SGI® InfiniteData™ Cluster for Hadoop®
Getting Started Guide

007-6313-001

Record of Revision

Version	Description
001	January 2014 Initial printing.

Contents

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
- 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

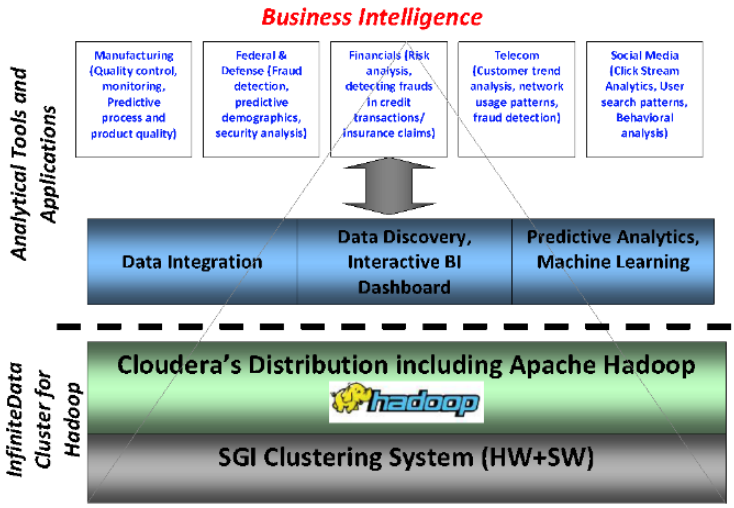

Figure 1-1 SGI Hadoop Business Intelligence Ecosystem

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

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

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, JobTracker	– 2x Intel Xeon Processor E5-2630 v2 (2.6 GHz , 6-core) – 8x 8GB 1.5v 1866MHz DIMMs (64GB memory) – 4x 3.5” 4TB 7200 rpm SATA 6Gb/s drives in RAID 10 configuration – 1x Dual-port 10GigE NIC – Redundant power supply
IDC3212-RP4	Compute/Slave nodes	DataNodes, TaskTrackers	– 2x Intel Xeon Processor E5-2630 v2 (2.6 GHz, 6-core) – 8x 8GB 1.5v 1866MHz DIMMs (64GB memory) – 12x 3.5” 4TB 7200 rpm SATA drives – 1x Dual-port 10GigE NIC
C1110-RP6		Application Node	– 2x Intel Xeon Processor E5-2680 v2 (2.8 GHz, 10-core) – 16x 8GB 1.5v 1866MHz DIMMs (128GB memory) – 4x 3.5” 4TB 7200 rpm SAS 6Gb/s drives in RAID 10 configuration – 1x Dual-port 10GigE NIC – Redundant power supply

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

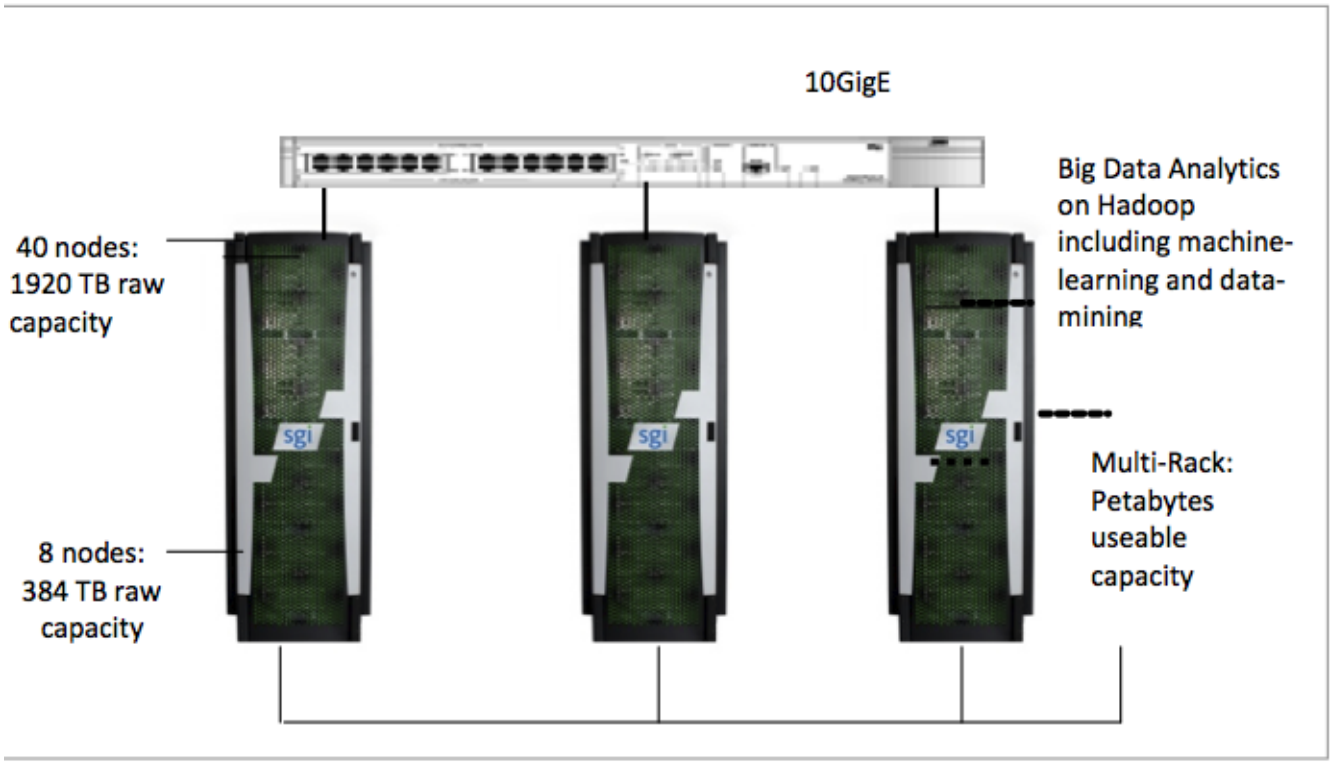

Figure 1-4 Data Capacity for Various Rack Configurations

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

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

Description	Image	U	Description
	1U Blank	45	
48port 10GigE Hadoop Data Switch (SPINE)		44	48port 10GigE Hadoop Data Switch (SPINE)
48port GigE SGIMC Management Switch		43	
Standby NameNode/SGIMC Head Node		42	
NameNode		41	
Application Node		40	
Jobtracker		39	
48port 10GigE Hadoop Data Switch		38	
48port 10GigE Hadoop Data Switch		37	
		36	
Data/TaskTracker Node (r01n35)		35	Data/TaskTracker Node (r01n36)
		34	
Data/TaskTracker Node (r01n33)		33	Data/TaskTracker Node (r01n34)
		32	
Data/TaskTracker Node (r01n31)		31	Data/TaskTracker Node (r01n32)
		30	
Data/TaskTracker Node (r01n29)		29	Data/TaskTracker Node (r01n30)
		28	
Data/TaskTracker Node (r01n27)		27	Data/TaskTracker Node (r01n28)
		26	
Data/TaskTracker Node (r01n25)		25	Data/TaskTracker Node (r01n26)
		24	
Data/TaskTracker Node (r01n23)		23	Data/TaskTracker Node (r01n24)
		22	
Data/TaskTracker Node (r01n21)		21	Data/TaskTracker Node ((r01n22)
		20	
Data/TaskTracker Node (r01n19)		19	Data/TaskTracker Node (r01n20)
		18	
Data/TaskTracker Node (r01n17)		17	Data/TaskTracker Node (r01n18)
		16	
Data/TaskTracker Node (r01n15)		15	Data/TaskTracker Node (r01n16)
		14	
Data/TaskTracker Node (r01n13)		13	Data/TaskTracker Node (r01n14)
		12	
Data/TaskTracker Node (r01n11)		11	Data/TaskTracker Node (r01n12)
		10	
Data/TaskTracker Node (r01n09)		9	Data/TaskTracker Node (r01n10)
		8	
Data/TaskTracker Node (r01n07)		7	Data/TaskTracker Node (r01n08)
		6	
Data/TaskTracker Node (r01n05)		5	Data/TaskTracker Node (r01n06)
		4	
Data/TaskTracker Node (r01n03)		3	Data/TaskTracker Node (r01n04)
		2	
Data/TaskTracker Node (r01n01)		1	Data/TaskTracker Node (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 JobTracker
- 1 Application node
- 36 DataNodes/TaskTracker nodes
- 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 GigE SGIMC Management Switch		44	
48port 10GigE Hadoop Data Switch		43	
48port 10GigE Hadoop Data Switch		42	
	1U Blank	41	
		40	
Data/TaskTracker Node (r02n39)		39	Data/TaskTracker Node (r02n40)
		38	
Data/TaskTracker Node (r02n37)		37	Data/TaskTracker Node (r02n38)
		36	
Data/TaskTracker Node (r02n35)		35	Data/TaskTracker Node (r02n36)
		34	
Data/TaskTracker Node (r02n33)		33	Data/TaskTracker Node (r02n34)
		32	
Data/TaskTracker Node (r02n31)		31	Data/TaskTracker Node (r02n32)
		30	
Data/TaskTracker Node (r02n29)		29	Data/TaskTracker Node (r02n30)
		28	
Data/TaskTracker Node (r02n27)		27	Data/TaskTracker Node (r02n28)
		26	
Data/TaskTracker Node (r02n25)		25	Data/TaskTracker Node (r02n26)
		24	
Data/TaskTracker Node (r02n23)		23	Data/TaskTracker Node (r02n24)
		22	
Data/TaskTracker Node (r02n21)		21	Data/TaskTracker Node ((r02n22)
		20	
Data/TaskTracker Node (r02n19)		19	Data/TaskTracker Node (r02n20)
		18	
Data/TaskTracker Node (r02n17)		17	Data/TaskTracker Node (r02n18)
		16	
Data/TaskTracker Node (r02n15)		15	Data/TaskTracker Node (r02n16)
		14	
Data/TaskTracker Node (r02n13)		13	Data/TaskTracker Node (r02n14)
		12	
Data/TaskTracker Node (r02n11)		11	Data/TaskTracker Node (r02n12)
		10	
Data/TaskTracker Node (r02n09)		9	Data/TaskTracker Node (r02n10)
		8	
Data/TaskTracker Node (r02n07)		7	Data/TaskTracker Node (r02n08)
		6	
Data/TaskTracker Node (r02n05)		5	Data/TaskTracker Node (r02n06)
		4	
Data/TaskTracker Node (r02n03)		3	Data/TaskTracker Node (r02n04)
		2	
Data/TaskTracker Node (r02n01)		1	Data/TaskTracker Node (r02n02)

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/TaskTracker nodes
- 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

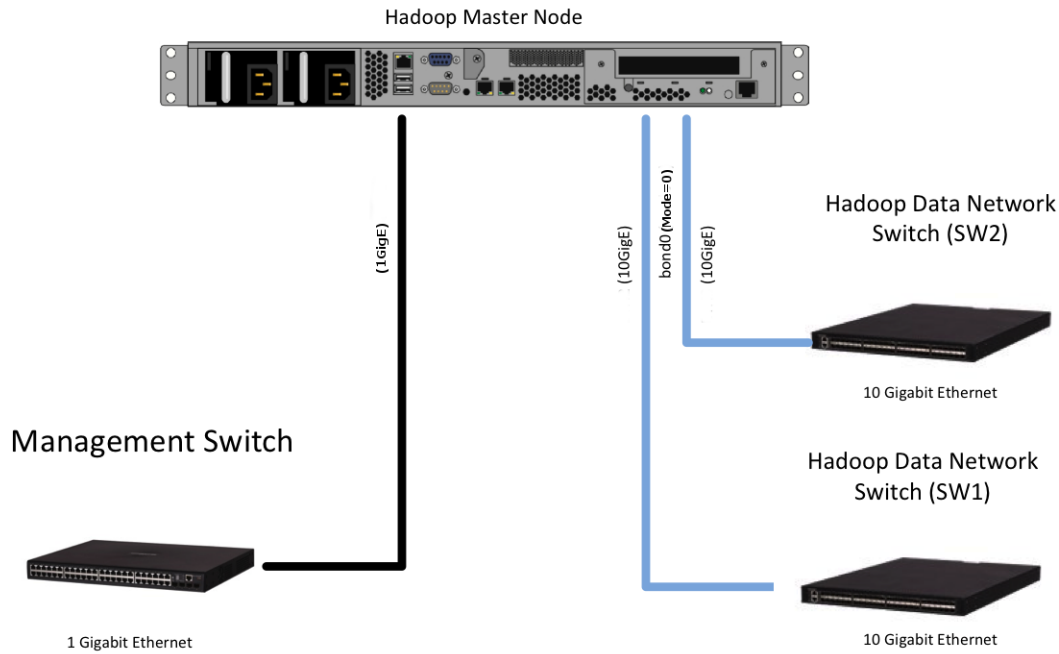

Figure 1-8 Network Topology—Master Node

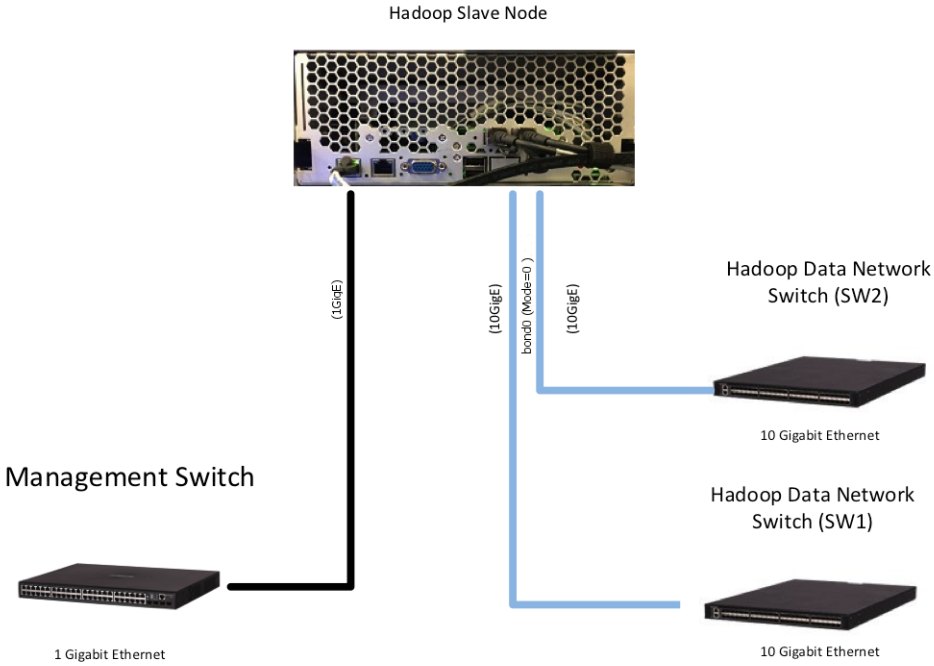

Figure 1-9 Network Topology—Slave Node

Rack Level for Single-Rack Configuration

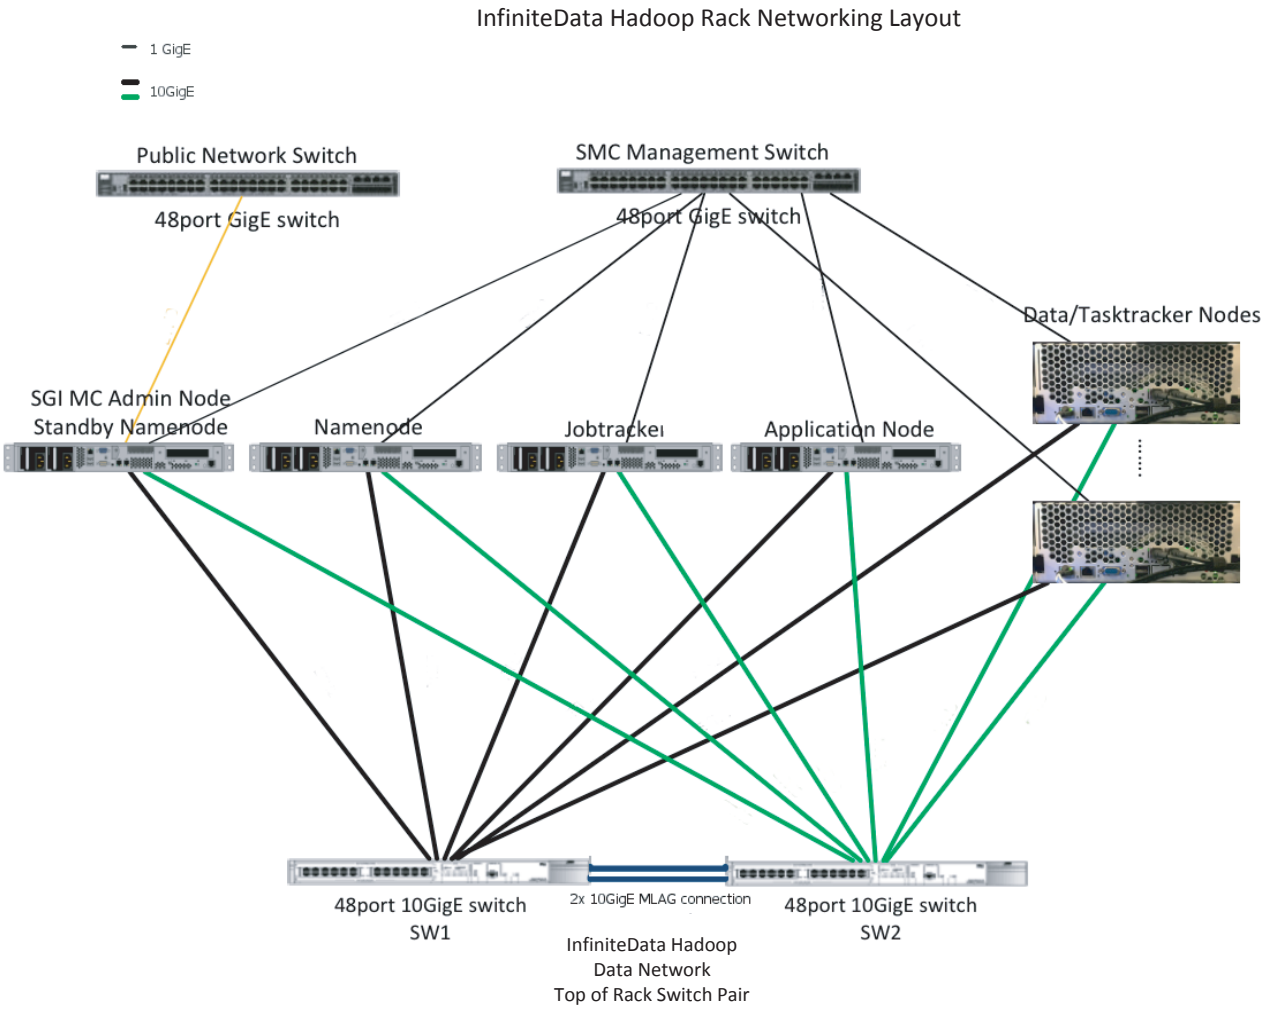

Figure 1-10 Network Topology—Rack Level for Single Rack

Multi-Rack Data Network

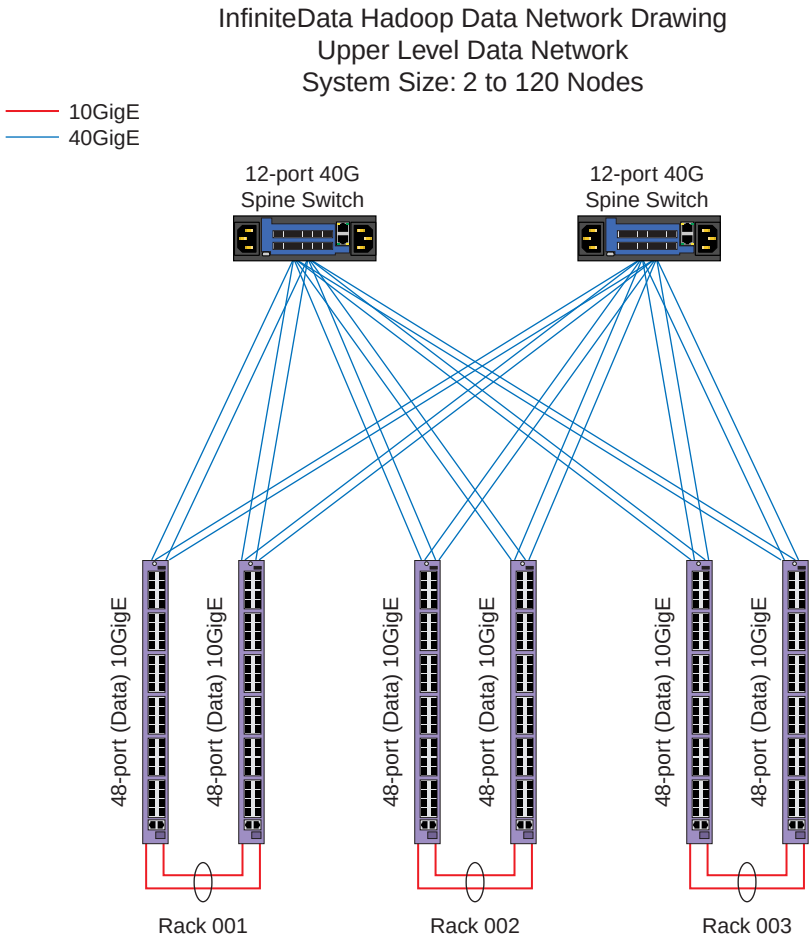

Figure 1-11 Network Topology—Multi-Rack Data Network (2 to 120 Nodes)

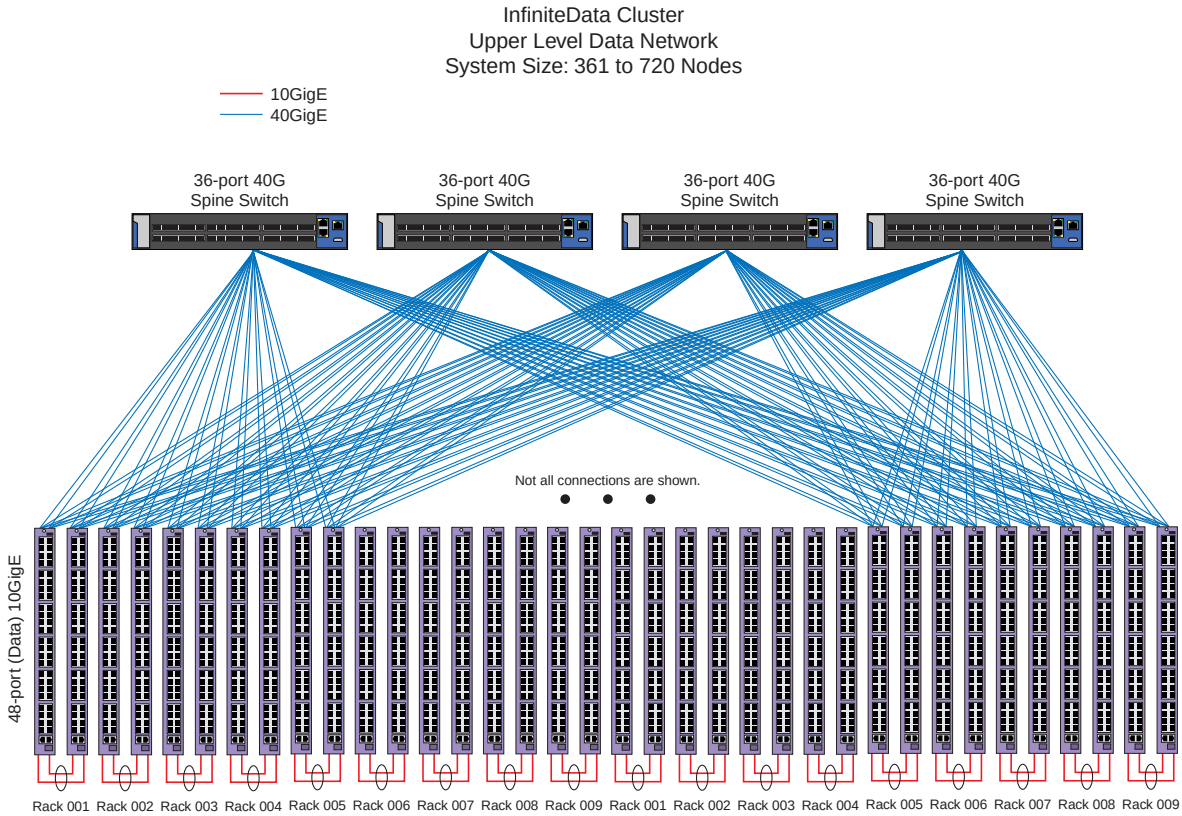

Figure 1-12 Network Topology—Multi-Rack Data Network (361 to 720 Nodes)

Management Network

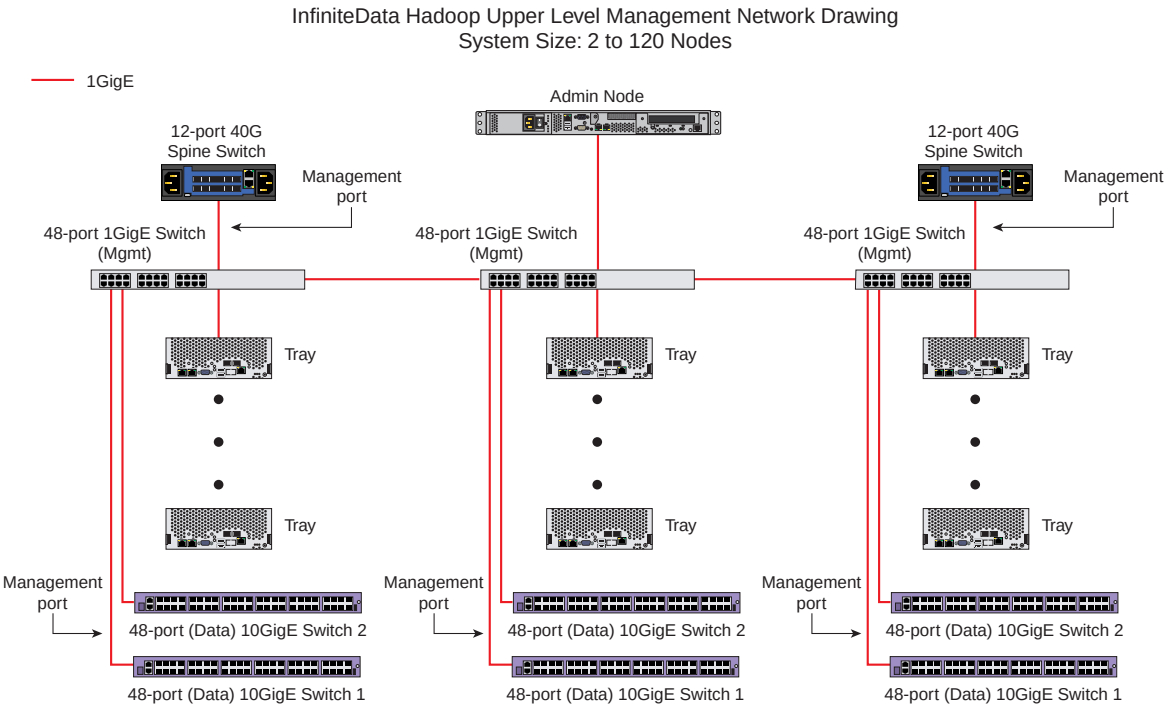

Figure 1-13 Network Topology—Management Network (2 to 120 Nodes)

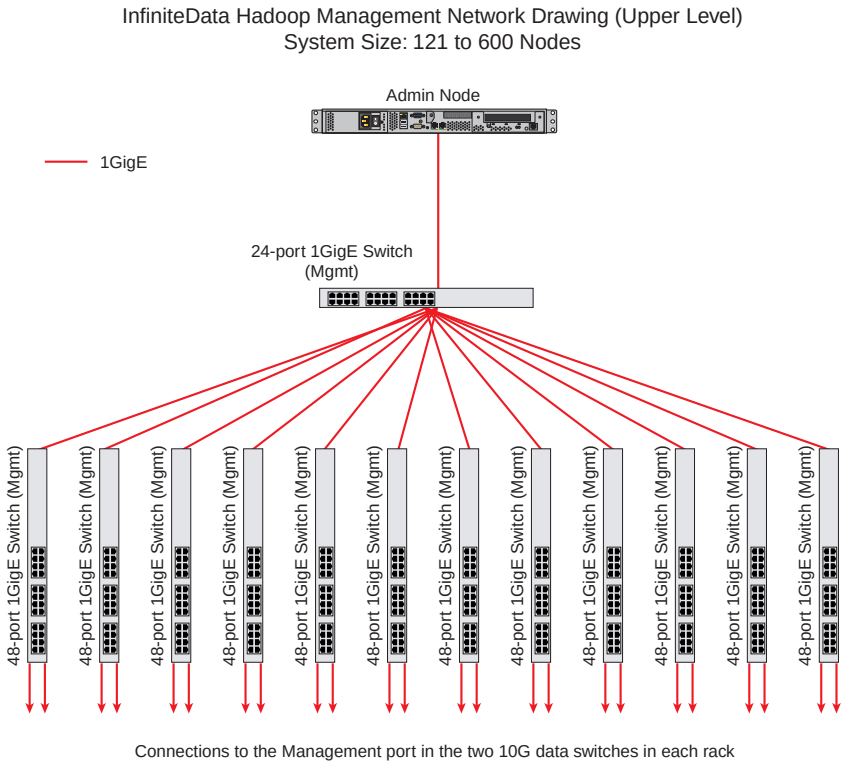

Figure 1-14 Network Topology—Management Network (121 to 600 Nodes)

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 4.x
- Cloudera Manager 4.x
- SGI Management Center 1.7

Figure 1-15 shows the SGI Hadoop software stack.

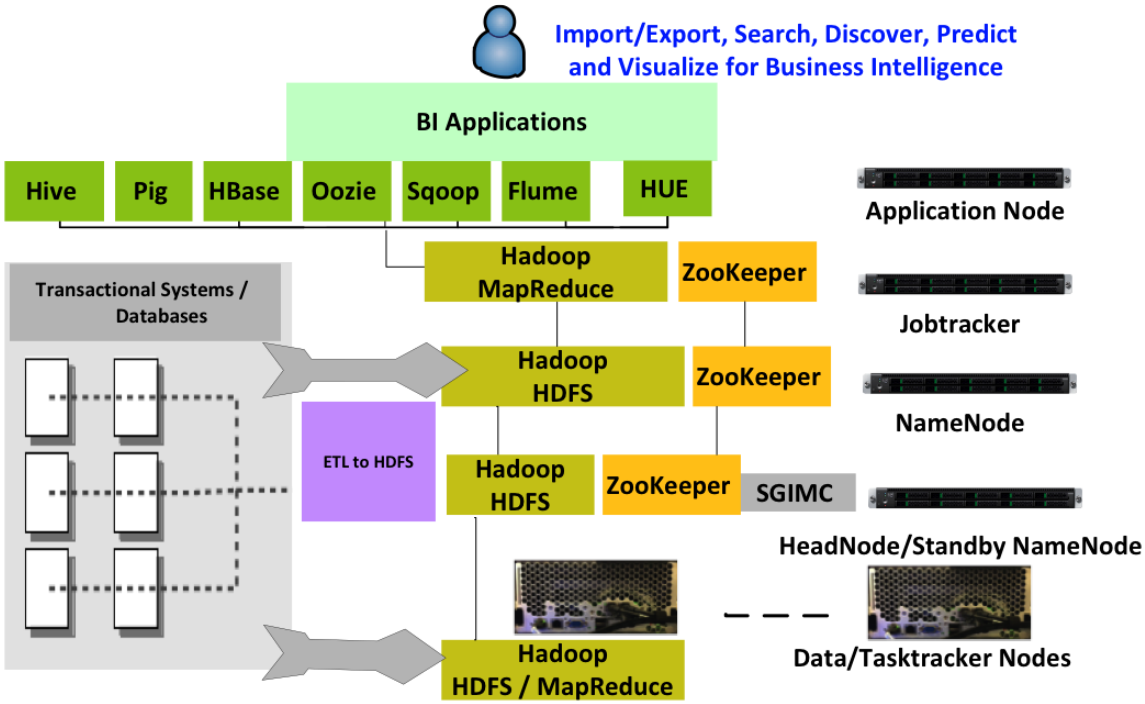

Figure 1-15 SGI Hadoop Software Stack

Cluster Startup

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

Accepting End-User License Agreements (EULAs)

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`
JobTracker	`sgi-jt`	`sgi-jt-data`
Application Node	`sgi-app`	`sgi-app-data`
DataNodes & TaskTrackers	`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-jt`
      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`.

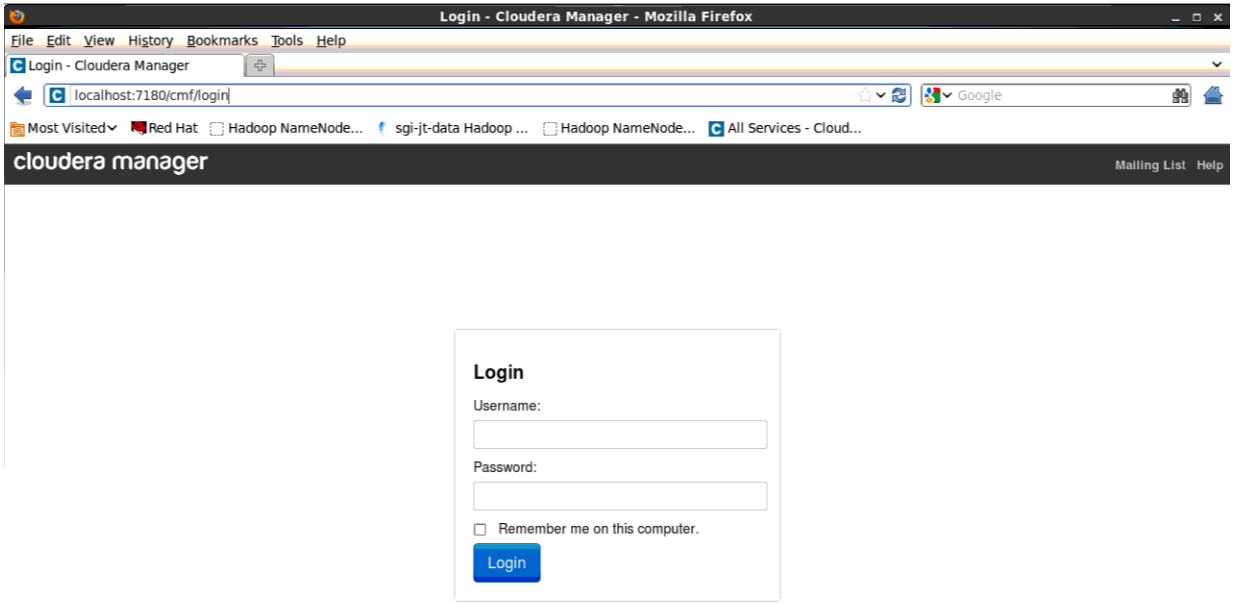

Figure 2-1 Cloudera Manager Login Screen

Starting Hadoop Cluster Services

After a successful login, the **All Services** screen, as shown in Figure 2-2, should appear. Start the Hadoop cluster services by clicking **Start** in the **Actions** list on the right side of the screen.

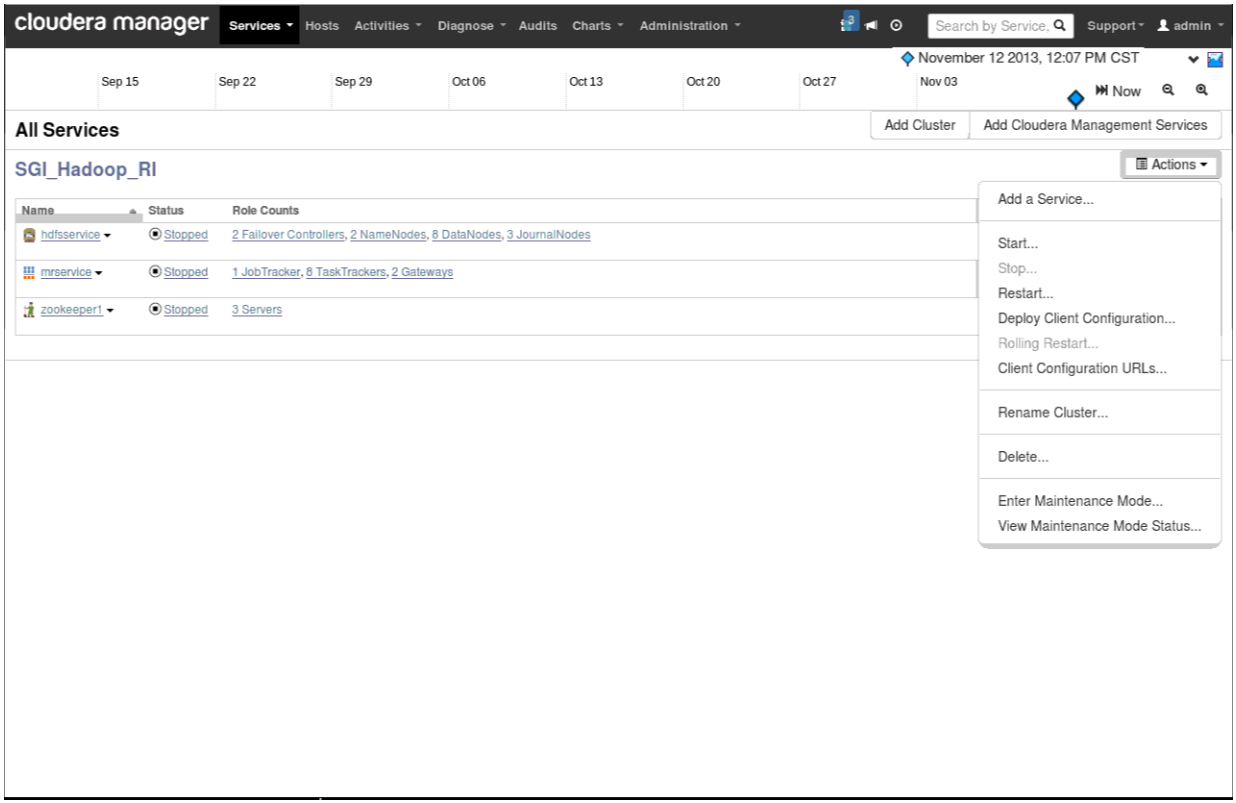

Figure 2-2 All Services Screen

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**.

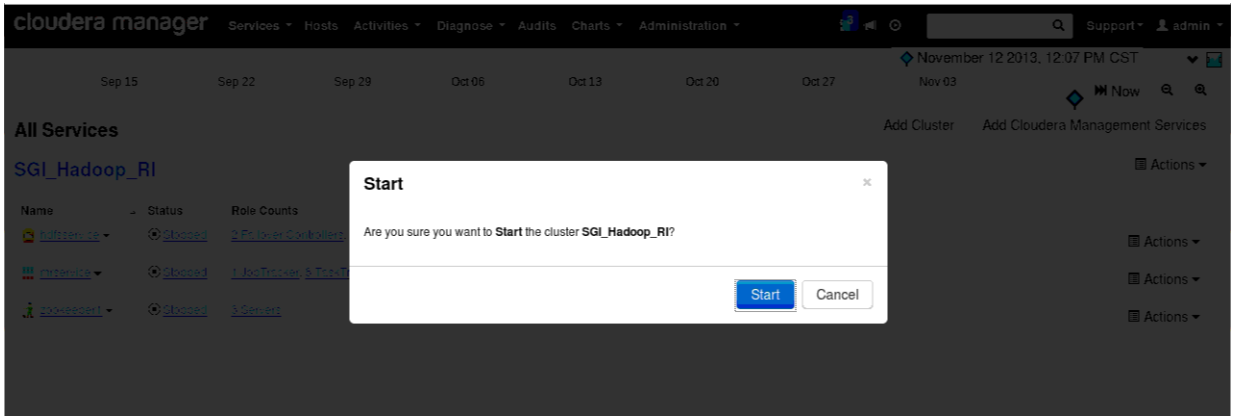

Figure 2-3 Starting the Hadoop Cluster

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.

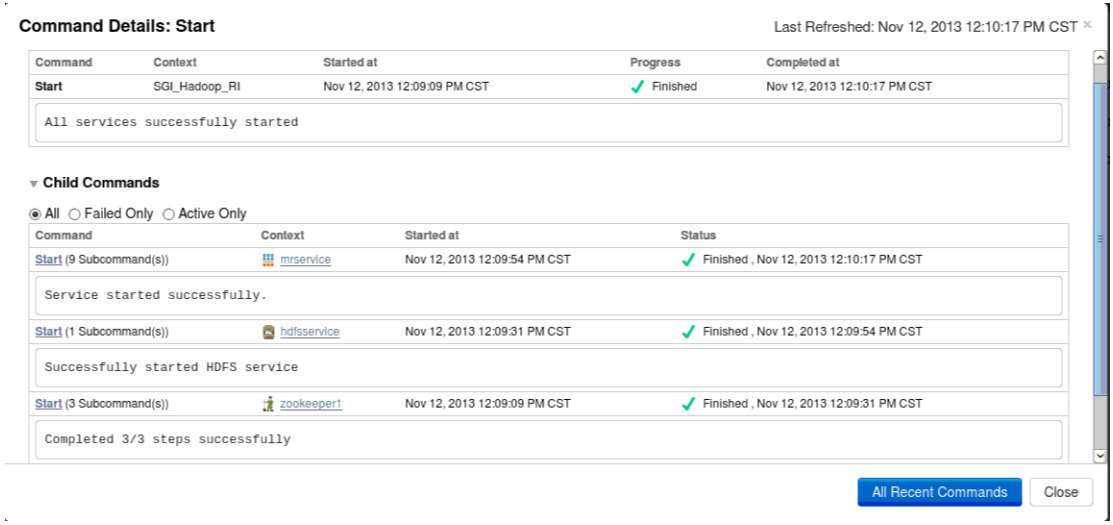

Figure 2-4 Successful Startup Details

Querying Hosts in the Cluster

To view all hosts running in the Hadoop cluster, click **Hosts** on the top bar of the window. Figure 2-5 shows the **All Hosts** screen.

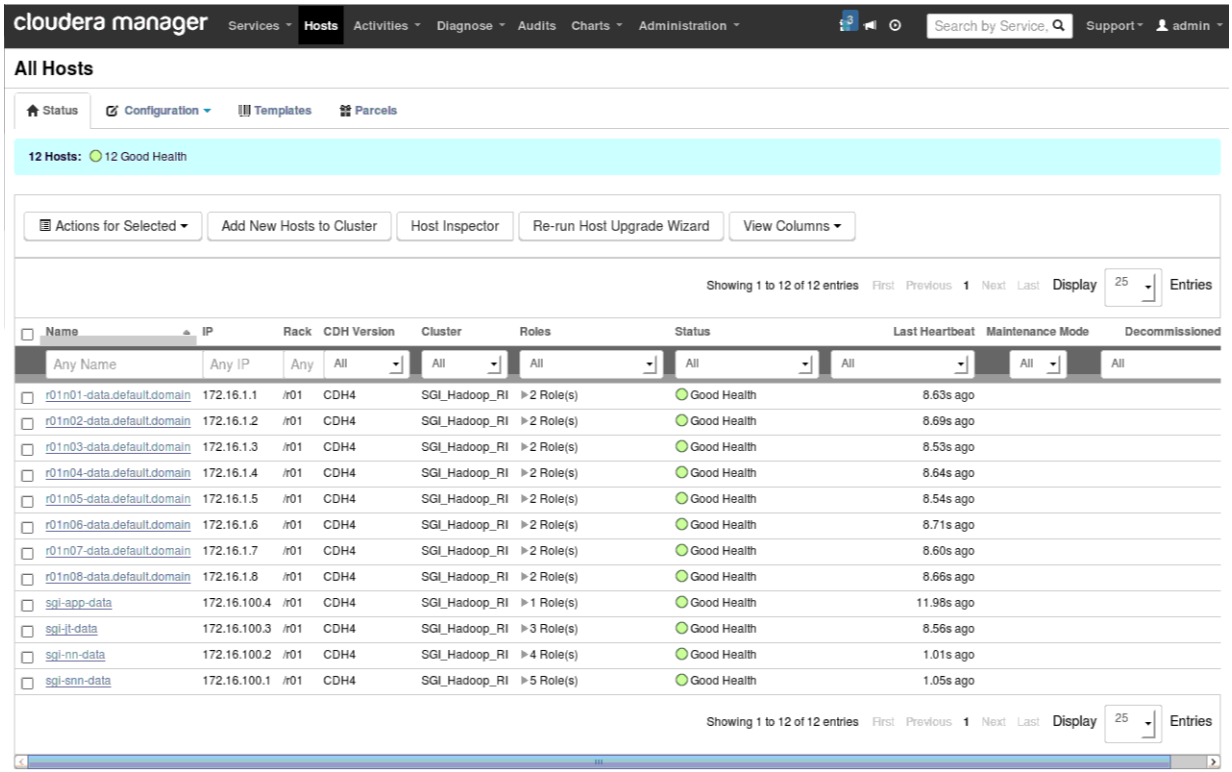

Figure 2-5 **All Hosts** Screen

Enabling Cloudera Manager Enterprise Features

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

Adminstration —> **License**

Figure 2-6 shows the **License** screen.

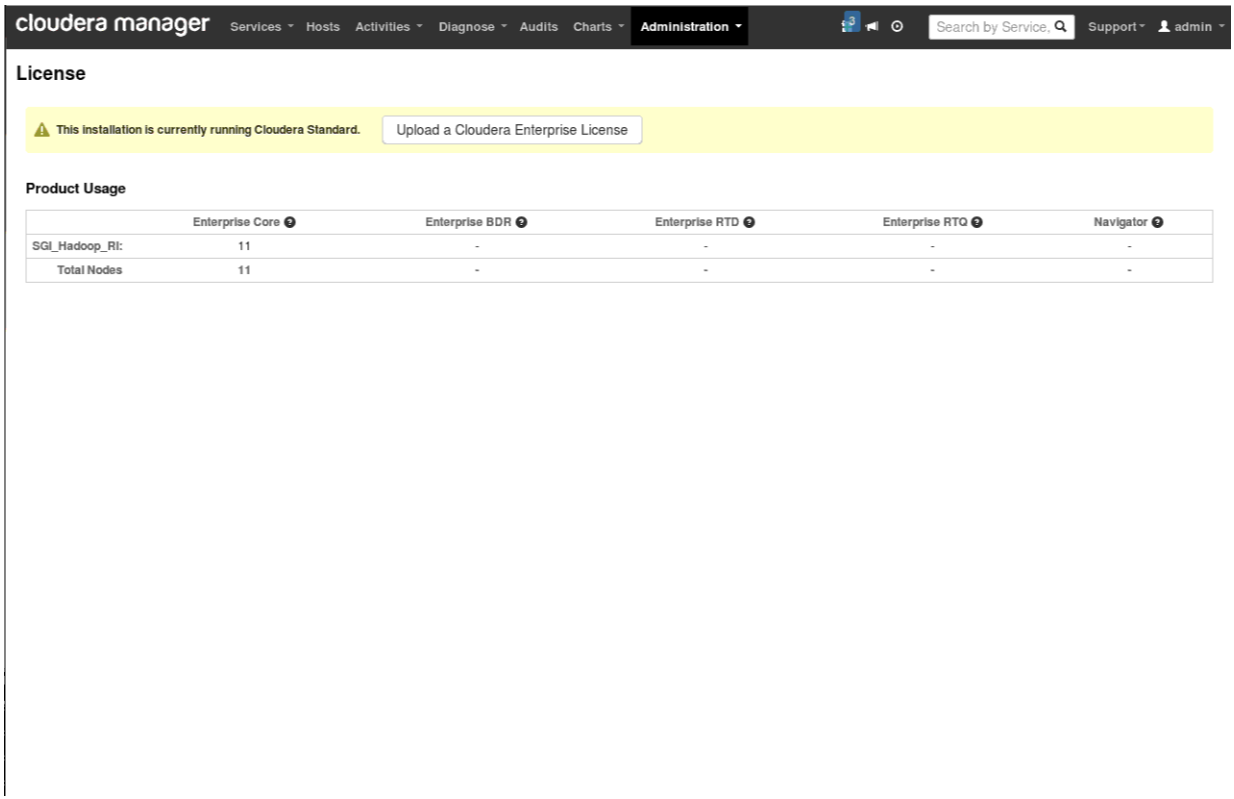

Figure 2-6 **License** Screen

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-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.*