



SGI® Hadoop® Based on Intel® Xeon®  
Processor E5 Family

Getting Started Guide

007-5875-001

---

## COPYRIGHT

© 2013 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, 46600 Landing Parkway, Fremont, CA 94538.

---

## TRADEMARKS AND ATTRIBUTIONS

Silicon Graphics, SGI, the SGI logo, 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. Datameer is a trademark of Datameer, Inc. 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. Kitenga is a registered trademark of Kitenga Inc. Pentaho (TM) is a registered trademark of Pentaho Corporation. Quantum4D is a registered trademark of Quantum4D, Inc. Red Hat and all Red Hat-based trademarks are trademarks or registered trademarks of Red Hat, Inc. in the United States and other countries.

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

---

## Record of Revision

<b>Version</b>	<b>Description</b>
001	January 2013 Initial printing.



---

# Contents

<b>About This Guide</b> . . . . .	<b>vii</b>
Audience. . . . .	vii
Related Publications. . . . .	viii
Product Support . . . . .	ix
Reader Comments . . . . .	x
<b>1 Overview</b> . . . . .	<b>1</b>
The GigE Implementation . . . . .	2
Hardware—GigE . . . . .	3
Servers . . . . .	3
Network Hardware. . . . .	5
Configurations—GigE . . . . .	6
Half-Rack . . . . .	7
Full-Rack (42U) . . . . .	8
Multi-Rack (Second Rack And Beyond) . . . . .	9
Network Topology—GigE . . . . .	10
Node Level. . . . .	11
Rack Level for Single-Rack Configuration . . . . .	12
Rack Level for Rack 1 in Multi-Rack Configuration . . . . .	13
Rack Level for Rack 2 (And Beyond) in Multi-Rack Configuration . . . . .	14
Inter-Rack Level . . . . .	15
The 10GigE Implementation . . . . .	16
Hardware—10GigE. . . . .	16
Servers . . . . .	16
Network Hardware. . . . .	18
Configurations—10GigE . . . . .	19
Half-Rack . . . . .	20
Full-Rack (46U) . . . . .	21

- Multi-Rack (Second Rack and Beyond) . . . . . 22
- Network Topology—10GigE . . . . . 23
  - Node Level . . . . . 23
  - Rack Level for Single-Rack Configuration. . . . . 24
  - Inter-Rack Level . . . . . 25
- Software . . . . . 26
- 2 Cluster Startup . . . . . 27**
  - Accepting End-User License Agreements (EULAs). . . . . 27
  - Configuring and Starting SGI Management Center . . . . . 28
  - Starting the Cluster for the First Time . . . . . 28
  - Re-Imaging the Server Nodes . . . . . 29

---

## About This Guide

This guide provides an overview of the SGI® Hadoop® Reference Implementations based on the Intel® Xeon® processor E5 family along with getting-started instructions for these implementations. 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 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 Rackable C2005 Server Family User's Guide* (007-5717-xxx)
- *SGI InfiniteStorage Server 3000 (ISS3000) User's Guide* (007-5721-xxx)
- *SGI Rackable C1110-RP6 System User Guide* (007-5843-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 online books, man pages, and other information.

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 documents whose access require a support contract. See “[Product Support](#)” on page ix.
- You can also view man pages by typing `man <title>` on a command line.

---

**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.)

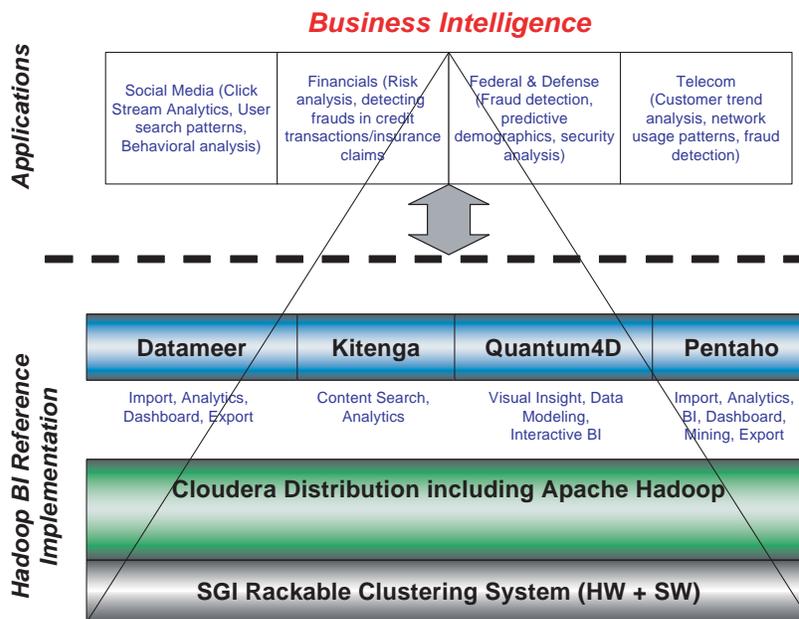
You can contact SGI in any of the following ways:

- 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.
- Send mail to the following address:

SGI  
Technical Publications  
46600 Landing Parkway  
Fremont, CA 94538

SGI values your comments and will respond to them promptly.

# 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
- Power optimization
- Capability of running business intelligence (BI) applications directly atop Hadoop (See Figure 1-1.)

There are two primary SGI Hadoop Reference Implementations based on the Intel® Xeon® Processor E5-2600 Series: GigE-based and 10GigE-based. This chapter describes these two implementations using the following topics:

- “The GigE Implementation” on page 2
- “The 10GigE Implementation” on page 16
- “Software” on page 26

## The GigE Implementation

This section describes the GigE implementation using the following topics:

- “Hardware—GigE” on page 3
- “Configurations—GigE” on page 6
- “Network Topology—GigE” on page 10

## Hardware—GigE

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

### Servers

The SGI Hadoop cluster employs the SGI Rackable™ C2005 and C1000 families of servers; a C2005 server and a C1110 server are shown in [Figure 1-2](#) and [Figure 1-3](#), respectively. This section describes the SGI servers that are used in the GigE-based SGI Hadoop cluster, their function in the Hadoop paradigm, and their specifications.



**Figure 1-2** An SGI Rackable C2005 Server



**Figure 1-3** An SGI Rackable C1110 Server

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

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

<b>SGI Server</b>	<b>Conventional Node Type</b>	<b>Hadoop Node Type</b>	<b>Specifications</b>
C2005-RP1 (half-depth) or C1110-RP6 (full-depth)	Master nodes	NameNode, Secondary NameNode, JobTracker	<ul style="list-style-type: none"> <li>– 2x Intel Xeon Processor E5-2630 (2.3 GHz , 6-core)</li> <li>– 8x 8GB 1.35v 1333MHz DIMMs (64GB memory)</li> <li>– 4x 2.5” 1TB 7200 rpm SATA 6Gb/s drives in RAID 10 configuration</li> <li>– 1x Dual-port 10GigE NIC</li> <li>– Redundant power supply</li> </ul>
C2005-RP1 (half-depth) or C1110-RP6 (full-depth)	Compute/Slave nodes	DataNodes, TaskTrackers	<ul style="list-style-type: none"> <li>– 2x Intel Xeon Processor E5-2630 (2.3 GHz, 6-core)</li> <li>– 8x 8GB 1.35v 1333MHz DIMMs (64GB memory)</li> <li>– 10x 2.5” 1TB 7200 rpm SATA 6Gb/s drives</li> </ul>
C2005-RP1 (half-depth) or C1110-RP6 (full-depth)		Application Node	<ul style="list-style-type: none"> <li>– 2x Intel Xeon Processor E5-2670 (2.6 GHz, 8-core)</li> <li>– 16x 8GB 1.35v 1333MHz DIMMs (128GB memory)</li> <li>– 4x 2.5” 1TB 7200 rpm SAS 6Gb/s drives in RAID 10 configuration</li> <li>– 1x Dual-port 10GigE NIC</li> <li>– Redundant power supply</li> </ul>

Table 1-2 describes the SGI Hadoop Reference Implementation with SGI GigE-based servers with the Intel Xeon Processor E5-2400 Series.

**Table 1-2** SGI Hadoop GigE-Based Servers–Intel Xeon Processor E5-2400 Series

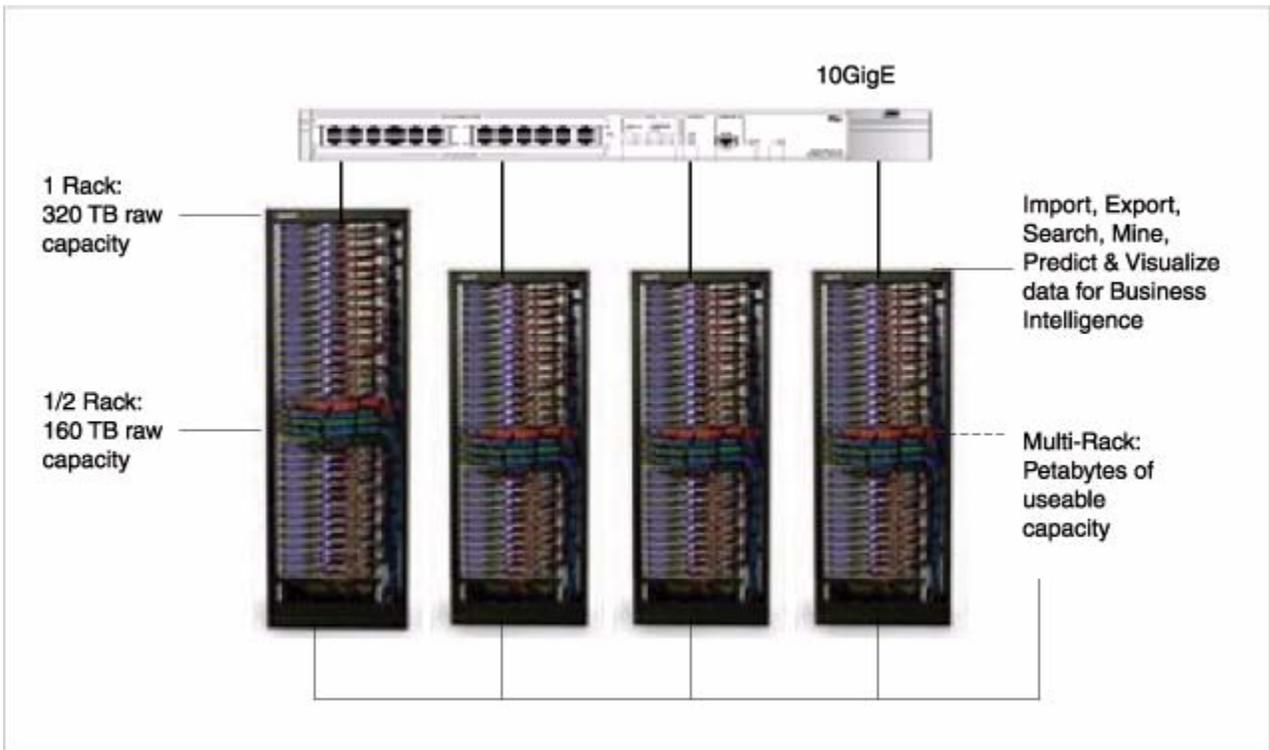
SGI Server	Conventional Node Type	Hadoop Node Type	Specifications
C2005-RN1 (half-depth) or C1110-RN3 (full-depth)	Master nodes	NameNode, Secondary NameNode, JobTracker	<ul style="list-style-type: none"> <li>– 2x Intel Xeon Processor E5-2420 (1.9 GHz, 6-core)</li> <li>– 6x 8GB 1.35v 1333MHz DIMMs (48GB memory)</li> <li>– 4x 2.5” 1TB 7200 rpm SATA 6Gb/s drives in RAID 10 configuration</li> <li>– 1x Dual-port 10GigE NIC</li> <li>– Redundant power supply</li> </ul>
C2005-RN1 (half-depth) or C1110-RN3 (full-depth)	Compute/Slave nodes	DataNodes, TaskTrackers	<ul style="list-style-type: none"> <li>– 2x Intel Xeon Processor E5-2420 (1.9 GHz, 6-core)</li> <li>– 6x 8GB 1.35v 1333MHz DIMMs (48GB memory)</li> <li>– 10x 2.5” 1TB 7200 rpm SATA 6Gb/s drives</li> </ul>
C2005-RP1 (half-depth) or C1110-RP6 (full-depth)		Application Node	<ul style="list-style-type: none"> <li>– 2x Intel Xeon Processor E5-2670 (2.6 GHz, 8-core)</li> <li>– 16x 8GB 1.35v 1333MHz DIMMs (128GB memory)</li> <li>– 4x 2.5” 1TB 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 two components:

- 2 LG-Ericsson ES-4550G 48-port GigE switches per rack
- 1 LG-Ericsson ES-5048XG 10-GigE spine switch

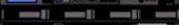
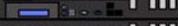
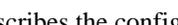
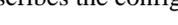
## Configurations—GigE



**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 half-rack, full-rack, and multi-rack configurations.

### Half-Rack

Side A		RackU	Side B	
Description	Image		Image	Description
48port GigE Hadoop Data Switch		42		48port GigE Hadoop Data Switch
48port GigE SGI MC Management Switch		41		48port 10GigE Hadoop Data Spine Switch
	2U Blank	40	2U Blank	
		39		
	3U Blank	38	3U Blank	
		37		
		36		
	3U Blank	35	3U Blank	
		34		
		33		
	3U Blank	32	3U Blank	
		31		
		30		
	3U Blank	29	3U Blank	
		28		
		27		
	3U Blank	26	3U Blank	
		25		
		24		
	3U Blank	23	3U Blank	
		22		
		21		
Secondary Name Node/ SGI-MC Headnode		20		Application Node
Name Node		19		Jobtracker
Data/TaskTracker Node		18		
Data/TaskTracker Node		17		
Data/TaskTracker Node		16		
Data/TaskTracker Node		15		
Data/TaskTracker Node		14		
Data/TaskTracker Node		13		
Data/TaskTracker Node		12		
Data/TaskTracker Node		11		
Data/TaskTracker Node		10		
Data/TaskTracker Node		9		
Data/TaskTracker Node		8		
Data/TaskTracker Node		7		
Data/TaskTracker Node		6		
Data/TaskTracker Node		5		
Data/TaskTracker Node		4		
Data/TaskTracker Node		3		
Data/TaskTracker Node		2		
Data/TaskTracker Node		1		

**Figure 1-5** Half-Rack Configuration

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

- 1 SGI Management Center node/Secondary NameNode
- 1 NameNode
- 1 JobTracker
- 1 Application node
- 16 DataNodes/TaskTracker nodes
- 2 48-port GigE stacked Hadoop data network switches
- 1 SGI Management Center network switch

Full-Rack (42U)

Side A		RackU	Side B	
Description	Image		Image	Description
48port GigE Hadoop Data Switch		42		48port GigE Hadoop Data Switch
48port GigE SGI MC Management Switch		41		48port 10GigE Hadoop Data Spine Switch
	2U Blank	40		
		39	2U Blank	
		38		
	2U Blank	37	2U Blank	
		36		
Secondary Name Node/ SGI-MC Headnode		35		Application Node
		34		
Name Node		33		Jobtracker
		32		
Data/TaskTracker Node		31		Data/TaskTracker Node
		30		
Data/TaskTracker Node		29		Data/TaskTracker Node
		28		
Data/TaskTracker Node		27		Data/TaskTracker Node
		26		
Data/TaskTracker Node		25		Data/TaskTracker Node
		24		
Data/TaskTracker Node		23		Data/TaskTracker Node
		22		
Data/TaskTracker Node		21		Data/TaskTracker Node
		20		
Data/TaskTracker Node		19		Data/TaskTracker Node
		18		
Data/TaskTracker Node		17		Data/TaskTracker Node
		16		
Data/TaskTracker Node		15		Data/TaskTracker Node
		14		
Data/TaskTracker Node		13		Data/TaskTracker Node
		12		
Data/TaskTracker Node		11		Data/TaskTracker Node
		10		
Data/TaskTracker Node		9		Data/TaskTracker Node
		8		
Data/TaskTracker Node		7		Data/TaskTracker Node
		6		
Data/TaskTracker Node		5		Data/TaskTracker Node
		4		
Data/TaskTracker Node		3		Data/TaskTracker Node
		2		
Data/TaskTracker Node		1		Data/TaskTracker Node

Figure 1-6 Full-Rack Configuration

Figure 1-6 describes the configuration of a full-rack configuration. The rack consists of the following:

- 1 SGI Management Center node/Secondary NameNode
- 1 NameNode
- 1 JobTracker
- 1 Application node
- 32 DataNodes/TaskTracker nodes
- 2 48-port GigE stacked Hadoop data network switches
- 1 SGI Management Center network switch

Multi-Rack (Second Rack And Beyond)

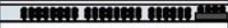
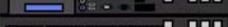
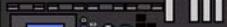
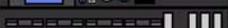
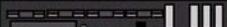
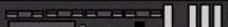
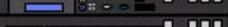
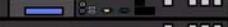
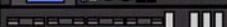
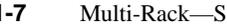
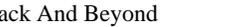
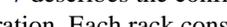
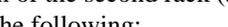
Side A		RackU	Side B	
Description	Image		Image	Description
48port GigE Hadoop Data Switch		36		48port GigE Hadoop Data Switch
48port GigE SGI MC Management Switch		35		
	2U Blank	34	3U Blank	
		33		
Data/TaskTracker Node		31		Data/TaskTracker Node
Data/TaskTracker Node		30		Data/TaskTracker Node
Data/TaskTracker Node		29		Data/TaskTracker Node
Data/TaskTracker Node		28		Data/TaskTracker Node
Data/TaskTracker Node		27		Data/TaskTracker Node
Data/TaskTracker Node		26		Data/TaskTracker Node
Data/TaskTracker Node		25		Data/TaskTracker Node
Data/TaskTracker Node		24		Data/TaskTracker Node
Data/TaskTracker Node		23		Data/TaskTracker Node
Data/TaskTracker Node		22		Data/TaskTracker Node
Data/TaskTracker Node		21		Data/TaskTracker Node
Data/TaskTracker Node		20		Data/TaskTracker Node
Data/TaskTracker Node		19		Data/TaskTracker Node
Data/TaskTracker Node		18		Data/TaskTracker Node
Data/TaskTracker Node		17		Data/TaskTracker Node
Data/TaskTracker Node		16		Data/TaskTracker Node
Data/TaskTracker Node		15		Data/TaskTracker Node
Data/TaskTracker Node		14		Data/TaskTracker Node
Data/TaskTracker Node		13		Data/TaskTracker Node
Data/TaskTracker Node		12		Data/TaskTracker Node
Data/TaskTracker Node		11		Data/TaskTracker Node
Data/TaskTracker Node		10		Data/TaskTracker Node
Data/TaskTracker Node		9		Data/TaskTracker Node
Data/TaskTracker Node		8		Data/TaskTracker Node
Data/TaskTracker Node		7		Data/TaskTracker Node
Data/TaskTracker Node		6		Data/TaskTracker Node
Data/TaskTracker Node		5		Data/TaskTracker Node
Data/TaskTracker Node		4		Data/TaskTracker Node
Data/TaskTracker Node		3		Data/TaskTracker Node
Data/TaskTracker Node		2		Data/TaskTracker Node
Data/TaskTracker Node		1		Data/TaskTracker Node

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:

- 32 DataNodes/TaskTracker nodes
- 2 48-port GigE stacked Hadoop data network switches
- 1 SGI Management Center network switch

## Network Topology—GigE

As described in [Table 1-3](#), the network topology of the SGI Hadoop cluster depends on its rack configuration.

**Table 1-3** Network Topology

Rack Configuration	Network Topology
Single-rack	The Master node servers are attached to the top-of-rack switches directly via 10-Gigabit Ethernet.
Multi-rack	A 10-Gigabit Ethernet aggregate spine switch is introduced into the networking topology. The Master node servers are attached to this spine switch directly.

This section illustrates the network topology from the most granular level (node level) to the top level (inter-rack level):

- [“Node Level” on page 11](#)
- [“Rack Level for Single-Rack Configuration” on page 12](#)
- [“Rack Level for Rack 1 in Multi-Rack Configuration” on page 13](#)
- [“Rack Level for Rack 2 \(And Beyond\) in Multi-Rack Configuration” on page 14](#)
- [“Inter-Rack Level” on page 15](#)

Node Level

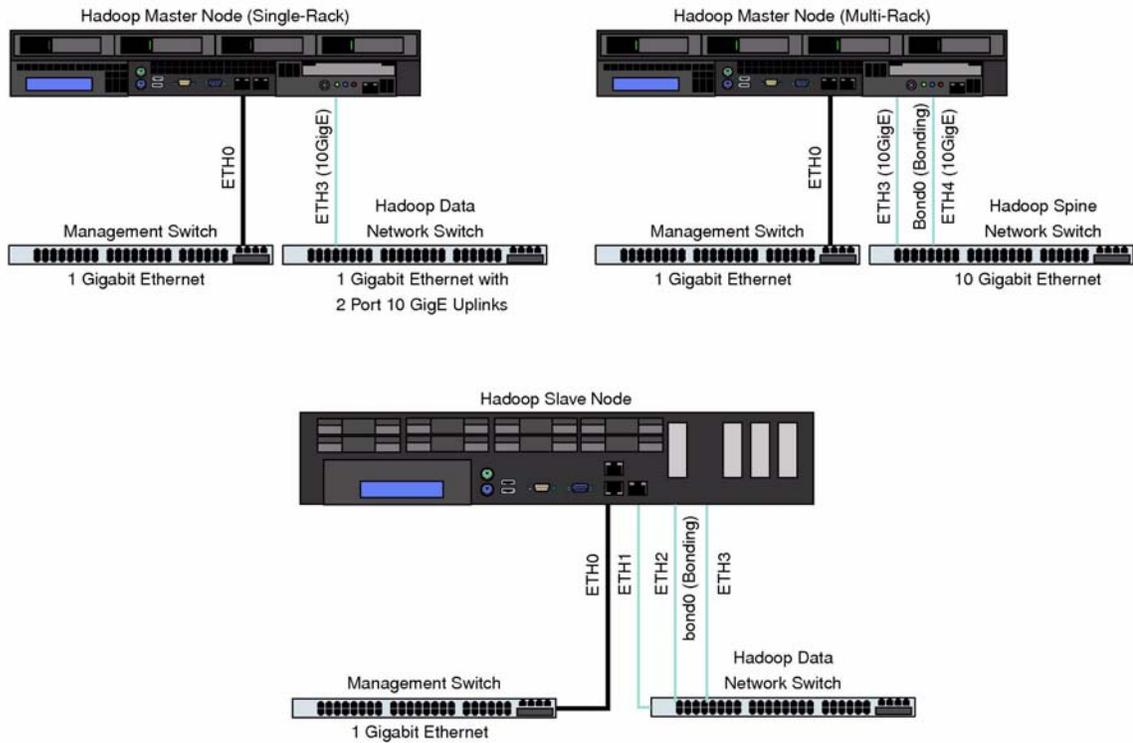


Figure 1-8 Network Topology—Node Level

Rack Level for Single-Rack Configuration

Hadoop Rack Networking Layout (Single Rack)

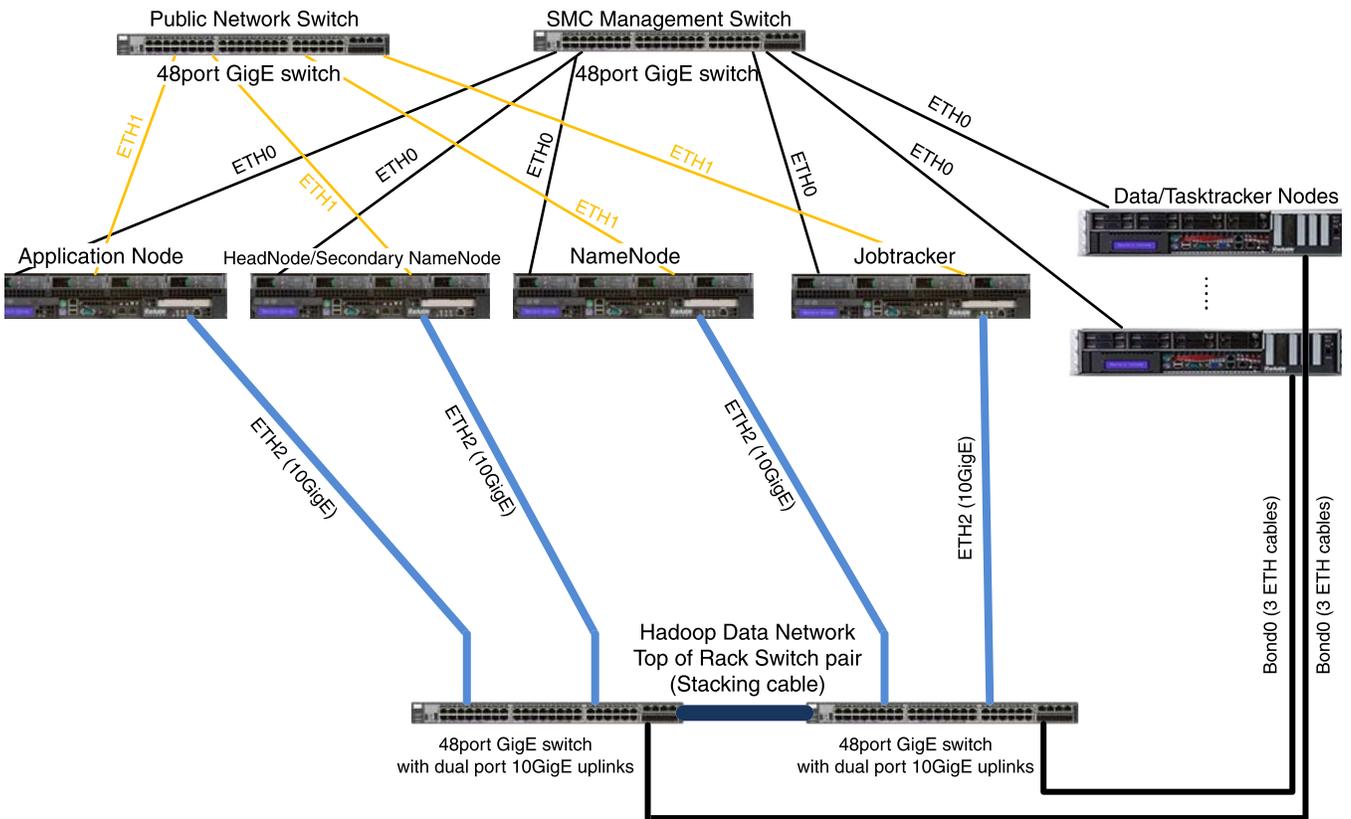


Figure 1-9 Network Topology—Rack Level for Single Rack

Rack Level for Rack 1 in Multi-Rack Configuration

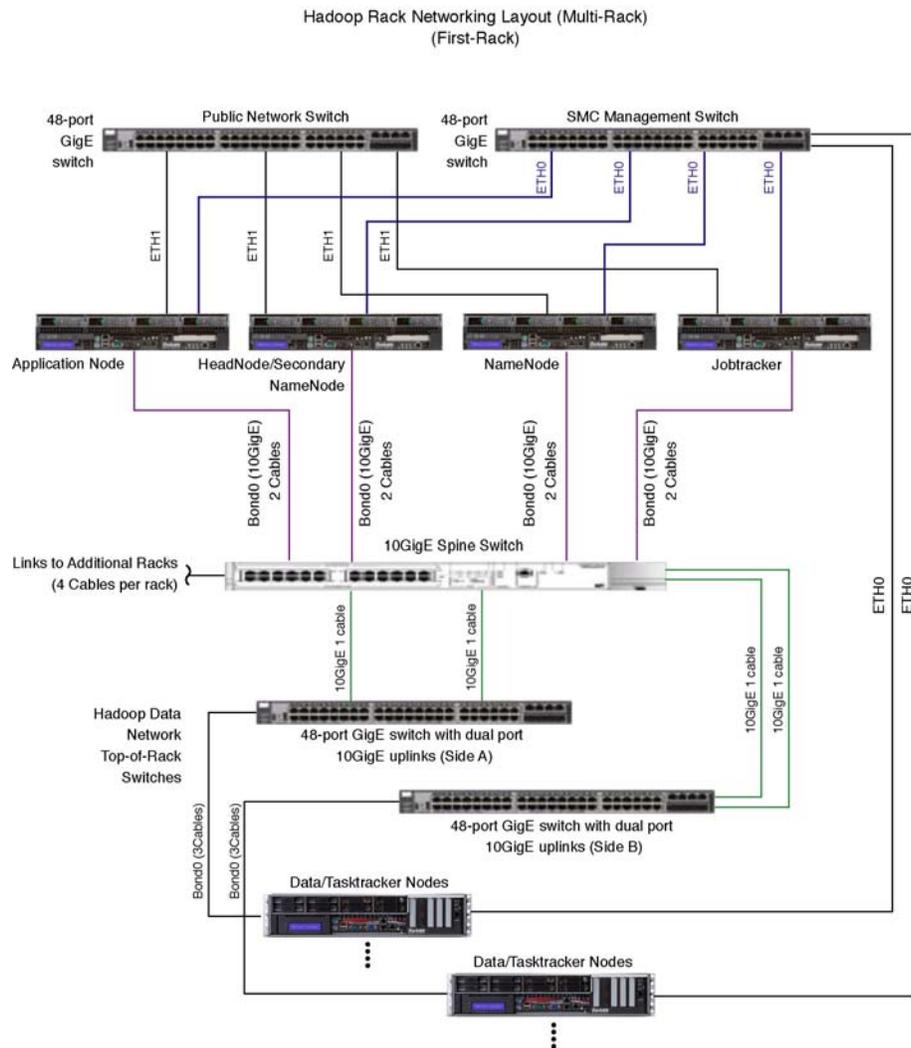
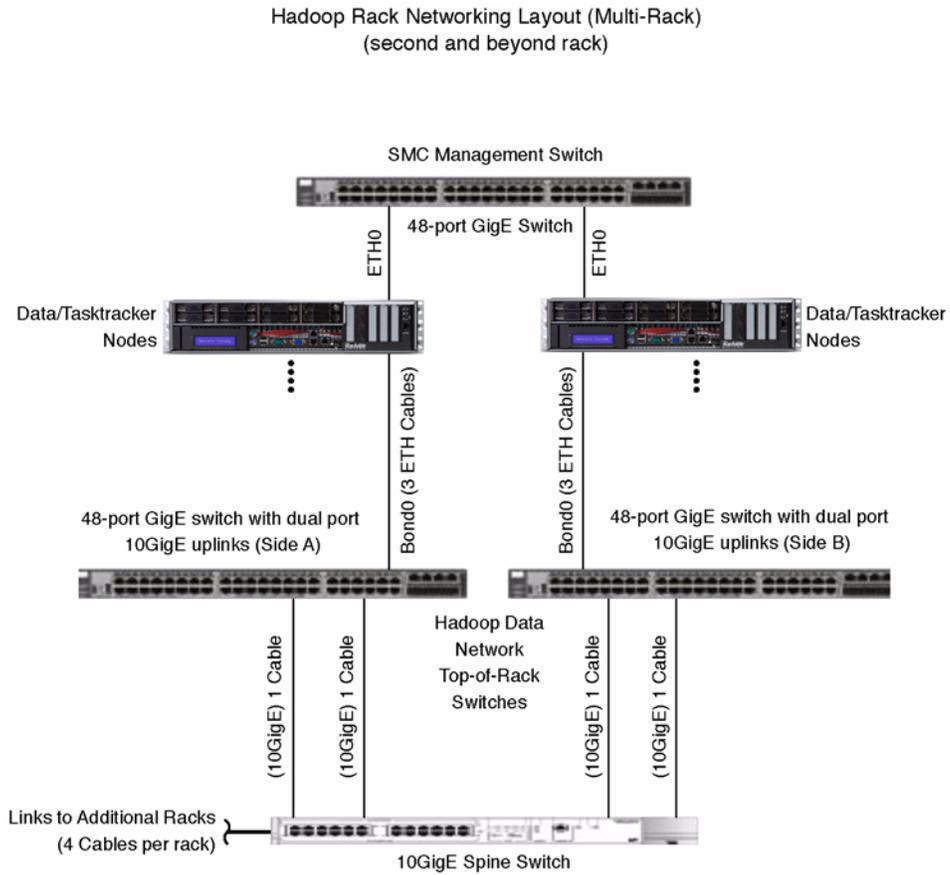


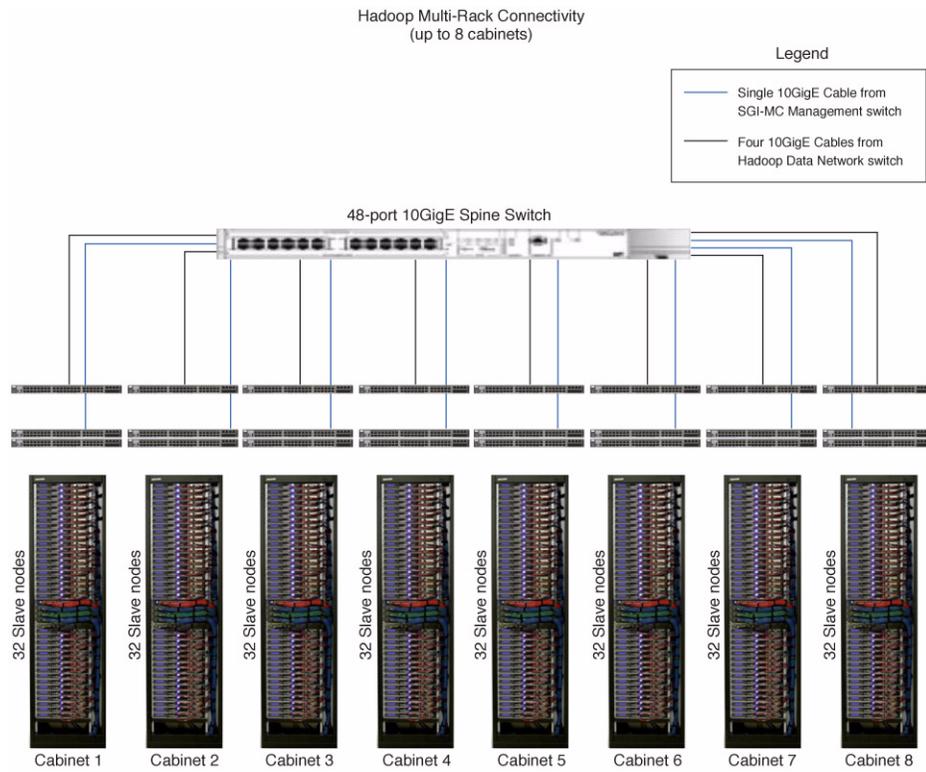
Figure 1-10 Network Topology—Rack Level for Rack 1 of Multi-Rack

### Rack Level for Rack 2 (And Beyond) in Multi-Rack Configuration



**Figure 1-11** Network Topology—Rack Level for Rack 2 (And Beyond) of Multi-Rack

Inter-Rack Level



**Figure 1-12** Network Topology—Inter-Rack Level

## The 10GigE Implementation

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

- “Hardware—10GigE” on page 16
- “Configurations—10GigE” on page 19
- “Network Topology—10GigE” on page 23

### Hardware—10GigE

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™ C2005 and ISS3012 servers; a C2005 server and an ISS3012 server are shown in [Figure 1-13](#) and [Figure 1-14](#), 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 1-13** An SGI Rackable C2005 Server



**Figure 1-14** An SGI ISS3012 Server

Table 1-4 describes the SGI Hadoop Reference Implementations with SGI 10GigE-based servers with the Intel Xeon Processor E5-2600 Series.

**Table 1-4** SGI Hadoop 10GigE-Based Half-Depth Servers–Intel Xeon Processor E5-2600 Series

SGI Server	Conventional Node Type	Hadoop Node Type	Specifications
C2005-RP1	Master nodes	NameNode, Secondary NameNode, JobTracker	<ul style="list-style-type: none"> <li>– 2x Intel Xeon Processor E5-2630 (2.3 GHz , 6-core)</li> <li>– 8x 8GB 1.35v 1333MHz DIMMs (64GB memory)</li> <li>– 4x 2.5” 1TB 7200 rpm SATA 6Gb/s drives in RAID 10 configuration</li> <li>– 1x Dual-port 10GigE NIC</li> <li>– Redundant power supply</li> </ul>
ISS3012-RP6	Compute/Slave nodes	DataNodes, TaskTrackers	<ul style="list-style-type: none"> <li>– 2x Intel Xeon Processor E5-2630 (2.3 GHz, 6-core)</li> <li>– 8x 8GB 1.35v 1333MHz DIMMs (64GB memory)</li> <li>– 12x 3.5” 3TB 7200 rpm SATA drives</li> <li>– 1x Dual-port 10GigE NIC</li> </ul>
C2005-RP1		Application Node	<ul style="list-style-type: none"> <li>– 2x Intel Xeon Processor E5-2670 (2.6 GHz, 8-core)</li> <li>– 16x 8GB 1.35v 1333MHz DIMMs (128GB memory)</li> <li>– 4x 2.5” 1TB 7200 rpm SAS 6Gb/s drives in RAID 10 configuration</li> <li>– 1x Dual-port 10GigE NIC</li> <li>– Redundant power supply</li> </ul>

Table 1-5 describes the SGI Hadoop Reference Implementation with SGI 10GigE-based servers with the Intel Xeon Processor E5-2400 Series.

**Table 1-5** SGI Hadoop 10GigE-Based Half-Depth Servers–Intel Xeon Processor E5-2400 Series

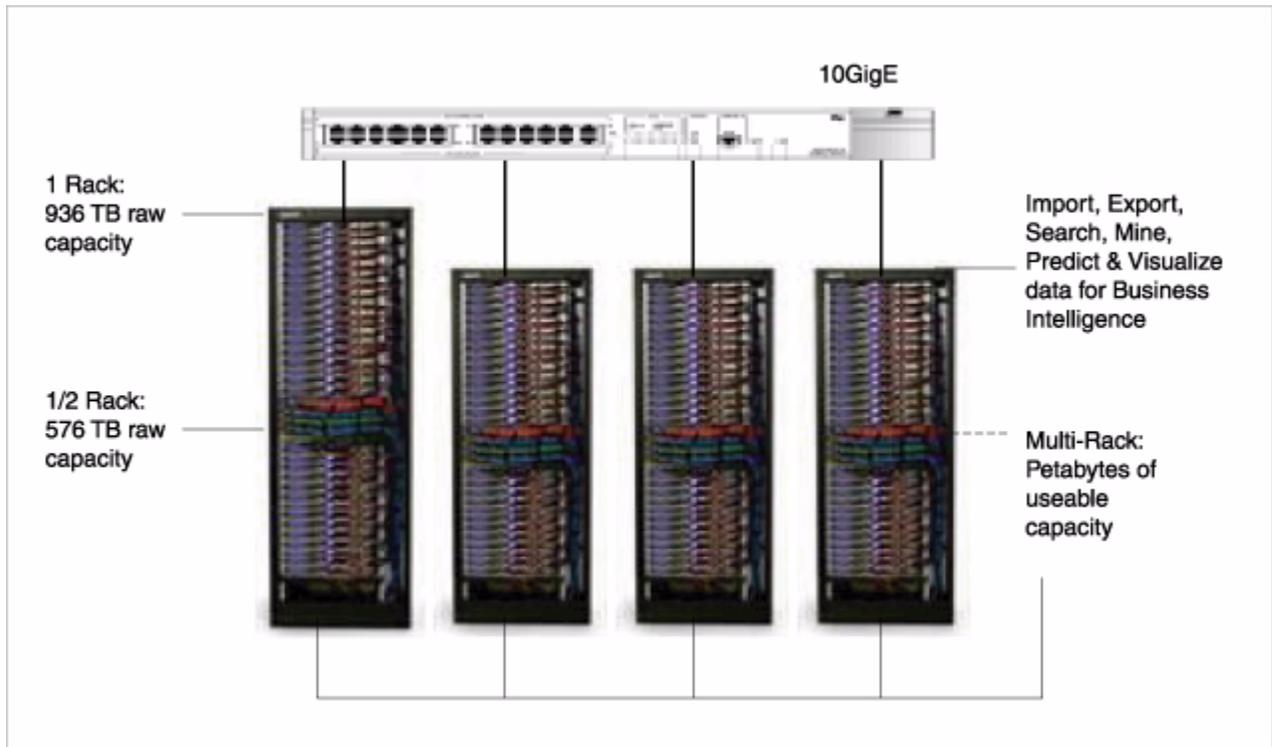
SGI Server	Conventional Node Type	Hadoop Node Type	Specifications
C2005-RN1	Master nodes	NameNode, Secondary NameNode, JobTracker	<ul style="list-style-type: none"> <li>– 2x Intel Xeon Processor E5-2420 (1.9 GHz, 6-core)</li> <li>– 6x 8GB 1.35v 1333MHz DIMMs (48GB memory)</li> <li>– 4x 2.5” 1TB 7200 rpm SATA 6Gb/s drives in RAID 10 configuration</li> <li>– 1x Dual-port 10GigE NIC</li> <li>– Redundant power supply</li> </ul>
ISS3012-RN1	Compute/Slave nodes	DataNodes, TaskTrackers	<ul style="list-style-type: none"> <li>– 2x Intel Xeon Processor E5-2420 (1.9 GHz, 6-core)</li> <li>– 6x 8GB 1.35v 1333MHz DIMMs (48GB memory)</li> <li>– 12x 3.5” 3TB 7200 rpm SATA drives</li> <li>– 1x Dual-port 10GigE NIC</li> </ul>
C2005-RP1		Application Node	<ul style="list-style-type: none"> <li>– 2x Intel Xeon Processor E5-2670 (2.6 GHz, 8-core)</li> <li>– 16x 8GB 1.35v 1333MHz DIMMs (128GB memory)</li> <li>– 4x 2.5” 1TB 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 two components:

- 2 LG-Ericsson ES-4550G 48-port GigE switches per rack
- 1 LG-Ericsson ES-5048XG 10-GigE spine switch

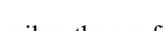
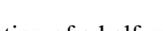
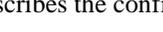
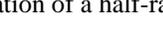
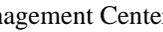
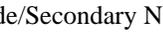
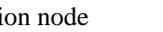
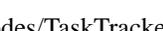
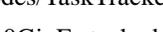
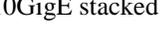
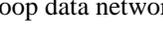
## Configurations—10GigE



**Figure 1-15** Data Capacity for Various Rack Configurations

The SGI Hadoop Cluster is available in single-rack and multi-rack configurations. [Figure 1-15](#) shows the range of data capacity for the configurations. This section describes the half-rack, full-rack, and multi-rack configurations.

## Half-Rack

Side A		RackU	Side B	
Description	Image		Image	Description
48port GigE SGI MC Management Switch	***** ***** ****	46		
48port 10GigE Hadoop Data Switch	***** *****	45	***** *****	48port 10GigE Hadoop Data Switch
	1U Blank	44		1U Blank
		43		
	3U Blank	42		3U Blank
		41		
		40		
	3U Blank	39		3U Blank
		38		
		37		
	3U Blank	36		3U Blank
		35		
		34		
	3U Blank	33		3U Blank
		32		
		31		
	3U Blank	30		3U Blank
		29		
		28		
Secondary/NameNode/SGIMC Head Node		27		Application Node
Name Node		26		Jobtracker
		25		
		24		
Data/TaskTracker Node		23		Data/TaskTracker Node
		22		
Data/TaskTracker Node		21		Data/TaskTracker Node
		20		
Data/TaskTracker Node		19		Data/TaskTracker Node
		18		
Data/TaskTracker Node		17		Data/TaskTracker Node
		16		
Data/TaskTracker Node		15		Data/TaskTracker Node
		14		
Data/TaskTracker Node		13		Data/TaskTracker Node
		12		
Data/TaskTracker Node		11		Data/TaskTracker Node
		10		
Data/TaskTracker Node		9		Data/TaskTracker Node
		8		
Data/TaskTracker Node		7		Data/TaskTracker Node
		6		
Data/TaskTracker Node		5		Data/TaskTracker Node
		4		
Data/TaskTracker Node		3		Data/TaskTracker Node
		2		
Data/TaskTracker Node		1		Data/TaskTracker Node

**Figure 1-16** Half-Rack Configuration

Figure 1-16 describes the configuration of a half-rack configuration. The rack consists of the following:

- 1 SGI Management Center node/Secondary NameNode
- 1 NameNode
- 1 JobTracker
- 1 Application node
- 16 DataNodes/TaskTracker nodes
- 2 48-port 10GigE stacked Hadoop data network switches
- 1 SGI Management Center network switch

Full-Rack (46U)

Side A	RackU	Side B
Description	Image	Description
48port GigE SGI MC Management Switch	46	
48port 10GigE Hadoop Data Switch	45	48port 10GigE Hadoop Data Switch
	44	
	43	
Secondary Name Node	42	Application Node
	41	
Name Node	40	Jobtracker
	39	
	38	
Data/TaskTracker Node	37	Data/TaskTracker Node
	36	
Data/TaskTracker Node	35	Data/TaskTracker Node
	34	
	33	
Data/TaskTracker Node	32	Data/TaskTracker Node
	31	
	30	
Data/TaskTracker Node	29	Data/TaskTracker Node
	28	
	27	
Data/TaskTracker Node	26	Data/TaskTracker Node
	25	
	24	
Data/TaskTracker Node	23	Data/TaskTracker Node
	22	
	21	
Data/TaskTracker Node	20	Data/TaskTracker Node
	19	
	18	
Data/TaskTracker Node	17	Data/TaskTracker Node
	16	
	15	
Data/TaskTracker Node	14	Data/TaskTracker Node
	13	
	12	
Data/TaskTracker Node	11	Data/TaskTracker Node
	10	
	9	
Data/TaskTracker Node	8	Data/TaskTracker Node
	7	
	6	
Data/TaskTracker Node	5	Data/TaskTracker Node
	4	
	3	
	2	
Data/TaskTracker Node	1	Data/TaskTracker Node

Figure 1-17 Full-Rack Configuration

Figure 1-17 describes the configuration of a full-rack configuration. The rack consists of the following:

- 1 SGI Management Center node/Secondary NameNode
- 1 NameNode
- 1 JobTracker
- 1 Application node
- 26 DataNodes/TaskTracker nodes
- 2 48-port 10GigE stacked Hadoop data network switches
- 1 SGI Management Center network switch

**Multi-Rack (Second Rack and Beyond)**

Side A		RackU	Side B	
Description	Image		Image	Description
48port GigE SGI MC Management Switch	***** ***** ****	46		
48port 10GigE Hadoop Data Switch	***** *****	45	***** *****	48port 10GigE Hadoop Data Switch
	2U Blank	44		
		43		
		42		
	3U Blank	41		
		40		
		39		
		38		
Data/TaskTracker Node		37		Data/TaskTracker Node
		36		
Data/TaskTracker Node		35		Data/TaskTracker Node
		34		
		33		
Data/TaskTracker Node		32		Data/TaskTracker Node
		31		
		30		
Data/TaskTracker Node		29		Data/TaskTracker Node
		28		
		27		
Data/TaskTracker Node		26		Data/TaskTracker Node
		25		
		24		
Data/TaskTracker Node		23		Data/TaskTracker Node
		22		
		21		
Data/TaskTracker Node		20		Data/TaskTracker Node
		19		
		18		
Data/TaskTracker Node		17		Data/TaskTracker Node
		16		
		15		
Data/TaskTracker Node		14		Data/TaskTracker Node
		13		
		12		
Data/TaskTracker Node		11		Data/TaskTracker Node
		10		
		9		
Data/TaskTracker Node		8		Data/TaskTracker Node
		7		
		6		
Data/TaskTracker Node		5		Data/TaskTracker Node
		4		
		3		
		2		
Data/TaskTracker Node		1		Data/TaskTracker Node

**Figure 1-18** Multi-Rack—Second Rack and Beyond

Figure 1-18 describes the configuration of the second rack (and subsequent racks) of a multi-rack configuration. Each rack consists of the following:

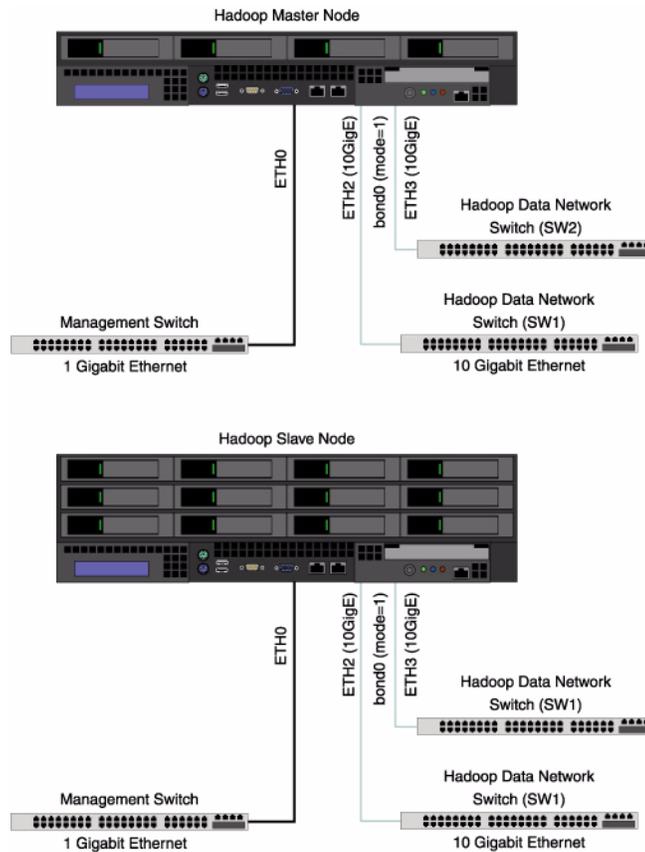
- 26 DataNodes/TaskTracker nodes
- 2 48-port GigE stacked Hadoop data network switches
- 1 SGI Management Center network switch

## Network Topology—10GigE

This section illustrates the network topology from the most granular level (node level) to the top level (inter-rack level):

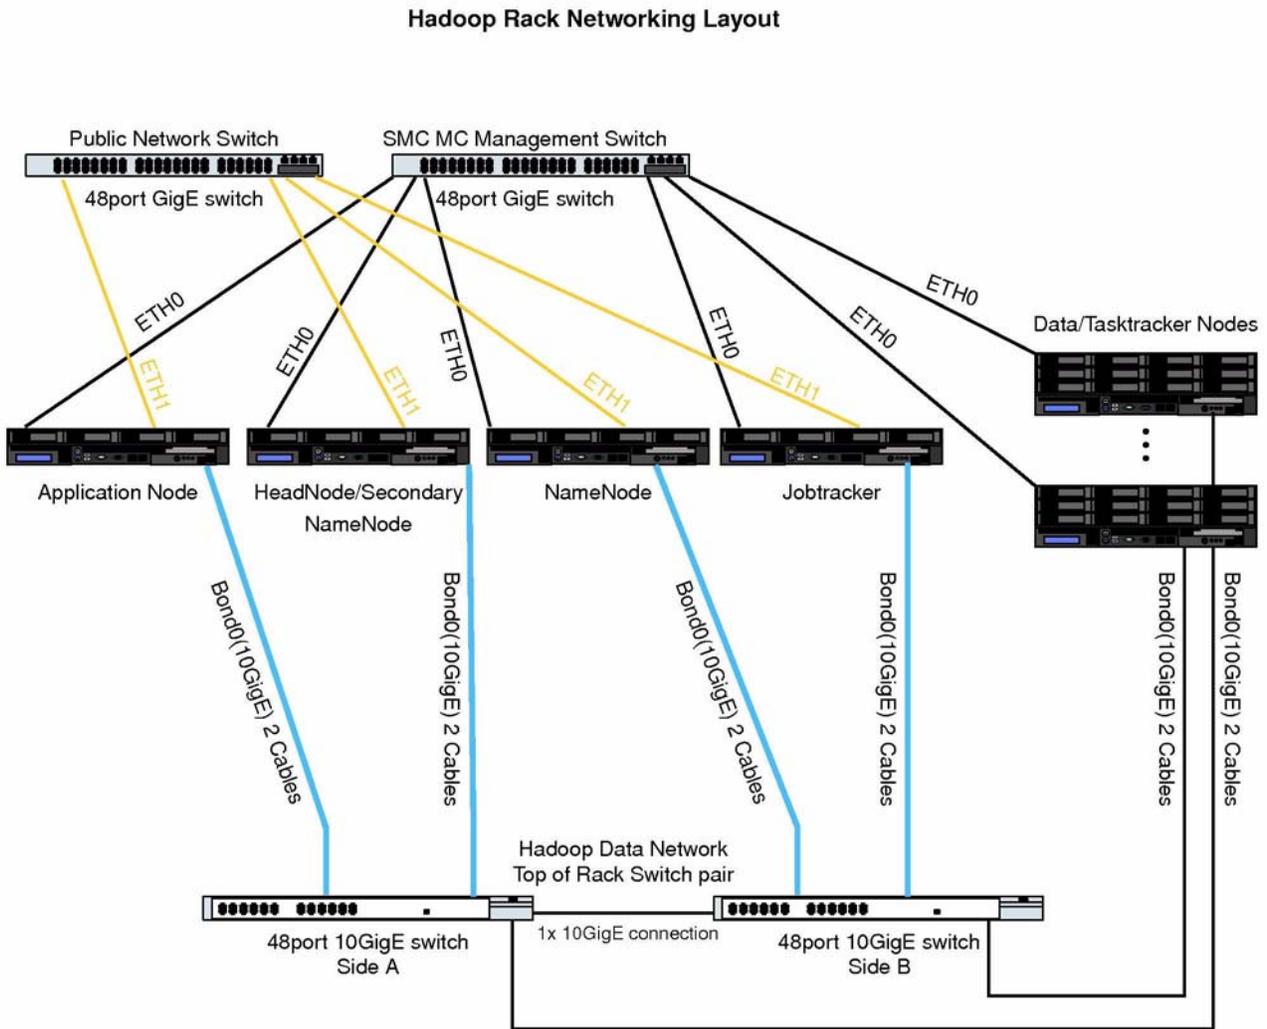
- “Node Level” on page 23
- “Rack Level for Single-Rack Configuration” on page 24
- “Inter-Rack Level” on page 25

### Node Level



**Figure 1-19** Network Topology—Node Level

### Rack Level for Single-Rack Configuration



**Figure 1-20** Network Topology—Rack Level for Single Rack

Inter-Rack Level

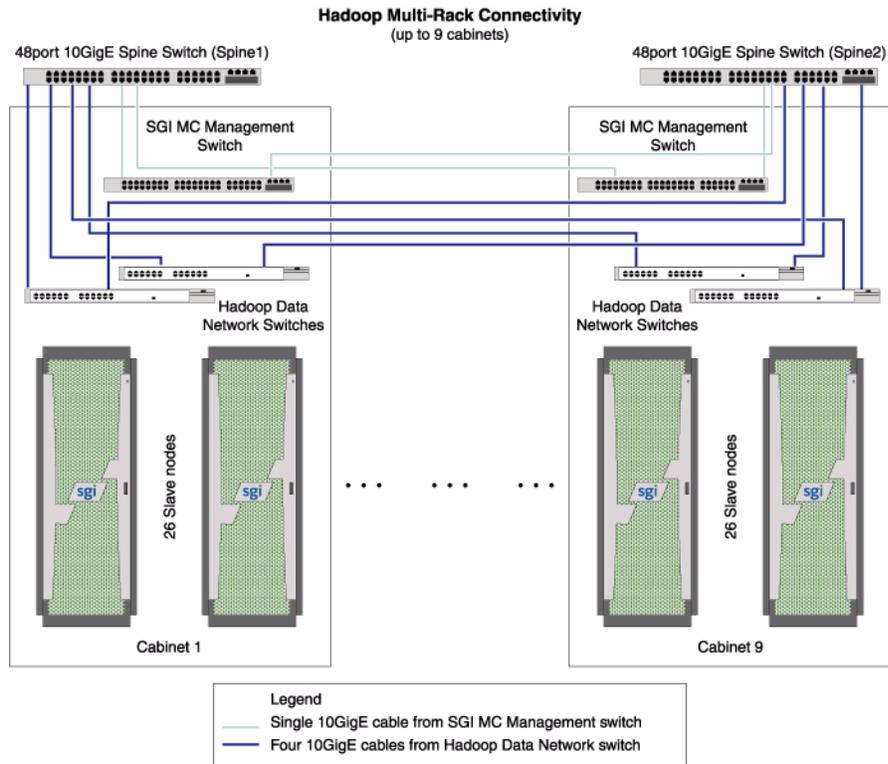


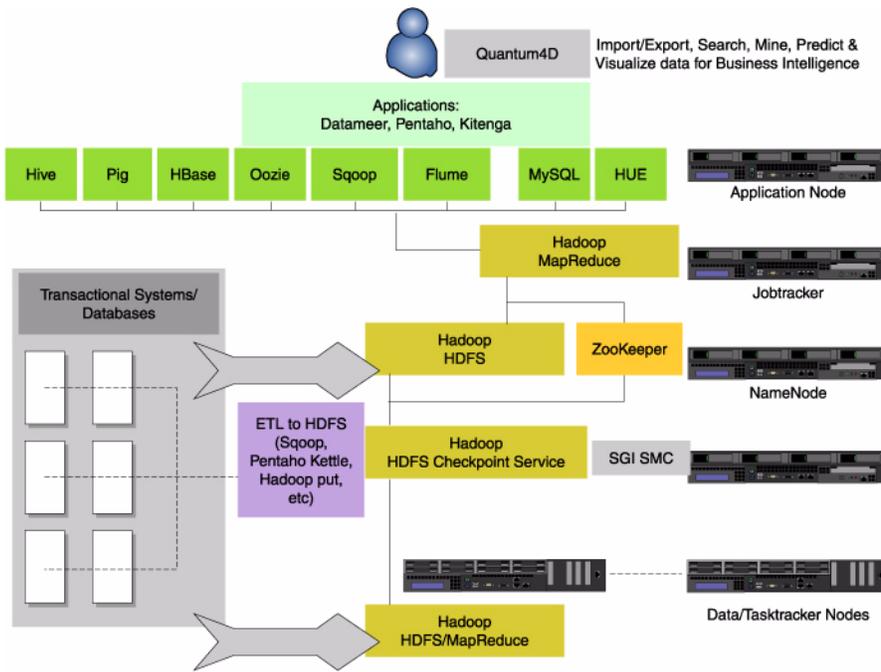
Figure 1-21 Network Topology—Inter-Rack Level

## Software

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

- Red Hat® Enterprise Linux (RHEL) 6 .2 (2.6.32-220.el6.x86\_64)
- Cloudera™ distribution Apache Hadoop 3 update 4 (Hadoop 0.20.2-cdh3u4)
- SGI Management Center 1.6.0
- An ecosystem of business intelligence applications software from ISVs like Kitenga®, Datameer™, Pentaho® and Quantum 4D® (See Figure 1-1 on page 1 and Figure 1-22.)

Figure 1-22 shows the SGI Hadoop software stack.



**Figure 1-22** 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 27
- “Configuring and Starting SGI Management Center” on page 28
- “Starting the Cluster for the First Time” on page 28
- “Re-Imaging the Server Nodes” on page 29

### 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

To configure and start the SGI Management Center to monitor the SGI Hadoop cluster, you will need to follow the instructions in the *SGI Management Center Quick Start Guide* and appropriately configure the Hadoop servers described in [Table 2-1](#).

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

Daemon	Hostname	Hadoop Data Network Hostname
NameNode	sgi-nn	sgi-nn-10ge
Secondary NameNode	sgi-snn	sgi-snn-10ge
JobTracker	sgi-jt	sgi-jt-10ge
Application Node	sgi-app	sgi-app-10ge
DataNodes & TaskTrackers	r[rack#]n[node#]	r[rack#]n[node#]-ge

## 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-jt`
    - iv. Compute/Slave nodes in the Compute group

Hadoop is configured to start once the servers have booted.

3. Use the web browser on the head node to log into the Hadoop web interfaces:
  - NameNode: `http://sgi-nn-10ge:50070`
  - Secondary NameNode: `http://sgi-snn-10ge:50090`
  - JobTracker: `http://sgi-jt-10ge:50030`
4. Verify that the cluster powered on correctly and that all slave nodes joined the Hadoop cluster.

Run the following command to verify the number of DataNodes match the expected slave node count:

```
# sudo -u hdfs hapoop dfsadmin -report
```

## 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-jt	Compute-Hadoop-Jobtracker
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*.

