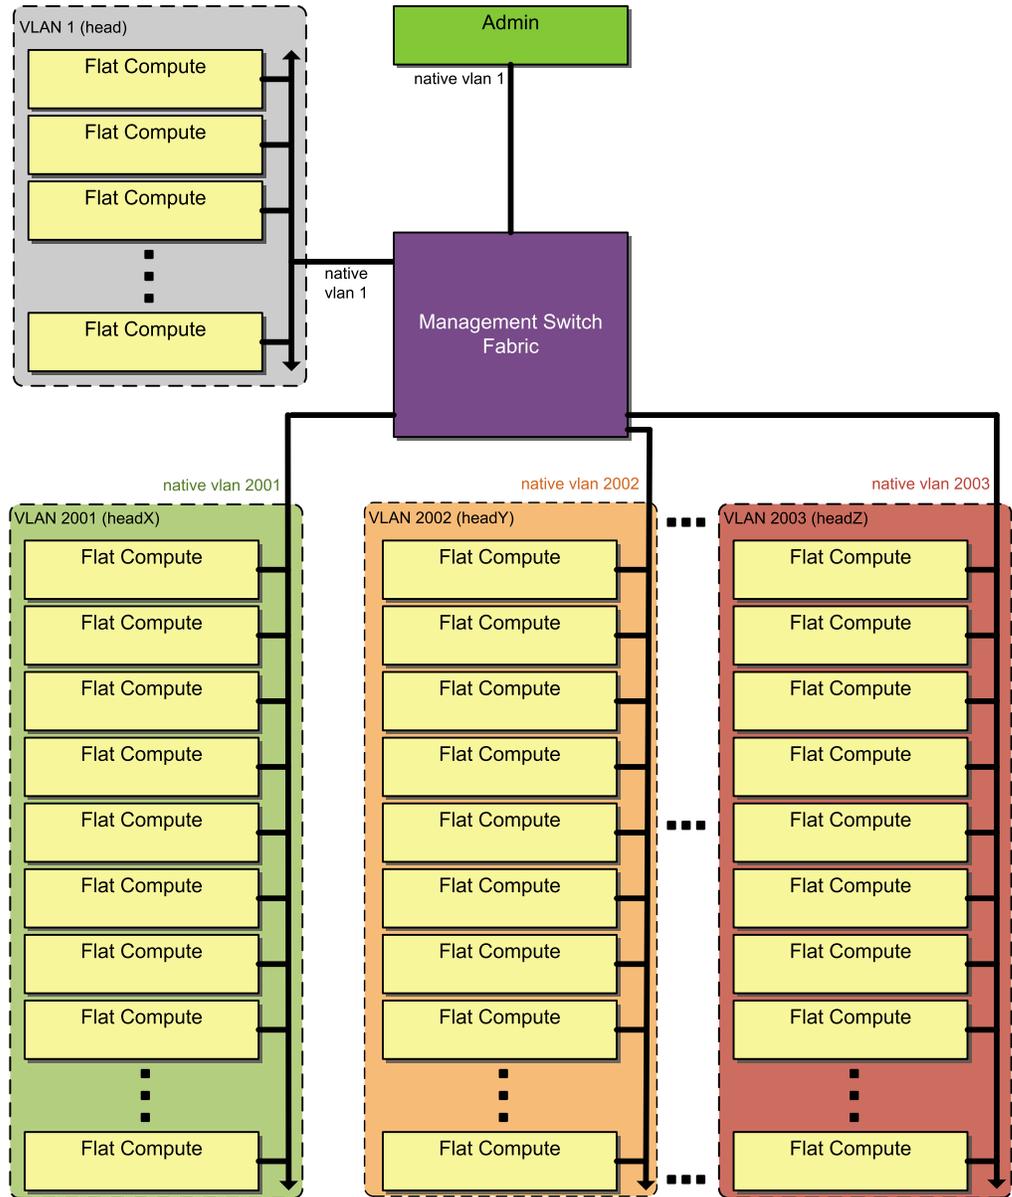


Figure 1-3 SGI Rackable Cluster Logical Representation

SGI ICE Networks

The system components in an SGI ICE cluster are attached to one or more of the following two internal networks:

- The data network.
- The management network.

The data network is designed for high-performance computing and bandwidth-intensive applications. This network is based on InfiniBand (IB) technology, and it facilitates communication to all SGI ICE compute nodes from the flat compute nodes. It connects the following:

- The SGI ICE compute nodes to each other. The InfiniBand network connects all the SGI ICE compute nodes (blades) to each other. The SGI ICE compute node InfiniBand connections are not part of the head network.
- The flat compute nodes to the SGI ICE compute nodes. One or two separate InfiniBand networks (or fabrics) segregate traffic within the SGI ICE system in a way that optimizes computing performance. When there are two InfiniBand networks, communication is segregated by InfiniBand interface, as follows:
 - `ib0`, which is used typically for Message Passing Interface (MPI) communication.
 - `ib1`, which is typically used for storage traffic.

The management network, also known as the head network, is designed for monitoring, provisioning, and other functions not covered by the data network. This Ethernet network is designed for communication between the admin node, the RLCs, and the flat compute nodes. These components communicate to each other directly within the head network. The head network connects the following nodes directly into the Ethernet switches:

- Admin Node
- RLCs
- Flat compute nodes
- Additional Ethernet Switches

The head network also includes several additional virtual local area networks (VLANs). Figure 1-4 on page 12 is a logical representation of the SGI ICE Ethernet

networks that shows how SGI ICE components are logically separated by these VLAN boundaries.

SGI ICE Cluster Management Network (Logical)

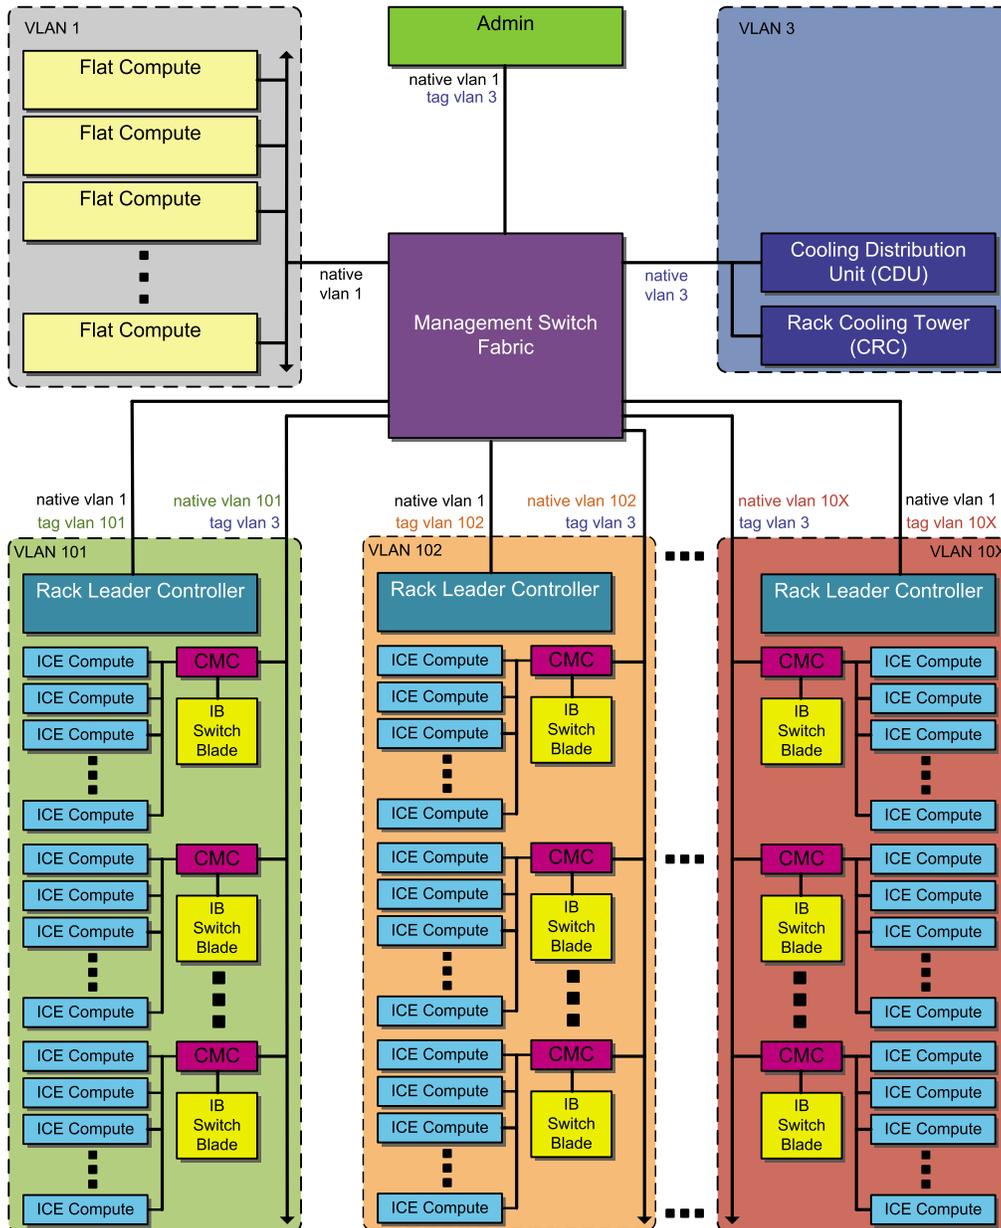


Figure 1-4 SGI ICE Networks

On an SGI ICE cluster, several virtual local area networks (VLANs) are configured. The following list shows the typical VLAN names and numbers:

- Head network. VLAN tag 1.
- Cooling network. VLAN tag 3. Clusters with MCells only.
- Rack 1 network. VLAN tag 101.
- Rack 2 network. VLAN tag 102.
- Rack *x* network. VLAN tag 10*x*.

Note: The head VLAN network must always be VLAN 1. Do not attempt to change VLAN 1's number. You can change the other VLAN numbers and names.

In Figure 1-4 on page 12, the head network is VLAN 1. The ports connected to the admin node and the flat compute nodes are in VLAN 1 natively.

The Ethernet Switches are configured with a VLAN for each RLC. This VLAN segregates management traffic so that communication between the SGI ICE compute nodes and their corresponding RLC is contained within that VLAN. Physically, the SGI ICE compute nodes are connected to a chassis management controller (CMC) and do not directly connect to the Ethernet switch. Instead, the CMCs connect directly to the Ethernet Switch. Only the RLC can communicate with the SGI ICE compute nodes and CMCs in its own logical rack.

Users can log into the admin node and into the flat compute nodes directly. If access to the RLC is required, users can log in directly to the admin node and then use the `ssh(1)` command to log into an RLC. The typical VLAN mapping for the Ethernet switches on each node is as follows:

Node type	VLANs
Admin	Native VLAN 1 Tagged VLAN 3
Flat compute	Native VLAN 1
RLC	Native VLAN 1 Tagged VLAN 10 <i>x</i> . VLAN created for each RLC.

CMCs	Native VLAN 10x. VLAN created for their corresponding RLC. Tagged VLAN 3.
Cooling equipment cooling distribution units (CDUs) and cooling rack controllers (CRCs)	Native VLAN 3. Clusters with MCells only.

Figure 1-5 on page 15 shows the physical networks of an SGI ICE cluster.

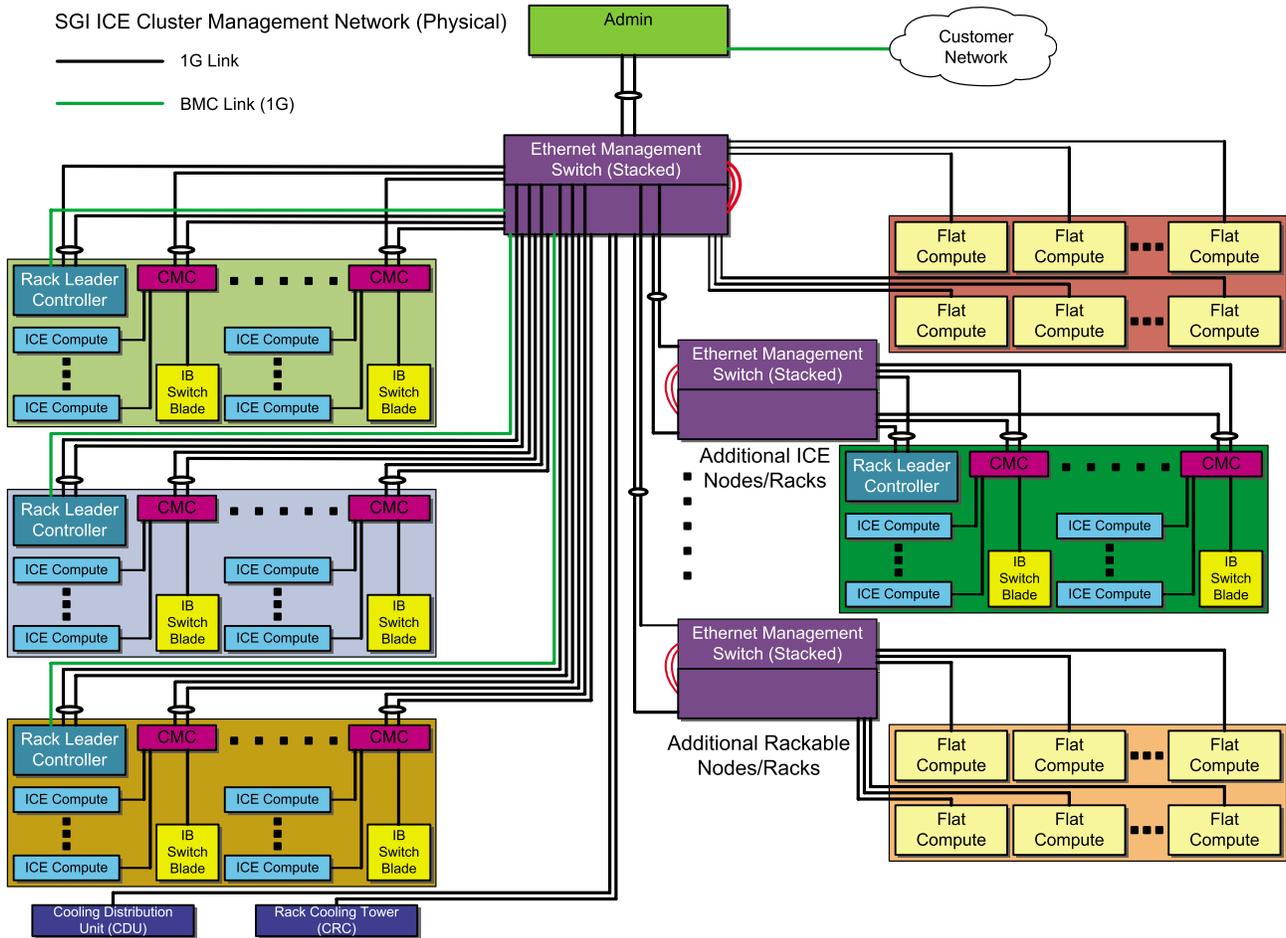


Figure 1-5 SGI ICE Cluster Physical Representation

Customizing a Factory-installed SGI Cluster (RHEL 6 and SLES 11)

This chapter contains the following topics:

- "About Customizing a Factory-installed Cluster" on page 17
- "Obtaining Information" on page 18
- "Changing the Password and Specifying Network Information" on page 19
- "Completing the Customization" on page 23
- "(Conditional) Pushing Changes to SGI ICE Compute Nodes" on page 26
- "Verifying the System and Configuring Additional Features" on page 28

About Customizing a Factory-installed Cluster

The process in this chapter explains how to configure an SGI cluster into your data center. This process assumes that you want to keep the factory configuration, and you want to customize only the networking and other site-specific characteristics. The process in this chapter assumes that the SGI factory configured the cluster with one of the following operating systems:

- Red Hat Enterprise Linux (RHEL) 6.x
- SLES 11 SPx

At the factory, SGI configured and tested the following:

- A factory-specified root password. One of the first steps in the configuration procedure is to change this root password on the admin node.
- Two slots. SGI configured the operating system that you ordered on slot 1. The operating system can be either Red Hat Enterprise Linux (RHEL), CentOS, or SLES. Slot 2 is blank.

The cluster supports a maximum of five slots. If you need more than the factory-configured two slots, you need to reconfigure the system. During the reconfiguration, you reinstall the operating system and perform many other tasks.

For the reconfiguration procedure, see Chapter 4, "Installing and Configuring an SGI Cluster System" on page 45.

- A serial-over-LAN connection on the admin node. This connection allows you to use the admin node as the system console. You can access the console by using the IPMItool's serial-over-LAN function.

The following procedures explain how to attach the cluster to your network and retain the factory-installed configuration:

- "Obtaining Information" on page 18
- "Changing the Password and Specifying Network Information" on page 19
- "Completing the Customization" on page 23
- "(Conditional) Pushing Changes to SGI ICE Compute Nodes" on page 26
- "Verifying the System and Configuring Additional Features" on page 28

Obtaining Information

Your configuration session can proceed more quickly if you gather some information before you start. When you perform the configuration, you update the factory-installed, system-wide root password, and you update the time zone. In addition, you provide information about your site network for the admin node's `eth0` network interface card (NIC).

The following procedure explains the information that you need to gather.

Procedure 2-1 To obtain information for the customization

1. Complete the following table:

Information Needed	Specifics for this Cluster
Factory-installed password	_____
Password for this system at your site	_____
Time zone	_____
IP address	_____
Netmask	_____

Hostname	
Default route/Gateway	
Fully qualified domain name (FQDN)	
House NTP server	
First house (site) DNS resolver IP address	
(Optional) Second house DNS resolver IP address	
(Optional) Third house DNS resolver IP address	
House (site) domain	
Cluster subdomain name	

2. Proceed to the following:

"Changing the Password and Specifying Network Information" on page 19

Changing the Password and Specifying Network Information

The following procedure explains how to change the password on the admin node and how to update the operating system configuration files with your site's networking information.

Procedure 2-2 To change the password and attach a factory-installed cluster to your site network

1. Use the console attached to the admin node, and log into the admin node as the root user.
2. Type the following command, and follow the prompts, to change the root password on the admin node and on all other nodes:

```
# cpasswd
```

If you need help, type the following command:

```
# cpasswd --h
```

If you do not have the current password, you can obtain the factory-installed password from your SGI representative.

For example:

```
admin node:~ # cpasswd
Enter new password:
Enter new password (again):
admin: updating /etc/shadow
rilead: updating /etc/shadow
service0: updating /etc/shadow
admin node:~ #
```

3. Change the system time zone.

This step is different, depending on the cluster's operating system, as follows:

- On RHEL platforms, type the following command:

```
# system-config-date
```

The `system-config-date` command starts a graphical user interface (GUI) tool. Within the GUI tool, change **only** the system time zone. The tool enables you to change other aspects of the configuration, but for this step, change only the system time zone.

- On SLES platforms, complete the following steps:

- Type the following command to start YAST:

```
# yast
```

- Select **System > Date and Time**.
- On the **Clock and Time Zone** page, select the correct setting.
- Click **Change**
- Click **Back**
- Exit YAST.

Note: Do not use this tool to change the NTP server, the time, or other configuration data.

4. Use a text editor, such as `vi` or `vim`, to open file `ifcfg-eth0`:

- On RHEL platforms, the path is as follows:

```
/etc/sysconfig/network-scripts/ifcfg-eth0
```

- On SLES platforms, the path is as follows:

```
/etc/sysconfig/network/ifcfg-eth0
```

5. Edit file `ifcfg-eth0` as follows:

- On RHEL platforms, add the `IPADDR` and `NETMASK` lines, and then add values appropriate for your site network. Also add a line that includes `ONBOOT=yes`.

For example:

```
IPADDR=128.162.244.88
NETMASK=255.255.255.0
ONBOOT=yes
```

- On SLES platforms, add the `IPADDR` and `NETMASK` lines, and then add values appropriate for your site network. Also add a line that includes `STARTMODE='onboot'`.

For example:

```
IPADDR='128.162.244.88'
NETMASK='255.255.255.0'
STARTMODE='onboot'
```

6. Save and close file `ifcfg-eth0`.

7. Create the networking configuration file.

This file contains data that is boot-persistent. A later step in the configuration procedure sets the hostname in a way that can persist until the next boot. The commands to create the networking file differ depending on your platform, as follows:

- On RHEL platforms, complete the following steps:

- Use a text editor to create the following file:

```
/etc/sysconfig/network
```

- Add the following three lines to file `/etc/sysconfig/network`:

```
NETWORKING=yes
HOSTNAME=admin_node_hostname
GATEWAY=gateway_IP_address
```

For *admin_node_hostname*, type the hostname you want to assign to the admin node.

For *gateway_IP_address*, type the IP address of the gateway for your house network.

For example:

```
NETWORKING=yes
HOSTNAME=my-system-admin
GATEWAY=128.162.244.1
```

– Save and close file `/etc/sysconfig/network`.

- On SLES platforms, complete the following steps:

– Create the following file:

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

– Add the following line to file `/etc/sysconfig/network/routes`:

```
default gateway - -
```

For *gateway*, type the IP address for the site gateway server.

For example:

```
default 100.100.100.101 - -
```

Note that SMC manages and rewrites everything below the `default gateway` line.

8. Use a text editor to open file `/etc/hosts`.
9. Add a line in the following format to file `/etc/hosts`:

```
admin_node_IP admin_node_FQDN admin_node_hostname
```

The variables in the preceding line are as follows:

- For *admin_node_IP*, type the IP address of the admin node.
- For *admin_node_FQDN*, type the fully qualified domain name (FQDN) of the admin node.
- For *admin_node_hostname*, type the hostname of the admin node.

For example, add the following line:

```
128.162.244.88 acme-admin.acme.usa.com acme-admin
```

10. Save and close file `/etc/hosts`.
11. Type the following command to set the admin node's hostname:

```
# hostname admin_node_hostname
```

For example:

```
# hostname acme-admin
```

For *admin_node_hostname*, type the hostname of the admin node. Make sure to type the hostname, which is the short name. Do not type the admin node's FQDN, which is the longer name.

12. Proceed to the following:
"Completing the Customization" on page 23

Completing the Customization

The following topic explains how to use the cluster configuration tool to add information about your site's network to the cluster database.

Procedure 2-3 To customize the cluster database

1. Type the following command to start the cluster configuration tool:

```
# /opt/sgi/sbin/configure-cluster
```

For general information about the cluster configuration tool, see "Configuring the Cluster" on page 79.

2. On the cluster configuration tool's main menu select **Configure the Time Client/Server (NTP)** and select **OK**.

The system guides you through the process to specify your house NTP server in file `/etc/ntp.conf`. This process differs, depending on your platform, as follows:

- On RHEL platforms, follow the the instructions that the cluster configuration tool presents to you.

- On SLES platforms, the cluster configuration tool opens a YAST menu. Follow the prompts in the YAST menu to set your NTP server.

3. On the cluster configuration tool's main menu select **Configure House DNS Resolvers** and select **OK**.

You can specify up to three house DNS resolvers.

4. Select **Quit** and select **OK** to log out from the cluster configuration tool.
5. Use the `cadmin` command, in the following format, to set the house (site) domain:

```
# cadmin --set-admin-domain site_domain
```

For *site_domain*, specify the full name of your house domain. For example, `usa.acme.com`.

6. Use the `cadmin` command, in the following format, to change the subdomain name for the cluster:

```
# cadmin --set-subdomain cluster_name
```

For *cluster_name*, specify the name of the system. For example, `ice.usa.acme.com`.

For more information about the `cadmin` command, type `cadmin -h` at the system prompt.

7. Type the following command to retrieve the name of the admin node:

```
# cadmin --show-hostname --node admin
```

8. Use the `cadmin` command, in the following format, to customize the admin node name:

```
# cadmin --set-hostname --node admin admin_node_name
```

For *admin_node_name*, specify the hostname you want to use for the cluster's admin node.

9. Type the following command to display the node images on the cluster's admin node at this time:

```
# cinstallman --show-images
```

```
Image Name                               BT VCS Compat_Distro
rhel6.6                                   0 1  rhel6
      3.10.0-229.el7.x86_64
```

```
ice-rhel6.6                0 1  rhel6
    3.10.0-229.el7.x86_64
lead-rhel6.6              0 1  rhel6
    3.10.0-229.el7.x86_64
```

10. Use the `cp(1)` command, in the following format, to set the time zone in the system images.

The format of this command is as follows:

```
cp /etc/localtime /var/lib/systemimager/images/image_name/etc
```

For *image_name*, type the name of one of the system images you retrieved with the `cinstallman` command. Type one `cp(1)` command for each master node system image. On an SGI ICE cluster that includes both SGI ICE compute nodes and flat compute nodes, you need three `cp(1)` commands, one for each image type. On an SGI Rackable cluster, you need one `cp(1)` command.

Example 1. On an SGI ICE cluster with RHEL, type the following commands, one for each master image:

```
# cp /etc/localtime /var/lib/systemimager/images/ice-rhel6.6/etc
# cp /etc/localtime /var/lib/systemimager/images/rhel6.6/etc
# cp /etc/localtime /var/lib/systemimager/images/lead-rhel6.6/etc
```

Example 2. On an SGI Rackable SLES cluster, type the following command to copy the compute node's master image:

```
# cp /etc/localtime /var/lib/systemimager/images/sles11/etc
```

11. Update the compute node images with your site's time zone information.

On all clusters, update the flat compute node images. On SGI ICE clusters, update the SGI ICE compute node images.

On an SGI ICE cluster, type the following command:

```
# pdcp -g leader /etc/localtime /etc/localtime
```

On an SGI cluster that includes flat compute nodes, type the following command:

```
# pdcp -g compute /etc/localtime /etc/localtime
```

12. Proceed to one of the following:

- If you have an SGI ICE cluster, proceed to the following:

"(Conditional) Pushing Changes to SGI ICE Compute Nodes" on page 26

- If you have an SGI Rackable cluster, proceed to the following:
"Verifying the System and Configuring Additional Features" on page 28

(Conditional) Pushing Changes to SGI ICE Compute Nodes

Complete the procedure in this topic if your cluster is an SGI ICE cluster.

Cluster admin nodes can host multiple forms of system images for each of the rack leader controller (RLC) nodes, SGI ICE compute nodes, and flat compute nodes. For example, you can have some production images and some test images, and you can push the images to the nodes as needed. "Completing the Customization" on page 23 explains how to update the master images for the RLC nodes and the compute nodes. In addition, that procedure explains how to push the updated the node images to the cluster's nodes.

In this topic's procedure, you push the updated SGI ICE compute node images to the SGI ICE compute nodes. The push action enables the SGI ICE compute nodes to run with the updated networking information that you configured in "Completing the Customization" on page 23.

Procedures that describe system operations, later on in this manual, refer back to this procedure because you need to update system images and push out new images as part of several system administration tasks.

The following procedure explains how to push SGI ICE compute node system images to compute nodes.

Procedure 2-4 To push software system images

1. (Conditional) Stop the SGI ICE compute nodes

Complete this step if you want to boot from NFS roots. Do not complete this step if you want to boot from `tmpfs` roots.

Type the following command:

```
# cpower node halt "r*i*n*"
```

The preceding command stops all the SGI ICE compute nodes. Use the preceding command when you need to push an updated image out to all the nodes.

2. (Optional) Provide information about the number of racks on your system.

Perform this step if you have a small system with fewer than eight IRUs per RLC.

The procedure pushes the updated SGI ICE compute image to all the SGI ICE compute nodes. This process can run for a long time on large systems. If you have a large number of IRUs, you need the system to perform expansions that enable you to change many SGI ICE compute nodes at a time. If you have fewer than eight IRUs per RLC, however, the expansions are not needed.

The following substeps explain how to prepare the system to work on a smaller number of SGI ICE compute nodes:

- Type the following command to retrieve the identifiers for the RLCs on your system:

```
# cnodes --leader
```

- Type the following command one or more times to suppress unnecessary processing:

```
cadmin --set-max-irus --node rlc_id number_of_racks
```

For *rlc_id*, specify the identifier for one of the RLCs in your system.

For *number_of_racks*, specify the number of IRUs associated with this RLC.

For example, the following command specifies that there is only one IRU associated with the RLC identified as *r1lead*:

```
# cadmin --set-max-irus --node r1lead 1
```

3. Use the *cimage* command, in the following format, to push the changes:

```
cimage --push-rack ice-image_name rack
```

For *image_name*, specify the name of the SGI ICE compute node image that you updated.

For *rack*, specify the nodes. To specify all SGI ICE compute nodes, specify *r** or *r*i*n**. To specify only selected nodes, specify *rxixnx*, and substitute specific integer numbers for the *x* characters.

For example, the following command pushes the time zone changes from "Completing the Customization" on page 23 to all the SGI ICE compute racks:

```
# cimage --push-rack ice-rhel6.6 r\*
```

4. Type the following command to power-up the SGI ICE compute nodes:

```
# cpower node on "*" 
```

5. Proceed to the following:

"Verifying the System and Configuring Additional Features" on page 28

Verifying the System and Configuring Additional Features

The following procedure explains where you can obtain information about how to configure additional features.

Procedure 2-5 To configure additional features

1. (Conditional) Verify that the blade daemons are running.

Perform the procedure in the following topic if you are attaching an SGI ICE cluster to your network.

"(Conditional) Verifying That the `smcrbd` daemon and the `icerackd` Daemon are Running" on page 143

Return to this step when finished.

2. Configure one or more of the following additional features, some of which are required:

- Mellanox OpenFabrics Enterprise Distribution (OFED) for Linux Software.

If you run SLES 11 SP3 on an SGI ICE cluster or on an SGI Rackable cluster with Mellanox Connect-IB adaptor cards, perform the procedure in the following topic:

"(Conditional) Enabling the Mellanox OpenFabrics Enterprise Distribution (OFED) for Linux Software" on page 167

The procedure explains how to update the OFED software, which is required for performance reasons.

- Networking features, such as network address translation (NAT). For information about how to configure optional features, see the following:
SGI Management Center (SMC) Administration Guide for Clusters.
- Array Services, the hardware event tracker (HET), or CPU frequency scaling. For information about these additional features, see one of the following:
 - Chapter 5, "Configuring Additional Features" on page 155
 - *SGI Foundation Software (SFS) User Guide*
- SGI Remote Services. This feature provides a secure connection to SGI Customer Support. For more information about SGI Remote Services, see the following website:

https://www.sgi.com/services/support/remote_7x24.html

Customizing a Factory-installed SGI Cluster (RHEL 7 and SLES 12)

This chapter contains the following topics:

- "About Customizing a Factory-installed Cluster " on page 31
- "Obtaining Information" on page 32
- "Changing the Password and Specifying Network Information" on page 33
- "Completing the Customization" on page 38
- "(Conditional) Pushing Changes to SGI ICE Compute Nodes" on page 41
- "Verifying the System and Configuring Additional Features" on page 43

About Customizing a Factory-installed Cluster

The process in this chapter explains how to configure an SGI cluster into your data center. This process assumes that you want to keep the factory configuration, and you want to customize only the networking and other site-specific characteristics. The process in this chapter assumes that the SGI factory configured the cluster with one of the following operating systems:

- Red Hat Enterprise Linux (RHEL) 7.1
- SLES 12

Your SGI cluster was configured and tested at the factory. At the factory, SGI configured the following:

- A factory-specified root password. One of the first steps in the configuration procedure is to change this root password on the admin node.
- Two slots. SGI configured the operating system that you ordered on slot 1. The operating system can be either Red Hat Enterprise Linux (RHEL), SLES, or CentOS. Slot 2 is blank.

The cluster supports a maximum of five slots. If you need more than the factory-configured two slots, you need to reconfigure the system. During the

reconfiguration, you reinstall the operating system and perform many other tasks. For the reconfiguration procedure, see Chapter 4, "Installing and Configuring an SGI Cluster System" on page 45.

- A serial-over-LAN connection on the admin node. This connection allows you to use the admin node as the system console. You can access the console by using the IPMItool's serial-over-LAN function.

The following procedures explain how to attach the cluster to your network and retain the factory-installed configuration:

- "Obtaining Information" on page 32
- "Changing the Password and Specifying Network Information" on page 33
- "Completing the Customization" on page 38
- "(Conditional) Pushing Changes to SGI ICE Compute Nodes" on page 41
- "Verifying the System and Configuring Additional Features" on page 43

Obtaining Information

Your configuration session can proceed more quickly if you gather some information before you start. When you perform the configuration, you update the factory-installed, system-wide root password, and you update the time zone. In addition, you provide information about your site network for the admin node's eth0 network interface card (NIC).

The following procedure explains the information that you need to gather.

Procedure 3-1 To obtain information for the customization

1. Complete the following table:

Information Needed	Specifics for this Cluster
Factory-installed password	_____
Password for this system at your site	_____
Time zone	_____
IP address	_____
Netmask	_____

Hostname	_____
Default route/Gateway	_____
Fully qualified domain name (FQDN)	_____
House NTP server	_____
First house (site) DNS resolver IP address	_____
(Optional) Second house DNS resolver IP address	_____
(Optional) Third house DNS resolver IP address	_____
House (site) domain	_____
Cluster subdomain name	_____

2. Proceed to the following:

"Changing the Password and Specifying Network Information" on page 33

Changing the Password and Specifying Network Information

The following procedure explains how to change the password on the admin node and how to update the operating system configuration files with your site's networking information.

Procedure 3-2 To change the password and attach a factory-installed cluster to your site network

1. Use the console attached to the admin node, and log into the admin node as the root user.
2. Type the following command to retrieve the MAC addresses of the network Ethernet cards on the admin node:

```
# ip addr
```

In the output, note the MAC address of eth0. You need this MAC address later in this procedure.

3. Use a text editor to open file `ifcfg-eth0`.

On RHEL systems, the path to this file is as follows:

`/etc/sysconfig/network-scripts/ifcfg-eth0`

4. Replace the contents of the `ifcfg-eth0` file with the following lines:

```
TYPE="Ethernet"
BOOTPROTO="none"
DEFROUTE="yes"
IPV4_FAILURE_FATAL="no"
IPV6INIT="no"
IPV6_AUTOCONF="yes"
IPV6_DEFROUTE="yes"
IPV6_PEERDNS="yes"
IPV6_PEERROUTES="yes"
IPV6_FAILURE_FATAL="no"
NAME="eth0"
DEVICE="eth0"
ONBOOT="yes"
IPADDR="128.162.233.35" # replace with your IP address for the admin node
PREFIX="24"
GATEWAY="128.162.233.1" # replace with your site gateway
DNS1="128.162.236.210" # replace with your site's primary DNS IP address
DNS2="198.149.32.11" # replace with your site's secondary DNS IP address
DOMAIN="americas.sgi.com engr.sgi.com sgi.com" # replace with your site domain
HWADDR="XX:XX:XX:XX:XX:XX" # replace with the MAC address of eth0
```

Note: SMC does not support IPV6 at this time. The following line is needed for this installation:

```
IPV6INIT="no"
```

You can remove the lines that start with `IPV6_` from the `ifcfg-eth0` file, or you can retain those lines for completeness. SMC ignores these lines during the installation.

5. Save and close file `ifcfg-eth0`.
6. Type the following command to create file `/etc/sysconfig/network` with no content:

```
# touch /etc/sysconfig/network
```

7. Use a text editor to open file `/etc/hosts`.

8. Add a line in the following format to file `/etc/hosts`:

```
admin_node_IP admin_node_FQDN admin_node_hostname
```

The variables in the preceding line are as follows:

- For *admin_node_IP*, type the IP address of the admin node.
- For *admin_node_FQDN*, type the fully qualified domain name (FQDN) of the admin node.
- For *admin_node_hostname*, type the hostname of the admin node.

For example, add the following line:

```
128.162.244.88 acme-admin.acme.usa.com acme-admin
```

9. Save and close file `/etc/hosts`.
10. Type the following command to set the admin node's hostname:

```
# hostnamectl set-hostname admin_node_hostname
```

For *admin_node_hostname*, type the hostname of the admin node. Make sure to type the hostname, which is the short name. Do not type the admin node's FQDN, which is the longer name.

11. Type the following commands to bring up `eth0` with the updated networking information:

```
# ifdown eth0
# ifup eth0
```

12. Type the following command to restart the name service cache daemon, `nscd`:

```
# systemctl restart nscd
```

13. Type the following command to restart the network daemon:

```
# systemctl restart network
```

14. Type the following command to retrieve the admin node's current time zone information:

```
# strings /etc/localtime | tail -1
CST6CDT,M3.2.0,M11.1.0
```

The previous output shows the admin node set to US Central time. If the output you see is not correct for this cluster, perform the following steps:

- Type the following command to change to the directory that contains the time zone configuration files:

```
# cd /usr/share/zoneinfo
```

- Select a file from that directory that describes the time zone for the admin node.
- Type the following commands to enable the new time zone configuration file.

For example:

```
# cp -l /usr/share/zoneinfo/time_zone_file /etc/localtime.$$  
# mv /etc/localtime.$$ /etc/localtime
```

For *time_zone_file*, type the name of the time zone file that you need from the /usr/share/zoneinfo directory.

For example, type the following commands to change the admin node's time zone to US Pacific time:

```
# cp -l /usr/share/zoneinfo/PST8PDT /etc/localtime.$$  
# mv /etc/localtime.$$ /etc/localtime  
mv: overwrite '/etc/localtime'? y
```

In the preceding example, the mv command prompts you to confirm your intent. Type **y**.

- Type the following command to confirm the time zone:

```
# strings /etc/localtime | tail -1  
PST8PDT,M3.2.0,M11.1.0
```

15. Type the following command, and follow the prompts, to change the root password on the admin node and on all other nodes:

```
# cpasswd
```

If you need help, type the following command:

```
# cpasswd --h
```

If you do not have the current password, you can obtain the factory-installed password from your SGI representative.

For example:

```
admin node:~ # cpasswd
Enter new password:
Enter new password (again):
admin: updating /etc/shadow
r1lead: updating /etc/shadow
service0: updating /etc/shadow
admin node:~ #
```

16. Change the system time zone.

This step is different, depending on the cluster's operating system, as follows:

- On RHEL platforms, complete the following steps:
 - Type the following command to display a list of time zones:

```
# timedatectl list-timezones
```
 - Type the following command to set the time zone:

```
# timedatectl set-timezone time_zone
```

For *time_zone*, specify one of the time zones from the `timedatectl list-timezones` command output.
- On SLES platforms, complete the following steps:
 - Type the following command to start YAST:

```
# yast
```
 - Select **System > Date and Time**.
 - On the **Clock and Time Zone** page, select the correct setting.
 - Click **Change**
 - Click **Back**
 - Exit YAST.

Note: Do not use this tool to change the NTP server, the time, or other configuration data.

17. Proceed to the following:

"Completing the Customization" on page 38

Completing the Customization

The following topic explains how to use the cluster configuration tool to add information about your site's network to the cluster database.

Procedure 3-3 To customize the cluster database

1. Type the following command to start the cluster configuration tool:

```
# /opt/sgi/sbin/configure-cluster
```

For general information about the cluster configuration tool, see "Configuring the Cluster" on page 79.

2. On the cluster configuration tool's main menu select **Configure the Time Client/Server (NTP)** and select **OK**.

The system guides you through the process to specify your house NTP server in file `/etc/ntp.conf`. This process differs, depending on your platform, as follows:

- On RHEL platforms, follow the the instructions that the cluster configuration tool presents to you.
 - On SLES platforms, the cluster configuration tool opens a YAST menu. Follow the prompts in the YAST menu to set your NTP server.
3. On the cluster configuration tool's main menu select **Configure House DNS Resolvers** and select **OK**.

You can specify up to three house DNS resolvers.

4. Select **Quit** and select **OK** to log out from the cluster configuration tool.

- Use the `cadmin` command, in the following format, to set the house (site) domain:

```
# cadmin --set-admin-domain site_domain
```

For *site_domain*, specify the full name of your house domain. For example, `usa.acme.com`.

- Use the `cadmin` command, in the following format, to change the subdomain name for the cluster:

```
# cadmin --set-subdomain cluster_name
```

For *cluster_name*, specify the name of the system. For example, `ice.usa.acme.com`.

For more information about the `cadmin` command, type `cadmin -h` at the system prompt.

- Type the following command to retrieve the name of the admin node:

```
# cadmin --show-hostname --node admin
```

- Use the `cadmin` command, in the following format, to customize the admin node name:

```
# cadmin --set-hostname --node admin admin_node_name
```

For *admin_node_name*, specify the hostname you want to use for the cluster's admin node.

- Type the following command to display the nodes images on the cluster's admin node at this time:

```
# cinstallman --show-images
Image Name                               BT VCS Compat_Distro
rhel7.1                                   0  1  rhel7
    3.10.0-229.el7.x86_64
ice-rhel7.1                               0  1  rhel7
    3.10.0-229.el7.x86_64
lead-rhel7.1                              0  1  rhel7
    3.10.0-229.el7.x86_64
```

- Use the `cp(1)` command to set the time zone in the system images.

The format of this command is as follows:

```
cp /etc/localtime /var/lib/systemimager/images/image_name/etc
```

For *image_name*, type the name of one of the system images you retrieved with the `cinstallman` command. Type one `cp(1)` command for each master node system image. On an SGI ICE cluster that includes both SGI ICE compute nodes and flat compute nodes, you need three `cp(1)` commands, one for each image type. On an SGI Rackable cluster, you need one `cp(1)` command.

Example 1. On an SGI ICE cluster with RHEL, type the following commands, one for each master image:

```
# cp /etc/localtime /var/lib/systemimager/images/ice-rhel7.1/etc
# cp /etc/localtime /var/lib/systemimager/images/rhel7.1/etc
# cp /etc/localtime /var/lib/systemimager/images/lead-rhel7.1/etc
```

Example 2. On an SGI Rackable SLES cluster, type the following command to copy the compute node's master image:

```
# cp /etc/localtime /var/lib/systemimager/images/sles12/etc
```

11. Update the RLC image and flat compute node images with your site's time zone information.

On an SGI ICE cluster, type the following command:

```
# pdcp -g leader /etc/localtime /etc/localtime
```

On an SGI cluster that includes flat compute nodes, type the following command:

```
# pdcp -g compute /etc/localtime /etc/localtime
```

12. Proceed to one of the following:

- If you have an SGI ICE cluster, proceed to the following:
"(Conditional) Pushing Changes to SGI ICE Compute Nodes" on page 41
- If you have an SGI Rackable cluster, proceed to the following:
"Verifying the System and Configuring Additional Features" on page 43

(Conditional) Pushing Changes to SGI ICE Compute Nodes

Cluster admin nodes can host multiple forms of system images for each of the rack leader controller (RLC) nodes, SGI ICE compute nodes, and flat compute nodes. For example, you can have some production images and some test images, and you can push the images to the nodes as needed. "(Conditional) Verifying That the `smcrbd` daemon and the `icerackd` Daemon are Running" on page 143 explains how to update the master images for the RLC nodes and the compute nodes. In addition, that procedure explains how to push the updated the node images to the cluster's nodes.

In this topic's procedure, you push the updated SGI ICE compute node images to the SGI ICE compute nodes. The push action enables the SGI ICE compute nodes to run with the updated networking information that you configured in "(Conditional) Verifying That the `smcrbd` daemon and the `icerackd` Daemon are Running" on page 143.

Procedures that describe system operations, later on in this manual, refer back to this procedure because you need to update system images and push out new images as part of several system administration tasks.

The following procedure explains how to push SGI ICE compute node system images to compute nodes.

Procedure 3-4 To push software system images

1. (Conditional) Stop the SGI ICE compute nodes

Complete this step if you want to boot from NFS roots. Do not complete this step if you want to boot from `tmpfs` roots.

Type the following command:

```
# cpower node halt "r*i*n"
```

The preceding command stops all the SGI ICE compute nodes. Use the preceding command when you need to push an updated image out to all the nodes.

2. (Optional) Provide information about the number of racks on your system.

Perform this step if you have a small system with fewer than eight IRUs per RLC.

The procedure pushes the updated SGI ICE compute image to all the SGI ICE compute nodes. This process can run for a long time on large systems. If you have a large number of IRUs, you need the system to perform expansions that enable you to change many SGI ICE compute nodes at a time. If you have fewer than eight IRUs per RLC, however, the expansions are not needed.

The following substeps explain how to prepare the system to work on a smaller number of SGI ICE compute nodes:

- Type the following command to retrieve the identifiers for the RLCs on your system:

```
# cnodes --leader
```

- Type the following command one or more times to suppress unnecessary processing:

```
cadmin --set-max-irus --node rlc_id number_of_racks
```

For *rlc_id*, specify the identifier for one of the RLCs in your system.

For *number_of_racks*, specify the number of IRUs associated with this RLC.

For example, the following command specifies that there is only one IRU associated with the RLC identified as *r1lead*:

```
# cadmin --set-max-irus --node r1lead 1
```

3. Use the `cimage` command, in the following format, to push the changes:

```
cimage --push-rack ice-image_name rack
```

For *image_name*, specify the name of the SGI ICE compute node image that you updated.

For *rack*, specify the nodes. To specify all SGI ICE compute nodes, specify `r*` or `r*i*n*`. To specify only selected nodes, specify `rxixnx`, and substitute specific integer numbers for the `x` characters.

For example, the following command pushes the time zone changes (from "Completing the Customization" on page 38) to all the SGI ICE compute racks:

```
# cimage --push-rack ice-rhel7.1 r\*
```

4. Type the following command to power-up the SGI ICE compute nodes:

```
# cpower node on ""
```

5. Proceed to the following:

"Verifying the System and Configuring Additional Features" on page 43

Verifying the System and Configuring Additional Features

The following procedure explains where you can obtain information about how to configure additional features.

Procedure 3-5 To configure additional features

1. (Conditional) Verify that the blade daemons are running.

Perform the procedure in the following topic if you are attaching an SGI ICE cluster to your network.

"(Conditional) Verifying That the `smcrbd` daemon and the `icerackd` Daemon are Running" on page 143

Return to this step when finished.

2. Configure one or more of the following additional features, some of which are required:

- Mellanox OpenFabrics Enterprise Distribution (OFED) for Linux Software.

If you run RHEL on an SGI ICE cluster or on an SGI Rackable cluster with Mellanox Connect-IB adaptor cards, perform the procedure in the following topic:

"(Conditional) Enabling the Mellanox OpenFabrics Enterprise Distribution (OFED) for Linux Software" on page 167

The procedure explains how to update the OFED software, which is required for performance reasons.

- Networking features, such as network address translation (NAT). For information about how to configure optional features, see the following:

SGI Management Center (SMC) Administration Guide for Clusters.

- Array Services, the hardware event tracker (HET), or CPU frequency scaling. For information about these additional features, see one of the following:

– Chapter 5, "Configuring Additional Features" on page 155

– *SGI Foundation Software (SFS) User Guide*

- SGI Remote Services. This feature provides a secure connection to SGI Customer Support. For more information about SGI Remote Services, see the following website:

https://www.sgi.com/services/support/remote_7x24.html

Installing and Configuring an SGI Cluster System

This chapter contains the following topics:

- "About Performing a New Installation and Configuring the Software on an SGI Cluster" on page 46
- "(Conditional) Migrating From SMC 1.7.x to SMC 3.x" on page 50
- "Planning the Image Installation Method" on page 53
- "Preparing to Install Software on a Cluster" on page 56
- "(Conditional) Setting a Static IP Address for the Baseboard Management Controller (BMC) in the Admin Node" on page 57
- "(Optional) Configuring a High Availability Admin Node or a High Availability Rack Leader Controller (RLC)" on page 59
- "Booting the System" on page 59
- "Configuring the Operating System on the Admin Node" on page 65
- "Configuring the Cluster" on page 79
- "(Conditional) Configuring External Domain Name Service (DNS) Servers" on page 93
- "Synchronizing the Software Repository, Installing Software Updates, Downloading Debugging Packages, and Cloning the Images" on page 94
- "(Conditional) Downloading the Intel Manycore Platform Software Stack (MPSS) Software and Creating Images" on page 97
- "Configuring the Switches" on page 109
- "Configuring the Cluster With the `discover` Command" on page 129
- "Verifying Power Operations and Configuring Power Management" on page 137
- "(Optional) Configuring a Backup Domain Name Service (DNS) Server" on page 146
- "(Conditional) Configuring the InfiniBand Subnetworks" on page 147

About Performing a New Installation and Configuring the Software on an SGI Cluster

SGI installs operating system software on each cluster system before factory shipment occurs. The topics in this chapter include the additional procedures that you need to complete in order to configure the system for your site.

If you want to completely reinstall the operating system and all other software, the topics in this chapter enable you to complete that task. For example, you might need to reinstall the operating system to meet site requirements or to recover a system in case of a disaster.

Figure 4-1 on page 47 depicts the software installation process.

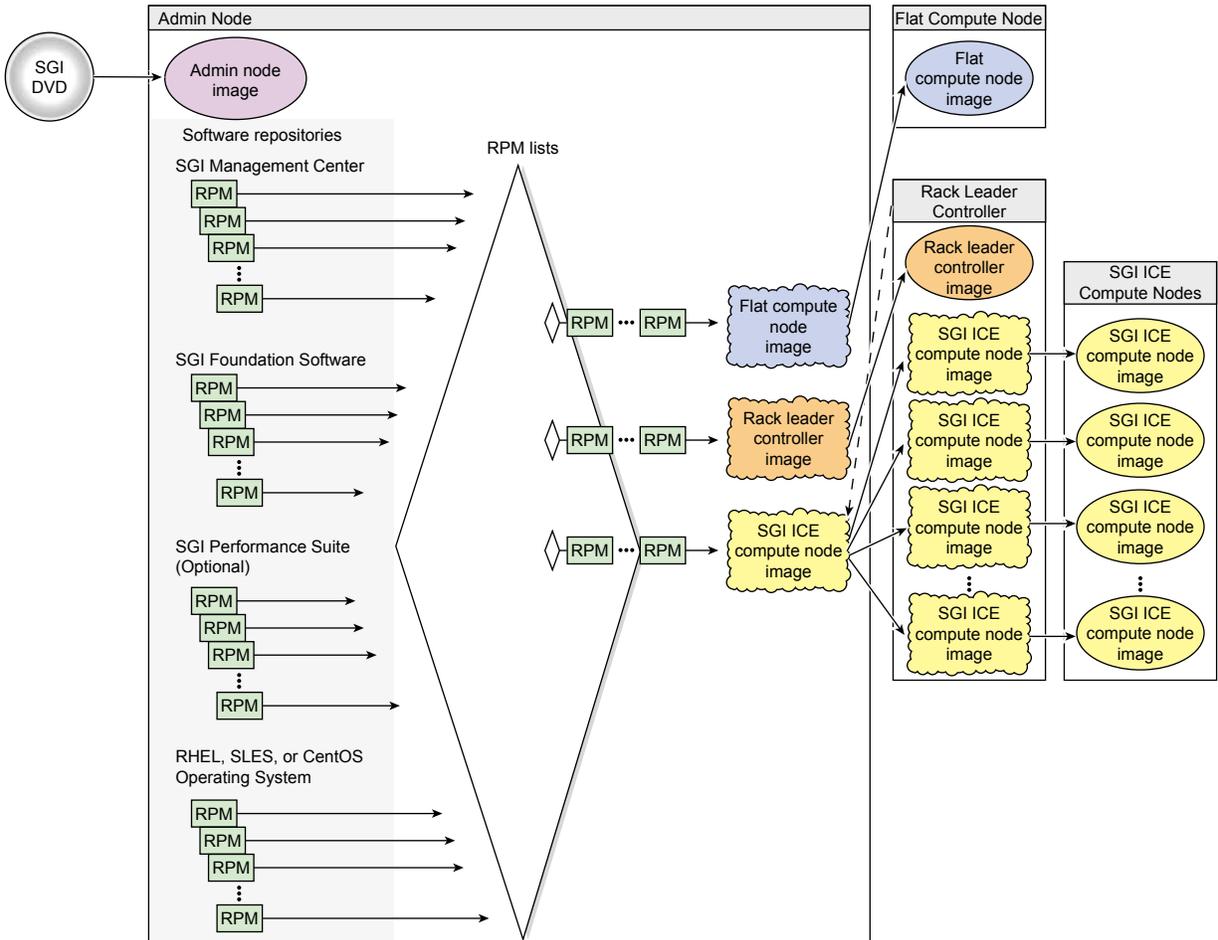


Figure 4-1 SGI Management Center Software Installation Process

Table 4-1 on page 48 shows the installation and configuration procedures to follow if you want to install the a cluster system from scratch. The cluster installation process is the same for SGI ICE clusters and SGI Rackable clusters. In the case of the SGI Rackable clusters, the SMC software omits the steps that install images on RLCs and on SGI ICE compute nodes. In this case, you reinstall the operating system on the nodes and configure everything yourself.

Table 4-1 SGI ICE System Installation and Configuration Process

Step	Task	See
1	(Conditional) Prepare the cluster to run the <code>migrate-sgi-mc</code> script.	"(Conditional) Migrating From SMC 1.7.x to SMC 3.x" on page 50
2	Plan the image installation method.	"Planning the Image Installation Method" on page 53
3	Prepare to install the cluster software.	"Preparing to Install Software on a Cluster" on page 56
4	(Conditional) Configure a static address for the baseboard management controller (BMC) on the admin node. Complete this step only if your site practices require a static IP on the BMC.	"(Conditional) Setting a Static IP Address for the Baseboard Management Controller (BMC) in the Admin Node" on page 57
5	(Optional) Configure a highly available admin node or a highly available rack leader controller (RLC).	"(Optional) Configuring a High Availability Admin Node or a High Availability Rack Leader Controller (RLC)" on page 59
6	Boot the system.	"Booting the System" on page 59
7	Install the operating system on the admin node. You can install the Red Hat Enterprise Server (RHEL), SLES, or CentOS operating system.	"Configuring the Operating System on the Admin Node" on page 65
8	Run the cluster configuration tool. Complete the initial cluster configuration tasks, which include the following: <ul style="list-style-type: none"> • Set up software repositories for required and optional software. • Install the admin node software. • Configure network settings. • Configure the NTP server. • Set up the initial admin node infrastructure. • Configure the house network DNS resolvers. 	"Configuring the Cluster" on page 79
9	(Conditional) Configure external domain name service (DNS). If you want to configure network address translation, you also need to configure an external DNS.	"(Conditional) Configuring External Domain Name Service (DNS) Servers" on page 93

Step	Task	See
10	Synchronize the repository updates, apply the latest patches to the newly installed software, and clone the images.	"Synchronizing the Software Repository, Installing Software Updates, Downloading Debugging Packages, and Cloning the Images" on page 94
11	(Conditional) Download the Intel Manycore Platform Software Stack (MPSS).	"(Conditional) Downloading the Intel Manycore Platform Software Stack (MPSS) Software and Creating Images" on page 97
12	(Conditional) Run the migration script	"(Conditional) Running the Migration Script and Editing the Cluster Definition File" on page 108
13	Configure the switches.	"Configuring the Switches" on page 109
14	Use the <code>discover</code> command to install and configure software on the rack leader controller and the flat compute nodes.	"Configuring the Cluster With the <code>discover</code> Command" on page 129
15	Configure power management.	"Verifying Power Operations and Configuring Power Management" on page 137
16	(Conditional) Verify that the blade daemons are running	"(Conditional) Verifying That the <code>smcrbd</code> daemon and the <code>icerackd</code> Daemon are Running" on page 143
17	(Optional) Configure a backup domain name service (DNS) server on a flat compute node.	"(Optional) Configuring a Backup Domain Name Service (DNS) Server" on page 146
17	Configure the InfiniBand subnetworks.	"(Conditional) Configuring the InfiniBand Subnetworks" on page 147
19	Configure optional features.	Chapter 5, "Configuring Additional Features" on page 155

(Conditional) Migrating From SMC 1.7.x to SMC 3.x

Complete the procedure in this topic if you want to upgrade an existing cluster from SMC 1.7.x to SMC 3.x.

The `migrate-sgi-mc` script ensures a smooth transition when you upgrade the SMC software on your cluster from SMC 1.7.x to SMC 3.x. The conversion process requires you to back up your current cluster to a network location, run the `migrate-sgi-mc` script, and install SMC 3.x.

The following procedures explain the steps that prepare your system for migration if you plan to install SMC 3.x as part of a cluster migration from SMC 1.7.x:

- "Preparing the Payloads for Migration" on page 50
- "Backing Up the Payloads and Running the Migration Script" on page 52

Preparing the Payloads for Migration

The following procedure explains how to verify the payloads and their kernels and how to export the kernels.

Procedure 4-1 To verify the SMC 1.7.x payloads and kernels

1. Use the `cd(1)` command to change to the directory that hosts the payload files, and use the `ls(1)` command to list the payload files:

For example:

```
# cd /opt/sgi/sgimc/imaging/root/payloads
# ls
Compute-SLES-11-3          payload-test-11-3          .vcs.profile
Compute-SLES-11-3-legacy  ss-test                    windows-cleint-test
Compute-SLES-11-3-migrate importedComputeSles11-3    .vcs.entries
```

2. Use the `cd(1)` command to change to one of the payload directories, and use the `vcs list` command to list the kernels.

For example:

```
# cd Compute-SLES-11-3
# vcs list
vcs - Version 1.4.0 Build 3
Copyright (c) 2006-2013 Silicon Graphics, Inc. All rights reserved.
=====
```

```
java vendor "Oracle Corporation"
java version "1.7.0_21"
OpenJDK Runtime Environment (build 1.7.0_21-b02)
OpenJDK 64-Bit Server VM (build 23.7-b01, mixed mode)
```

```
=====
payloads
```

```
-----
Centos6.6Brian (2)
Compute-SLES-11-3 (3)
Compute-SLES-11-3-legacy (1)
importedComputeSles11-3 (1)
payload-test-11-3 (1)
ss-test (1)
```

```
kernels
```

```
-----
Compute-SLES-11-3 (1)
RAID-SLES-11-3 (1)
centos6.6brian (1)
importedComputeSles11-3 (1)
```

```
images
```

```
-----
Compute-SLES-11-3 (4)
RAID (1)
centos6.6Brian (1)
importedComputeSles11-3 (1)
```

The preceding output shows that payload `Compute-SLES-11-3` includes kernels. In addition, the kernels are checked in. This payload is ready to be checked out. The next step explains what to do if the payload does not include checked-in kernels.

3. (Conditional) Use the SMC interface to add a kernel to the payload, and check in the kernel.

Complete this step if the preceding command shows that there are no kernels in this particular payload.

Before you migrate to SMC 3.x, all payloads must have kernels, and all kernels must be checked in.

4. Change to the directory in which you want to store the checked-out payload files.

This can be any directory on the SGI cluster.

For example, type the following command to change to the `.../root` directory:

```
# cd /opt/sgi/sgimc/imaging/root
```

5. Use the `vcs checkout` command, in the following format, to check out the payload to the `.../root` directory:

```
vcs checkout -R payloads -M:payload_name -r:1
```

For *payload_name*, specify the name of the payload.

For example:

```
# vcs checkout -R payloads -M:Compute-SLES-11-3 -r:1
```

6. Proceed as follows:

- If there are additional payloads you need to examine and check out, repeat the preceding steps for the additional payloads.
- If the output shows that this particular payload and kernels are checked into VCS, and this is the last or only payload you need to examine, proceed to the following:

"Backing Up the Payloads and Running the Migration Script" on page 52

Backing Up the Payloads and Running the Migration Script

The following procedure explains how to back up the current cluster configuration to a network location and how to run the `dbix` script.

The `dbix` script exports all SMC database entries into a file.

Procedure 4-2 To back up the SMC 1.7.x configuration to a network location and export the database entries

1. Back up the SGI cluster to a network location.

The migration script performs a major transformation on your cluster. In particular, make sure that the following are included in the backup:

- The entire `/opt/sgi/sgimc/etc` directory.

- All payloads.
2. Use the `cd(1)` command to change to the directory that hosts the exported payloads.

For example, in this procedure, this directory is
`/opt/sgi/sgimc/imaging/root`.

3. Type the following command to export all SMC database entries:

```
# dbix -x > dbix.txt
```

4. Copy file `dbix.txt` to a backup location on your network.
5. Proceed to the following:

"Planning the Image Installation Method" on page 53

Planning the Image Installation Method

The SMC `discover` command installs software images on the nodes and facilitates adding nodes to a cluster. You use the `discover` command during the initial installation, and you can use the `discover` command again later if you want to reconfigure a node's network settings or you want to update the cluster after a hardware equipment change.

SGI supports three different file transfer methods for use during installation. These methods are `rsync` (default), `UDPCast`, and `BitTorrent`. The `BitTorrent` method is supported for legacy clusters.

The fastest image installation method for your cluster depends on the cluster's topology. Before you begin the installation, familiarize yourself with the image transport and installation methods and make sure that your installation plan uses the method that is most appropriate for your cluster. Your site network configuration can also affect the speed at which the `discover` command can push software to nodes.

The following procedure explains how to determine the image installation method that is most appropriate for your cluster.

Procedure 4-3 To plan the installation method

1. Consider using the cluster definition file.

Perform the following steps if the cluster has a working slot at this time:

- Type the following command to generate a cluster definition file:

```
discover --show-configfile > file_name
```

For *file_name*, type a name for the cluster definition file.

- Save the cluster definition file in a safe place on a computer off of the cluster.

If you have a working cluster definition file, you can supply the file as input to the `discover` command and to the `configure-cluster` command. The file enables you to complete the installation and configuration process complete more quickly. Without a cluster definition file, you need to power on and power off each component during the configuration process. The cluster definition file supplies the information that you would typically define by using the menus in the cluster configuration tool. When you specify a cluster configuration file as input to the `configure-cluster` command or to the `discover` command, the command reads in the options from the file and implements them in the cluster.

If you plan to reinstall the software on a new cluster that you just received from SGI, you can obtain the cluster definition file used in the manufacturing process from your SGI representative.

For more information about the cluster definition file, see the following:

"About the Cluster Definition File" on page 110

2. Determine the number and type of nodes that need to be imaged.

When you run the `discover` command during system installation, only the flat compute nodes and the rack leader controllers (RLCs) receive software images. The SGI ICE compute nodes receive their images directly from their RLC, so you do not need to consider the number of SGI ICE compute nodes in this calculation. Count the number of nodes as follows:

- If you have five or six nodes, the default transport method, `rsync`, is appropriate. For example, if you have three RLCs and two flat compute nodes, you can use the default method. You do not need to consider the number of SGI ICE compute nodes that are associated with each RLC. You do not need to edit your cluster definition file, nor do you need to plan for any additional command line options for the `discover` command.

You do not need to complete the rest of this procedure. Proceed to the following:

"Preparing to Install Software on a Cluster" on page 56

- If you have more than five or six nodes, consider using the UDPcast transport method. If you have hundreds of flat compute nodes, you most definitely need to consider using UDPcast.
3. To use UDPcast, either plan to provide additional arguments to the `discover` command when you run it (later in the installation process) or edit the cluster definition file at this time.

If you specify options on the `discover` command line, those options override those that appear in the configuration file. If you prefer to specify the UDPcast transport on the `discover` command line, plan to include the `udpcast` argument.

For example, if you have three RLCs and 200 flat compute nodes, you can specify the following command:

```
# discover --leaderset 1,3,transport=udpcast --nodeset 1,200,transport=udpcast --configfile myfile --all
```

If you want to edit the cluster definition file at this time, complete the following steps:

- Obtain a copy of the cluster definition file from your sales representative or generate one by typing the following command:

```
discover --show-configfile > filename
```

For *filename*, specify the output file name.

- Open the cluster definition file from within a text editor.
- Search in the file for each block of text that describes a node. Each node block begins with the keyword `temponame=`. For example, the following text block describes one of the RLCs:

```
temponame=r1lead, mgmt_bmc_net_name=head-bmc, mgmt_bmc_net_macs=00:25:90:58:8b:75,  
mgmt_net_name=head, mgmt_net_macs=00:25:90:58:8a:94/00:25:90:58:8a:95, redundant_mgmt_network=yes,  
switch_mgmt_network=yes, mic=0, dhcp_bootfile=grub2, conserver_logging=yes, conserver_ondemand=no,  
console_device=ttyS1
```

- At the end of each node definition block, add the following:

```
, transport=udpcast
```

- Save and close the file.
4. Proceed to the following:
"Preparing to Install Software on a Cluster" on page 56

Preparing to Install Software on a Cluster

The following procedure explains the information you need to obtain before you begin working with the cluster. Your installation session can proceed more quickly if you gather information before you begin.

Procedure 4-4 To prepare for an installation

1. Contact your site's network administrator to obtain network information.

Obtain the information to use when you configure the baseboard management controller (BMC) of the admin node. Your network administrator can provide an IP address, a hostname, or a fully qualified domain name (FQDN) for each of the following addresses:

- (Optional) The current IP address of the BMC on the admin node. You can set the BMC address from a serial console if you do not have this information.
- The address you want to set for the BMC.
- The netmask you want to set for the BMC.
- The default gateway you want to set for the BMC.

Obtain the following information to use when you configure the network for the SGI ICE system:

- Hostname
- Domain name
- IP address
- Netmask
- Default route
- Root password

Obtain the following information about your site's house network:

- IP addresses of the domain name servers (DNSs)
2. (Optional) Obtain the configuration file for your cluster from your SGI representative.

The configuration file contains system data, for example, the MAC address information for the nodes. If you have these addresses, the node discovery process can complete more quickly. For more information about the cluster definition file, see the following:

"Planning the Image Installation Method" on page 53

(Conditional) Setting a Static IP Address for the Baseboard Management Controller (BMC) in the Admin Node

Perform the procedure in this topic if one of the following is true:

- Your site practices require a static IP address for the BMC.
- You want to configure a high availability admin node. In this case, perform this topic's procedure on the BMCs on each of the two admin nodes.

When you set the IP address for the BMC on the admin node, you ensure access to the admin node when the site DHCP server is inaccessible.

The following procedures explain how to set a static IP address.

Procedure 4-5 Method 1 — To change from the BIOS

1. Use the BIOS documentation for the admin node.

Procedure 4-6 Method 2 — To change the IP address from the admin node.

1. Log into the admin node as the root user.
2. Type the following command to retrieve the current network settings:

```
# ipmitool lan print 1
```

3. In the output from the preceding command, look for the IP Address Source line and the IP Address line.

For example:

```
IP Address Source      : DHCP Address
IP Address             : 128.162.244.59
```

Note the IP address in this step and decide whether or not this IP address is acceptable. The rest of this procedure explains how to keep this IP address or to set a different static IP address.

4. Type the following command to specify that you want the BMC to have a static IP address:

```
# ipmitool lan set 1 ipsrc static
```

This step specifies that the IP address on the BMC is a static IP address, and this step sets the IP address to the IP address that is currently assigned to the BMC. If you want to set the IP address to a different IP address, proceed to the following step. If the current IP address is acceptable, you do not need to perform the next step.

5. (Optional) Set a different IP address.

Complete this step if you want to set the static IP address to be different from the IP address that is set currently.

Type `ipmitool` commands in the following format:

```
ipmitool lan set 1 ipaddr ip_addr
ipmitool lan set 1 netmask netmask
ipmitool lan set 1 defgw gateway
```

The arguments are as follows:

Argument	Specification
<i>ip_addr</i>	The IP address you want to assign to the BMC.
<i>netmask</i>	The netmask you want to assign to the BMC.
<i>gateway</i>	The gateway you want to assign to the BMC.

For example, you can type the following commands to set the IP address to 100.100.100.100:

```
# ipmitool lan set 1 ipaddr 100.100.100.100
# ipmitool lan set 1 netmask 255.255.255.0
```

```
# ipmitool lan set 1 defgw 128.162.244.1
```

6. Proceed to one of the following:

- If you want to configure a high availability admin node, proceed to the following:
"(Optional) Configuring a High Availability Admin Node or a High Availability Rack Leader Controller (RLC)" on page 59
- If you want to configure a traditional admin node, proceed to the following:
"Booting the System" on page 59

(Optional) Configuring a High Availability Admin Node or a High Availability Rack Leader Controller (RLC)

SGI supports the ability to configure the admin node and rack leader controllers (RLCs) as highly available nodes in an SGI ICE cluster. If you want to enable high availability (HA) on the admin node or on the RLCs, contact your SGI representative.

Booting the System

You can configure the cluster to boot from one, two (default), three, four, or five *slots*. A slot consists of all the partitions related to a Linux installation.

A multiple-slot disk layout is also called a *cascading dual-root layout* or a *cascading dual-boot layout*. The installer creates the same disk layout on all nodes. Each slot includes the following:

- A `/boot` partition.
- A `/`, or root, partition.
- A `/boot/efi` partition. A slot includes this partition only if the node is an EFI node.

When you insert an SMC operating system installation disk and power-on the admin node, you can select a boot method from the GNU GRUB menu. If you select **Install: Wipe Our and Start Over: Prompted**, which is the default, the installer creates two slots and writes the initial installation to slot 1. After the system is installed, you cannot change the number of slots without destroying the data on the disks.

After you install a multislot cluster, you can boot the cluster with the operating system of your choice. This ability might be useful if you ever want to test an operating system or other software because you can roll back an upgrade completely.

The following are some other characteristics of single-boot systems and multiple-boot systems:

Multiple-boot

You can install different operating systems, or different operating system versions, into different slots. Note that if you have an SGI ICE cluster, the admin node and the RLCs must have the same operating system installed.

RLCs and flat compute nodes boot from their own disk. Data is retained in the master boot record (MBR).

RLC and flat compute node software is reinstalled from the admin node.

As you increase the number of slots, you decrease the amount of disk space per slot. SGI recommends a minimum of 100 GB per slot.

If all slots on your cluster are running either SGI Tempo 2.9.0 (or later) or SMC 3.0 (or later), then your cluster uses the partition layout designed for the SGI Tempo 2.9.0 and later releases. If you upgraded your cluster, it is possible that you have the legacy partition layout on one or more slots. For information about partitions, including those for legacy partitions, see the following:

Appendix D, "Partition Layout Information" on page 223

Note: SMC supports both EFI BIOS and legacy x86_64 BIOS. If you are not sure which BIOS your cluster supports, contact your sales representative.

The following procedure explains how to boot the system and begin the installation.

Procedure 4-7 To boot the system

1. Power-on the admin node.

Single-boot

You can install only one operating system for the entire cluster.

RLCs and flat compute nodes boot from the boot partition in the slot that is currently configured as the boot slot. Only the admin node retains data in the MBR.

Software on the RLCs and flat compute nodes is reinstalled over the network.

A single slot uses all available disk space.

As Figure 4-2 on page 62 shows, the power-on button is on the right of the admin node.

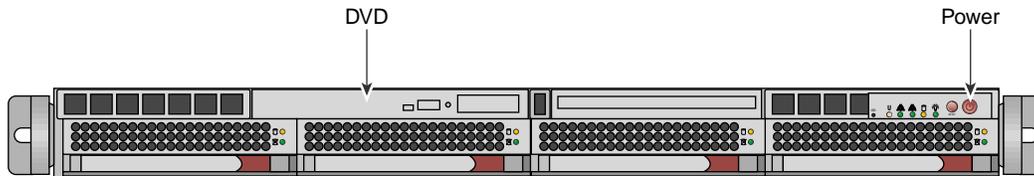


Figure 4-2 Admin Node Power On Button and DVD Drive

2. Insert the SGI Admin Node Autoinstallation DVD into the DVD drive on the admin node.
3. Use the arrow keys to select one of the boot options, press Enter, and monitor the installation.

On the GNU GRUB boot menu, the options are as follows:

- **Display Instructions**

Select this option if you want information about custom boot parameters. This option displays information and returns to the boot menu.

Each boot option has a set of default behaviors. In addition to the default behaviors, you can specify custom boot parameters if you select one of the following options:

- **Install: Install to Designated Slot**
- **Install: Wipe Out and Start Over: Prompted**
- **Install: Custom, type 'e' to edit kernel parameters**

If you think you might want to specify one or more custom boot parameters, for example, `console=`, select the **Display Instructions** option and familiarize yourself with these parameters before you select an actionable parameter.

- **Install: Install to Designated Slot**

Select this option if you have an open slot on your cluster, and you want to install an operating system in the slot. If you select this option, only the open slot is affected. All other slots remain as configured.

- **Install: Wipe Out and Start Over: Prompted**

Select this option if you want to reinstall the cluster. This options destroys all information currently on the cluster. The installer partitions the admin node with the specified number of slots, and the installer writes the initial installation to the designated slot. In the factory, SGI configures systems by default with two slots, and the installation is in slot 1. For example, for an initial installation, select this option.

- **Rescue: Prompted**

Select this option to create a troubleshooting environment.

- **Install: Custom, type 'e' to edit kernel parameters**

Select this option if you want to perform a custom installation. This option lets you supply all boot options as command line parameters, as opposed to specifying boot options in response to the system prompts that the other methods offer. More information is available in **Display instructions**. SGI recommends this option only for users with installation experience.

All the options launch you into an installation dialog, and at the end of the dialog, the final question asks you to confirm your choices. In this way, you have the chance to cancel your choices and return to the GNU GRUB boot menu to start over. The following are some of the installation dialog prompts that appear when you select a boot option:

- **Enter number of slots to allow space for: (1-5):**

Type 1, 2, 3, 4, or 5, and press Enter.

This dialog question appears only if you select **Install: Wipe Out and Start Over: Prompted** from the GNU GRUB menu.

- **Enter which slot to install to:**

Type 1, 2, 3, 4, or 5, and press Enter.

This dialog question appears only if you select **Install: Install to Designated Slot** from the GNU GRUB menu.

- **Destructively bypass sanity checks? (y/n):**

If you type **y** and press Enter, the installer proceeds without checking to see if there is any data in the partition.

If you type **n** and press Enter, the installer checks to see if there is data in the partition before proceeding.

- **Additional parameters (like console=, etc):**

If you want to specify any additional boot parameters, type them in a comma-separated list and press Enter.

For information about the boot parameters that are available, select Display Instructions from the GNU GRUB menu and press Enter.

- **OK to proceed? (y/n):**

If you type **y** and press Enter, the boot proceeds.

If you type **n** and press enter, the menu returns you to the main GNU GRUB menu.

The installation can take several minutes.

4. Remove the operating system installation DVD.

5. At the # prompt, type **reboot**.

This is the first boot from the admin node's hard disk.

6. (Optional) Suppress log messages.

If you want to suppress the admin node's log message output to the screen during the boot, edit file `/etc/syscontrol.conf` and add the following line to the top of the file (line 1):

```
kernel.printk = 2 4 1 7
```

In the preceding `kernel.printk` line, the spaces between the numbers 2 4 1 7 are Tab characters.

7. Proceed to the following:

"Configuring the Operating System on the Admin Node" on page 65

Configuring the Operating System on the Admin Node

The cluster systems support the Red Hat Enterprise Linux (RHEL) operating system, CentOS, and the SLES and operating system. SMC supports the following operating systems on the admin node:

- Red Hat Enterprise Linux (RHEL) 7.1 and RHEL 6.6
- SLES 11 SP3

Use one of the following procedures to install your operating system software on the admin node:

- "Configuring RHEL 7.x on the Admin Node" on page 65
- "Configuring RHEL 6.x on the Admin Node" on page 69
- "Configuring SLES 11 on the Admin Node" on page 74

Note: In SGI documentation, you can assume that feature descriptions for RHEL platforms also pertain to CentOS platforms unless otherwise noted.

Configuring RHEL 7.x on the Admin Node

This procedure in this topic describes how to configure Red Hat Enterprise Linux 7.x on the admin node.

Procedure 4-8 To install RHEL 7.x on an SGI ICE admin node

1. Use one of the following methods to connect to the admin node:
 - Through the intelligent platform management interface (IMPI) tool
 - Through the console attached to the cluster
 - Through an attached, separate keyboard, video display terminal, and mouse
2. Use a text editor, such as `vi` or `vim`, to open file `/etc/hosts`.
3. Add a line in the following format to file `/etc/hosts`:

```
admin_node_IP admin_node_FQDN admin_node_hostname
```

The variables in the preceding line are as follows:

- For *admin_node_IP*, type the IP address of the admin node.
- For *admin_node_FQDN*, type the fully qualified domain name (FQDN) of the admin node.
- For *admin_node_hostname*, type the hostname of the admin node.

For example, add the following line:

```
128.162.244.88 acme-admin.acme.usa.com acme-admin
```

4. Save and close file `/etc/hosts`.
5. Type the following command to set the admin node's hostname:

```
# hostnamectl set-hostname admin_node_hostname
```

For *admin_node_hostname*, type the hostname of the admin node. Make sure to type the hostname, which is the short name. Do not type the admin node's FQDN, which is the longer name.

6. Type the following command to create file `/etc/sysconfig/network` with no content:

```
# touch /etc/sysconfig/network
```

7. Use a text editor to open file `ifcfg-eth0`.

On RHEL 7.x systems, the path to this file is as follows:

```
/etc/sysconfig/network-scripts/ifcfg-eth0
```

8. Replace the contents of the `ifcfg-eth0` file with the following lines:

```
TYPE="Ethernet"  
BOOTPROTO="none"  
DEFROUTE="yes"  
IPV4_FAILURE_FATAL="no"  
IPV6INIT="no"  
IPV6_AUTOCONF="yes"  
IPV6_DEFROUTE="yes"  
IPV6_PEERDNS="yes"  
IPV6_PEERROUTES="yes"  
IPV6_FAILURE_FATAL="no"  
NAME="eth0"  
DEVICE="eth0"
```

```

ONBOOT="yes"
IPADDR="128.162.233.35"           # replace with your IP address for the admin node
PREFIX="24"
GATEWAY="128.162.233.1"         # replace with your site gateway
DNS1="128.162.236.210"         # replace with your site's primary DNS IP address
DNS2="198.149.32.11"          # replace with your site's secondary DNS IP address
DOMAIN="americas.sgi.com engr.sgi.com sgi.com" # replace with your site domain
HWADDR="XX:XX:XX:XX:XX:XX"     # replace with the MAC address of eth0

```

Note: SMC does not support IPV6 at this time. The following line is needed for this installation:

```
IPV6INIT="no"
```

You can remove the lines that start with `IPV6_` from the `ifcfg-eth0` file, or you can retain those lines for completeness. SMC ignores these lines during the installation.

9. Save and close file `ifcfg-eth0`.
10. Type the following commands to bring up `eth0` with the updated networking information:

```
# ifdown eth0
# ifup eth0
```

11. Type the following command to restart the name service cache daemon, `nscd`:

```
# systemctl restart nscd
```

12. Type the following command to retrieve the admin node's current time zone information:

```
# strings /etc/localtime | tail -1
CST6CDT,M3.2.0,M11.1.0
```

The previous output shows the admin node set to US Central time. If the output you see is not correct for this cluster, perform the following steps:

- Type the following command to change to the directory that contains the time zone configuration files:

```
# cd /usr/share/zoneinfo
```

- Select a file from that directory that describes the time zone for the admin node.
- Type the following commands to enable the new time zone configuration file.

For example:

```
# /bin/cp -l /usr/share/zoneinfo/time_zone_file /etc/localtime.$$  
# /bin/mv /etc/localtime.$$ /etc/localtime
```

For *time_zone_file*, type the name of the time zone file that you need from the /usr/share/zoneinfo directory.

For example, type the following commands to change the admin node's time zone to US Pacific time:

```
# /bin/cp -l /usr/share/zoneinfo/PST8PDT /etc/localtime.$$  
# /bin/mv /etc/localtime.$$ /etc/localtime
```

- Type the following command to confirm the time zone:

```
# strings /etc/localtime | tail -1  
PST8PDT,M3.2.0,M11.1.0
```

13. (Conditional) Edit file /etc/ntp.conf to direct requests to the network time protocol (NTP) server at your site.

Complete the following steps if you want to direct requests to your site's NTP server instead of to the public time servers of the pool.ntp.org project:

- Use a text editor to open file /etc/ntp.conf.
- Insert a pound character (#) into column 1 of each of each line that includes rhel.pool.ntp.org.

Note: Do not edit or remove entries that serve the cluster networks.

- At the end of the file, add a line that points to your site's NTP server.

The following is an example of a correctly edited file:

```
# Use public servers from the pool.ntp.org project.  
# Please consider joining the pool (http://www.pool.ntp.org)  
# server 0.rhel.pool.ntp.org
```

```
# server 1.rhel.pool.ntp.org
# server 2.rhel.pool.ntp.org
server ntp.mycompany.com
```

The preceding output has been truncated at the right for inclusion in this guide.

- Type the following command to restart the NTP server:

```
# systemctl restart ntpd
```

14. (Conditional) Type a tilde character (~) and then a period character (.) to exit from the IPMI tool.

Complete this step if you connected to the system through the IPMI tool.

15. (Optional) Configure the system so that you can perform the installation from a VGA screen and can perform later operations from a serial console.

If you want to enable this capability, perform the following steps:

- Use a text editor to open file `/boot/grub/menu.lst`.
- Search the file for the word `kernel` at the beginning of a line.
- Add the following to the `kernel` line: `console=type`.

For example:

```
kernel /boot/vmlinuz-2.6.16.56-0.12-smp root=/dev/disk/by-label/sgiroot console=ttyS1,38400n8
splash=silent showopts
```

- Add the `console=type` parameter to the end of every `kernel` line. By default, this is set to `ttyS1,38400n8`. You might have `ttys2`, for example.

Later, if you want to access the admin node from only a VGA, you can remove the `console=` parameters.

16. Proceed to the following:

"Configuring the Cluster" on page 79

Configuring RHEL 6.x on the Admin Node

This procedure in this topic describes how to configure RHEL 6.x on the admin node.

Procedure 4-9 To install RHEL 6.x on an SGI ICE admin node

1. Use one of the following methods to connect to the admin node:
 - Through the intelligent platform management interface (IMPI) tool
 - Through the console attached to the cluster
 - Through a separate keyboard, video display terminal, and mouse
2. Use a text editor, such as `vi` or `vim`, to open file `/etc/sysconfig/network-scripts/ifcfg-eth0`.
3. Add lines for the `IPADDR`, `NETMASK`, and `NETWORK` values appropriate for your site network to file `/etc/sysconfig/network-scripts/ifcfg-eth0`.

For example:

```
IPADDR=128.162.244.88
NETMASK=255.255.255.0
NETWORK=128.162.244.0
```

4. Save and close file `/etc/sysconfig/network-scripts/ifcfg-eth0`.
5. Use a text editor to create file `/etc/sysconfig/network`.
6. Add the following three lines to file `/etc/sysconfig/network`:

```
NETWORKING=yes
HOSTNAME=admin_node_hostname
GATEWAY=gateway_IP_address
```

For *admin_node_hostname*, type the hostname you want to assign to the admin node.

For *gateway_IP_address*, type the IP address of the gateway for your house network.

For example:

```
NETWORKING=yes
HOSTNAME=my-system-admin
GATEWAY=128.162.244.1
```

7. Save and close file `/etc/sysconfig/network`.
8. Use a text editor to open file `/etc/hosts`.

9. Add a line in the following format to file `/etc/hosts`:

```
admin_node_IP admin_node_FQDN admin_node_hostname
```

The variables in the preceding line are as follows:

- For *admin_node_IP*, type the IP address of the admin node.
- For *admin_node_FQDN*, type the fully qualified domain name (FQDN) of the admin node.
- For *admin_node_hostname*, type the hostname of the admin node.

For example, add the following line:

```
128.162.244.88 my-system-admin.domain-name.mycompany.com my-system-admin
```

10. Save and close file `/etc/hosts`.
11. Type the following command to set the admin node hostname:

```
# hostname admin_node_hostname
```

For *admin_node_hostname*, type the hostname of the admin node.

For example:

```
# hostname my-system-admin
```

12. Use a text editor to create file `/etc/resolv.conf`.
13. Add lines to file `/etc/resolv.conf` that specify the search domain and the domain name service (DNS) servers at your site.

Later in the configuration process, when you run the cluster configuration tool, the tool uses the DNS servers you specify in this step for its defaults.

Specify lines with the following format:

```
search search_domain  
nameserver name_server_IP  
nameserver name_server_IP
```

The following is an example `resolv.conf` file:

```
search mydomain.com  
nameserver 192.168.0.1  
nameserver 192.168.0.25
```

14. Type the following `nscd(8)` command to force the invalidation of the name service cache daemon:

```
# nscd -i hosts
```

15. Type the following commands, in the order shown, to restart services:

```
# /etc/init.d/network restart
# /etc/init.d/rpcbind start
# /etc/init.d/nfslock start
```

16. Type the following command to retrieve the admin node's current time zone information:

```
# strings /etc/localtime | tail -1
CST6CDT,M3.2.0,M11.1.0
```

The previous output shows the admin node set to US Central time. If the output you see is not correct for this cluster, perform the following steps:

- Type the following command to change to the directory that contains the time zone configuration files:

```
# cd /usr/share/zoneinfo
```

- Select a file from that directory that describes the time zone for the admin node.
- Type the following commands to enable the new time zone configuration file.

For example:

```
# /bin/cp -l /usr/share/zoneinfo/time_zone_file /etc/localtime.$$
# /bin/mv /etc/localtime.$$ /etc/localtime
```

For *time_zone_file*, type the name of the time zone file that you need from the `/usr/share/zoneinfo` directory.

For example, type the following commands to change the admin node's time zone to US Pacific time:

```
# /bin/cp -l /usr/share/zoneinfo/PST8PDT /etc/localtime.$$
# /bin/mv /etc/localtime.$$ /etc/localtime
```

- Type the following command to confirm the time zone:

```
# strings /etc/localtime | tail -1
PST8PDT,M3.2.0,M11.1.0
```

17. (Conditional) Edit file `/etc/ntp.conf` to direct requests to the network time protocol (NTP) server at your site.

Complete the following steps if you want to direct requests to your site's NTP server instead of to the public time servers of the `pool.ntp.org` project:

- Use a text editor to open file `/etc/ntp.conf`.
- Insert a pound character (#) into column 1 of each of each line that includes `rhel.pool.ntp.org`.

Note: Do not edit or remove entries that serve the cluster networks.

- At the end of the file, add a line that points to your site's NTP server.

The following is an example of a correctly edited file:

```
# Use public servers from the pool.ntp.org project.
# Please consider joining the pool (http://www.pool.ntp.org)
# server 0.rhel.pool.ntp.org
# server 1.rhel.pool.ntp.org
# server 2.rhel.pool.ntp.org
server ntp.mycompany.com
```

The preceding output has been truncated at the right for inclusion in this guide.

- Type the following command to restart the NTP server:

```
# /etc/init.d/ntpd restart
```

18. (Conditional) Type a tilde character (~) and then a period character (.) to exit from the IPMI tool.

Complete this step if you connected to the system through the IPMI tool.

19. (Optional) Configure the system so that you can perform the installation from a VGA screen and can perform later operations from a serial console.

If you want to enable this capability, perform the following steps:

- Use a text editor to open file `/boot/grub/menu.lst`.
- Search the file for the word `kernel` at the beginning of a line.
- Add the following to the kernel line: `console=type`.

For example:

```
kernel /boot/vmlinuz-2.6.16.56-0.12-smp root=/dev/disk/by-label/sgiroot console=ttyS1,38400n8 splash=silent showopts
```

- Add the `console=type` parameter to the end of every kernel line. By default, this is set to `ttyS1,38400n8`. You might have `ttys2`, for example.

Later, if you want to access the admin node from only a VGA, you can remove the `console=` parameters.

20. Proceed to the following:

"Configuring the Cluster" on page 79

Configuring SLES 11 on the Admin Node

The SLES YAST interface enables you to install the SLES operating system on a cluster. To navigate the YAST modules, use key combinations such as the following:

- The Tab key moves the cursor forward, and the Shift + Tab keys move the cursor backward.
- The arrow keys move the cursor up, down, left, and right.
- To use shortcuts, press the Alt key + the highlighted letter.
- Press Enter to complete or confirm an action.
- Press Ctrl + L to refresh the screen.

For more information about navigation, see Appendix A, "YAST Navigation" on page 201.

The following procedure explains how to use YAST to install SLES 11 on a cluster.

Procedure 4-10 To install SLES 11 on an SGI ICE admin node

1. Connect to the admin node by one of the following methods:

- Through the intelligent platform management interface (IMPI) tool
 - Through the console attached to the cluster
 - Through a separate keyboard, video display terminal, and mouse
2. On the **Language and Keyboard Layout** screen, complete the following steps:
 - Select your language
 - Select your keyboard layout
 - Select **Next**.
 3. On the **Welcome** screen, select **Next**.
 4. On the **Hostname and Domain Name** screen, complete the following steps:
 - Type the hostname for this cluster.
 - Type the domain name.
 - Clear the box next to **Change Hostname via DHCP**. The box appears with an x in it by default, but you need to clear this box.
 - Select **Assign Hostname to Loopback IP**. Put an x in this box.
 - Select **Next**.
 5. On the **Network Configuration** screen, complete the following steps:
 - Select **Change**. A pop-up window appears.
 - On the pop-up window, choose **Network Interfaces**.
 6. On the **Network Settings** screen, complete the following steps:
 - Highlight the first network interface card that appears underneath **Name**.
 - Select **Edit**.
 7. On the **Network Card Setup** screen, specify the admin node's house/public network interface.

Figure 4-3 on page 76 shows the **Network Card Setup** screen.

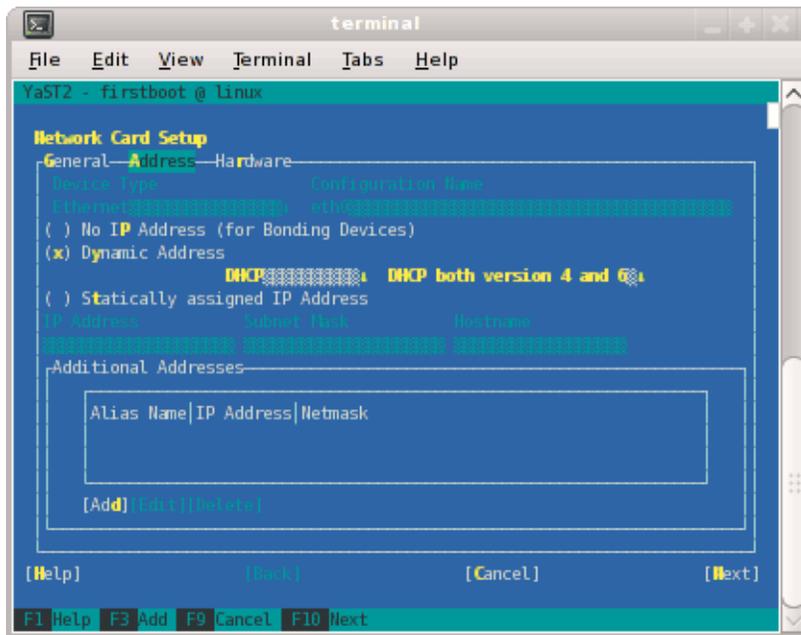


Figure 4-3 Network Card Setup Screen

Complete the following steps:

- Select **Statically Assigned IP Address**. SGI recommends a static IP address, not DHCP, for the admin node.
- In the **IP Address** field, type the system's IP address.
- In the **Subnet Mask** field, type the system's subnet mask.
- In the **Hostname** field, type the system's fully qualified domain name (FQDN). SGI requires you to type an FQDN, not the system's shorter hostname, into this field. For example, type `mysystem-admin.mydomainname.com`. Failure to supply an FQDN in this field causes the `configure-cluster` command to fail.
- Select **Next**.

You can specify the default route, if needed, in a later step.

8. On the **Network Settings** screen, complete the following steps:

- Select **Hostname/DNS**.
- In the **Hostname** field, type the system's fully qualified domain name (FQDN).
- In the **Domain Name** field, type the domain name for your site.
- Put an **x** in the box next to **Assign Hostname to Loopback IP**.
- In the **Name Servers and Domain Search List**, type the name servers for your house network.
- Back at the top of the screen, select **Routing**.

The **Network Settings > Routing** screen appears.

- In the **Default Gateway** field, type your site's default gateway.
- Select **OK**.

9. On the **Network Configuration** screen, click **Next**.

The **Saving Network Configuration** screen appears and saves your configuration.

10. On the **Clock and Time Zone** screen, complete the following steps:

- Select your region.
- Select your time zone.
- (Optional) In the **Hardware Clock Set To** field, choose **Local Time** or accept the default of **UTC**.
- Select **Next**.

This step synchronizes the time in the BIOS hardware with the time in the operating system. Your choice depends on how the BIOS hardware clock is set. If the clock is set to GMT, which corresponds to UTC, your system can rely on the operating system to switch from standard time to daylight savings time and back automatically.

11. On the **Password for System Administrator "root"** screen, complete the following steps:

- In the **Password for root User** field, type the password you want to use for the root user.

This password becomes the root user's password for all the nodes on the ICE system. These nodes are as follows:

- admin node
- Flat compute nodes
- Rack leader controller (RLC) (Optional)
- SGI ICE compute nodes (blades) (Optional)
- In the **Confirm password** field, type the root user's password again.
- In the **Test Keyboard Layout** field, type a few characters.

For example, if you specified a language other than English, type a few characters that are unique to that language. If these characters appear in this plain text field, you can use these characters in passwords safely.

- Select **Next**.
12. On the **User Authentication Method** screen, select one of the authentication methods and select **Next**.

Typically, users accept the default (**Local**).

13. On the **New Local User** screen, create additional user accounts or select **Next**.

If you do not create additional users, select **Yes** on the **Empty User Login** warning pop-up window, and select **Next**.

14. On the **Installation Completed** screen, select **Finish**.
15. Type a tilde character (~) and then a period character (.) to exit from the IPMI tool.
16. Log into the admin node, open file `/etc/hosts` within a text editor, and verify that the admin node's fully qualified domain name (FQDN) and hostname are entered correctly.

For example, the following `/etc/hosts` file entry contains the correct data in the three required fields and is correct for an admin node with an IP address of 100.100.100.100, an FQDN of `mysystem-admin.mydomain.com`, and a hostname of `mysystem-admin`:

```
100.100.100.100      mysystem-admin.mydomain.com      mysystem-admin
```

Make sure that the `/etc/hosts` file on the admin node contains the required information. If it does not, edit the `/etc/hosts` file to contain the three required fields as the preceding example shows.

17. Confirm that the system is working as expected.

If necessary, restart YAST to correct settings.

18. (Optional) Configure the system so that you can perform the installation from a VGA screen and can perform later operations from a serial console.

If you want to enable this capability, perform the following steps:

- Use a text editor to open file `/boot/grub/menu.lst`.
- Search the file for the word `kernel` at the beginning of a line.
- Add the following to the kernel line: `console=type`.

For example:

```
kernel /boot/vmlinuz-2.6.16.56-0.12-smp root=/dev/disk/by-label/sgiroot console=ttyS1,38400n8 splash=silent showopts
```

- Add the `console=type` parameter to the end of every kernel line. By default, this is set to `ttyS1,38400n8`. You might have `ttys2`, for example.

Later, if you want to access the admin node from only a VGA, you can remove the `console=` parameters.

19. Proceed to the following:

"Configuring the Cluster" on page 79

Configuring the Cluster

Configuring the cluster includes the following actions:

- Creating repositories for software installation files and updates.
- Installing the admin node's cluster software.
- Configuring the cluster subdomain and examine other network settings. The cluster subdomain is likely to be different from the `eth0` domain on the admin node itself.

- Configuring the NTP server.
- Installing the cluster's software infrastructure. This step can take 30 minutes.
- Configuring the house network's DNS resolvers.

The following procedure explains how to use either the cluster configuration tool or the cluster definition file to configure the cluster:

Procedure 4-11 To configure the cluster

1. Locate your site's SGI software distribution DVDs or verify the path to your site's online software repository.

You can install the software from either physical media or from an ISO on your network.

2. From the VGA screen, or through an `ssh` connection, log into the admin node as the root user.

SGI recommends that you run the cluster configuration tool either from the VGA screen or from an `ssh` session to the admin node. Avoid running the `configure-cluster` command from a serial console.

3. Use either Method 1 or Method 2 to configure the cluster.

Method 1 — Using the Cluster Configuration Tool — is as follows:

- Type the following command to start the cluster configuration tool:

```
# /opt/sgi/sbin/configure-cluster
```

- Proceed to the following step:

Procedure 4-11, step 4 on page 81

Method 2 — Using the Cluster Definition File — is as follows:

- Type `crepo` commands to create repositories for each of the following software's initial installation packages and for updates:
 - The operating system software, either RHEL, SLES, or CentOS
 - SGI Foundation Software
 - SGI Management Center
 - (Optional) SGI Performance Suite

Use the `crepo` command in the following format:

```
crepo --add rpm_repo_directory --custom rpm_repo_name
```

The variables in this command are as follows:

- For *rpm_repo_directory*, specify the full path to the directory that contains the RPM files.

If you have hard media mounted in the admin node's DVD drive, specify the path to that media. If you have the software for the operating system and the SGI packages in an ISO file on your network, specify the path to the files on your network.

- For *rpm_repo_name*, create a name for the image. You can specify the same name for both *rpm_repo_directory* and *rpm_repo_name*.

After SMC builds the image, the `cinstallman --show-images` command returns the *rpm_repo_name* in the Image Name column of its output. Below the image name, the output also shows the kernels used for each image.

For example, type the following commands:

```
# crepo --add /tmp/sles11sp3 --custom spes11sp3
# crepo --add /tmp/sfs --custom sfs
# crepo --add /tmp/smc --custom smc
# crepo --add /tmp/sps --custom sps
```

- Type the following command to define the cluster according to the content in the cluster definition file:

```
# /opt/sgi/sbin/configure-cluster --configfile path
```

For *path*, specify the path to the configuration file.

- Proceed to the following step:

Procedure 4-11, step 35 on page 91

4. On the cluster configuration tool's **Initial Configuration Check** screen, select **OK** on the initial window.

Figure 4-4 on page 82 shows the initial window.

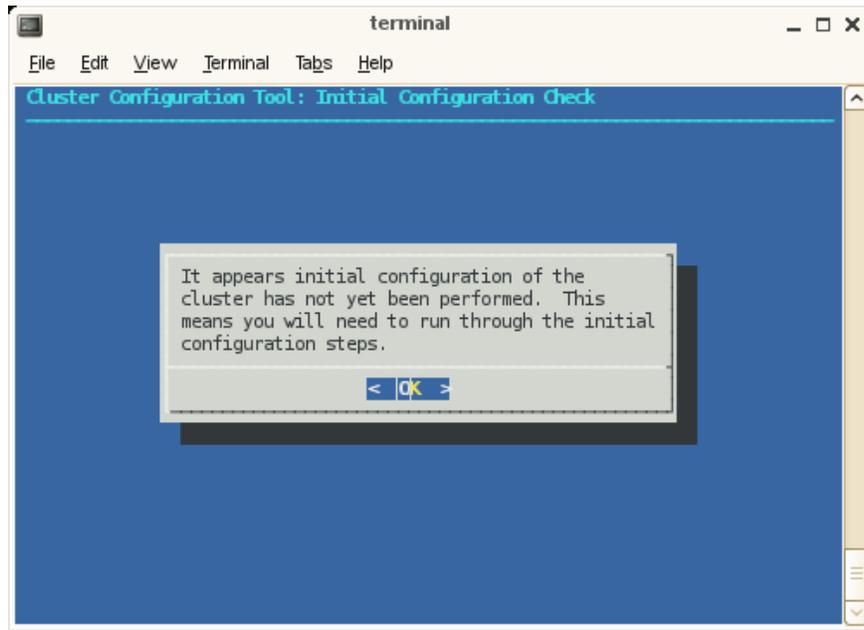


Figure 4-4 Initial Configuration Check Screen

The cluster configuration tool recognizes a configured cluster. If you start the tool on a configured SGI ICE system, it opens into the **Main Menu**.

5. On the **Initial Cluster Setup** screen, select **OK** on the screen.

Figure 4-5 on page 83 shows the window.

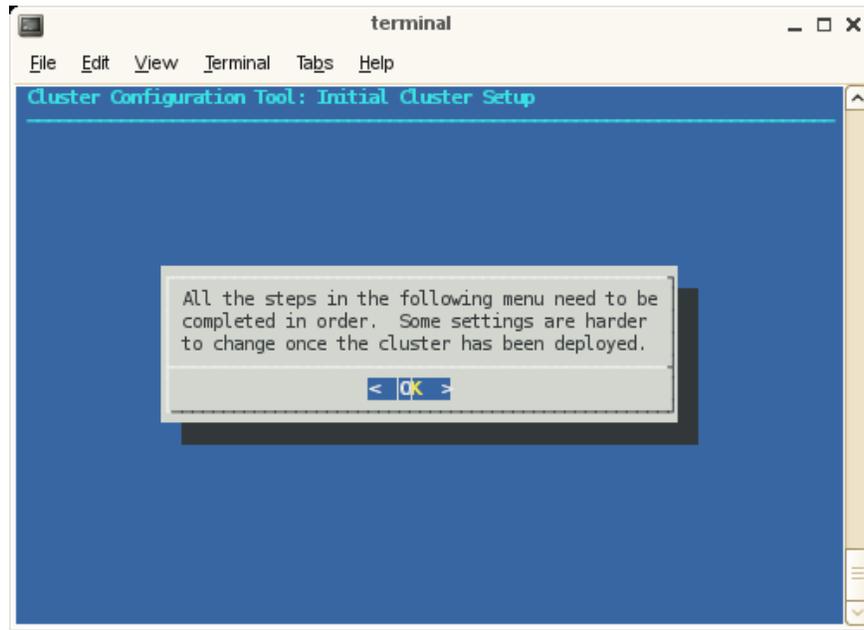


Figure 4-5 Initial Cluster Setup Screen with the initial screen

6. On the **Initial Cluster Setup** screen, select **R Repo Manager: Set Up Software Repos**, and click **OK**.

Figure 4-6 on page 84 shows the **Initial Cluster Setup** screen with the task menu. This procedure guides you through the tasks you need to perform for each of the menu selections on the **Initial Cluster Setup** screen.

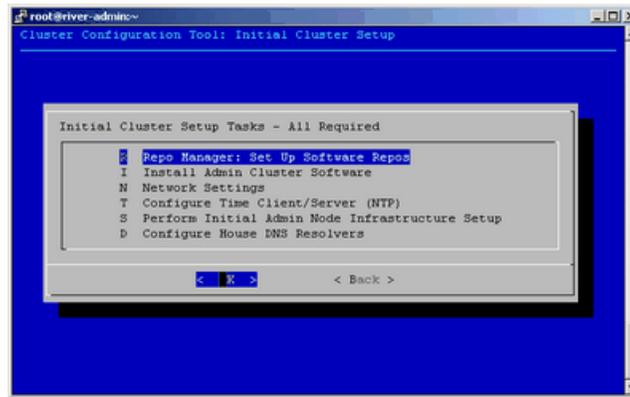


Figure 4-6 Initial Cluster Setup Tasks Screen

The next few steps create software repositories for the initial installation packages and for updates. You need to create repositories for the following software:

- The operating system software, either RHEL or SLES
- SGI Foundation Software
- SGI Management Center
- (Optional) SGI Performance Suite

The menu system prompts you to insert hard media or specify a path for some of the preceding software, so locate your system disks before you proceed.

7. On the **One or more ISOs were embedded on the ...** screen, select **Yes**.
8. On the **Repositories are created ...** screen, press **Enter**.
9. On the **You will now be prompted to add additional media ...** screen, select **OK**.
10. On the **Would you like to register media with Tempo? ...** screen, select **Yes**.
11. On the **Please either insert the media in your DVD drive ...** screen, select either **Insert DVD** or **Use Custom path/url**.

Proceed as follows:

- To install the software from DVDs, perform the following steps:

- Insert a DVD.
- Select **Mount inserted DVD**.
- On the **Media registered successfully with crepo ...** screen, select **OK**, and eject the DVD.
- On the **Would you like to register media with Tempo? ...** screen, select **Yes** if you have more software that you need to register.

If you select **Yes**, repeat the preceding tasks in this sequence for the next DVD.

If you select **No**, proceed to the next step.

- To install the software from a network location, perform the following steps:
 - Select **Use custom path/URL**.
 - On the **Please enter the full path to the mount point or the ISO file ...** screen, type the full path in *server_name:path_name/iso_file* format. This field also accepts a URL or an NFS path. Select **OK** after typing the path.
 - On the **Media registered successfully with crepo ...** screen, select **OK**.
 - On the **Would you like to register media with Tempo? ...** screen, select **Yes** if you have more software that you need to register.

If you select **Yes**, repeat the preceding tasks in this sequence for the next DVD.

If you select **No**, proceed to the next step.

12. Repeat the following steps until all software is installed:

- Procedure 4-11, step 10 on page 84
- Procedure 4-11, step 11 on page 84

If you plan to configure SGI MPT and run SGI MPT programs, make sure to install SGI-Accelerate and SGI-MPI from the SGI Performance Suite.

13. On the **Initial Cluster Setup Tasks** screen, select **I Install Admin Cluster Software**, and select **OK**.

This step installs the cluster software that you wrote to the repositories.

14. On the **Initial Cluster Setup Tasks** screen, select **N Network Settings**, and select **OK**.

15. (Conditional) Create a routed management network.

Complete this step if you have at least 300–500 flat compute nodes in your cluster. If you have more than 500 flat compute nodes, consider creating more than one routed management network.

Complete the following steps:

- On the **Cluster Network Settings** screen, select **A Add Subnet**, and select **OK**
- On the **Select network type** screen, press the space bar to move the asterisk (*) up to the first line. This action selects the upper line, and the line now looks like this:

```
( *) 1 mgmt/mgmt-bmc
```

- Select **OK**.
- On the **Insert network name, subnet, netmask, bmc subnet and bmc network** screen, type in the information to define the routed management network. Use the arrow keys to move from field to field on this screen. The information you need to enter is as follows:

Field	Information
name	A unique name for this network. For example, head2.
subnet	The network IP address (start of the range) for the nodes on the routed management network.
netmask	The network mask for the nodes on routed management network.
bmc subnet	The network IP address (start of the range) for the node BMCs on the routed management network.
bmc netmask	The network mask for the node BMCs on the routed management network.

- On the **Network name ...** screen, verify that this is the information you specified for the routed management network, and select **OK**.
- On the **Network name-bmc ...** screen, verify that this is the information you specified for the node BMC network, and select **OK**.

16. On the **Cluster Network Settings** screen, select **S List and Adjust Subnet Addresses**, and select **OK**.
17. On the **Warning: Changing the subnet IP addresses ...** screen, click **OK**.
18. Review the settings on the **Subnet Network Addresses** screen, and modify these settings only if absolutely necessary.

Figure 4-7 on page 87 shows the **Subnet Network Addresses** screen. This screen displays the default networks and netmasks that reside within the cluster.

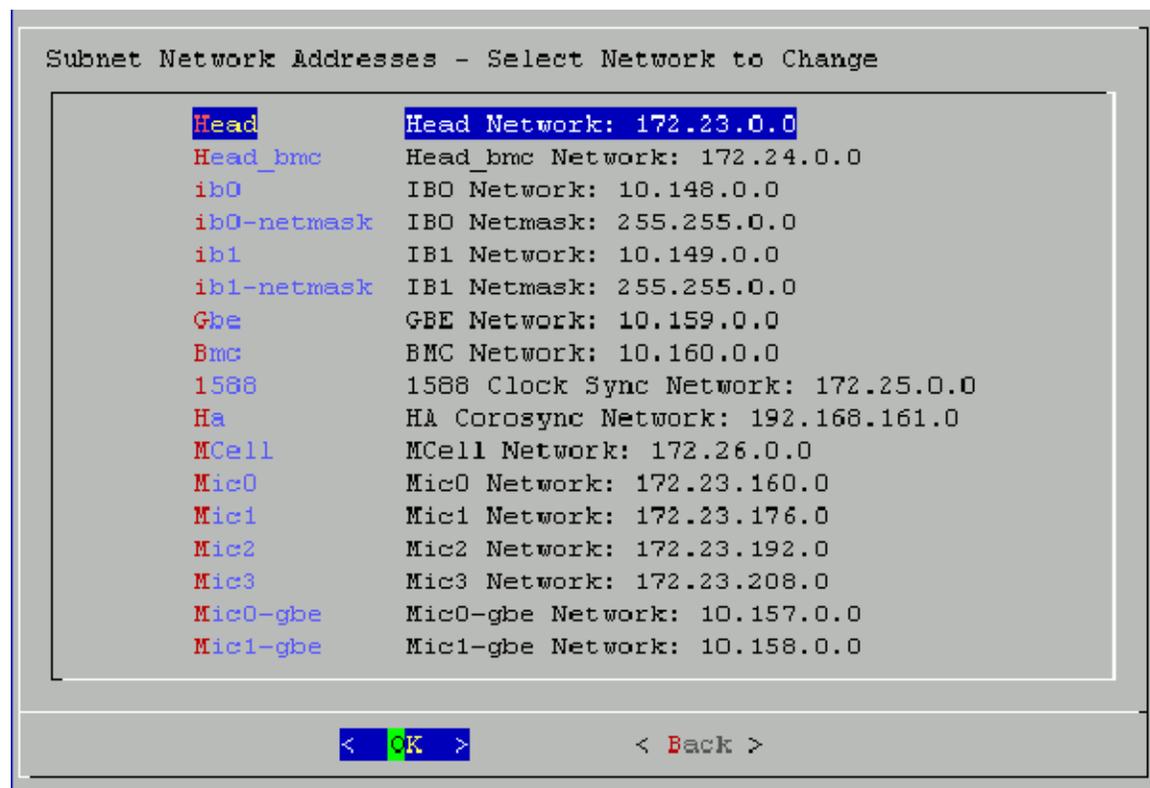


Figure 4-7 Subnet Network Addresses Screen

If you accept the defaults, select **OK**.

If you do not accept the defaults, you can change the network settings. For example, it is possible that your site has existing networks or conflicting network requirements. For additional information about the IP address ranges, see Appendix B, "Subnetwork Information" on page 203. Complete the following steps if you need to change the network settings:

- Highlight the setting you want to change, and select **OK**.
- Type in a new IP address, and select **OK**.
- Press Enter.

On the **Update Subnet Addresses** screen, the **Head Network** field shows the admin node's IP address. SGI recommends that you do not change the IP address of the admin node or rack leader controllers (RLCs) if at all possible. You can change the IP addresses of the InfiniBand network (**IB0** and **IB1**) to match the IP requirements of the house network, and then select **OK**.

19. On the **Cluster Network Settings** screen, select **D Configure Cluster Domain Name**, and select **OK**.
20. On the **Please enter the domain name for this cluster** pop-up window, type the domain name, and select **OK**.

The domain you type becomes a subdomain to your house network.

For example, type `ice.americas.sgi.com`.

21. On the **Cluster Network Settings** screen, select **Back**.
22. On the **Initial Cluster Setup** screen, select **T Configure Time Client/Server (NTP)**, and select **OK**.
23. Configure your NTP server.

On the subsequent screens, you set the admin node as the time server to the cluster. For this step, the installer screens differ on RHEL platforms and SLES platforms.

On RHEL platforms, complete the following step:

- On the **A new ntp.conf has been put in to position ...** screen, select **OK**.

On SLES platforms, complete the following steps:

- On the **A new ntp.conf has been put in to position ...** screen, select **OK**.

- Use the YAST interface and the SLES documentation to guide you through the NTP configuration.
 - On the **This procedure will replace your ntp configuration file ...** screen, select **Yes**.
24. On the **Initial Cluster Setup Tasks** menu, select **S Perform Initial Admin Node Infrastructure Setup**, and select **OK**.
 25. On the **A script will now perform the initial cluster ...** screen, select **OK**.

This step runs a series of scripts that configure the admin node. The scripts also create the root images for the RLCs, SGI ICE compute nodes, and flat compute nodes. The scripts run for approximately 30 minutes. At the end, the script issues a line that includes **install-cluster completed** in its output.

The final output of the script is as follows:

```
/opt/sgi/sbin/create-default-sgi-images Done!
```

The output of the `mksiimage` commands are stored in a log file at the following location:

```
/var/log/cinstallman
```

26. On the **Initial Cluster Setup Complete** window, select **OK**.
27. On the **One or more ISOs were embedded on the admin install DVD and copied to ...**, screen, select **OK**.

Depending on what you have installed, this screen might not appear.

28. On the **Initial Cluster Setup** menu, select **D Configure House DNS Resolvers**, and select **OK**.

Figure 4-8 on page 90 shows the **Configure House DNS Resolvers** screen.

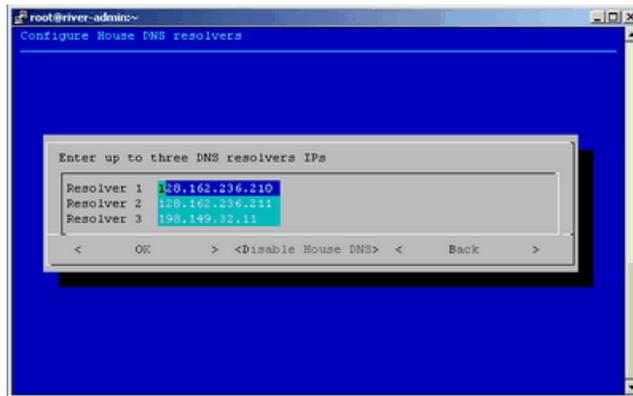


Figure 4-8 Configure House DNS Resolvers Screen

The system autopopulates the values on the **Configure House DNS Resolvers** screen to match the DNS specifications on the admin node. The DNS resolvers you specify here enable the flat compute nodes to resolve host names on your network. You can set the DNS resolvers to the same name servers used on the admin node itself.

Perform one of the following actions:

- To accept these settings, select **OK**, and then select **Yes**.
- To change the settings, type in different IP addresses, select **OK**, and then select **Yes**.
- To disable house network resolvers, select **Disable House DNS**.

On the **Setting DNS Forwarders to ...** screen, select **Yes**.

29. On the **Initial Cluster Setup** screen, select **Back**.

This action returns you to the cluster configuration tool main menu.

30. On the **Main Menu**, select **S Configure Switch Management Network (optional)**, and select **OK**.

The switch management network enables the Ethernet switch to control all VLANs and trunking.

31. On the pop-up window that appears, make sure that **Y yes** is selected, and select **OK**.

Figure 4-9 on page 91 shows the selection pop-up window:

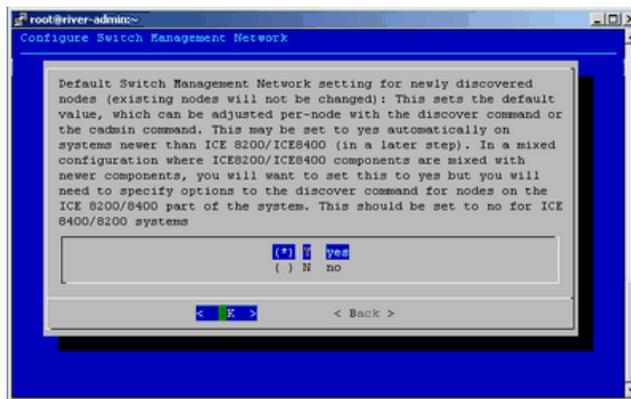


Figure 4-9 Configure Switch Management Network screen

32. (Conditional) On the **Main Menu**, select **N Configure MCell Network (optional)**, and select **OK**.

Complete this step if your SGI ICE system contains MCells.

33. (Conditional) On the screen that appears, select **Y yes**, and select **OK**.

Complete this step if your SGI ICE system contains MCells.

34. Select **Quit**.

35. Type the `cattr list -g` command to verify the features you configured with the cluster configuration tool.

Example: The following output is generated on an SGI ICE cluster with MCells. If your system does not include MCells, the `mcell_network` value should display `no`. The output is as follows:

```
# cattr list -g
global
cluster_domain      : smc.americas.sgi.com
tempo_dhcp_option   : 149
```

```
head_vlan           : 1
mcell_vlan          : 3
rack_vlan_start     : 101
rack_vlan_end       : 1100
mgmt_vlan_start     : 2001
mgmt_vlan_end       : 2500
redundant_mgmt_network : yes
switch_mgmt_network : yes
mcell_network       : yes
discover_skip_switchconfig : no
max_rack_irus       : 4
mic                 : 0
blademond_scan_interval : 120
dhcp_bootfile       : grub2
udpcast_min_receivers : 1
udpcast_min_wait    : 10
udpcast_max_wait    : 10
udpcast_max_bitrate : 900m
udpcast_rexmit_hello_interval : 0
udpcast_mcast_rdv_addr : 224.0.0.1
my_sql_replication  : yes
conserver_logging   : yes
conserver_ondemand  : no
edns_udp_size       : 512
replication_file     : mysql-bin.000005
replication_position : 9103
```

Note: On an SGI Rackable cluster, the `cattr` output is similar to the preceding example output, but the output contains fewer fields.

If you need to respecify any global values, start the cluster configuration tool again, and correct your specifications. To start the cluster configuration tool, type the following command:

```
# /opt/sgi/sbin/configure-cluster
```

36. Proceed to one of the following:

- To configure one or more external Domain Name Service (DNS) servers, proceed to "(Conditional) Configuring External Domain Name Service (DNS) Servers" on page 93.

- To synchronize the software repository, install updates, and clone the images, proceed to "Synchronizing the Software Repository, Installing Software Updates, Downloading Debugging Packages, and Cloning the Images" on page 94.

(Conditional) Configuring External Domain Name Service (DNS) Servers

Perform the procedure in this section if you want to enable network address translation (NAT) gateways for the cluster. A later procedure explains how to configure NAT as a service on a flat compute node. If you want to enable NAT, perform the procedure in this topic at this time.

When external DNS and NAT are enabled, the host names for the SGI ICE compute nodes (blades) in the cluster resolve through external DNS servers. The SGI ICE compute nodes need to be able to reach your house network.

Note: You cannot configure this feature after you run the `discover` command. If you attempt to configure this feature after you run the `discover` command, the IP addresses assigned previously on the configured nodes remain.

The following procedure explains how to configure external DNS servers.

Procedure 4-12 To configure external DNS servers

1. Obtain a large block of IP addresses from your network administrator.

This feature requires you to reserve a block of IP addresses on your house network. If you want to use external DNS servers, all nodes on the InfiniBand networks, both the `ib0` and `ib1` networks are included. The external DNS is enabled to provide addresses for all rack leader controllers (RLCs), all SGI ICE compute nodes, and all flat compute nodes.

2. Through an `ssh` connection, log into the admin node as the root user.
3. Type the following command to start the cluster configuration tool:

```
# /opt/sgi/sbin/configure-cluster
```

4. Select **E Configure External DNS Masters (optional)**, and select **OK**.
5. On the **This option configures SGI Tempo to look up the IP addresses for the InfiniBand networks from external DNS servers ...** screen, select **Yes**.

6. On the **Enter up to five external DNS master IPs** screen, type the IP addresses of up to five external DNS servers on your house network, and select **OK**.
7. On the **Setting external DNS masters to *ip_addr***, select **Yes**.
8. Proceed to "Synchronizing the Software Repository, Installing Software Updates, Downloading Debugging Packages, and Cloning the Images" on page 94.

Synchronizing the Software Repository, Installing Software Updates, Downloading Debugging Packages, and Cloning the Images

The following procedure explains how to update the software in the repositories that you created with the cluster configuration tool. The following procedure assumes that the cluster has a connection to the internet. If you need to perform this procedure on a secure cluster, you need to modify this procedure. For a secure system, obtain the software updates from SGI Supportfolio manually and use the `crepo` command to install the software manually.

Procedure 4-13 To update the software

1. Through an `ssh` connection, log into the admin node as the root user.
2. Type the following command to retrieve information about the network interface card (NIC) bonding method on the admin node:

```
# cadmin --show-mgmt-bonding --node admin
```

If bonding has been set appropriately, the command returns `802.3ad`.

If the command does not return `802.3ad`, type the following commands to set the bonding appropriately and reboot the system:

```
# cadmin --set-mgmt-bonding --node admin 802.3ad
# reboot
```

3. Type the following command to retrieve the new images from SGI SupportFolio and the operating system vendor:

```
# sync-repo-updates
```

For RHEL-based systems, make sure the system is subscribed as `rhel-x86_64--server-6`.

This step requires that the system be connected to the internet. Contact your SGI representative if this update method is not acceptable for your site.

4. (Optional) Download debugging packages and install them into the images.

Both RHEL and SLES offer debugging packages that you can download. For example, you need these additional packages if you want to run the `crash(8)` utility. Use the following information to download this additional software:

- On RHEL platforms, see the following website:

<https://access.redhat.com/solutions/9907>

- On SLES 12 platforms, see your SGI representative for assistance.

- On SLES 11 platforms, see the following website:

<https://www.novell.com/support/kb/doc.php?id=3074997>

5. Type the `cinstallman --show-images` command to retrieve the image names.

For example:

```
# cinstallman --show-images
Image Name                               BT VCS Compat_Distro
rhel7.1                                   0  1  rhel7
    3.10.0-229.el7.x86_64
ice-rhel7.1                               0  1  rhel7
    3.10.0-229.el7.x86_64
lead-rhel7.1                              0  1  rhel7
    3.10.0-229.el7.x86_64
```

The preceding output includes a line for the MPSS images if you downloaded MPSS software for Intel Many Integrated Core Architecture (Intel MIC Architecture) based products and created an image.

6. (Optional) Clone the images.

Complete this step if you want to back up the current images before they are installed.

Type the following command:

```
cinstallman --create-image --clone --source src_image_name --image image
```

For *src_image_name*, specify the name of the source image. For example:
lead-rhel7.1.

For *image*, specify a file name for the copied file (the clone). For example:
lead-rhel7.1.backup

Example 1: The following commands create backup copies of the current installation images:

```
# cinstallman --create-image --clone --source ice-compute-rhel7.1 --image ice-compute-rhel7.1.backup
# cinstallman --create-image --clone --source rhel7.1 --image rhel7.1.backup
# cinstallman --create-image --clone --source lead-rhel7.1 --image lead-rhel7.1.backup
```

Example 2: In this example, the commands create backup copies of the current installation images and tag the backup copies as source-controlled copies. The commands assume that there are multiple versions of the source image that exist at this time. The commands copy revision 2 of the source image to the backup.

```
# cinstallman --create-image --clone --source ice-compute-rhel7.1 --rev 2 --image ice-compute-rhel7.1.backup
# cinstallman --create-image --clone --source rhel7.1 --rev 2 --image rhel7.1.backup
# cinstallman --create-image --clone --source lead-rhel7.1 --rev 2 --image lead-rhel7.1.backup
```

7. Type a series of `cinstallman --update-image` commands to install the software images on the nodes.

For each *image*, specify the software package you want to install on each type of node.

For example, to install the packages shown in Procedure 4-13, step 5 on page 95, type the following commands:

```
# cinstallman --update-image --image ice-compute-rhel7.1
# cinstallman --update-image --image rhel7.1
# cinstallman --update-image --image lead-rhel7.1
```

8. Proceed to one of the following:
 - If your cluster contains Intel® Many Integrated Core Architecture (Intel MIC Architecture) devices, proceed to "(Conditional) Downloading the Intel Manycore Platform Software Stack (MPSS) Software and Creating Images" on page 97.
 - If your cluster does not contain MIC devices, proceed to "Configuring the Switches" on page 109.

(Conditional) Downloading the Intel Manycore Platform Software Stack (MPSS) Software and Creating Images

Perform the procedures in this topic if nodes in your cluster are equipped with Intel Many Integrated Core Architecture (Intel MIC Architecture) based products. The Intel Many Integrated Core (MIC) devices are part of the Intel Manycore Platform Software Stack that runs on the Intel Xeon Phi Coprocessors found on SGI ICE compute nodes and flat compute nodes.

Intel Corporation provides software for its Intel MIC architecture products, and you need to download this software for use on your SGI cluster. The MPSS package that you download contains the software packages for the MIC devices on the SGI cluster nodes. The procedures in this topic explain how to download the RPMs from the Intel Corporation website and how to create images for the nodes that are equipped with MIC devices.

Your system might have MIC devices on SGI ICE compute nodes, on the flat compute nodes, or both. If your cluster is equipped with MIC devices, each compute blade includes one or two IP addresses for each device. Only one cable connects each compute blade to the network, but each MIC device requires its own, unique IP address on your network.

Complete the procedures that are appropriate for your hardware configuration.

The following procedures explain how to obtain and deploy the MPSS software from Intel Corporation:

- "Downloading the MPSS Software From the Intel Corporation" on page 97
- "Creating Images for the SGI ICE Compute Nodes That Include MIC Devices" on page 99
- "Creating Images for the Flat Compute Nodes That Include MIC Devices" on page 103

Downloading the MPSS Software From the Intel Corporation

Complete the procedure in this topic if you have any MIC devices on your cluster.

The following procedure explains how to download the MPSS software from Intel Corporation.

Procedure 4-14 To download the MPSS package

1. Open a browser, and navigate to the following website:

<http://software.intel.com/mic-developer>

2. Click the **Tools & Downloads** tab.
3. Click the **Software Drivers: Intel Manycore Platform Software Stack (Intel MPSS)** link.

Follow the instructions on the website to download the Linux software version for your operating system platform. The download comes in the form of a `tar(1)` file.

4. Use the instructions from Intel to build the RPM files that you need.

A later procedure explains how to transfer these files to the cluster and build new images.

5. Plan how to specify the number MIC devices per node to SMC.

The next major task in the installation is creating node images that contain the MPSS software. However, at this time, SGI recommends that you plan how to specify the MIC devices to SMC. The procedure called "Configuring the Switches" on page 109 is the actual procedure that includes specifying the cluster configuration.

The `mic=number` parameter specifies the quantity of MIC devices on a node. The minimum number of devices is 1. The maximum number of devices is 16. For example, if the cluster contains four MIC devices per node, specify `mic=4` in the parameter string that defines the node.

You can specify this parameter in the cluster definition file (recommended) or on the `discover` command line, as follows:

- The cluster definition file defines characteristics for each cluster node. If you use a cluster definition file, specify a `mic=number` parameter in the list of parameters for each node.
- The `discover` command accepts node characteristics. If you do not have a cluster definition file, plan to specify the `mic=number` parameter on the `discover` command line.

Make sure to specify the correct number of MIC devices. If your cluster definition file contains `mic=0`, or if you do not specify a `mic=` parameter on the `discover` command, SMC assumes that the cluster contains no MIC devices.

6. Proceed to one of the following topics:

- "Creating Images for the SGI ICE Compute Nodes That Include MIC Devices" on page 99
- "Creating Images for the Flat Compute Nodes That Include MIC Devices" on page 103

Plan to perform both of the preceding procedures if you have MIC devices on both SGI ICE compute nodes and on flat compute nodes.

Creating Images for the SGI ICE Compute Nodes That Include MIC Devices

The following procedure explains how to create SGI ICE compute node images that include MIC device software.

Procedure 4-15 To create compute node images for SGI ICE compute nodes with MIC devices

1. On the admin node, use the `mkdir(1)` command, in the following format, to create a directory for the RPM repository:

```
mkdir -p /tftpboot/intel/mpss_repository_directory
```

For *mpss_repository_directory*, type a name for the directory that is to contain the MPSS repository. For convenience, make sure to include an identifier for the MPSS release level you downloaded.

For example:

```
# mkdir -p /tftpboot/intel/mpss_u3-2.1.6720-19
```

2. Use operating system commands to copy the RPM files you downloaded to the `/tftpboot/intel/mpss_repository_directory` directory on the admin node.

For example, use `cp(1)`, `ftp(1)`, `rsync(1)`, `scp(1)`, or another method.

3. (Conditional) Rebuild the MPSS modules.

Complete this step

4. Use the `crepo` command, in the following format, to specify a custom repository for the MPSS RPMs:

```
crepo --add rpm_repo_directory -custom rpm_repo_name
```

The variables in this command are as follows:

- For *rpm_repo_directory*, specify the full path to the directory that contains the RPM files.
- For *rpm_repo_name*, create a name for the image. You can specify the same name for both *rpm_repo_directory* and *rpm_repo_name*. After SMC builds the image, the `cinstallman --show-images` command returns this label in the **Image Name** column of its output.

For example:

```
# crepo --add /tftpboot/intel/mpss_u3-2.1.6720-19 --custom mpss_u3-2.1.6720-19
```

5. Type the following command to confirm that the MPSS image is in the correct repository:

```
# crepo --show
```

6. Use the `crepo` command, in the following format, to add the custom repository to the generated RPM list in `/etc/opt/sgi/rpmlists`:

```
crepo --select rpm_repo_name
```

For *rpm_repo_name*, create a name for the image. You can specify the same name for both *rpm_repo_directory* and *rpm_repo_name*. After SMC builds the image, the `cinstallman --show-images` command returns this label in the **Image Name** column of its output. This is the same *rpm_repo_name* that you specified in the following step:

Procedure 4-15, step 4 on page 99

For example, the following command adds the custom repository and displays the content of the repository:

```
# crepo --select mpss_u3-2.1.6720-19
Selecting: mpss_u3-2.1.6720-19
Updating: /etc/opt/sgi/rpmlists/generated-ice-rhel7.1.rpmlist
Updating: /etc/opt/sgi/rpmlists/generated-lead-rhel7.1.rpmlist
Updating: /etc/opt/sgi/rpmlists/generated-rhel7.1.rpmlist
```

7. Type the following command to confirm that you selected the new repository that contains the MPSS RPMs:

```
# crepo --show
* mpss_u3-2.1.6720-19 : /tftpboot/intel/mpss_u3-2.1.6720-19
* Red-Hat-Enterprise-Linux-7.1 : /tftpboot/distro/rhel7.1
```

- * SGI-MPI-1.10-rhel7 : /tftpboot/sgi/SGI-MPI-1.10-rhel7
- * SGI-Management-Center-3.1-rhel7 : /tftpboot/sgi/SGI-Management-Center-3.0-rhel7
- * SGI-Foundation-Software-2.12-rhel7 : /tftpboot/sgi/SGI-Foundation-Software-2.11-rhel7

The asterisk character (*) in column 1 indicates an image that is selected.

8. Type the following command to display the images that are available for installation on the compute nodes:

```
# cinstallman --show-images
Image Name                               BT VCS Compat_Distro
rhel7.1                                   0  1  rhel7
    3.10.0-229.el7.x86_64
ice-rhel7.1                               0  1  rhel7
    3.10.0-229.el7.x86_64
lead-rhel7.1                              0  1  rhel7
    3.10.0-229.el7.x86_64
```

9. Use the `cinstallman` command, in the following format, to clone the current operating system image:

```
cinstallman --create-image --clone --source current_image --image new_image
```

The variables in this command are as follows:

- For *current_image*, type the name of the operating system image you want to use that is on the system right now. Choose one that appears in the output from the `cinstallman --show-images` command in the **Image Name** list. For example, `ice-rhel7.1`.
- For *new_image*, type a new name for the operating system image that is to include the MPSS file RPMs. SGI recommends that you include information in the new name that can enable you to identify this image as the operating system image that includes MPSS software. For example, `ice-compute-rhel7.1-mic-6720-19` identifies the new image as a RHEL image that contains a revision of the MPSS MIC software.

For example:

```
# cinstallman --create-image --clone --source ice-rhel7.1 --image ice-rhel7.1-mic-6720-19
```

10. Type the following command to display the images and confirm that the new image appears in the list:

```
# cinstallman --show-images
Image Name                               BT VCS Compat_Distro
lead-rhel7.1                              0  1  rhel7
    3.10.0-229.el7.x86_64
rhel7.1                                    0  1  rhel7
    3.10.0-229.el7.x86_64
ice-compute-rhel7.1                       1  1  rhel7
    3.10.0-229.el7.x86_64
ice-compute-rhel7.1-mic-6720-19          1  1  rhel7
    3.10.0-229.el7.x86_64
```

11. Use the `cinstallman` command, in the following format, to install the `sgi-mic-compute` package and the MPSS RPMs:

```
cinstallman --yum-image --image image_name install sgi-mic-compute
```

For *image_name*, specify the *new_image* that you created in the following step:

Procedure 4-15, step 9 on page 101

For example:

```
# cinstallman --yum-image --image ice-rhel7.1-mic-6720-19 install sgi-mic-compute
```

12. (Conditional) Enable the SLES operating system to load the MPSS package. Complete this step if you are installing the MPSS packages on a SLES platform. Complete the following steps:

- Open file `/etc/modprobe.d/unsupported-modules` with a text editor.
- Add the following line at the end of the file:

```
allow_unsupported_modules 1
```

- Save and close the file.

13. Type the following command to display all the images:

```
# cimage --list-images
image: ice-rhel7.1
```

```
kernel: 2.6.32-358.el6.x86_64
image: ice-rhel7.1-mic-6720-19
kernel: 2.6.32-358.el6.x86_64
```

The preceding output shows the newly installed image, `ice-rhel7.1-mic-6720-19`.

14. Use the `cimage` command, in the following format, to set the default image for the compute nodes:

```
cimage --set-default --file_system ice-rhel7.1-mic-6720-19 kernel
```

The variables in this command are as follows:

- For `file_system`, type either `nfs` or `tmpfs`, according to your site practice.
- For `kernel`, type the kernel associated with the new image, as shown in the output from the `cimage --list-images` command in the following step:

Procedure 4-15, step 13 on page 102

For example:

```
# cimage --set-default --tmpfs --ice-rhel7.1-mic-6720-19 2.6.32-358.el6.x86_64
```

15. Type the following command to confirm that the new image is the default image:

```
# cimage --show-default
ice-rhel7.1-stout7-mic-6720-15 2.6.32-358.el6.x86_64 tmpfs
```

16. Proceed to one of the following:

- If your cluster contains MIC devices on the flat compute nodes, proceed to "Creating Images for the Flat Compute Nodes That Include MIC Devices" on page 103.
- If your cluster does not contain MIC devices on the flat compute nodes, proceed to "Configuring the Switches" on page 109.

Creating Images for the Flat Compute Nodes That Include MIC Devices

The following procedure explains how to create images for flat compute nodes that include MIC device software.

Procedure 4-16 To create software images for flat compute nodes with MIC devices

1. On the admin node, use the `mkdir(1)` command, in the following format, to create a directory for the RPM repository:

```
mkdir -p /tftpboot/intel/mpss_repository_directory
```

For *mpss_repository_directory*, type a name for the directory that is to contain the MPSS repository. For convenience, make sure to include an identifier for the MPSS release level you downloaded.

For example:

```
# mkdir -p /tftpboot/intel/mpss_u3-2.1.6720-19
```

2. Use operating system commands to copy the RPM files you downloaded to the `/tftpboot/intel/mpss_repository_directory` directory on the admin node.

For example, use `cp(1)`, `ftp(1)`, `rsync(1)`, `scp(1)`, or another method.

3. (Conditional) Rebuild the MPSS modules.

Complete this step if you are installing a kernel that is updated from your operating system's base kernel.

The Intel Corporation distributes modules that support the RHEL and SLES base distributions. If you are installing an updated kernel, however, you need to rebuild the modules for the updated kernel.

Use an editor to open the following file, and follow the instructions from Intel Corporation regarding how to build the correct modules for your kernel:

```
/tftpboot/intel/mpss-version/mpss-version/docs/readme.txt
```

For *version*, specify the version that matches your MPSS level.

4. Use the `crepo` command, in the following format, to specify a custom repository for the MPSS RPMs:

```
crepo --add rpm_repo_directory -custom rpm_repo_name
```

The variables in this command are as follows:

- For *rpm_repo_directory*, specify the full path to the directory that contains the RPM files.

- For *rpm_repo_name*, create a name for the image. You can specify the same name for both *rpm_repo_directory* and *rpm_repo_name*.

After SMC builds the image, the `cinstallman --show-images` command returns this label in the **Image Name** column of its output.

For example:

```
# crepo --add /tftpboot/intel/mpss_u3-2.1.6720-19 --custom mpss_u3-2.1.6720-19
```

5. Type the following command to confirm that the MPSS image is in the correct repository:

```
# crepo --show
```

6. Use the `crepo` command, in the following format, to add the custom repository to the generated RPM list in `/etc/opt/sgi/rpmlists`:

```
crepo --select rpm_repo_name
```

For *rpm_repo_name*, create a name for the image. You can specify the same name for both *rpm_repo_directory* and *rpm_repo_name*. After the image is built, the `cinstallman --show-images` command returns this label in the **Image Name** column of its output. This is the same *rpm_repo_name* that you specified in the following step:

Procedure 4-16, step 4 on page 104

For example, the following command adds the custom repository and displays the content of the repository:

```
# crepo --select mpss_u3-2.1.6720-19
Selecting: mpss_u3-2.1.6720-19
Updating: /etc/opt/sgi/rpmlists/generated-ice-rhel7.1.rpmlist
Updating: /etc/opt/sgi/rpmlists/generated-lead-rhel7.1.rpmlist
Updating: /etc/opt/sgi/rpmlists/generated-rhel7.1.rpmlist
```

7. Type the following command to confirm that you selected the new repository that contains the MPSS RPMs:

```
# crepo --show
* mpss_u3-2.1.6720-19 : /tftpboot/intel/mpss_u3-2.1.6720-19
* Red-Hat-Enterprise-Linux-7.1 : /tftpboot/distro/rhel7
* SGI-MPI-1.10-rhel7 : /tftpboot/sgi/SGI-MPI-1.10-rhel7
* SGI-Management-Center-3.1-rhel7 : /tftpboot/sgi/SGI-Management-Center-3.1-rhel7
```

```
* SGI-Foundation-Software-2.12-rhel7 : /tftpboot/sgi/SGI-Foundation-Software-2.12-rhel7
```

The asterisk character (*) in column 1 indicates an image that is selected.

8. Type the following command to display the images that are available for installation on the flat compute nodes:

```
# cinstallman --show-images
Image Name                               BT VCS Compat_Distro
rhel7.1                                   0  1  rhel7
    3.10.0-229.el7.x86_64
ice-rhel7.1                               0  1  rhel7
    3.10.0-229.el7.x86_64
lead-rhel7.1                              0  1  rhel7
    3.10.0-229.el7.x86_64
```

9. Use the `cinstallman` command, in the following format, to clone the current operating system image:

```
cinstallman --create-image --clone --source current_image --image new_image
```

The variables in this command are as follows:

- For *current_image*, type the name of the operating system image you want to use. Choose one that appears in the output from the `cinstallman --show-images` command in the **Image Name** list. For example, `rhel7.1`.
- For *new_image*, type a new name for the operating system image that is to include the MPSS file RPMs. SGI recommends that you include information in the new name that can enable you to identify this image as the operating system image that includes MPSS software. For example, `rhel7.1-mic-6720-19` identifies the new image as a RHEL image that contains a revision of the MPSS MIC software.

For example:

```
# cinstallman --create-image --clone --source rhel7.1 --image rhel7.1-mic-6720-19
```

10. Type the following command to display the images and confirm that the new image appears in the list:

```
# cinstallman --show-images
Image Name                               BT VCS Compat_Distro
lead-rhel7.1                             0  1  rhel7.1
    2.6.32-504.el6.x86_64
```

```

rhel7.1          0 1  rhel7.1
                2.6.32-504.el6.x86_64
ice-rhel7.1     1 1  rhel7.1
                2.6.32-504.el6.x86_64
rhel7.1-mic-6720-19 1 1  rhel7.1
                2.6.32-504.el6.x86_64
    
```

11. Use the `cinstallman` command, in the following format, to install the `sgi-mic-service` package and the MPSS RPMs:

```
cinstallman --yum-image --image image_name install sgi-mic-service
```

For *image_name*, specify the *new_image* that you created in the following step:

Procedure 4-16, step 9 on page 106

For example:

```
# cinstalexallman --yum-image --image rhel7.1-mic-6720-19 install sgi-mic-service
```

12. (Conditional) Enable the SLES operating system to load the MPSS package.

Complete this step if you are installing the MPSS packages on a SLES platform.

Complete the following steps:

- Open file `/etc/modprobe.d/unsupported-modules` with a text editor.
- Add the following line at the end of the file:

```
allow_unsupported_modules 1
```

- Save and close the file.

13. Type the following command to display all the images:

```

# cimage --list-images
image: ice-rhel7.1
      kernel: 2.6.32-358.el6.x86_64
image: ice-rhel7.1-mic-6720-19
      kernel: 2.6.32-358.el6.x86_64
image: rhel7.1-mic-6720-19
      kernel: 2.6.32-358.el6.x86_64
    
```

The preceding output shows the newly installed image, `rhel7.1-mic-6720-19`.

14. Use the `cimage` command, in the following format, to set the default image for the flat compute nodes:

```
cimage --set-default --file_system rhel7.1-mic-6720-19 kernel
```

The variables in this command are as follows:

- For *file_system*, type either `nfs` or `tmpfs`, according to your site practice.
- For *kernel*, type the kernel associated with the new image, as shown in the output from the `cimage --list-images` command in the following step:

Procedure 4-16, step 13 on page 107

For example:

```
# cimage --set-default --tmpfs --rhel7.1-mic-6720-19 2.6.32-358.el6.x86_64
```

15. Type the following command to confirm that the new image is the default image:

```
# cimage --show-default  
rhel7.1-stout7-mic-6720-15 2.6.32-358.el6.x86_64 tmpfs
```

16. Proceed to one of the following:

- If you are installing SMC as part of a migration upgrade from SMC 1.7.x, proceed to the following:
"(Conditional) Running the Migration Script and Editing the Cluster Definition File" on page 108
- If you are installing SMC on a new cluster or as part of a refresh, proceed to the following:
"Configuring the Switches" on page 109

(Conditional) Running the Migration Script and Editing the Cluster Definition File

Perform the procedure in this topic only if you are installing the SMC software as part of a migration upgrade from SMC 1.7.x.

You need to use separate `migrate-sgi-mc` commands to convert each SMC 1.7.x payload file.

The following procedure explains how to run the migration script.

Procedure 4-17 To run the migration script

1. On the admin node, type the following command:

```
# migrate-sgi-mc
```

When the script runs, it prompts you for the image location and performs the following tasks:

- Copies the payloads from the specified location into SMC 3.x environment.
 - Clean up the image to remove unnecessary packages.
 - Installs packages required by SMC 3.x.
 - Registers the new image with the SMC database.
 - Generates a configuration template, `smc.txt`, for use with the `discover` command.
2. Repeat the preceding step for each payload file that you want to convert.
 3. Open file `smc.txt` from within a text editor.

File `smc.txt` is a cluster definition file that you can use as input to the `discover` command. Read through the instructions in "Configuring the Switches" on page 109 and familiarize yourself with the content of this generated file. Make sure to edit the file as necessary.
 4. Proceed to the following:

"Configuring the Switches" on page 109

Configuring the Switches

SGI clusters have both management switches and InfiniBand (IB) switches. The individual switches are paired into *switch stacks*, and there are two switches per stack. In each stack, the top switch is typically the master switch, and the bottom switch is typically the slave switch. Although a switch stack actually includes two switches, most documentation refers to a switch stack as a *switch*.

An SGI cluster is equipped with the following types of switches:

- Spine switches. A spine switch is the primary management switch or the primary IB switch. There is one primary management switch and one primary IB switch.
- Leaf switches. A leaf switch is a secondary management switch or a secondary IB switch. There can be many leaf switches configured as part of a cluster system.

The `discover` command initializes and configures the system components for the cluster. The switch configuration procedures explain how to use the `discover` command to configure the cluster's management switches first. After you configure the management switches, if you have MCells, you configure the cooling equipment on the MCell network's switch ports.

The cluster configuration requires that the same IP address be assigned to the cluster's head gateway and to the first management switch, usually `mgmtswitch0`. The procedures in this topic assume that you want to use the default IP address, which is `172.23.255.254`, for both components, but the procedures include example commands that show how to configure an alternate IP address.

Proceed to the following to familiarize yourself with the cluster definition file's purpose:

"About the Cluster Definition File" on page 110

About the Cluster Definition File

A cluster definition file specifies media access control (MAC) addresses, IP addresses, node roles, hostnames, and other information for the cluster components. You can specify the cluster definition file as input to the `discover` command and to the `configure-cluster` command. Cluster configuration can proceed much more quickly if you have a cluster definition file. Without the file, you need to power cycle each component manually. For new clusters, you can obtain a cluster definition file from your SGI representative. For clusters that are configured with at least one working slot, you can type the following command to generate a cluster configuration file:

```
discover --show-configfile > file_name
```

You can write the cluster definition file to any *file_name* and store the cluster definition file in any directory.

Example 1. This example cluster definition file is for an SGI ICE cluster that includes one SGI ICE compute rack and several flat compute nodes. The following information highlights some characteristics of this cluster:

- The **temponame** field and the **hostname1** field appear in bold print in this example. The **temponame** field can contain the hostname of the node or it can contain a label for the role of the component in the cluster. The **temponame** field is used by SMC internal operations and for flat compute nodes is always **service** followed by a number. The **hostname1** field defines the hostname for the component, and it is this hostname that users need to specify when they want to log into the node. For example, if you configure user services on the two of the flat compute nodes in this cluster, users can log into the cluster by logging into **n0** or **n1**.
- The flat compute nodes can have any name. This example file uses the default names. This example shows two ranges. One range includes only **n0** and **n1**. The other range starts at **n101** and includes **n102**, **n103**, and so on.

Note: Not all flat nodes are shown in the example file.

The file is as follows:

```
[discover]
temponame=r1lead, mgmt_bmc_net_name=head-bmc, mgmt_bmc_net_macs=00:25:90:58:8b:75,
mgmt_net_name=head, mgmt_net_macs=00:25:90:58:8a:94/00:25:90:58:8a:95,
redundant_mgmt_network=yes, switch_mgmt_network=yes, mic=0, dhcp_bootfile=grub2,
conserver_logging=yes, conserver_ondemand=no, console_device=ttyS1
temponame=service0, mgmt_bmc_net_name=head-bmc, mgmt_bmc_net_macs=00:25:90:58:7d:7f,
mgmt_net_name=head, mgmt_net_macs=00:25:90:58:7d:32/00:25:90:58:7d:33, hostname1=n0,
redundant_mgmt_network=yes, switch_mgmt_network=yes, mic=0, dhcp_bootfile=grub2,
conserver_logging=yes, conserver_ondemand=no, root_type=disk, console_device=ttyS1
temponame=service1, mgmt_bmc_net_name=head-bmc, mgmt_bmc_net_macs=00:25:90:58:96:a2,
mgmt_net_name=head, mgmt_net_macs=00:25:90:58:96:54/00:25:90:58:96:55, hostname1=n1
redundant_mgmt_network=yes, switch_mgmt_network=yes, mic=0, dhcp_bootfile=grub2,
conserver_logging=yes, conserver_ondemand=no, root_type=disk, console_device=ttyS1
temponame=service101, mgmt_bmc_net_name=head-bmc, mgmt_bmc_net_macs=00:1E:67:2C:53:92,
mgmt_net_name=head, mgmt_net_macs=00:1E:67:2C:53:8E/00:1e:67:2c:53:8f, hostname1=n101,
redundant_mgmt_network=yes, switch_mgmt_network=yes, mic=0, dhcp_bootfile=ipxe,
conserver_logging=yes, conserver_ondemand=no, root_type=disk, console_device=ttyS0
temponame=service102, mgmt_bmc_net_name=head-bmc, mgmt_bmc_net_macs=00:1E:67:2C:58:AF,
mgmt_net_name=head, mgmt_net_macs=00:1E:67:2C:58:AB/00:1e:67:2c:58:ac, hostname1=n102,
redundant_mgmt_network=yes, switch_mgmt_network=yes, mic=0, dhcp_bootfile=ipxe,
conserver_logging=yes, conserver_ondemand=no, root_type=disk, console_device=ttyS0
temponame=service103, mgmt_bmc_net_name=head-bmc, mgmt_bmc_net_macs=00:1E:67:2C:54:E2,
mgmt_net_name=head, mgmt_net_macs=00:1E:67:2C:54:DE/00:1e:67:2c:54:df, hostname1=n103,
```

4: Installing and Configuring an SGI Cluster System

```
redundant_mgmt_network=yes, switch_mgmt_network=yes, mic=0, dhcp_bootfile=ipxe,  
conserver_logging=yes, conserver_ondemand=no, root_type=disk, console_device=ttyS0
```

```
.  
.
.
```

```
temponame=mgmtsw0, mgmt_net_name=head, mgmt_net_macsb4:0e:dc:38:6b:17, net=head/head-  
bmc, ice=yes, type=spine
```

```
temponame=mgmtsw1, mgmt_net_name=head, mgmt_net_macsb4:0e:dc:38:6b:18, net=head/head-  
bmc, ice=no, type=leaf
```

```
[dns]
```

```
cluster_domain=acme.americas.sgi.com
```

```
nameserver1=137.38.225.5
```

```
nameserver2=137.38.31.248
```

```
[attributes]
```

```
dhcp_bootfile=grub2
```

```
udpcast_rexmit_hello_interval=0
```

```
udpcast_min_receivers=1
```

```
head_vlan=1
```

```
mcell_network=yes
```

```
udpcast_min_wait=10
```

```
my_sql_replication=yes
```

```
redundant_mgmt_network=yes
```

```
max_rack_irus=16
```

```
udpcast_max_bitrate=900m
```

```
udpcast_max_wait=10
```

```
udpcast_mcast_rdv_addr=224.0.0.1
```

```
rack_vlan_end=1100
```

```
switch_mgmt_network=yes
```

```
mcell_vlan=3
```

```
mic=0
```

```
conserver_logging=yes
```

```
rack_vlan_start=101
```

```
conserver_ondemand=no
```

```
blademonnd_scan_interval=120
```

```
[networks]
```

```
name=private, subnet=172.26.0.0, netmask=255.255.0.0
```

```
name=public, subnet=137.38.82.0, netmask=255.255.255.0, gateway=137.38.82.254
```

```
name=head, type=mgmt, vlan=1, subnet=172.23.0.0, netmask=255.255.0.0, gateway=172.23.255.254
```

```
name=head-bmc, type=mgmt-bmc, vlan=1, subnet=172.24.0.0, netmask=255.255.0.0
name=mcell-net, type=cooling, subnet=172.26.0.0, netmask=255.255.0.0
name=ha-net, type=ha, subnet=192.168.161.0, netmask=255.255.255.0
name=ib-0, type=ib, subnet=10.148.0.0, netmask=255.255.0.0
name=ib-1, type=ib, subnet=10.149.0.0, netmask=255.255.0.0
name=gbe, type=lead-mgmt, subnet=10.159.0.0, netmask=255.255.0.0, rack_netmask=255.255.252.0
name=bmc, type=lead-bmc, subnet=10.160.0.0, netmask=255.255.0.0, rack_netmask=255.255.252.0
gateway=172.99.255.254
```

Example 2. This example cluster definition file is for an SGI Rackable cluster with 100 flat compute nodes. For simplicity's sake, the example file shows only two flat compute services nodes and the management switches. The following information highlights some characteristics of this cluster:

- The information in the `temponame` field defines the role for each of the two flat compute nodes in this cluster. The content of the `temponame` field and the `hostname1` field can be identical; in other words, you can use the node's hostname as its `temponame`.

The content of the `temponame` field for each flat compute node is `servicen`, where n is a number from 1 through 101.

- The `hostname1` field defines the hostname that users need to specify when they want to log into a node. The text in the `hostname1` field is the text that appears in the output for most SMC commands when the command generates output.
- The cluster definition file specifies a multicast installation that uses `udpcast` transport for the flat compute nodes, `service1` and `service101`.
- The top-level switch, `mgmtsw0`, is defined as spine switch; this switch is always connected to the admin node. Switch `mgmtsw1` is defined as a leaf switch and is connected to the spine switch, `mgmtsw0`.
- The definition for both switches includes `ice=no` because this cluster has no SGI ICE components.

The file is as follows:

```
[discover]
temponame=service1, mgmt_bmc_net_name=head-bmc, mgmt_bmc_net_macs=00:25:90:1A:6D:3E,
mgmt_net_name=head, mgmt_net_macs=00:25:90:1A:AC:D0/00:25:90:1a:ac:d1, hostname1=mycluster1,
redundant_mgmt_network=yes, switch_mgmt_network=yes, mic=0, dhcp_bootfile=grub2,
conserver_logging=yes, conserver_ondemand=no, root_type=disk,
console_device=ttyS1,transport=udpcast
```

4: Installing and Configuring an SGI Cluster System

...

```
temponame=service101, mgmt_bmc_net_name=head-bmc, mgmt_bmc_net_macs=00:1E:67:2C:53:92,  
mgmt_net_name=head, mgmt_net_macs=00:1E:67:2C:53:8E, hostname1=n101,  
redundant_mgmt_network=yes, switch_mgmt_network=yes, mic=0, dhcp_bootfile=ipxe,  
conserver_logging=yes, conserver_ondemand=no, root_type=disk,  
console_device=ttyS0,transport=udpcast  
temponame=mgmtsw0, mgmt_net_name=head,mgmt_net_macs=00:26:F3:C3:7A:40, net=head/head-  
bmc, ice=no, type=spine  
temponame=mgmtsw1, mgmt_net_name=head,mgmt_net_macs=00:04:96:97:C0:78,  
net=head/head-bmc, ice=no, type=leaf
```

```
[dns]  
cluster_domain=smc-default.americas.sgi.com  
nameserver1=128.162.236.210  
nameserver2=128.162.236.211  
nameserver3=198.149.32.11
```

```
[attributes]  
dhcp_bootfile=grub2  
udpcast_min_receivers=1  
head_vlan=1  
mcell_network=yes  
udpcast_min_wait=10  
my_sql_replication=yes  
redundant_mgmt_network=yes  
max_rack_irus=16  
udpcast_max_bitrate=900m  
udpcast_max_wait=10  
rack_vlan_end=1100  
switch_mgmt_network=yes  
mcell_vlan=3  
mic=0  
conserver_logging=yes  
rack_vlan_start=101  
conserver_ondemand=no  
blademonnd_scan_interval=120
```

```
[networks]  
name=private, subnet=172.26.0.0, netmask=255.255.0.0
```

```

name=public, subnet=128.162.243.0, netmask=255.255.255.0, gateway=128.162.243.1
name=head, type=mgmt, vlan=1, subnet=172.23.0.0, netmask=255.255.0.0, gateway=172.23.255.254
name=head-bmc, type=mgmt-bmc, vlan=1, subnet=172.24.0.0, netmask=255.255.0.0
name=mcell-net, type=cooling, subnet=172.26.0.0, netmask=255.255.0.0
name=ha-net, type=ha, subnet=192.168.161.0, netmask=255.255.255.0
name=ib-0, type=ib, subnet=10.148.0.0, netmask=255.255.0.0
name=ib-1, type=ib, subnet=10.149.0.0, netmask=255.255.0.0
name=gbe, type=lead-mgmt, subnet=10.159.0.0, netmask=255.255.0.0, rack_netmask=255.255.252.0
name=bmc, type=lead-bmc, subnet=10.160.0.0, netmask=255.255.0.0, rack_netmask=255.255.252.0
    
```

The following list explains the procedures you need to follow to configure the switches:

Procedure	Circumstance
"Verifying the Switch Cabling" on page 115	All switch configuration circumstances. Perform this procedure regardless of the presence of a cluster definition file or MCells.
"Configuring Management Switches With a Cluster Definition File" on page 118	If you have a cluster definition file.
"Configuring Management Switches Without a Cluster Definition File" on page 122	If you do not have a cluster definition file.
"(Conditional) Configuring the Cooling Racks and Cooling Distribution Units (CDUs) on the MCell Network's Switch Ports" on page 126	(Conditional) If you have MCells. This extra procedure configures the MCell switches separately from the rest of the cluster switches.

Verifying the Switch Cabling

The following figures show example switch cabling. Depending on the switch configuration procedure you use, you might have to plug cables out from and into switch ports during the configuration process. Regardless of the procedure you use, the cables on your switches eventually become cabled as shown in the following figures.

Figure 4-10 on page 116 shows a switch stack with two switches. In this switch stack, the two switches constitute the spine switch stack. One is the master switch and the other is the slave switch.

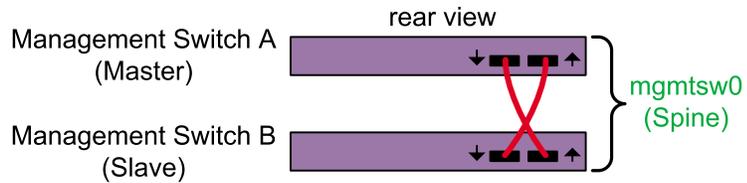


Figure 4-10 Spine Switch Stack With Two Switches

Figure 4-11 on page 116 shows a switch stack with multiple switches. The first two switches constitute the spine switch stack, and the other switches constitute the secondary switch stack.

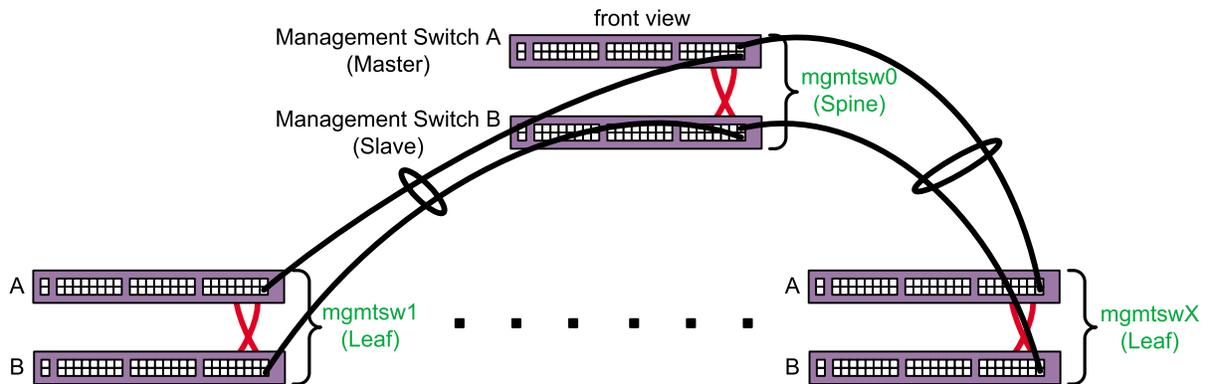


Figure 4-11 Switch Stack With Multiple Switches

The following procedure explains how to inspect your switches and prepare for the configuration procedure.

Procedure 4-18 To verify your switches

1. Visually inspect your system.

Note the types of switches you have and their identifiers. At a minimum, you have one spine switch stack. The admin node connects to the master switch in the spine switch stack. You might have additional leaf switch stacks.

Within each stack, each switch is labeled *MSWxx*. In the spine switch stack, the master switch is labeled *MSW0A*, and the slave switch is labeled *MSW0B*. In the first leaf switch stack, the master switch is labeled *MSW1A*, and the slave switch is labeled *MSW1B*. The *A* and *B* on the switch labels identify the master switch and slave switch in the stack. The `switchconfig set` command operates on a switch stack (not just one switch), so you need to note only the characters on the label that precede the *A* and *B* when you provide information to the `switchconfig set` command. Use the following table to determine the value you need to use for *switch* on the `switchconfig` commands:

<i>switch</i>	Label
<code>mgmtsw0</code>	<i>MSW0A</i> or <i>MSW0B</i>
<code>mgmtsw1</code>	<i>MSW1A</i> or <i>MSW1B</i>
<code>mgmtsw2</code>	<i>MSW2A</i> or <i>MSW2B</i>
<code>mgmtsw3</code>	<i>MSW3A</i> or <i>MSW3B</i>
<code>mgmtsw4</code>	<i>MSW4A</i> or <i>MSW4B</i>
<code>mgmtsw5</code>	<i>MSW5A</i> or <i>MSW5B</i>
<code>mgmtsw6</code>	<i>MSW6A</i> or <i>MSW6B</i>
<code>mgmtsw7</code>	<i>MSW7A</i> or <i>MSW7B</i>
<code>mgmtsw8</code>	<i>MSW8A</i> or <i>MSW8B</i>
<code>mgmtsw9</code>	<i>MSW9A</i> or <i>MSW9B</i>

2. Make sure that only the admin node is plugged in and that all rack circuit breakers are powered off.

If you have a cluster definition file, connect all other nodes and switches to a power source, but do not power them on. That is, make sure that all chassis management controllers (CMCs) on the rack leader controllers (RLCs), all RLCs, all SGI ICE compute nodes, all flat compute nodes, all switches, and so on, are not powered on.

If you do not have a cluster definition file, unplug all nodes and switches other than the admin node. The admin node should be the only component that is plugged in.

3. (Conditional) Remove the cascade cables that connect the slaves switches together.
Complete this step only if you have two or more switch stacks.

Locate the cascading cables that connect the switch stacks to each other, and unplug the cascading cable end from the lower switch in the neighboring switch stack. On the front of the switch stack, unplug the cascading cables that plug into the ports on the slave switches. When you unplug one end of each cable, you prevent a switching loop.

Note: Do not unplug the stacking cables in the rear of the switch stack. The installation procedure instructs you to plug or unplug only the cascading cables in the front of the switch stack.

For example, if you have one leaf switch stack, locate the cascading cable that runs from the slave switch on the spine switch to the slave switch on the leaf switch. Unplug that cable from the port on the slave switch.

4. Proceed as follows:
 - If you have a cluster definition file, proceed to the following:
"Configuring Management Switches With a Cluster Definition File" on page 118
 - If you do not have a cluster definition file, proceed to the following:
"Configuring Management Switches Without a Cluster Definition File" on page 122

Configuring Management Switches With a Cluster Definition File

The following procedure explains how to configure your switches when you have each switch's MAC information in a cluster definition file.

Procedure 4-19 To configure switches — with a cluster definition file

1. Through an `ssh` connection, log in as root to the admin node, and write the cluster definition file to a location on the admin node.

For example, write it to `/var/tmp/config_file`.

2. Plug in all the switches.
3. (Conditional) Specify a site-specific IP address for the head gateway.

Complete this step if you need to set a site-specific, nondefault IP address for the spine switch stack (head gateway).

By default, SGI sets the IP address for the spine switch stack (head gateway) to 172.23.0.254. Type the `cadmin` command in the following format to set a site-specific IP address:

```
cadmin -set-head-gateway IP_addr
```

For *IP_addr*, specify your site-specific IP address.

4. Use the `discover` command, in the following format, to configure the spine switch stack (the switches labeled MSW0A and MSW0B), which is attached to the admin node:

```
discover --mgmtswitch 0 --configfile path
```

For *path*, type the full path to the location of the cluster definition file.

For example:

```
# discover --mgmtswitch 0 --configfile /var/tmp/config_file
```

This step assigns an IP address to the spine switch stack. The spine switch stack becomes the head gateway.

5. (Conditional) Type additional `discover` commands for each secondary switch stack.

Complete this step if you have more than one switch.

The formats for these additional commands are as follows:

```
discover --mgmtswitch num --configfile path
```

The arguments are as follows:

Argument	Specification
-----------------	----------------------

<i>num</i>	The identifier for the switch stack.
------------	--------------------------------------

<i>path</i>	The full path to the location of the cluster definition file.
-------------	---

For example:

```
# discover --mgmtswitch 1 --configfile /var/tmp/mac_file
```

6. (Conditional) On the front of the switch stack, plug the cascading cables into the ports on the slave switches.

Complete this step only if you have two or more switch stacks.

This step is the opposite of the following step, in which you unplugged the cascading cables from the ports on the slave switches:

Procedure 4-18, step 3 on page 118

7. Type the following command to retrieve information about the switches that you discovered, and examine the output for errors:

```
# cnodes --mgmt-switch
```

8. Use the `ssh(1)` command to open a console window to the admin node.

9. Type the following command, and monitor the power-on process in the admin node's console window:

```
# tail -f /var/log/cmcdetected.log
```

10. Flip the power breakers on the cluster's CMCs, one rack at a time.

The `cmcdetected` daemon runs on the admin node. It configures the top level switches so that the CMCs are on the appropriate rack VLAN. After you power on rack one, the `cmcdetected` daemon detects the rack and adds the rack to the switch. After the CMCs for rack one appear on the switch, power on the CMCs for rack two.

11. Use the `switchconfig` command, in the following format, to list the CMCs associated with each switch:

```
switchconfig list -s switch
```

For *switch*, specify the system ID for a switch, for example `mgmtsw0`, `mgmtsw1`, and so on.

Issue one of these commands for each switch in your configuration and examine the output. If the output includes all the CMCs in your SGI ICE system, you can proceed with your configuration. If the output contains errors or does not include all expected CMCs, contact your SGI representative for troubleshooting information.

12. Save the configuration to the nonvolatile memory (flash) on the switches.

Note: This step is very important. In the event of a power outage or other interruption, the switch stack boots with the saved configuration.

Type the `switchconfig` command in the following format:

```
switchconfig save_running_config -s mgmtsw0[,mgmtsw1,mgmtsw2,...]
```

Include the parameters `mgmtsw1`, `mgmtsw2`, and so on, only if there are switches in addition to the spine switch (`mgmtsw0`).

13. Use the `switchconfig` command, in the following format, to back up the switch configuration to a file on the admin node:

```
switchconfig pull_switch_config -s switch_ID -f file [--debug]
```

The arguments are as follows:

Argument	Specification
-----------------	----------------------

<i>switch_ID</i>	The switch system ID. For example, this could be <code>mgmtsw0</code> , <code>mgmtsw1</code> , <code>mgmtsw2</code> , and so on. The output from the <code>cnodes</code> command includes this information.
------------------	---

<i>file</i>	The name of the file to receive the switch configuration information. The command writes the file to the <code>/tftpboot/<i>file</i>.cfg</code> . If your <i>file</i> specification ends in <code>.cfg</code> , the command does not append another <code>.cfg</code> string to the file name.
-------------	--

The `--debug` parameter is optional. When specified, the command writes debugging information to `/var/log/switchconfig`.

For example, the following command writes the configuration file for `mgmtsw0` to file `/tftpboot/mgmtsw0_startup1.cfg` on the admin node:

```
# switchconfig pull_switch_config -s mgmtsw0 -f mgmtsw0_startup1 --debug
```

Issue one of these commands per switch.

In the future, if you need to replace the switch, you can save configuration time if you push out this configuration file from the admin node to the new switch.

14. After all management switches have been configured, proceed as follows:
 - If you have MCells, proceed to the following:

"(Conditional) Configuring the Cooling Racks and Cooling Distribution Units (CDUs) on the MCell Network's Switch Ports" on page 126

- If you do not have MCells, proceed to the following:

"Configuring the Cluster With the `discover` Command" on page 129

Configuring Management Switches Without a Cluster Definition File

The following procedure explains how to configure your switches when you do not have the switch information in a cluster definition file.

Procedure 4-20 To configure switches — without a cluster definition file

1. Through an `ssh` connection, log in to the admin node as the root user.
2. (Conditional) Specify a site-specific IP address for the head gateway.

Complete this step if you need to set a site-specific, nondefault IP address for the spine switch stack (head gateway).

By default, SGI sets the IP address for the spine switch stack (head gateway) to 172.23.0.254. Type the following command to set a site-specific IP address:

```
cadadmin -set-head-gateway IP_addr
```

For *IP_addr*, specify your site-specific IP address.

3. Type the `discover` command in one of the following formats to configure the spine switch:

- On SGI ICE platforms, type the following `discover` command:

```
# discover --mgmtswitch 0,ice=yes,type=spine
```

- On SGI Rackable platforms, type the following `discover` command:

```
# discover --mgmtswitch 0,ice=no,type=spine
```

Notes:

- When you specify the `ice=yes` parameter, you indicate to the `discover` command that the cluster includes rack leader controllers and is, therefore, an SGI ICE cluster.

- The `type=type` parameter specifies whether the switch being configured is the spine switch or one of the leaf switches. If you do not specify a `type` parameter, and you are configuring a management switch, the `discover` command uses the link layer discovery protocol (LLDP) to attempt to determine the switch that is directly connected to the admin node.
4. When prompted, connect the spine switch stack to a power source.

To complete this step, plug in the master switch and then the slave switch so that the entire spine switch stack is powered up. In this way, the master switch boots just a few seconds before the slave switch.

The `discover` command configures the MAC address of the switch after you connect the spine switch stack to a power source.

5. (Conditional) Plug in the switch when prompted, and type a `discover` command to configure each leaf switch stack.

Complete this step if you have leaf switch stacks.

Complete the following steps:

- a. Plug in the switch stack when the system prompts you to do so.
- b. Type a `discover` command, in the following format, for your platform::

- On SGI ICE platforms, type one or more `discover` commands in the following format:

```
# discover --mgmtswitch num,ice=yes,type=leaf
```

For *num*, type the identifier for the switch.

For example:

```
# discover --mgmtswitch 1,ice=yes,type=leaf
```

- On SGI Rackable platforms, type one or more `discover` commands in the following format:

```
# discover --mgmtswitch num,ice=no,type=leaf
```

For *num*, type the identifier for the switch.

For example:

```
# discover --mgmtswitch 1,ice=no,type=leaf
```

6. (Conditional) On the front of the switch stack, plug the cascading cables into the ports on the slave switches.

Complete this step only if you have leaf switch stacks.

This step is the opposite of the following step, in which you unplugged the cascading cables from the ports on the slave switches:

Procedure 4-18, step 3 on page 118

7. Type the following command to retrieve information about the switches that you configured, and examine the output for errors:

```
# cnodes --mgmtswitch
```

8. Use the `ssh(1)` command to open a console window to the admin node.

9. Type the following command, and monitor the power-on process in the admin node's console window:

```
# tail -f /var/log/cmcdetected.log
```

10. (SGI ICE clusters only) Flip the power breakers on the cluster's CMCs, one rack at a time.

The `cmcdetected` daemon runs on the admin node. It configures the top level switches so that the CMCs are on the appropriate rack VLAN. After you power on rack one, the `cmcdetected` daemon detects the rack and adds the rack to the switch. After the CMCs for rack one appear on the switch, power on the CMCs for rack two.

11. (SGI ICE clusters only) Use the `switchconfig` command, in the following format, to list the CMCs associated with each switch:

```
switchconfig list -s switch
```

For *switch*, specify the system ID for a switch, for example `mgmtsw0`, `mgmtsw1`, and so on.

Issue one of these commands for each switch in your configuration and examine the output. If the output includes all the CMCs in your cluster, you can proceed with your configuration. If the output contains errors or does not include all expected CMCs, contact your SGI representative for troubleshooting information.

12. Save the configuration to the nonvolatile memory (flash) on the switches.

Note: This step is very important. In the event of a power outage or other interruption, the switch stack boots with the saved configuration.

Type the `switchconfig` command in the following format:

```
# switchconfig save_running_config -s mgmtsw0[,mgmtsw1,mgmtsw2,...]
```

Include the parameters `mgmtsw1`, `mgmtsw2`, and so on, only if there are switches in addition to the spine switch (`mgmtsw0`).

13. For each switch you configured, use the `switchconfig` command, in the following format, to back up the switch configuration to a file on the admin node:

```
switchconfig pull_switch_config -s switch_ID -f file [--debug]
```

The arguments are as follows:

Argument	Specification
<i>switch_ID</i>	The switch system ID. For example, this could be <code>mgmtsw0</code> , <code>mgmtsw1</code> , <code>mgmtsw2</code> , and so on. The output from the <code>cnodes</code> command includes this information.
<i>file</i>	The name of the file to receive the switch configuration information. The command writes the file to the <code>/tftpboot/<i>file</i>.cfg</code> . If your <i>file</i> specification ends in <code>.cfg</code> , the command does not append another <code>.cfg</code> string to the file name.

The `--debug` parameter is optional. When specified, the command writes debugging information to `/var/log/switchconfig`.

For example, the following command writes the configuration file for `mgmtsw0` to file `/tftpboot/mgmtsw0_startup1.cfg` on the admin node:

```
switchconfig pull_switch_config -s mgmtsw0 -f mgmtsw0_startup1 --debug
```

Issue one of these commands per switch.

In the future, if you need to replace a switch, you can save configuration time if you push out this configuration file from the admin node to the new switch.

14. Use the `discover` command, in the following format, to save the MAC addresses to a cluster definition file:

```
# discover --show-configfile > path
```

For *path*, type the full path to the location of the cluster definition file. For example, `/var/tmp/mac_file`.

You can use this cluster definition file if you ever have to configure the switches again.

15. After all management switches have been configured, proceed as follows:
 - If you have MCells, proceed to the following:

"(Conditional) Configuring the Cooling Racks and Cooling Distribution Units (CDUs) on the MCell Network's Switch Ports" on page 126
 - If you do not have MCells, proceed to the following:

"Configuring the Cluster With the `discover` Command" on page 129

(Conditional) Configuring the Cooling Racks and Cooling Distribution Units (CDUs) on the MCell Network's Switch Ports

Perform the procedure in this topic if you have an SGI ICE cluster that includes MCells.

A cluster contains CDUs and cooling rack controllers (CRCs). The CDUs and CRCs have statically assigned IP addresses. These IP addresses are critical to associating the individual rack units (IRUs) with specific CDUs or CRCs. For information about these IP addresses, see the following:

Appendix C, "SGI ICE MCell Network IP Addresses" on page 217

The following procedure explains how to configure the switches attached to the MCells.

Procedure 4-21 To configure MCell switches

1. Gather information about the MCell switches in your cluster.

Visually inspect your system. Note the switches identifiers, and note the port identifiers.

2. Log in as the root user to the admin node.

3. Type following command to retrieve information about the virtual local area networks (VLANs) that are configured at this time:

```
# cattr list -g mcell_vlan
global
  mcell_vlan          : 3
```

The preceding output shows that the MCell VLAN is VLAN 3.

4. Use the `switchconfig set` command, in the following format, to configure the ports on which the CDUs and the CRCs are connected to the MCells:

```
switchconfig set -b none -d vlan_num -p ports -s switch
```

Type an individual `switchconfig set` command for each switch on the cluster network.

The arguments are as follows:

Argument	Specification
<i>vlan_num</i>	The VLAN number of the MCell network. For <i>vlan_num</i> , use the output from the <code>cattr list</code> command as shown earlier in this procedure. The default is 3, and SGI recommends that you do not change this value. This argument appears in two places in the <code>switchconfig</code> command.
<i>ports</i>	Specify the target ports. The command configures both the target ports and the corresponding redundant ports.
<i>switch</i>	The ID number of the management switch to which the CDU or CDC is attached. For example: <code>mgmtsw0</code> .

To determine this value, you need to visually inspect the switch, as follows:

- Locate each CDU or CDC. The following are example labels for CDUs: DU01, DU02, and so on.
- Follow the cable that connects each CDU or CDC to a switch. The following is an example label for a cable that connects each CDU to a switch: DU01-LAN1 | 101MSW0A-36.

- Note the label on the switch. Make sure that the labels on the cables correspond to the labels on the switch ports.

Example 1. The following command configures VLAN 3 on management switch 0 for target ports 1/31, 1/32, and 1/33 and for redundant ports 2/31, 2/32, and 2/33:

```
switchconfig set -b none -d 3 -p 1/31,1/32,1/33 -s mgmtsw0
```

Example 2. The following command configures VLAN 3 on management switch 0 for target ports 2/31, 2/32, and 2/33 and for redundant ports 1/31, 1/32, and 1/33:

```
switchconfig set -b none -d 3 -p 2/31,2/32,2/33 -s mgmtsw0
```

Note: If you make a mistake in your configuration, you can disable the ports from the VLANs you configured. The following example command removes the configuration of VLAN 3 from the target ports and the redundant ports:

```
switchconfig unset -p 1/31,1/32,1/33 -s mgmtsw0
```

-
5. Repeat the following step for each CDU and each CRC attached to your system:

Procedure 4-21, step 4 on page 127

If you encounter errors, issue a `switchconfig set` command again.

6. Save the configuration to the nonvolatile memory (flash) on the switches.

Note: This step is very important. In the event of a power outage or other interruption, the switch stack boots with the saved configuration.

Type the `switchconfig` command in the following format:

```
switchconfig save_running_config -s mgmtsw0[,mgmtsw1,mgmtsw2,...]
```

Include the parameters `mgmtsw1`, `mgmtsw2`, and so on, only if there are switches in addition to the spine switch (`mgmtsw0`).

7. Type the following command to back up the switch configuration to a file on the admin node:

```
switchconfig pull_switch_config -s switch_ID -f file
```

The arguments are as follows:

Argument	Specification
<i>switch_ID</i>	The switch system ID. For example, this could be <code>mgmtsw0</code> , <code>mgmtsw1</code> , <code>mgmtsw2</code> , and so on. The output from the <code>cnodes</code> command includes this information.
<i>file</i>	The name of the file to receive the switch configuration information. The command writes the file to the <code>/tftpboot/<i>file</i>.cfg</code> . If your <i>file</i> specification ends in <code>.cfg</code> , the command does not append another <code>.cfg</code> string to the file name.

For example, the following command writes the configuration file for `mgmtsw0` to file `/tftpboot/mgmtsw0_startup1.cfg` on the admin node:

```
switchconfig pull_switch_config -s mgmtsw0 -f mgmtsw0_startup1 [--debug]
```

The `--debug` parameter is optional. When specified, the command writes debugging information to `/var/log/switchconfig`.

In the future, if you need to replace the switch, you can save configuration time if you push out this configuration file from the admin node to the new switch.

8. After all switches have been configured, proceed to the following:

"Configuring the Cluster With the `discover` Command" on page 129

Configuring the Cluster With the `discover` Command

The `discover` command finds and configures all non-admin nodes and all external switches. If you have a cluster definition file, this procedure can complete more quickly. The procedure in this topic includes configuration steps that explain how to complete the procedure both with and without a cluster definition file.

The following procedure explains how to configure the RLCs (if present), the SGI ICE compute nodes (if present), the flat compute nodes, and the external switches.

Procedure 4-22 To configure the nodes and switches

1. Visually inspect your cluster and note the labels on the nodes.

RLCs are numbered starting with 1. For example, `r1lead`, `r2lead`, and so on.

SGI ICE compute nodes are numbered starting with 0. The numbering depends on the RLC number and on the IRU number within the RLC. For example, the first blade on RLC 1, IRU 1 is numbered as `r1i0n0`, and if there are eight IRUs in the rack, the last blade on the last IRU of RLC 1 is numbered `r1i0n7`.

Flat compute nodes are numbered starting with 0. For example, `n0`, `n1`, `n2`, and so on.

2. Check the power cords on all nodes, as follows:

- If you have a cluster definition file, make sure all nodes are plugged in.

Do not power-on the nodes at this time. When the node is plugged in and connected to a power source, the baseboard management controller (BMC) is started, and that is all that is required at this time.

- If you do not have a cluster definition file, make sure that all nodes are unplugged from their power sources.

3. Through an `ssh` connection, log into the admin node as the root user.
4. Type the following command to retrieve the option code that is in use:

```
# cadmin --show-dhcp-option
```

The nodes determine the integrated Ethernet devices by accepting DHCP leases that belong only to the cluster. Cluster systems use DHCP option code 149 by default. In rare situations, a house network DHCP server could be configured to use this option code. In this case, nodes that are connected to the house network can mistake a house DHCP server as belonging to the cluster's DHCP server, which can lead to an installation failure. Change this option code only if absolutely necessary.

To change the `dhcp` option code number used for this operation, type a command such as the following:

```
# cadmin --set-dhcp-option 150
```

This command sets the DHCP option code to 150.

5. (Conditional) Plug in all the racks and all the flat compute nodes.

Complete this step if you have a cluster definition file.

A cluster definition file contains information about the cluster, including the MAC addresses for the nodes. If you use a cluster definition file, the configuration

process can complete more quickly. Contact your SGI representative to find out if a cluster definition file is available. For more information about cluster definition files, see "Configuring the Switches" on page 109.

6. Use one or more `discover` commands to configure the cluster nodes.

To retrieve online information about the `discover` command, type `discover --h` while logged into a cluster.

For `discover` command examples, see one of the following:

- "discover Command Example 1" on page 133
- "discover Command Example 2" on page 134
- "discover Command Example 3" on page 136
- "discover Command Example 4" on page 136
- "discover Command Example 5" on page 136
- "discover Command Example 6" on page 137

7. (Conditional) When prompted to do so by the system, plug in each individual rack or flat compute node.

Complete this step if you did **not** use a cluster definition file as input to the `discover` command.

The system prompt for this action is as follows:

```
At this time, please turn on the power to this compute node.
Do not turn the system on.
```

The blue light on each component turns on when configuration is complete.

You can use the `console(1)` command if you want to watch the installation progress. The sessions are also logged.

8. Type the following commands to update the configuration files:

```
# update-configs
```

9. Type the following commands to save the configuration to the nonvolatile memory (NVM) on the switches:

```
switchconfig save_running_config -s mgmtsw0[,mgmtsw1,mgmtsw2,...]
```

Include the parameters `mgmtsw1`, `mgmtsw2`, and so on, only if there are switches in addition to management switch 0.

10. (Conditional) Confirm the status of the MIC devices.

Complete this step if your cluster has MIC devices.

Complete the following steps:

- Type the following command to make sure that the MIC devices came online:

```
# pdsh -g leader pdsh -g ice-compute 'micctrl --status'
```

The following example output shows that the MIC devices came online correctly:

```
***** rack_1 *****
***** rack_1 *****
----- rli0n6-----
mic0: online (mode: linux image: /lib/firmware/mic/uos.img)
mic1: online (mode: linux image: /lib/firmware/mic/uos.img)
----- rli0n7-----
mic0: online (mode: linux image: /lib/firmware/mic/uos.img)
mic1: online (mode: linux image: /lib/firmware/mic/uos.img)
----- rli0n8-----
mic0: online (mode: linux image: /lib/firmware/mic/uos.img)
mic1: online (mode: linux image: /lib/firmware/mic/uos.img)
```

- Type the following command to verify the IP addresses and MTU size of the MIC devices:

```
# pdsh -g leader pdsh -g ice-compute 'micctrl --config | grep -E \"MIC IP|Bits|MtuSize\"'
```

The following example output shows that all the MIC devices have addresses on the 10.157.1.0/24 and the 10.158.1.0/24 networks and that the MIC devices use the correct MTU size of 9000:

```
***** rack_1 *****
***** rack_1 *****
----- rli0n4-----
MIC IP:    10.157.1.6
Net Bits:  24
MtuSize:   9000
MIC IP:    10.158.1.6
Net Bits:  24
```

```

        MtuSize: 9000
----- rli0n5-----
        MIC IP: 10.157.1.7
        Net Bits: 24
        MtuSize: 9000
        MIC IP: 10.158.1.7
        Net Bits: 24
        MtuSize: 9000
----- rli0n6-----
        MIC IP: 10.157.1.8
        Net Bits: 24
        MtuSize: 9000
        MIC IP: 10.158.1.8
        Net Bits: 24
        MtuSize: 9000
----- rli0n7-----
        MIC IP: 10.157.1.9
        Net Bits: 24
        MtuSize: 9000
        MIC IP: 10.158.1.9
        Net Bits: 24
        MtuSize: 9000
----- rli0n8-----
        MIC IP: 10.157.1.10
        Net Bits: 24
        MtuSize: 9000
        MIC IP: 10.158.1.10
        Net Bits: 24
        MtuSize: 9000
    
```

11. Proceed to the following:

"Verifying Power Operations and Configuring Power Management" on page 137

discover Command Example 1

If you have a cluster definition file, use the following `discover` command:

```
discover --configfile path_to_CDF --all
```

For *path_to_CDF*, specify the full path to your cluster definition file. The `--all` parameter on this command directs the system to rediscover all nodes and components in the cluster definition file.

discover Command Example 2

If you do not have a cluster definition file, or if you want to configure only selected nodes, use a `discover` command with parameters that specify each node. If you have an SGI ICE cluster, make sure to specify the `--leader` parameter or the `--leaderset` parameter to configure the rack leader controllers (RLCs). For example, use the `discover` command in the following format:

```
discover [--leader[set] ID[,mic=mic_num]] --node[set] specs[,mic=mic_num]  
[,dhcp_bootfile=ipxe][--configfile cluster_def_file][hostname1=name]
```

The arguments are as follows:

Argument	Specification
----------	---------------

ID	<p>Used only for SGI ICE clusters.</p> <p>Specifications used to configure the SGI ICE compute nodes, including the ID number(s) for the rack(s) that you want to configure.</p> <p>For example, if you want to configure one rack for an SGI ICE cluster, specify <code>--leader</code> and the system ID number that corresponds to that rack.</p> <p>If you want to configure a range of racks, specify <code>--leaderset</code>, the starting system ID number, a comma (<code>,</code>), and the ending system ID number.</p> <p>Examples:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;"><code>--leader 2</code></td> <td>Configures RLC 2. The RLC is configured with a temponame of <code>r2lead</code>.</td> </tr> <tr> <td><code>--leaderset 1,3</code></td> <td>Configures RLCs 1, 2, and 3. The RLCs are configured with the following temponames: <code>r1lead</code>, <code>r2lead</code>, <code>r3lead</code>.</td> </tr> </table>	<code>--leader 2</code>	Configures RLC 2. The RLC is configured with a temponame of <code>r2lead</code> .	<code>--leaderset 1,3</code>	Configures RLCs 1, 2, and 3. The RLCs are configured with the following temponames: <code>r1lead</code> , <code>r2lead</code> , <code>r3lead</code> .
<code>--leader 2</code>	Configures RLC 2. The RLC is configured with a temponame of <code>r2lead</code> .				
<code>--leaderset 1,3</code>	Configures RLCs 1, 2, and 3. The RLCs are configured with the following temponames: <code>r1lead</code> , <code>r2lead</code> , <code>r3lead</code> .				
mic_num	<p>Specify the number of Intel® Many Integrated Core (MIC) devices that reside on each node. By default, the <code>discover</code> command assumes zero (0). If you have MIC devices on any SGI ICE compute nodes, specify</p>				

the number you have, which can be 1 or 2. For flat compute nodes, *mic_num* can be 1, 2, 3, or 4.

Specify the *,mic=mic_num* parameter only if your cluster includes MIC devices.

specs

Specifications used to configure the flat compute nodes. After the cluster is completely configured, you can configure services, for example DNS or Lustre, on one or more of these nodes.

To configure only one node, specify the *--node* parameter. To configure a series of nodes, specify the *--nodeset* parameter.

Examples:

```
--node 2
```

Configures flat compute node 2.

The node is configured with a *temponame* of *service2* and a *hostname* of *n2*.

```
--nodeset 1,3
```

Configures flat compute nodes 1, 2, and 3.

The nodes are configured with the following *temponames*: *service1*, *service2*, and *service3*.

The nodes are configured with the following *hostnames*: *n1*, *n2*, and *n3*.

```
--nodeset 200,100,hostname1=snXXX
```

Configures flat compute nodes 200 through 299.

The nodes are configured with the following *temponames*: *service200*, *service201*, and so on up to *service299*.

The nodes are configured with the following *hostnames*: *sn200*, *sn201*, and so on up to *sn299*.

name

The *hostname* that you want to assign to the nodes.

For example, if you specify *--hostname1=newclusterXXX* all nodes configured with this *discover* command are named with *newcluster* plus a numeric suffix. That is, the command results in

nodes named `newcluster0`, `newcluster1`, `newcluster2`, `newcluster3`, and so on.

Use the `dhcp_bootfile=ipxe` parameter in troubleshooting situations. If you already issued a `discover` command and one or more nodes failed to boot, specify the `dhcp_bootfile=ipxe` parameter, which directs the server boot agent to load iPXE rather than GRUB version 2. When this parameter is used, the iPXE software loads GRUB version 2.

discover Command Example 3

The following command uses a cluster definition file to configure rack 1 and flat compute node 0:

```
# discover --configfile myconfigfile --leader 1 --node 0
```

This command does not specify a cluster definition file, so when the command runs, you need to power cycle each component on and off.

discover Command Example 4

If you have one rack of SGI ICE compute nodes and one flat compute node, type the following command:

```
# discover --leader 1 --node 0
```

This command does not specify a cluster definition file, so when the command runs, you need to power cycle each component on and off.

discover Command Example 5

If you have five racks of SGI ICE compute nodes and three flat compute nodes, type the following command:

```
# discover --leaderset 1,5 --nodeset 1,3
```

This command does not specify a cluster definition file, so when the command runs, you need to power cycle each component on and off.

discover Command Example 6

If you have one rack of SGI ICE compute nodes, one flat compute node, two MIC devices attached to each blade on the SGI ICE compute node rack, and two MIC devices attached to the flat compute node, type the following command:

```
# discover --leader 1,mic=2 --node 0,mic=2,image=mic_serv_image
```

For *mic_serv_image*, specify the label (name) of the flat compute node image that you created in "Creating Images for the Flat Compute Nodes That Include MIC Devices" on page 103.

This command does not specify a cluster definition file, so when the command runs, you need to power cycle each component on and off.

Verifying Power Operations and Configuring Power Management

The SGI power management service provides the following features:

- Power and energy measurement.
- A power limiting capability that you can configure for the entire cluster, for specific racks or rack sets, or for individual nodes within the cluster.

SGI supports the power management feature on the following platforms:

- SGI ICE XA and SGI ICE X series systems
- SGI Rackable systems that support the Intel Xeon E3-2600 Series and the Intel Xeon E5-2600 Series processor platforms

SGI does not support power management on SGI Altix ICE 8400 or SGI Altix ICE 8200 systems.

For more information about power operation, see the following:

SGI Management Center (SMC) Administration Guide for Clusters

The following procedure explains how to determine if the power management service is configured correctly and how to set an initial system-wide limit.

Note: Several lines in the output examples in the following procedures have been wrapped and indented for inclusion in this documentation.

Procedure 4-23 To verify the installation of the power management software and configure an initial power management threshold

1. Log into the admin node as the root user, and type the following command to verify that the correct RPMs are present:

```
# rpm -qa | grep sgi-ta
sgi-tacli-3.1.x-sgi712r7.sles11sp3
sgi-talib-3.1.x-sgi712r6.sles11sp3
sgi-tasac-3.1.x-sgi712r12.sles11sp3
```

The output shows that the three packages you need are present on the admin node. These package names begin with `sgi-tacli-3.1.x`, `sgi-talib-3.1.x`, and `sgi-tasac-3.1.x`. This is correct for the SMC 3.1.x releases. The rest of the package name differs depending on the operating system platform. These packages are correct for the SMC 3.1.x releases on a SLES 11 SP3 platform.

Contact your SGI representative if packages seem to be missing or incorrect.

2. Verify that the power daemon is running.

This command differs depending on your platform, as follows:

- On RHEL 7.1 platforms and SLES 12 platforms, type the following command:

```
# systemctl status sgi-ta
```

- On RHEL 6.x platforms and SLES 11 SPx platforms, type the following command:

```
# service sgi-ta status
```

Example 1. The following output shows that the daemon is running as expected on a RHEL 7.1 system:

```
# systemctl status sgi-ta
sgi-ta.service - LSB: SGI-TA Tempo Twisted server for SAC
   Loaded: loaded (/etc/rc.d/init.d/sgi-ta)
   Active: active (exited) since Mon 2015-06-29 14:34:29 CDT; 2 weeks 6 days ago
```

Example 2. The following output shows that the daemon is running as expected on a SLES 11 SP3 system:

```
# service sgi-ta status
Checking for service sgi-ta                                done
```

If the daemon is not running, start the daemon. Use one of the following commands:

- On RHEL 7.1 platforms and SLES 12 platforms, type the following command:

```
# systemctl start sgi-ta
```

- On RHEL 6.x and SLES 11 SPx platforms, type the following command:

```
# service sgi-ta start
```

3. Use a text editor to open the log file for the `sacmain.tac` daemon on the admin node, and verify that the log entries indicate a running daemon.

The log file resides in `/var/log/sgi-ta`.

Example 1. The following log file entries show that the `sacmain.tac` daemon is running as expected:

```
2015-04-13 14:08:32+0000 [-] Starting factory 2015-04-13 14:49:46+0000 [-] clustermap.populate
- unused Status table entry:1:should_not_be_installed
2015-04-13 14:49:46+0000 [-] clustermap.populate - unused Status table entry:2:should_be_installed
2015-04-13 14:49:46+0000 [-] clustermap.populate - unused Status table entry:3:undefined_1
2015-04-13 14:49:46+0000 [-] clustermap.populate - unused Status table entry:4:undefined_2
2015-04-13 14:49:46+0000 [-] clustermap.populate - unused Status table entry:5:undefined_3
2015-04-13 14:49:46+0000 [-] clustermap.populate - unused Status table entry:6:undefined_4
2015-04-13 14:49:46+0000 [-] clustermap.populate - unused Status table entry:7:undefined_5
2015-04-13 14:49:46+0000 [-] clustermap.populate - unused Status table entry:8:finished
2015-04-13 14:49:46+0000 [-] clustermap.populate - status tables loaded
2015-04-13 14:49:46+0000 [-] clustermap.populate - cluster data loaded
2015-04-13 14:49:46+0000 [-] clustermap.populate - cluster config loaded
2015-04-13 14:49:46+0000 [-] clustermap.populate - sac 'icicle' loaded
2015-04-13 14:49:46+0000 [-] clustermap.populate - service nodes loaded :3
2015-04-13 14:49:46+0000 [-] clustermap.populate - rack data loaded. 2 racks
2015-04-13 14:49:46+0000 [-] clustermap.populate - rack leaders loaded.
```

Example 2. The following log file entries show the RLCs connecting and are also to be expected:

```
2015-04-13 14:56:43+0000 [sacore.LFactory] build leader protocol for IPv4Address(TCP, '172.23.0.3', 43448)
2015-04-13 14:56:43+0000 [LBroker,0,172.23.0.3] leader connection made. protocol:2015-04-13 14:56:43+0000
[LBroker,0,172.23.0.3] registered leader reference for r2lead
2015-04-13 14:57:27+0000 [sacore.LFactory] build leader protocol for IPv4Address(TCP, '172.23.0.2', 45822)
2015-04-13 14:57:27+0000 [LBroker,1,172.23.0.2] leader connection made. protocol:2015-04-13 14:57:27+0000
```

```
[LBroker,1,172.23.0.2] registered leader reference for r1lead
```

If the log file entries show Traceback activity, the daemon might not be running correctly. If you see Traceback entries, and you need help to interpret them, contact your SGI representative.

4. (Conditional) Use the `ssh(1)` command to log into an RLC, and verify that the correct RPMs are present on each RLC.

Complete this step if your cluster is an SGI ICE cluster.

Type the following commands:

```
# ssh r1lead
# rpm -qa | grep sgi-ta
sgi-talib-3.1.0-sgi712r6.sles11sp3
sgi-tarlc-3.1.0-sgi712r3.sles11sp3
```

The output shows that the two packages you need are present on the RLC. These package names begin with `sgi-talib-3.1.0` and `sgi-tarlc-3.1.0`. This is correct for the SMC 3.1 release. These packages are correct for SMC 3.1 on a SLES 11 SP3 platform.

Make sure to repeat this step on each RLC node in the cluster.

Contact your SGI representative if packages seem to be missing or incorrect.

5. Verify that the power service daemon, `sgi-ta`, is running on each RLC.

This command differs depending on your platform, as follows:

- On RHEL 7.1 platforms and SLES 12 platforms, type the following command:

```
# systemctl status sgi-ta
```

- On RHEL 6.x platforms and SLES 11 SPx platforms, type the following command:

```
# service sgi-ta status
```

Example 1. The following output shows that the daemon is running as expected on a RHEL 7.1 system:

```
# systemctl status sgi-ta
sgi-ta.service - LSB: SGI-TA Tempo Twisted server for SAC
   Loaded: loaded (/etc/rc.d/init.d/sgi-ta)
```

Active: active (exited) since Mon 2015-06-29 14:34:29 CDT; 2 weeks 6 days ago

Example 2. The following output shows that the daemon is running as expected on a SLES 11 SP3 system:

```
# service sgi-ta status
Checking for service sgi-ta                               done
```

If the daemon is not running, start the daemon. Use one of the following commands:

- On RHEL 7.1 platforms and SLES 12 platforms, type the following command:

```
# systemctl start sgi-ta
```

- On RHEL 6.x and SLES 11 SPx platforms, type the following command:

```
# service sgi-ta start
```

Make sure to repeat this step on each RLC node in the cluster.

6. (Conditional) Use a text editor to open the log file for the power management daemon on each RLC node, and verify that the log entries indicate a running daemon.

Complete this step if your cluster is an SGI ICE cluster.

The log file resides in `/var/log/sgi-ta`.

Example 1. The following output shows that the `rlcmain.tac` daemon is running as expected:

```
2015-04-13 14:08:32+0000 [-] Starting factory 2015-04-13 14:49:46+0000 [-] clustermap.populate
- unused Status table entry:1:should_not_be_installed
2015-04-13 14:49:46+0000 [-] clustermap.populate - unused Status table entry:2:should_be_installed
2015-04-13 14:49:46+0000 [-] clustermap.populate - unused Status table entry:3:undefined_1
2015-04-13 14:49:46+0000 [-] clustermap.populate - unused Status table entry:4:undefined_2
2015-04-13 14:49:46+0000 [-] clustermap.populate - unused Status table entry:5:undefined_3
2015-04-13 14:49:46+0000 [-] clustermap.populate - unused Status table entry:6:undefined_4
2015-04-13 14:49:46+0000 [-] clustermap.populate - unused Status table entry:7:undefined_5
2015-04-13 14:49:46+0000 [-] clustermap.populate - unused Status table entry:8:finished
2015-04-13 14:49:46+0000 [-] clustermap.populate - status tables loaded
2015-04-13 14:49:46+0000 [-] clustermap.populate - cluster data loaded
2015-04-13 14:49:46+0000 [-] clustermap.populate - cluster config loaded
2015-04-13 14:49:46+0000 [-] clustermap.populate - sac 'icicle' loaded
```

```
2015-04-13 14:49:46+0000 [-] clustermap.populate - service nodes loaded :3
2015-04-13 14:49:46+0000 [-] clustermap.populate - rack data loaded. 2 racks
2015-04-13 14:49:46+0000 [-] clustermap.populate - rack leaders loaded.
```

Example 2. Traceback messages in a log file can show a problem that occurred and then resolved, thus eliminating the problem. The following log entries include a Traceback message containing information about a temporary connection failure from which the cluster recovered:

```
2015-04-13 14:57:11+0000 [-] Log opened.
2015-04-13 14:57:11+0000 [-] twistd 8.2.0 (/usr/bin/python 2.6.8) starting up.
2015-04-13 14:57:11+0000 [-] reactor class: twisted.internet.epollreactor.EPollReactor.
2015-04-13 14:57:11+0000 [-] tlib.theart.HeartbeatMonitorProtocol starting on 1319
2015-04-13 14:57:11+0000 [-] Starting protocol 2015-04-13 14:57:11+0000 [-] ra proxy start
2015-04-13 14:57:11+0000 [-] SAC connection attempt failed:[Failure instance: Traceback (failure with no frames):
  : No route to host: 101: Network is unreachable.
    ] - continue trying
2015-04-13 14:57:27+0000 [Broker,client] SAC connected. get root object.
2015-04-13 14:57:27+0000 [Broker,client] SAC root reference obtained. attempt registration
2015-04-13 14:57:27+0000 [Broker,client] SAC registration success. sending heartbeat state immediately
2015-04-13 14:57:27+0000 [Broker,client] heartbeat state sent
```

Make sure to repeat this step for each daemon on each RLC node in the cluster.

7. Configure an initial power management threshold.

The power management hardware on compute nodes must be initialized in order for power readings to function. This is one-time step to be performed now.

Type the following command:

```
# mpower system set_limit 1000W
```

The command applies an initial 1000W limit to every SGI ICE compute node and every flat compute node. The 1000W value is higher than any compute node's maximum power draw. As a result, there is no power limit, but the power measurement averaging can now function.

8. Proceed to one of the following:

- "(Conditional) Verifying That the smcrbd daemon and the icerackd Daemon are Running" on page 143
- "(Optional) Configuring a Backup Domain Name Service (DNS) Server" on page 146

- "(Conditional) Configuring the InfiniBand Subnetworks" on page 147

(Conditional) Verifying That the `smcrbd` daemon and the `icerackd` Daemon are Running

Perform the procedure in this topic if you have an SGI ICE cluster. You do not need to perform this procedure if you have an SGI Rackable cluster.

The following procedure explains how to make sure that the `smcrbd` daemon and the `icerackd` daemon are running properly.

Procedure 4-24 To verify the daemons

1. Log into the admin node as the root user, and make sure that the `smcrbd` daemon is running.

This command differs depending on your platform, as follows:

- On RHEL 7.1 platforms and SLES 12 platforms, type the following command:

```
# systemctl status smcrbd
```

- On RHEL 6.x platforms and SLES 11 SPx platforms, type the following command:

```
# service smcrbd status
```

Note: The lines in some of the following examples are wrapped for inclusion in this documentation.

Example 1. The following example shows the daemon running as expected on a RHEL 7.1 system:

```
# systemctl status smcrbd
smcrbd.service - LSB: Automatic blade discovery daemon
   Loaded: loaded (/etc/rc.d/init.d/smcrbd)
   Active: active (running) since Mon 2015-06-29 14:34:35 CDT; 2 weeks 6 days ago
   CGroup: /system.slice/smcrbd.service
           +-2062 /usr/bin/python /usr/bin/twistd -l /var/log/smcrbd.log --pidfile /var/run/smcrbd.pid
           -y /opt/sgi/lib/smcrbd
```

Example 2. The following example shows the daemon running as expected on a SLES 11 SPx system:

```
# service smcrbd status
Checking for service smcrbd
```

Example 3. The following example shows the daemon running as expected on a RHEL 6.x system:

```
# service smcrbd status
Checking for service smcrbd [ OK ]
```

If the daemon is not running, start the daemon. Use one of the following commands:

- On RHEL 7.1 platforms and SLES 12 platforms, type the following command:

```
# systemctl start smcrbd
```

- On RHEL 6.x and SLES 11 SPx platforms, type the following command:

```
# service smcrbd start
```

2. Use a text editor to open file `/var/log/smcrbd.log`, which is the log file for the `smcrbd` daemon on the admin node, and verify that the log entries indicate a running daemon.

For example, the following log file entries show that the `smcrbd-log` daemon is running as expected:

```
2015-04-23 09:04:42-0500 [Broker,1,172.23.0.2] called remotely to add cmc: 1 1 08:00:69:16:53:fe
2015-04-23 09:04:42-0500 [Broker,1,172.23.0.2] getCmcNode()
2015-04-23 09:04:42-0500 [-] addCmc: rack = 1, iru = 1, mac = 08:00:69:16:53:fe
2015-04-23 09:04:42-0500 [-] rack = 1, iru = 1, name = rli1c, hostname = rli1c
2015-04-23 09:04:42-0500 [-] CMC already in database
2015-04-23 09:07:21-0500 [Broker,1,172.23.0.2] called remotely to add cmc: 1 0 08:00:69:16:55:15
2015-04-23 09:07:21-0500 [Broker,1,172.23.0.2] getCmcNode()
2015-04-23 09:07:21-0500 [-] addCmc: rack = 1, iru = 0, mac = 08:00:69:16:55:15
2015-04-23 09:07:21-0500 [-] rack = 1, iru = 0, name = rli0c, hostname = rli0c
2015-04-23 09:07:21-0500 [-] CMC already in database
```

If the log file entries show traceback activity, the daemon might not be running correctly. If you see traceback entries, and you need help to interpret them, contact your SGI representative.

3. Use the `ssh(1)` command to log into an RLC.

For example, for `r1lead`, type the following command:

```
# ssh r1lead
Last login: Wed Apr 15 18:43:16 2015 from admin
```

4. Verify that the `icerackd` daemon is running.

This command differs depending on your platform, as follows:

- On RHEL 7.1 platforms and SLES 12 platforms, type the following command:

```
# systemctl status icerackd
```

- On RHEL 6.x platforms and SLES 11 SPx platforms, type the following command:

```
# service icerackd status
```

For example, the following output shows that the daemon is running as expected:

```
Checking for service icerackd done
```

If the daemon is not running, start the daemon. Use one of the following commands:

- On RHEL 7.1 platforms and SLES 12 platforms, type the following command:

```
# systemctl start icerackd
```

- On RHEL 6.x and SLES 11 SPx platforms, type the following command:

```
# service icerackd start
```

Make sure to verify that the daemon is running on each RLC node in the cluster.

5. On each RLC node, use a text editor to open file `/var/log/icerackd.log`, which is the log file for the power management daemon, and use the `ps(1)` command to verify that the log entries indicate a running daemon.

For example, the following output shows that the `icerackd` daemon is running as expected:

```
2015-04-22 19:41:14-0500 [-] pcap heard: when=(1429749674, 581382) sz=334B
2015-04-22 19:41:14-0500 [-] CMC: mac = 08:00:69:16:55:15, rack=1, id=0
2015-04-22 19:41:14-0500 [-] calling remote add_cmc(1, 0, 08:00:69:16:55:15)
```

```
2015-04-22 19:41:14-0500 [Broker,client] server reported:  {'mac': '08:00:69:16:55:15', 'cmc': '0', 'rack': '1'}
2015-04-22 19:41:15-0500 [-] pcap heard: when=(1429749675, 714942) sz=334B
2015-04-22 19:41:15-0500 [-] CMC: mac = 08:00:69:16:53:fe, rack=1, id=1
2015-04-22 19:41:15-0500 [-] calling remote add_cmc(1, 1, 08:00:69:16:53:fe)
2015-04-22 19:41:15-0500 [Broker,client] server reported:  {'mac': '08:00:69:16:53:fe', 'cmc': '1', 'rack': '1'}
```

Make sure to repeat this step for each daemon on each RLC node in the cluster.

6. Proceed to one of the following:

- "(Optional) Configuring a Backup Domain Name Service (DNS) Server" on page 146
- "(Conditional) Configuring the InfiniBand Subnetworks" on page 147

(Optional) Configuring a Backup Domain Name Service (DNS) Server

Typically, the DNS on the admin node provides name services for the cluster. When you configure a backup DNS, however, the cluster can use a flat compute node as a secondary DNS server if the admin node is down, being serviced, or is otherwise not available. You can configure a backup DNS only after you run the `discover` command to configure the cluster. This is an optional feature.

The following procedure explains how to configure a flat compute node to act as a DNS.

Procedure 4-25 To enable a backup DNS

1. Through an `ssh` connection, log into the admin node as the root user.
2. Type the following command to retrieve a list of available flat compute nodes:

```
# cnodes --compute
```

The flat compute node you want to use as a backup DNS must be configured in the system already. That is, you must have run the `discover` command to configure the flat compute node.

3. Type the following command to start the cluster configuration tool:

```
# /opt/sgi/sbin/configure-cluster
```

4. On the **Main Menu** screen, select **B Configure Backup DNS Server (optional)**, and select **OK**.

5. On screen that appears, type the identifier for the flat compute node that you want to designate as the backup DNS, and select **OK**.

For example, you could configure flat compute node `n101` as the host for the backup DNS server.

To disable this feature, select **Disable Backup DNS** from the same menu and select **Yes** to confirm your choice.

6. (Conditional) Proceed to the following:

"(Conditional) Configuring the InfiniBand Subnetworks" on page 147

(Conditional) Configuring the InfiniBand Subnetworks

Perform the procedures in this topic as follows:

- If you have an SGI ICE cluster. On this platform, you can configure the InfiniBand subnetwork either on a rack leader controller (RLC) or on a flat compute node.
- If you have an SGI Rackable cluster, and you want to configure the InfiniBand subnetwork on one of the flat compute nodes. Some InfiniBand switches on SGI Rackable clusters come configured for an InfiniBand subnetwork. You can perform the procedure in this topic if your switch is not preconfigured for InfiniBand or if you prefer to configure this service on a flat compute node.

The InfiniBand network on the cluster uses Open Fabrics Enterprise Distribution (OFED) software. For information about OFED, see the following website:

<http://www.openfabrics.org>

For more information about the InfiniBand fabric implementation on SGI clusters, see the *SGI Management Center (SMC) Administration Guide for Clusters*.

Each cluster has two InfiniBand fabric network cards, `ib0` and `ib1`. Each subnetwork has a subnet manager, which runs on an RLC or on a flat compute node.

The following procedure explains how to configure the master and the standby components and how to verify the configuration:

- "Configuring the InfiniBand Subnetworks" on page 148

Configuring the InfiniBand Subnetworks

The following procedure explains how to configure the InfiniBand subnetwork master and standby components on an SGI ICE cluster or on an SGI Rackable cluster.

Procedure 4-26 To configure the InfiniBand subnetworks

1. Through an `ssh` connection, log into the admin node as the root user.
2. Type the following command to disable InfiniBand switch monitoring:

```
# cattr set disableIbSwitchMonitoring true
```

The system sometimes issues InfiniBand switch monitoring errors before the InfiniBand network has been fully configured. The preceding command disables InfiniBand switch monitoring.

3. Use one of the following methods to access the InfiniBand network configuration tool:

- Type the following command to start the cluster configuration tool:

```
# /opt/sgi/sbin/configure-cluster
```

After the cluster configuration tool starts, select **F Configure InfiniBand Fabric**, and select **OK**.

- Type the following command to start the InfiniBand management tool:

```
# tempo-configure-fabric
```

Both of the preceding methods lead you to the same InfiniBand configuration page. On the InfiniBand configuration pages, **Quit** takes you to the previous screen.

4. Select **A Configure InfiniBand ib0**, and select **Select**.
5. On the **Configure InfiniBand** screen, select **A Configure Topology**, and select **Select**.
6. On the **Topology** screen, select the topology your system uses, and select **Select**.

The menu selections are as follows:

- **H HYPERCUBE**
- **E EHYPERCUBE** (Enhanced Hypercube)

- **F FAT TREE**
 - **G BFTREE**
7. On the **Configure InfiniBand** screen, select **B Master / Standby**, and select **Select**.
 8. On the **Master / Standby** screen, type the component identifiers for the master (primary) and the standby (backup, secondary) subnetwork, and select **Select**.

Example 1. On an SGI ICE cluster, if you have only one rack leader controller (RLC), type `r1lead` in the **MASTER** field, and leave the **STANDBY** field blank. If you have more than one RLC, specify different RLCs in the **MASTER** and **STANDBY** fields.

Figure 4-12 on page 149 shows a completed screen.

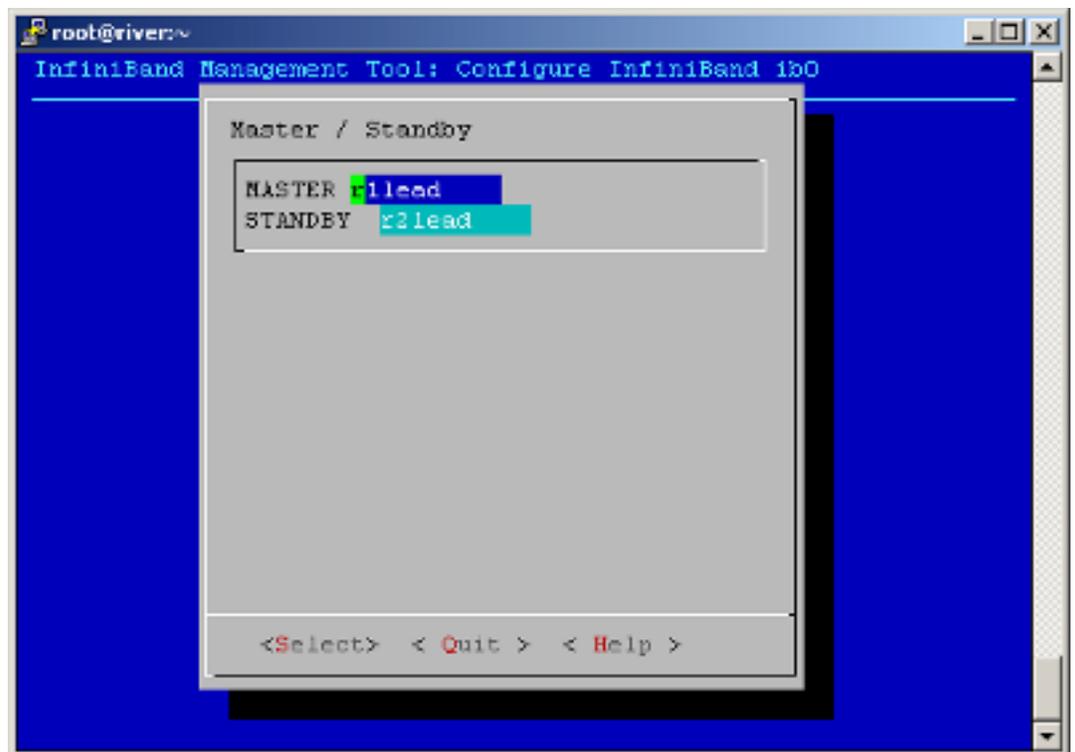


Figure 4-12 Completed InfiniBand (ib0) Master / Standby Screen

Example 2. On an SGI Rackable cluster, type `n1` in the **MASTER** field, and type `n101` in the **STANDBY** field.

9. On the **Configure InfiniBand** screen, select **Commit**.

Wait for the confirmatory messages to appear in the window before you continue.

The next few steps repeat the preceding steps, but this time you configure the `ib1` interface.

10. On the InfiniBand Management Tool main menu screen, select **B Configure InfiniBand ib1**, and select **Select**.
11. On the **Configure InfiniBand** screen, select **A Configure Topology**, and select **Select**.
12. On the **Topology** screen, select the topology your system uses, and select **Select**.
Select the topology that exists on your system. The menu selections are as follows:
 - **H HYPERCUBE**
 - **E EHYPERCUBE** (Enhanced Hypercube)
 - **F FAT TREE**
 - **G BFTREE**
13. On the **Configure InfiniBand** screen, select **B Master / Standby**, and select **Select**.
14. On the **Master / Standby** screen, type the component identifiers for the master (primary) and the standby (backup, secondary) subnetwork, and select **Select**.

Example 1. On an SGI ICE cluster, if you have only one rack leader controller (RLC), type `r1lead` in the **MASTER** field, and leave the **STANDBY** field blank. If you have two RLCs, you can flip the specifications for `ib1`. For example, assume that for `ib0`, you specified **MASTER** as `r1lead` and **STANDBY** as `r2lead`. For `ib1`, you can specify **MASTER** as `r2lead` and **STANDBY** as `r1lead`. If you have three or more RLCs, specify different RLCs in the **MASTER** and **STANDBY** fields.

Example 2. On an SGI Rackable cluster, type `n101` in the **MASTER** field, and type `n1` in the **STANDBY** field.

15. On the **Configure InfiniBand** screen, select **Commit**.

Wait for the confirmatory messages to appear in the window before you continue.

16. On the InfiniBand Management Tool main menu screen, select **C Administer Infiniband ib0**, and select **Select**.
 17. On the **Administer InfiniBand** screen, select **Start**, and select **Select**.
 18. On the **Start SM master_ib0 on ib0 succeeded!** screen, select **OK**.
 19. Select **Quit** to return to the InfiniBand Management Tool main menu screen.
- The next few steps repeat the preceding steps, but this time you start the `ib1` interface.
20. On the InfiniBand Management Tool main menu screen, select **D Administer Infiniband ib1**, and select **Select**.
 21. On the **Administer InfiniBand** screen, select **Start**, and select **Select**.
 22. On the **Start SM master_ib1 on ib1 succeeded!** screen, select **OK**.
 23. On the **Administer InfiniBand** screen, select **Status**, and select **Select**.

The **Status** option returns information similar to the following:

```
Master SM
Host = rllead
Guid = 0x0002c9030006938b
Fabric = ib0
Topology = hypercube
Routing Engine = dor
OpenSM = running
```

24. Wait for the status messages to stop, and press `Enter`.
 25. Select **Quit** on the menus that follow to exit the configuration tool.
 26. Use the `ssh(1)` command to change to one of the leader nodes.
- For example:
- ```
ssh rllead
```
27. Use the `ibstatus` command to retrieve the status information for this node.

For example:

```
ibstatus
Infiniband device 'mlx4_0' port 1 status:
```

```
default gid: fec0:0000:0000:0000:0002:c903:00f3:5311
base lid: 0x1
sm lid: 0x1
state: 4: ACTIVE
phys state: 5: LinkUp
rate: 56 Gb/sec (4X FDR)
link_layer: InfiniBand

Infiniband device 'mlx4_0' port 2 status:
default gid: fec0:0000:0000:0001:0002:c903:00f3:5312
base lid: 0x2a
sm lid: 0x2a
state: 4: ACTIVE
phys state: 5: LinkUp
rate: 56 Gb/sec (4X FDR)
link_layer: InfiniBand
```

The output shows the status as `ACTIVE` on both ports, which is correct.

28. Use the `ibhosts` command to retrieve the status of the compute nodes, and verify that each node you configured is connected to the RLC.

For example:

```
rllead:~ # ibhosts
Ca : 0x0002c9030032bd50 ports 1 "rli0n8 HCA-1"
Ca : 0x0002c9030014a630 ports 1 "rli0n7 HCA-1"
Ca : 0x0002c9030014b140 ports 1 "rli0n6 HCA-1"
Ca : 0x0002c9030018cf00 ports 1 "rli0n5 HCA-1"
Ca : 0x0002c9030018cfa0 ports 1 "rli0n13 HCA-1"
Ca : 0x0002c9030018ce90 ports 1 "rli0n15 HCA-1"
Ca : 0x0002c9030014a610 ports 1 "rli0n14 HCA-1"
Ca : 0x0002c9030014b110 ports 1 "rli0n16 HCA-1"
Ca : 0x0002c9030018d170 ports 1 "rli0n17 HCA-1"
```

The output shows each node connected, as expected, to the leader.

## Configuring Additional Features

There are several additional features that you might need to configure, depending on your platform and configuration. For example, if you have an SGI ICE XA cluster

with a SLES 11 SP3 platform, you need to upgrade your Mellanox OFED software for performance reasons.

The information in the following chapter explains how to configure additional features that you might want or need:

Chapter 5, "Configuring Additional Features" on page 155



## Configuring Additional Features

This chapter includes the following topics:

- "Configuring Array Services for MPI Programs" on page 155
- "(Conditional) Enabling the Mellanox OpenFabrics Enterprise Distribution (OFED) for Linux Software" on page 167
- "Troubleshooting Configuration Changes" on page 174

### About Additional Features

This chapter includes information about additional features, such as Array Services, that you might want to configure after your cluster is configured.

SGI includes additional optional software in the SGI Foundation Software (SFS) package. For information about the hardware event tracker (HET), CPU frequency scaling, and other SFS software packages, see the following:

*SGI Foundation Software (SFS) User Guide*

### Configuring Array Services for MPI Programs

You can configure compute nodes into an array. After you configure a set of compute nodes into an array, the Array Services software can perform authentication and coordination functions when MPI programs are running. For more information about MPI, see the *SGI MPI and SGI SHMEM User Guide*

You cannot include an SGI ICE admin node or any rack leader controller (RLC) nodes in an array. You can include an SGI Rackable admin node in an array.

For general Array Services configuration information, see the following:

- `arrayconfig(8)`, which describes how to use the `arrayconfig(8)` command to configure Array Services.
- `arrayconfig_tempo(8)`, which describes Array Services configuration characteristics that are specific to clusters.

The procedures in the following topics assume that you want to create new a master image for the compute nodes and a new master image for the flat compute nodes that you configured with user services. After you create these images, you can push out the new images to the compute nodes and to the compute services nodes. The alternative is to configure Array Services directly on the nodes themselves, but this method leaves you with an Array Services configuration that is overwritten the next time someone pushes new software images to the cluster's nodes.

---

**Note:** The procedures in the following topics assume that you want to install Array Services on an SGI ICE cluster. The steps you need to complete for an SGI Rackable cluster are very similar with the major difference being that you do not need to perform steps that pertain to racks. On an SGI Rackable cluster, the steps in the following procedures that pertain to racks and rack leader controllers (RLCs) do not need to be performed.

---

The following procedures explain how to configure Array Services.

- "Planning the Configuration" on page 156
- "Preparing the Images" on page 159
- "Power Cycling the Nodes and Pushing Out the New Images" on page 166

## Planning the Configuration

The following procedure explains preliminary steps that you can take to plan your array and the security you want to enable.

**Procedure 5-1** To plan the array

1. Verify that the SGI Performance Suite software is installed on the cluster.
2. Log into the admin node as the root user.
3. Use the `cnodes` command to display a list of available nodes, and decide which nodes you want to include in the array.

For example:

- To gather information about flat compute nodes, type the following command:

```
cnodes --compute
```

- To gather information about SGI ICE compute nodes, type the following command:

```
cnodes --ice-compute
```

The command shows all the compute nodes, including those that might be configured as compute services nodes at this time.

4. Use the `cinstallman` command to display a list of the available system images, and decide which images you want to edit.

For example, the following output is for an example SGI ICE cluster that is running in production mode:

```
cinstallman --show-images
Image Name BT VCS Compat_Distro
lead-sles11sp3 0 1 sles11sp3 # Default RLC node image
 3.0.101-0.47.52-default
 3.0.76-0.11-default
lead-sles11sp3.prod1 0 1 sles11sp3 # Site's production RLC image
 3.0.101-0.47.52-default
 3.0.76-0.11-default
sles11sp3 0 1 sles11sp3 # Default flat compute node image
 3.0.101-0.47.52-default
 3.0.76-0.11-default
sles11sp3.prod1 0 1 sles11sp3 # Site's production compute services node image
 3.0.101-0.47.52-default
 3.0.76-0.11-default
ice-sles11sp3 0 1 sles11sp3 # Default SGI ICE compute node image
 3.0.101-0.47.52-default
 3.0.76-0.11-default
ice-sles11sp3.prod1 0 1 sles11sp3 # Site's production SGI ICE compute node image
 3.0.101-0.47.52-default
 3.0.76-0.11-default
```

The output includes image `sles11sp3.prod1`, which is the image installed on a flat compute node that is configured as a compute services node. Image `sles11sp3.prod1` is based on image `sles11sp3`, but it can include software to support user logins and a backup DNS server.

All system images are stored in `/var/lib/systemimager/images`.

The preceding output shows the original, factory-shipped system images for the RLCs, the flat compute nodes, and the SGI ICE compute nodes. These original files are as follows:

- `lead-sles11sp3`
- `sles11sp3`
- `ice-sles11sp3`

The output also shows customized images for this SGI ICE cluster. These file names end in `.prod1`, for production use, and are as follows:

- `lead-sles11sp3.prod1`
- `sles11sp3.prod1`. This is the image that resides on the flat compute services node.
- `ice-sles11sp3.prod1`

For each of these images, the associated kernel is `3.0.76-0.11-default`.

The examples in this Array Services configuration procedure add the Array Services information to the customized, production images with the `.prod1` suffix.

5. Decide what kind of security you want to enable.

Array Services includes its own authentication and security, but if your site requires additional security, you can configure MUNGE security, which the installation includes. Your security choices are as follows:

- `munge` on all the nodes you want to include in the array. Configures additional security provided by MUNGE. The installation process installs MUNGE by default. If you decide to use MUNGE, the SGI MPT configuration process explains how to enable MUNGE at the appropriate time.
- `none` on the compute services nodes and `none` on the compute nodes **OR** `noremote` on the compute services nodes and `none` on the the compute nodes. These specifications have the following effects:
  - When you specify `none` on all the nodes you want to include in the array, all authentication is disabled.
  - When you specify `noremote` on the compute services nodes and specify `none` on the compute nodes, users must run their jobs directly from the

compute services nodes. In this case, users cannot submit SGI MPI jobs remotely.

- `simple` (default). Generates hostname/key pairs by using either the OpenSSL, `rand` command, 64-bit values (if available) or by using `$RANDOM` Bash facilities.

6. Proceed to the following:

"Preparing the Images" on page 159

## Preparing the Images

Before you can create images that include Array Services, you need to copy, or *clone*, the production system images your system is using at this time.

The following procedure explains how to prepare the images.

**Procedure 5-2** To prepare the system images

1. Log into the admin node as the root user, and use two `cinstallman` commands, in the following format, to clone (1) one of the images that resides on a compute services node and (2) one of the images that reside on a compute node:

```
cinstallman --create-image --clone --source existing_image --image new_image
```

For *existing\_image*, specify the name of one of the existing images.

For *new\_image*, specify the new name for that to want to give to the image.

For example, the following commands copy the first-generation production images to new, second-generation production images:

```
cinstallman --create-image --clone --source sles11sp3.prod1 --image sles11sp3.prod2
cinstallman --create-image --clone --source ice-sles11sp3.prod1 --image ice-sles11sp3.prod2
```

2. Type the following command to change to the system images directory:

```
cd /var/lib/systemimager/images
```

3. (Optional) Use the `cp(1)` command to copy the MUNGE key from the new compute services node image to the new compute node image.

Complete this step if you want to configure the additional security that MUNGE provides.

The MUNGE key resides in `/etc/munge/munge.key` and must be identical on all the nodes that you want to include in the array. The copy command is as follows:

```
cp /var/lib/systemimager/images/new_computeservices_image/etc/munge/munge.key \
/var/lib/systemimager/images/new_compute_image/etc/munge/munge.key
```

For `new_computeservices_image`, specify the name of the new compute services node image you created.

For `new_compute_image`, specify the name of the new compute node image you created.

For example:

```
cp /var/lib/systemimager/images/sles11sp3.prod2/etc/munge/munge.key \
/var/lib/systemimager/images/ice-sles11sp3.prod2/etc/munge/munge.key
```

---

**Note:** The commands and formats in this step use the backslash (\) continuation character.

---

4. Use the following `cinstallman` command to install the new image on the compute services node:

```
cinstallman --assign-image --node hostname(s) --image new_computeservices_image --kernel kernel
```

For `hostname`, specify the hostname or hostnames of the compute services node that you want users to log into when they log into the array.

For `new_computeservices_image`, specify the name of the new image you created.

For `kernel`, specify the Linux kernel that you want to be associated with this image.

For example, the following command installs the new image on node `n1`:

```
cinstallman --assign-image --node n1 --image sles11sp3.prod2 --kernel 3.0.76-0.11-default
```

5. Use the `ssh(1)` command to log into the compute services node from which you expect users to run SGI MPI programs.

For example, log into `n1`.

6. Use the `arrayconfig(8)` command to configure the compute service node(s) and compute nodes into an array.

The `arrayconfig(8)` command creates the following files on the compute service node to which you are logged in:

- `/etc/array/arrayd.conf`
- `/etc/array/arrayd.auth`

Type the `arrayconfig(8)` command in the following format:

```
/usr/sbin/arrayconfig -a arrayname -f -m -A method node node ...
```

For *arrayname*, specify a name for the array. The default is `default`.

For *method*, specify `munge`, `none`, or `simple`. A later step explains how to specify `noremove` for a compute services node.

For each *node*, specify a list of node IDs.

**Example 1.** To specify that array `myarray` use MUNGE security and include all compute service nodes and all compute nodes, type the following command:

```
/usr/bin/arrayconfig -a myarray -f -m -A munge $(cnodes --compute --ice-compute)
```

**Example 2.** To specify that array `yourarray` use no security, include one compute service node, and include all compute nodes, type the following command:

```
/usr/bin/arrayconfig -a yourarray -f -m -A none n0 $(cnodes --ice-compute)
```

7. Proceed to the following:

"Configuring the Authentication Files in the New System Images on the Admin Node" on page 161

## Configuring the Authentication Files in the New System Images on the Admin Node

Complete one of the following procedures, based upon whether you want to permit remote access to the compute services node:

- If you specified `-A munge` or `-A simple` for authentication

OR

If you specified `-A none` for authentication, and you want to permit users to log into a compute services node remotely to submit MPI programs, proceed to the following:

"Permitting Remote Access to the Compute Services Node" on page 162

- If you specified `-A none` for authentication, and you want to prevent users from logging into a compute services node remotely to submit MPI programs, proceed to the following:

"Preventing Remote Access to the Service Node" on page 163

### Permitting Remote Access to the Compute Services Node

The following procedure assumes that you want to permit job queries and commands on the compute services node. It explains how to copy the array daemon files to the admin node.

**Procedure 5-3** To permit remote access to the compute services node

1. Log into one of the compute services nodes as the root user.
2. Copy the `arrayd.auth` file and the `arrayd.conf` files from the compute services node to the new compute services node image on the admin node.

Type the following command:

```
scp /etc/array/arrayd.* \
admin:/var/lib/systemimager/images/computeservices_image/etc/arrayd.*
```

For *computeservices\_image*, specify the compute services node image on the admin node.

Type this command all on one line. Note that the command in this step uses a backslash (\) character to continue the command to the following line.

For example:

```
scp /etc/array/arrayd.* \
admin:/var/lib/systemimager/images/sles11sp3.prod2/etc/arrayd.*
```

3. Copy the `arrayd.auth` file and the `arrayd.conf` files from the compute services node to the new compute node image on the admin node.

Type the following command:

```
scp /etc/array/arrayd.* \
admin:/var/lib/systemimager/images/compute_image/etc/arrayd.*
```

For *compute\_image*, specify the SGI ICE compute node image on the admin node.

Type this command all on one line. Note that the command in this step uses a backslash (\) character to continue the command to the following line.

For example:

```
scp /etc/array/arrayd.* \
admin:/var/lib/systemimager/images/ice-sles11sp3.prod2/etc/arrayd.*
```

4. Proceed to the following:

"Distributing Images to all the Nodes in the Array" on page 164

### Preventing Remote Access to the Service Node

You can prevent a compute services node from receiving any requests from other computers on the network. In this case, the compute services node can send requests to all remote nodes, but it does not listen on TCP port 5434 for any incoming requests. Complete the procedure in this topic if this is the behavior your site requires.

The following procedure explains how to configure the Array Services files to prevent remote access and how to copy the array daemon files to the admin node.

**Procedure 5-4** To prevent remote access to the compute services node

1. Log into one of the compute services nodes as the root user.
2. Open the following file with a text editor:

```
/etc/array/arrayd.auth
```

3. Type the following, all on one line:

```
AUTHENTICATION NOREMOTE
```

4. Save and close the file.

The file should contain only the one line.

5. Type the following command to copy `/etc/array/arrayd.auth` and `/etc/array/arrayd.conf` from the compute services node to the new compute services node image on the admin node:

```
scp /etc/array/arrayd.* \
admin:/var/lib/systemimager/images/computeservices_image/etc/arrayd.*
```

For example:

```
scp /etc/array/arrayd.* \
admin:/var/lib/systemimager/images/sles11sp3.prod2/etc/arrayd.*
```

6. Log into the admin node as the root user.
7. Create file  
`/var/lib/systemimager/images/compute_image/etc/array/arrayd.auth.`

For example:

```
vi /var/lib/systemimager/images/sles11sp3.prod2/etc/array/arrayd.auth
```

8. Add the following all on one line:

```
AUTHENTICATION NONE
```

9. Save and close the file.

The file should contain only the one line.

10. Type the following command to copy the `/etc/array/arrayd.conf` file to the SGI ICE compute nodes:

```
scp /etc/array/arrayd.conf \
admin:/var/lib/systemimager/images/compute_image/etc/arrayd.conf
```

For example:

```
scp /etc/array/arrayd.conf \
admin:/var/lib/systemimager/images/ice-sles11sp3.prod2/etc/arrayd.conf
```

11. Proceed to the following:

"Distributing Images to all the Nodes in the Array" on page 164

## Distributing Images to all the Nodes in the Array

The following procedure explains how to assign the new compute services node image and the new compute node image to the compute services nodes and to the compute nodes.

**Procedure 5-5** To assign images

1. Log into the admin node as the root user.

2. Use one or more `cinstallman` commands in the following format to assign the new compute services node image to the compute service nodes:

```
cinstallman --assign-image --node hostname --image new_computeservice_image --kernel kernel
```

For *hostname*, specify the hostname for one or more of the compute services nodes. In the cluster definition file, this is the name that appears in the `hostname1` field.

If you want to specify more than one hostname, and your nodes are similarly named, use the `*` wildcard character to represent a string of identical characters in this field. For example, if your hostnames are `n1`, `n2`, `n3`, and `n57`, specify `n*` in this field if you want to specify all compute services nodes.

For *new\_service\_image*, specify the name of the new compute services node image you created.

For *kernel*, specify the Linux kernel that you want to be associated with this image.

Example 1. The following command assigns the new compute services node image to all compute services nodes:

```
cinstallman --assign-image --node n* --image sles11sp3.prod2 --kernel 3.0.76-0.11-default
```

Example 2. The following command assigns the new compute services node image to `n101`:

```
cinstallman --assign-image --node n101 --image sles11sp3.prod2 --kernel 3.0.76-0.11-default
```

3. Use one or more `cinstallman` commands in the following format to assign the new image to the compute services nodes the next time you boot the compute services nodes:

```
cinstallman --next boot --image hostname
```

For *hostname*, specify the hostname for one or more of the compute services nodes. In the cluster definition file, this is the name that appears in the `hostname1` field. You can use the `*` wildcard character in this field if specifying more than one hostname.

For example:

```
cinstallman --next-boot image --node n101
```

4. Proceed to the following:

"Power Cycling the Nodes and Pushing Out the New Images" on page 166

## Power Cycling the Nodes and Pushing Out the New Images

The following procedure explains how to install the new images on the compute services nodes and the compute nodes that you want to include in the array.

**Procedure 5-6** To propagate the images

1. Type the following command to reboot the compute services nodes and the compute nodes:

```
cpower node reboot "*"
```

2. Use one or more `cpower` commands in the following format to power off the compute nodes that you want to reimage:

```
cpower node off "hostname"
```

For *hostname*, specify the hostname(s) of the compute nodes.

You can use wildcard characters if you have many compute nodes.

For example, the following command powers off all the compute nodes on an SGI ICE cluster:

```
cpower node off "r*i*n*"
```

Issue as many `cpower` commands as needed.

3. (Conditional) Use the `cimage` command to push the new compute node image out to the compute nodes.

Complete this step only on SGI ICE clusters. You do not need to complete this step on SGI Rackable clusters.

Use the following format:

```
cimage --push-rack new_compute_image rack
```

For *new\_compute\_image*, specify the name of the new compute node image you created.

For *rack*, specify the IDs of the racks in which the compute nodes reside.

For example, the following command pushes the compute node images to all compute nodes in all racks:

```
cimage --push-rack ice-sles11sp3.prod2 r*
```

4. Use one or more `cpower` command in the following format to start the compute nodes:

```
cpower node on "hostname"
```

For *hostname*, specify the hostnames of the compute nodes.

You can use wildcard characters if you have many compute nodes.

For example, the following command powers on all compute nodes:

```
cpower node on "r*i*n*"
```

Issue as many `cpower` commands as needed.

## (Conditional) Enabling the Mellanox OpenFabrics Enterprise Distribution (OFED) for Linux Software

Perform the procedure in this topic if you use the SLES 11 SP3 operating system on one of the following SGI cluster platforms:

- SGI ICE XA
- SGI Rackable with Mellanox Connect-IB adaptor cards

On other platforms, this procedure is optional because the OFED software in the operating system includes the functionality that is found in the Mellanox OFED 2.4-1.0.0 software package.

The following procedure explains how to enable Mellanox's OFED 2.4-1.0.0 software package and remove the OFED software included by default in the operating system software on your cluster.

**Procedure 5-7** To enable Mellanox OFED

1. Log into the admin node as the root user.
2. Log into SGI Supportfolio.
3. Use the `scp(1)` command to copy the appropriate Mellanox software packages to the `/tmp` directory on the admin node.

You can find information about the Mellanox software package at the following website:

[https://support.sgi.com/content\\_request/521490/index.html](https://support.sgi.com/content_request/521490/index.html)

The packages are in `tar(1)` format. Download the following files:

- The RPM list files:

```
sms-mlnxofed-24104-rpmlist-files.tgz
```

- An operating-system-specific package. Download the package that corresponds to the operating system on the compute node or SGI ICE compute node of your cluster. For example:

```
mlnx-ofed-2.4-1.0.0-rhel71-rpms.tgz
mlnx-ofed-2.4-1.0.0-rhel66-rpms.tgz
mlnx-ofed-2.4-1.0.0-rhel65-rpms.tgz
mlnx-ofed-2.4-1.0.0-sles12-rpms.tgz
mlnx-ofed-2.4-1.0.0-sles11sp3-rpms.tgz
```

4. Use the `mkdir(1)` command in the following format to create a repository on the admin node:

```
mkdir -p /tftpboot/mlnx
```

5. Type the following command to unpack the `tar(1)` file to the directory you created:

```
tar vxzf mlnx-ofed-2.4-1.0.0-os_version-rpms.tgz -C /tftpboot/mlnx
```

For *os\_version*, specify `rhel71`, `rhel66`, `sles12`, `sles11sp3`, or similar, to match what you created with the `mkdir(1)` command.

6. Use the `crepo` command in the following format to create a repository for the operating-system-specific package:

```
crepo --add /tftpboot/mlnx/dir --custom MLNX_OFED-2.4-1.0.0-distro
```

For *dir*, specify one of the following:

```
mlnx-ofed-2.4-1.0.0-rhel7.1
mlnx-ofed-2.4-1.0.0-rhel6.6
mlnx-ofed-2.4-1.0.0-sles11sp3
mlnx-ofed-2.4-1.0.0-sles12
```

For *distro*, specify `rhel71`, `rhel66`, `sles12`, or `sles11sp3`.

7. Unpack the RPM list file.

For example:

```
tar vxzf tempo-mlnxofed-24104-rpmlist-files.tgz -C /opt/sgi/share/rpmlists
service-mlnxofed-24104-rhel71.rpmlist
compute-mlnxofed-24104-rhel71.rpmlist
leader-mlnxofed-24104-rhel71.rpmlist
service-mlnxofed-24104-rhel66.rpmlist
compute-mlnxofed-24104-rhel66.rpmlist
leader-mlnxofed-24104-rhel66.rpmlist
service-mlnxofed-24104-sles12.rpmlist
compute-mlnxofed-24104-sles12.rpmlist
leader-mlnxofed-24104-sles12.rpmlist
service-mlnxofed-24104-sles11sp3.rpmlist
compute-mlnxofed-24104-sles11sp3.rpmlist
leader-mlnxofed-24104-sles11sp3.rpmlist
```

8. Use the `cinstallman` command to clone the images as appropriate for your operating system.

The `cinstallman` commands can take a few minutes to run.

Example 1. Type the following commands to create RHEL 6.6 images:

```
cinstallman --create-image --clone --source lead-rhel6.6 --image lead-rhel6.6.mlnx24104
cinstallman --create-image --clone --source ice-rhel6.6 --image ice-rhel6.6.mlnx24104
cinstallman --create-image --clone --source rhel6.6 --image rhel6.6.mlnx24104
```

Example 2. Type the following commands to create SLES 11 SP3 images:

```
cinstallman --create-image --clone --source ice-sles11sp3 --image ice-sles11sp3.mlnx24104
cinstallman --create-image --clone --source sles11sp3 --image sles11sp3.mlnx24104
cinstallman --create-image --clone --source lead-sles11sp3 --image lead-sles11sp3.mlnx24104
```

9. Unselect the distro repository.

This step ensures that you avoid naming conflicts when you add the new software.

Example 1. On RHEL 6.6, type the following command:

```
crepo --unselect Red-Hat-Enterprise-Linux-6.6
```

Example 2. On SLES 11 SP3, type the following command:

```
crepo --unselect SUSE-Linux-Enterprise-Server-11-SP3
```

10. Select the newly created `MLNX_OFED` repository.

Example 1. On RHEL 6.6, type the following command:

```
crepo --select MLNX_OFED-2.4-1.0.0-rhel6.6
```

Example 2. On SLES 11 SP3, type the following command:

```
crepo --select MLNX_OFED-2.4-1.0.0-sles11sp3
```

11. (Conditional) Remove unneeded packages.

Complete this step only if your cluster runs the RHEL 6.6 operating system on its RLCs or the SLES 11 SP3 operating system on its RLCs.

The `cinstallman` command that removes packages from the RLC image is as follows:

```
cinstallman --yum-image --image RLC_image_name remove pkg
```

The previous command includes the following variables:

|                   |                                                                            |
|-------------------|----------------------------------------------------------------------------|
| <i>image_name</i> | The name of one of the RLC images. For example:<br>lead-rhel6.6.mlnx24104. |
| <i>pkg</i>        | The name of the package you want to remove.                                |

Remove the following packages:

- Remove `libmlx5.x86_64` from the RHEL 6.6 RLC image.

For example:

```
cinstallman --yum-image --image lead-rhel6.6.mlnx24104 remove libmlx5.x86_64
```

- Remove `dapl` from the RHEL 6.6 RLC image and the SLES 11 SP3 RLC image.

For example:

```
cinstallman --yum-image --image lead-rhel6.6.mlnx24104 remove dapl
```

12. Use the `cinstallman` command to remove the operating system's OFED package from each of the newly created images.

The format of the `cinstallman` command that you need to use is as follows:

```
cinstallman --yum-image --image node_image remove pkg_name
```

The previous command includes the following variables:

*node\_image* Specify one of the node names from the *node\_image* column of the following table.

*pkg\_name*

Specify one of the package names from the *pkg\_name* column of the following table.

| Operating System | Operating System Image | <i>node_image</i>        | <i>pkg_name</i>         |
|------------------|------------------------|--------------------------|-------------------------|
| RHEL 7.1         | RLC                    | lead-rhel7.1.mlnx24104   | sgi-lead-distro-ofed    |
|                  | SGI ICE compute node   | ice-rhel7.1.mlnx24104    | sgi-compute-distro-ofed |
|                  | Flat compute node      | rhel7.1.mlnx24104        | sgi-service-distro-ofed |
| RHEL 6.6         | RLC                    | lead-rhel6.6.mlnx24104   | sgi-lead-distro-ofed    |
|                  | SGI ICE compute node   | ice-rhel6.6.mlnx24104    | sgi-compute-distro-ofed |
|                  | Flat compute node      | rhel6.6.mlnx24104        | sgi-service-distro-ofed |
| SLES 12          | RLC                    | lead-sles12.mlnx24104    | sgi-lead-distro-ofed    |
|                  | SGI ICE compute node   | ice-sles12.mlnx24104     | sgi-compute-distro-ofed |
|                  | Flat compute node      | sles12.mlnx24104         | sgi-service-distro-ofed |
| SLES 11 SP3      | RLC                    | lead-sles11sp3.mlnx24104 | sgi-lead-distro-ofed    |
|                  | SGI ICE compute node   | ice-sles11sp3.mlnx24104  | sgi-compute-distro-ofed |
|                  | Flat compute node      | sles11sp3.mlnx24104      | sgi-service-distro-ofed |

For example, type the following commands to remove the SLES 11 SP3 distro's OFED package from the existing images:

```
cinstallman --yum-image --image sles11sp3.mlnx24104 remove sgi-service-distro-ofed
cinstallman --yum-image --image ice-sles11sp3.mlnx24104 remove sgi-compute-distro-ofed
cinstallman --yum-image --image lead-sles11sp3.mlnx24104 remove sgi-lead-distro-ofed
```

- Using the appropriate \*.rpm<sub>list</sub> files, install the Mellanox OFED packages into the newly created images.

Earlier in this procedure, you used the `tar(1)` command to unpack the compressed file into the repository you created. Within that directory, the files are as follows:

- For RHEL, the files you need to install into the images are as follows:

```
compute-mlnxofed-24104-rhel6.6.rpmlist
leader-mlnxofed-24104-rhel6.6.rpmlist
```

```
service-mlnxofed-24104-rhel6.6.rpm
```

- For SLES, the files you need to install into the images are as follows:

```
compute-mlnxofed-24104-sles11sp3.rpm
leader-mlnxofed-24104-sles11sp3.rpm
service-mlnxofed-24104-sles11sp3.rpm
```

The following format shows the `cinstallman` commands you can use to install the Mellanox OFED package into the default RHEL images:

```
cinstallman --refresh-image --image leader_image_name \
--rpm /opt/sgi/share/rpmlists/mlnx-ofed/lead-mlnxofed-24104-rhel7.1.rpm
cinstallman --refresh-image --image ice_compute_image_name --rpm \
/opt/sgi/share/rpmlists/mlnx-ofed/compute-mlnxofed-24104-rhel7.1.rpm
cinstallman --refresh-image --image flat_compute_image_name --rpm \
/opt/sgi/share/rpmlists/mlnx-ofed/service-mlnxofed-24104-rhel7.1.rpm
```

The following format shows the `cinstallman` commands you can use to install the Mellanox OFED package into the default SLES images:

```
cinstallman --refresh-image --image leader_image_name \
--rpm /opt/sgi/share/rpmlists/mlnx-ofed/lead-mlnxofed-24104-sles12.rpm
cinstallman --refresh-image --image ice_compute_image_name \
--rpm /opt/sgi/share/rpmlists/mlnx-ofed/compute-mlnxofed-24104-sles12.rpm
cinstallman --refresh-image --image flat_compute_image_name \
--rpm /opt/sgi/share/rpmlists/mlnx-ofed/service-mlnxofed-24104-sles12.rpm
```

14. Use the `crepo` command to restore the repository selections.

Complete the following steps:

- Use the `crepo --unselect` command to unselect the Mellanox OFED repository.

The format of this command is as follows:

```
crepo --unselect m_ofed
```

For *m\_ofed*, specify the name of the Mellanox OFED software package for your operating system.

For example:

```
crepo --unselect MLNX_OFED-2.4-1.0.4-rhel6.6
Unselecting: MLNX_OFED-2.4-1.0.4-rhel6.6
```

- Use the `crepo --select` command to select the operating system distribution's repository.

The format of this command is as follows:

```
crepo --select os_distro
```

For *os\_distro*, specify the name of the operating system distribution.

For example:

```
crepo --select Red-Hat-Enterprise-Linux-6.6
Selecting: Red-Hat-Enterprise-Linux-6.6
Updating: /etc/opt/sgi/rpmlists/generated-rhel6.6.rpmlist
Updating: /etc/opt/sgi/rpmlists/generated-ice-rhel6.6.rpmlist
Updating: /etc/opt/sgi/rpmlists/generated-lead-rhel6.6.rpmlist
Updating: /etc/opt/sgi/rpmlists/generated-admin-rhel6.6.rpmlist
```

15. Use the `cimage` command with its `--push-rack` argument to put the new image on the RLCs.

For example:

```
cimage --push-rack ice-sles11sp3.mlnx24104 r*
```

16. Use the `cimage` command with its `--set` argument to assign the image to the compute nodes.

For example:

```
cimage --set ice-sles11sp3.mlnx24104 3.0.101-0.40-default r*
```

17. Use the `cimage` command to assign the new images to the nodes.

For example, to assign the new images to the `r1lead` RLC node and flat compute node `service0` for SLES 11 SP3, type the following:

```
cinstallman --assign-image --image lead-sles11sp3.mlnx24104 --kernel kernel_version --node r1lead
cinstallman --assign-image --image sles11sp3.mlnx24104 --kernel kernel_version --node service0
```

For *kernel\_version*, specify one of the kernels shown for this image in the output from the `cinstallman --show-images` command on your cluster.

18. Use the `cinstallman` command to reimage the nodes with the new images that contain the upgraded Mellanox software.

For example, to reimage the `r1lead` RLC node and flat compute node `service0`, type the following:

```
cinstallman --next-boot image --node r1lead
cinstallman --next-boot image --node service0
```

19. Boot the nodes affected by this change.

For example, type the following command to boot flat compute node `service0`.

```
ssh service0
reboot
```

## Troubleshooting Configuration Changes

If a configuration change does not affect the cluster in the intended manner, try one of the following approaches:

- Edit the node image on the admin node. For example, you can reconfigure the image for the compute nodes that you use for user services on the admin node and reimage the compute services nodes with that new image.
- Edit the node customization scripts. For example, the SGI ICE compute node update scripts reside on the admin node in the `/opt/sgi/share/per-host-customization/global` directory.

## Troubleshooting

This chapter covers the following topics:

- "About Troubleshooting" on page 176
- "Using the `switchconfig` Command" on page 176
- "SGI ICE Compute Nodes Are Taking Too Long To Boot (SGI ICE Clusters Only)" on page 178
- "Verify the Bonding Mode on the Rack Leader Controller (RLC) (SGI ICE Clusters Only)" on page 179
- "`cimage --push-rack` Pushes Too Many (or Too Few) Expansions (SGI ICE Clusters Only)" on page 182
- "Cannot ping the CMCs from the Rack Leader Controller (RLC) (SGI ICE Clusters Only)" on page 182
- "Restarting the `blademon` Daemon (SGI ICE Clusters Only)" on page 185
- "Log Files" on page 185
- "CMC `slot_map` / `blademon` Debugging Hints" on page 186
- "Resolving CMC Slot Map Ordering Problems" on page 187
- "In `tmpfs` Mode, File Has Date in the Future Warnings" on page 189
- "Ensuring Hardware Clock Has the Correct Time" on page 189
- "Troubleshooting a Rack Leader Controller (RLC) With Misconfigured Switch Information (SGI ICE Clusters Only)" on page 190
- "Switch Wiring Rules" on page 192
- "Admin Node `eth2` Link in the Bond is Down" on page 193
- "Installing SGI Tempo Versions Older than 2.9.0" on page 193
- "Booting Nodes With iPXE After an Upgrade" on page 194

## About Troubleshooting

This chapter provides answers to some common problems users encounter when installing or upgrading a cluster. It includes diagnosis and troubleshooting information.

## Using the `switchconfig` Command

The `switchconfig` command displays switch settings and enables you to configure switches.

To retrieve help output online, which includes examples, type the following:

```
switchconfig set --help
```

Unless you want to retrieve help output, the following parameters are required:

- `--default-vlan` | `-d` *default\_vlan*
- `--switches` | `-s` *hostname\_or\_IP*[, *hostname\_or\_IP*][, ...]
- `--macs` | `m` *mac\_addr*[, *mac\_addr*][, ...] or `--ports` | `-p` *port\_num*[, *port\_num*][, ...]

The `switchconfig` command's subcommands are as follows:

| Subcommand             | Action                                                                   |
|------------------------|--------------------------------------------------------------------------|
| <code>list</code>      | Prints the current settings.                                             |
| <code>set</code>       | Assigns VLAN(s) and various settings for one or more MAC addresses.      |
| <code>unset</code>     | Returns associated port(s) to default settings.                          |
| <code>list_ip</code>   | Displays the IP address configured for the switch.                       |
| <code>set_ip</code>    | Adds an IP for the VLAN on the switch in order to route the MIC traffic. |
| <code>unset_ip</code>  | Removes the IP for the VLAN on the switch.                               |
| <code>list_ospf</code> | Displays OSPF configured for the switch.                                 |
| <code>set_ospf</code>  | Sets OSPF in order to route the MIC traffic between the switches.        |

|                                     |                                                                                                                                |
|-------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|
| <code>unset_ospf</code>             | Disables OSPF and all of the network statements that go along with it.                                                         |
| <code>list_mtu</code>               | Displays the assigned MTU value for all ports in the switch stack.                                                             |
| <code>set_mtu</code>                | Assigns all ports in the switch stacks to MTU value.                                                                           |
| <code>list_default_gateway</code>   | Displays the default gateway assigned to the switch network.                                                                   |
| <code>set_default_gateway</code>    | Sets default gateway for the switch network. Usually, the gateway is a service node with a connection to the customer network. |
| <code>unset_default_gateway</code>  | Unsets the default gateway for the switch network.                                                                             |
| <code>list_route</code>             | Displays the routing table.                                                                                                    |
| <code>set_snmp_community</code>     | Sets SNMP community settings.                                                                                                  |
| <code>list_snmp_community</code>    | Lists SNMP community settings.                                                                                                 |
| <code>save_running_config</code>    | Saves current running configuration to switch NV memory.                                                                       |
| <code>reset_factory_defaults</code> | Sets the switch back to the factory default configuration and reboots.                                                         |
| <code>sanity_check</code>           | Runs a test to sanity check the switch configuration.                                                                          |
| <code>change_password</code>        | Changes the password for the admin account on the switch.                                                                      |
| <code>push_switch_config</code>     | Pushes the start-up switch configuration file to the switch specified from the TFTP server and loads the switch.               |
| <code>pull_switch_config</code>     | Pulls the switch configuration file from the switch and saves to the TFTP server.                                              |

find

Takes one or more MAC addresses and maps them to the management switch that it is physically connected to using FDB/LLDP.

## SGI ICE Compute Nodes Are Taking Too Long To Boot (SGI ICE Clusters Only)

If the SGI ICE compute nodes on an SGI ICE cluster are taking a long time to boot, perform the following:

- See "Verify the Bonding Mode on the Rack Leader Controller (RLC) (SGI ICE Clusters Only)" on page 179 to verify that the compute nodes have the proper bonding setup.
- Verify the rack leader controller (RLC) has a MegaRAID controller. 144 nodes do not boot well with 106x controllers, for example. You can verify this with `lspci` command.

To verify the MegaRAID battery is working and charged, perform the following:

```
/opt/MegaRAID/MegaCli/MegaCli64 -ShowSummary -a0
```

You should see 'Status : Healthy' under 'BBU' (BBU = Battery Backup Unit).

---

**Note:** If this is the first time the node has booted up, it takes several hours for the BBU to be charged.

---

- Verify cache is set to write-back, as follows:

```
/opt/MegaRAID/MegaCli/MegaCli64 -LDGetProp -Cache -LALL -a0
```

---

**Note:** Never force write-back on if bad BBU (`-CachedBadBBU`) as data loss happens with an orderly shutdown that includes a power off.

---

When you see the output: `Cache Policy:WriteBack`, write-back is enabled.

To enable the write-back policy, perform the following:

```
/opt/MegaRAID/MegaCli/MegaCli64 -LDSetProp -NoCachedBadBBU -LALL -a0
```

## Verify the Bonding Mode on the Rack Leader Controller (RLC) (SGI ICE Clusters Only)

The redundant management network (RMN) is configured by default. To verify the bonding mode, log into an RLC and type the following command:

```
rlllead:~ # cat /proc/net/bonding/bond0
Ethernet Channel Bonding Driver: v3.5.0 (November 4, 2008)
```

```
Bonding Mode: IEEE 802.3ad Dynamic link aggregation
Transmit Hash Policy: layer2+3 (2)
MII Status: up
MII Polling Interval (ms): 100
Up Delay (ms): 0
Down Delay (ms): 0
```

```
802.3ad info
LACP rate: slow
Aggregator selection policy (ad_select): stable
Active Aggregator Info:
 Aggregator ID: 1
 Number of ports: 2
 Actor Key: 17
 Partner Key: 4
 Partner Mac Address: b4:0e:dc:37:4f:a7
```

```
Slave Interface: eth0
MII Status: up
Link Failure Count: 1
Permanent HW addr: 00:25:90:38:e5:22
Aggregator ID: 1
```

```
Slave Interface: eth1
MII Status: up
Link Failure Count: 0
Permanent HW addr: 00:25:90:38:e5:23
Aggregator ID: 1
```

By default, RLCs use link aggregation control protocol (LACP) (IEEE 802.3ad Dynamic link aggregation) as the bonding mode. Alternatively, you can set the bonding mode to active-backup, which means that the redundant management

network is (potentially) disabled and only one link is active. For more information, see the following:

<https://www.kernel.org/doc/Documentation/networking/bonding.txt>

Use the cluster configuration tool's **Configure Redundant Management Network** option to turn on the redundant management network (RMN) system for nodes being discovered going forward.

Set the redundant management networking mode on, as follows:

```
cadadmin --enable-redundant-mgmt-network --node r1lead
```

Set the bonding mode per node, as follows:

```
cadadmin --set-mgmt-bonding --node r1lead 802.3ad
```

You need to reboot the system.

The `/proc/net/bonding/bond0` file, should show the bonding mode with link aggregation configured, as follows:

```
Bonding Mode: IEEE 802.3ad Dynamic link aggregation
```

The number of ports should be the following:

```
Number of ports: 2
```

2 is the correct value for an RMN configuration. If the number is 1, it mean the trunk has not formed. The likely causes for this are, as follows:

- The Ethernet cable is not connected to top level switch. From the RLC, use the `/sbin/ethtool` on `eth0` and `eth1` to verify the link is present, as follows:

```
r1lead:~ # /sbin/ethtool eth0
Settings for eth0:
 Supported ports: [TP]
 Supported link modes: 10baseT/Half 10baseT/Full
 100baseT/Half 100baseT/Full
 1000baseT/Full
 Supports auto-negotiation: Yes
 Advertised link modes: 10baseT/Half 10baseT/Full
 100baseT/Half 100baseT/Full
 1000baseT/Full
```

```

Advertised auto-negotiation: Yes
Speed: 1000Mb/s
Duplex: Full
Port: Twisted Pair
PHYAD: 1
Transceiver: internal
Auto-negotiation: on
Supports Wake-on: umbg
Wake-on: g
Current message level: 0x00000003 (3)
Link detected: yes

```

- The Ethernet cable is connected, but linking is wrong. When the `/sbin/ethtool` command output shows the link speed as 100 Mb/s due to a bad cable the trunk leg is rejected.
- The top level Ethernet switch is misconfigured: Perhaps the `switchconfig` tool did not get this port configured properly. You can either log in to the switch to try to diagnose, or try the following procedure:

1. Find the MAC address of the `r1lead` bond interface, as follows:

```

r1lead:~ # ip addr show bond0
6: bond0: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP
 link/ether 00:30:48:34:2b:5d brd ff:ff:ff:ff:ff:ff
 inet 172.23.0.7/16 brd 172.23.255.255 scope global bond0
 inet 172.24.0.1/16 brd 172.24.255.255 scope global bond0:bmc
 inet 10.42.0.1/24 brd 10.42.0.255 scope global bond0:discover
 inet6 fe80::225:90ff:fe38:e522/64 scope link
 valid_lft forever preferred_lft forever

```

2. From the admin node, run the `switchconfig list --switches mgmtsw0` command to list the MAC addresses trunks from the switches, as follows:

```

sys-admin:~ # switchconfig list --switches mgmtsw0
Current MAC/port configuration:

Switch Identifier: mgmtsw0 IP Address: 172.23.0.6

 MAC Port Trunk default-VLAN allowed-VLANs

00-25-90-3F-16-C4 1/6 1 1 1(u)
00-30-48-F7-84-65 1/48 1 1 1(u)

```

|                   |     |     |     |              |
|-------------------|-----|-----|-----|--------------|
| 00-25-90-38-E5-22 | 1/5 | 1   | 1   | 1(u), 101(t) |
| 00-25-90-38-E5-23 | 1/5 | 1   | 1   | 1(u), 101(t) |
| 00-25-90-38-E5-22 | 1/5 | 1   | 101 | 1(u), 101(t) |
| 00-25-90-38-85-BC | 1/7 | 2   | 1   | 1(u)         |
| 00-25-90-38-85-BD | 1/7 | 2   | 1   | 1(u)         |
|                   |     | ... |     |              |

If the RLC `r1lead` bond interface MAC address shows up in the `Port` column and not the `Trunk` column, the switch is not configured correctly.

- To properly configure the switch, from the admin node, perform a command similar to the following:

```
switchconfig set -s mgmtsw0 -v num=1 -v num=101,tag=tagged -b lacp -d 1 -m 00:25:90:38:E5:30
```

This replaces 101 with the proper VLAN number. 101 for rack group 1, 102 for rack group 2, and so on.

- ssh onto the `r1lead` and verify that the RLC shows `Number of ports: 2`.

## **cimage --push-rack Pushes Too Many (or Too Few) Expansions (SGI ICE Clusters Only)**

When you perform `cimage --push-rack` (or when `blademon` calls `discovery-rack`), it creates read/write expansions for each compute node.

Use the `configure-cluster` GUI **Configure Default Max Rack IRU Setting** option to set the default number of individual rack units (IRUs) supported by a rack leader controller (RLC). Set this value to the number of CMCs that each RLC serves. The default is 8. When you change the value, it affects only future node discoveries.

You can change the setting per-node with the `cadmin` command, as follows:

```
sys-admin:~ # cadmin --set max-rack-irus --node r1lead 8
```

## **Cannot ping the CMCs from the Rack Leader Controller (RLC) (SGI ICE Clusters Only)**

If this is an RLC with a brand new, never-before-discovered top level switch (or set of switches), the `cmcdetectd` daemon will see CMCs asking for IP addresses on the

HEAD network. It configures the top level switch(es) so that the CMCs are on the appropriate rack VLAN. Make sure `cmcdetectd` is running, restart if needed.

You can diagnose this some by running the `tcpdump` command looking for DHCP requests. The requests should be seen on the RLC and not the admin node. For example, type the following command from `r1lead`:

```
/usr/sbin/tcpdump -i bond0 -s600 -nn -vv -e -t -l -p broadcast and src port 68 and dst port 67
tcpdump: listening on bond0, link-type EN10MB (Ethernet), capture size 600 bytes
00:25:90:3f:16:c4 > ff:ff:ff:ff:ff:ff, ethertype IPv4 (0x0800), length 590: (tos 0x0, ttl 64, id 0, offset 0, \
flags [none], proto UDP (17), length 576) 0.0.0.0.68 > 255.255.255.255.67:
[udp sum ok] BOOTP/DHCP, Request from 00:25:90:3f:16:c4, length 548, xid 0x8b8d332a, Flags [none] (0x0000)
 Client-IP 172.24.0.2
 Client-Ethernet-Address 00:25:90:3f:16:c4
 Vendor-rfc1048 Extensions
 Magic Cookie 0x63825363
 DHCP-Message Option 53, length 1: Request
 ...
```

If the switch was previously discovered but you are reinstalling the system or discovering a new root slot, `cmcdetectd` will not detect any CMC DHCP requests on HEAD. In this case, you need to be sure to run the cluster configuration tool, and set **Configure Switch Management Network** to `yes`. Note that changing `configure-cluster` only takes effect for nodes discovered in the future. If you have an existing RLC already discovered, you will need to run a command like the following:

```
cadmin --enable-switch-mgmt-network --node r1lead
```

After rebooting the RLC, make sure that the `ip addr` command shows `vlan101` as an interface and not `vlan1` or `vlan2` interfaces, as follows:

```
r1lead:~ # ip addr
...
13: vlan101@bond0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state
 link/ether 00:30:48:f2:7e:a2 brd ff:ff:ff:ff:ff:ff
 inet 192.168.160.1/22 brd 192.168.160.255 scope global vlan101
 inet 10.159.0.1/22 brd 10.159.3.255 scope global vlan101:bge
 inet6 fe80::255:90ff:fe38:e522/64 scope link
 valid_lft forever preferred_lft forever
```

Confirm `dhcpcd` is running on the RLC. If `dhcpcd` is not running, CMCs will not get their IP addresses. Check for errors starting `dhcpcd`. If `blademon`d failed to create the

`ice.conf dhcpd` configuration file (`/etc/dhcpd.conf.d`), see "Restarting the `blademon` Daemon (SGI ICE Clusters Only)" on page 185.

Verify proper CMC configuration. The CMC is configured for its rack number and slot number. If they are not configured correctly, multiple CMCs can be configured the same way resulting in problems. This can also result in the `ice.conf dhcp` configuration file being corrupted. You may need a USB serial cable to fix the CMCs if this is the case.

One troubleshooting approach is to run `tcpdump` on the RLC, as follows:

```
/usr/sbin/tcpdump -i bond0 -s600 -nn -vv -e -t -l -p broadcast and src port 68 and dst port 67
```

Watch the DHCP requests over several minutes. If you see the same Client Identifier being requested by more than one MAC address, you are in a situation where the CMCs are not configured correctly.

Verify that the RLC is properly configured in the switch (see "Troubleshooting a Rack Leader Controller (RLC) With Misconfigured Switch Information (SGI ICE Clusters Only)" on page 190).

Confirm the wiring rules. See "Switch Wiring Rules" on page 192.

If you moved some CMCs from one RLC number to another and you already adjusted the rack and slot number in the CMC, The switch likely does not know about the changes. The CMCs are likely in the wrong VLAN, potentially a VLAN that is no longer in use. For example, if you had the CMCs served by the `r3lead` RLC but decided to decommission `r3lead` and move the CMCs to `r1lead` instead this situation could arise. In this case, the switch must be reconfigured. Use the `switchconfig` command to configure the ports connected to those CMCs for head. The admin node `cmcdetectd` daemon will move them to the correct ultimate location.

You need to know the MACs of the CMC embedded Linux for this, so perhaps record this when you change the slot/rack number in the CMC. **Hint:** `dbdump` may still have the information depending on how you removed the RLC.

An example command is, as follows:

```
switchconfig set -v num=1 -b manual -d 1 -m 08:00:69:16:51:49 --switches mgmtsw0
```

If you have more than one management switch, then list them in a comma-separated-list for `--switches`.

In a non-redundant-management configuration (switches not stacked), if the `dhcpd` daemon shows DHCP requests from the CMC but the CMC remains unpingable, it could be that both CMC-0 and CMC-1 are connected and linked. This breaks the wiring rules. When we are **not** wired for redundant management networking, only CMC-0 should be connected.

When **not** wired for redundant management networking (when switches are not stacked), do not connect CMC-1.

## Restarting the `blademon` Daemon (SGI ICE Clusters Only)

From the rack leader controller (RLC), perform the following steps:

1. Stop the daemon:

```
r1lead:~ # service blademon stop
```

2. Remove `/etc/dhcpd.conf.d/ice.conf` or `/etc/dhcp/dhcpd.conf.d/ice.conf`:

```
rm ice.conf dhcpd.conf
```

3. Remove `slot_map`:

```
r1lead:~ # rm /var/opt/sgi/lib/blademon/slot_map
```

4. Start the daemon:

```
r1lead:~ # service blademon start
```

## Log Files

All of the log files reside in the `/var/log` directory. In addition to the `messages` log file and in some cases `dhcpd` file on the rack leader controller (RLC), here are some interesting `/var/log` directory log files:

- `/var/log/discover-rack`

On the admin node, the `discover-rack` call is facilitated by `blademon` when new nodes are found. This log will often show problems with discovering nodes.

- `/var/log/blademon`d

On the RLCs, this shows the `blademon`d daemon actions. This includes showing when blade changes are found and it also shows its call to `discover-rack`, and so on. If there are CMC communication issues, they will often be noticed in this log.

- `/var/log/cmcdetectd.log`

On the admin node, `cmcdetectd` logs its actions as it configures the switches for CMCs in the system. Watch for progress or errors here.

- `/var/log/switchconfig.log`

On the admin node, there is a `switchconfig` command line tool. This tool is largely used by the `discover` command as nodes are discovered. Its actions are logged in to this log file. If RLC VLANs are not functioning properly, check the `switchconfig` log file.

## CMC slot\_map / blademon

 Debugging Hints

This section describes what to do when the `blademon`d daemon cannot find a system blade, as follows:

- Can you ping the CMCs? See "Cannot ping the CMCs from the Rack Leader Controller (RLC) (SGI ICE Clusters Only)" on page 182.
- If the CMCs are pingable, verify that they have a valid slot map. If the slot map returned by the CMC is missing entries, then `blademon`d cannot function properly. It operates on information passed to it by the CMC. Some commands to run from the rack leader controller (RLC) are, as follows:

- Dump the slot map from each CMC to your screen:

```
r1lead:~ # /opt/sgi/lib/dump-cmc-slot-tables
```

- Query an individual slot map:

```
r1lead:~ # echo STATUS | netcat r1i0c 4502
```

---

**Note:** In some software distributions, `netcat` is `nc`.

---

If the CMCs are pingable and the CMCs have valid slot maps, then you can focus on how `blademon`d is functioning.

You can turn on debug mode in the `blademon`d daemon by sending it a `SIGUSR1` signal from the RLC, as follows:

```
kill -USR1 pid
```

To turn debug mode off, send it another `SIGUSR1` signal. You should see a message in the `blademon`d log about debug mode being enabled or disabled.

The `blademon`d daemon maintains the slot map at `/var/opt/sgi/lib/blademon`d/slot\_map on the RLCs. This appears as `/var/opt/sgi/lib/blademon`d/slot\_map.*rack\_number* on the admin node.

For a `blademon`d `--help` statement, `ssh` onto the `r1lead` RLC, as follows:

```
[root@admin ~]# ssh r1lead
Last login: Tue Jan 17 13:21:34 2012 from admin
[root@r1lead ~]#
[root@r1lead ~]# /opt/sgi/lib/blademon --help
Usage: blademon [OPTION] ...
```

Discover CMCs and blades managed by CMCs.

Note: This daemon normally takes no arguments.

```
--help Print this usage and exit.
--debug Enable debug mode (also can be enabled by setting CM_DEBUG)
--fakecmc Development only: Discover fake CMCs instead of real ones
--scan-once Initialize, scan for blades, set blades up. Do not daemonize.
 Do not keep looping - do one pass and exit.
```

## Resolving CMC Slot Map Ordering Problems

If there are `ssh(1)` key failures or if the compute node hosts seem to be BMCs, it is possible that there are problems with the CMC slot map might be corrupted.

The CMC maintains a cache file that records which MACs are BMC MACs and which are host MACs. It uses this information, combined with switch port location information in the embedded Broadcom switch, to generate the slot map used by the `blademon`d daemon.

In certain situations, for example, a CMC reflash, might remove the cache file but leave CMC power active. In this situation, the CMC does not know which MACs on a given embedded switch port are host and which are BMC and gets the order randomly incorrect. It then caches the incorrect order. To fix this for each CMC, turn the power off with `pfctl`, zero out the MAC cache file, and reset each CMC. Then have `blademon` start over from scratch (see "Restarting the `blademon` Daemon (SGI ICE Clusters Only)" on page 185).

The following procedure explains these steps.

**Procedure 6-1** To fix a corrupted CMC slot map

1. `ssh` as root to the rack leader controller (RLC), as follows:

```
sys-admin:~ # ssh rlead
Last login: Thu Jan 26 13:57:53 2012 from admin
rlead:~ #
```

2. Disable the `blademon` daemon.

On RHEL 7.1 systems and SLES 12 systems, type the following command:

```
rlead:~ # systemctl stop blademon
```

On RHEL 6.x systems and SLES 11 SPx system, type the following command:

```
rlead:~ # service blademon stop
```

3. Turn off IRU power for each CMC using the `PDSH` command, as follows:

```
PDSH_SSH_ARGS_APPEND="-F /root/.ssh/cmc_config" pdsh -g cmc pfctl off
```

4. Zero out the slot map cache file, as follows:

```
PDSH_SSH_ARGS_APPEND="-F /root/.ssh/cmc_config" pdsh -g cmc cp /dev/null /work/net/broadcom_mac_addr_cache
```

5. Reboot the CMC, as follows:

```
PDSH_SSH_ARGS_APPEND="-F /root/.ssh/cmc_config" pdsh -g cmc reboot
```

6. Restart `blademon` from scratch, see "Restarting the `blademon` Daemon (SGI ICE Clusters Only)" on page 185.

## In `tmpfs` Mode, File Has Date in the Future Warnings

If you boot a compute node with `tmpfs`, part of the process transfers a root tarball using multicast. This tarball is then expanded. If you see hundreds of "file X has a time in the future" messages, it likely means your hardware clock is not set to system time properly (see "Ensuring Hardware Clock Has the Correct Time" on page 189).

## Ensuring Hardware Clock Has the Correct Time

Some software distributions do not synchronize the system time to the hardware clock as expected. As a result, the hardware clock may not get synchronized with the system time as it should. At shut down, the system time is copied to the hardware clock, but sometimes this does not happen.

To set all the compute node hardware clocks up properly, perform the following:

- Make sure the admin node and rack leader controller (RLC) have the correct time
- Make sure the admin node and RLCs are synchronized with `ntp`. An admin node can show a message like the following:

```
ntp[20489]: synchronized to 128.162.244.1, stratum 2
```

- An RLC might show a message like the following:

```
20 Jan 22:54:14 ntpd[16831]: synchronized to 172.23.0.1, stratum 3
```

- Make sure the compute nodes have the correct time. They use `ntp` broadcast packets but still will display this:

```
20 Jan 23:05:16 ntpd[4925]: synchronized to 192.168.159.1, stratum 4
```

You can also use a command like the following and view the output:

```
sys-admin:~ # pdsh -g leader pdsh -g compute date
```

- Issue the following command to set the hardware clock to the system clock, as follows:

```
sys-admin:~ # pdsh -g leader pdsh -g hwclock --systohc
```

- You can run the `hwclock` without options to confirm the current hardware clock time, as follows:

```
sys-admin:~ # hwclock
Thu 26 Jan 2012 10:57:27 PM CST -0.750431 seconds
```

## Troubleshooting a Rack Leader Controller (RLC) With Misconfigured Switch Information (SGI ICE Clusters Only)

Normally, as you discover RLCs, `switchconfig` is called automatically and the switch ports associated with the RLC are configured in the special way needed for RLCs, as follows:

- Default VLAN 1
- Accept rack VLAN packets tagged (rack 1 vlan is vlan101)
- Link Aggregation is the bonding mode between the two ports associated with the RLC

If an RLC is moved in the switch or if `switchconfig` failed during discovery for some reason, you can run `switchconfig` by hand to configure the switch, as follows:

1. Certain switch wires rules must be followed in switch configuration, see "Switch Wiring Rules" on page 192.
2. Make sure all management switches are reachable from the admin node.
3. Find the MAC addresses associated with the RLC interfaces. You can do this by running the following command on the RLC in question:

```
rllead:~ # cat /proc/net/bonding/bond0
Ethernet Channel Bonding Driver: v3.5.0 (November 4, 2008)
```

```
Bonding Mode: IEEE 802.3ad Dynamic link aggregation
Transmit Hash Policy: layer2+3 (2)
MII Status: up
MII Polling Interval (ms): 100
Up Delay (ms): 0
Down Delay (ms): 0
```

```
802.3ad info
LACP rate: slow
Aggregator selection policy (ad_select): stable
```

```
Active Aggregator Info:
 Aggregator ID: 1
 Number of ports: 2
 Actor Key: 17
 Partner Key: 4
 Partner Mac Address: b4:0e:dc:37:4f:a7
```

```
Slave Interface: eth0
MII Status: up
Link Failure Count: 0
Permanent HW addr: 00:25:90:38:e5:22
Aggregator ID: 1
```

```
Slave Interface: eth1
MII Status: up
Link Failure Count: 0
Permanent HW addr: 00:25:90:38:e5:23
Aggregator ID: 1
```



**Caution:** Because bonded interfaces are in play, you cannot get both MAC addresses from using the `ip` command. The `ip` command shows the same MAC address for `eth0` and `eth1` if redundant management networking is enabled.

---

4. Determine which management switches are present, as follows:

```
r1lead:~ # cnodes -mgmtsw
mgmtsw0
```

5. When you have the list of management switches and the MAC addresses of the RLCs, run a command similar to the following:

```
switchconfig set --vlan num=1 --vlan num=101,tag=tagged --bonding=802.3ad --default-vlan 1 /
--macs 00:e0:ed:0a:f2:0d,00:e0:ed:0a:f2:0e --switches mgmtsw0,mgmtsw
```

This replaces the MACs and management switches with the proper ones. It replaces the 101 with the VLAN for the rack, normally "100 + rack number" so rack 1 is 101, rack 2 102.

## Switch Wiring Rules

This section is mainly of interest to clusters that have a redundant management network setup (stacked pairs of switches) or larger systems that have switch stacks cascaded from the top level switch.

When discovering cascaded switches, it is impossible to know the connected switch ports of all trunks in advance. So when discovering cascaded switches, you can only start with one cable for discovering, then add the second one later on.

When trunks are configured, it is often hard to find the MAC address of both legs of the trunk. This is because the trunked connection just uses one MAC for the connection. Therefore, you need to rely on rules that infer the second port's connection based on the first port.

Some simple wiring rules are, as follows:

- In a redundant management network (RMN) configuration, when connecting admin nodes, rack leader controllers (RLCs), compute services nodes, and CMCs, you must always use the same port number for the same node in both switches in the stack. In other words:
  - If you connect `r1lead eth0` to switch A, port 43, then you must connect `r1leadeth1` to switch B, port 43.
  - Likewise, if you connect CMC `r1i0c CMC-0` port to switch A, port 2, then `r1i0c CMC-1` port must go to switch B port 2.
- When adding cascaded switch stacks, all switch stacks must cascade from the primary switch stack. In other words, there is always only, at most, one switch hop.
- When discovering cascaded switches pairs in an RMN setup, observe the following:
  - If you are connecting switch stack 1, switch A, port 48 to switch stack 2, then you must connect the second trunked connection to stack 2, switch B, port 48.
  - Until the cascaded switch stack is discovered, you must leave one trunk leg unplugged temporarily to prevent looping.
  - The `discover` command will tell you when it is safe to plug in the second leg of the trunk. This avoids circuit loops.

## Admin Node eth2 Link in the Bond is Down

A problem occasionally occurs, especially in SGI XE 270 admin nodes, where the active-backup or 802.3ad bonded bond0 interface contains an Ethernet eth2 interface that is down/not linked. To verify this, perform the following:

- Check the Ethernet port of the add-in card and confirm that it is lit.
- Confirm that the add-in card connection to the management switches is using port 0 with port 1 not connected (so not miswired).
- If you look at `/proc/net/bonding/bond0` file, you can confirm that eth2 is the link that is down.
- Use the `/sbin/ethtool eth2` command and confirm that the `Link detected:` is no.
- Run the commands `ip link set eth up` and then run the `/sbin/ethtool eth3` command to determine if the link detected is yes.

In this scenario, it is likely that the eth2/eth3 interfaces have been swapped. Another clue is that if eth2 (look at `/proc/net/bonding/bond0` since the bond enforces the same MAC address for all bonded members) has a MAC address that is larger than the MAC address of eth3 (as seen by `ip addr show eth3`).

To correct this situation, edit the `/etc/udev/rules.d/70-persistent-net.rules` file and swap the MACs associated with eth2 and eth3 in the file.

When you reboot the system, the admin node comes back up with eth2 and eth3 properly ordered.

## Installing SGI Tempo Versions Older than 2.9.0

After you upgrade or install SGI Tempo 2.9.0 on any slot, the boot manager is changed to GRUB version 2. At this point, you can no longer install SGI Tempo versions earlier than 2.9.0 on any slot. The procedure in this topic explains how to install an earlier SGI Tempo version. For information about booting an SGI Tempo 2.9.0 system, see the following:

"Booting the System" on page 59

The following procedure explains how to install a version of SGI Tempo that is earlier than 2.9.0. After you run this procedure, the boot system is again GRUB version 1. Subsequently, if you install SGI Tempo 2.9.0, the boot system changes again to GRUB version 2.

**Procedure 6-2** To install an SGI Tempo version that is earlier than SGI Tempo 2.9.0

1. Log into the slot upon which you installed SGI Tempo 2.9.0 as the root user.
2. Type the following command to revert to the GRUB version 1 boot manager on the admin node:

```
/opt/sgi/lib/revert-admin-to-legacy-grub
```

3. Use the SGI DVD to install the older version of SGI Tempo that you want to use.

## Booting Nodes With iPXE After an Upgrade

If a node fails to boot after an upgrade, you might need to specify that iPXE load first and that iPXE load GRUB version 2. From the admin node, type the following command to specify the nodes:

```
cadmin --set-dhcp-bootfile --node node_ID ipxe
```

For *node\_ID*, specify the identifier of the node that did not boot. For example, for a compute node, specify its hostname.

To verify whether a node is enabled to load iPXE first, type the following command:

```
cadmin --show-dhcp-bootfile
```

## Miniroot Operations

The following information can help you with troubleshooting a suspected miniroot kernel problem:

- "About the SMC Miniroot" on page 195
- "Entering Rescue Mode" on page 196

## About the SMC Miniroot

The SMC miniroot is a small Linux environment based on the same RPM repositories that generated the root image itself. SMC uses the miniroot to install the software and to boot the following nodes over the cluster network:

- Rack leader controller (RLC) nodes
- Flat compute nodes
- SGI ICE compute nodes

The miniroot is a small, bootable file system. It includes kernel modules such as disk drivers, Ethernet drivers, and other software. These software drivers are associated with a specific kernel number. As new driver updates become available, the operating systems distribute additional kernels. At least one kernel needs to be associated with a specific node image, but there can be more than one kernel associated with a specific node image.

When SMC boots a node, SMC uses the images that reside in the admin node's image repository. Because the nodes boot over the network, it is important that the images in the admin node's image repository include the correct kernels. That is, it is important that the following be identical:

- The kernel in the on-node image. This is the installed image that resides on the node while the node is running.
- The kernel in the node image repository on the admin node. There can be multiple node images for a single node type in the image repository.
- The kernels in the kernel repository on the admin node. There can be multiple kernels in the kernel repository, and the `cinstallman` command includes a kernel from this repository when it builds a node image. These kernels reside in the `/tftpboot/boot` directory on the admin node.

When you use the `cinstallman` command to update images, SMC ensures synchronization between the on-node image and the image in the admin node's image repository.

If you change an on-node image manually, without using the `cinstallman` command, and you subsequently boot the node, then either the boot will fail or there will be a mismatch between the kernel loaded over the network versus the kernel and associated modules in the image itself. This mismatch can lead to a node that boots but has no network, for example. Therefore, it is important that all the images in the image repository on the admin node contain the on-node images with all the kernels

in use. If you update any images manually, use the following `cinstallman` command (1) to synchronize the kernels and the `initrd` daemon in the images and (2) to write a copy of the kernel to the `/tftpboot/boot` directory for future use when performing network boots:

```
cinstallman --update-kernels --image image_name
```

## Entering Rescue Mode

If you need to go in to miniroot rescue mode, you can do this by ensuring that the `rescue=1` kernel command line argument is added to the kernel command line list. For example, to accomplish this, use a command similar to the following:

```
cadmin --set-kernel-extra-params --node n1 rescue=1 cinstallman --refresh-netboot --node n1
```

To remove `rescue=1`, use the following command:

```
cadmin --unset-kernel-extra-params --node n1 cinstallman --refresh-netboot --node n1
```

## Booting an SGI Flat Compute Node on an Installed Cluster

When you boot a flat compute node on an installed and configured SGI ICE cluster system, the cluster completes a series of steps. For such a node, you can assume the following node characteristics:

- During the configuration process, the `discover` command was run. The `discover` command detected the node successfully, so the node is configured.
- The cluster-wide DHCP server is aware of the node's existence.
- At least one node image is available for this node in the admin node's image repository, and an image is assigned to the node.
- The GRUB version 2 network configuration files are available.
- The node has been installed, but now you want to boot it.
- The BIOS hardware on the node is configured for network booting. This is the factory-defined state.

The following topics describe the events that occur in each of the boot phases:

- "Phase 1 — Initiating the Boot" on page 197

- "Phase 2 — Loading the Node's Kernel" on page 198
- "Phase 3 — Loading the Miniroot" on page 200
- "Phase 4 — Starting the Operating System on the Node" on page 200

---

**Note:** The boot process is different on an SGI ICE cluster that has not been configured when compared to an SGI ICE cluster that has been configured.

---

### Phase 1 — Initiating the Boot

The following are the events that occur when you initiate the boot:

1. From the admin node, type the following command:

```
power node on node
```

For example:

```
power node on n0
```

2. The node turns on.
3. The factory-defined boot process starts. The boot proceeds as follows:
  - The BIOS boots from the network interface card (NIC).
  - The ROM on the NIC activates the PXE protocol to boot over the network.
  - The ROM sends a DHCP request for the node's network information. The node's IP address is a static address.

In this case, the admin node receives the DHCP request. If this were an SGI ICE compute node that was managed by a rack leader controller (RLC), the RLC receives the request.

The DHCP configuration files reside in the following directory:

```
/etc/dhcp/dhcpd.conf.d/ice.conf
```

Include files also reside in the `dhcpd.conf.d` directory.

4. The admin node responds to the DHCP request by sending DHCP packets that include the following:
  - The node's network configuration.

- The location of the GRUB version 2 boot loader.

The loader resides on the admin node in one of the following locations:

- On legacy platforms, this location is as follows:

```
/tftpboot/boot/grub2/i386-pc/core.0
```

- On EFI boot systems, this location is as follows:

```
/tftpboot/boot/grub2/i386-pc/core.efi
```

5. The NIC loads GRUB version 2.
6. The admin node directs the node to boot GRUB version 2.
7. The NIC starts.

## Phase 2 — Loading the Node's Kernel

In Phase 2, GRUB version 2 starts up and loads the kernel onto the node. All requests and all transfers are done using TFTP.

The final steps in this topic's process are for GRUB version 2 to load the node's kernel and the operating system's `initrd` files. These files reside on the admin node in the following directories:

```
/tftpboot/boot/image_name/kernel
/tftpboot/boot/image_name/initrd
```

The steps are as follows:

1. GRUB version 2 sends a TFTP request to the admin node for a configuration file.
2. The admin node responds to GRUB version 2 by sending file `/tftpboot/boot/grub2/grub.cfg`, which includes instructions that explain how to load kernel files and `initrd` files.
3. The `grub.cfg` file requests the node-specific configuration file.

The configuration file resides on the admin node in the following location, which includes the node's IP address:

```
/tftpboot/boot/grub2/tempo/ip_addr.cfg
```

For example, the configuration file might reside in

```
/tftpboot/boot/grub2/tempo/172.23.0.2.cfg
```

The node's configuration file includes the following information:

- The kernel and the kernel parameters that the node needs.
- The initrd file the node needs.
- The image assigned to the node. The image itself resides in the following location:

`/var/lib/systemimager/images/images`

- Instructions for booting the node. The file also includes instructions for installing the node, but GRUB version 2 uses these installation instructions only when the node is installed.
- All setup by Tempo.
- Port number identification information. The node's miniroot image resides on a specific port. The miniroot images reside in the following directory:

`/var/lib/systemimager/miniroot/squeezed/images`

The following is an example `ip_addr.cfg` file:

```
#
Boot the SGI miniroot with appropriate options (via the distro initrd)
kernel="sles11sp3.24104/vmlinuz-3.0.101-0.47.52-default"
initrd="sles11sp3.24104/initrd-3.0.101-0.47.52-default"
append="ROOTFS=disk IMAGE_PENDING=0 IMAGE=sles11sp3.24104 SLOT=1 console=ttyS1,115200n8
MONITOR_SERVER=172.23.0.1 MONITOR_CONSOLE=yes intel_idle.max_cstate=1 processor.max_cstate=1
TRANSPORT=rsync TTL=1 MCAST_RDV_ADDR=224.0.0.1 FLAMETHROWER_DIRECTORY_PORTBASE=9000
START_TIMEOUT=30 RECEIVE_TIMEOUT=5 crashkernel=256M"
gfxpayload=text
pxechain=0
```

---

**Note:** In the preceding example, lines have been wrapped for inclusion in this documentation.

---

4. GRUB version 2 loads the kernel with options.
5. The kernel requests the operating system distribution's `initrd` file.

### Phase 3 — Loading the Miniroot

In Phase 3, GRUB 2 loads the miniroot. The steps are as follows:

1. GRUB 2 runs the operating system distribution's `initrd` file on the node, starts the `udpcast` to obtain the miniroot, and runs additional SGI scripts associated with the node image.
2. The `udpcast` command transfers the node's miniroot from the miniroot's assigned port number to the node.
3. GRUB version2 unpacks the miniroot, which arrives on the node as a `tar` file.
4. The node's miniroot begins operating.

### Phase 4 — Starting the Operating System on the Node

In Phase 4, the miniroot starts running processes on the node and the operating system takes over. The steps are as follows:

1. The miniroot finds the root and boot file systems.
2. The miniroot mounts the root and boot file systems.
3. The operating system distribution's startup scripts start to run.

## YAST Navigation

The following list shows SLES YAST navigation key sequences:

| <b>Key</b>          | <b>Action</b>                                                                                                                                                    |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Tab                 |                                                                                                                                                                  |
| Alt + Tab           |                                                                                                                                                                  |
| Esc + Tab           |                                                                                                                                                                  |
| Shift + Tab         |                                                                                                                                                                  |
|                     | Moves you from label to label or from list to list.                                                                                                              |
| Ctrl + L            | Refreshes the screen.                                                                                                                                            |
| Enter               | Starts a module from a selected category, runs an action, or activates a menu item.                                                                              |
| Up arrow            | Changes the category. Selects the next category up.                                                                                                              |
| Down arrow          | Changes the category. Selects the next category down.                                                                                                            |
| Right arrow         | Starts a module from the selected category.                                                                                                                      |
| Shift + right arrow |                                                                                                                                                                  |
| Ctrl + A            |                                                                                                                                                                  |
|                     | Scrolls horizontally to the right. Useful in screens if use of the <code>left arrow</code> key would otherwise change the active pane or current selection list. |
| Alt + <i>letter</i> |                                                                                                                                                                  |
| Esc + <i>letter</i> |                                                                                                                                                                  |
|                     | Selects the label or action that begins with the <i>letter</i> you select. Labels and selected fields in the display contain a highlighted <i>letter</i> .       |
| Exit                | Quits the YAST interface.                                                                                                                                        |



## Subnetwork Information

This chapter contains the following topics:

- "About the Cluster Subnetworks" on page 203
- "SGI Rackable Subnetworks" on page 204
- "SGI ICE VLANs and Subnetworks" on page 208

### About the Cluster Subnetworks

Cluster hardware components can be connected to multiple networks. Within those networks, the components use a variety of virtual local area networks (VLANs) to logically separate network communication between specific sets of hardware. VLANs allow the sharing of physical hardware, both switches and routers, between all hosts on a network while maintaining predefined communication boundaries. You can use the reference material in this appendix section when you reconfigure or debug a cluster network.

SGI Rackable clusters and SGI ICE clusters have a very similar network design, but the SGI ICE clusters typically include special hardware that requires some additional network VLANs on the cluster network.

SGI Management Center (SMC) commands enable you to perform some procedures on only one component or on a range of similar components. Addressing methods differ depending on the command, component, the VLAN (or VLANs) in which the component resides, and whether or not the component has an IP address that is externally available.

The tables in this appendix section use the following terms:

- **Component.** The name of the component that you typically use in speech or in writing. For example: admin node, flat compute node, and so on.
- **SMC name.** The system-wide unique identifier for the component as it appears in the `hostname1` field in the cluster definition file.

## SGI Rackable Subnetworks

The following topics provide system-wide, rack-specific information about the SGI Rackable subnetworks:

- "SGI Rackable Component Naming Conventions" on page 204
- "SGI Rackable Address Ranges and VLANs for Management Software and Application Software" on page 205
- "SGI Rackable System Management Network Topology Overview" on page 207

## SGI Rackable Component Naming Conventions

Table B-1 on page 205 explains how to specify components when you run administrative and user commands. *X*, *Y*, and *Z* are always integer numbers. The Examples column shows the hostnames that you might specify when you use the `ssh(1)` command to connect to a component. The names in the Examples column are those that could appear in the cluster definition file. To generate a copy of the cluster definition file, log in as the root user and type the following command:

```
discover --show-configfile > out_file
```

**Table B-1** Naming Conventions

| Component                          | SMC Name                                                            | Examples                                                                                                                                                 |
|------------------------------------|---------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Admin node                         | N/A                                                                 | rackable1<br>sleet                                                                                                                                       |
| Ethernet management switch         | mgmtswX                                                             | mgmtsw0, typically the spine switch<br>mgmtsw100, typically a leaf switch                                                                                |
| SGI Rackable flat compute node     | serviceX / nX<br>serviceX-mic0 / nX-mic0<br>serviceX-mic1 / nX-mic1 | service0, the first flat compute node<br>service100, the 100th flat compute node<br>n0, the first flat compute node<br>n100, the 100th flat compute node |
| SGI Rackable flat compute node BMC | serviceX-BMC / nX-bmc                                               | service0-bmc<br>service100-bmc<br>n0-bmc<br>n100-bmc                                                                                                     |
| InfiniBand (IB) switch             | rXi YsZ-bmc                                                         | r1i4s1-bmc, rack 1, IRU 4, switch 1                                                                                                                      |
| Ethernet data switch               | rXdataswY                                                           | r1datasw1, rack 1, data switch 1<br>r100datasw2, the rack 100, data switch 2                                                                             |

## SGI Rackable Address Ranges and VLANs for Management Software and Application Software

The tables in this topic show the system-wide IP address ranges that the cluster management software uses and defines. The following notes pertain to these tables:

- All SGI Rackable clusters have an Ethernet-based network. The SMC software uses this network to handle administrative tasks and monitor different elements that exist inside of an SGI Rackable cluster. This network is often referred to as the *management network* or the *management fabric*.
- All SGI Rackable compute nodes and their baseboard management controllers (BMCs) are included in the same VLAN, but to avoid cross communication, the nodes and BMCs use a different IP address range so they cannot communicate with each other.
- SGI Rackable clusters include all the flat compute nodes in a single VLAN (head/head-bmc). This VLAN uses internet group management protocol (IGMP)

to handle all multicast join/leave requests for the provisioning protocol used, which is UDPCast.

- When the UDPCast protocol is used as the transport medium, SGI Rackable compute nodes generate an IGMP `join` packet for a specific multicast group that the switch recognizes and forwards the appropriate multicast traffic from the admin node to the compute node. This activity completes the multicast transfer of data.
- Almost all SGI Rackable clusters have a separate fabric used for computation, storage, MPI, and so on. This is often referred to as the *data network* or *data fabric*.
- The data network can be configured as either Infiniband or a separate Ethernet network, as follows:
  - If the network is Infiniband, SMC assigns a static IP addresses to interfaces `ib0`, and optionally `ib1`, for dual-plane Infiniband fabrics.
  - If the network is Ethernet, SMC assigns the Ethernet interfaces a unique name in the format of `sgimac_addr`. For example: `sgi001122aabbcc`.

This is done as an attempt to not have the Ethernet interfaces interfere with the standard `eth0/eth1` interfaces used on the management fabric. Manual intervention is required to assign IP addresses to these Ethernet interfaces or to configure them for bonding.

Table B-2 on page 207 shows the system-wide IP address ranges that the cluster management software uses. The following notes pertain to this table:

- The `head_bmc` network is a separate IP subnetwork.
- MIC devices, if present, are included on the `head` node network.

Table B-2 on page 207 and Table B-3 on page 207 show the address ranges and VLANs that the management software and the application software use.

**Table B-2** SGI Rackable System-wide Ethernet VLAN Information

| Subnetwork<br>VLAN<br>Number | Subnetwork<br>Name | IP Range / Subnet<br>Mask | Nodes in the Subnetwork                  |
|------------------------------|--------------------|---------------------------|------------------------------------------|
| 1                            | head               | 172.23.0.0/16             | Admin Node<br>Flat compute nodes<br>MICs |
| 1                            | head-bmc           | 172.24.0.0/16             | Admin Node BMC<br>Flat compute node BMCs |

Table B-3 on page 207 shows SGI Rackable InfiniBand fabric information.

**Table B-3** SGI Rackable InfiniBand Fabric Information

| Subnetwork Name | IP Range /<br>Subnetwork Mask | Nodes                                            |
|-----------------|-------------------------------|--------------------------------------------------|
| ib0             | 10.148.0.0/16                 | Admin node (optional)<br>Flat compute IB<br>MICs |
| ib1             | 10.149.0.0/16                 | Admin node (optional)<br>Flat compute IB<br>MICs |

## SGI Rackable System Management Network Topology Overview

The system management network for an SGI Rackable cluster is configured in a fashion that is very similar to SGI ICE clusters, but there are a few subtle differences. The following list explains these differences:

- The spine-to-leaf topology is the standard topology in SGI Rackable clusters. If your cluster requires switch redundancy, you can configure stacked Ethernet switches. Such a configuration might not always be required or desired. If redundant stacked switches are part of the configuration, the following connections are doubled and configured for link bonding: the admin node connection to the spine switch and the inter-switch links.

- Most SGI Rackable compute nodes share a BMC port with interface eth0. Typically, this single port is cabled to a top-of-rack (TOR) leaf switch with a single 1G cable.
- Typically, the leaf switch is connected to the spine switch with a single 1G cable. If the Ethernet switches are stacked, dual 1G cables are used and are configured for bonding. If link oversubscription is a concern (47:1), the 1G Ethernet leaf switches can connect to a 10G Ethernet spine switch via a 10G connection or a dual-10G connection with the correct models of switches.
- The admin node connects to the spine switch through either a single 1G cable or dual-1G cables. If dual-1G cables are used, the software creates a logically bonded interface between the two links. If the spine switch is a 10G Ethernet switch, this connection requires a 10G connection or dual-10G connection.

## SGI ICE VLANs and Subnetworks

The following topics provide system-wide, rack-specific information about the SGI ICE VLANs and subnetworks:

- "SGI ICE Component Naming Conventions" on page 208
- "SGI ICE Address Ranges and VLANs for Management Software and Application Software" on page 210
- "SGI ICE System Control Configuration" on page 213

## SGI ICE Component Naming Conventions

Table B-4 on page 209 explains how to specify components when you run administrative and user commands. *X*, *Y*, and *Z* are always positive integer numbers. The examples represent the hostnames that a user might specify when they use the `ssh(1)` command to connect to a component. The names in the Examples column are those that could appear in the cluster definition file. To generate a copy of the cluster definition file, log in as the root user and type the following command:

```
discover --show-configfile > out_file
```

**Table B-4** Naming Conventions

| Component                                 | SMC Name                                                            | Examples                                                                                                                                                                                 |
|-------------------------------------------|---------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Admin node                                | admin                                                               | admin<br>icex1<br>mysiteicex<br>sleet                                                                                                                                                    |
| Ethernet management switch                | mgmtswX                                                             | mgmtsw0, typically the spine switch<br>mgmtsw100, typically a leaf switch                                                                                                                |
| Rack leader controller (RLC)              | rXlead                                                              | r1lead, the RLC on the first rack<br>r2lead, the RLC on the second rack                                                                                                                  |
| RLC baseboard management controller (BMC) | rXlead-bmc                                                          | r1lead-bmc<br>r2lead-bmc                                                                                                                                                                 |
| SGI ICE compute node                      | rXiYnZ<br>rXiYnZ-eth<br>rXiYnZ-mic0<br>rXiYnZ-mic1                  | r1i3n10, rack 1, IRU 3, node 10<br>r33i6n3, rack 33, IRU 6, node 3                                                                                                                       |
| SGI ICE compute node BMC                  | rXiYnZ-bmc                                                          | r1i3n10-bmc, rack 1, IRU 3, node 10 BMC                                                                                                                                                  |
| SGI Rackable flat compute node            | serviceX / nX<br>serviceX-mic0 / nX-mic0<br>serviceX-mic1 / nX-mic1 | service0, the first SGI Rackable compute node<br>service100, the 100th SGI Rackable compute node<br>n0, the first SGI Rackable compute node<br>n100, the 100th SGI Rackable compute node |
| SGI Rackable flat compute node BMC        | serviceX-BMC / nX-bmc                                               | service0-bmc<br>service100-bmc<br>n0-bmc<br>n100-bmc                                                                                                                                     |
| InfiniBand switch                         | rXiYsZ-bmc                                                          | r1i4s1-bmc, rack 1, IRU 4, switch 1                                                                                                                                                      |
| Chassis management controller (CMC)       | rXiYc                                                               | r1i1c, rack 1, IRU 1 CMC<br>r22i8c, rack 22, IRU 8, CMC                                                                                                                                  |

## **SGI ICE Address Ranges and VLANs for Management Software and Application Software**

The tables in this topic show the system-wide IP address ranges that the cluster management software uses. The following notes pertain to these tables:

- The nodes and their baseboard management controllers (BMCs) are included in the same VLAN, but to avoid cross-communication, the nodes and the BMCs use different IP ranges and subnet masks.
- The MCell (cooling) VLAN uses an 802.1Q tag of 3 and is passed on all inter-switch connections, chassis management controller (CMC) connections, and the admin node connection
- The rack leader controllers (RLCs) are in the head VLAN (untagged), which allow communication to the admin node. The RLCs also have a rack-specific VLAN tagged interface to communicate with their respective CMCs.
- The CMCs are included only in their respective rack-specific VLAN, not the head VLAN. This means that the CMCs and SGI ICE compute nodes can only communicate with their defined RLC.
- The InfiniBand (IB) switches are included only in their respective rack-specific VLAN, not the head VLAN. Physically, the IB switch BMCs are located behind the CMCs. Logically, the IB switches obtain an IP address in the BMC's IP range within that rack-specific VLAN.

Table B-5 on page 211, Table B-6 on page 211, and Table B-7 on page 212 show the address ranges and VLANs that the management software and the application software use.

**Table B-5** SGI ICE System-wide Ethernet VLAN Information

| Subnetwork<br>VLAN<br>Number | Subnetwork<br>Name | IP Range / Subnet<br>Mask | Nodes in the Subnetwork                              |
|------------------------------|--------------------|---------------------------|------------------------------------------------------|
| 1                            | head               | 172.23.0.0/16             | Admin Node<br>RLCs<br>Flat compute nodes<br>MICs     |
| 1                            | head-bmc           | 172.24.0.0/16             | Admin Node BMC<br>RLC BMCs<br>Flat compute node BMCs |
| 3                            | mcell-net          | 172.26.0.0/16             | Admin node<br>Cooling CRCs<br>Cooling CDU<br>CMCs    |

Table B-6 on page 211 shows the per-rack IP address ranges that the cluster management software uses in the rack VLANs.

**Table B-6** SGI ICE Rack-specific Ethernet VLAN Information

| Subnetwork<br>VLAN<br>Number | Subnetwork<br>Name | IP Range /<br>Subnetwork<br>Mask | Nodes in the Subnetwork                                                                                |
|------------------------------|--------------------|----------------------------------|--------------------------------------------------------------------------------------------------------|
| 101                          | vlan101:gbe        | 10.159.0.0/22                    | Rack 1's RLCs<br>Rack 1's SGI ICE compute nodes                                                        |
| 101                          | vlan101:bmc        | 10.160.0.0/22                    | Rack 1's RLC BMC<br>Rack 1's SGI ICE compute node<br>BMCs<br>Rack 1's IB switches<br>Rack 1's CMC BMCs |
| 102                          | vlan102:gbe        | 10.159.4.0/22                    | Rack 2's RLCs<br>Rack 2's SGI ICE compute nodes                                                        |

B: Subnetwork Information

---

| Subnetwork<br>VLAN<br>Number | Subnetwork<br>Name | IP Range /<br>Subnetwork<br>Mask | Nodes in the Subnetwork                                                                                |
|------------------------------|--------------------|----------------------------------|--------------------------------------------------------------------------------------------------------|
| 102                          | vlan102:bmc        | 10.160.4.0/22                    | Rack 2's RLC BMC<br>Rack 2's SGI ICE compute node<br>BMCs<br>Rack 2's IB switches<br>Rack 2's CMC BMCs |
| XXX                          | vlanXXX:gbe        | 10.159.X.0/22                    | Rack X's RLCs<br>Rack X's SGI ICE compute nodes                                                        |
| XXX                          | vlanXXX:bmc        | 10.160.X.0/22                    | Rack X's RLC BMC<br>Rack X's SGI ICE compute node<br>BMCs<br>Rack X's IB switches<br>Rack X's CMC BMCs |

Table B-7 on page 212 shows the system-wide IP address ranges for cluster application software. Only the RLCs that provide InfiniBand subnetwork services need to connect.

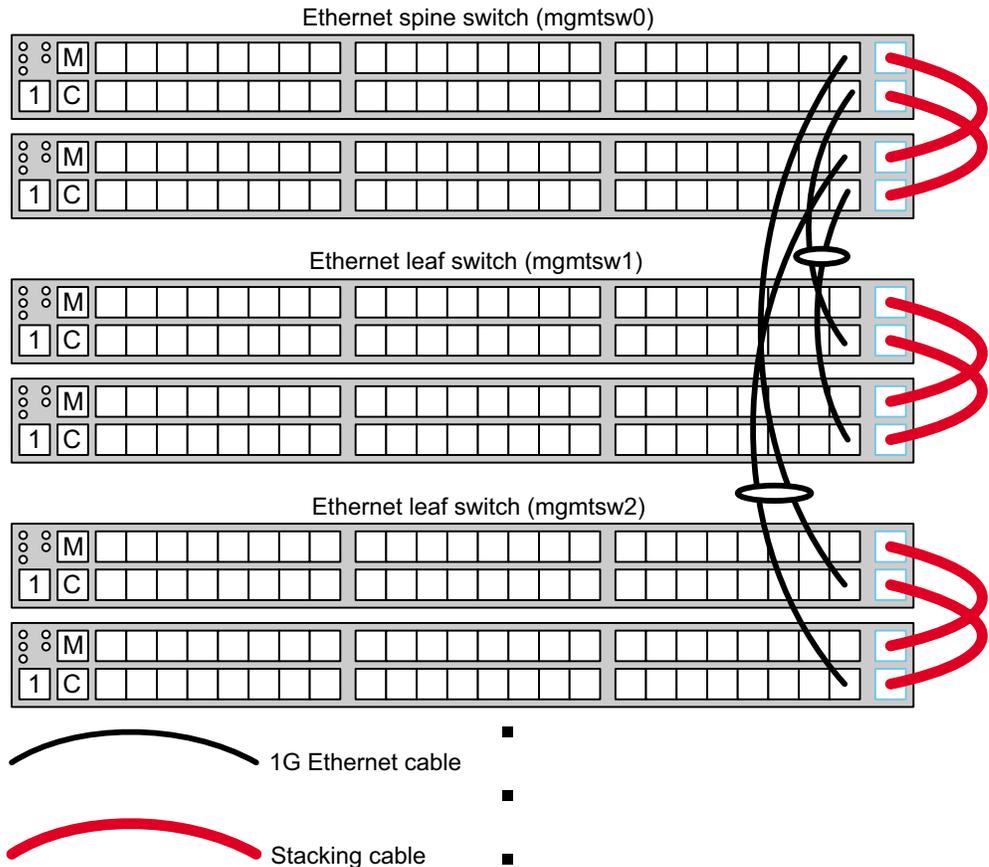
**Table B-7** SGI ICE InfiniBand Fabric Information

| Subnetwork Name | IP Range /<br>Subnetwork Mask | Nodes                                                                                           |
|-----------------|-------------------------------|-------------------------------------------------------------------------------------------------|
| ib0             | 10.148.0.0/16                 | Admin node (optional)<br>RLCs (Optional)<br>SGI ICE compute nodes<br>Flat compute nodes<br>MICs |
| ib1             | 10.149.0.0/16                 | Admin node (optional)<br>RLCs (Optional)<br>SGI ICE compute nodes<br>Flat compute nodes<br>MICs |

## **SGI ICE System Control Configuration**

You can configure the system management network for an SGI ICE cluster either as a redundant management network or as a nonredundant management network, as follows:

- A redundant management network is the default configuration. The physical topology is as follows:
  - A pair of stacked Ethernet switches that act as the spine switch. These switches consist of two physical switches logically configured to act as a single switch. This configuration facilitates high availability and increased throughput.
  - Additional Ethernet switches, which act as leaf switches. These additional switches are also paired, stacked Ethernet switches. The spine-to-leaf connections connect each physical switch in the spine switch to each physical switch in the leaf switches. Figure B-1 on page 214 shows these connections.



**Figure B-1** Spine Switch and Leaf Switches

For redundancy, the admin node, the RLCs, the CMCs, and any flat compute nodes attach to the additional switch with two cables, creating dual links. The cooling distribution units (CDUs) and the cooling rack controllers (CRCs) do not have any additional links or cables.

You can customize the dual links to use different modes of standard Linux bonding, but the default bonding settings are as follows:

- Admin node: `active-backup`. `Mode=1`.

- RLCs: link aggregation control protocol (LACP) or 802.3ad. Mode=4  
xmit\_hash\_policy=layer2+3.
- Flat compute nodes: active-backup. Mode=1.
- CMCs: static port-channels. No Linux mode. CMCs use switch bonding.

When the `discover` command runs, the Ethernet switches configure the matching ports and bonding methods to match the bonding settings for each of the different node types.

- In a nonredundant management network configuration, the topology is the same as the redundant management network, but the secondary switches are not present. In this configuration, the SGI ICE components have only a single connection into the system management network.



---

## SGI ICE MCell Network IP Addresses

If you need to troubleshoot the MCell cooling equipment on an SGI ICE cluster, you need to know the IP addresses of the cooling rack controllers (CRCs) and cooling distribution units (CDUs) so that you can type a `ping(8)` command to the component. Each piece of equipment bears a label with its equipment number.

For CRCs, the IP address is `172.26.128.number`.

For CDUs, the IP address is `172.26.144.number`.

Table C-1 on page 217 shows the IP addresses for CRCs and CDUs.

**Table C-1** MCell Network Associations

| Rack | Cooling Rack Controllers (CRCs) | Cooling Distribution Unit (CDUs) |
|------|---------------------------------|----------------------------------|
| 1    | 172.26.128.1                    | 172.26.144.1                     |
| 2    | 172.26.128.1                    | 172.26.144.1                     |
| 3    | 172.26.128.2                    | 172.26.144.1                     |
| 4    | 172.26.128.2                    | 172.26.144.1                     |
| 5    | 172.26.128.3                    | 172.26.144.2                     |
| 6    | 172.26.128.3                    | 172.26.144.2                     |
| 7    | 172.26.128.4                    | 172.26.144.2                     |
| 8    | 172.26.128.4                    | 172.26.144.2                     |
| 9    | 172.26.128.5                    | 172.26.144.3                     |
| 10   | 172.26.128.5                    | 172.26.144.3                     |
| 11   | 172.26.128.6                    | 172.26.144.3                     |
| 12   | 172.26.128.6                    | 172.26.144.3                     |
| 13   | 172.26.128.7                    | 172.26.144.4                     |
| 14   | 172.26.128.7                    | 172.26.144.4                     |
| 15   | 172.26.128.8                    | 172.26.144.4                     |

C: SGI ICE MCell Network IP Addresses

---

| Rack | Cooling Rack Controllers (CRCs) | Cooling Distribution Unit (CDUs) |
|------|---------------------------------|----------------------------------|
| 16   | 172.26.128.8                    | 172.26.144.4                     |
| 17   | 172.26.128.9                    | 172.26.144.5                     |
| 18   | 172.26.128.9                    | 172.26.144.5                     |
| 19   | 172.26.128.10                   | 172.26.144.5                     |
| 20   | 172.26.128.10                   | 172.26.144.5                     |
| 21   | 172.26.128.11                   | 172.26.144.6                     |
| 22   | 172.26.128.11                   | 172.26.144.6                     |
| 23   | 172.26.128.12                   | 172.26.144.6                     |
| 24   | 172.26.128.12                   | 172.26.144.6                     |
| 25   | 172.26.128.13                   | 172.26.144.7                     |
| 26   | 172.26.128.13                   | 172.26.144.7                     |
| 27   | 172.26.128.14                   | 172.26.144.7                     |
| 28   | 172.26.128.14                   | 172.26.144.7                     |
| 29   | 172.26.128.15                   | 172.26.144.8                     |
| 30   | 172.26.128.15                   | 172.26.144.8                     |
| 31   | 172.26.128.16                   | 172.26.144.8                     |
| 32   | 172.26.128.16                   | 172.26.144.8                     |
| 33   | 172.26.128.17                   | 172.26.144.9                     |
| 34   | 172.26.128.17                   | 172.26.144.9                     |
| 35   | 172.26.128.18                   | 172.26.144.9                     |
| 36   | 172.26.128.18                   | 172.26.144.9                     |
| 37   | 172.26.128.19                   | 172.26.144.10                    |
| 38   | 172.26.128.19                   | 172.26.144.10                    |
| 39   | 172.26.128.20                   | 172.26.144.10                    |
| 40   | 172.26.128.20                   | 172.26.144.10                    |
| 41   | 172.26.128.21                   | 172.26.144.11                    |
| 42   | 172.26.128.21                   | 172.26.144.11                    |

| Rack | Cooling Rack Controllers (CRCs) | Cooling Distribution Unit (CDUs) |
|------|---------------------------------|----------------------------------|
| 43   | 172.26.128.22                   | 172.26.144.11                    |
| 44   | 172.26.128.22                   | 172.26.144.11                    |
| 45   | 172.26.128.23                   | 172.26.144.12                    |
| 46   | 172.26.128.23                   | 172.26.144.12                    |
| 47   | 172.26.128.24                   | 172.26.144.12                    |
| 48   | 172.26.128.24                   | 172.26.144.12                    |
| 49   | 172.26.128.25                   | 172.26.144.13                    |
| 50   | 172.26.128.25                   | 172.26.144.13                    |
| 51   | 172.26.128.26                   | 172.26.144.13                    |
| 52   | 172.26.128.26                   | 172.26.144.13                    |
| 53   | 172.26.128.27                   | 172.26.144.14                    |
| 54   | 172.26.128.27                   | 172.26.144.14                    |
| 55   | 172.26.128.28                   | 172.26.144.14                    |
| 56   | 172.26.128.28                   | 172.26.144.14                    |
| 57   | 172.26.128.29                   | 172.26.144.15                    |
| 58   | 172.26.128.29                   | 172.26.144.15                    |
| 59   | 172.26.128.30                   | 172.26.144.15                    |
| 60   | 172.26.128.30                   | 172.26.144.15                    |
| 61   | 172.26.128.31                   | 172.26.144.16                    |
| 62   | 172.26.128.31                   | 172.26.144.16                    |
| 63   | 172.26.128.32                   | 172.26.144.16                    |
| 64   | 172.26.128.32                   | 172.26.144.16                    |
| 65   | 172.26.128.33                   | 172.26.144.17                    |
| 66   | 172.26.128.33                   | 172.26.144.17                    |
| 67   | 172.26.128.34                   | 172.26.144.17                    |
| 68   | 172.26.128.34                   | 172.26.144.17                    |
| 69   | 172.26.128.35                   | 172.26.144.18                    |

C: SGI ICE MCell Network IP Addresses

---

| Rack | Cooling Rack Controllers (CRCs) | Cooling Distribution Unit (CDUs) |
|------|---------------------------------|----------------------------------|
| 70   | 172.26.128.35                   | 172.26.144.18                    |
| 71   | 172.26.128.36                   | 172.26.144.18                    |
| 72   | 172.26.128.36                   | 172.26.144.18                    |
| 73   | 172.26.128.37                   | 172.26.144.19                    |
| 74   | 172.26.128.37                   | 172.26.144.19                    |
| 75   | 172.26.128.38                   | 172.26.144.19                    |
| 76   | 172.26.128.38                   | 172.26.144.19                    |
| 77   | 172.26.128.39                   | 172.26.144.20                    |
| 78   | 172.26.128.39                   | 172.26.144.20                    |
| 79   | 172.26.128.40                   | 172.26.144.20                    |
| 80   | 172.26.128.40                   | 172.26.144.20                    |
| 81   | 172.26.128.41                   | 172.26.144.21                    |
| 82   | 172.26.128.41                   | 172.26.144.21                    |
| 83   | 172.26.128.42                   | 172.26.144.21                    |
| 84   | 172.26.128.42                   | 172.26.144.21                    |
| 85   | 172.26.128.43                   | 172.26.144.22                    |
| 86   | 172.26.128.43                   | 172.26.144.22                    |
| 87   | 172.26.128.44                   | 172.26.144.22                    |
| 88   | 172.26.128.44                   | 172.26.144.22                    |
| 89   | 172.26.128.45                   | 172.26.144.23                    |
| 90   | 172.26.128.45                   | 172.26.144.23                    |
| 91   | 172.26.128.46                   | 172.26.144.23                    |
| 92   | 172.26.128.46                   | 172.26.144.23                    |
| 93   | 172.26.128.47                   | 172.26.144.24                    |
| 94   | 172.26.128.47                   | 172.26.144.24                    |
| 95   | 172.26.128.48                   | 172.26.144.24                    |
| 96   | 172.26.128.48                   | 172.26.144.24                    |

| Rack | Cooling Rack Controllers (CRCs) | Cooling Distribution Unit (CDUs) |
|------|---------------------------------|----------------------------------|
| 97   | 172.26.128.49                   | 172.26.144.25                    |
| 98   | 172.26.128.49                   | 172.26.144.25                    |
| 99   | 172.26.128.50                   | 172.26.144.25                    |
| 100  | 172.26.128.50                   | 172.26.144.25                    |



## Partition Layout Information

This appendix section includes the following topics:

- "About the Partition Layout on SGI Clusters" on page 223
- "About the Current Release's Partition Layout" on page 225
- "About the Legacy Partition Layout" on page 228

### About the Partition Layout on SGI Clusters

As of the SGI Tempo 2.9.0 release, the partition layout scheme changed for SGI cluster systems. The SMC 3.0 and later releases use the partition layout scheme that SGI Tempo 2.9.0 introduced. This topic explains how the new partition layout scheme affects SGI clusters.

If you upgrade an SGI ICE cluster to an SGI Tempo 2.9.0 or later release or to an SMC 3.0 or later release, it is possible for your system to have the following mixture of slots:

- Slots with the legacy, SGI Tempo pre-2.9.0 partition scheme
- Slots with the SGI Tempo 2.9.0 or later or the SMC 3.0 or later partitioning scheme

The SGI Tempo 2.9.0 release introduced a partition layout that uses the GUID partition table (GPT) and a new boot system, which is GRUB version 2. In previous SGI Tempo versions, the partition layout was the MSDOS layout, and the boot system was GRUB version 1. These changes affect system operations in the following ways:

- If your cluster was originally installed with one of the following, then the software on your cluster uses the new partition layout:
  - SGI Tempo 2.9.0 or later
  - SMC 3.0. or later

In this case, you can use the `cadmin` command, as shown in the SGI documentation, to manage all the slots.

- If you upgrade a slot from an SGI Tempo release that is earlier than the SGI Tempo 2.9.0 release, note the following:

- If you upgrade any slot to the current SGI Tempo or SMC release, the admin node converts the boot loader installed in the master boot record (MBR) to GRUB version 2. This boot loader chooses a slot when you boot the admin node.
- If you do not upgrade all the slots to SGI Tempo 2.9.0 or later or to SMC 3.0 or later, the only administrative action you can perform on the slots with the older SGI Tempo software is to boot the slot. You can boot the slots that host older versions of SGI Tempo at the console or by selecting the slot on the GRUB version 2 boot menu. Because the admin node MBR is now upgraded to GRUB version 2, you can no longer use the `cadmin` command to manage any slots that host older versions of SGI Tempo.
- You cannot install SGI Tempo releases older than SGI Tempo 2.9.0 on any of the system's slots. If you perform any future, from-scratch installations, these installations must be to SGI Tempo 2.9.0 or later or to SMC 3.0 or later. You can, however, use the `updatetempo` upgrade script to upgrade slots to releases other than SGI Tempo 2.9.0 or SMC 3.0.

Disks with the new GPT layout can exist on the system alongside disks with the MSDOS partition layout. On a cluster system newly installed with SGI Tempo 2.9.0 or later or with SMC 3.0 or later, the software configures the GPT layout on all system disks and all disks that you clear, except for the SGI ICE compute nodes, if present. After you upgrade one or more slots to SGI Tempo 2.9.0 or later or to SMC 3.0 or later, if you add new system disks or you clear existing system disks, the software configures the GPT layout on system disks that you add and on all disks that you clear, except for the compute nodes.

- For the admin node, rack leader controller (RLC), and flat compute nodes, when you install SGI Tempo 2.9.0 or later or SMC 3.0 or later with blank system disks, or with system disks that have been cleared, the software uses the new GPT layout. If the disks have the old MSDOS partition table format, the software continues to use the old MSDOS partition table format until you manually clear the disks and reinstall (if desired).

The following topics explain the partition layout for slots installed with the current release and for slots installed with legacy software:

- "About the Current Release's Partition Layout" on page 225
- "About the Legacy Partition Layout" on page 228

## About the Current Release's Partition Layout

The following topics show the partition layout that the current release creates:

- "Partition Layout for a One-slot Cluster" on page 225
- "Partition Layout for a Two-slot Cluster (Default)" on page 226
- "Partition Layout for a Five-slot Cluster" on page 227

### Partition Layout for a One-slot Cluster

Table D-1 on page 225 shows the partition layout for a one-slot cluster. This layout yields one boot partition. If you configure a single-slot system and later decide to add another partition, the addition process destroys all the data on your system.

**Table D-1** Partition Layout for a Single-boot Cluster

| Partition | File System Type                                        | File System Label | Notes                                                                                                            |
|-----------|---------------------------------------------------------|-------------------|------------------------------------------------------------------------------------------------------------------|
| 1         | Ext3                                                    | sgidata           | Contains slot information. On the admin node, contains GRUB version 2 data for choosing root slots at boot time. |
| 2         | swap                                                    | sgiswap           | Swap partition.                                                                                                  |
| 3-10      | N/A                                                     | N/A               | N/A                                                                                                              |
| 11        | Ext3                                                    | sgiboot           | Slot 1 /boot partition.                                                                                          |
| 12-20     | N/A                                                     | N/A               | N/A                                                                                                              |
| 21        | VFAT                                                    | sgiefi            | /boot/efi partition. EFI BIOS clusters only. On x86_64 BIOS clusters, this partition is unused.                  |
| 22-30     | N/A                                                     | N/A               | N/A                                                                                                              |
| 31        | Ext3 on admin node and flat compute nodes. XFS on RLCs. | sgiroot           | Slot 1 / partition.                                                                                              |

## Partition Layout for a Two-slot Cluster (Default)

Table D-2 on page 226 shows the partition layout for a two-slot cluster. This layout yields two boot partitions.

**Table D-2** Partition Layout for a Dual-boot Cluster (Default Layout)

| Partition | File System Type                                     | File System Label | Notes                                                                                                            |
|-----------|------------------------------------------------------|-------------------|------------------------------------------------------------------------------------------------------------------|
| 1         | Ext3                                                 | sgidata           | Contains slot information. On the admin node, contains GRUB version 2 data for choosing root slots at boot time. |
| 2         | swap                                                 | sgiswap           | Swap partition.                                                                                                  |
| 3-10      | N/A                                                  | N/A               | N/A                                                                                                              |
| 11        | Ext3                                                 | sgiboot           | Slot 1 /boot partition.                                                                                          |
| 12        | Ext3                                                 | sgiboot2          | Slot 2 /boot partition.                                                                                          |
| 13-20     | N/A                                                  | N/A               | N/A                                                                                                              |
| 21        | VFAT                                                 | sgiefi            | Slot 1 /boot/efi partition. EFI BIOS clusters only. On x86_64 BIOS clusters, this partition is unused.           |
| 22        | VFAT                                                 | sgiefi2           | Slot 2 /boot/efi partition. EFI BIOS clusters only. On x86_64 BIOS clusters, this partition is unused.           |
| 23-30     | N/A                                                  | N/A               | N/A                                                                                                              |
| 31        | Ext3 on admin node, flat compute nodes. XFS on RLCs. | sgiroot           | Slot 1 / partition.                                                                                              |
| 32        | Ext3 on admin node, flat compute nodes. XFS on RLCs. | sgiroot2          | Slot 2 / partition.                                                                                              |

## Partition Layout for a Five-slot Cluster

Table D-3 on page 227 shows the partition layout for a five-slot cluster. This layout yields five boot partitions.

**Table D-3** Partition Layout for a Quintuple-boot Cluster

| Partition | File System Type | File System Label | Notes                                                                                                       |
|-----------|------------------|-------------------|-------------------------------------------------------------------------------------------------------------|
| 1         | Ext3             | sgidata           | Contains slot information. On the admin node, contains GRUB version 2 for choosing root slots at boot time. |
| 2         | swap             | sgiswap           | Swap partition.                                                                                             |
| 3-10      | N/A              | N/A               | N/A                                                                                                         |
| 11        | Ext3             | sgiboot           | Slot 1 /boot partition.                                                                                     |
| 12        | Ext3             | sgiboot2          | Slot 2 /boot partition.                                                                                     |
| 13        | Ext3             | sgiboot3          | Slot 3 /boot partition.                                                                                     |
| 14        | Ext3             | sgiboot4          | Slot 4 /boot partition.                                                                                     |
| 15        | Ext3             | sgiboot5          | Slot 5 /boot partition.                                                                                     |
| 16-20     | N/A              | N/A               | N/A                                                                                                         |
| 21        | VFAT             | sgiefi            | Slot 1 /boot/efi partition. EFI BIOS clusters only. On x86_64 BIOS clusters, this partition is unused.      |
| 22        | VFAT             | sgiefi2           | Slot 2 /boot/efi partition. EFI BIOS clusters only. On x86_64 BIOS clusters, this partition is unused.      |
| 23        | VFAT             | sgiefi3           | Slot 3 /boot/efi partition. EFI BIOS clusters only. On x86_64 BIOS clusters, this partition is unused.      |

| Partition | File System Type                                     | File System Label | Notes                                                                                                  |
|-----------|------------------------------------------------------|-------------------|--------------------------------------------------------------------------------------------------------|
| 24        | VFAT                                                 | sgiefi4           | Slot 4 /boot/efi partition. EFI BIOS clusters only. On x86_64 BIOS clusters, this partition is unused. |
| 25        | VFAT                                                 | sgiefi5           | Slot 5 /boot/efi partition. EFI BIOS clusters only. On x86_64 BIOS clusters, this partition is unused. |
| 26-30     | N/A                                                  | N/A               | N/A                                                                                                    |
| 31        | Ext3 on admin node, flat compute nodes. XFS on RLCs. | sgiroot           | Slot 1 / partition.                                                                                    |
| 32        | Ext3 on admin node, flat compute nodes. XFS on RLCs. | sgiroot2          | Slot 2 / partition.                                                                                    |
| 33        | Ext3 on admin node, flat compute nodes. XFS on RLCs. | sgiroot3          | Slot 3 / partition.                                                                                    |
| 34        | Ext3 on admin node, flat compute nodes. XFS on RLCs. | sgiroot4          | Slot 4 / partition.                                                                                    |
| 35        | Ext3 on admin node, flat compute nodes. XFS on RLCs. | sgiroot5          | Slot 5 / partition.                                                                                    |

## About the Legacy Partition Layout

The following topics show the legacy partition layouts for various slot configurations:

- "Legacy Partition Layout for a One-slot SGI ICE Cluster" on page 229

- "Legacy Partition Layout for a Two-slot SGI ICE Cluster" on page 229
- "Legacy Partition Layout for a Five-slot SGI ICE Cluster" on page 230

### Legacy Partition Layout for a One-slot SGI ICE Cluster

Table D-4 on page 229 shows the legacy partition layout for a one-slot SGI ICE cluster. This layout yields one boot partition.

**Table D-4** Partition Layout for a Single-boot SGI ICE Cluster

| Partition | File System Type                                   | File System Label | Notes                                                           |
|-----------|----------------------------------------------------|-------------------|-----------------------------------------------------------------|
| 1         | ext3                                               | sgiboot           | /boot partition.                                                |
| 2         | -                                                  | N/A               | Extended partition. Makes logicals out of the rest of the disk. |
| 5         | swap                                               | sgiswap           | Swap partition.                                                 |
| 6         | ext3 on admin node, service nodes.<br>XFS on RLCs. | sgiroot           | / partition.                                                    |

### Legacy Partition Layout for a Two-slot SGI ICE Cluster

Table D-5 on page 229 shows the legacy partition layout for a two-slot SGI ICE cluster. This layout yields two boot partitions.

**Table D-5** Partition Layout for a Dual-boot SGI ICE Cluster (Default Layout)

| Partition | File System Type | File System Label | Notes                                |
|-----------|------------------|-------------------|--------------------------------------|
| 1         | swap             | sgiswap           | Partition layout for multiple slots. |
| 2         | ext3             | sgidata           | SGI data partition.                  |

| Partition | File System Type                                   | File System Label | Notes                                                           |
|-----------|----------------------------------------------------|-------------------|-----------------------------------------------------------------|
| 3         | -                                                  | N/A               | Extended partition. Makes logicals out of the rest of the disk. |
| 5         | ext3                                               | sgiboot           | Slot 1 /boot partition.                                         |
| 6         | ext3 on admin node, service nodes.<br>XFS on RLCs. | sgiroot           | Slot 1 / partition.                                             |
| 7         | ext3                                               | sgiboot           | Slot 2 /boot partition.                                         |
| 8         | ext3 on admin node, service nodes.<br>XFS on RLCs. | sgiroot           | Slot 2 / partition.                                             |

### Legacy Partition Layout for a Five-slot SGI ICE Cluster

Table D-6 on page 230 shows the legacy partition layout for a five-slot SGI ICE cluster. This layout yields five boot partitions.

**Table D-6** Legacy Partition Layout for a Quintuple-boot SGI ICE Cluster

| Partition | File System Type | File System Label | Notes                                                           |
|-----------|------------------|-------------------|-----------------------------------------------------------------|
| 1         | swap             | sgiswap           | Partition layout for multiple slots.                            |
| 2         | ext3             | sgidata           | SGI data partition.                                             |
| 3         | -                | N/A               | Extended partition. Makes logicals out of the rest of the disk. |
| 5         | ext3             | sgiboot           | Slot 1 /boot partition.                                         |
| 6         | ext3 or XFS      | sgiroot           | Slot 1 / partition.                                             |
| 7         | ext3             | sgiboot           | Slot 2 /boot partition.                                         |

| Partition | File System Type | File System Label | Notes                   |
|-----------|------------------|-------------------|-------------------------|
| 8         | ext3 or XFS      | sgiroot           | Slot 2 / partition.     |
| 9         | ext3             | sgiboot           | Slot 3 /boot partition. |
| 10        | ext3 or XFS      | sgiroot           | Slot 3 / partition.     |
| 11        | ext3             | sgiboot           | Slot 4 /boot partition. |
| 12        | ext3 or XFS      | sgiroot           | Slot 4 / partition.     |
| 13        | ext3             | sgiboot           | Slot 5 /boot partition. |
| 14        | ext3 or XFS      | sgiroot           | Slot 5 / partition.     |



## Specifying Configuration Attributes

This appendix includes the following topics:

- "About Configuration Attributes" on page 233
- "UDPCast Options" on page 234
- "VLAN and General Network Options" on page 239
- "Console Server Options" on page 243
- "Miscellaneous Options" on page 244

### About Configuration Attributes

SGI cluster configuration information can be specified in several ways. For example:

- When you configure the cluster for the first time, you can provide configuration information in the cluster definition file or you can provide information by responding to the prompts in the online cluster configuration tool.
- When you add nodes to a cluster, you can specify node attributes as parameters to the `discover` command that you use to configure the nodes.
- When you use the `caadmin` command you set and apply an attribute.
- When you use the `cattr` command, you set an attribute.

SMC supports several global cluster attributes, and some attributes can be specified in more than one way. The following topics describe each attribute, show each attribute's default value, show additional other accepted values or ranges of values, and show the commands or files in which you can specify the value:

- "UDPCast Options" on page 234
- "VLAN and General Network Options" on page 239
- "Console Server Options" on page 243
- "Miscellaneous Options" on page 244

## UDPCast Options

The following configuration attributes control UDPCast operations:

- "edns\_udp\_size" on page 234
- "udpcast\_max\_bitrate" on page 234
- "udpcast\_max\_wait" on page 235
- "udpcast\_mcast\_rdv\_addr" on page 235
- "udpcast\_min\_receivers" on page 236
- "udpcast\_min\_wait" on page 236
- "udpcast\_rexmit\_hello\_interval" on page 237
- "udpcast\_ttl" on page 238

### **edns\_udp\_size**

Specifies the `edns-udp-size` option in `/etc/named.conf`. This value is the default packet size, in bytes. This is the packet size that remote servers can receive.

Default = 512.

Values = any positive integer number.

Accepted by:

- `catrr` command

### **udpcast\_max\_bitrate**

Specifies the maximum numbers of bits that are conveyed or processed per second, expressed as a number followed by a unit of measure, such as `m`.

Default = 900m.

Values = any positive integer number followed by a unit of measure. The default unit of measure is `m` (megabytes). For the complete list of units of measure, see the `udp-sender(1)` man page.

Accepted by:

- Cluster definition file
- `cadmin` command
- `cattr` command
- `discover` command

### `udpcast_max_wait`

Specifies the greatest amount of time that can elapse between when the first client node connects and any other client nodes connect. Clients that connect after this time has elapsed receive their software in a subsequent broadcast.

Default = 10.

Values = any positive integer number.

Accepted by:

- Cluster definition file
- `cadmin` command
- `cattr` command
- `discover` command

### `udpcast_mcast_rdv_addr`

Specifies the UDPcast rendezvous (RDV) address. Used for senders and receivers to find each other.

Note that the admin node default address and the global (rack leader controller (RLC) node) default address are different. If you change the global setting, which is used by leaders, also make the following changes:

- Adjust `--set-udpcast-mcast-rdv-addr`.
- Use the `cimage` command to push an image and initiate changes on the RLCs.

Default for the admin node = 239.0.0.1.

Default for the global (leaders) = 224.0.0.1.

Values = any valid IP address.

Accepted by:

- Cluster configuration tool
- Cluster definition file
- `cadmin` command
- `cattr` command
- `discover` command

#### **udpcast\_min\_receivers**

Specifies the minimum number of receiver nodes for UDPcast.

Default = 1.

Values = any positive integer number.

Accepted by:

- Cluster definition file
- `cadmin` command
- `cattr` command
- `discover` command

#### **udpcast\_min\_wait**

Specifies the minimum amount of time that the system waits, while allowing clients to connect, before the software broadcast begins. This is the time between when the first client node connects and any other client nodes connect. The UDPcast distributes the software to all clients that connect during this interval.

Default = 10.

Values = any positive integer number.

Accepted by:

- Cluster definition file
- `cadmin` command
- `cattr` command
- `discover` command

### `udpcast_rexmit_hello_interval`

Specifies the frequency with which the UDP sender transmits `hello` packets.

---

**Note:** The admin node has a different default than the RLCs.

---

The defaults are as follows:

- For the admin node, the default is 5000 (5 seconds).
- For the RLCs, the default is 0.

Values = any positive integer number.

Accepted by:

- Cluster definition file
- `cadmin` command
- `cattr` command
- `discover` command

The `--rexmit-hello-interval` setting is especially important when the rendezvous (RDV) address is not 224.0.0.1. The admin node, for example, defaults to 239.0.0.1 for UDP sender processes.

When a UDP receiver process starts for an RDV address other than 224.0.0.1, the operating system sends an IGMP packet that is detected by the Ethernet switch. The Ethernet switch then updates its tables with this information, thus allowing the multicast packets to properly route through the switch. The problem is that sometimes the UDP receiver sends its connection packet before the switch has had a chance to update the switch routing. If the request packet is not detected by the UDP sender on the admin node, perhaps because the packet was sent before the switch was set up to pass the packet, the UDP receiver could wait forever for a UDPcast

stream. The `udpcast_rexmit_hello_interval` value configures the UDP sender to send a HELLO packet at regular intervals and configures UDP receivers to respond to the packet. This way, even if the UDP receiver request is missed, the UDP receiver sends a fresh request after seeing a HELLO packet from the UDP sender.

By default, SGI sets the `udpcast_rexmit_hello_interval` value to 5000 (5 seconds) for UDP senders running on the admin node. On SGI ICE clusters, the RLC node UDP senders are set to 0 (disabled) because it is not normally needed when 224.0.0.1 is the RDV address and there are no VLANs being crossed. If you need to change the RDV address used by RLC nodes to serve `tmpfs` SGI ICE compute nodes, set the `udpcast_remixit_hello_interval` value so you can avoid the situation described in this topic.

You can use the `cadmin` command, and the following parameters, to display or change this value: `--show-udpcast-rexmit-hello-interval`, `--set-udpcast-rexmit-hello-interval`. The RLCs use the global value when serving `tmpfs` SGI ICE compute nodes. The admin node uses its value when it serves flat compute and RLC nodes using the UDPCast transport mechanism.

For more information, see the information about the `--rexmit-hello-interval` on the `udp-sender(1)` man page.

## `udpcast_ttl`

Sets the UDPCast time to live (TTL), which specifies the number of VLAN boundaries a request can cross.

---

**Note:** The admin node has a different default than the RLCs.

---

The defaults are as follows:

- For the admin node, the default is 2. The admin nodes serves the RLCs and flat compute nodes.
- For RLCs, the default is 1. RLCs serve only the ice compute nodes.

When `udpcast_ttl=1`, the request cannot cross a VLAN boundary. When `udpcast_ttl=2`, the request can cross one VLAN boundary. If your site has routed management networks, a data transmission might have to cross from one VLAN to another. If your site has no routed management networks, or if your site policy requires, you can set `udpcast_ttl=1` for both the RLCs and the admin node.

Values = any positive integer number.

Accepted by:

- Cluster definition file
- `cattr` command
- `discover` command

## VLAN and General Network Options

The following configuration attributes control the VLAN configuration and other aspects of the cluster network:

- "head\_vlan" on page 239
- "mcell\_network" on page 240
- "mcell\_vlan" on page 240
- "mgmt\_vlan\_end" on page 240
- "mgmt\_vlan\_start" on page 241
- "rack\_vlan\_end" on page 241
- "rack\_vlan\_start" on page 242
- "redundant\_mgmt\_network" on page 242
- "switch\_mgmt\_network" on page 242

### head\_vlan

Specifies the number of the head network VLAN. SGI recommends that you do not change this value.

Default = 1.

Range =  $1 \leq arg \leq 4096$ .

Accepted by:

- Cluster definition file

- `cattr` command

### **mcell\_network**

Specifies whether the cluster includes MCells. This value must be set to `yes` when the cluster includes MCell cooling equipment. This value can be set to `yes` or `no` for clusters that do not include MCells.

Values = `yes` (default) or `no`.

Accepted by:

- Cluster configuration tool
- Cluster definition file
- `cadmin` command
- `cattr` command

### **mcell\_vlan**

Specifies the cooling network VLAN. SGI recommends that you do not change this value.

Default = 3.

Range =  $1 \leq \text{arg} \leq 4096$ .

Accepted by:

- Cluster configuration tool
- Cluster definition file
- `cattr` command

### **mgmt\_vlan\_end**

Specifies the last flat compute node rack VLAN. Use caution when changing this value. Take care not to overlap other VLAN settings.

Default = 2500.

Range = 1 <= *arg* <= 4096.

Accepted by:

- Cluster configuration tool
- Cluster definition file
- `cattr` command

### **mgmt\_vlan\_start**

Specifies the first flat compute node rack VLAN. Use caution when changing this value. Take care not to overlap other VLAN settings.

Default = 2001.

Range = 1 <= *arg* <= 4096.

Accepted by:

- Cluster configuration tool
- Cluster definition file
- `cattr` command

### **rack\_vlan\_end**

Specifies the last SGI ICE rack VLAN. Use caution when changing this value. Take care not to overlap other VLAN settings.

Default = 1100.

Range = 1 <= *arg* <= 4096.

Accepted by:

- Cluster configuration tool
- Cluster definition file
- `cattr` command

### **rack\_vlan\_start**

Specifies the first SGI ICE rack VLAN. Use caution when changing this value. Take care not to overlap other VLAN settings.

Default = 101.

Range = 1 <= *arg* <= 4096.

Accepted by:

- Cluster configuration tool
- Cluster definition file
- `cattr` command

### **redundant\_mgmt\_network**

Specifies the default setting for the redundant management network. If no value is supplied to the `discover` command at configuration time, the installer populates the node(s) attribute(s) with this value.

Values = `yes` (default) or `no`.

Accepted by:

- Cluster configuration tool
- Cluster definition file
- `cadmin` command
- `cattr` command
- `discover` command

### **switch\_mgmt\_network**

Specifies the default setting for the switch management network. If no value is supplied to the `discover` command at configuration time, the installer populates the node(s) attribute(s) with this value. Must be set to `no` when discovering 8200/8400 racks.

Values = *yes* (default) or *no*.

Accepted by:

- Cluster configuration tool
- Cluster definition file
- `cadmin` command
- `cattr` command
- `discover` command

## Console Server Options

The admin node and the rack leader controller (RLC) nodes manage other nodes. On an SGI Rackable cluster, the admin node manages flat compute nodes. On an SGI ICE system the admin node manages RLCs and flat compute nodes, and the RLCs manage SGI ICE compute nodes.

On the management nodes (the admin node and the RLCs), there are files in the `/var/log/consoles` directory for each node that the admin node or the RLC manages. The files contain log information from the baseboard management controller (BMC) on each node under the admin node's or RLC's control.

The console server options let you control the quantity and frequency of log information that is collected. SMC logs BMC output to the `/var/log/consoles` directory. In the `/var/log/consoles` directory, there is a file for each node in the cluster. If you tune the console server options, you can limit the amount of traffic between the console and the cluster. Set these options if you need to minimize network contention.

The following options control conserver operations:

- "conserver\_logging" on page 243
- "conserver\_ondemand" on page 244

### `conserver_logging`

Specifies conserver logging. If set to *yes*, the conserver logs messages to the console via IPMItool. This feature uses some network bandwidth.

Values = `yes` (default) or `no`.

Accepted by:

- Cluster definition file
- `cadmin` command
- `cattr` command
- `discover` command

### **conserver\_ondemand**

Specifies `conserver` logging frequency. When set to `no`, logging is enabled all the time. When set to `yes`, logging is enabled only when someone is connected.

Values = `yes` or `no` (default).

Accepted by:

- Cluster definition file
- `cadmin` command
- `cattr` command
- `discover` command

## **Miscellaneous Options**

The following are miscellaneous options that control the cluster:

- "`blademon`\_scan\_interval" on page 245
- "`cluster_domain`" on page 245
- "`dhcp_bootfile`" on page 246
- "`discover_skip_switchconfig`" on page 247
- "`max_rack_irus`" on page 247
- "`mic`" on page 248

- "my\_sql\_replication" on page 248

### **blademond\_scan\_interval**

Specifies how often the blade monitor detects changes in the blades. For example, if you are changing blades or doing other blade maintenance, you could set this option to a very high value so the daemon does not run during the maintenance period.

Specifies the sleep time for the `blademond` daemon. The daemon waits the specified number of seconds in between checking if the CMC slot maps have changed.

Default = 120.

Values = can be 0 or any positive integer.

Accepted by:

- Cluster definition file
- `cadmin` command
- `cattr` command

### **cluster\_domain**

Specifies the cluster domain name. SGI recommends that users change this value.

Default = `smc-default.americas.sgi.com`.

Values = must be a standard domain name.

Accepted by:

- Cluster definition file
- Cluster configuration tool
- `cadmin` command
- `cattr`

## dhcp\_bootfile

Specifies whether to load iPXE or GRUB2 first when a node is first configured. By default, `dhcp_bootfile=grub2`, which means that GRUB2 loads first. If you specify `dhcp_bootfile=ipxe`, however, the server boot agent loads iPXE instead of GRUB2, and then iPXE loads GRUB2.

In some cases, a node can fail to boot over the network with the default settings. For example, if a node hangs when it tries to load the kernel and `initrd` during the boot from its system disk, modify the DHCP boot file setting to load iPXE first and then have iPXE load GRUB2. You might need to specify iPXE for just certain nodes. You can use the `discover` or `cadmin` command to specify the new boot order, as follows:

- To use the `discover` command and configure `service0` with an iPXE boot first, include the following on the command line:

```
discover --service 0,dhcp_bootfile=ipxe other options
```

For *other options*, specify any additional options you need for the configuration. For example, you can specify a cluster configuration file that contains many other specifications. The `service 0,dhcp_bootfile=ipxe` argument on this `discover` command line overrides the default booting option in the cluster configuration file.

- To use the `cadmin` command and specify that iPXE boot first on `service0`, include the following on the command line:

```
cadmin --set-dhcp-bootfile --node service0 ipxe
```

Another useful `cadmin` command option is `--show-dhcp-bootfile`, which returns the boot file specification. For example:

```
cadmin --set-dhcp-bootfile --node service0
grub2
```

Values = `grub2` (default) or `ipxe`.

Accepted by:

- Cluster definition file
- `cadmin` command
- `cattr` command
- `discover` command

**discover\_skip\_switchconfig**

Signals the installer to omit the switch configuration steps. When set to `yes`, the installer does not configure the switches. Set this to `yes` when you want to perform a quick configuration change, but you do not need to update the switch configuration. This value is not saved in the cluster definition file, but it can be specified there.

Values = `yes` or `no` (default).

Accepted by:

- Cluster definition file
- `cadmin` command
- `cattr` command
- `discover` command

**max\_rack\_irus**

Specifies the maximum number of IRUs in the cluster. When you use the `discover` command, the installer autopopulates the database with this value during the rack configuration. This value is not saved in the cluster definition file, but it can be specified there.

Default = 8.

Range =  $1 \leq arg \leq 16$ .

Accepted by:

- Cluster definition file
- Cluster configuration tool
- `cadmin` command
- `cattr` command
- `discover` command

### **mic**

Specifies the default number of Intel Phi (MIC) devices on each compute node. The default is 4 or fewer for flat compute nodes. The default is 2 or fewer for SGI ICE compute nodes. This value is not saved in the cluster definition file, but it can be specified there.

Values = 0 (default), 1, 2, 3, 4.

Accepted by:

- Cluster definition file
- `cattr` command
- `discover` command

### **my\_sql\_replication**

Specifies if MySQL replication is disabled or enabled on any node. If set to `yes`, MySQL replication is enabled. Unless otherwise specified, this value is set on each node when the `discover` command runs. If you want to enable MySQL replication for any node, you need to make sure that MySQL replication is set to `yes` on the admin node.

For more information about MySQL replication, see the following:

*SGI Management Center (SMC) Administration Guide for Clusters*

Values = `yes` (default) or `no`.

Accepted by:

- Cluster configuration tool
- Cluster definition file
- `cadmin` command
- `cattr` command
- `discover` command

---

## Index

### C

Configure backup DNS server, 146

### D

DHCP option code, 130

dhcp options

    changing, 130

disabling InfiniBand switch monitoring, 147

### I

InfiniBand configuration, 147

    disabling InfiniBand switch monitoring, 147

initial configuration of a RHEL 6 admin node, 69

initial configuration of a RHEL 7 admin node, 65

### O

overview, 1

### S

system management software, 1

system overview, 1

### T

troubleshooting

    frequently asked questions, 176

    initial system setup, 176