



SGI® InfiniteData™ Cluster for Hadoop®  
Getting Started Guide

007-6313-002

---

## COPYRIGHT

© 2014 Silicon Graphics International Corp. All rights reserved; provided portions may be copyright in third parties, as indicated elsewhere herein. No permission is granted to copy, distribute, or create derivative works from the contents of this electronic documentation in any manner, in whole or in part, without the prior written permission of SGI.

---

## LIMITED RIGHTS LEGEND

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

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

---

## TRADEMARKS AND ATTRIBUTIONS

Silicon Graphics, SGI, the SGI logo, InfiniteData, Rackable, and Supportfolio are trademarks or registered trademarks of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries worldwide.

Cloudera is a trademark of Cloudera Inc. in the USA and other countries. Firefox is a registered trademark of the Mozilla Foundation. Hadoop is a registered trademark of Apache Software Foundation. Intel and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. Java is a registered trademark of Oracle and/or one of its affiliates. Red Hat and all Red Hat-based trademarks are trademarks or registered trademarks of Red Hat, Inc. in the United States and other countries.

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

---

## Record of Revision

<b>Version</b>	<b>Description</b>
001	January 2014 Initial printing.
002	June 2014 Updated to include support of CDH 5 software and other miscellaneous changes.



---

# Contents

	<b>About This Guide</b>	<b>vii</b>
	Audience	vii
	Related Publications	viii
	Product Support	ix
	Reader Comments	x
<b>1</b>	<b>Overview</b>	<b>1</b>
	The 10GigE Implementation	2
	Hardware	2
	Servers	2
	Network Hardware	4
	Configurations	5
	Full-Rack (46U)	6
	Multi-Rack (First and Subsequent Racks)	7
	Network Topology	9
	Node Level	9
	Rack Level for Single-Rack Configuration	11
	Multi-Rack Data Network	12
	Management Network	14
	Software	16
<b>2</b>	<b>Cluster Startup</b>	<b>17</b>
	Accepting End-User License Agreements (EULAs)	17
	Configuring and Starting SGI Management Center	18
	Starting the Cluster for the First Time	18
	Accessing Cloudera Manager	19
	Starting Hadoop Cluster Services	20
	Querying Hosts in the Cluster	23
	Enabling Cloudera Manager Enterprise Features	24
	Re-Imaging the Server Nodes	25



---

## About This Guide

This guide provides an overview of the SGI® Hadoop® Reference Implementation based on the SGI® InfiniteData™ Cluster platform along with getting-started instructions for this implementation. This guide consists of the following chapters:

- Chapter 1, “Overview,” provides an overview of the SGI Hadoop solution.
- Chapter 2, “Cluster Startup,” describes licensing and Hadoop specifics for configuring cluster management and monitoring.

## Audience

This guide is written for the system administrators of the Hadoop cluster and developers. The guide assumes the reader is familiar with clusters, the Hadoop technology, and business intelligence applications.

## Related Publications

The following SGI documents are relevant to your Hadoop solution:

- *SGI InfiniteData Cluster Hardware User Guide* (007-6308-xxx)
- *SGI Rackable C1110-RP6 System User Guide* (007-5843-xxx)
- *SGI Management Center Quick Start Guide* (007-5672-xxx)
- *SGI Management Center (SMC) Installation and Configuration* (007-5643-xxx)
- *SGI Management Center (SMC) System Administrator's Guide* (007-5642-xxx)
- *SGI InfiniteStorage Server 3000 (ISS3000) User's Guide* (007-5721-xxx)

You can obtain SGI documentation in the following ways:

- Refer to the SGI Technical Publications Library (TPL) at <http://docs.sgi.com>. Various formats are available. The TPL contains the most recent and most comprehensive set of books and man pages.

To get the latest revision of a document on the TPL, use the core publication number as your search string. For example, use 007-1234 as your search string to get the latest version of the document with part number 007-1234-xxx.

- Refer to the SGI Supportfolio™ webpage for release notes and other documents whose access require a support contract. See [“Product Support” on page ix](#).

---

**Note:** For information about third-party system components, see the documentation provided by the manufacturer/supplier.

---



## Product Support

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

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

Be sure to have the following information before you call Technical Support:

- Product serial number
- Product model name and number
- Applicable error messages
- Add-on boards or hardware
- Third-party hardware or software
- Operating system type and revision level

## Reader Comments

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

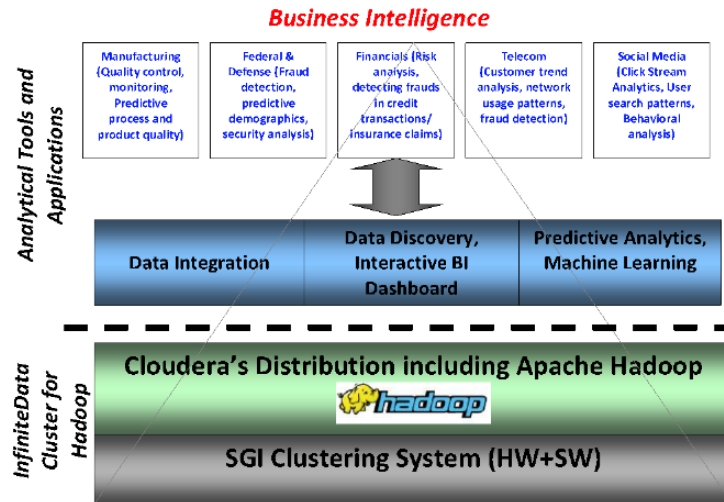
You can contact SGI in any of the following ways:

- Send e-mail to the following address: [techpubs@sgi.com](mailto: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.

# Overview



**Figure 1-1** SGI Hadoop Business Intelligence Ecosystem

The SGI Hadoop Reference Implementations provide pre-defined and pre-certified Hadoop solutions with these features:

- Pre-defined and pre-certified configurations
- High performance
- High availability
- Power optimization
- Capability of running business intelligence (BI) applications directly atop Hadoop (See Figure 1-1.)

This SGI Hadoop Reference Implementation is 10GigE-based and uses the Intel® Xeon® E5-2600 v2 Processor Series. This chapter describes this implementation using the following topics:

- “The 10GigE Implementation” on page 2
- “Software” on page 16

## The 10GigE Implementation

This section describes the 10GigE implementation using the following topics:

- “Hardware” on page 2
- “Configurations” on page 5
- “Network Topology” on page 9

## Hardware

This section describes the hardware used in the 10GigE-based implementation: first, the servers and then the network hardware.

### Servers

The 10GigE-based SGI Hadoop cluster employs SGI Rackable™ C1110 and SGI InfiniteData Cluster 3212 (IDC3212) servers; a C1110 server and an IDC3212 server are shown in [Figure 1-2](#) and [Figure 1-3](#), respectively. This section describes the SGI servers that are used in the 10GigE-based SGI Hadoop cluster, their function in the Hadoop paradigm, and their specifications.



**Figure 1-2** An SGI Rackable C1110 Server



**Figure 1-3** An SGI IDC3212 Server

Table 1-1 describes the SGI Hadoop Reference Implementations with SGI 10GigE-based servers with the Intel Xeon Processor E5-2600 v2 Series.

**Table 1-1** SGI Hadoop 10GigE-Based Half-Depth Servers–Intel Xeon Processor E5-2600 v2 Series


SGI Server	Conventional Node Type	Hadoop Node Type	Specifications
C1110-RP6	Master nodes	NameNode, Standby NameNode, ResourceManager	<ul style="list-style-type: none"> <li>– 2x Intel Xeon Processor E5-2630 v2 (2.6 GHz , 6-core)</li> <li>– 8x 8GB 1.5v 1866MHz DIMMs (64GB memory)</li> <li>– 4x 3.5” 4TB 7200 rpm SATA 6Gb/s drives in RAID 10 configuration</li> <li>– 1x Dual-port 10GigE NIC</li> <li>– Redundant power supply</li> </ul>
IDC3212-RP4	Compute/Slave nodes	DataNodes, NodeManagers	<ul style="list-style-type: none"> <li>– 2x Intel Xeon Processor E5-2630 v2 (2.6 GHz, 6-core)</li> <li>– 8x 8GB 1.5v 1866MHz DIMMs (64GB memory)</li> <li>– 12x 3.5” 4TB 7200 rpm SATA drives</li> <li>– 1x Dual-port 10GigE NIC</li> </ul>
C1110-RP6		Application Node	<ul style="list-style-type: none"> <li>– 2x Intel Xeon Processor E5-2680 v2 (2.8 GHz, 10-core)</li> <li>– 16x 8GB 1.5v 1866MHz DIMMs (128GB memory)</li> <li>– 4x 3.5” 4TB 7200 rpm SAS 6Gb/s drives in RAID 10 configuration</li> <li>– 1x Dual-port 10GigE NIC</li> <li>– Redundant power supply</li> </ul>

## Network Hardware

The network hardware consists of the following components:

- 1 Edge-corE ECS4610-50T 48-port GigE switch per rack
- 2 Extreme Networks Summit X670v 10-GigE switches per rack
- Mellanox SX1012 or Mellanox SX1036 40 GigE spine switches (quantity dependent on number of racks)

## Configurations



**Figure 1-4** Data Capacity for Various Rack Configurations

The SGI Hadoop Cluster is available in single-rack and multi-rack configurations. [Figure 1-4](#) shows the range of data capacity for the configurations. This section describes the full-rack and multi-rack configurations.

Full-Rack (46U)

InfiniteData Cluster Server Layout

sgl	Description	Image	RackU	Description
		2U Blank	45	
			44	
	48port 10GigE SGIMC Management Switch		43	
	Standby NameNode/SGIMC Head Node		42	
	NameNode		41	
	ResourceManager Node		40	
	Application Node		39	
	48port 10GigE Hadoop Data Switch		38	
	48port 10GigE Hadoop Data Switch		37	
			36	
	DataNode/NodeManager (r01n35)		35	DataNode/NodeManager (r01n36)
			34	
	DataNode/NodeManager (r01n33)		33	DataNode/NodeManager (r01n34)
			32	
	DataNode/NodeManager (r01n31)		31	DataNode/NodeManager (r01n32)
			30	
	DataNode/NodeManager (r01n29)		29	DataNode/NodeManager (r01n30)
			28	
	DataNode/NodeManager (r01n27)		27	DataNode/NodeManager (r01n28)
			26	
	DataNode/NodeManager (r01n25)		25	DataNode/NodeManager (r01n26)
			24	
	DataNode/NodeManager (r01n23)		23	DataNode/NodeManager (r01n24)
			22	
	DataNode/NodeManager (r01n21)		21	DataNode/NodeManager (r01n22)
			20	
	DataNode/NodeManager (r01n19)		19	DataNode/NodeManager (r01n20)
			18	
	DataNode/NodeManager (r01n17)		17	DataNode/NodeManager (r01n18)
			16	
	DataNode/NodeManager (r01n15)		15	DataNode/NodeManager (r01n16)
			14	
	DataNode/NodeManager (r01n13)		13	DataNode/NodeManager (r01n14)
			12	
	DataNode/NodeManager (r01n11)		11	DataNode/NodeManager (r01n12)
			10	
	DataNode/NodeManager (r01n09)		9	DataNode/NodeManager (r01n10)
			8	
	DataNode/NodeManager (r01n07)		7	DataNode/NodeManager (r01n08)
			6	
	DataNode/NodeManager (r01n05)		5	DataNode/NodeManager (r01n06)
			4	
	DataNode/NodeManager (r01n03)		3	DataNode/NodeManager (r01n04)
			2	
	DataNode/NodeManager (r01n01)		1	DataNode/NodeManager (r01n02)

**Figure 1-5** Full-Rack Configuration

Figure 1-5 describes a full-rack configuration. The rack consists of the following:

- 1 SGI Management Center node/Standby NameNode
- 1 NameNode
- 1 ResourceManager
- 1 Application node
- 36 DataNodes/NodeManagers
- 2 48-port 10GigE stacked Hadoop data network switches
- 1 SGI Management Center network switch



## Multi-Rack (First and Subsequent Racks)

InfiniteData Cluster Server Layout First Rack




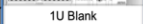
































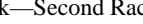

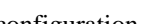
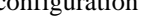
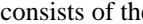



SGI	Description	Image	U	Description
		1U Blank	45	
	48port 10GigE Hadoop Data Switch (SPINE)		44	48port 10GigE Hadoop Data Switch (SPINE)
	48port 10GigE SGIMC Management Switch		43	
	Standby NameNode/SGIMC Head Node		42	
	NameNode		41	
	ResourceManager Node		40	
	Application Node		39	
	48port 10GigE Hadoop Data Switch		38	
	48port 10GigE Hadoop Data Switch		37	
			36	
	DataNode/NodeManager (r01n35)		35	DataNode/NodeManager (r01n36)
			34	
	DataNode/NodeManager (r01n33)		33	DataNode/NodeManager (r01n34)
			32	
	DataNode/NodeManager (r01n31)		31	DataNode/NodeManager (r01n32)
			30	
	DataNode/NodeManager (r01n29)		29	DataNode/NodeManager (r01n30)
			28	
	DataNode/NodeManager (r01n27)		27	DataNode/NodeManager (r01n28)
			26	
	DataNode/NodeManager (r01n25)		25	DataNode/NodeManager (r01n26)
			24	
	DataNode/NodeManager (r01n23)		23	DataNode/NodeManager (r01n24)
			22	
	DataNode/NodeManager (r01n21)		21	DataNode/NodeManager (r01n22)
			20	
	DataNode/NodeManager (r01n19)		19	DataNode/NodeManager (r01n20)
			18	
	DataNode/NodeManager (r01n17)		17	DataNode/NodeManager (r01n18)
			16	
	DataNode/NodeManager (r01n15)		15	DataNode/NodeManager (r01n16)
			14	
	DataNode/NodeManager (r01n13)		13	DataNode/NodeManager (r01n14)
			12	
	DataNode/NodeManager (r01n11)		11	DataNode/NodeManager (r01n12)
			10	
	DataNode/NodeManager (r01n09)		9	DataNode/NodeManager (r01n10)
			8	
	DataNode/NodeManager (r01n07)		7	DataNode/NodeManager (r01n08)
			6	
	DataNode/NodeManager (r01n05)		5	DataNode/NodeManager (r01n06)
			4	
	DataNode/NodeManager (r01n03)		3	DataNode/NodeManager (r01n04)
			2	
	DataNode/NodeManager (r01n01)		1	DataNode/NodeManager (r01n02)

**Figure 1-6** Multi-Rack—First Rack

Figure 1-6 shows the first rack of a multi-rack configuration. The rack consists of the following:

- 1 SGI Management Center node/Standby NameNode
- 1 NameNode
- 1 ResourceManager
- 1 Application node
- 36 DataNodes/NodeManagers
- 2 48-port 10GigE stacked Hadoop data network switches
- 2 40GigE network spine switches
- 1 SGI Management Center network switch

InfiniteData Cluster Server Layout Second Rack

	Description	Image	RackU	Description
	Optional Additional Spine Switches	1U Blank	45	Optional Additional Spine Switches
	48port 10GigE SGIMC Management Switch		44	
	48port 10GigE Hadoop Data Switch (SPINE)		43	
	48port 10GigE Hadoop Data Switch (SPINE)		42	
		1U Blank	41	
			40	
	DataNode/NodeManager (r01n39)		39	DataNode/NodeManager (r01n40)
			38	
	DataNode/NodeManager (r01n37)		37	DataNode/NodeManager (r01n38)
			36	
	DataNode/NodeManager (r01n35)		35	DataNode/NodeManager (r01n36)
			34	
	DataNode/NodeManager (r01n33)		33	DataNode/NodeManager (r01n34)
			32	
	DataNode/NodeManager (r01n31)		31	DataNode/NodeManager (r01n32)
			30	
	DataNode/NodeManager (r01n29)		29	DataNode/NodeManager (r01n30)
			28	
	DataNode/NodeManager (r01n27)		27	DataNode/NodeManager (r01n28)
			26	
	DataNode/NodeManager (r01n25)		25	DataNode/NodeManager (r01n26)
			24	
	DataNode/NodeManager (r01n23)		23	DataNode/NodeManager (r01n24)
			22	
	DataNode/NodeManager (r01n21)		21	DataNode/NodeManager (r01n22)
			20	
	DataNode/NodeManager (r01n19)		19	DataNode/NodeManager (r01n20)
			18	
	DataNode/NodeManager (r01n17)		17	DataNode/NodeManager (r01n18)
			16	
	DataNode/NodeManager (r01n15)		15	DataNode/NodeManager (r01n16)
			14	
	DataNode/NodeManager (r01n13)		13	DataNode/NodeManager (r01n14)
			12	
	DataNode/NodeManager (r01n11)		11	DataNode/NodeManager (r01n12)
			10	
	DataNode/NodeManager (r01n09)		9	DataNode/NodeManager (r01n10)
			8	
	DataNode/NodeManager (r01n07)		7	DataNode/NodeManager (r01n08)
			6	
	DataNode/NodeManager (r01n05)		5	DataNode/NodeManager (r01n06)
			4	
	DataNode/NodeManager (r01n03)		3	DataNode/NodeManager (r01n04)
			2	
	DataNode/NodeManager (r01n01)		1	DataNode/NodeManager (r01n02)

**Figure 1-7** Multi-Rack—Second Rack and Beyond

Figure 1-7 describes the configuration of the second rack (and subsequent racks) of a multi-rack configuration. Each rack consists of the following:


- 40 DataNodes/NodeManagers
- 2 48-port 10GigE stacked Hadoop data network switches
- 1 SGI Management Center network switch

## Network Topology


This section illustrates the network topology from the most granular level (node level) to the top level (multi-rack level) and the topology of the management network:

- “Node Level” on page 9
- “Rack Level for Single-Rack Configuration” on page 11
- “Multi-Rack Data Network” on page 12
- “Management Network” on page 14

### Node Level

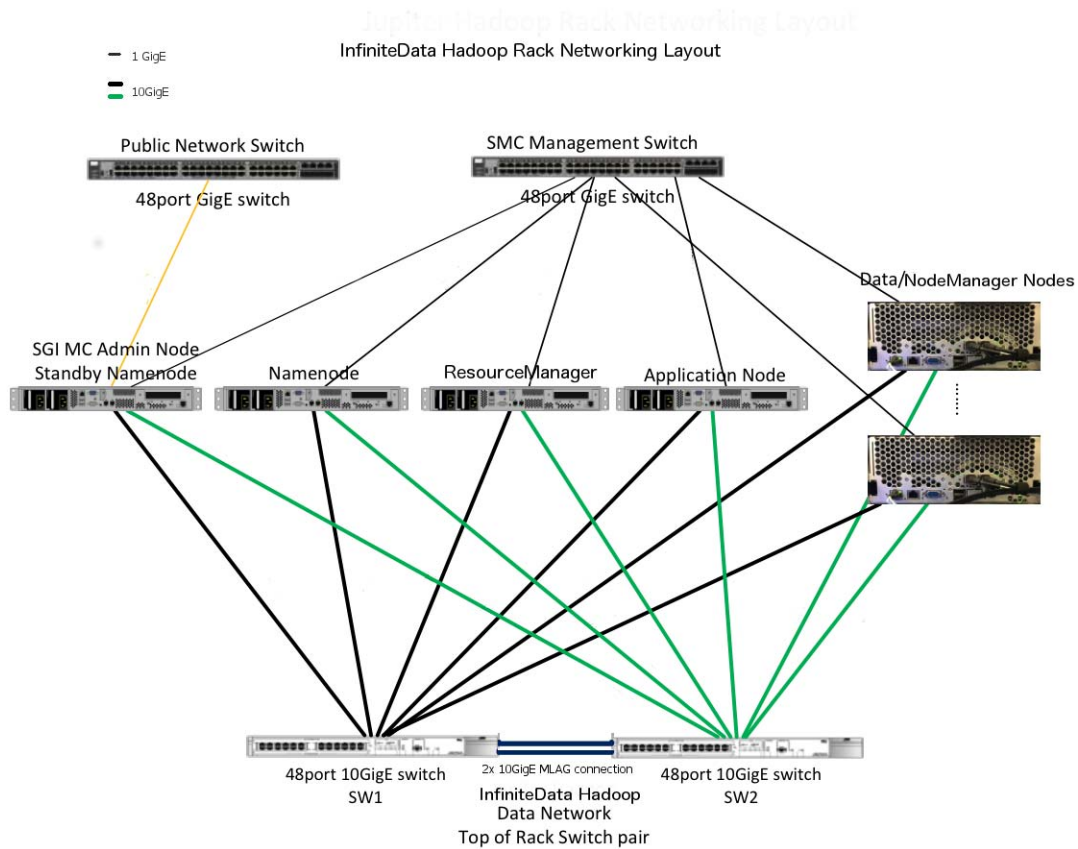


**Figure 1-8** Network Topology—Master Node



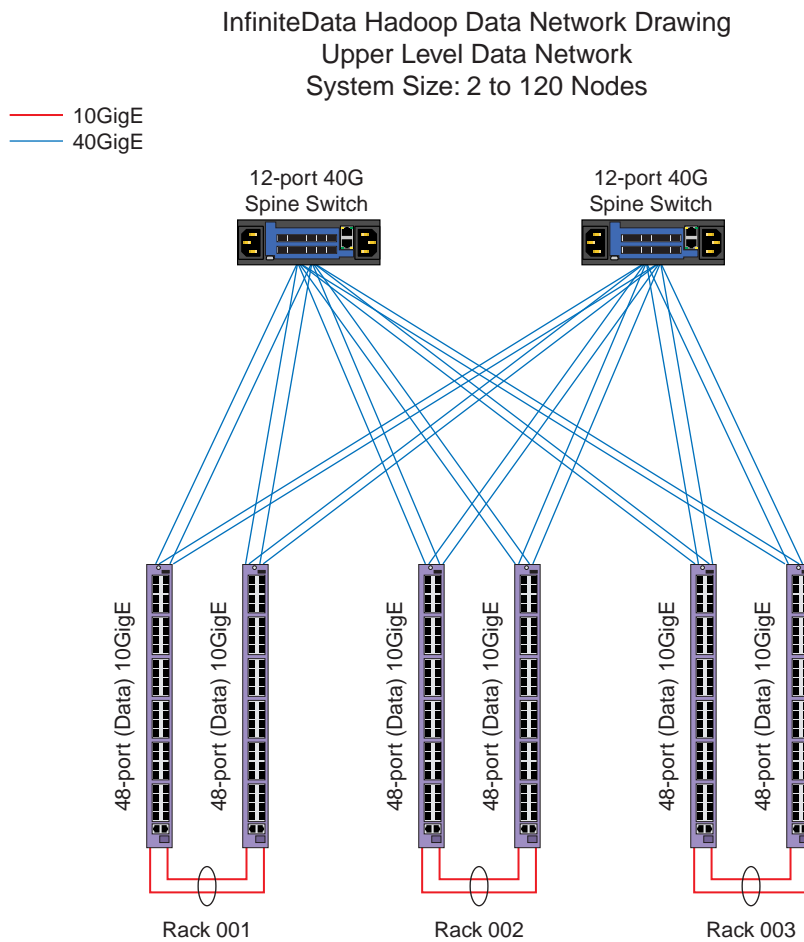
**Figure 1-9** Network Topology—Slave Node

### Rack Level for Single-Rack Configuration

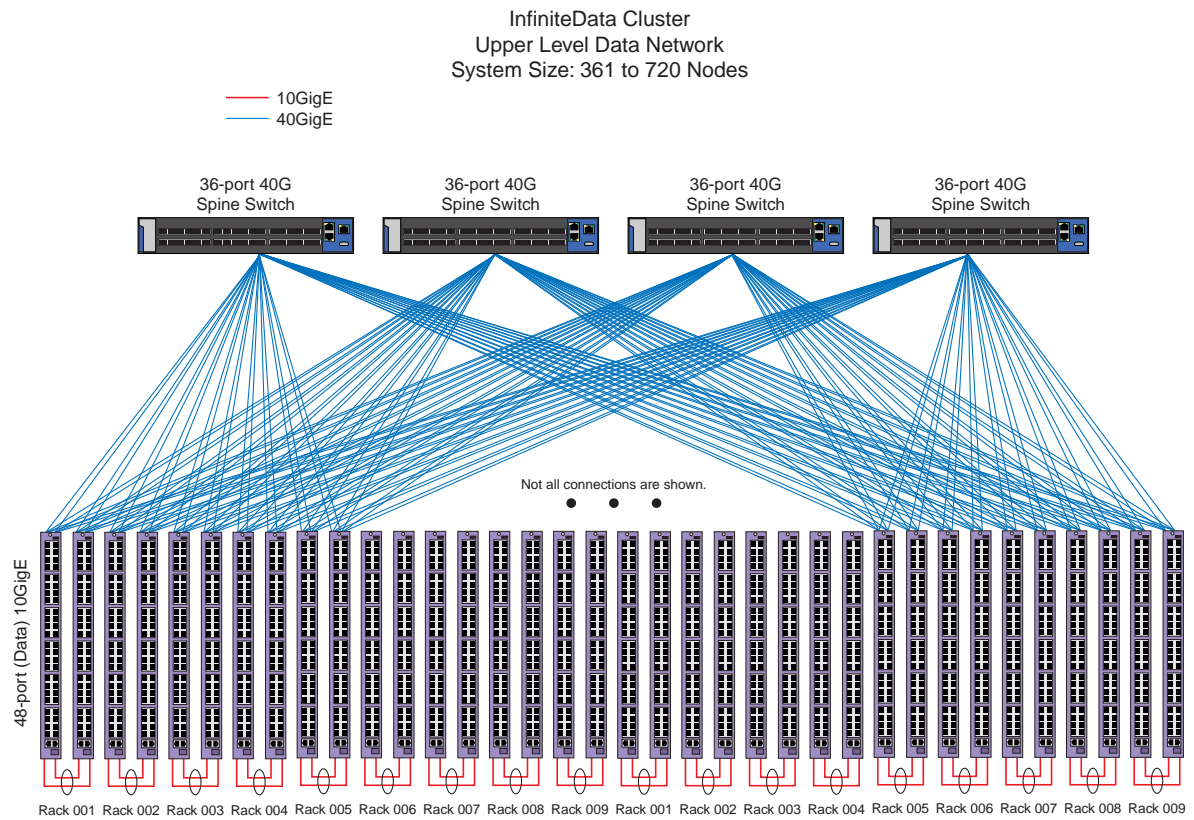


**Figure 1-10** Network Topology—Rack Level for Single Rack

### Multi-Rack Data Network

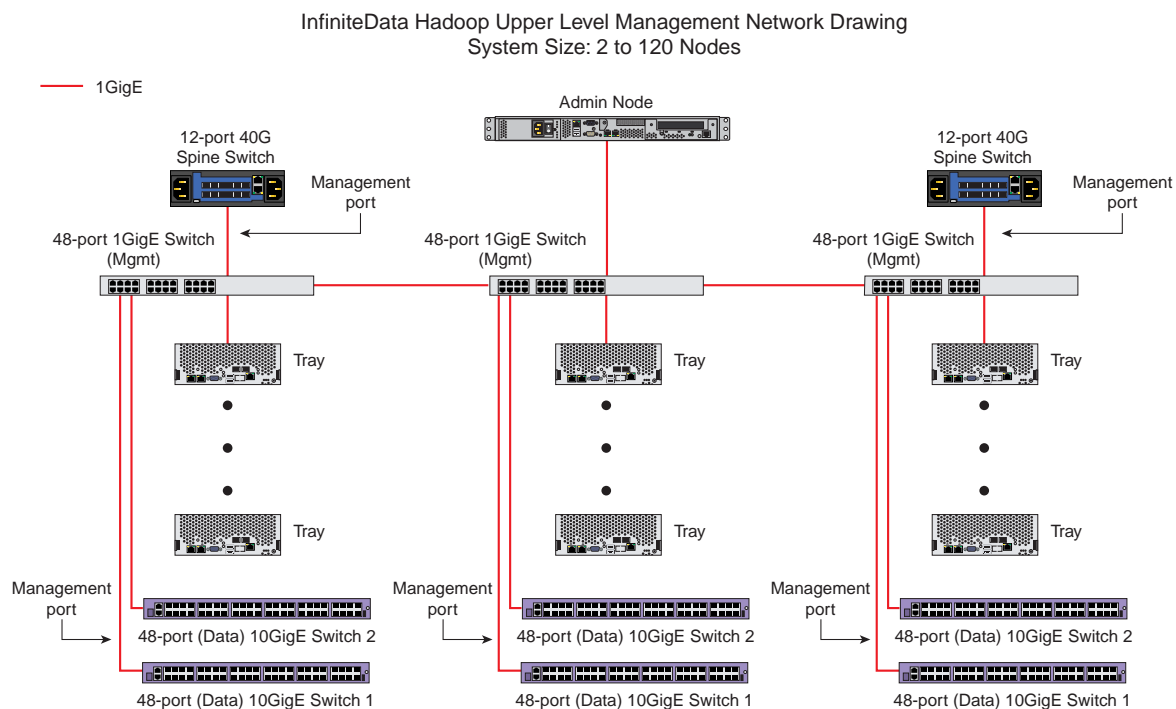


**Figure 1-11** Network Topology—Multi-Rack Data Network (2 to 120 Nodes)



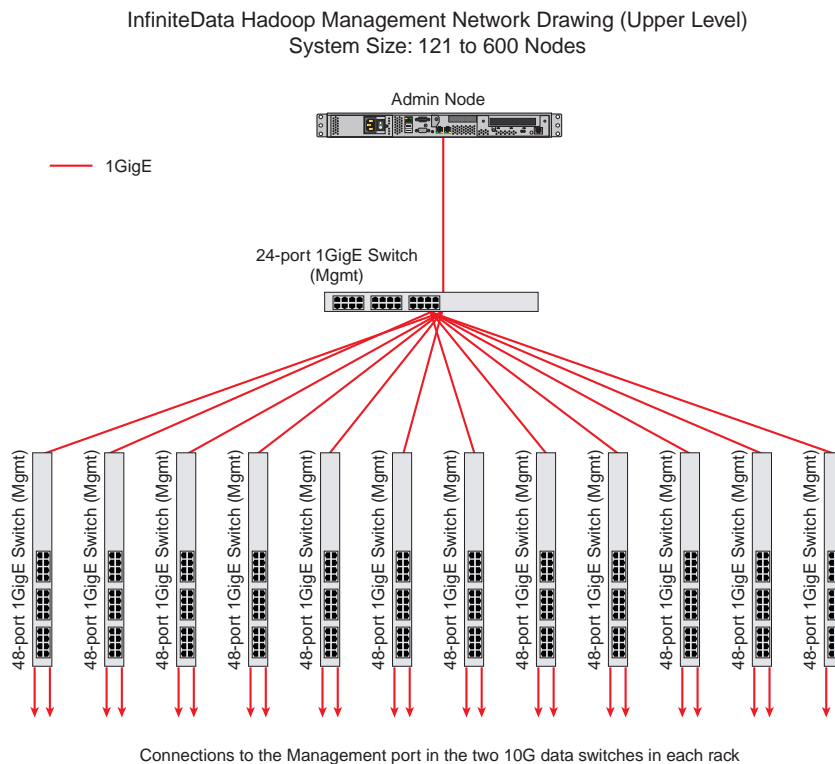
**Figure 1-12** Network Topology—Multi-Rack Data Network (361 to 720 Nodes)

## Management Network



**Figure 1-13** Network Topology—Management Network (2 to 120 Nodes)





**Figure 1-14** Network Topology—Management Network (121 to 600 Nodes)

## Software

The software stack for the SGI Hadoop solution consists of the following components:

- Red Hat® Enterprise Linux (RHEL) 6.x
- Cloudera™ distribution Apache Hadoop 5.x
- Cloudera Manager 5.x
- SGI Management Center 1.7

Figure 1-15 shows the SGI Hadoop software stack.

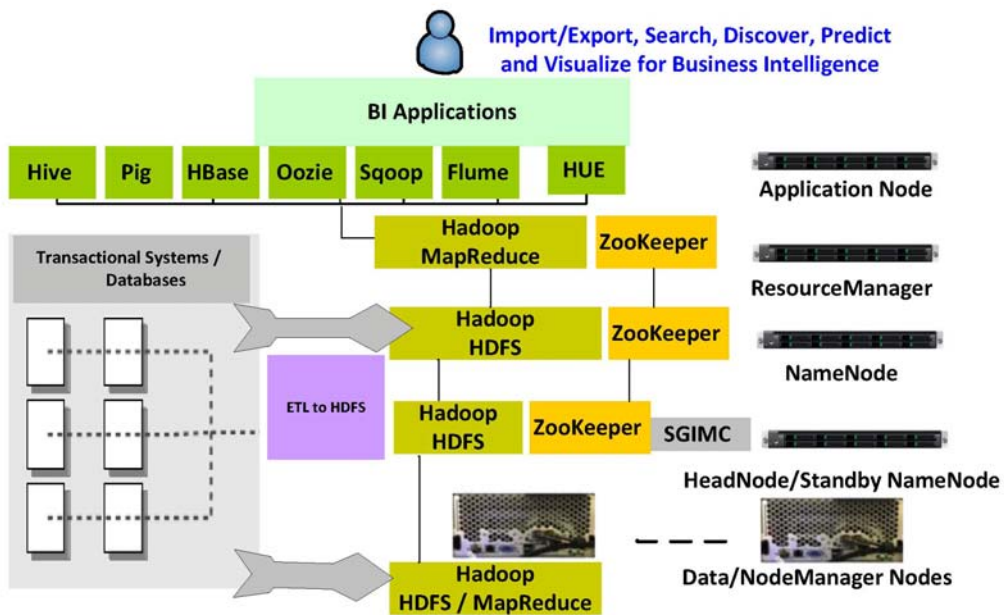


Figure 1-15 SGI Hadoop Software Stack

## Cluster Startup

This chapter describes the broad steps for starting the SGI Hadoop cluster:

- “Accepting End-User License Agreements (EULAs)” on page 17
- “Configuring and Starting SGI Management Center” on page 18
- “Starting the Cluster for the First Time” on page 18
- “Accessing Cloudera Manager” on page 19
- “Starting Hadoop Cluster Services” on page 20
- “Querying Hosts in the Cluster” on page 23
- “Enabling Cloudera Manager Enterprise Features” on page 24
- “Re-Imaging the Server Nodes” on page 25

### Accepting End-User License Agreements (EULAs)

The SGI Hadoop solution contains third-party software whose end-user license agreements you must read and accept. One such product is the Java® Distribution Kit (JDK). The JDK copyright and third-party license agreement can be found on any of the cluster nodes in directory `/usr/share/doc/java-1.6.0-sun-devel-1.6.0.25`. Read and accept the conditions in the license agreement.

If you get trial versions of business intelligence applications, they also will require you to accept their EULAs.

## Configuring and Starting SGI Management Center

You will use the SGI Management Center to perform the conventional platform management functions (power control, environmental monitoring, provisioning, etc.) for the Hadoop cluster. To configure and start the SGI Management Center, you will need to follow the instructions in the *SGI Management Center Quick Start Guide* and configure the Hadoop servers as specified in [Table 2-1](#).

**Table 2-1** Hostnames for SGI Hadoop Servers

Daemon	Hostname	Hadoop Data Network Hostname
NameNode	sgi-nn	sgi-nn-data
Standby NameNode	sgi-snn	sgi-snn-data
ResourceManager	sgi-rm	sgi-rm-data
Application Node	sgi-app	sgi-app-data
DataNodes & NodeManagers	r[rack#]n[node#]	r[rack#]n[node#]-data

## Starting the Cluster for the First Time

Use the following steps to start the SGI Hadoop cluster the first time.

1. Power on the head node of the cluster.
2. Use SGI Management Center to start the nodes in the cluster.
  - a. Log in as `root`.
  - b. Start the SGI Management Center with the following command:
 

```
# mgrclient
```
  - c. Within the Management GUI, select the nodes to start, right-click, and select **Power > On**.
  - d. Start the nodes in the following order:
    - i. `sgi-app`
    - ii. `sgi-nn`
    - iii. `sgi-rm`
    - iv. Compute/Slave nodes in the Compute group

## Accessing Cloudera Manager

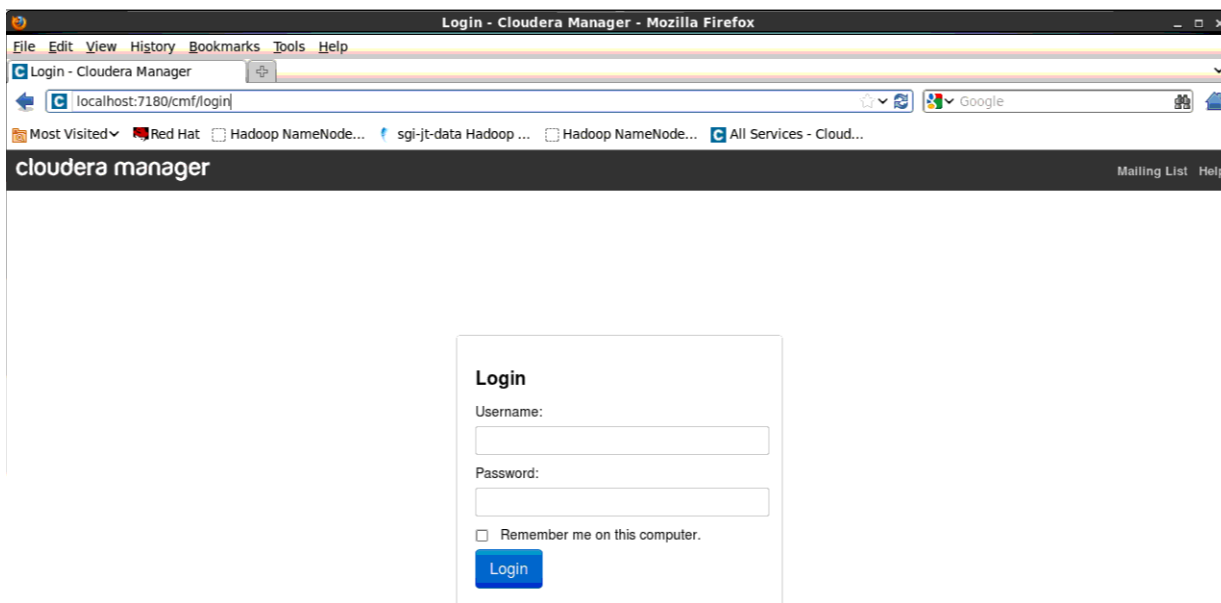
You will use Cloudera Manager for the application management functions of the Hadoop cluster. To access Cloudera Manager, do the following:

1. Open the web browser on the cluster head node.
2. Enter the URL `http://localhost:7180` to access the Cloudera Manager or use the Firefox® bookmark for the Cloudera Manager.

The login screen, shown in [Figure 2-1](#), should appear.

3. Enter your Cloudera Manager login username and password.

The default is `admin/admin`.



**Figure 2-1** Cloudera Manager Login Screen

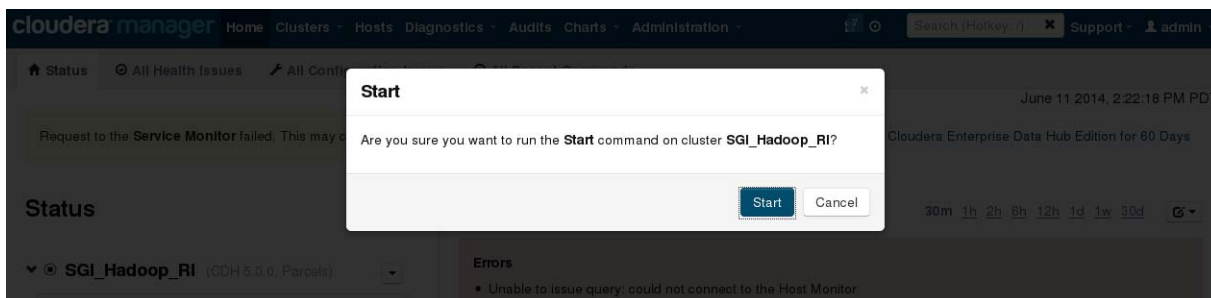
## Starting Hadoop Cluster Services

After a successful login, the **Cloudera Manager** home page appears. Start the Hadoop cluster by clicking **Start** in the **SGI\_Hadoop\_RI** Actions drop down box, as shown in [Figure 2-2](#).

The screenshot displays the Cloudera Manager web interface. The top navigation bar includes 'Home', 'Clusters', 'Hosts', 'Diagnostics', 'Audits', 'Charts', and 'Administration'. A search bar and user profile 'admin' are also visible. The main content area is divided into 'Status' and 'Charts' sections. In the 'Status' section, the 'SGI\_Hadoop\_RI' cluster is selected, and its 'Actions' menu is open, showing options like 'Start', 'Stop', 'Restart', and 'Rolling Restart'. The 'Charts' section shows an 'Errors' alert: 'Unable to issue query: the Host Monitor is not running'. Below this, there are three chart placeholders for 'Cluster CPU', 'Cluster Disk IO', and 'Cluster Network IO', all displaying 'QUERY ERROR'.

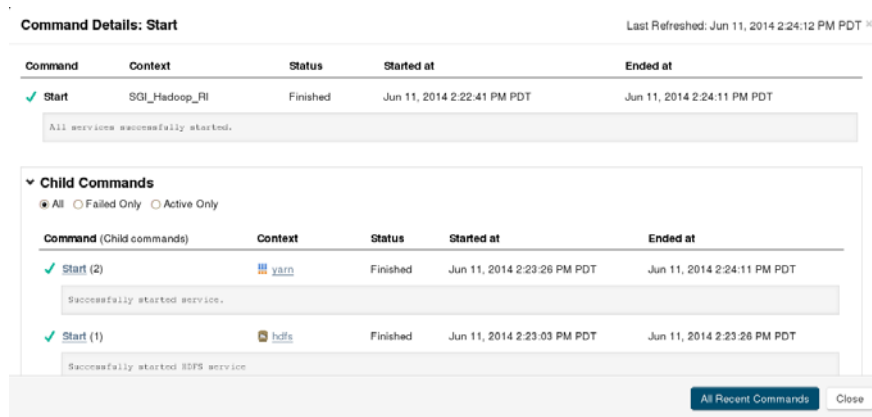
**Figure 2-2** Starting the Hadoop Cluster

After you initialize the startup of cluster services, Cluster Manager prompts you to start the Hadoop cluster, as shown in [Figure 2-3](#). Select **Start**.



**Figure 2-3** Starting the Hadoop Cluster

After you select **Start**, Cloudera Manager will display the **Command Details** screen to show the status of the action, as shown in [Figure 2-4](#).



**Figure 2-4** Successful Startup Details

After the SGI\_Hadoop\_RI cluster is started, click **Start** in the **mgmt Actions** drop-down box (See Figure 2-5.) to start the Cloudera Management Services.

The screenshot displays the Cloudera Manager web interface. The top navigation bar includes 'Home', 'Clusters', 'Hosts', 'Diagnostics', 'Audits', 'Charts', and 'Administration'. A search bar and user profile 'admin' are also visible. The main content area is divided into two columns: 'Status' and 'Charts'.  
In the 'Status' column, the 'SGI\_Hadoop\_RI' cluster is shown with a red circle icon and 'Unknown Health'. Below it, a table lists services: hdfs, yarn, and zookeeper, all with red status indicators. Underneath, the 'Cloudera Management Service' is listed with a blue 'mgmt' icon. A dropdown menu for 'mgmt Actions' is open, showing 'Start', 'Stop', and 'Restart' options.  
In the 'Charts' column, an 'Errors' section contains a message: 'Unable to issue query: the Host Monitor is not running'. Below this, three chart placeholders are visible: 'Cluster CPU', 'Cluster Disk IO', and 'Cluster Network IO', each displaying a 'QUERY ERROR' message.

**Figure 2-5** Starting Cloudera Management Services



## Querying Hosts in the Cluster

To view all hosts running in the Hadoop cluster, click **Hosts** on the top bar of the window. Figure 2-6 shows the **All Hosts** screen.

**8 Hosts:** 8 Good Health

Showing 1 to 8 of 8 entries

Name	IP	Rack	CDH Version	Cluster	Roles	Status	Last Heartbeat	Maintena
<a href="#">r05n13-data.default.domain</a>	172.16.5.13	/r05	CDH 5	SGI_Hadoop_RI	▶ 2 Role(s)	Good Health	7.97s ago	
<a href="#">r05n14-data.default.domain</a>	172.16.5.14	/r05	CDH 5	SGI_Hadoop_RI	▶ 2 Role(s)	Good Health	7.95s ago	
<a href="#">r05n16-data.default.domain</a>	172.16.5.16	/r05	CDH 5	SGI_Hadoop_RI	▶ 2 Role(s)	Good Health	8s ago	
<a href="#">r05n17-data.default.domain</a>	172.16.5.17	/r05	CDH 5	SGI_Hadoop_RI	▶ 2 Role(s)	Good Health	8.02s ago	
<a href="#">sqi-app-data.default.domain</a>	172.16.5.15	/r01	CDH 5	SGI_Hadoop_RI	▶ 1 Role(s)	Good Health	7.92s ago	
<a href="#">sqi-it-data.default.domain</a>	172.16.5.12	/r01	CDH 5	SGI_Hadoop_RI	▶ 4 Role(s)	Good Health	606ms ago	
<a href="#">sqi-nn-data.default.domain</a>	172.16.5.11	/r01	CDH 5	SGI_Hadoop_RI	▶ 4 Role(s)	Good Health	456ms ago	
<a href="#">sqi-snn-data.default.domain</a>	172.16.11.2	/r01	CDH 5	SGI_Hadoop_RI	▶ 9 Role(s)	Good Health	527ms ago	

Showing 1 to 8 of 8 entries

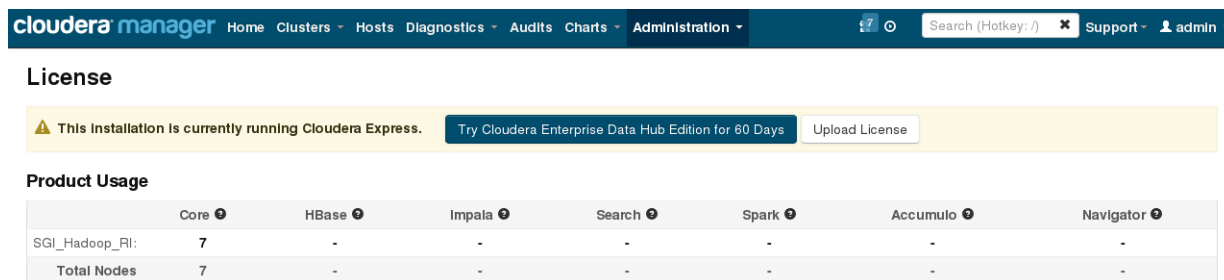
Figure 2-6 All Hosts Screen

## Enabling Cloudera Manager Enterprise Features

If you have purchased the Cloudera Enterprise license, you can enter the key from the **License** screen:

**Administration** —> **License**

Figure 2-7 shows the **License** screen.



**License**

⚠ This installation is currently running Cloudera Express. [Try Cloudera Enterprise Data Hub Edition for 60 Days](#) [Upload License](#)

**Product Usage**

	Core	HBase	Impala	Search	Spark	Accumulo	Navigator
SGL_Hadoop_RL:	7	-	-	-	-	-	-
Total Nodes	7	-	-	-	-	-	-

**Figure 2-7** License Screen

## Re-Imaging the Server Nodes

In SGI Management Center, there are compute images for each node type. [Table 2-2](#) shows the mapping. Re-provision the nodes with the compute images as needed.

**Table 2-2** Compute Images for SGI Hadoop Servers

Node Name	Image Name
sgi-nn	Compute-Hadoop-Namenode
sgi-rm	Compute-Hadoop-ResourceManager
sgi-app	Compute-Hadoop-App
r[rack#]n[node#]	Compute-Hadoop-Slave

To provision a node, do the following:

1. Select the appropriate node.
2. Right-click.
3. Select **Provision** > *compute-image-for-node*.

---

**Note:** Proceed to step 4 only if you are using Cloudera Manager 5.x.

---

4. After you have re-imaged the node and it is back online, copy file r[rack#]n[node#] (same as the node name) on the headnode to the node itself. The following are the pertinent directories:

Headnode directory:

```
/opt/sgi/Factory-Install/Integration/Hadoop/CM-UUIDs
```

Node directory: /var/lib/cloudera-scm-agent/uuid

Example:

```
scp /opt/sgi/Factory-Install/Integration/Hadoop/CM-UUIDs/r01n01
root@r01n01:/var/lib/cloudera-scm-agent/uuid/r01n01
```

