



Cisco UCS Common Platform Architecture (CPA) for Big Data with Cloudera

Last Updated: October 25, 2013



Building Architectures to Solve Business Problems

About the Authors



Raghunath Nambiar

Raghunath Nambiar, Strategist, Data Center Solutions, Cisco Systems

Raghunath Nambiar is a Distinguished Engineer at Cisco's Data Center Business Group. His current responsibilities include emerging technologies and big data strategy.



Greg Rahn

Greg Rahn, Solutions Architect, Partner Engineering Group, Cloudera, Inc.

Greg Rahn is a Solutions Architect in the Partner Engineering Group at Cloudera. His focus is helping Cloudera's partners optimize their hardware platforms for Hadoop.



Manan Trivedi

Manankumar Trivedi, Performance Engineer, Data Center Solution, Cisco Systems

Manan is a member of the solution engineering team focusing on big data infrastructure and performance. He holds masters of science degree from Stratford University.

Acknowledgments

The authors acknowledge Scott Armstrong, Karthik Kulkarni, and Ashwin Manjunatha for their contributions in developing this document.

About Cisco Validated Design (CVD) Program

The CVD program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information visit:

<http://www.cisco.com/go/designzone>

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0809R)

© 2013 Cisco Systems, Inc. All rights reserved.

Cisco UCS Common Platform Architecture (CPA) for Big Data with Cloudera

Overview

Today's enterprise must store and analyze massive amounts of unstructured data to uncover crucial insights that can lead to competitive advantage. Cisco and Cloudera offer high-performance infrastructure for big data analytics that is cost effective, flexible, and scalable.

While the competitive pressure on enterprises vastly increases, the amount of data being ingested and managed has exploded and is accelerating quickly. At the same time, the need for timely and more accurate analytics has also increased. As a result, the need for a cost-effective, flexible, and scalable infrastructure to store and process data has never been greater. Cisco and Cloudera have partnered to deliver tested and certified Hadoop infrastructure solutions and ongoing support that help take the time and risk out of deploying Hadoop. These solutions provide a comprehensive, enterprise-class platform for Hadoop applications powered by Cloudera Enterprise, tested by Cisco and certified by the Cloudera Certified Technology program to streamline deployment and reduce risk.

Audience

This document describes the architecture and deployment procedures of Cloudera Enterprise on the Cisco UCS CPA for Big data. The intended audience of this document includes, but is not limited to, sales engineers, field consultants, professional services, IT managers, partner engineering and customers who want to deploy Cloudera Enterprise on the Cisco UCS CPA for Big Data.

Big Data and Apache Hadoop

The volume, variety, and velocity of unstructured data coming from a profusion of inter-connected devices are unprecedented. "Big Data" refers to data that just doesn't fit easily into traditional relational models because it's often a mix of structured and unstructured data, it comes in too fast, and it's too expensive to store in a way that's accessible. The ability to leverage big data requires a new type of data



Corporate Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

Copyright © 2013 Cisco Systems, Inc. All rights reserved.

management platform that can adequately capture and extract value from all of it. Apache Hadoop is the open source framework that lets organizations mine the insights of new and emerging types of information, a capability that simply did not exist before.

Unlike relational databases that are designed exclusively for SQL, the operations that need to be performed on this “Big data” are as diverse as the data itself. SQL is no longer the only desired method to prosecute data. The power of Hadoop is the ability to bring multiple computation frameworks to a single pool of data (within a single set of system resources), depending on what needs to be done: batch processing/transformations, interactive SQL, search, machine learning, statistical computation, and others. Previously doing each of these things required copying data sets into separate specialized systems, which added cost and complexity while diluting data integrity.

Moving beyond its roots in Web 2.0 technology, Apache Hadoop is rapidly emerging as an essential enterprise platform. Consumer and commercial industries are all finding applications for big data analytics, particularly as they are faced with new challenges in today’s web and social content and interaction models. But taking advantage of Hadoop is not simple – it’s a complex distributed system comprised of a dozen different open source projects.

Together, Cisco and Cloudera are well positioned to help organizations exploit the valuable business insights in all their data, regardless of whether it’s structured, semi structured or unstructured. Cloudera is the leading provider of enterprise-grade Hadoop infrastructure software and services, and the leading contributor to the Apache Hadoop project overall. Cloudera provides an enterprise-ready Hadoop-based solution known as Cloudera Enterprise, which includes their market leading open source Hadoop distribution (CDH), their comprehensive management system (Cloudera Manager), and technical support. Cisco has been the leader in networking for decades, providing proven solutions that meet critical business requirements. Cisco UCS C-Series Rack-Mount Servers based on Intel® Xeon® processors complete these offerings, delivering an integrated Hadoop infrastructure.

Cisco UCS Common Platform Architecture (CPA) for Big Data

The Cloudera Hadoop Reference Configuration is based on [Cisco UCS Common Platform Architecture \(CPA\) for Big Data](#), a highly scalable architecture designed to meet a variety of scale-out application demands with seamless data integration and management integration capabilities built using the following components:

- **Cisco UCS 6200 Series Fabric Interconnects:** The Cisco UCS 6200 Series Fabric Interconnects are a core part of Cisco UCS, providing both network connectivity and management capabilities across Cisco UCS 5100 Series Blade Server Chassis as well as Cisco UCS C-Series Rack-Mount Servers. Typically deployed in redundant pairs, the Fabric Interconnects offer line-rate, low-latency, lossless, 10 Gigabit Ethernet connectivity and unified management with Cisco UCS Manager in a highly available management domain.
- **Cisco UCS 2200 Series Fabric Extenders:** Cisco UCS 2200 Series Fabric Extenders behave as remote line cards for a parent switch and provide a highly scalable and extremely cost-effective unified server-access platform.
- **Cisco UCS C-Series Rack Mount Servers:** Cisco UCS C240 M3 Rack-Mount Servers are 2-socket servers based on Intel Xeon E-2600 series processors and supporting up to 768 GB of main memory. 24 Small Form Factor (SFF) disk drives are supported in performance optimized option and 12 Large Form Factor (LFF) disk drives are supported in capacity option, along with 4 Gigabit Ethernet LAN-on-motherboard (LOM) ports.

- **Cisco UCS Virtual Interface Cards (VICs):** Unique to Cisco, Cisco UCS Virtual Interface Cards incorporate next-generation converged network adapter (CNA) technology from Cisco, and offer dual 10-Gbps ports designed for use with Cisco UCS C-Series Rack-Mount Servers. Optimized for virtualized networking, these cards deliver high performance and bandwidth utilization and support up to 256 virtual devices.
- **Cisco UCS Manager:** Cisco UCS Manager resides within the Cisco UCS 6200 Series Fabric Interconnects. It makes the system self-aware and self-integrating, managing all of the system components as a single logical entity. Cisco UCS Manager can be accessed through an intuitive graphical user interface (GUI), a command-line interface (CLI), or an XML application-programming interface (API). Cisco UCS Manager uses service profiles to define the personality, configuration, and connectivity of all resources within Cisco UCS, radically simplifying provisioning of resources so that the process takes minutes instead of days. This simplification allows IT departments to shift their focus from constant maintenance to strategic business initiatives.

Solution Overview

The current version of the Cisco UCS CPA for Big Data offers two options depending on the compute and storage requirements:

- **High Performance Configuration**—Offers balance of compute power with IO bandwidth optimized for price/performance is built using C240 M3 rack servers powered by two Intel Xeon E5-2665 processors (16 cores) with 256 GB of memory and 24 1TB SFF disk drives.
- **High Capacity Configuration**—Optimized for low cost per terabyte, is built using C240 M3 rack servers powered by two Intel Xeon E5-2640 processors (12 cores) with 128GB of memory and 12 3TB LFF disk drives.

The solutions are offered in single rack and multiple rack scale. The single rack configuration consists of two Cisco UCS 6296UP 96-port Fabric Interconnects (supports up to 10 racks, 160 servers), two Cisco Nexus 2232PP 10GigE Fabric Extenders and 16 Cisco UCS C240 M3 Rack-Mount Servers (High-Performance Configuration or High-Capacity Configuration). Each server in the configuration connects to the unified fabric through two active-active 10-GigE links using a Cisco UCS VNIC.

Multi-rack configurations include two Cisco Nexus 2232PP fabric extenders and 16 Cisco UCS C240 M3 Rack-Mount Servers for every additional rack.




Note

This document provides Cisco Validated Design of deploying Cloudera Enterprise on single rack High Performance Configuration.

The configuration consists of:

- Two Cisco UCS 6296UP Fabric Interconnects
- Two Cisco Nexus 2232PP Fabric Extenders
- Sixteen Cisco UCS C240M3 Rack-Mount Servers
- Two vertical PDUs, country specific
- Cisco R42610 standard 42U Rack

Figure 1 Cisco UCS CPA for Big Data Rack Front and Rear View



The High Performance Rack and High Capacity Rack configurations are available through the [Cisco SmartPlay program](#) in single SKU bundles. PDUs and Rack is not included in the bundle.



Note

Contact your Cisco representative for country specific PDU information.

Rack and PDU Configuration

The configuration consists of two vertical PDUs, two Cisco UCS 6296UP Fabric Interconnects, two Cisco Nexus 2232PP Fabric Extenders and sixteen Cisco UCS C240M3 Servers are connected to each of the vertical PDUs for redundancy; thereby, ensuring availability during power source failure.

The rack configuration is shown in [Table 1](#).

Table 1 Rack Configuration Details

Cisco 42U Racks	Components
42	Cisco UCS FI 6296UP
41	
40	Cisco UCS FI 6296UP
39	
38	Cisco Nexus FEX 2232PP
37	Cisco Nexus FEX 2232PP

Table 1 **Rack Configuration Details**

Cisco 42U Racks	Components
36	Unused
35	Unused
34	Unused
33	Unused
32	Cisco UCS C240M3
31	
30	Cisco UCS C240M3
29	
28	Cisco UCS C240M3
27	
26	Cisco UCS C240M3
25	
24	Cisco UCS C240M3
23	
22	Cisco UCS C240M3
21	
20	Cisco UCS C240M3
19	
18	Cisco UCS C240M3
17	
16	Cisco UCS C240M3
15	
14	Cisco UCS C240M3
13	
12	Cisco UCS C240M3
11	
10	Cisco UCS C240M3
9	
8	Cisco UCS C240M3
7	
6	Cisco UCS C240M3
5	
4	Cisco UCS C240M3
3	
2	Cisco UCS C240M3
1	

Server Configuration and Cabling

The C240 M3 rack server is equipped with Intel Xeon E5-2665 processors, 256 GB of memory, Cisco UCS Virtual Interface Card 1225 Cisco, Cisco LSI MegaRAID SAS 9266-8i - storage controller, and 24 x 1TB 7.2K SATA disk drives.

Figure 3 illustrates the physical connectivity of Cisco UCS C240M3 Servers to Cisco Nexus 2232PP Fabric Extenders and Cisco UCS 6296UP Fabric Interconnects.

Figure 2 Cisco Hardware Connectivity





Figure 3 shows the ports of the Cisco Nexus 2232PP Fabric Extender connecting the Cisco UCS C240M3 Servers. 16 Cisco UCS C220M3 Servers are used in the rack configuration offered by Cisco.

Figure 3 Connectivity Diagram of Cisco Nexus 2232PP FEX and Cisco UCS C220M3 Servers



Note Cisco UCS Manager version used for this deployment is UCS 2.1(1e).

For more information on configuring single-wire management, see:

http://www.cisco.com/en/US/docs/unified_computing/ucs/c-series_integration/ucsm2.1/b_UCSM2-1_C-Integration_chapter_010.html

For more information on physical connectivity illustrations and cluster setup, see:

http://www.cisco.com/en/US/docs/unified_computing/ucs/c-series_integration/ucsm2.1/b_UCSM2-1_C-Integration_chapter_010.html#reference_FE5B914256CB4C47B30287D2F9CE3597

Software Distributions and Versions

The software distributions required versions are:

- [Cloudera Enterprise, page 12](#)
- [Red Hat Enterprise Linux \(RHEL\), page 12](#)

Cloudera Enterprise

Cloudera software for Cloudera Distribution for Apache Hadoop is v4.x (CDH4). For more information on Cloudera, see:

www.cloudera.com.

Red Hat Enterprise Linux (RHEL)

The operating system supported is Red Hat Enterprise Linux Server 6.2. For more information on the Linux support, see:

www.redhat.com.

Table 2 provides the software versions tested and validated for this model.

Table 2 Software Version Details

Layer	Components	Version or Release
Compute	Cisco UCS C240M3	1.4.7cc
Network	Cisco UCS 6296UP	UCS 2.1(1e)
	Cisco Nexus 2232PP	5.1(3)N2(2.11a)
	Cisco UCS VIC 1225 Firmware	2.1(1e)
	Cisco UCS VIC 1225 Driver	2.1.1.41
Storage	LSI 9266-8i Firmware	23.7.0-0039
	LSI 9266-8i Driver	06.504.01.00

Table 2 **Software Version Details**

Layer	Components	Version or Release
Software	Red Hat Enterprise Linux Server	6.2 (x86_64)
	Cisco UCSM	2.1(1e)
	CDH	4.1.3 (x86_64)
	Cloudera Manager	4.5 (x86_64)

Fabric Configuration

This section provides details for configuring a fully redundant, highly available configuration for a FlexPod Select for Hadoop. Follow these steps to configure Cisco 6296UP Fabric Interconnect.

1. Configure FI A
2. Configure FI B
3. Connect to IP address of FI A using web browser. Launch Cisco UCS Manger
4. Edit the chassis discovery policy.
5. Enable server and Uplink Ports
6. Create pools and polices for service profile template.
7. Create SP template, 16 profiles
8. Start discover process
9. Associate to server

Performing an Initial Setup of Cisco UCS 6296UP Fabric Interconnects

Follow these steps for initial setup of the Cisco UCS 6296 Fabric Interconnects:

Cisco UCS 6296 FI A

1. Connect to the console port on the first Cisco UCS 6296 Fabric Interconnect.
2. At the configuration method prompt, enter **console**.
3. If asked to either do a new setup or restore from backup, enter **setup** to continue.
4. Enter **y** to continue to set up a new fabric interconnect.
5. Enter **y** to enforce strong passwords.
6. Enter the password for the admin user.
7. Enter the same password again to confirm the password for the admin user.
8. When asked if this fabric interconnect is part of a cluster, enter **y** to continue.
9. Enter **A** for the switch fabric.
10. Enter the cluster name for the system name.
11. Enter the Mgmt0 IPv4 address for management port on the fabric interconnect.
12. Enter the Mgmt0 IPv4 subnet mask for the management port on the fabric interconnect.

13. Enter the IPv4 address of the default gateway.
14. Enter the cluster IPv4 address.
15. To configure DNS, enter **y**.
16. Enter the DNS IPv4 address.
17. Enter **y** to set up the default domain name.
18. Enter the default domain name.
19. Review the settings that were printed to the console, and if they are correct, enter **yes** to save the configuration.
20. Wait for the login prompt to make sure the configuration is saved successfully.

Cisco UCS 6296UP FI B

1. Connect to the console port on the second Cisco UCS 6296 Fabric Interconnect.
2. At the configuration method prompt, enter **console**.
3. The installer detects the presence of the partner fabric interconnect and adds this fabric interconnect to the cluster. Enter **y** to continue the installation.
4. Enter the admin password for the first fabric interconnects.
5. Enter the Mgmt0 IPv4 address for the management port on the subordinate fabric interconnect.
6. Enter **y** to save the configuration.
7. Wait for the login prompt to make sure the configuration is saved successfully.

For more information on configuring Cisco UCS 6200 Series Fabric Interconnect, see:

http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/gui/config/guide/2.0/b_UCSM_GUI_Configuration_Guide_2_0_chapter_0100.html

Logging into Cisco UCS Manager

Follow these steps to log into Cisco UCS Manager:

1. Open a Web browser and type the Cisco UCS 6296UP Fabric Interconnect cluster address.
2. Select the Launch link to download the Cisco UCS Manager Software.
3. If a Security Alert dialog box appears, click **Yes** to accept the security certificate and continue.
4. In the Cisco UCS Manager launch page, click **Launch UCS Manager**.
5. When prompted, enter admin for the user name and enter the administrative password and click **Login** to log in to the Cisco UCS Manager GUI.

Upgrade Cisco UCS Manager Software to Version 2.1(1e)

This document assumes the use of UCS 2.1(1e). For more information on upgrading the software version to Cisco UCS 2.0 release, see:

http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/upgrading/from2.0/to2.1/b_UpgradingCiscoUCSFrom2.0To2.1.pdf

This link provides you information on upgrading Cisco UCS Manager software and Cisco UCS 6296 Fabric Interconnect software to version 2.1(1e).

**Note**

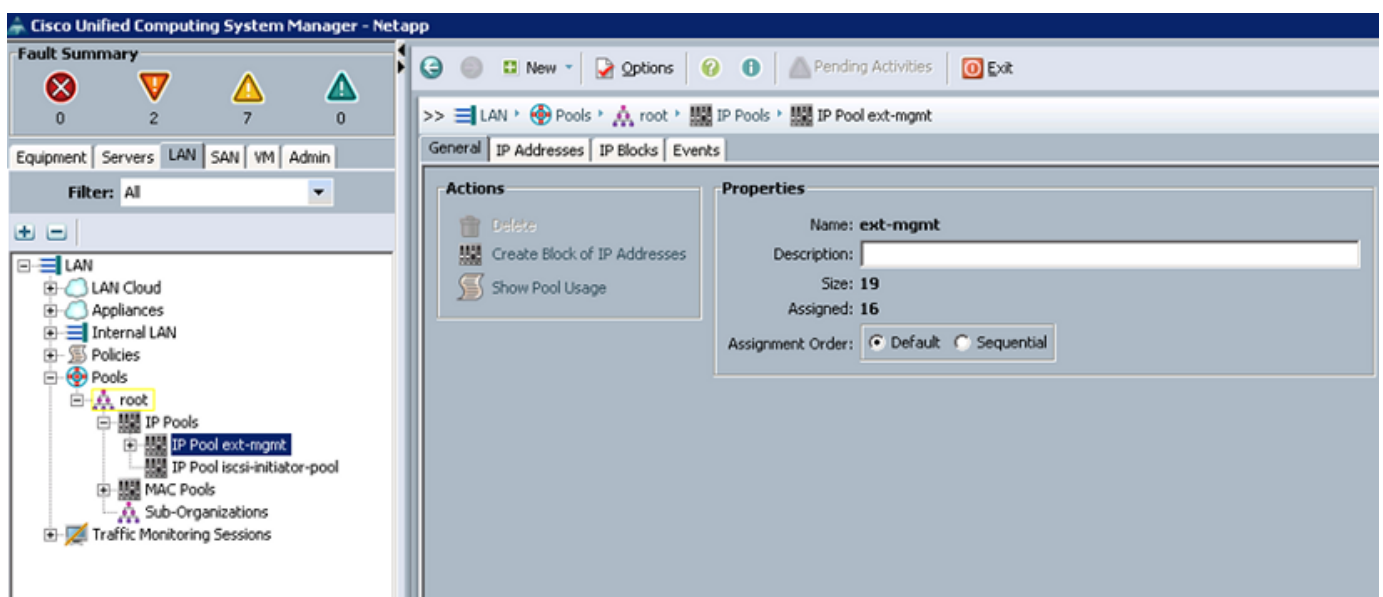
Make sure the Cisco UCS C-Series version 2.1(1e) software bundle is loaded on the Fabric Interconnects.

Adding a Block of IP Addresses for KVM Console

Follow these steps to create a block of KVM IP addresses for server access in the Cisco UCS Manager GUI:

1. Select the LAN tab at the top in the left pane in the UCSM GUI.
2. Select **Pools > IP Pools > IP Pool ext-mgmt** as shown in [Figure 4](#).

Figure 4 Management IP Pool in Cisco UCS Manager



3. Right-click the IP Pool ext-mgmt.
4. Select Create Block of IP Addresses. Create Block of IP Address window appears as shown in [Figure 5](#).

Figure 5 *Creating a Block of IP Addresses*

The screenshot shows a dialog box titled "Create Block of IP Addresses". It contains several input fields: "From" (0.0.0.0), "Size" (1), "Subnet Mask" (255.255.255.0), "Default Gateway" (0.0.0.0), "Primary DNS" (0.0.0.0), and "Secondary DNS" (0.0.0.0). The "From" and "Default Gateway" fields have red error indicators, suggesting they are not valid. There are "OK" and "Cancel" buttons at the bottom right.

5. Enter the starting IP address of the block and number of IPs needed as well as the subnet and gateway information.

Figure 6 *Entering the Block of IP Addresses*

The screenshot shows the same dialog box as Figure 5, but with different values: "From" (10.29.160.10), "Size" (25), "Subnet Mask" (255.255.255.0), "Default Gateway" (10.29.160.1), "Primary DNS" (0.0.0.0), and "Secondary DNS" (0.0.0.0). The "From" and "Default Gateway" fields now have small blue question mark icons, indicating they are valid. There are "OK" and "Cancel" buttons at the bottom right.

6. Click **OK** to create the IP block.
7. Click **OK** in the confirmation message box.

Editing the Chassis Discovery Policy

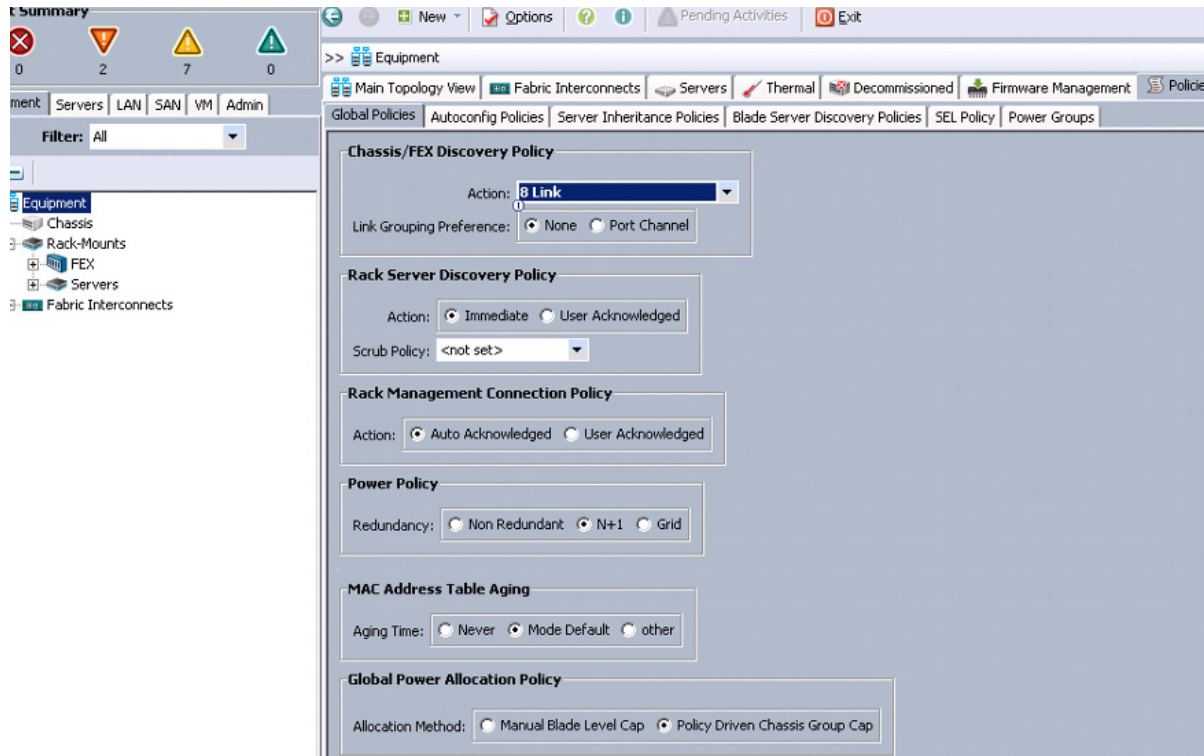
Setting the discovery policy now will simplify the addition of Cisco UCS B-Series Chassis in the future and additional fabric extenders for further C-Series connectivity.

To modify the chassis discovery policy, follow these steps:

1. Navigate to the Equipment tab in the left pane in the UCSM GUI.

- In the right pane, select the Policies tab.
- Under Global Policies, change the Chassis Discovery Policy to 8-link as shown in [Figure 7](#).

Figure 7 Editing the Chassis Discovery Policy



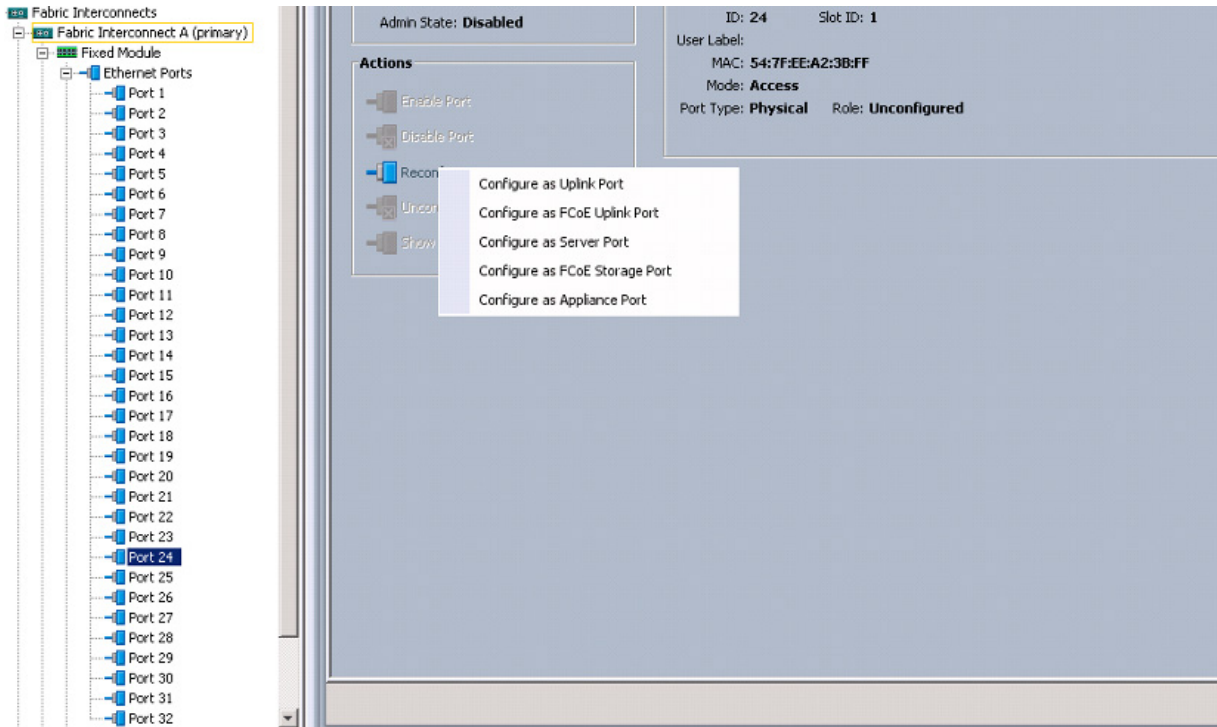
- Click **Save Changes** in the bottom right corner in the Cisco UCSM GUI.
- Click **OK**.

Enabling Server and Uplink Ports

To enable the server ports and uplink ports, follow these steps:

- Select the Equipment tab on the top left corner in the left pane in the UCSM GUI.
- Select **Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module**.
- Expand the Unconfigured Ethernet Ports.
- Select the number of ports that are connected to the Cisco Nexus 2232PP FEXs (8 per FEX), right-click them, and select **Reconfigure > Configure as a Server Port** as shown in [Figure 8](#).

Figure 8 **Enabling Server Ports**



5. Select port 1 that is connected to the Cisco Catalyst 2960-S switches, right-click them, and select **Reconfigure > Configure as Uplink Port**.
6. Select Show Interface and select 10GB for Uplink Connection.
7. A pop-up window appears to confirm your selection. Click **Yes**, then **OK** to continue.
8. Select **Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module**.
9. Expand the Unconfigured Ethernet Ports.
10. Select the number of ports that are connected to the Cisco Nexus 2232 FEXs (8 per FEX), right-click them, and select **Reconfigure > Configure as Server Port**.
11. A pop-up window appears to confirm your selection. Click **Yes**, then **OK** to continue.
12. Select port 1 that is connected to the Cisco Catalyst 2960-S switches, right-click and select **Reconfigure > Configure as Uplink Port**.
13. Select Show Interface and select 10GB for Uplink Connection.
14. A pop-up window appears to confirm your selection. Click **Yes**, then **OK** to continue.

Creating MAC Address Pools

Follow these steps to configure the necessary MAC address pools in the Cisco UCS Manager GUI:

1. Select the LAN tab in the left pane in the UCSM GUI.
2. Select **Pools > root**.
3. Right-click the MAC Pools under the root organization.
4. Select Create MAC Pool to create the MAC address pool.
5. Enter nosh for the name of the MAC pool.
6. (Optional) enter a description of the MAC pool.
7. Click **Next**.
8. Click **Add**.
9. Specify a starting MAC address.
10. Specify a size of the MAC address pool sufficient to support the available server resources. See [Figure 10](#), [Figure 11](#), and [Figure 12](#).

Figure 10 Specifying the First MAC Address and Size

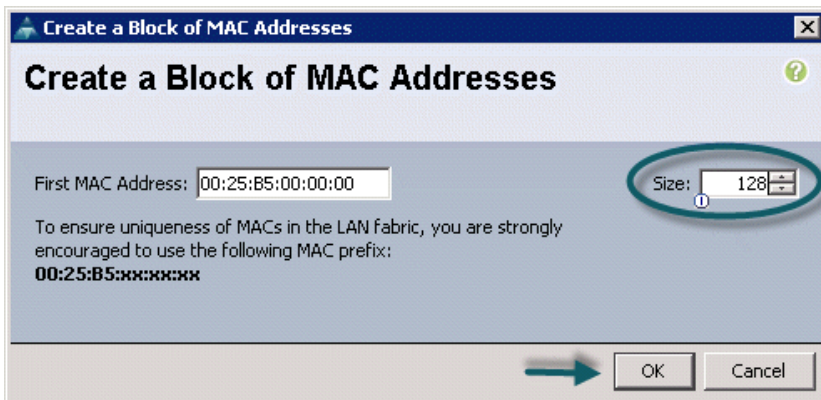
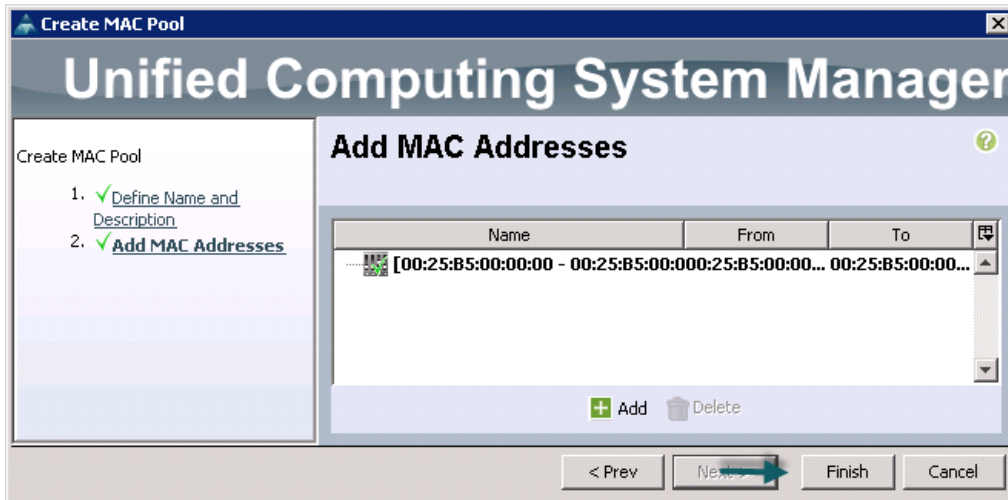
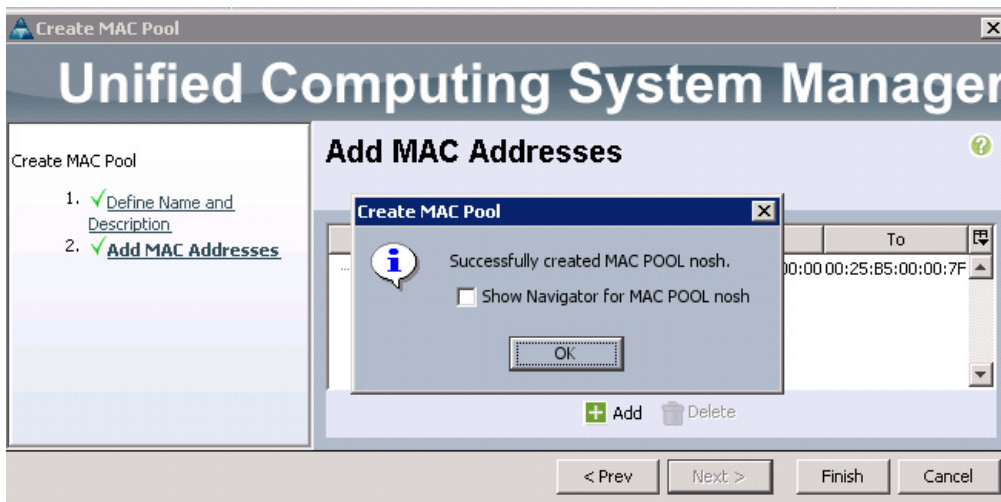


Figure 11 Adding a Range of MAC Addresses



11. Click **OK**.
12. Click **Finish**.
13. Click **OK** in the success message box.

Figure 12 Confirming Newly Added MAC Pool



Configuring VLANs

VLANs are configured as shown in [Table 3](#).

Table 3 *VLAN Configuration*

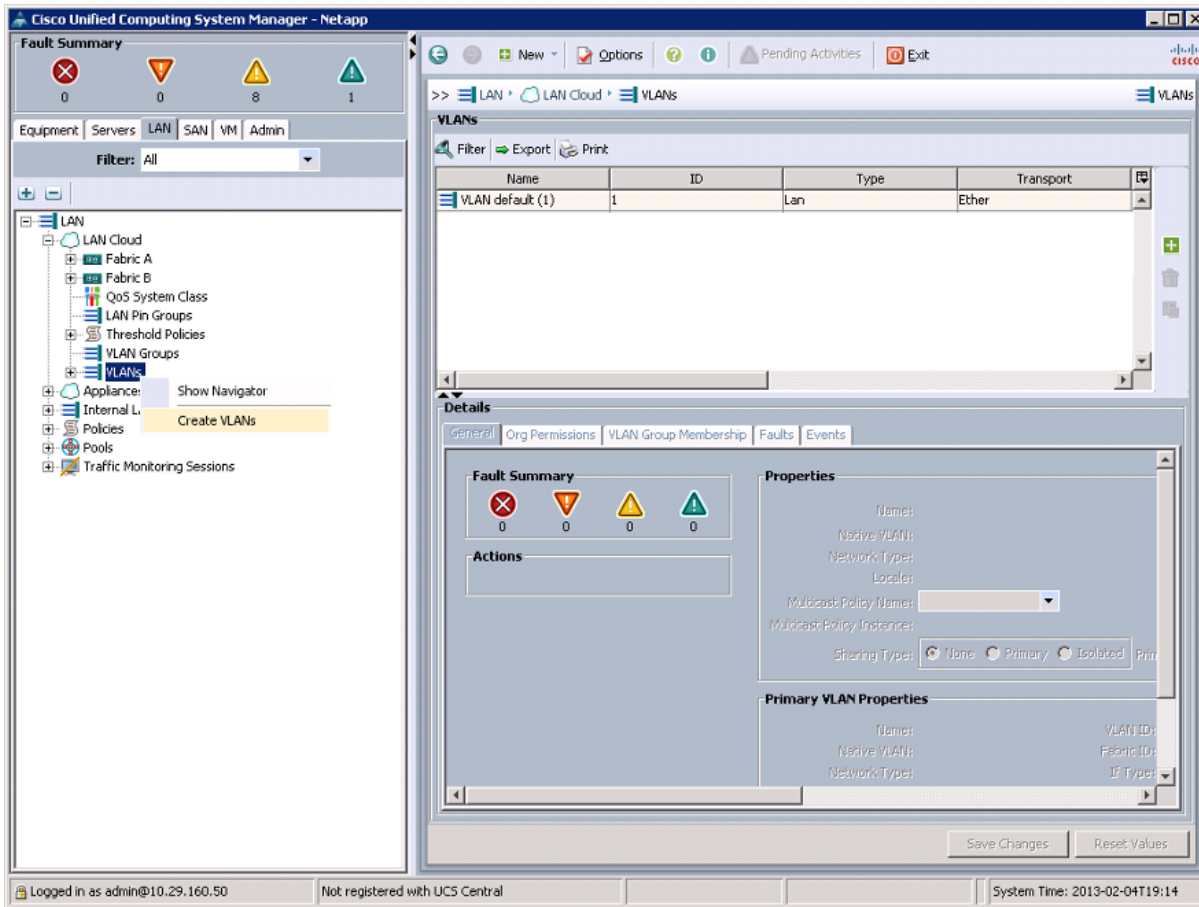
VLAN	Fabric	NIC Port	Function	Failover
vlan160_mgmt	A	eth0	Management, User connectivity	Fabric failover to B
vlan12_HDFS	B	eth1	Hadoop	Fabric failover to A
vlan11_DATA	A	eth2	SAN/NAS access, ETL	Fabric failover to B

All of the VLANs created need to be trunked to the upstream distribution switch connecting the fabric interconnects. For this implementation vlan160_mgmt is configured for management access and user connectivity, vlan12_HDFS is configured for Hadoop interconnect traffic and vlan11_DATA is configured for optional SAN/NAS access, heavy ETL etc.

Follow these steps to configure VLANs in the Cisco UCS Manager GUI:

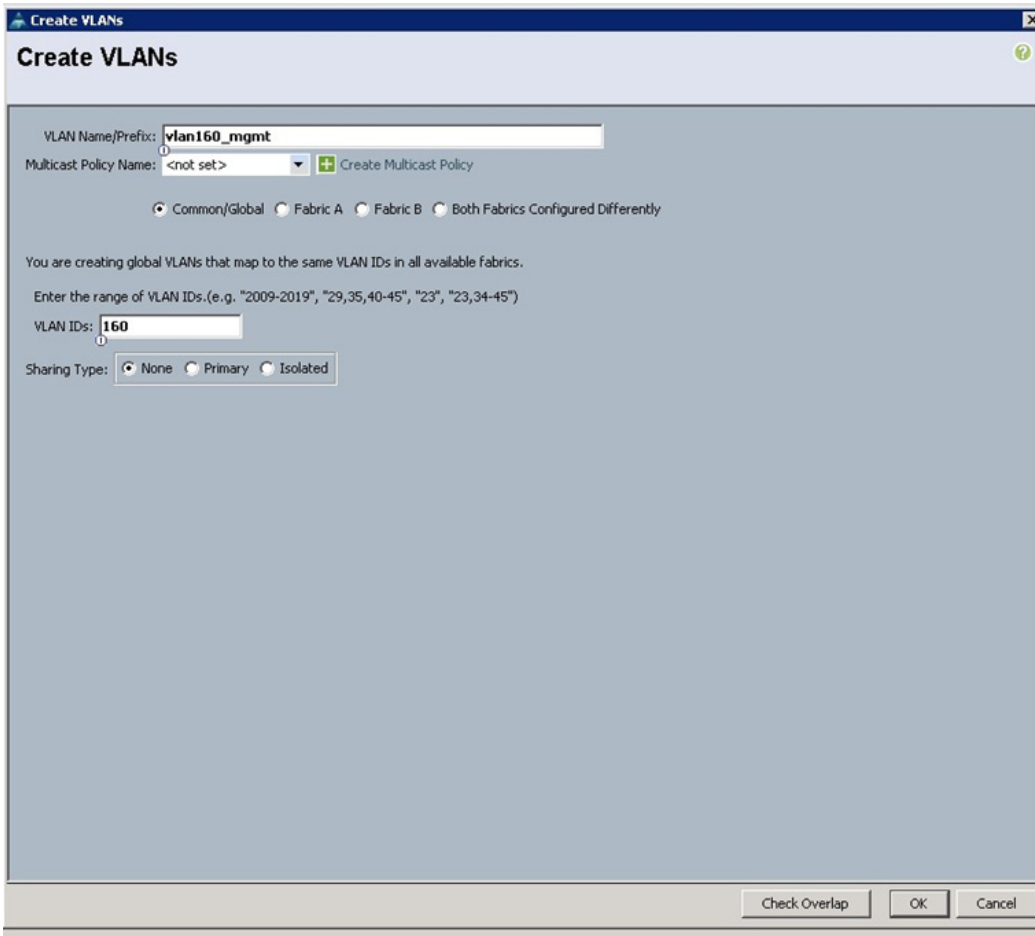
1. Select the LAN tab in the left pane in the UCSM GUI.
2. Select **LAN > VLANs**.
3. Right-click the VLANs under the root organization.
4. Select Create VLANs to create the VLAN.

Figure 13 Creating VLANs



5. Enter vlan160_mgmt for the VLAN Name.
6. Select Common/Global for vlan160_mgmt.
7. Enter 160 on VLAN IDs of the Create VLAN IDs.

Figure 14 Creating VLAN for Fabric A



8. Click **OK** and then, click **Finish**.
9. Click **OK** in the success message box.
10. Select the LAN tab in the left pane again.
11. Select **LAN > VLANs**.
12. Right-click the VLANs under the root organization.
13. Select Create VLANs to create the VLAN.
14. Enter vlan11_DATA for the VLAN Name.
15. Select Common/Global for vlan11_DATA.
16. Enter 11 on VLAN IDs of the Create VLAN IDs.

Figure 15 **Creating VLAN for Fabric B**

Create VLANs

VLAN Name/Prefix:

Multicast Policy Name:

Common/Global
 Fabric A
 Fabric B
 Both Fabrics Configured Differently

You are creating global VLANs that map to the same VLAN IDs in all available fabrics.

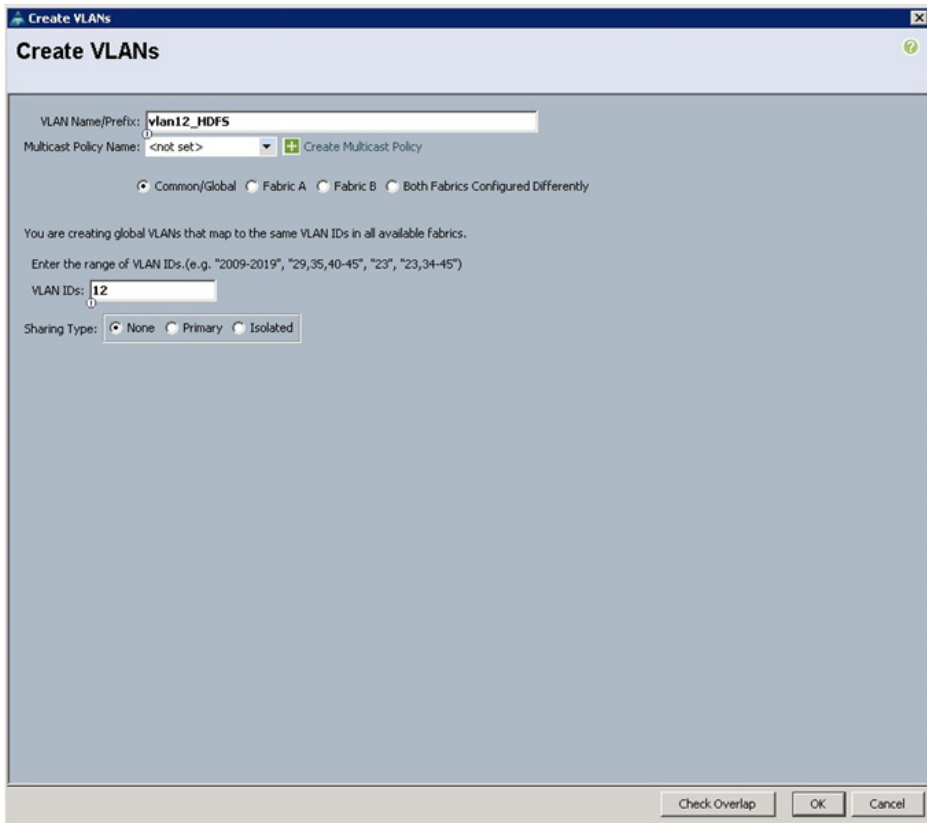
Enter the range of VLAN IDs.(e.g. "2009-2019", "29,35,40-45", "23", "23,34-45")

VLAN IDs:

Sharing Type: None
 Primary
 Isolated

17. Click **OK** and then, click **Finish**.
18. Click **OK** in the success message box.
19. Select the LAN tab in the left pane again.
20. Select **LAN > VLANs**.
21. Right-click the VLANs under the root organization.
22. Select Create VLANs to create the VLAN.
23. Enter vlan12_HDFS for the VLAN Name.
24. Select Common/Global for the vlan12_HDFS.
25. Enter 12 on VLAN IDs of the Create VLAN IDs.

Figure 16 Creating Global HDFS VLAN



26. Click **OK** then click **Finish**.

Creating Server Pool

A server pool contains a set of servers. These servers typically share the same characteristics. Those characteristics can be their location in the chassis, or an attribute such as server type, amount of memory, local storage, type of CPU, or local drive configuration. You can manually assign a server to a server pool, or use server pool policies and server pool policy qualifications to automate the assignment.

Follow these steps to configure the server pool in the Cisco UCS Manager GUI:

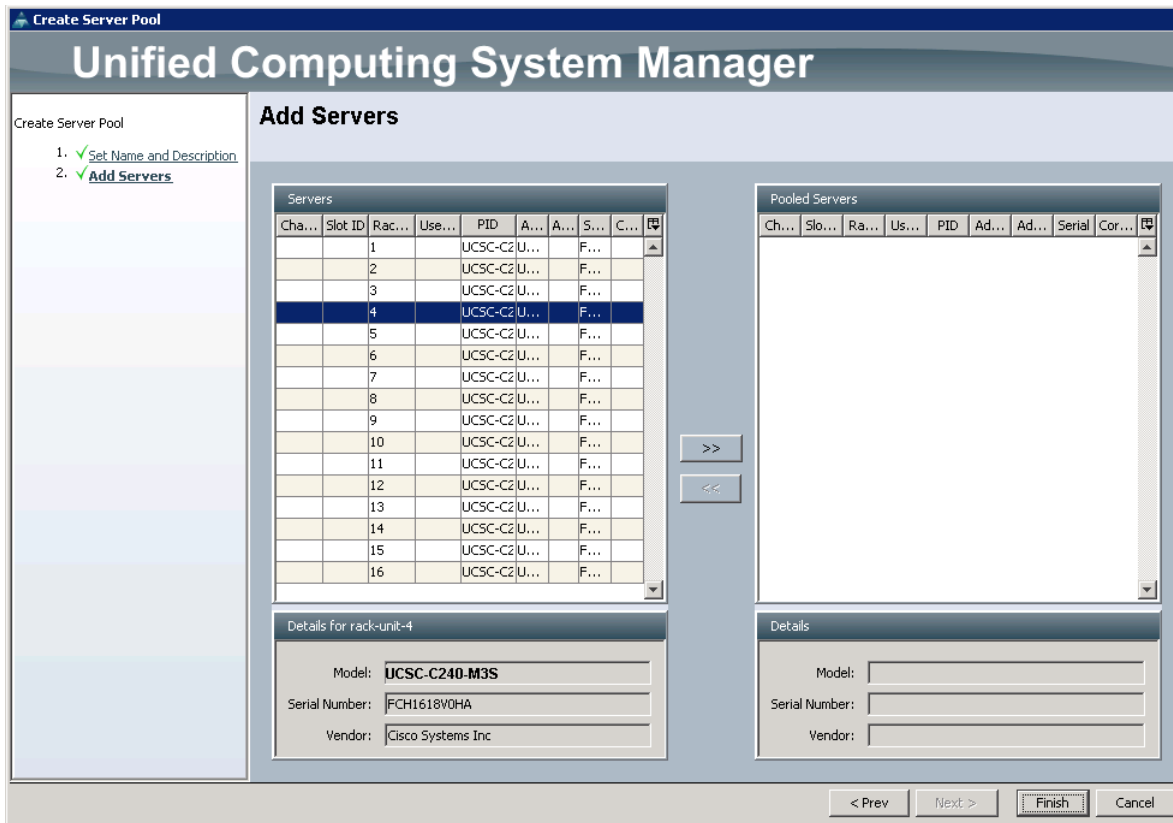
1. Select the Servers tab in the left pane in the Cisco UCS Manager GUI.
2. Select **Pools > root**.
3. Right-click the Server Pools.
4. Select Create Server Pool.
5. Enter ucs for the Server Pool Name.
6. (Optional) enter a description for the organization.

Figure 17 **Creating Server Pool**

The screenshot shows the 'Create Server Pool' wizard in the Unified Computing System Manager. The window title is 'Create Server Pool' and the main header is 'Unified Computing System Manager'. The wizard is in the 'Set Name and Description' step, which is indicated by a green checkmark in the left-hand navigation pane. The navigation pane shows two steps: '1. Set Name and Description' (checked) and '2. Add Servers'. The main area contains a 'Name' field with the text 'ucsd' and a 'Description' field. At the bottom, there are four buttons: '< Prev', 'Next >', 'Finish', and 'Cancel'.

7. Click **Next** to add servers.
8. Select all the Cisco UCS C220M3 servers to be added to the nosh server pool. Click >> to add them to the pool.
9. Click **Finish**.
10. Click **OK** and then click **Finish**.

Figure 18 Adding Server Pool



Creating Policies for Service Profile Template

Creating Host Firmware Package Policy

Firmware management policies allow the administrator to select the corresponding packages for a given server configuration. These often include adapter, BIOS, board controller, FC adapters, HBA option ROM, and storage controller properties.

Follow these steps create a firmware management policy for a given server configuration in the Cisco UCS Manager GUI:

1. Select the Servers tab in the left pane in the UCSM GUI.
2. Select **Policies > root**.
3. Right-click Host Firmware Packages.
4. Select Create Host Firmware Package.
5. Enter ucs as the Host firmware package name.
6. Select **Simple** radio button to configure the Host Firmware package.
7. Select the appropriate Rack package that you have.
8. Click **OK** to complete creating the management firmware package.

- Click **OK**.

Figure 19 **Creating Host Firmware Package**

Creating QoS Policies

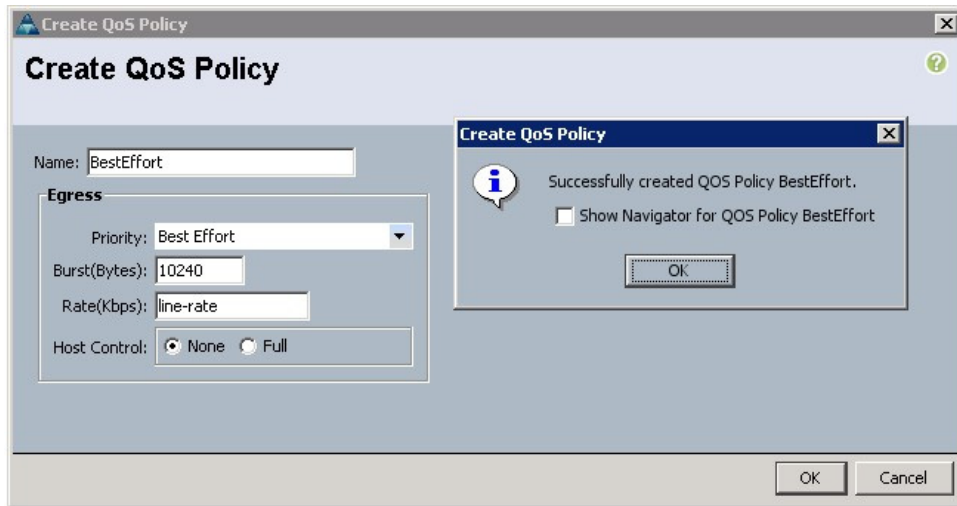
Follow these steps to create QoS policy for a given server configuration in the Cisco UCS Manager GUI:

BestEffort Policy

- Select the LAN tab in the left pane in the UCSM GUI.
- Select **Policies > root**.
- Right-click QoS Policies and select Create QoS Policy.
- Enter BestEffort as the name of the policy.
- Select Best Effort for Priority from the drop down menu.
- Keep the Burst (Bytes) field as default, which is 10240.
- Keep the Rate (Kbps) field as default, which is line-rate.

8. Make sure the Host Control radio button is **None**.
9. Click **OK**.

Figure 20 *Creating QoS Policy - BestEffort*

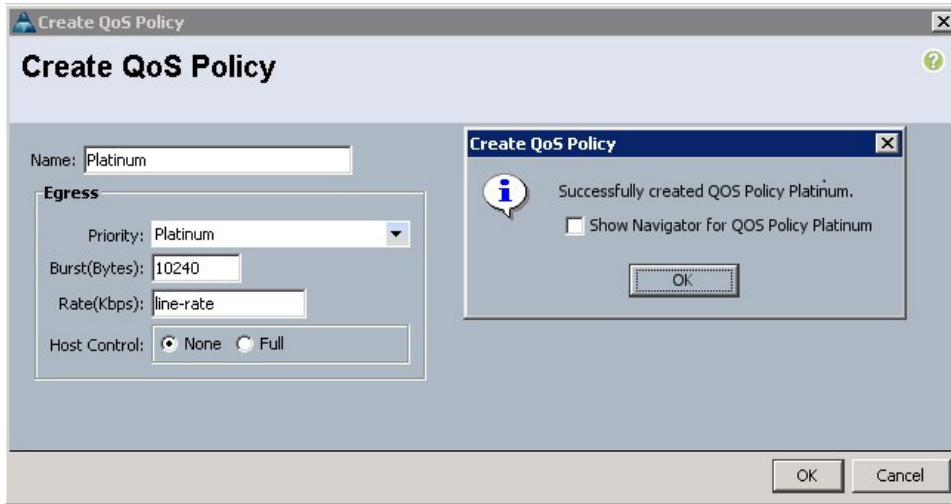


10. In the pop-up window, click **OK** to complete the QoS policy creation.

Platinum Policy

1. Select the LAN tab in the left pane in the UCSM GUI.
2. Select **Policies > root**.
3. Right-click QoS Policies and select Create QoS Policy.
4. Enter Platinum as the name of the policy.
5. Select Platinum for Priority from the drop down menu.
6. Keep the Burst (Bytes) field as default, which is 10240.
7. Keep the Rate (Kbps) field as default, which is line-rate.
8. Make sure the Host Control radio button is **None**.
9. Click **OK**.
10. In the pop-up window, click **OK** to complete the QoS policy creation.

Figure 21 Creating QoS Policy - Platinum

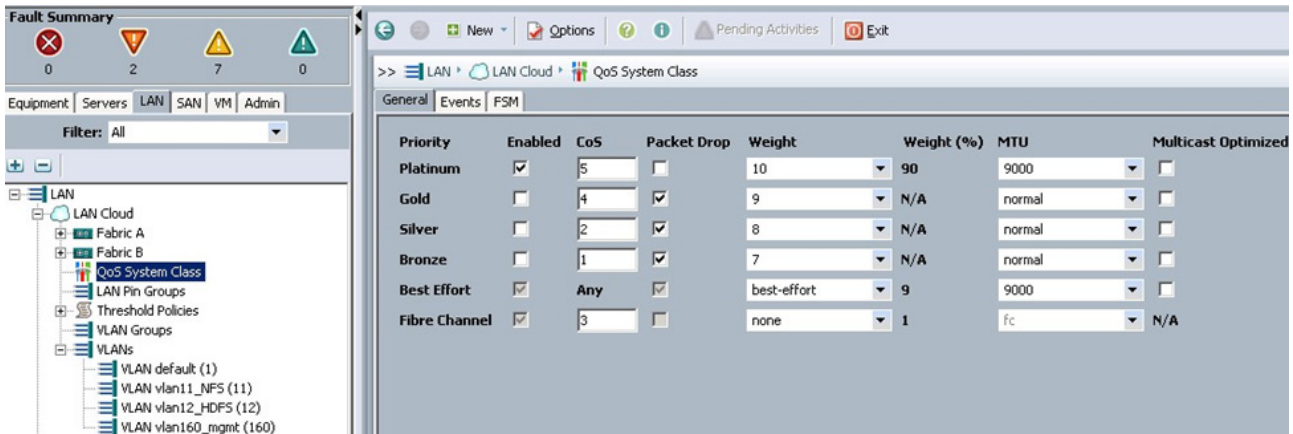


Setting Jumbo Frames

These steps provide details for setting Jumbo frames and enabling the quality of service in the Cisco UCS Fabric:

1. Select the Servers tab in the left pane in the UCSM GUI.
2. Select **LAN Cloud > QoS System Class**.
3. In the right pane, select the General tab.
4. In the Platinum row, enter 9000 for MTU.
5. In the Best Effort row, enter 9000 for MTU.
6. Check the Enabled check box next to Platinum.

Figure 22 Setting Jumbo Frame in Cisco UCS Fabric



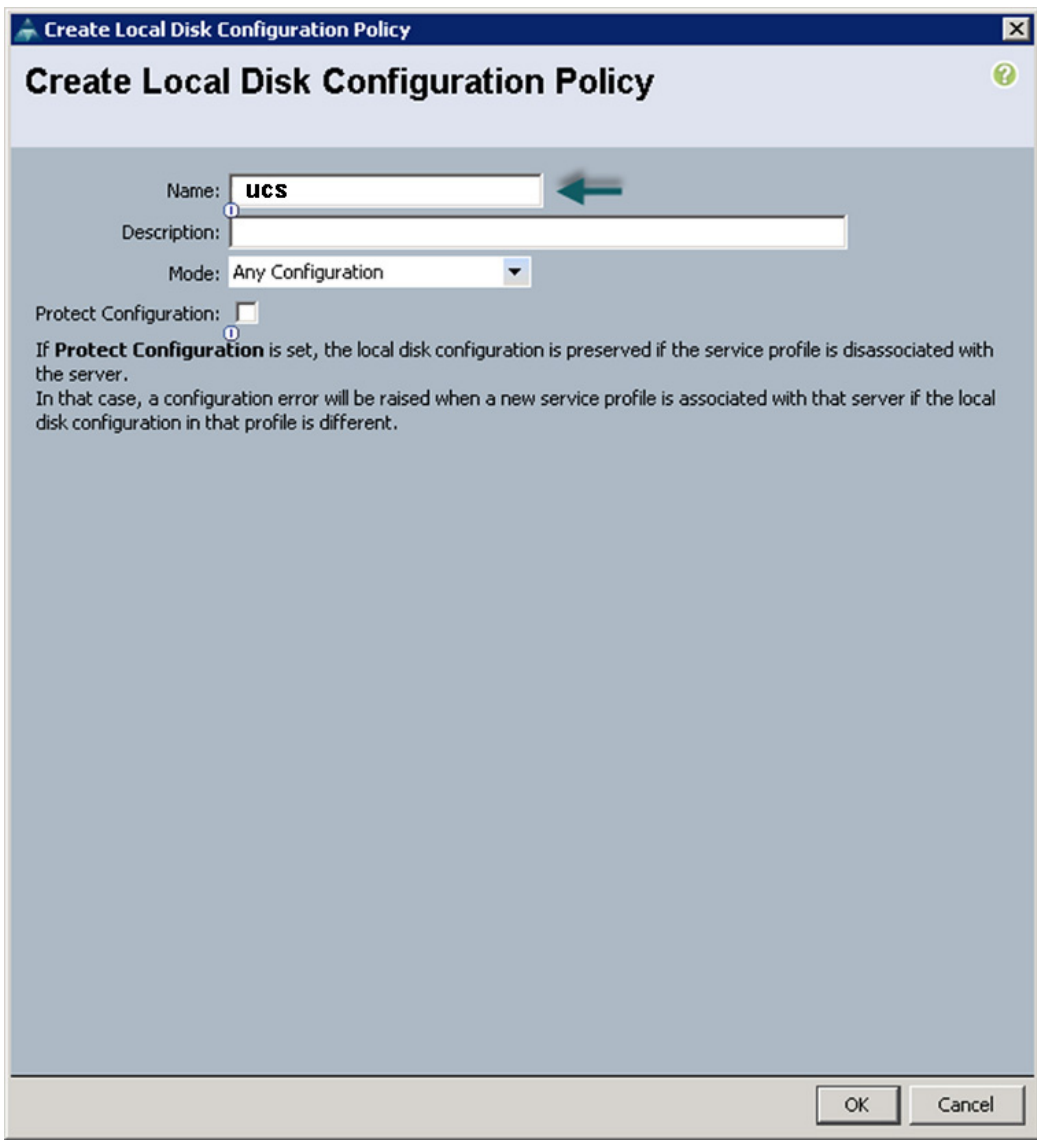
7. Click **Save Changes**.
8. Click **OK**.

Create a Local Disk Configuration Policy

Follow these steps to create local disk configuration in the Cisco UCS Manager GUI:

1. Select the Servers tab in the left pane in the UCSM GUI.
2. Select **Policies > root**.
3. Right-click Local Disk Config Policies.
4. Select Create Local Disk Configuration Policy.
5. Enter ucs as the local disk configuration policy name.
6. Change the Mode to Any Configuration. Uncheck the Protect Configuration check box.

Figure 23 *Configuring Local Disk Policy*



7. Click **OK** to create the Local Disk Configuration Policy.

8. Click **OK**.

Create a Server BIOS Policy

The BIOS policy feature in Cisco UCS automates the BIOS configuration process.

The traditional method of setting the BIOS is manual and often error-prone. By creating a BIOS policy and assigning the policy to a server or group of servers, you can have the transparency in BIOS settings and configuration.

Follow these steps to create a server BIOS policy in the Cisco UCS Manager GUI:

1. Select the Servers tab in the left pane in the UCSM GUI.
2. Select **Policies > root**.
3. Right-click BIOS Policies.
4. Select Create BIOS Policy.
5. Enter ucs as the BIOS policy name.
6. Change the BIOS settings as per [Figure 24](#), [Figure 25](#), [Figure 26](#), and [Figure 27](#).
7. Click **Finish** to complete creating the BIOS policy.
8. Click **OK**.

Figure 24 *Creating BIOS Policy*

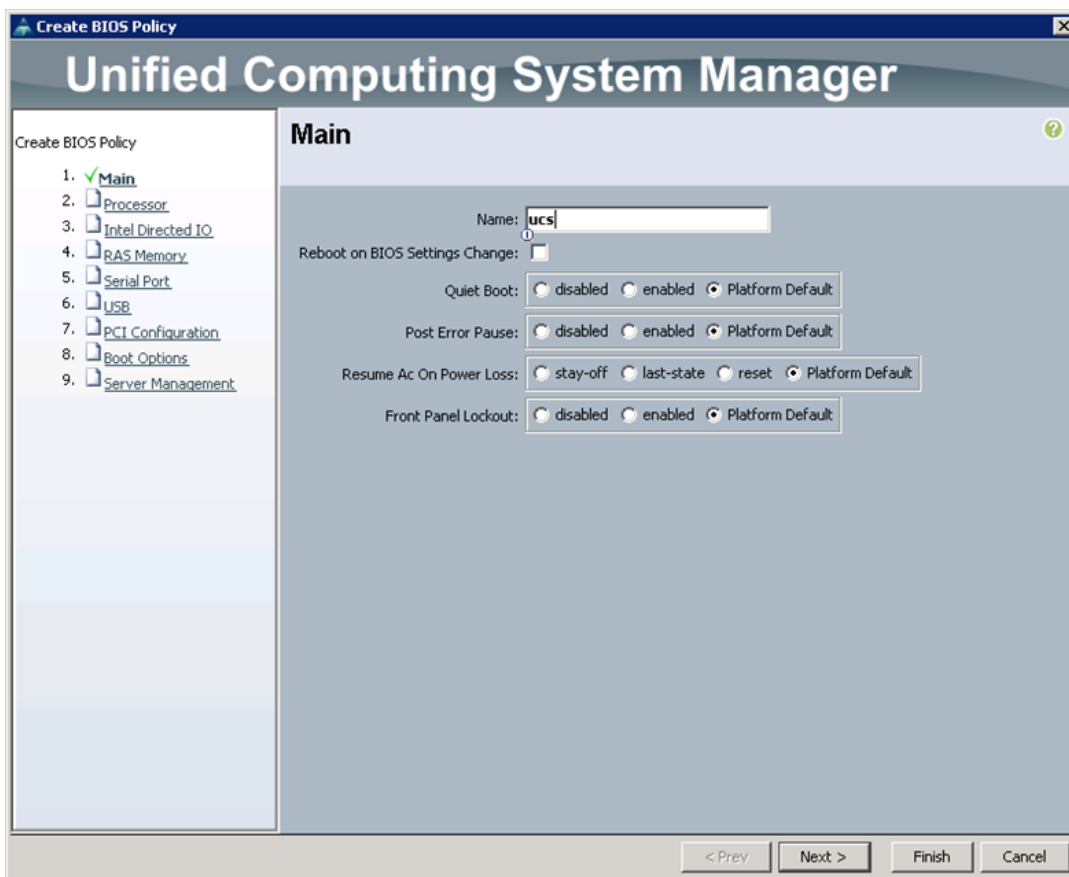


Figure 25 Processor Settings

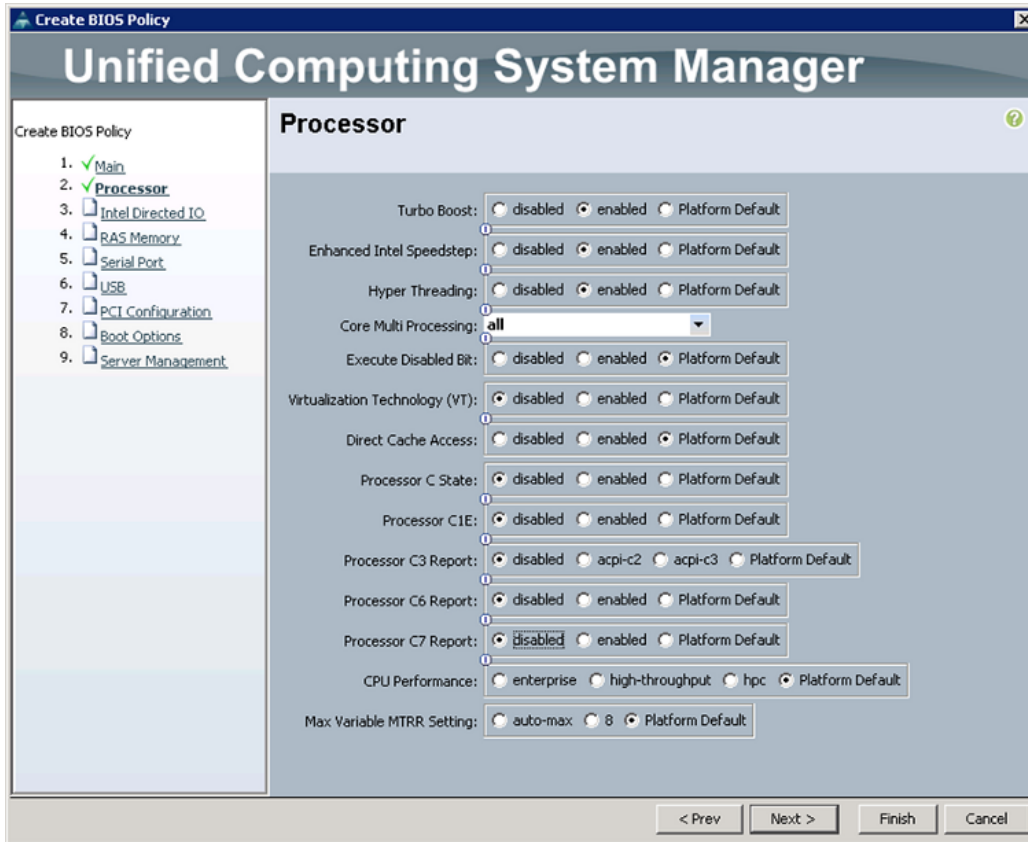
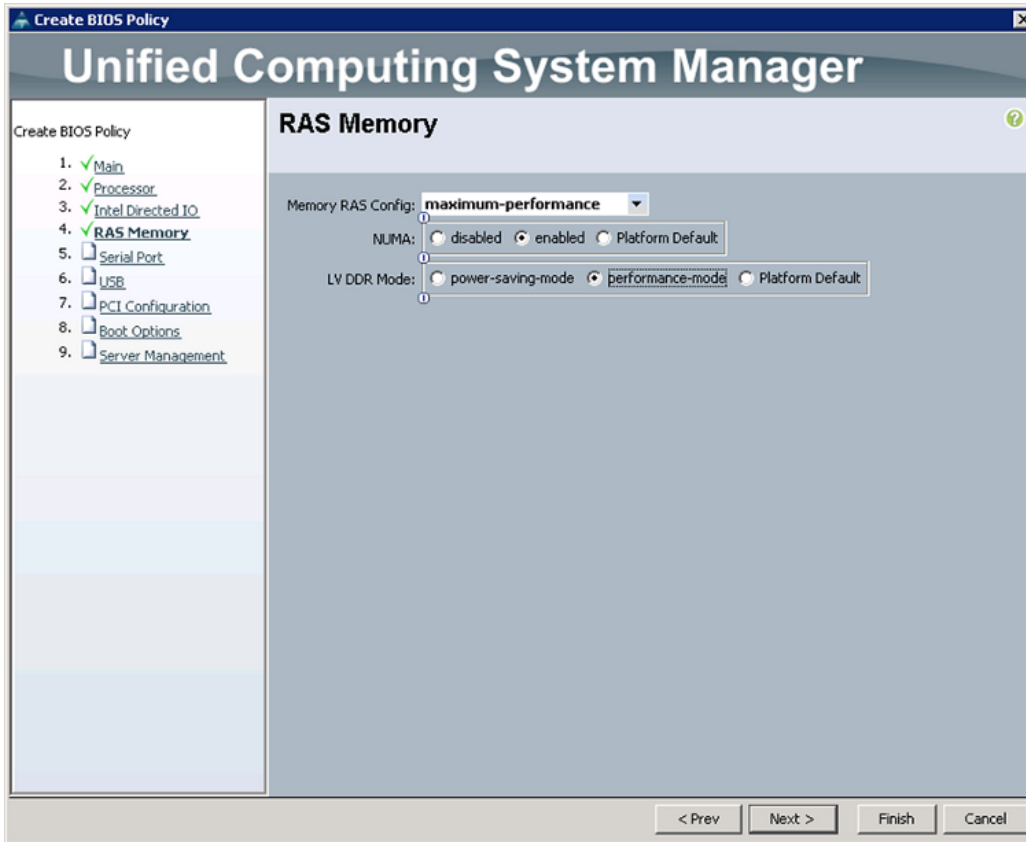


Figure 26 Intel Direct IO Settings



Figure 27 Memory Settings

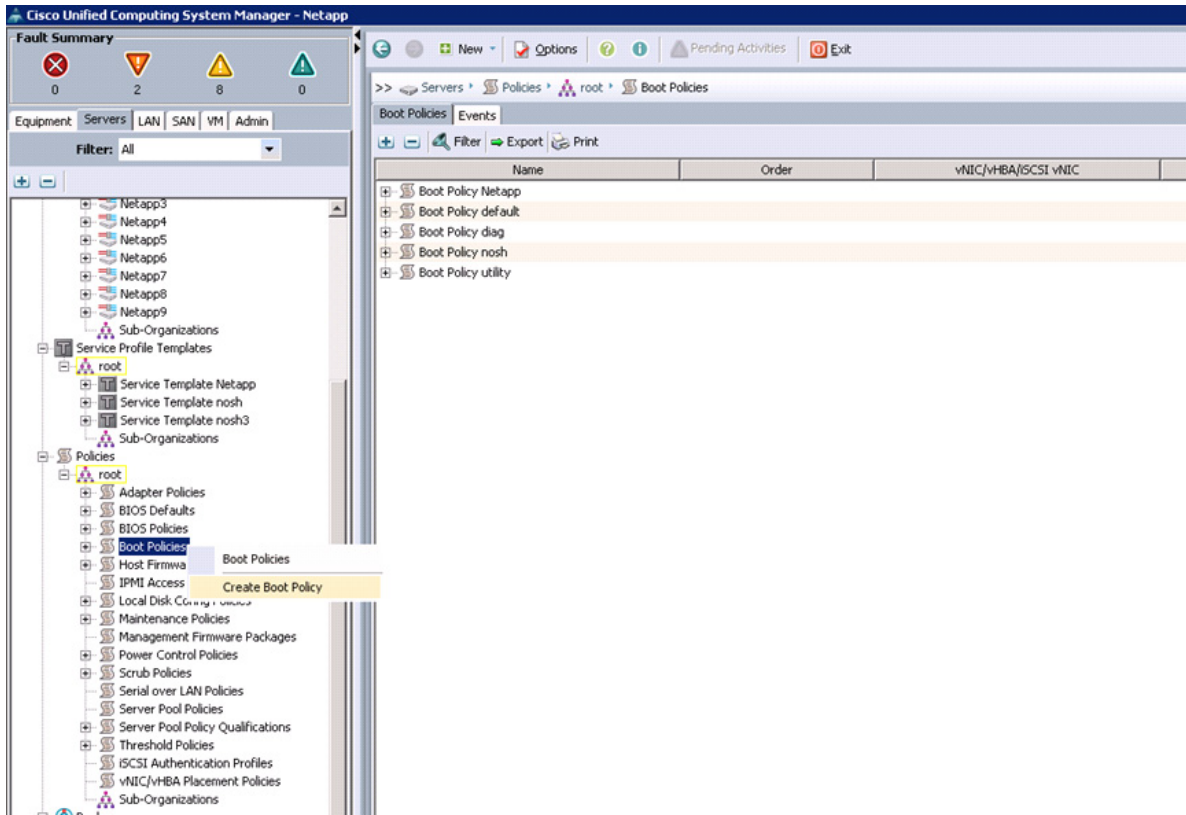


Creating Boot Policy

Follow these steps to create boot policies within the Cisco UCS Manager GUI:

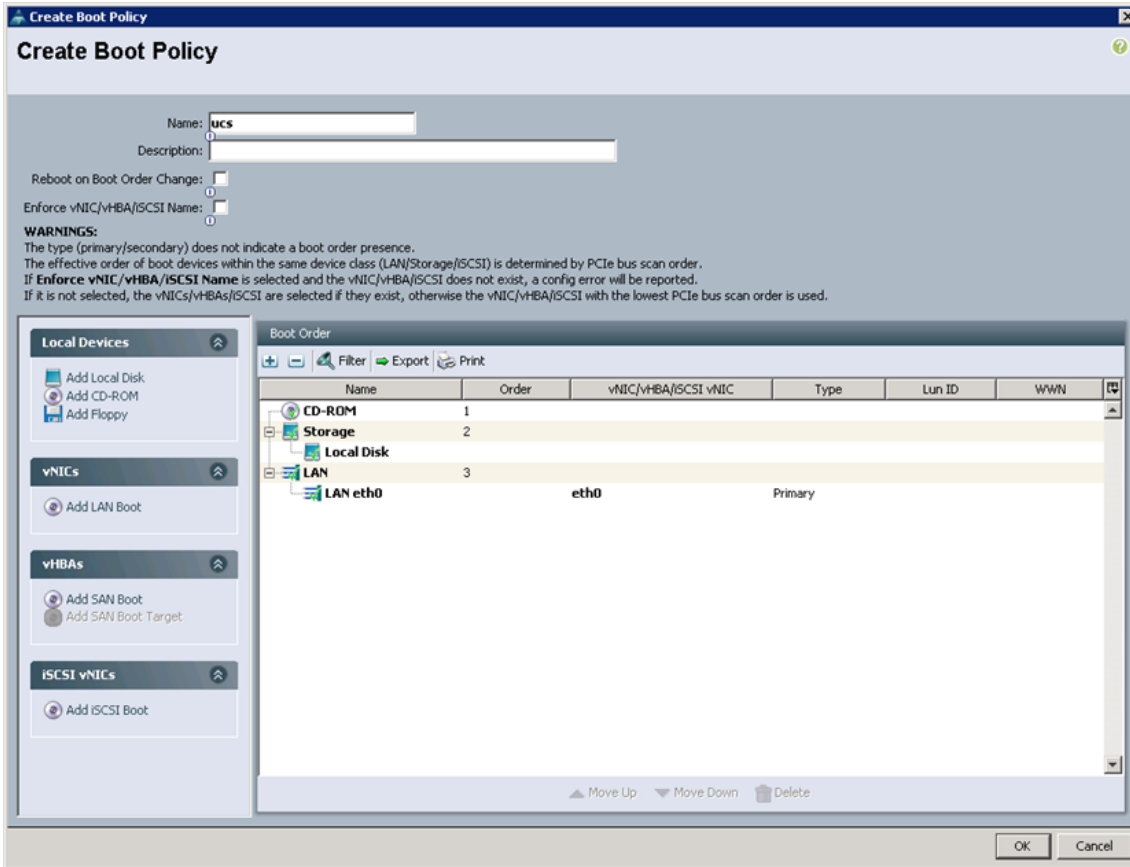
1. Select the Servers tab in the left pane in the UCSM GUI.
2. Select **Policies > root**.
3. Right-click the Boot Policies.
4. Select Create Boot Policy.

Figure 28 Creating Boot Policy



5. Enter ucs as the boot policy name.
6. (Optional) enter a description for the boot policy.
7. Keep the Reboot on Boot Order Change check box unchecked.
8. Expand Local Devices and select Add CD-ROM.
9. Expand Local Devices and select Add Local Disk.
10. Expand vNICs and select Add LAN Boot and enter eth0.
11. Click **OK** to add the Boot Policy.
12. Click **OK**.

Figure 29 Creating Boot Order




Creating Service Profile Template

To create a service profile template, follow these steps:

1. Select the Servers tab in the left pane in the UCSM GUI.
2. Select **Policies > root**.
3. Right-click root.
4. Select Create Service Profile Template.

Figure 30 Creating Service Profile Template

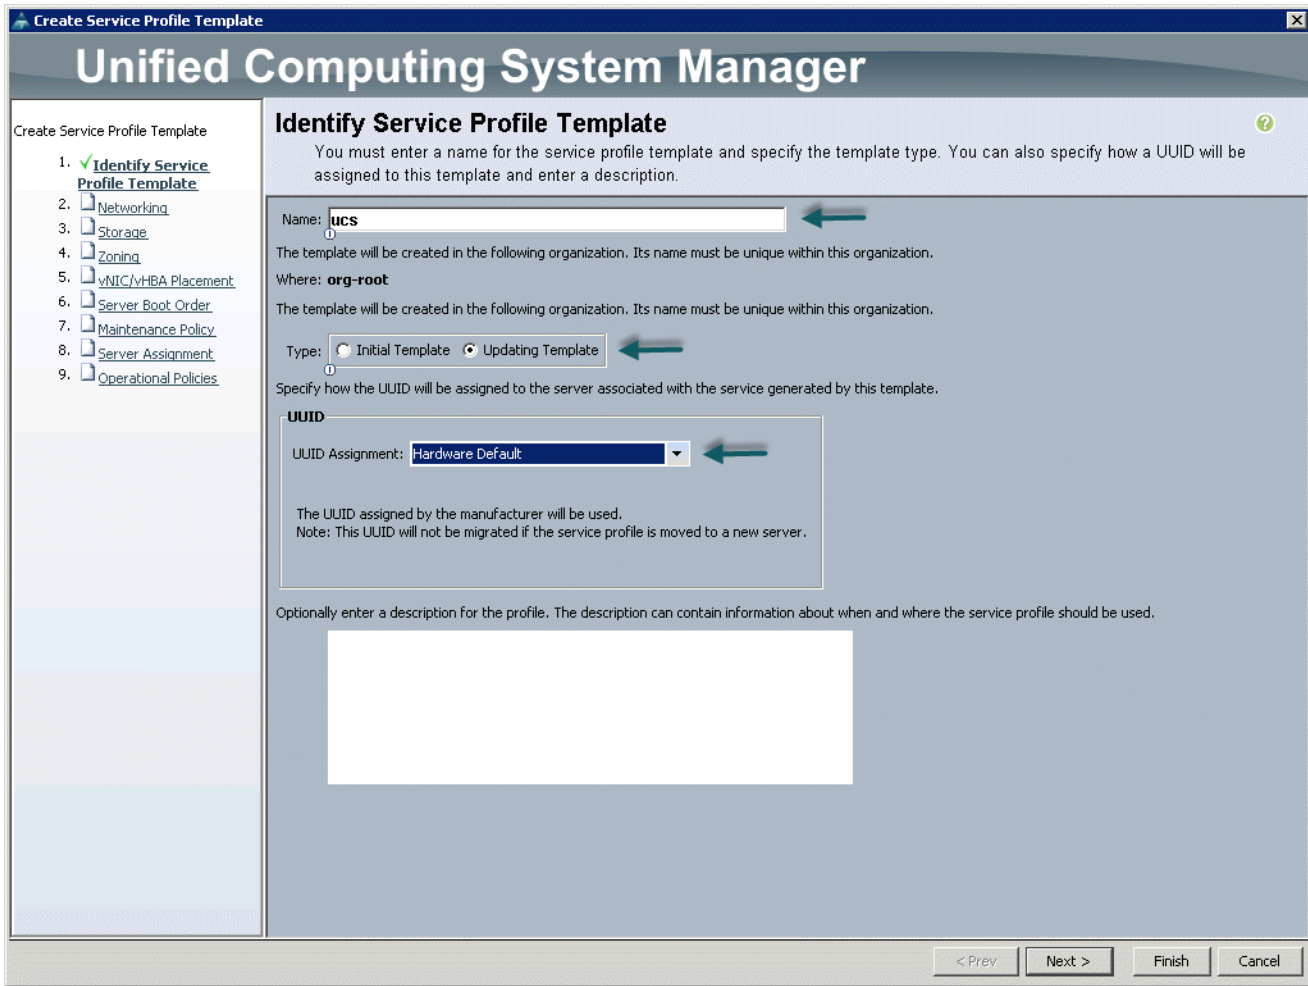


5. The Create Service Profile Template window appears.

The following steps provide the detailed configuration procedure used to create a service profile template:

- a. Name the service profile template as ucs. Select the **Updating Template** radio button.
- b. In the UUID section, select Hardware Default as the UUID pool.

Figure 31 Identifying Service Profile Template



c. Click **Next** to continue to the next section.

Configuring Network Settings for the Template

In the Networking window, follow these steps to create vNICs:

1. Keep the Dynamic vNIC Connection Policy field as default.
2. Select the **Expert** radio button for the option How would you like to configure LAN connectivity?
3. Click **Add** to add a vNIC to the template.

Figure 32 Adding vNICs

Unified Computing System Manager

Create Service Profile Template

- Identify Service Profile Template
- Networking**
- Storage
- Zoning
- vNIC/vHBA Placement
- Server Boot Order
- Maintenance Policy
- Server Assignment
- Operational Policies

Networking
Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy: Select a Policy to use (no Dynamic vNIC Policy by default) + Create Dynamic vNIC Connection Policy

How would you like to configure LAN connectivity? Simple **Expert** Use Connectivity Policy

Click **Add** to specify one or more vNICs that the server should use to connect to the LAN.

Name	MAC Address	Fabric ID	Native VLAN
Delete + Add			

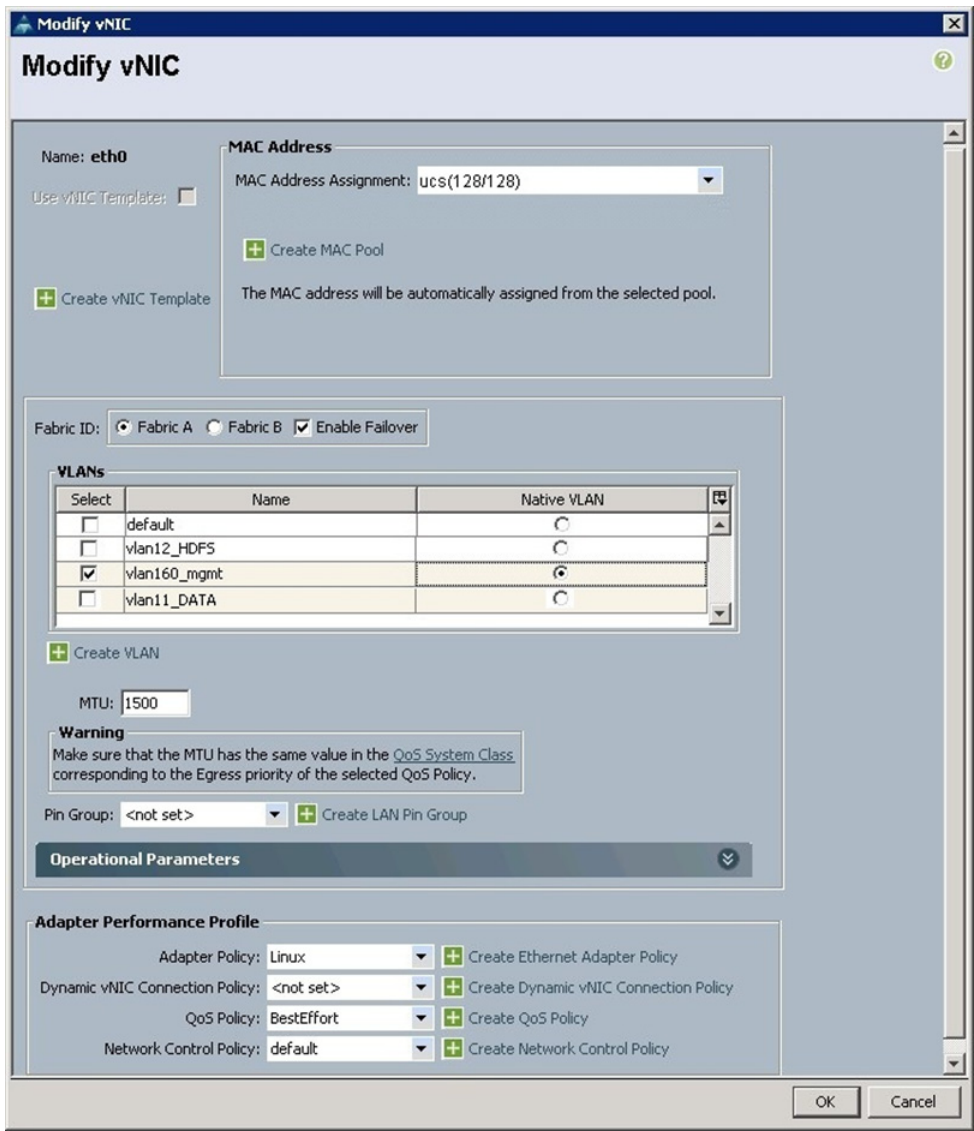
Click **Add** to specify one or more iSCSI vNICs that the server should use.

Name	Overlay vNIC Name	iSCSI Adapter Policy	MAC Address
+ Add Delete Modify			

< Prev Next > Finish Cancel

- The Create vNIC window displays. Name the vNIC as eth0.
- Select nosh in the Mac Address Assignment pool.
- Select the **Fabric A** radio button and check the Enable failover check box for the Fabric ID.
- Check the vlan160_mgmt check box for VLANs and select the **Native VLAN** radio button.
- Select MTU size as 1500.
- Select adapter policy as Linux.
- Keep the Dynamic vNIC connection policy as <no set>.
- Select QoS Policy as BestEffort.
- Keep the Network Control Policy as Default.
- Click **OK**.

Figure 33 Creating Management vNIC



14. Click **Add** to add another vNIC to the template.
15. The Create vNIC window appears. Name the vNIC as eth1.
16. Select ucs in the Mac Address Assignment pool.
17. Select the **Fabric B** radio button and check the Enable failover check box for the Fabric ID.
18. Check the vlan12_HDFS check box for VLANs and select the **Native VLAN** radio button for Native VLAN.
19. Select MTU size as 9000.
20. Select Adapter Policy as Linux.
21. Keep the Dynamic vNIC Connection Policy as <not set>.
22. Select QoS Policy to Platinum.
23. Keep the Network Control Policy as Default.

24. Click **OK**.

Figure 34 Configuring vNIC eth1

Modify vNIC

Name: eth1

Use vNIC Template:

MAC Address

MAC Address Assignment: ucs(128/128)

+ Create MAC Pool

The MAC address will be automatically assigned from the selected pool.

+ Create vNIC Template

Fabric ID: Fabric A Fabric B Enable Failover

Select	Name	Native VLAN
<input type="checkbox"/>	default	<input type="radio"/>
<input checked="" type="checkbox"/>	vlan12_HDFS	<input checked="" type="radio"/>
<input type="checkbox"/>	vlan160_mgmt	<input type="radio"/>
<input type="checkbox"/>	vlan11_DATA	<input type="radio"/>

+ Create VLAN

MTU: 9000

Warning
Make sure that the MTU has the same value in the QoS System Class corresponding to the Egress priority of the selected QoS Policy.

Pin Group: <not set> + Create LAN Pin Group

Operational Parameters

Adapter Performance Profile

Adapter Policy: Linux + Create Ethernet Adapter Policy

Dynamic vNIC Connection Policy: <not set> + Create Dynamic vNIC Connection Policy

QoS Policy: Platinum + Create QoS Policy

Network Control Policy: default + Create Network Control Policy

OK Cancel

25. Click **Add** to add another vNIC to the template.

26. The Create vNIC window appears. Name the vNIC as eth2.

27. Select ucs in the Mac Address Assignment pool.

28. Select the **Fabric A** radio button and check the Enable failover check box for the Fabric ID.

29. Check the vlan11_DATA check box for VLANs and select the **Native VLAN** radio button.

30. Select MTU size as 9000.

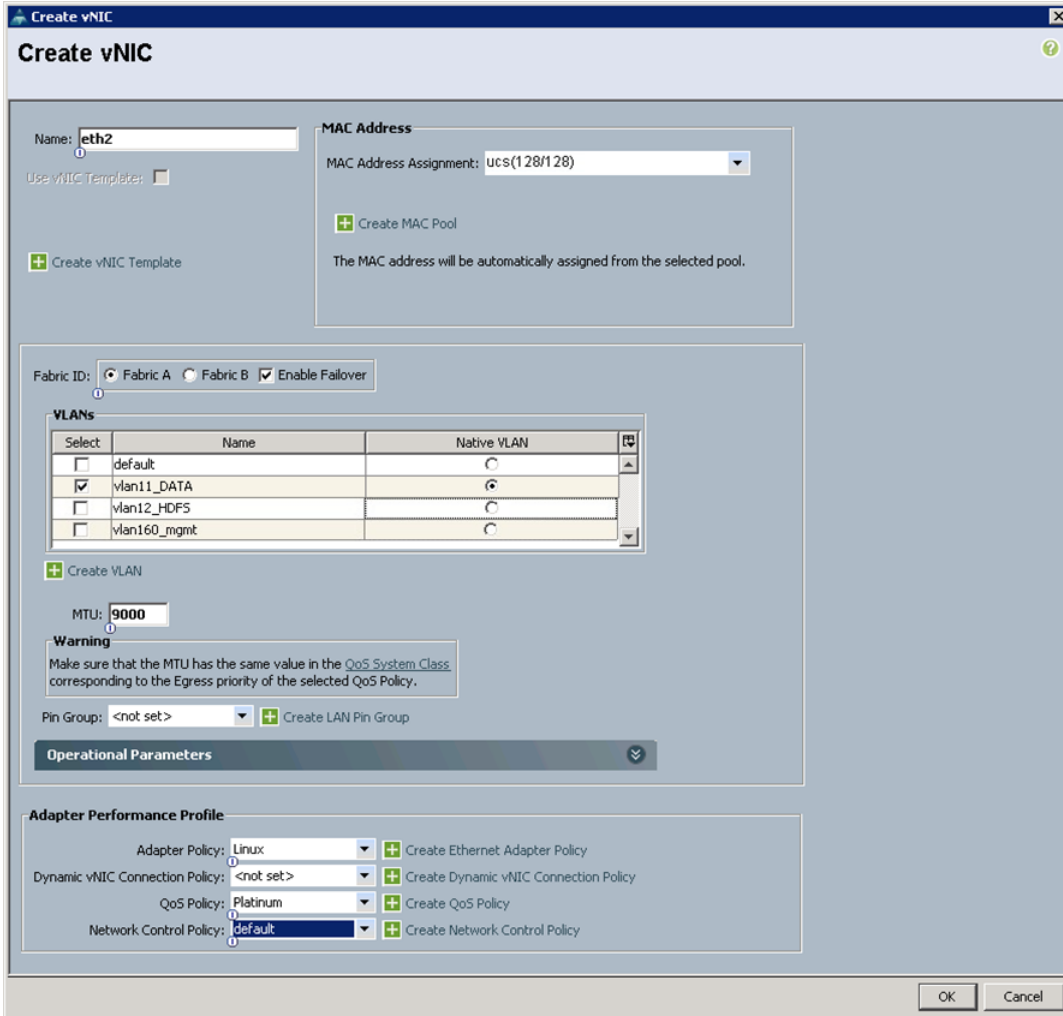
31. Select adapter policy as Linux.

32. Keep the Dynamic vNIC Connection Policy as <no set>.

33. Select QoS Policy as Platinum.

34. Keep the Network Control Policy as Default.
35. Click **OK**.

Figure 35 **Configuring vNIC eth2**



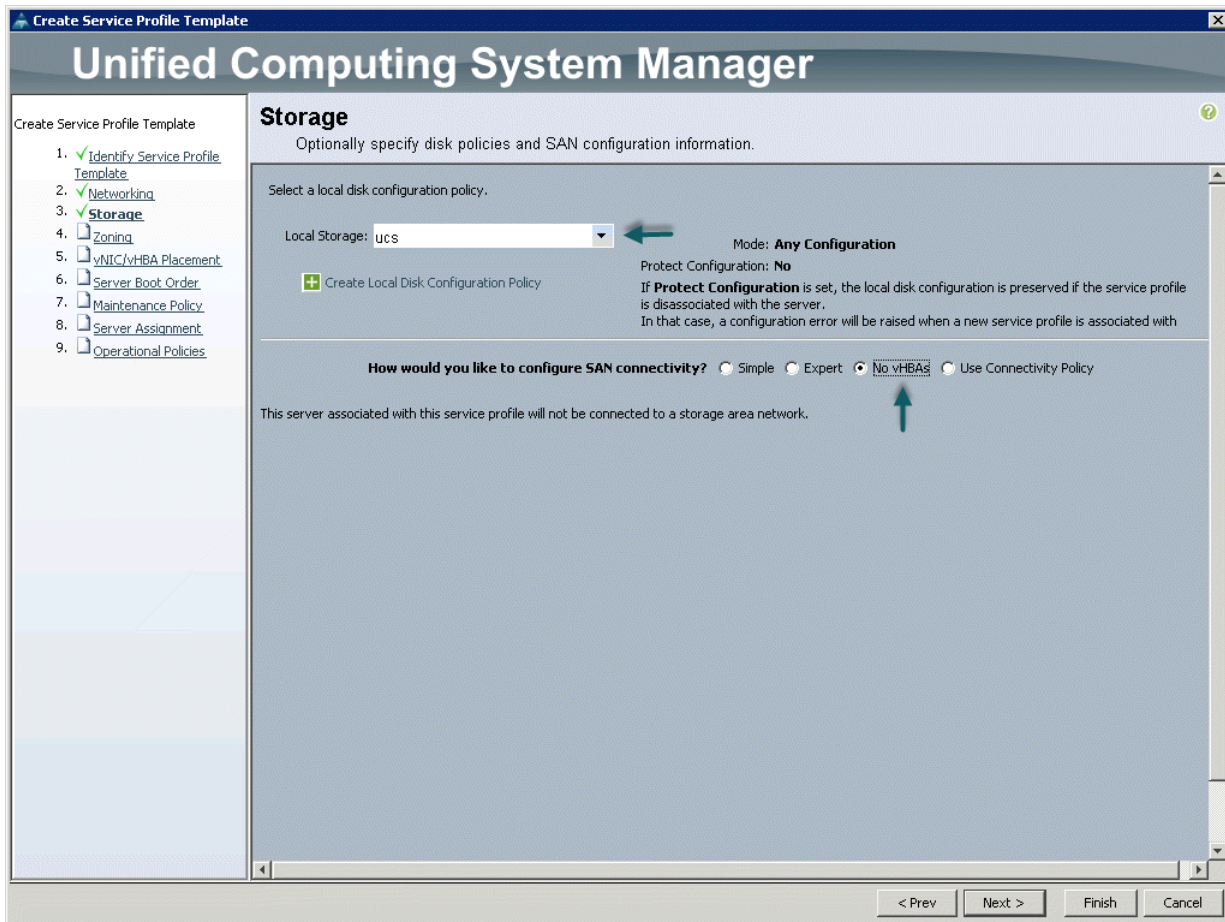
36. Click **Next** to continue to the next section.

Configuring Storage Policy for the Template

In the Storage window, follow these steps to configure storage:

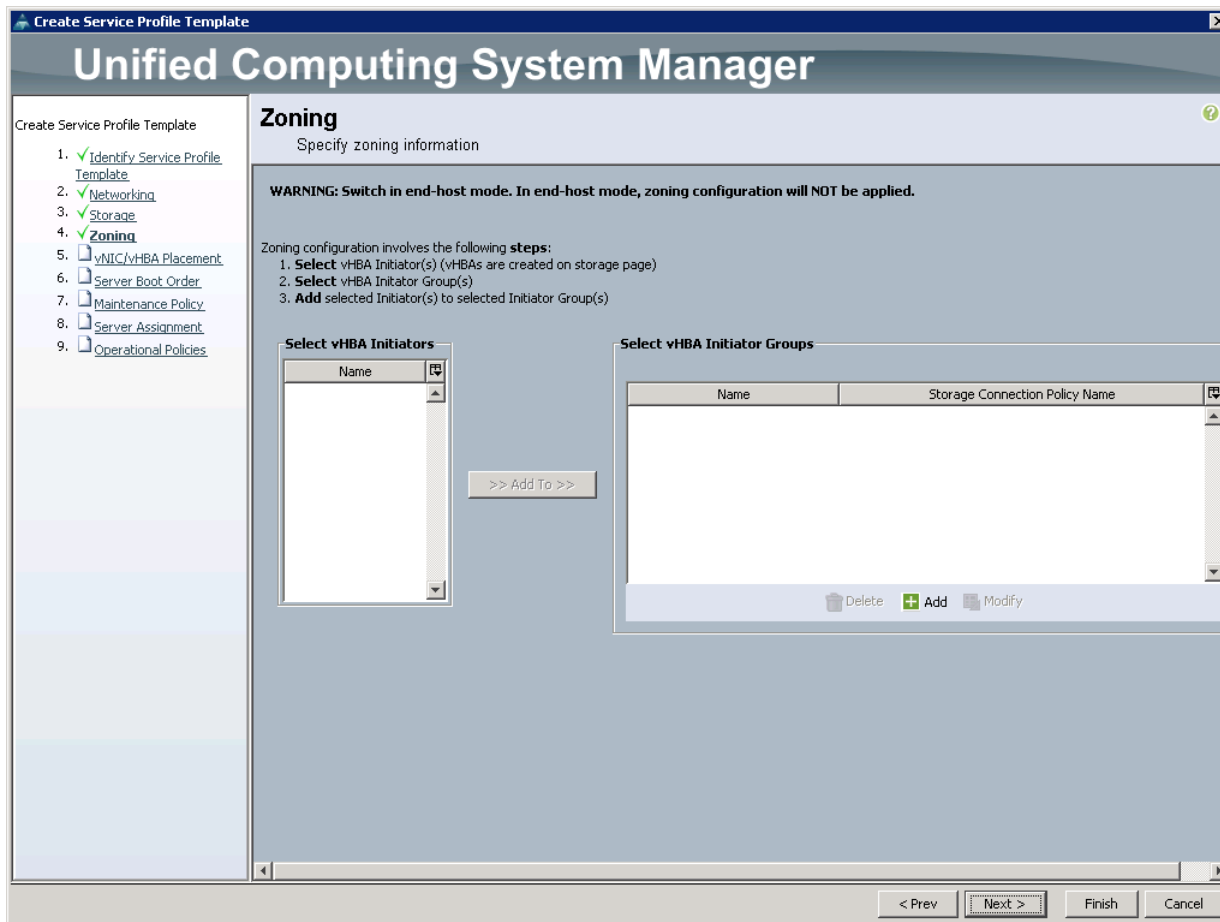
1. Select ucs for the local disk configuration policy.
2. Select the **No vHBAs** radio button for the option How would you like to configure SAN connectivity?

Figure 36 Storage Settings



3. Click **Next** to continue to the next section.
4. Click **Next** in the Zoning Window to go to the next section.

Figure 37 Configure Zoning




Configuring vNIC/vHBA Placement for the Template

In the vNIC/vHBA Placement Section, follow these steps to configure placement policy:

1. Select the Default Placement Policy option for Select Placement field.
2. Select eth0, eth1, and eth2 assign the vNICs in the following order:
 - eth0
 - eth1
 - eth2
3. Review the table to make sure that all of the vNICs were assigned in the appropriate order.

Figure 38 Creating vNIC and vHBA Policy



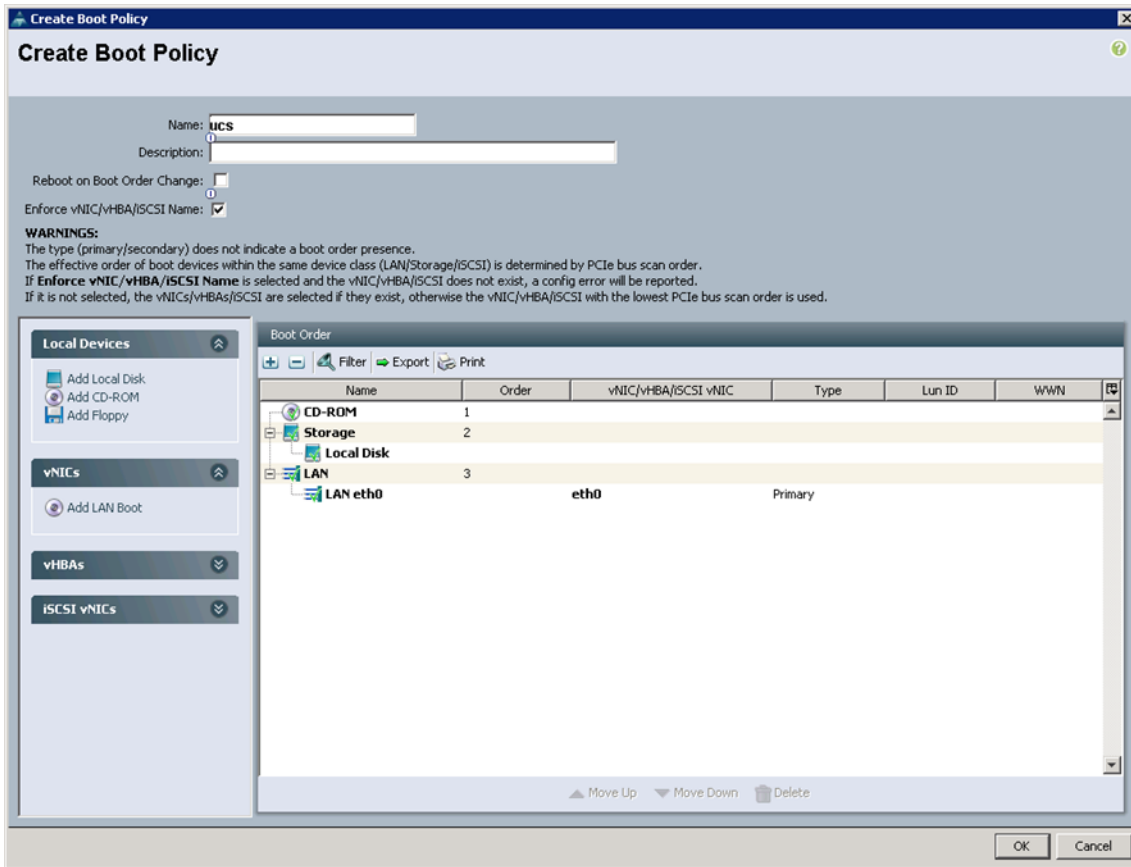
4. Click **Next** to continue to the next section.

Configuring Server Boot Order for the Template

In the Server Boot Order Section, follow these steps to set the boot order for servers:

1. Select **ucs** for the Boot Policy Name field.
2. Check the **Reboot on Boot Order Change** check box.
3. Check the **Enforce vNIC/vHBA/iSCSI Name** check box.
4. Review the table to make sure that all of the boot devices were created and identified. Verify that the boot devices are in the correct boot sequence.
5. Click **OK**.

Figure 39 Creating Boot Policy



6. Click Next to continue to the next section.

Configuring Maintenance Policy for the Template

In the Maintenance Policy window, follow these steps to apply maintenance policy:

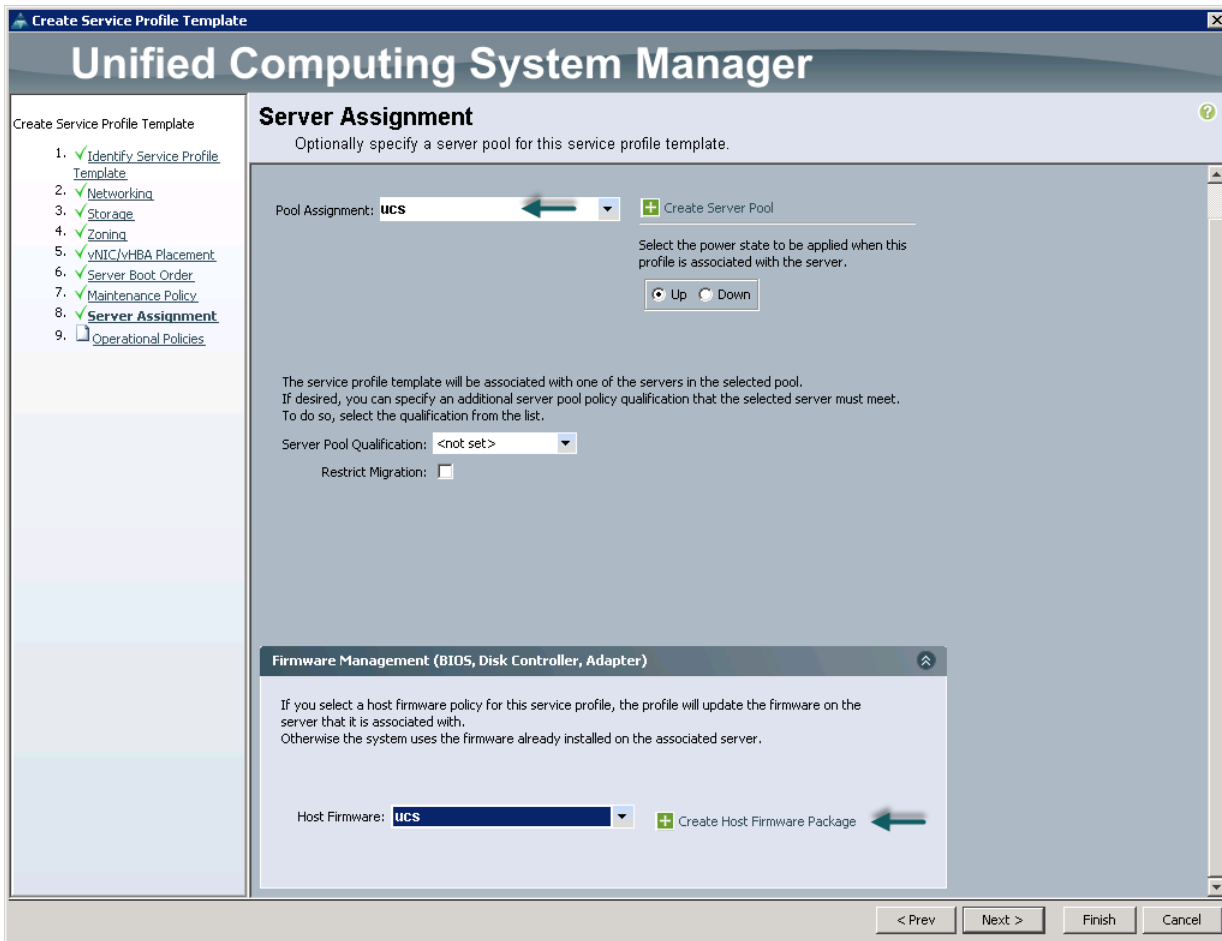
1. Keep the Maintenance Policy at no policy used by default.
2. Click Next to continue to the next section.

Configuring Server Assignment for the Template

In the Server Assignment window, follow these steps to assign servers to the pool:

1. Select ucs for the Pool Assignment field.
2. Keep the Server Pool Qualification field at default.
3. Select nosh for the Host Firmware Package.

Figure 40 Server Assignment



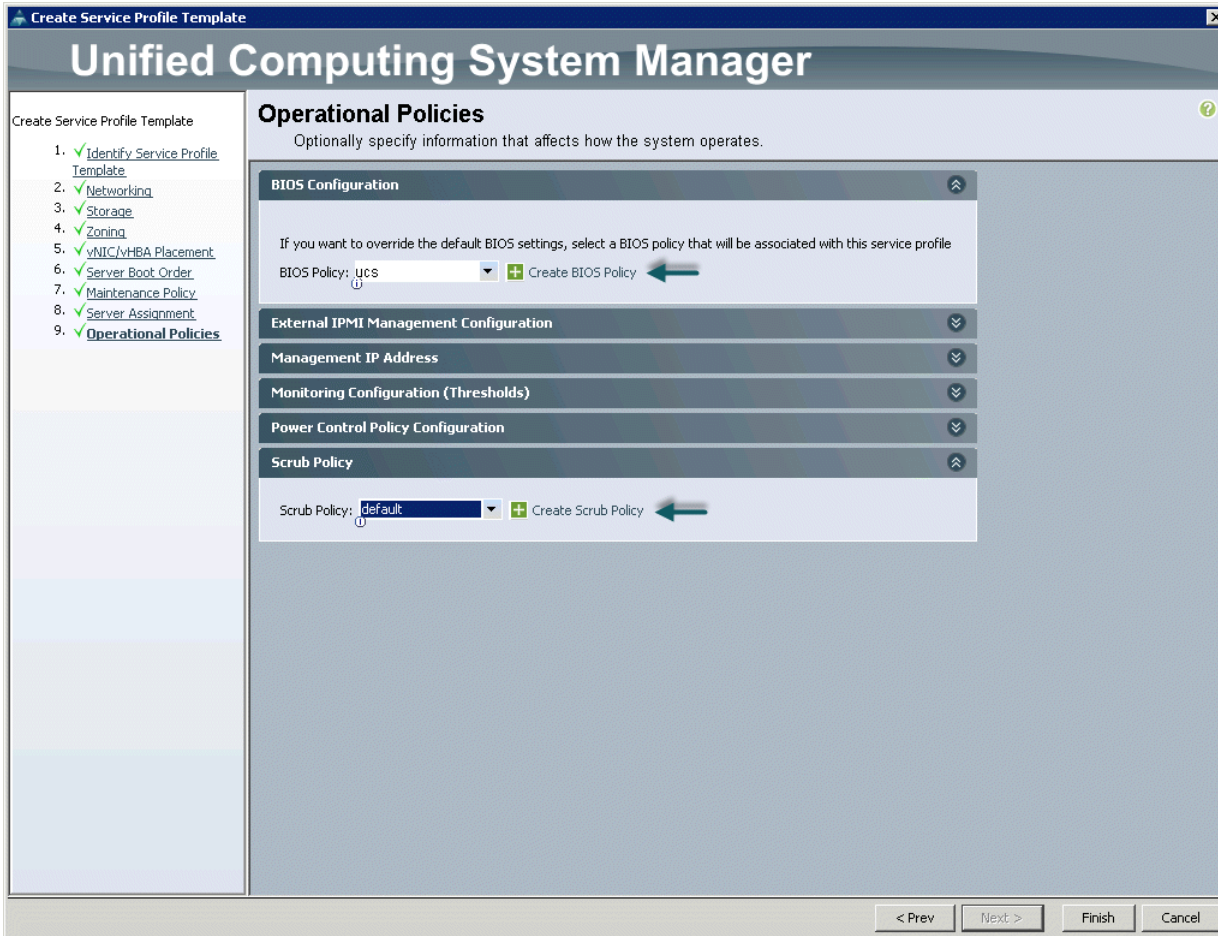
4. Click **Next** to continue to the next section.

Configuring Operational Policies for the Template

In the Operational Policies window, follow these steps:

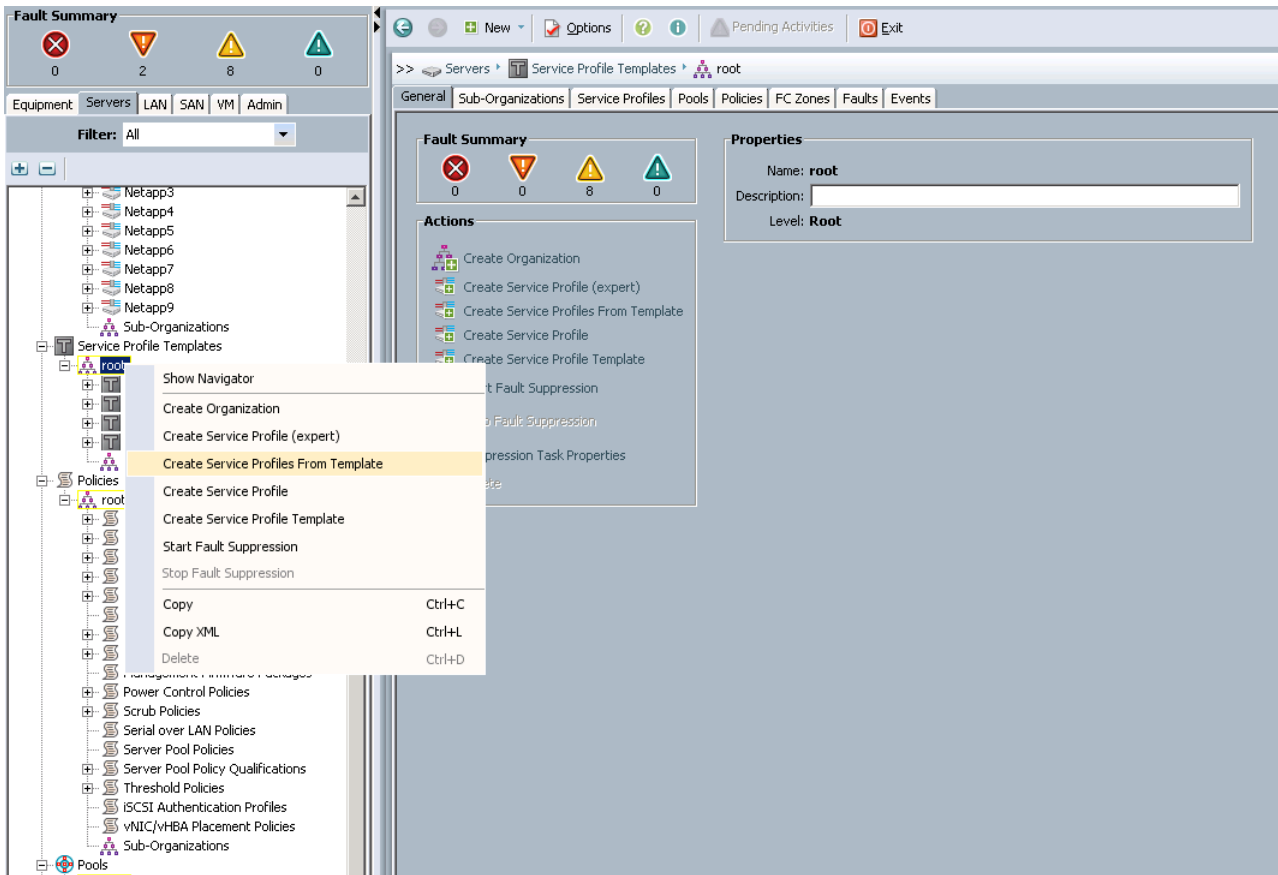
1. Select **ucs** in the BIOS Policy field.

Figure 41 Creating Operational Policies



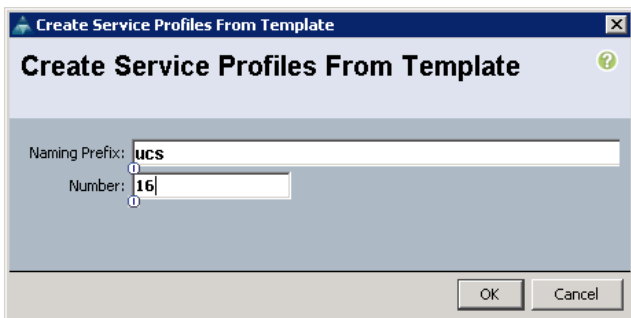
2. Click **Finish** to create the Service Profile template.
 3. Click **OK** in the pop-up window to exit the wizard.
- Select the Servers tab in the left pane in the UCSM GUI.
1. Select **Service Profile Templates > root**.
 2. Right-click the root.
 3. Select Create Service Profile Template.

Figure 42 *Creating Service Profile*



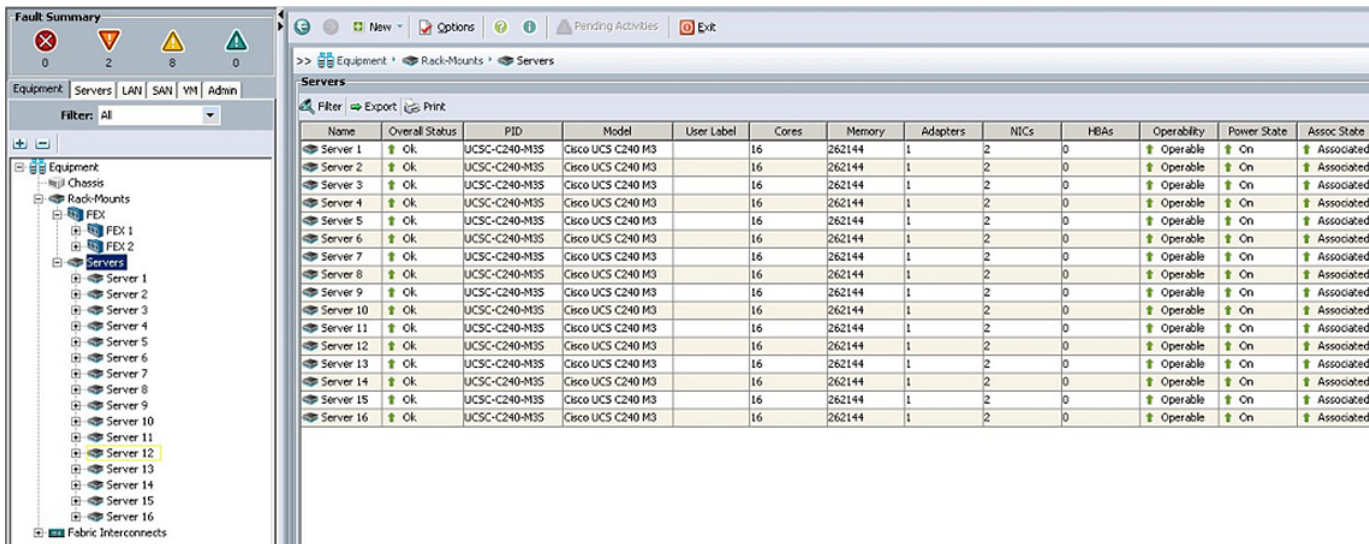
4. The Create Service Profile from Template window appears.

Figure 43 *Creating Service Profile from Template*



5. Now connect the power cable to the servers.
6. Servers will be discovered by UCS Manager.
7. Association of Service Profile will take place automatically.
8. The final Cisco UCS Manager window is shown in [Figure 44](#).

Figure 44 UCS Manager Showing Sixteen Nodes



Configuring Disk Drives for OS

As mentioned above, the focus of this CVD is the High Performance Configuration featuring 24 1TB SFF disk drives. The disk drives are configured as individual RAID0 volumes with 1MB strip size. Read ahead cache is enabled and write cache is enabled while battery is present. The first disk drive is used for operating system and remaining 23 disk drives are using for HDFS as described in the following sections.



Note

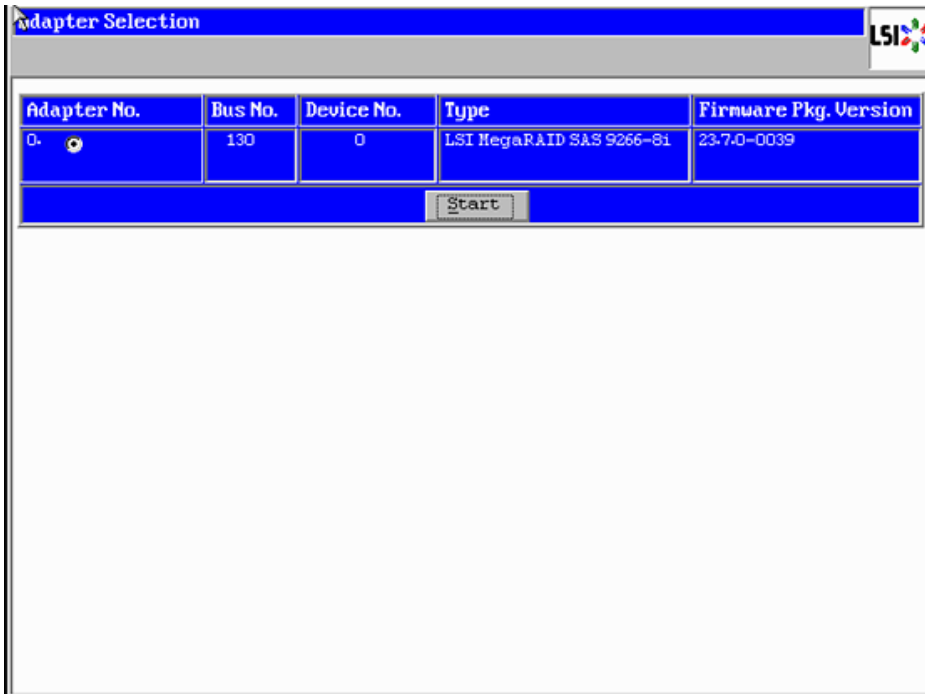
In the case of High Capacity Configuration featuring 12 3TB LFF disk drives, the disk drives are configured as individual RAID0 volumes with 1MB strip size. Read ahead cached is enabled and write cache is enabled while battery is present. Two partitions of 1TB and 2TB are created on the first disk drive, the 1TB partition is used for operating system and the 2TB partition is used for HDFS along with disk drives 2 through 12.

There are several ways to configure RAID: using LSI WebBIOS Configuration Utility embedded in the MegaRAID BIOS, booting DOS and running MegaCLI commands, using Linux based MegaCLI commands, or using third party tools that have MegaCLI integrated. For this deployment, the first disk drive is configured using LSI WebBIOS Configuration Utility and rest is configured using Linux based MegaCLI commands after the OS is installed.

Follow these steps to create RAID0 on the first disk drive to install the operating system:

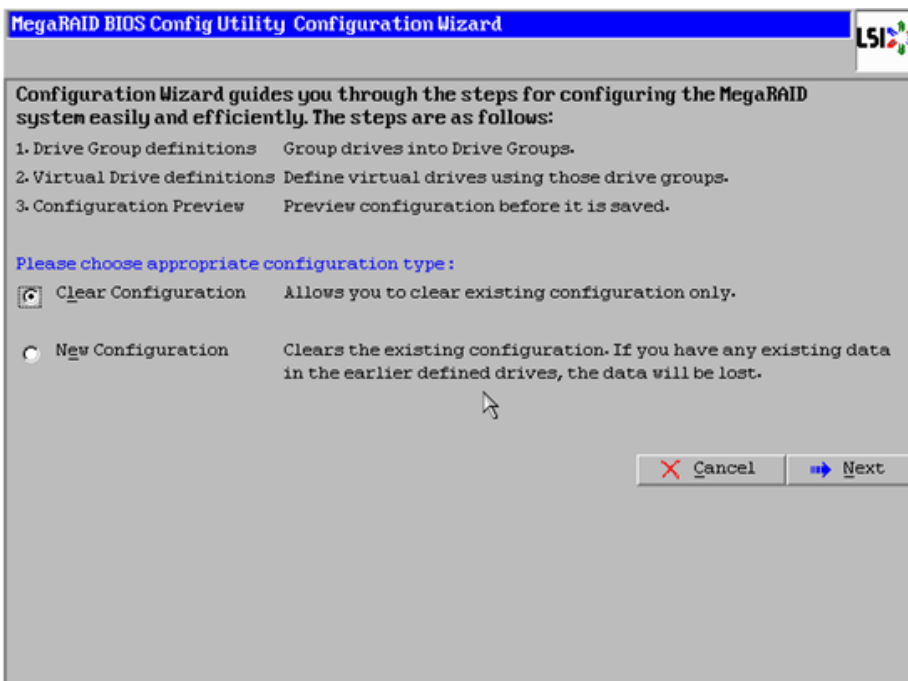
1. Once the server has booted and the MegaRAID Controller has been detected, the following will appear on the screen:
 - Press <Ctrl><H> for WebBIOS.
 - Press Ctrl+H immediately.
 - The Adapter Selection window appears.
2. Click **Start** to continue.

Figure 45 RAID Configuration for LSI MegaRAID SAS Controllers



3. Click **Configuration Wizard**.
4. In the configure wizard window, select the configuration type as **Clear Configuration** and click **Next** to clear the existing configuration.

Figure 46 Clearing Existing Configuration



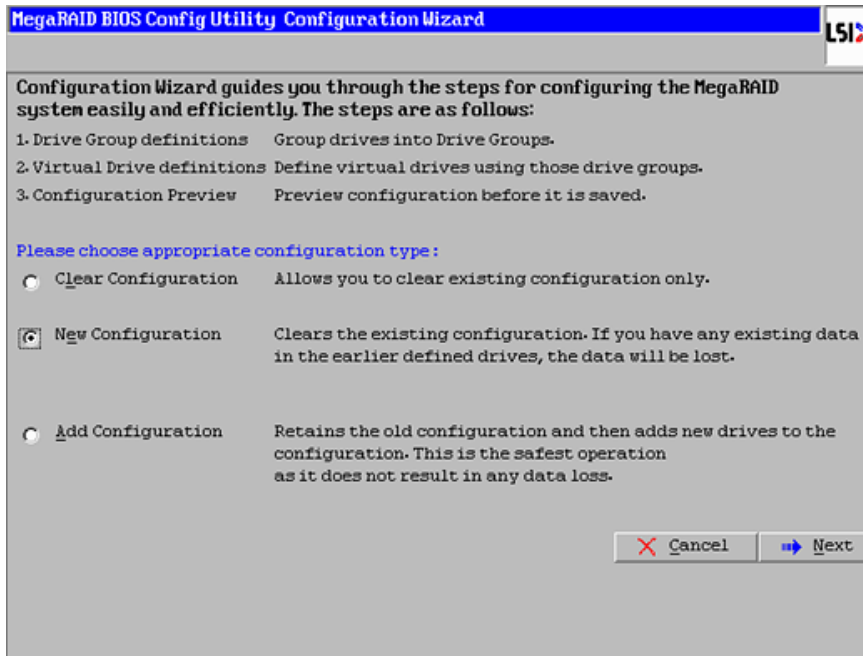
5. Click **Yes** when asked to confirm clear configuration.
6. In the Physical View, make sure all the drives are Unconfigured Good.
7. Click **Configuration Wizard**.

Figure 47 *Confirming Clearance of the previous Configuration on the Controller*



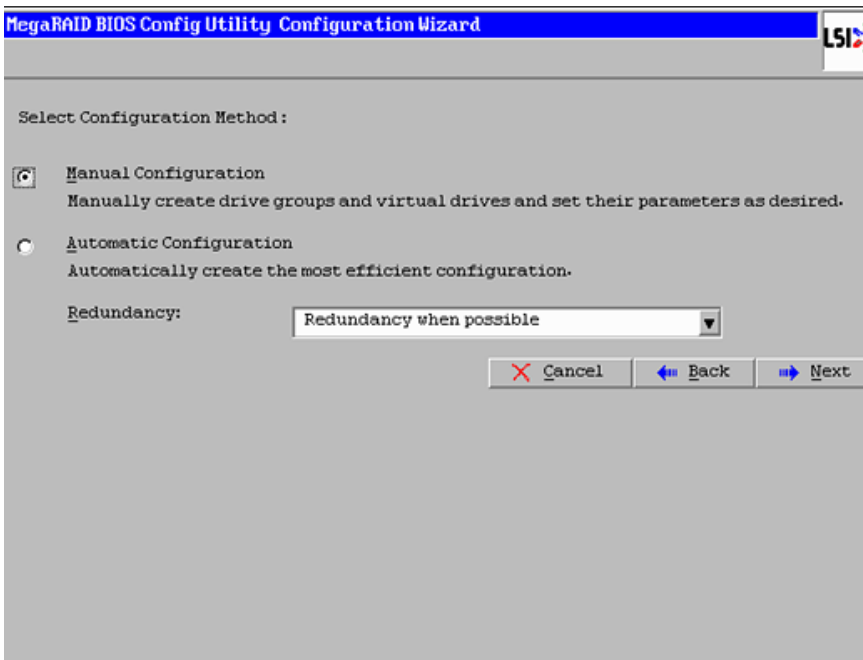
8. In the configure wizard window, select the configuration type as New Configuration and click **Next**.

Figure 48 **Selecting a New Configuration**



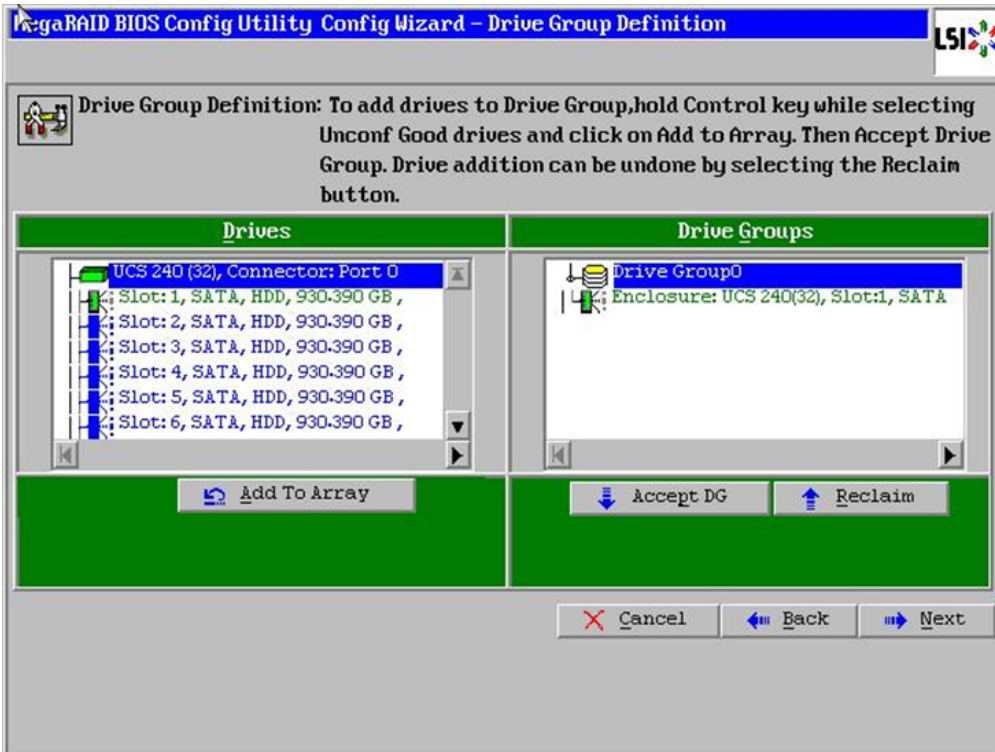
9. Select the configuration method to be Manual Configuration to have control over all attributes of the new storage configuration such as drive groups, virtual drives, and to set their parameters.
10. Click Next.

Figure 49 **Selecting Manual Configuration**



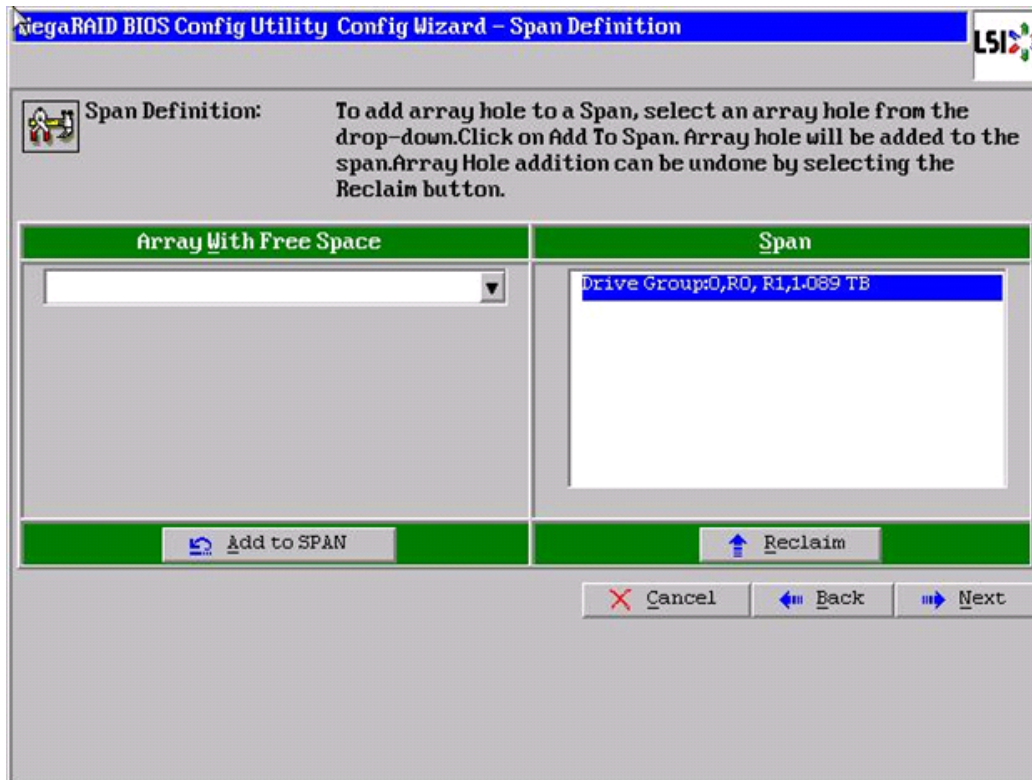
11. The Drive Group Definition window appears. In this window select the two drives to create drive groups.
12. Click **Add to Array** to move the drives to a proposed drive group configuration in the Drive Groups pane. Click **Accept DG** and then, click **Next**.

Figure 50 Moving Drives to Drive Groups



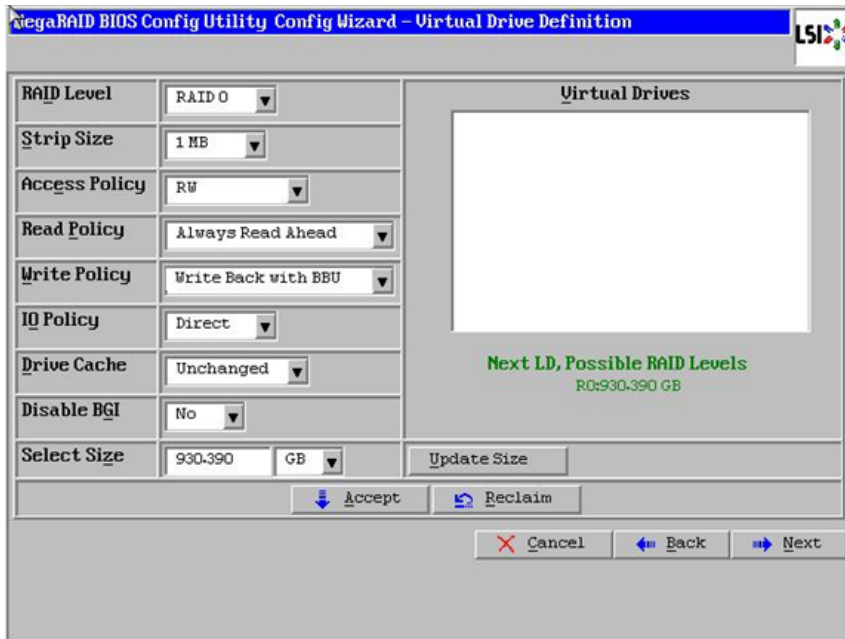
13. In the Span definitions Window, Click **Add to SPAN** and then, click **Next**.

Figure 51 Adding Arrayhole to Span



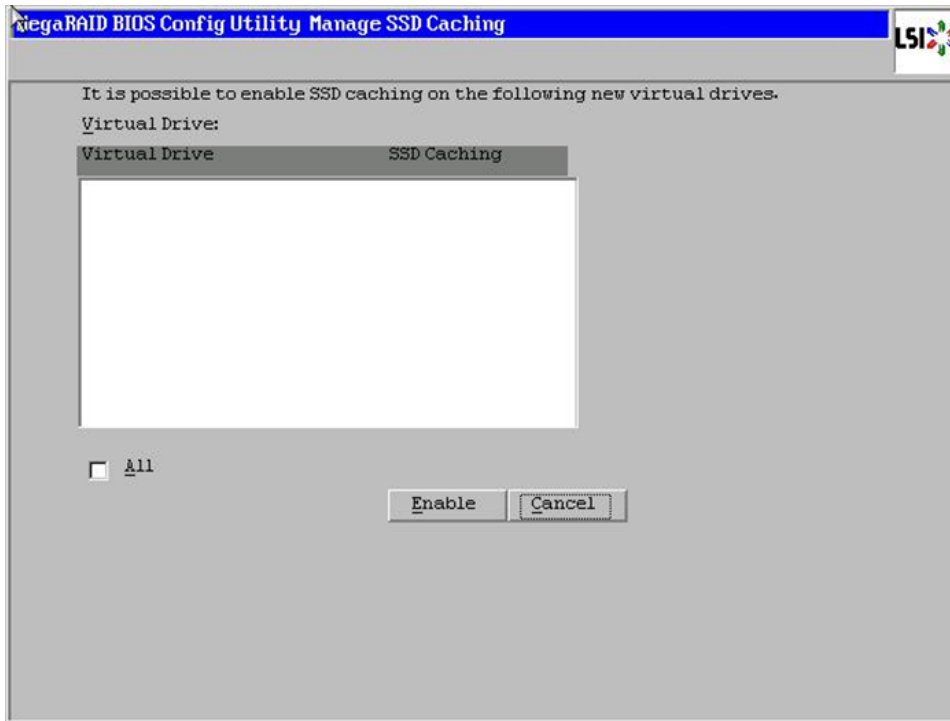
14. In Virtual Drive definitions window, follow these steps to configure read normal and write through modes:
 - a. Click **UpdateSize**.
 - a. Change Strip Size to 1MB. A larger strip size produces higher read performance.
 - b. From the read Policy drop down list, choose Always Read Ahead.
 - c. From the Write Policy drop down list, choose Write Back with BBU.
 - d. Make Sure RAID Level is set to RAID0.
 - e. Click **Accept** to accept the changes to the virtual drive definitions.
 - f. Click **Next**.

Figure 52 Defining Virtual Drive



15. After you finish the virtual drive definitions, click **Next**. The Configuration Preview window appears showing VD0.
16. Review the virtual drive configuration in the Configuration Preview window and click **Accept** to save the configuration.
17. Click **Yes** to save the configuration.
18. In the managing SSD Caching Window, Click **Cancel**.

Figure 53 *SSD Caching on the Created Virtual Drive*



19. Click **Yes**. When asked to confirm to initialize.


Figure 54 *Confirmation to Initialize*



20. Set VD0 as the Boot Drive and click **Go**.
21. Click **Home**.

22. Review the Configuration and Click **Exit**.

Figure 55 *Setting the Virtual Drive as Boot Drive*



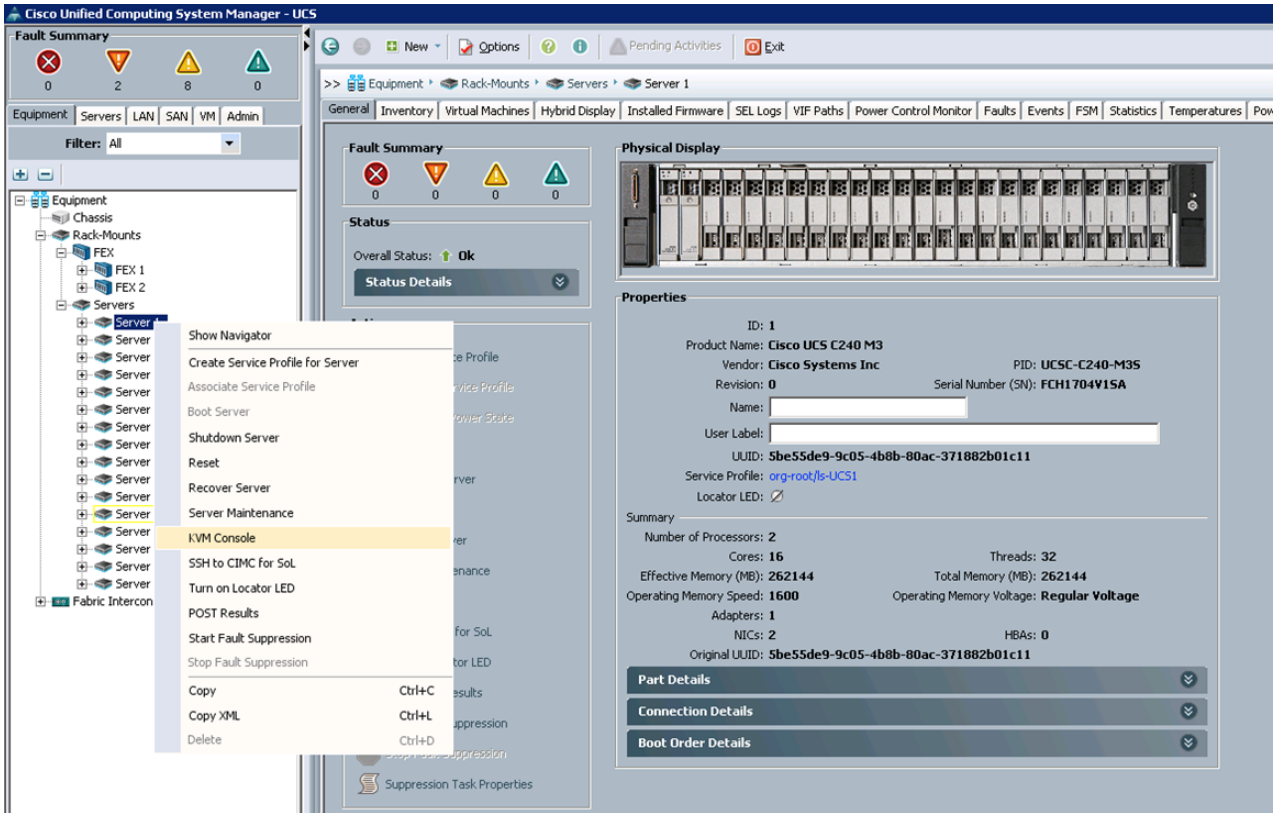
Installing Red Hat Enterprise Linux Server 6.2 using KVM

There are multiple methods to install Red Hat Linux operating system. The installation procedure described in this deployment guide uses KVM console and virtual media from Cisco UCS Manager.

To open the KVM console, follow these steps:

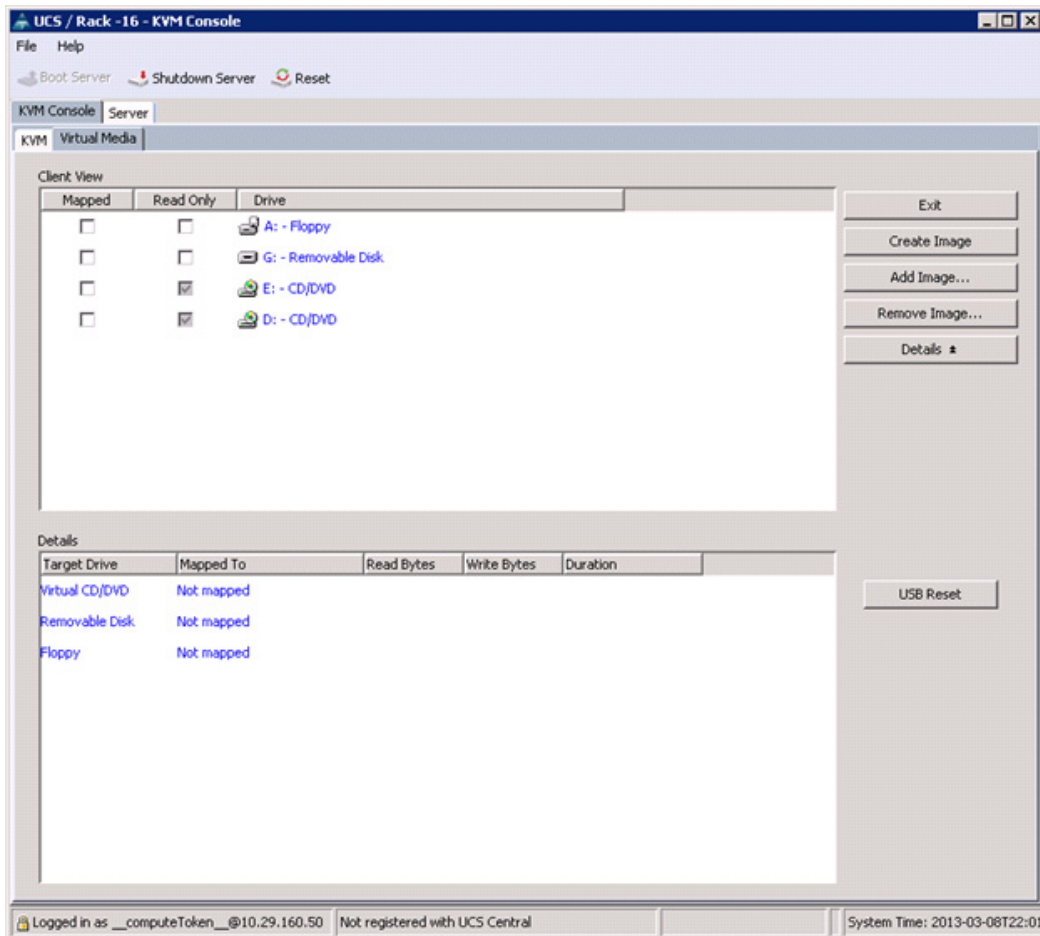
1. Log in to the Cisco UCS 6296 Fabric Interconnect and launch the Cisco UCS Manager application.
2. Select Equipment tab.
3. In the navigation pane expand Rack-mount and then Servers.
4. Right-click on the server and select KVM Console.

Figure 56 Launching KVM Console



5. In the KVM window, select the Virtual Media tab.

Figure 57 Adding ISO Image




6. Click **Add Image** in the window that appeared.
7. Browse to the Red Hat Enterprise Linux Server 6.2 installer ISO image file.



Note The Red Hat Enterprise Linux 6.2 DVD is assumed to be on the client machine.

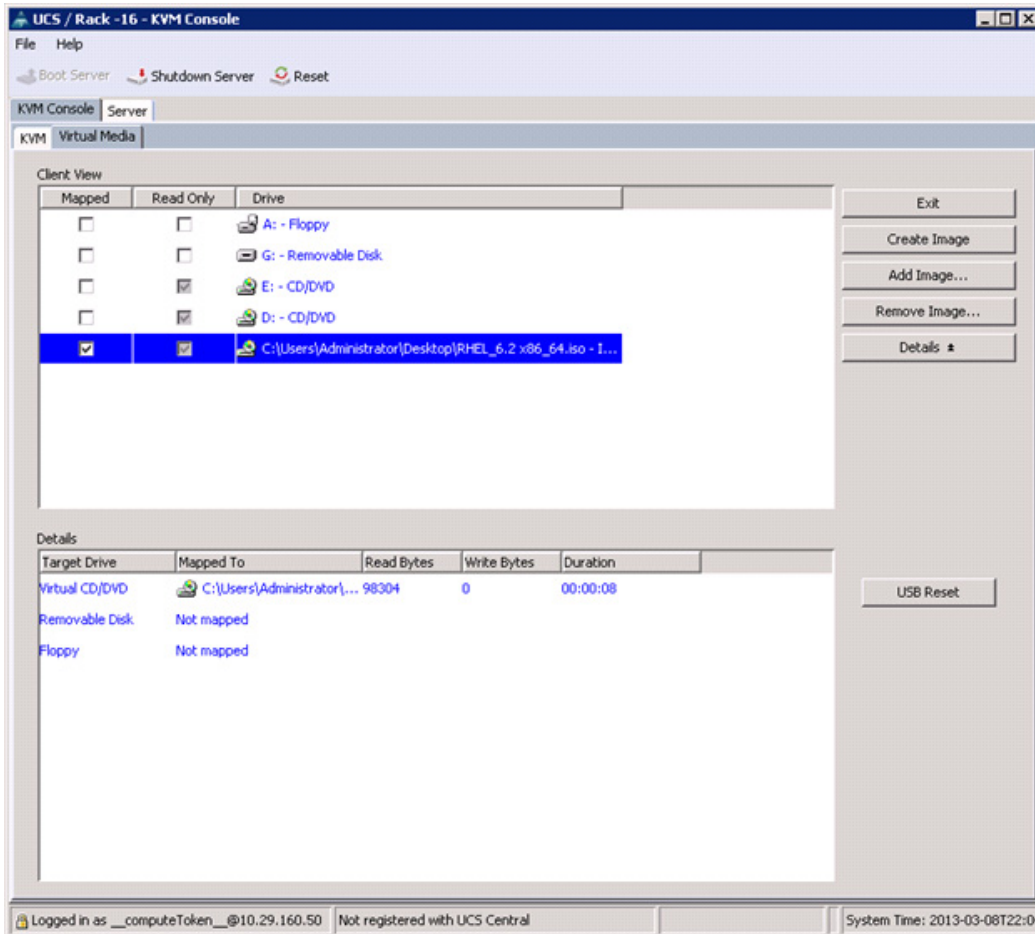
8. Click **Open** to add the image to the list of virtual media.

Figure 58 **Selecting the Red Hat Enterprise Linux ISO Image**



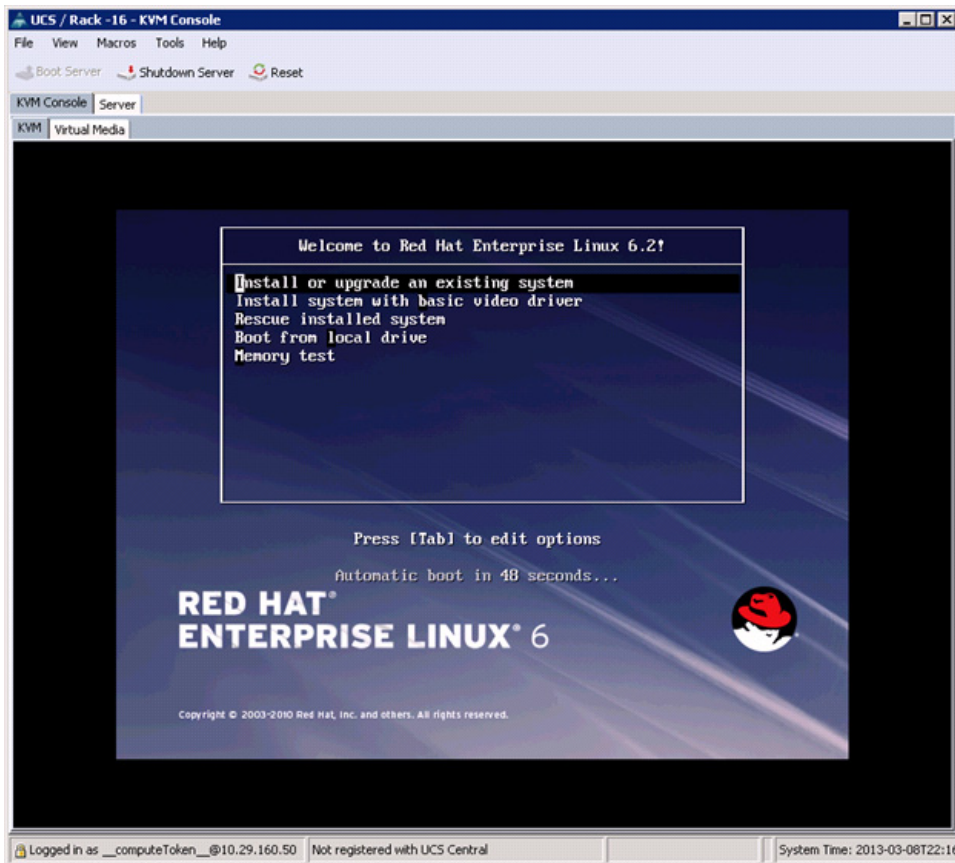
9. Check the check box for Mapped, next to the entry corresponding to the image you just added.

Figure 59 Mapping the ISO Image



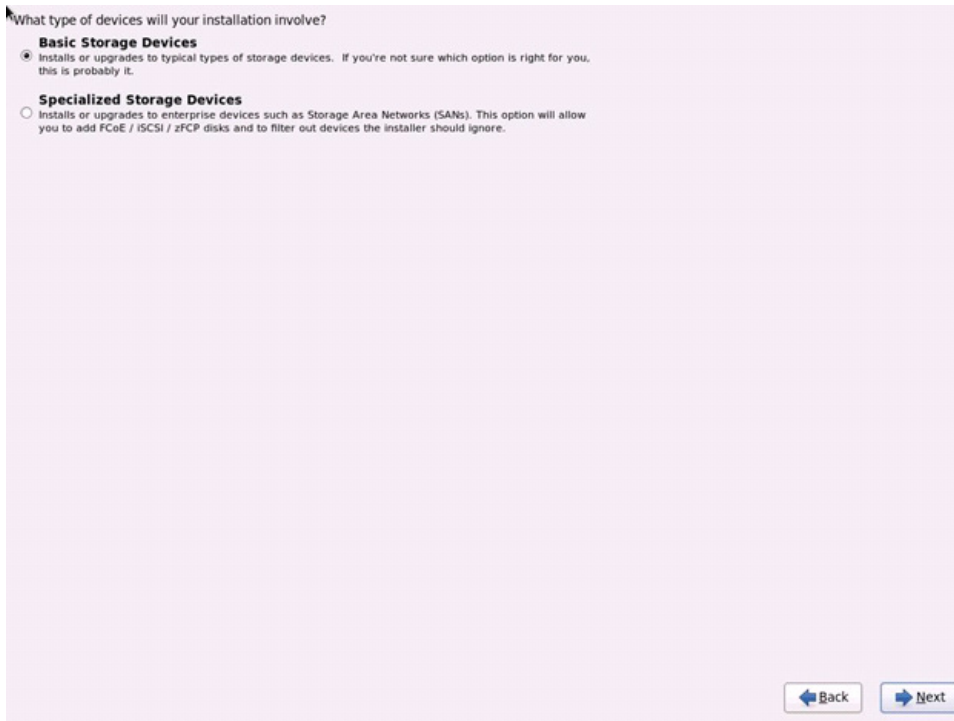
10. In the KVM window, select the KVM tab to monitor during boot.
11. In the KVM window, select the Boot Server button in the upper left corner.
12. Click **OK**.
13. Click **OK** to reboot the system.
14. On reboot, the machine detects the presence of the Red Hat Enterprise Linux Server 6.2 install media.
15. Select the Install or Upgrade an Existing System option.

Figure 60 **Selecting the RHEL Installation Option**



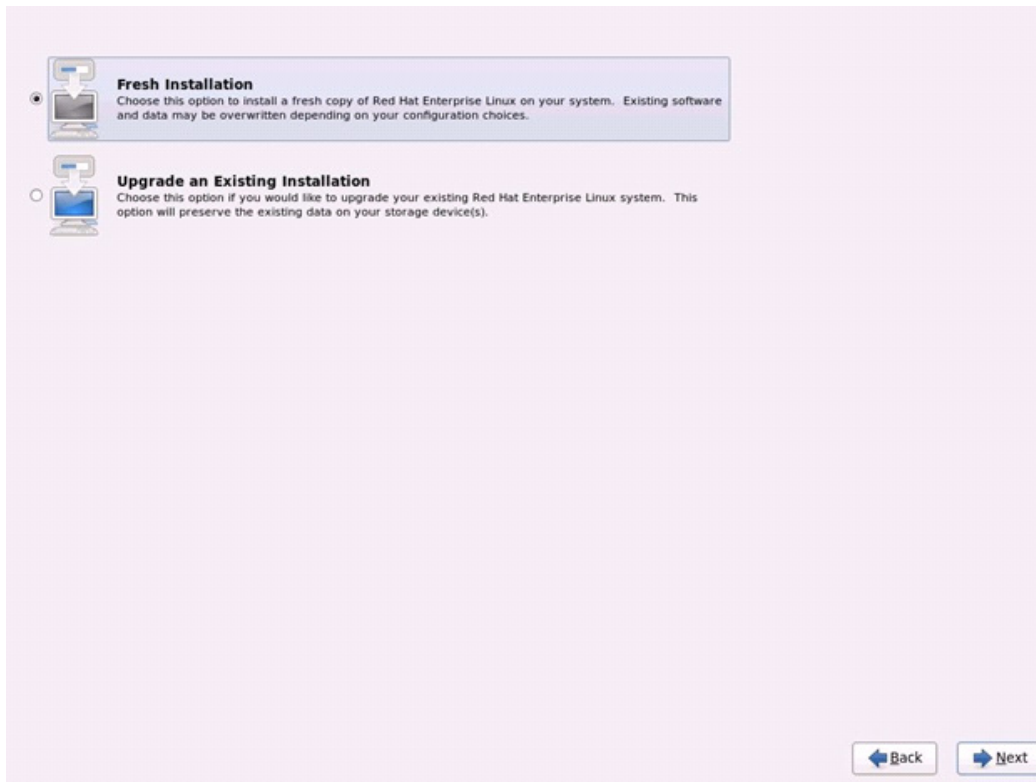
16. Skip the Media test as we are installing from ISO Image, click **Next** to continue.
17. Select Language for the Installation and click **Next**.
18. Select Basic Storage Devices and click **Next**.

Figure 61 **Selecting Storage Device Type**



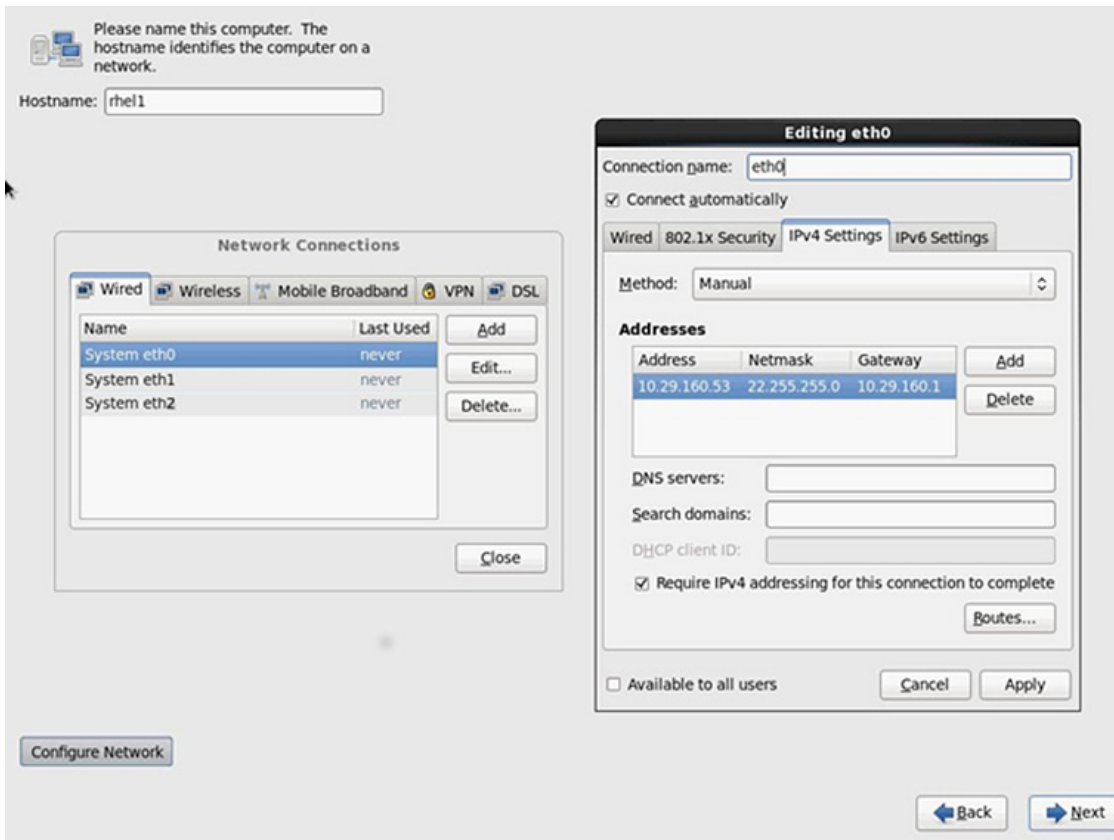
19. Select Fresh Installation and click **Next**.

Figure 62 **Selecting Installation Type**



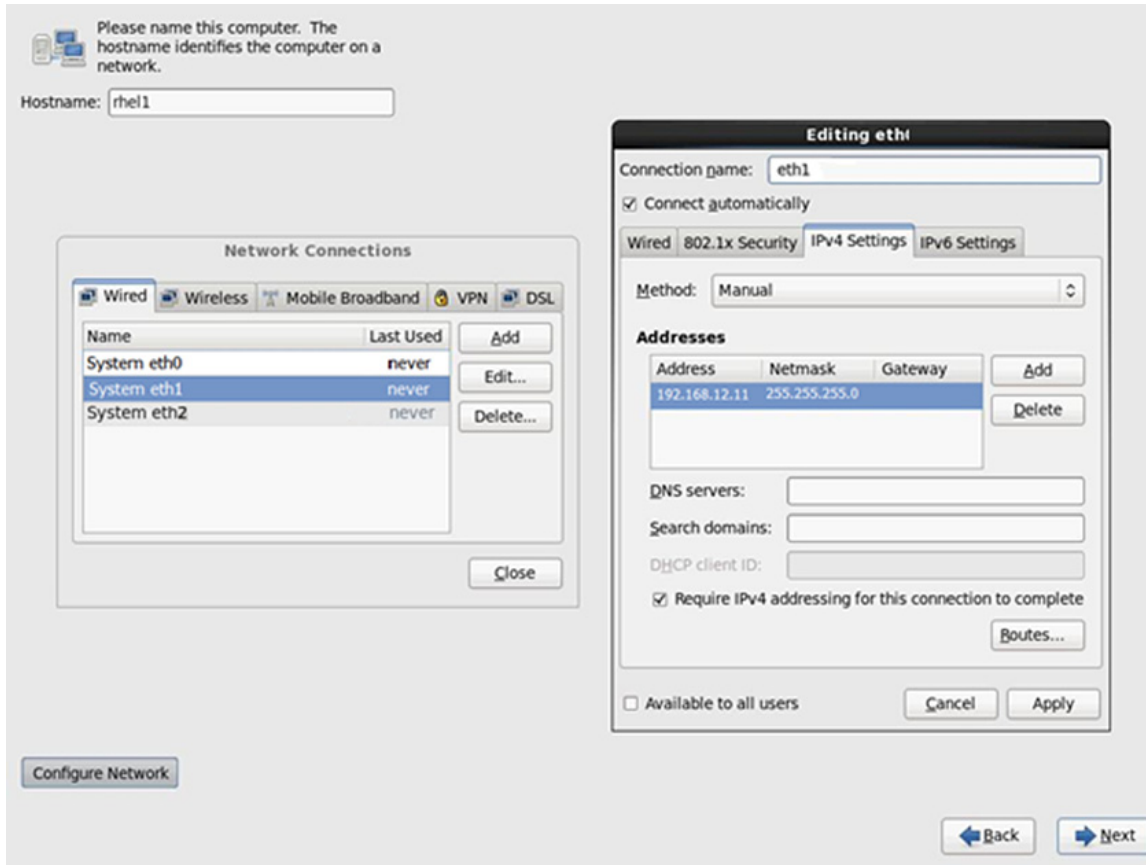
20. Enter the Host name of the server and click **Configure Network**.
21. Network Connections window appears.
22. In the Network Connections window, Select the tab Wired.
23. Select the interface System eth0, and click **Edit**.
24. Editing System eth0 window appears.
25. Check the Connect automatically check box.
26. For the field Method, select Manual from the drop down list.
27. Click **Add** and enter IP Address, Netmask and Gateway.
28. For this solution we have used the following:
 - a. IP Address: 10.29.160.53
 - b. Netmask: 255.255.255.0
 - c. Gateway: 10.29.160.1
29. (Optional) Add DNS servers.
30. Click **Apply**.

Figure 63 Configuring Network for eth0



31. Repeat the steps 24 to 30 for system eth1 with the following:
 - a. IP Address: 192.168.12.11
 - b. Netmask: 255.255.255.0

Figure 64 **Configuring Network for eth1**



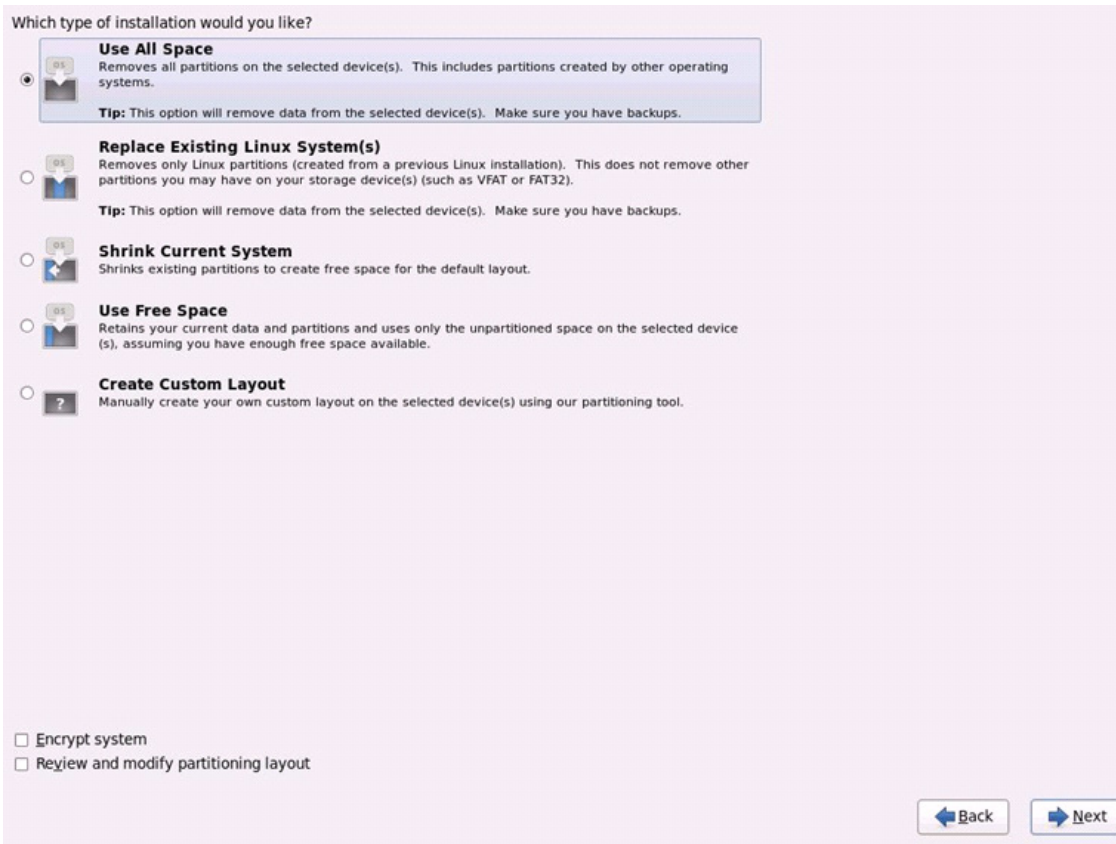
32. Repeat the steps 24 to 30 for system eth2 with the following:
- IP Address: 192.168.11.11
 - Netmask: 255.255.255.0



Note [Table 4](#) lists the IP addresses of the cluster nodes.

- Select the Appropriate Time Zone and click **Next**.
- Enter the root Password and click **Next**.
- Select Use All Space and Click **Next**.

Figure 65 *Selecting RHEL Install Type*



36. Select an appropriate boot drive.
37. Click **Write changes to the disks** and then, click **Next**.

Figure 66 *Writing Partitioning Options into the Disk*



38. Select Basic Server Installation and Click **Next**.

Figure 67 **Selecting RHEL Installation Option**

The default installation of Red Hat Enterprise Linux is a basic server install. You can optionally select a different set of software now.

Basic Server

Database Server

Web Server

Identity Management Server

Virtualization Host

Desktop

Software Development Workstation

Minimal

Please select any additional repositories that you want to use for software installation.

High Availability

Load Balancer

Red Hat Enterprise Linux

...

You can further customize the software selection now, or after install via the software management application.

Customize later Customize now

39. After the installer is finished loading, press **Enter** to continue with the install.

Figure 68 *Installation Process in Progress*



40. Once the installation is complete, reboot the system.
 Repeat steps (step1 to 40) to install Red Hat Linux on Servers 2 through 16.



Note

The OS installation and configuration of the nodes that is mentioned above can be automated through PXE boot or third party tools.

The hostnames and their corresponding IP addresses are shown in [Table 4](#).

Table 4 *Host Names and IP Addresses*

Host Name	eth0	eth1	eth2
rhel1	10.29.160.53	192.168.12.11	192.168.11.11
rhel2	10.29.160.54	192.168.12.12	192.168.11.12
rhel3	10.29.160.55	192.168.12.13	192.168.11.13
rhel4	10.29.160.56	192.168.12.14	192.168.11.14
rhel5	10.29.160.57	192.168.12.15	192.168.11.15
rhel6	10.29.160.58	192.168.12.16	192.168.11.16
rhel7	10.29.160.59	192.168.12.17	192.168.11.17
rhel8	10.29.160.60	192.168.12.18	192.168.11.18
rhel9	10.29.160.61	192.168.12.19	192.168.11.19
rhel10	10.29.160.62	192.168.12.20	192.168.11.20
rhel11	10.29.160.63	192.168.12.21	192.168.11.21

Table 4 *Host Names and IP Addresses*

Host Name	eth0	eth1	eth2
rhel12	10.29.160.64	192.168.12.22	192.168.11.22
rhel13	10.29.160.65	192.168.12.23	192.168.11.23
rhel14	10.29.160.66	192.168.12.24	192.168.11.24
rhel15	10.29.160.67	192.168.12.25	192.168.11.25
rhel16	10.29.160.68	192.168.12.26	192.168.11.26

Post OS Install Configuration

Choose one of the nodes of the cluster or a separate node as Admin Node for management such as installing Cloudera Manager (to install Hadoop), parallel shell, creating a local Red Hat repo and others. In this document, we use rhel1 for this suppose.

Setting Up Password-less Login

To manage all of the clusters nodes from the admin node we need to setup password-less login. It assists in automating common tasks with Parallel-SSH (pssh) and shell-scripts without having to use passwords.

Once Red Hat Linux is installed across all the nodes in the cluster, follow the steps below in order to enable password less login across all the nodes.

1. Login to the Admin Node (rhel1)
ssh 10.29.160.53
2. Run the **ssh-keygen** command to create both public and private keys on the admin node.

Figure 69 *Create Public and Private Keys*

```
[root@rhel1 ~]# ssh-keygen
Generating public/private rsa key pair.
Enter file in which to save the key (/root/.ssh/id_rsa):
Created directory '/root/.ssh'.
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /root/.ssh/id_rsa.
Your public key has been saved in /root/.ssh/id_rsa.pub.
The key fingerprint is:
ab:4e:78:10:54:81:4e:04:8d:af:4f:a4:b2:c4:bb:88 root@rhel1
The key's randomart image is:
+--[ RSA 2048 ]-----+
|  .=ooo.          |
|  ..+            |
|   +             |
|  .+.           |
| .+.  S         |
|.oo .o .        |
|.O.O. o .       |
|+. .o .         |
|E.. .o          |
+-----+

```

- Then run the following command from the admin node to copy the public key `id_rsa.pub` to all the nodes of the cluster. **ssh-copy-id** appends the keys to the remote-host's `.ssh/authorized_key`.

```
for IP in {53..68}; do echo -n "$IP -> "; ssh-copy-id -i ~/.ssh/id_rsa.pub
10.29.160.$IP; done
```

Enter **yes** for Are you sure you want to continue connecting (yes/no)?

Enter the password of the remote host.

Installing and Configuring Parallel Shell

Parallel-SSH

Parallel SSH is used to run commands on several hosts at the same time. It takes a file of hostnames and a bunch of common ssh parameters as parameters, executes the given command in parallel on the nodes specified.

- From the system that is connected to the Internet, download pssh.

```
wget https://parallel-ssh.googlecode.com/files/pssh-2.3.1.tar.gz
```

Figure 70 **Command to Download pssh**

```
[root@redhat ~]# wget https://parallel-ssh.googlecode.com/files/pssh-2.3.1.tar.gz
--2013-04-24 05:39:42-- https://parallel-ssh.googlecode.com/files/pssh-2.3.1.tar.gz
Resolving parallel-ssh.googlecode.com... 74.125.129.82, 2607:f8b0:400e:c02::52
Connecting to parallel-ssh.googlecode.com|74.125.129.82|:443... connected.
HTTP request sent, awaiting response... 200 OK
Length: 23427 (23K) [application/x-gzip]
Saving to: âpssh-2.3.1.tar.gz.1â

100%[=====]
2013-04-24 05:39:43 (240 KB/s) - âpssh-2.3.1.tar.gz.1â
```

```
scp pssh-2.3.1.tar.gz rhell:/root
```

- Copy `pssh-2.3.1.tar.gz` to the Admin Node

```
ssh rhell
tar xzf pssh-2.3.1.tar.gz
cd pssh-2.3.1
python setup.py install
```

Figure 71 Command to Copy pssh to Admin Node

```
[root@redhat ~]# scp pssh-2.3.1.tar.gz rhell:/root
The authenticity of host 'rhell (10.29.160.53)' can't be established.
RSA key fingerprint is 25:15:c9:7d:e0:db:78:2c:0d:ce:e5:2d:e3:e2:5e:44.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'rhell' (RSA) to the list of known hosts.
root@rhell's password:
pssh-2.3.1.tar.gz
[root@redhat ~]# ssh rhell
root@rhell's password:
Last login: Wed Apr 24 09:06:38 2013 from 10.29.160.90
[root@rhell ~]# tar xzf pssh-2.3.1.tar.gz
[root@rhell ~]# cd pssh-2.3.1
[root@rhell pssh-2.3.1]# python setup.py install
running install
running build
running build_py
running build_scripts
running install_lib
running install_scripts
changing mode of /usr/bin/pslurp to 755
changing mode of /usr/bin/pnuke to 755
changing mode of /usr/bin/prsync to 755
changing mode of /usr/bin/pscp to 755
changing mode of /usr/bin/pssh-askpass to 755
changing mode of /usr/bin/pssh to 755
running install_data
running install_egg_info
Removing /usr/lib/python2.6/site-packages/pssh-2.3.1-py2.6.egg-info
Writing /usr/lib/python2.6/site-packages/pssh-2.3.1-py2.6.egg-info
```

3. Extract and Install pssh on the Admin node.
4. Create a host file containing the IP addresses of all the nodes in the cluster. This file is passed as a parameter to pssh to identify the nodes to run the commands on.



Note You can have multiple files based on roles such as datanodes, zookeepernodes, allnodes etc

```
vi /root/allnodes
# This file contains ip address of all nodes of the cluster
#used by parallel-shell (pssh). For Details man pssh
10.29.160.53
10.29.160.54
10.29.160.55
10.29.160.56
10.29.160.57
10.29.160.58
10.29.160.59
10.29.160.60
10.29.160.61
10.29.160.62
10.29.160.63
10.29.160.64
10.29.160.65
10.29.160.66
10.29.160.67
10.29.160.68
```

Configuring /etc/hosts

Follow these steps to create the host file across all the nodes in the cluster:

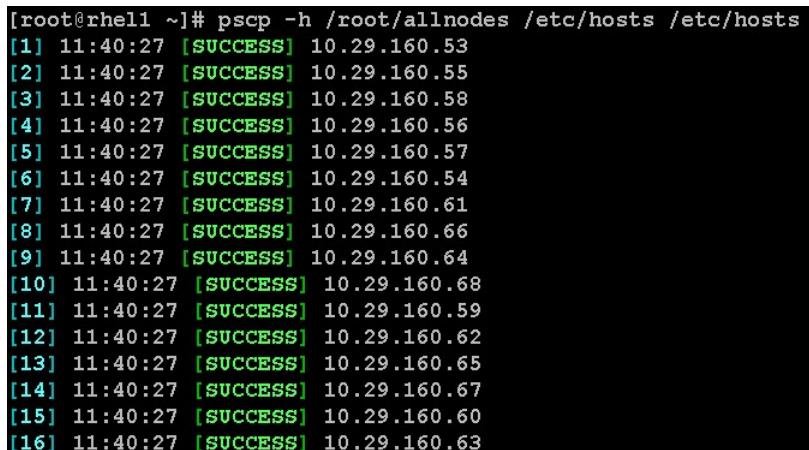
1. Populate the host file with IP addresses and corresponding hostnames on the Admin node (rhel1).

```
vi /etc/hosts
127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4
::1 localhost localhost.localdomain localhost6 localhost6.localdomain6
10.29.160.53 rhel1
10.29.160.54 rhel2
10.29.160.55 rhel3
10.29.160.56 rhel4
10.29.160.57 rhel5
10.29.160.58 rhel6
10.29.160.59 rhel7
10.29.160.60 rhel8
10.29.160.61 rhel9
10.29.160.62 rhel10
10.29.160.63 rhel11
10.29.160.64 rhel12
10.29.160.65 rhel13
10.29.160.66 rhel14
10.29.160.67 rhel15
10.29.160.68 rhel16
```

2. Deploy /etc/hosts from the admin node (rhel1) to all the nodes via the following pscp command:

```
pscp -h /root/allnodes /etc/hosts /etc/hosts
```

Figure 72 Command to Deploy /etc/hosts to All the Nodes



```
[root@rhel1 ~]# pscp -h /root/allnodes /etc/hosts /etc/hosts
[1] 11:40:27 [SUCCESS] 10.29.160.53
[2] 11:40:27 [SUCCESS] 10.29.160.55
[3] 11:40:27 [SUCCESS] 10.29.160.58
[4] 11:40:27 [SUCCESS] 10.29.160.56
[5] 11:40:27 [SUCCESS] 10.29.160.57
[6] 11:40:27 [SUCCESS] 10.29.160.54
[7] 11:40:27 [SUCCESS] 10.29.160.61
[8] 11:40:27 [SUCCESS] 10.29.160.66
[9] 11:40:27 [SUCCESS] 10.29.160.64
[10] 11:40:27 [SUCCESS] 10.29.160.68
[11] 11:40:27 [SUCCESS] 10.29.160.59
[12] 11:40:27 [SUCCESS] 10.29.160.62
[13] 11:40:27 [SUCCESS] 10.29.160.65
[14] 11:40:27 [SUCCESS] 10.29.160.67
[15] 11:40:27 [SUCCESS] 10.29.160.60
[16] 11:40:27 [SUCCESS] 10.29.160.63
```

Create Local Redhat Repo

If your infrastructure node and your cluster nodes have Internet access, you may be able to skip this section.

To create a repository using RHEL DVD or ISO on the admin node (in this deployment rhel1 is used for this purpose), create a directory with all the required RPMs, run the createrepo command and then publish the resulting repository on a website.

1. On the Admin node (rhel1) create a directory that would contain the repository.

```
mkdir -p /var/www/html/rhelrepo
```

2. Copy the contents of the Red Hat DVD to /var/www/html/rhelrepo
3. Alternatively, if you have access to a Red Hat ISO Image, Copy the ISO file to the admin node (rhel1).

```
scp rhel-server-6.2-x86_64-dvd.iso rhel1:/root
```

Here we assume you have the Red Hat ISO file located in your present working directory.

```
mkdir -p /mnt/rheliso
```

```
mount -t iso9660 -o loop /root/rhel-server-6.2-x86_64-dvd.iso /mnt/rheliso/
```

4. Next, copy the contents of the ISO to the /var/www/html/rhelrepo directory

```
cp -r /mnt/rheliso/* /var/www/html/rhelrepo
```

5. Now on rhel1 create a .repo file to enable the use of the yum command.

```
vi /var/www/html/rhelrepo/rheliso.repo
[rhel6.2]
name=Red Hat Enterprise Linux 6.2
baseurl=http://10.29.160.53/rhelrepo
gpgcheck=0
enabled=1
```



Note Based on this repo file yum requires httpd to be running on rhel1 for other nodes to access the repository. Steps to install and configure httpd are in the following section.

6. Copy the rheliso.repo to all the nodes of the cluster.

```
pscp -h /root/allnodes /var/www/html/rhelrepo/rheliso.repo /etc/yum.repos.d/
```

Figure 73 Command to Copy RHEL ISO to All the Nodes

```
[root@rhel1 ~]# pscp -h /root/allnodes /var/www/html/rhelrepo/rheliso.repo
o /etc/yum.repos.d/
[1] 15:00:09 [SUCCESS] 10.29.160.57
[2] 15:00:09 [SUCCESS] 10.29.160.54
[3] 15:00:09 [SUCCESS] 10.29.160.53
[4] 15:00:09 [SUCCESS] 10.29.160.56
[5] 15:00:09 [SUCCESS] 10.29.160.58
[6] 15:00:09 [SUCCESS] 10.29.160.55
[7] 15:00:09 [SUCCESS] 10.29.160.60
[8] 15:00:09 [SUCCESS] 10.29.160.59
[9] 15:00:09 [SUCCESS] 10.29.160.65
[10] 15:00:09 [SUCCESS] 10.29.160.64
[11] 15:00:09 [SUCCESS] 10.29.160.61
[12] 15:00:09 [SUCCESS] 10.29.160.67
[13] 15:00:09 [SUCCESS] 10.29.160.62
[14] 15:00:09 [SUCCESS] 10.29.160.63
[15] 15:00:09 [SUCCESS] 10.29.160.66
[16] 15:00:09 [SUCCESS] 10.29.160.68
```

7. To make use of repository files on rhel1 without httpd, edit the baseurl of repo file /etc/yum.repos.d/rheliso.repo to point repository location in the file system.

```
vi /etc/yum.repos.d/rheliso.repo
[rhel6.2]
name=Red Hat Enterprise Linux 6.2
baseurl=file:///var/www/html/rhelrepo
```

```

    gpgcheck=0
    enabled=1

```

8. `pssh -h /root/allnodes "yum clean all"`.

Figure 74 *Running Yum to Clean All the Nodes*

```

[root@rhell1 ~]# pssh -h /root/allnodes "yum clean all"
[1] 12:14:09 [SUCCESS] 10.29.160.55
[2] 12:14:09 [SUCCESS] 10.29.160.53
[3] 12:14:09 [SUCCESS] 10.29.160.57
[4] 12:14:09 [SUCCESS] 10.29.160.54
[5] 12:14:09 [SUCCESS] 10.29.160.62
[6] 12:14:09 [SUCCESS] 10.29.160.59
[7] 12:14:09 [SUCCESS] 10.29.160.56
[8] 12:14:09 [SUCCESS] 10.29.160.58
[9] 12:14:09 [SUCCESS] 10.29.160.61
[10] 12:14:09 [SUCCESS] 10.29.160.65
[11] 12:14:09 [SUCCESS] 10.29.160.60
[12] 12:14:09 [SUCCESS] 10.29.160.68
[13] 12:14:09 [SUCCESS] 10.29.160.63
[14] 12:14:09 [SUCCESS] 10.29.160.64
[15] 12:14:10 [SUCCESS] 10.29.160.66
[16] 12:14:10 [SUCCESS] 10.29.160.67

```

9. Creating the Red Hat Repository Database.

Install the `createrepo` package. Use it to regenerate the repository database(s) for the local copy of the RHEL DVD contents. Then purge the yum caches.

```

yum -y install createrepo
cd /var/www/html/rhelrepo
createrepo .
yum clean all

```

Figure 75 *Purging Yum Caches*

```

[root@rhell1 rhelrepo]# createrepo .
368/3596 - Packages/pygobject2-doc-2.20.0-5.el6.x86_64.rpm
iso-8859-1 encoding on Ville Skyttä <ville.skytta@iki.fi> - 2.8.2-2
3596/3596 - Packages/lohit-bengali-fonts-2.4.3-6.el6.noarch.rpm
Saving Primary metadata
Saving file lists metadata
Saving other metadata

```

Upgrading LSI driver

The latest LSI driver is required for performance and bug fixes. The latest drivers can be downloaded from the link below:

<http://software.cisco.com/download/release.html?mdfid=284296254&flowid=31743&softwareid=283853158&release=1.5.1&relind=AVAILABLE&rellifecycle=&reltype=latest>

In the ISO image, the required driver `kmod-megaraid_sas-v06.504.01.00.rpm` can be located at `ucs-cxxx-drivers.1.5.1\Linux\Storage\LSI\92xx\RHEL\RHEL6.2`

Figure 76 Command to Download and Transfer Drivers to the Admin Node

```
[root@redhat ~]# scp kmod* rhel1:/root/
kmod-megaraid_sas-debug-v06.504.01.00_ 100% 306KB 306.4KB/s 00:00
kmod-megaraid_sas-v06.504.01.00_rhel6. 100% 302KB 301.5KB/s 00:00
```

From a node connected to the Internet, download and transfer `kmod-megaraid_sas-v06.504.01.00.rpm` to `rhel1` (admin node). Install the rpm on all nodes of the cluster using the following pssh commands. For this example the rpm is assumed to be in present working directory of `rhel1`.

```
pscp -h /root/allnodes kmod-megaraid_sas-v06.504.01.00_rhel6.2-2.x86_64.rpm /root/
```

Figure 77 Copy rpm on All the Nodes

```
[root@rhel1 ~]# pscp -h /root/allnodes kmod-megaraid_sas-v06.504.01.00_rhel6.2-2.x86_64.rpm /root/
[1] 15:46:54 [SUCCESS] 10.29.160.53
[2] 15:46:54 [SUCCESS] 10.29.160.64
[3] 15:46:54 [SUCCESS] 10.29.160.55
[4] 15:46:54 [SUCCESS] 10.29.160.56
[5] 15:46:54 [SUCCESS] 10.29.160.60
[6] 15:46:54 [SUCCESS] 10.29.160.58
[7] 15:46:54 [SUCCESS] 10.29.160.59
[8] 15:46:54 [SUCCESS] 10.29.160.54
[9] 15:46:54 [SUCCESS] 10.29.160.57
[10] 15:46:54 [SUCCESS] 10.29.160.61
[11] 15:46:54 [SUCCESS] 10.29.160.63
[12] 15:46:54 [SUCCESS] 10.29.160.66
[13] 15:46:54 [SUCCESS] 10.29.160.62
[14] 15:46:54 [SUCCESS] 10.29.160.65
[15] 15:46:54 [SUCCESS] 10.29.160.67
[16] 15:46:54 [SUCCESS] 10.29.160.68
```

```
pssh -h /root/allnodes "rpm -ivh kmod-megaraid_sas-v06.504.01.00_rhel6.2-2.x86_64.rpm"
```

Figure 78 Install rpm on All the Nodes

```
[root@rhel1 ~]# pssh -h /root/allnodes "rpm -ivh kmod-megaraid_sas-v06.504.01.00_rhel6.2-2.x86_64.rpm"
[1] 15:49:11 [SUCCESS] 10.29.160.53
[2] 15:49:13 [SUCCESS] 10.29.160.67
[3] 15:49:13 [SUCCESS] 10.29.160.54
[4] 15:49:13 [SUCCESS] 10.29.160.58
[5] 15:49:13 [SUCCESS] 10.29.160.62
[6] 15:49:13 [SUCCESS] 10.29.160.60
[7] 15:49:13 [SUCCESS] 10.29.160.65
[8] 15:49:13 [SUCCESS] 10.29.160.57
[9] 15:49:13 [SUCCESS] 10.29.160.61
[10] 15:49:13 [SUCCESS] 10.29.160.66
[11] 15:49:13 [SUCCESS] 10.29.160.64
[12] 15:49:13 [SUCCESS] 10.29.160.56
[13] 15:49:13 [SUCCESS] 10.29.160.55
[14] 15:49:14 [SUCCESS] 10.29.160.59
[15] 15:49:14 [SUCCESS] 10.29.160.63
[16] 15:49:16 [SUCCESS] 10.29.160.68
```

Configuring NTP

The Network Time Protocol (NTP) is used to synchronize the time of all the nodes within the cluster using the admin node. The Network Time Protocol daemon (ntpd) sets and maintains the system time of day in synchronism with the timeserver located in the admin node. Configuring NTP is critical for any Hadoop Cluster.

1. Configure `/etc/ntp.conf` on the admin node with the following contents:

```
vi /etc/ntp.conf
driftfile /var/lib/ntp/drift
restrict 127.0.0.1
restrict -6 ::1
server 127.127.1.0
fudge 127.127.1.0 stratum 10
includefile /etc/ntp/crypto/pw
keys /etc/ntp/keys
```

2. Create `/root/ntp.conf` on the admin node and copy it to all nodes.

```
vi /root/ntp.conf
server 10.29.160.53
driftfile /var/lib/ntp/drift
restrict 127.0.0.1
restrict -6 ::1
includefile /etc/ntp/crypto/pw
keys /etc/ntp/keys
```

3. Copy `ntp.conf` file from the admin node to `/etc` of all the nodes by executing the following command in the admin node (rhell).

```
for SERVER in {54..68}; do scp /root/ntp.conf 10.29.160.$SERVER:/etc/ntp.conf; done
```



Note Do not use `pssh /root/allnodes` command without editing the host file `allnodes` as it overwrites `/etc/ntp.conf` from the admin node.

Figure 79 Command to Copy `ntp.conf` to All the Nodes

```
[root@rhell ~]# for SERVER in {54..68}; do scp /root/ntp.conf 10.29.160.$SERVER:/etc/ntp.conf; done
ntp.conf 100% 142 0.1KB/s 00:00
ntp.conf 100% 142 0.1KB/s 00:00
ntp.conf 100% 142 0.1KB/s 00:00
ntp.conf 100% 142 0.1KB/s 00:00
ntp.conf 100% 142 0.1KB/s 00:00
ntp.conf 100% 142 0.1KB/s 00:00
ntp.conf 100% 142 0.1KB/s 00:00
ntp.conf 100% 142 0.1KB/s 00:00
ntp.conf 100% 142 0.1KB/s 00:00
ntp.conf 100% 142 0.1KB/s 00:00
ntp.conf 100% 142 0.1KB/s 00:00
ntp.conf 100% 142 0.1KB/s 00:00
ntp.conf 100% 142 0.1KB/s 00:00
ntp.conf 100% 142 0.1KB/s 00:00
ntp.conf 100% 142 0.1KB/s 00:00
ntp.conf 100% 142 0.1KB/s 00:00
ntp.conf 100% 142 0.1KB/s 00:00
```

4. Restart NTP daemon on all the nodes.

```
pssh -h /root/allnodes "service ntpd restart"
```


Figure 80 Command to Restart NTP Daemon on All the Nodes

```
[root@rhell1 ~]# pssh -h /root/allnodes "service ntpd restart"
[1] 13:38:55 [SUCCESS] 10.29.160.54
[2] 13:38:55 [SUCCESS] 10.29.160.53
[3] 13:38:55 [SUCCESS] 10.29.160.56
[4] 13:38:55 [SUCCESS] 10.29.160.57
[5] 13:38:55 [SUCCESS] 10.29.160.55
[6] 13:38:55 [SUCCESS] 10.29.160.58
[7] 13:38:55 [SUCCESS] 10.29.160.60
[8] 13:38:55 [SUCCESS] 10.29.160.59
[9] 13:38:55 [SUCCESS] 10.29.160.64
[10] 13:38:55 [SUCCESS] 10.29.160.62
[11] 13:38:55 [SUCCESS] 10.29.160.61
[12] 13:38:55 [SUCCESS] 10.29.160.66
[13] 13:38:55 [SUCCESS] 10.29.160.63
[14] 13:38:55 [SUCCESS] 10.29.160.65
[15] 13:38:55 [SUCCESS] 10.29.160.67
[16] 13:38:55 [SUCCESS] 10.29.160.68
```

5. Ensure that the NTP daemon restarts after the reboot.

```
pssh -h /root/allnodes "chkconfig ntpd on"
```

Figure 81 Command to Check NTP Daemon Status

```
[root@rhell1 ~]# pssh -h /root/allnodes "chkconfig ntpd on"
[1] 13:52:55 [SUCCESS] 10.29.160.54
[2] 13:52:55 [SUCCESS] 10.29.160.55
[3] 13:52:55 [SUCCESS] 10.29.160.57
[4] 13:52:55 [SUCCESS] 10.29.160.56
[5] 13:52:55 [SUCCESS] 10.29.160.60
[6] 13:52:55 [SUCCESS] 10.29.160.61
[7] 13:52:55 [SUCCESS] 10.29.160.58
[8] 13:52:55 [SUCCESS] 10.29.160.53
[9] 13:52:55 [SUCCESS] 10.29.160.59
[10] 13:52:55 [SUCCESS] 10.29.160.63
[11] 13:52:55 [SUCCESS] 10.29.160.62
[12] 13:52:55 [SUCCESS] 10.29.160.64
[13] 13:52:55 [SUCCESS] 10.29.160.65
[14] 13:52:55 [SUCCESS] 10.29.160.67
[15] 13:52:55 [SUCCESS] 10.29.160.66
[16] 13:52:55 [SUCCESS] 10.29.160.68
```

Installing httpd

Follow these steps to install httpd on the admin node to host repositories.

1. The Red Hat repository is hosted using HTTP on the admin node.; this machine is accessible by all the hosts in the cluster.

```
yum -y install httpd
```

2. Add ServerName and make the necessary changes to the server configuration file.

```
vi /etc/httpd/conf/httpd.conf
ServerName 10.29.160.53:80
```

Figure 82 *Editing Server Configuration File*

```
#ServerName www.example.com:80
ServerName 10.29.160.53:80
#
# UseCanonicalName: Determines how Apache constructs self-referencing
# URLs and the SERVER_NAME and SERVER_PORT variables.
# When set "Off", Apache will use the Hostname and Port supplied
# by the client. When set "On", Apache will use the value of the
# ServerName directive.
#
UseCanonicalName Off
```

3. Ensure httpd is able to read the repofiles

```
chcon -R -t httpd_sys_content_t /var/www/html/rhelrepo
```

4. Start httpd

```
service httpd start
chkconfig httpd on
```

Installing xfsprogs

Install xfsprogs on all the nodes for xfs filesystem.

```
pssh -h /root/allnodes "yum -y install xfsprogs"
```

Figure 83 *Command to Install xfsprogs*

```
[root@rhell ~]# pssh -h /root/allnodes "yum -y install xfsprogs"
[1] 12:26:34 [SUCCESS] 10.29.160.57
[2] 12:26:35 [SUCCESS] 10.29.160.56
[3] 12:26:35 [SUCCESS] 10.29.160.53
[4] 12:26:35 [SUCCESS] 10.29.160.59
[5] 12:26:35 [SUCCESS] 10.29.160.61
[6] 12:26:35 [SUCCESS] 10.29.160.63
[7] 12:26:35 [SUCCESS] 10.29.160.54
[8] 12:26:35 [SUCCESS] 10.29.160.62
[9] 12:26:35 [SUCCESS] 10.29.160.66
[10] 12:26:35 [SUCCESS] 10.29.160.60
[11] 12:26:35 [SUCCESS] 10.29.160.68
[12] 12:26:35 [SUCCESS] 10.29.160.58
[13] 12:26:35 [SUCCESS] 10.29.160.64
[14] 12:26:35 [SUCCESS] 10.29.160.55
[15] 12:26:35 [SUCCESS] 10.29.160.65
[16] 12:26:35 [SUCCESS] 10.29.160.67
```

Configuring Disk Drives for Data

In the section titled Configuring RAID on First Disk Drive for OS describes the steps to configure the first disk drive for the operating system. Remaining disk drives can also be configured similarly or using MegaCli as described below.

1. From the LSI website: www.lsi.com/support/Pages/Download-Results.aspx?keyword=9266-8i download MegaCli and its dependencies and transfer to Admin node.

```

scp /root/MegaCli64 rhell:/root/
scp /root/Lib_Utils-1.00-08.noarch.rpm rhell:/root/
scp /root/Lib_Utils2-1.00-01.noarch.rpm rhell:/root/

```

- Copy all three files to all the nodes using the following commands:

```
pscp -h /root/allnodes /root/MegaCli64 /root/
```

Figure 84 Command to Copy MegaCli

```

[root@rhell ~]# pscp -h /root/allnodes /root/MegaCli64 /root/
[1] 13:00:40 [SUCCESS] 10.29.160.53
[2] 13:00:40 [SUCCESS] 10.29.160.61
[3] 13:00:40 [SUCCESS] 10.29.160.58
[4] 13:00:40 [SUCCESS] 10.29.160.62
[5] 13:00:40 [SUCCESS] 10.29.160.56
[6] 13:00:40 [SUCCESS] 10.29.160.57
[7] 13:00:40 [SUCCESS] 10.29.160.66
[8] 13:00:40 [SUCCESS] 10.29.160.59
[9] 13:00:40 [SUCCESS] 10.29.160.60
[10] 13:00:40 [SUCCESS] 10.29.160.55
[11] 13:00:40 [SUCCESS] 10.29.160.68
[12] 13:00:40 [SUCCESS] 10.29.160.54
[13] 13:00:40 [SUCCESS] 10.29.160.63
[14] 13:00:40 [SUCCESS] 10.29.160.64
[15] 13:00:40 [SUCCESS] 10.29.160.65
[16] 13:00:40 [SUCCESS] 10.29.160.67

```

```
pscp -h /root/allnodes /root/Lib_Utils* /root/
```

Figure 85 Command to Copy MegaCli Dependencies

```

[root@rhell ~]# pscp -h /root/allnodes /root/Lib_Utils* /root/
[1] 13:01:26 [SUCCESS] 10.29.160.53
[2] 13:01:26 [SUCCESS] 10.29.160.58
[3] 13:01:26 [SUCCESS] 10.29.160.59
[4] 13:01:26 [SUCCESS] 10.29.160.60
[5] 13:01:26 [SUCCESS] 10.29.160.67
[6] 13:01:26 [SUCCESS] 10.29.160.63
[7] 13:01:26 [SUCCESS] 10.29.160.61
[8] 13:01:26 [SUCCESS] 10.29.160.57
[9] 13:01:26 [SUCCESS] 10.29.160.54
[10] 13:01:26 [SUCCESS] 10.29.160.56
[11] 13:01:26 [SUCCESS] 10.29.160.62
[12] 13:01:26 [SUCCESS] 10.29.160.55
[13] 13:01:26 [SUCCESS] 10.29.160.64
[14] 13:01:26 [SUCCESS] 10.29.160.66
[15] 13:01:26 [SUCCESS] 10.29.160.65
[16] 13:01:26 [SUCCESS] 10.29.160.68

```

- Run the following command to install the rpms on all the nodes:

```
pssh -h /root/allnodes "rpm -ivh Lib_Utils*"
```

Figure 86 Command to Install rpm on All the Nodes

```

[root@rhel1 ~]# pssh -h /root/allnodes "rpm -ivh Lib_Utils*"
[1] 13:02:05 [SUCCESS] 10.29.160.64
[2] 13:02:05 [SUCCESS] 10.29.160.62
[3] 13:02:05 [SUCCESS] 10.29.160.57
[4] 13:02:05 [SUCCESS] 10.29.160.66
[5] 13:02:05 [SUCCESS] 10.29.160.58
[6] 13:02:05 [SUCCESS] 10.29.160.59
[7] 13:02:05 [SUCCESS] 10.29.160.54
[8] 13:02:05 [SUCCESS] 10.29.160.67
[9] 13:02:05 [SUCCESS] 10.29.160.60
[10] 13:02:05 [SUCCESS] 10.29.160.65
[11] 13:02:05 [SUCCESS] 10.29.160.56
[12] 13:02:05 [SUCCESS] 10.29.160.55
[13] 13:02:05 [SUCCESS] 10.29.160.63
[14] 13:02:05 [SUCCESS] 10.29.160.61
[15] 13:02:05 [SUCCESS] 10.29.160.68
[16] 13:02:05 [SUCCESS] 10.29.160.53

```

- Issue the following command from the admin node to create the virtual drives with RAID 0 configurations on all the nodes (rhel1-16).

```

pssh -h /root/allnodes ". /MegaCli64 -cfgeachdskraid0 WB RA direct NoCachedBadBBU
strpsz1024 -a0"
WB: Write back
RA: Read ahead
Direct: Reads are not buffered in cache memory
NoCachedBadBBU: Do not write cache when the BBU is bad
strpsz1024: Strip Size of 1024K

```



Note

The command above will not override existing configurations. To clear and reconfigure existing configurations refer to Embedded MegaRAID Software Users Guide available at www.lsi.com.

Configuring File System

Follow these steps to configure the file system for CDH:

- On the Admin node, create a file containing the following script.

To create partition tables and file systems on the local disks supplied to each of the nodes, run the following script as the root user on each node.

```

vi /root/driveconf.sh
#!/bin/bash
disks_count=`lsblk -id | grep sd | wc -l`
if [ $disks_count -eq 24 ]; then
    echo "Found 24 disks"
else
    echo "Found $disks_count disks. Expecting 24. Exiting.."
    exit 1
fi
[[ "-x" == "${1}" ]] && set -x && set -v && shift 1
for X in /sys/class/scsi_host/host*/scan
do
    echo '- - -' > ${X}
done
for X in /dev/sd?
do
    echo $X

```

```

if [[ -b ${X} && ` /sbin/parted -s ${X} print quit|/bin/grep -c boot` -ne 0 ]]
then
echo "${X} bootable - skipping."
continue
else
Y=${X##*/}1
/sbin/parted -s ${X} mklabel gpt quit
/sbin/parted -s ${X} mkpart 1 6144s 100% quit
/sbin/mkfs.xfs -f -q -l size=65536b, lazy-count=1, su=256k -d sunit=1024, swidth=6144
-r extsize=256k -L ${Y} ${X}1
(( $? )) && continue
/bin/mkdir -p /CDH/${Y}
(( $? )) && continue
/bin/mount -t xfs -o allocsize=128m, noatime, nobarrier, nodiratime ${X}1
/CDH/${Y}
(( $? )) && continue
echo "LABEL=${Y} /CDH/${Y} xfs allocsize=128m, noatime, nobarrier, nodiratime 0 0" >>
/etc/fstab
fi
done

```

2. Run the following command to copy driveconf.sh to all the nodes:

```
pscp -h /root/allnodes /root/driveconf.sh /root/
```

Figure 87 Command to Copy driveconf.sh

```

[root@rhell ~]# pscp -h /root/allnodes /root/driveconf.sh /root/
[1] 16:12:42 [SUCCESS] 10.29.160.53
[2] 16:12:42 [SUCCESS] 10.29.160.60
[3] 16:12:42 [SUCCESS] 10.29.160.57
[4] 16:12:42 [SUCCESS] 10.29.160.59
[5] 16:12:42 [SUCCESS] 10.29.160.58
[6] 16:12:42 [SUCCESS] 10.29.160.54
[7] 16:12:42 [SUCCESS] 10.29.160.56
[8] 16:12:42 [SUCCESS] 10.29.160.61
[9] 16:12:42 [SUCCESS] 10.29.160.55
[10] 16:12:42 [SUCCESS] 10.29.160.63
[11] 16:12:42 [SUCCESS] 10.29.160.65
[12] 16:12:42 [SUCCESS] 10.29.160.62
[13] 16:12:42 [SUCCESS] 10.29.160.64
[14] 16:12:42 [SUCCESS] 10.29.160.66
[15] 16:12:42 [SUCCESS] 10.29.160.67
[16] 16:12:42 [SUCCESS] 10.29.160.68

```

3. Run the following command from the admin node to run the script across all the other nodes:

```
pssh -h /root/allnodes "./driveconf.sh"
```

Figure 88 Running `driveconfig.sh` on All the Nodes

```
[root@rhell ~]# pssh -h /root/allnodes "./driveconf.sh"
[1] 16:15:24 [SUCCESS] 10.29.160.67
[2] 16:15:24 [SUCCESS] 10.29.160.54
[3] 16:15:24 [SUCCESS] 10.29.160.63
[4] 16:15:24 [SUCCESS] 10.29.160.66
[5] 16:15:24 [SUCCESS] 10.29.160.65
[6] 16:15:24 [SUCCESS] 10.29.160.62
[7] 16:15:24 [SUCCESS] 10.29.160.61
[8] 16:15:24 [SUCCESS] 10.29.160.60
[9] 16:15:24 [SUCCESS] 10.29.160.59
[10] 16:15:24 [SUCCESS] 10.29.160.58
[11] 16:15:24 [SUCCESS] 10.29.160.57
[12] 16:15:24 [SUCCESS] 10.29.160.64
[13] 16:15:25 [SUCCESS] 10.29.160.56
[14] 16:15:25 [SUCCESS] 10.29.160.55
[15] 16:15:25 [SUCCESS] 10.29.160.53
[16] 16:15:35 [SUCCESS] 10.29.160.68
```

Prerequisites for CDH

Following prerequisites are required for installing CDH:

- [Disable SELinux, page 86](#)
- [Disabling iptables, page 87](#)
- [Download and Configure Cloudera Repo, page 88](#)
- [Oracle JDK Installation, page 92](#)

Disable SELinux

Cloudera Hadoop Installation requires all the nodes to disable SELinux. This will completely disable all SELinux functions including file and process labeling. This can be done by editing `/etc/selinux/config` and changing the SELINUX line to SELINUX=disabled.

```
pssh -h /root/allnodes "sed -i 's/enforcing/disabled/g' /etc/selinux/config "
pssh -h /root/allnodes "setenforce 0"
```



Note

The above command may fail if SELinux is already disabled.

Figure 89 Command to Disable SELinux

```
[root@rhell ~]# pssh -h /root/allnodes "sed -i 's/enforcing/disabled/g' /etc/selinux/config"
[1] 14:07:40 [SUCCESS] 10.29.160.53
[2] 14:07:40 [SUCCESS] 10.29.160.54
[3] 14:07:40 [SUCCESS] 10.29.160.57
[4] 14:07:40 [SUCCESS] 10.29.160.55
[5] 14:07:40 [SUCCESS] 10.29.160.56
[6] 14:07:40 [SUCCESS] 10.29.160.59
[7] 14:07:40 [SUCCESS] 10.29.160.58
[8] 14:07:40 [SUCCESS] 10.29.160.63
[9] 14:07:40 [SUCCESS] 10.29.160.61
[10] 14:07:40 [SUCCESS] 10.29.160.60
[11] 14:07:40 [SUCCESS] 10.29.160.66
[12] 14:07:40 [SUCCESS] 10.29.160.67
[13] 14:07:40 [SUCCESS] 10.29.160.62
[14] 14:07:40 [SUCCESS] 10.29.160.65
[15] 14:07:40 [SUCCESS] 10.29.160.64
[16] 14:07:40 [SUCCESS] 10.29.160.68
```

Disabling iptables

To disable the Linux firewall, run the following commands on all nodes:

```
pssh -h /root/allnodes "service iptables stop"
```

Figure 90 Command to Disable Linux Firewall

```
[root@rhell ~]# pssh -h /root/allnodes "service iptables stop"
[1] 14:13:25 [SUCCESS] 10.29.160.54
[2] 14:13:25 [SUCCESS] 10.29.160.55
[3] 14:13:25 [SUCCESS] 10.29.160.57
[4] 14:13:25 [SUCCESS] 10.29.160.59
[5] 14:13:25 [SUCCESS] 10.29.160.56
[6] 14:13:25 [SUCCESS] 10.29.160.62
[7] 14:13:25 [SUCCESS] 10.29.160.60
[8] 14:13:25 [SUCCESS] 10.29.160.66
[9] 14:13:25 [SUCCESS] 10.29.160.61
[10] 14:13:25 [SUCCESS] 10.29.160.63
[11] 14:13:25 [SUCCESS] 10.29.160.67
[12] 14:13:25 [SUCCESS] 10.29.160.58
[13] 14:13:25 [SUCCESS] 10.29.160.53
[14] 14:13:25 [SUCCESS] 10.29.160.68
[15] 14:13:25 [SUCCESS] 10.29.160.65
[16] 14:13:25 [SUCCESS] 10.29.160.64
```

```
pssh -h /root/allnodes "chkconfig iptables off"
```

Figure 91 Command to Check the Linux Firewall Status

```
[root@rhell1 ~]# pssh -h /root/allnodes "chkconfig iptables off"
[1] 14:13:25 [SUCCESS] 10.29.160.54
[2] 14:13:25 [SUCCESS] 10.29.160.55
[3] 14:13:25 [SUCCESS] 10.29.160.57
[4] 14:13:25 [SUCCESS] 10.29.160.59
[5] 14:13:25 [SUCCESS] 10.29.160.56
[6] 14:13:25 [SUCCESS] 10.29.160.62
[7] 14:13:25 [SUCCESS] 10.29.160.60
[8] 14:13:25 [SUCCESS] 10.29.160.66
[9] 14:13:25 [SUCCESS] 10.29.160.61
[10] 14:13:25 [SUCCESS] 10.29.160.63
[11] 14:13:25 [SUCCESS] 10.29.160.67
[12] 14:13:25 [SUCCESS] 10.29.160.58
[13] 14:13:25 [SUCCESS] 10.29.160.53
[14] 14:13:25 [SUCCESS] 10.29.160.68
[15] 14:13:25 [SUCCESS] 10.29.160.65
[16] 14:13:25 [SUCCESS] 10.29.160.64
```

Download and Configure Cloudera Repo

CDH can be installed in many ways. Method 1 demonstrates the installation of CDH4 using Cloudera Manager when all the hosts of the cluster have an Internet connection. Method 2 demonstrates the installation of CDH4 if the cluster has no connectivity to the Internet.

Figure 92 Downloading and Executing Cloudera Manager

```
# curl -O http://archive.cloudera.com/cm4/installer/latest/cloudera-manager-installer.bin

% Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
           Dload  Upload   Total     Spent    Left  Speed
100 489k 100 489k    0     0 255k      0  0:00:01  0:00:01  --:--:-- 258k
# chmod +x cloudera-manager-installer.bin
# ./cloudera-manager-installer.bin
```

- **Method 1**—The easiest way to install the CDH4 is through Cloudera Manager. If all the hosts in the cluster are connected to the Internet, download the latest version of Cloudera Manager from the Cloudera website to the admin node and execute it.
- **Method 2**—If all the hosts are not connected to the internet we use local repository to install CDH4. We first need to create a directory with all the required rpms, run the createrepo command, and then publish the resulting repository on a website. We use the Admin node to host the repository.

Run the following commands from the admin node:

1. Create the directories within the admin node.

```
mkdir -p /var/www/html/clouderarepo/
```

2. Download the Cloudera Software Repos.

From a host connected to the Internet, download the Cloudera Software Repo as shown below and transfer it to the admin node.

3. Download Cloudera Manager Repository from the system connected to the Internet.

```
mkdir -p /tmp/clouderarepo/
cd /tmp/clouderarepo/
```

```
wget http://archive.cloudera.com/cm4/redhat/6/x86_64/cm/cloudera-manager.repo
```

```
reposync --config=./cloudera-manager.repo --repoid=cloudera-manager
```


Figure 93 Downloading Cloudera Manager Repository

```
[root@redhat clouderarepo]# wget http://archive.cloudera.com/cm4/redhat/6/x86_64/cm/cloudera-manager.repo
--2013-04-24 08:58:22-- http://archive.cloudera.com/cm4/redhat/6/x86_64/cm/cloudera-manager.repo
Resolving archive.cloudera.com... 184.73.217.71
Connecting to archive.cloudera.com|184.73.217.71|:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 275 [text/plain]
Saving to: âcloudera-manager.repo.1â

100% [=====>] 275      --.-K/s   in 0s

2013-04-24 08:58:22 (54.6 MB/s) - âcloudera-manager.repo.1â

[root@redhat clouderarepo]# reposync --config=./cloudera-manager.repo --repoid=cloudera-manager
```

4. Download the Cloudera Manager Installer.

```
cd /tmp/clouderarepo/
wget
http://archive.cloudera.com/cm4/installer/latest/cloudera-manager-installer.bin
```

Figure 94 Downloading Cloudera Manager Installer

```
[root@redhat clouderarepo]# wget http://archive.cloudera.com/cm4/installer/latest/cloudera-manager-installer.bin
--2013-04-24 09:04:40-- http://archive.cloudera.com/cm4/installer/latest/cloudera-manager-installer.bin
Resolving archive.cloudera.com... 184.73.217.71
Connecting to archive.cloudera.com|184.73.217.71|:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 501701 (490K) [application/octet-stream]
Saving to: âcloudera-manager-installer.binâ

5% [====>] 27,199
```

5. Download the CDH4 Repository

```
cd /tmp/clouderarepo/
wget http://archive.cloudera.com/cdh4/redhat/6/x86_64/cdh/cloudera-cdh4.repo
reposync --config=./cloudera-cdh4.repo --repoid=cloudera-cdh4
```

Figure 95 Downloading CDH4 Repository

```
[root@redhat clouderarepo]# wget http://archive.cloudera.com/cdh4/redhat/6/x86_64/cdh/cloudera-cdh4.repo
--2013-04-24 09:12:58-- http://archive.cloudera.com/cdh4/redhat/6/x86_64/cdh/cloudera-cdh4.repo
Resolving archive.cloudera.com... 184.73.217.71
Connecting to archive.cloudera.com|184.73.217.71|:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 232 [text/plain]
Saving to: âcloudera-cdh4.repo.1â

100% [=====>] 232      --.-K/s   in 0s

2013-04-24 09:12:58 (50.6 MB/s) - âcloudera-cdh4.repo.1â

[root@redhat clouderarepo]# reposync --config=./cloudera-cdh4.repo --repoid=cloudera-cdh4
```

6. Download Impala Repository.

```
cd /tmp/clouderarepo/
wget
http://archive.cloudera.com/impala/redhat/6/x86_64/impala/cloudera-impala.repo
reposync --config=./cloudera-impala.repo --repoid=cloudera-impala
```

Figure 96 *Downloading Impala Repository*

```
[root@redhat clouderarepo]# wget http://beta.cloudera.com/impala/redhat/6/x86_64/impala/cloudera-impala.repo
--2013-04-24 09:27:00-- http://beta.cloudera.com/impala/redhat/6/x86_64/impala/cloudera-impala.repo
Resolving beta.cloudera.com... 54.243.131.33
Connecting to beta.cloudera.com|54.243.131.33|:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 199 [text/plain]
Saving to: âcloudera-impala.repo.1â

100%[=====>] 199          --.-K/s   in 0s

2013-04-24 09:27:00 (38.4 MB/s) - âcloudera-impala.repo.1â

[root@redhat clouderarepo]# reposync --config=./cloudera-impala.repo --repoid=cloudera-impala
```

7. Copy the repository directory to the admin node

```
scp -r /tmp/clouderarepo/ rhell1:/var/www/html
```

Figure 97 *Copying Cloudera Repository to the Admin Node*

```
[root@redhat clouderarepo]# scp -r /tmp/clouderarepo/ rhell1:/var/www/html/
root@rhell1's password:
cloudera-cdh4.repo 100% 232 0.2KB/s 00:00
hadoop-mapreduce-2.0.0+960-1.cdh4.2.1.p0.9.e16.x86_64.rpm 100% 9891KB 9.7MB/s 00:00
hue-beeswax-2.2.0+194-1.cdh4.2.1.p0.8.e16.x86_64.rpm 100% 654KB 654.3KB/s 00:00
hue-impala-2.2.0+194-1.cdh4.2.1.p0.8.e16.x86_64.rpm 100% 19KB 19.0KB/s 00:00
hadoop-2.0.0+960-1.cdh4.2.1.p0.9.e16.x86_64.rpm 100% 16MB 15.8MB/s 00:01
hadoop-libhdfs-2.0.0+960-1.cdh4.2.1.p0.9.e16.x86_64.rpm 100% 26KB 26.0KB/s 00:00
hadoop-hdfs-datanode-2.0.0+960-1.cdh4.2.1.p0.9.e16.x86_64.rpm 100% 5224 5.1KB/s 00:00
hadoop-doc-2.0.0+960-1.cdh4.2.1.p0.9.e16.x86_64.rpm 100% 4096KB 4.0MB/s 00:00
hadoop-hdfs-secondarynamenode-2.0.0+960-1.cdh4.2.1.p0.9.e16.x86_64.rpm 100% 4964 4.9KB/s 00:00
hue-jobs-2.2.0+194-1.cdh4.2.1.p0.8.e16.x86_64.rpm 100% 208KB 207.8KB/s 00:00
hue-base-0.94.2+218-1.cdh4.2.1.p0.8.e16.x86_64.rpm 100% 36MB 36.5MB/s 00:00
hue-about-2.2.0+194-1.cdh4.2.1.p0.8.e16.x86_64.rpm 100% 16KB 16.0KB/s 00:00
hue-base-thrift-0.94.2+218-1.cdh4.2.1.p0.8.e16.x86_64.rpm 100% 5316 5.2KB/s 00:00
hue-help-2.2.0+194-1.cdh4.2.1.p0.8.e16.x86_64.rpm 100% 14KB 14.2KB/s 00:00
```

8. On admin node (rhell) run create repo command.

```
cd /var/www/html/clouderarepo/

createrepo --baseurl http://10.29.160.53/clouderarepo/cloudera-manager/
/var/www/html/clouderarepo/cloudera-manager

createrepo --baseurl http://10.29.160.53/clouderarepo/cloudera-cdh4
/var/www/html/clouderarepo/cloudera-cdh4

createrepo --baseurl http://10.29.160.53/clouderarepo/cloudera-impala
/var/www/html/clouderarepo/cloudera-impala
```

Figure 98 Running `createrepo` Command on the Admin Node

```
[root@rhell clouderarepo]# createrepo --baseurl http://10.29.160.53/clouderarepo
/cloudera-manager/ /var/www/html/clouderarepo/cloudera-manager
7/7 - RPMS/x86_64/cloudera-manager-agent-4.5.1-1.cm451.p0.294.x86_64.rpm
Saving Primary metadata
Saving file lists metadata
Saving other metadata
[root@rhell clouderarepo]# createrepo --baseurl http://10.29.160.53/clouderarepo
/cloudera-cdh4 /var/www/html/clouderarepo/cloudera-cdh4
77/77 - RPMS/noarch/mahout-0.7+15-1.cdh4.2.1.p0.8.e16.noarch.rpm
Saving Primary metadata
Saving file lists metadata
Saving other metadata
[root@rhell clouderarepo]# createrepo --baseurl http://10.29.160.53/clouderarepo
/cloudera-impala /var/www/html/clouderarepo/cloudera-impala
6/6 - RPMS/noarch/bigtop-utils-0.4+300-1.cdh4.0.1.p0.1.e16.noarch.rpm
Saving Primary metadata
Saving file lists metadata
Saving other metadata
```



Note Visit <http://10.29.160.53/clouderarepo> to verify the files.

9. Create the Cloudera Manager repo file with following contents:

```
vi /var/www/html/clouderarepo/cloudera-manager/cm.repo
[cloudera-manager]
name=Cloudera Manager
baseurl=http://10.29.160.53/clouderarepo/cloudera-manager/
gpgcheck = 0
```

10. Create the Hadoop repo file with following contents:

```
vi /var/www/html/clouderarepo/cloudera-cdh4/ch.repo
[cloudera-cdh4]
name=Cloudera's Distribution for Hadoop, Version 4
baseurl= http://10.29.160.53/clouderarepo/cloudera-cdh4
gpgcheck = 0
```

11. Create the Cloudera Impala repo file with following contents:

```
vi /var/www/html/clouderarepo/cloudera-impala/ci.repo
[cloudera-impala]
name=Impala
baseurl= http://10.29.160.53/clouderarepo/cloudera-impala
gpgcheck = 0
```

Copy the file `cm.repo`, `ch.repo`, `ci.repo` into `/etc/yum.repos.d/` on the admin node to enable it to find the packages that are locally hosted.

```
cp /var/www/html/clouderarepo/cloudera-manager/cm.repo /etc/yum.repos.d/
cp /var/www/html/clouderarepo/cloudera-cdh4/ch.repo /etc/yum.repos.d/
cp /var/www/html/clouderarepo/cloudera-impala/ci.repo /etc/yum.repos.d/
```

Figure 99 Copying `.repo` Files on the Admin Node

```
[root@rhell yum.repos.d]# cp /var/www/html/clouderarepo/cloudera-manager/cm.repo
etc/yum.repos.d/
[root@rhell yum.repos.d]# cp /var/www/html/clouderarepo/cloudera-cdh4/ch.repo /et
c/yum.repos.d/
[root@rhell yum.repos.d]# cp /var/www/html/clouderarepo/cloudera-impala/ci.repo /
etc/yum.repos.d/
```

The contents of the repo files are:

Copy the file cm.repo, ch.repo, ci.repo to /etc/yum.repos.d/ of all the nodes of the cluster from the admin node.

```
pscp -h /root/allnodes /etc/yum.repos.d/c* /etc/yum.repos.d/
```

Figure 100 Copying .repo Files on All the Nodes

```
[root@rhell yum.repos.d]# ls
ch.repo ci.repo cm.repo rheliso.repo
[root@rhell yum.repos.d]# cat cm.repo
[cloudera-manager]
name=Cloudera Manager
baseurl=http://10.29.160.53/clouderarepo/cloudera-manager/
gpgcheck = 0
[root@rhell yum.repos.d]# cat ch.repo
[cloudera-cdh4]
name=Cloudera's Distribution for Hadoop, Version 4
baseurl= http://10.29.160.53/clouderarepo/cloudera-cdh4
gpgcheck = 0
[root@rhell yum.repos.d]# cat cm.repo
[cloudera-manager]
name=Cloudera Manager
baseurl=http://10.29.160.53/clouderarepo/cloudera-manager/
gpgcheck = 0
[root@rhell yum.repos.d]# cat rheliso.repo
[rhel6.2]
name=Red Hat Enterprise Linux 6.2
baseurl=http://10.29.160.53/rhelrepo
gpgcheck=0
enabled=1
```

Oracle JDK Installation

From the system connected to the Internet, using a web browser, download the latest version of Java JDK 6 and copy it to /root/ on rhell

```
scp jdk-6u37-linux-x64-rpm.bin rhell:/root/
```

From rhell, copy jdk-6u37-linux-x64-rpm.bin to all nodes of the cluster.

```
pscp -h /root/allnodes /root/jdk-6u37-linux-x64-rpm.bin /root/
```

Install JDK6 on all the nodes by logging into each node and executing the following command from /root/.

```
sh jdk-6u37-linux-x64-rpm.bin -noregister
```

Figure 101 Verify Java Installation on All the Nodes

```
[root@rhell ~]# java -version
java version "1.6.0_37"
Java(TM) SE Runtime Environment (build 1.6.0_37-b06)
Java HotSpot(TM) 64-Bit Server VM (build 20.12-b01, mixed mode)
```



Note

Uninstall all versions of java if there is any conflict during installation.

Installing Cloudera Manager

Cloudera Manager, an end to end management application, is used to install and configure CDH. During CDH Installation, Cloudera Manager's Wizard will help to install Hadoop services on all nodes using the following procedure:

- Discovery of the cluster nodes
- Configure the Cloudera parcel or package repositories
- Install Hadoop, Cloudera Manager Agent (CMA) and Impala on all the cluster nodes.
- Install the Oracle JDK if it is not already installed across all the cluster nodes.
- Assign various services to nodes.
- Start the Hadoop services.

Follow these steps to install Cloudera Manager:

1. Change the permission of Cloudera Manager Installer on the admin node.

```
cd /var/www/html/clouderarepo
chmod 777 cloudera-manager-installer.bin
```

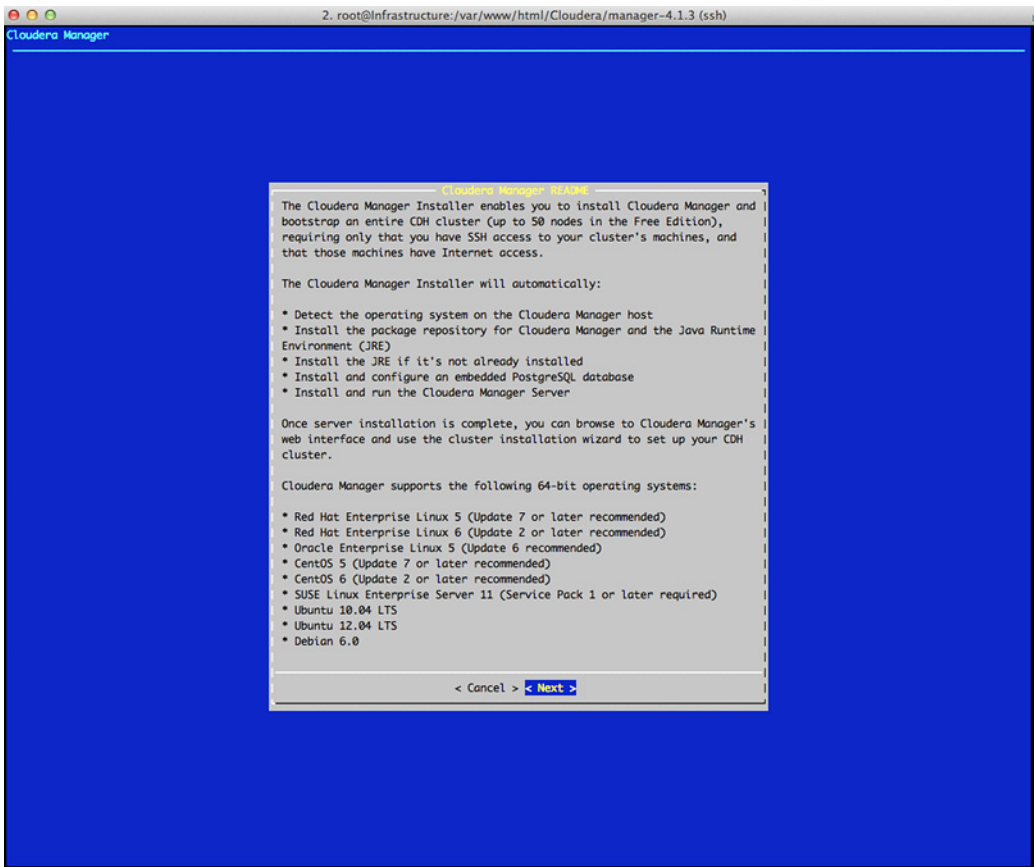
Figure 102 **Starting Cloudera Manager Installer**

```
[root@rhel1 cloudera-manager]# ls
cloudera-manager-installer.bin  cm.repo  repodata  RPMS
[root@rhel1 cloudera-manager]# ./cloudera-manager-installer.bin
```


2. Execute the following command in the admin node (rhel1) to start Cloudera Manager Installer.

```
cd /var/www/html/clouderarepo/cloudera-manager
./cloudera-manager-installer.bin
```

3. This displays the Cloudera Manager Read Me file. Click **Next**.


Figure 103 *Cloudera Manager Installer*

4. Click **Next** in the End User License agreement page.
5. Click **Yes** in the license agreement confirmation page.
6. Click **Yes** in the Oracle Binary Code License Agreement for the Java SE Platform Products page.

Figure 104 *Accepting the Oracle Binary Code License Agreement*


7. Wait for the installer to install the packages needed for Cloudera Manager.

Figure 105 **Installation In Progress**




8. Save the url displayed `http://10.29.160.53:7180`. You will need this url to access Cloudera Manager. If you are unable to connect to the server, check to see if iptables and SELinux are disabled.

Figure 106 Cloudera Manager URL



9. Click OK.

Figure 107 Cloudera Manager is Installed on the Cluster




10. Once the installation of Cloudera Manager is complete. Install CDH4 using the Cloudera Manager web interface.

Installing Cloudera Enterprise Core (CDH4)

To install Cloudera Enterprise Core, follow these steps:

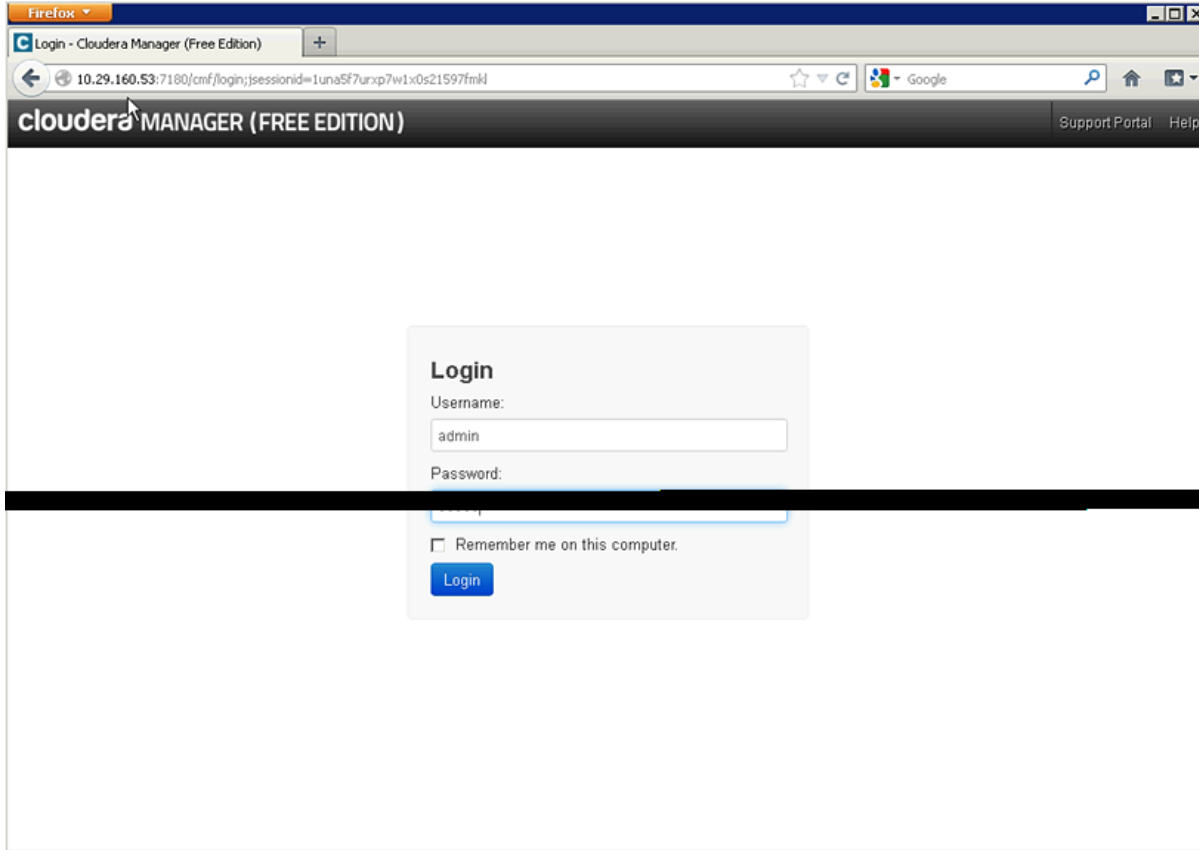
1. Access the Cloudera Manager using the URL displayed by the Installer, <http://10.29.160.53:7180>.

Figure 108 Starting Cloudera Manager



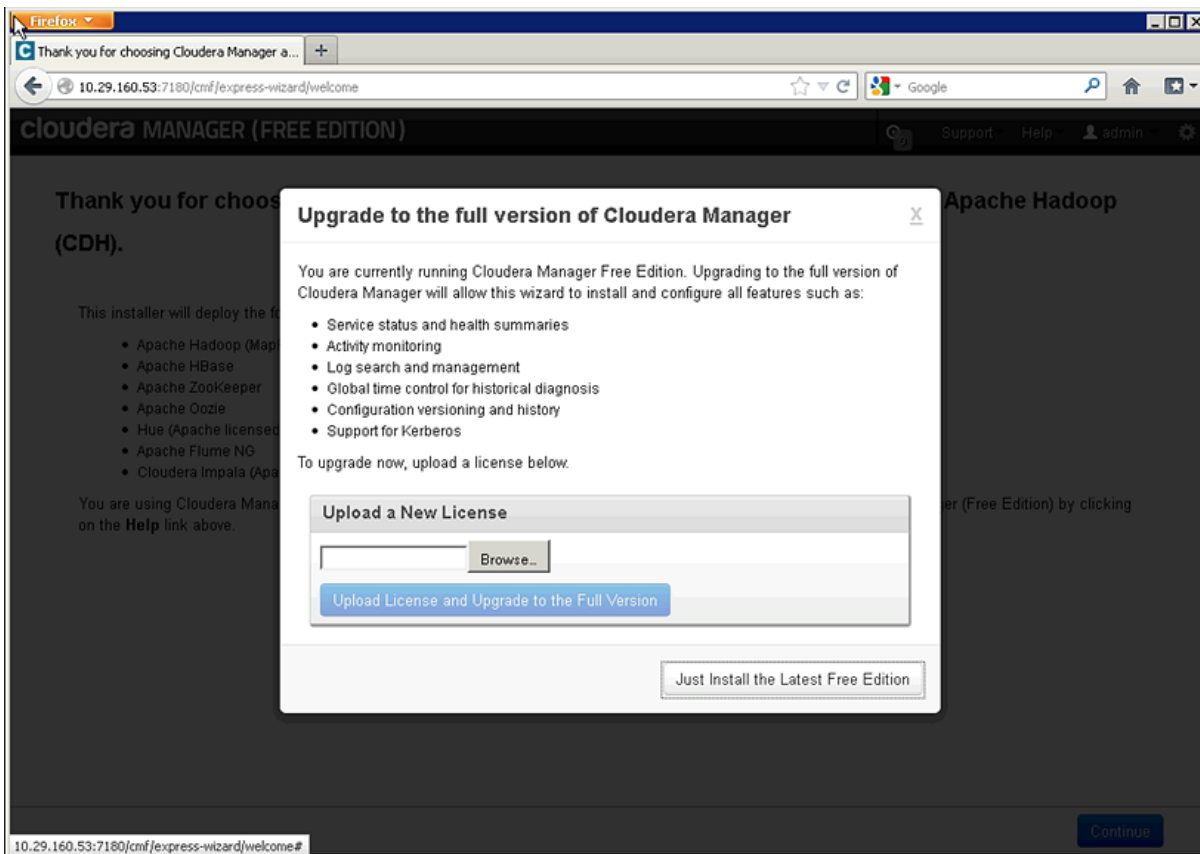
2. Login to the Cloudera Manager. Enter "admin" for both the Username and Password fields.

Figure 109 Cloudera Manager Login page



3. Click Just Install the Latest Free Edition.


Figure 110 Installing Cloudera Manager



4. Click **Continue** in the confirmation page.
5. Specify the hosts that are part of the cluster using their IP addresses or hostname. The figure below shows use of a pattern to specify ip-address range.
192.168.12.[11-26]
6. After the IP addresses are entered, click **Search**.

Figure 111 Searching for Cluster Nodes

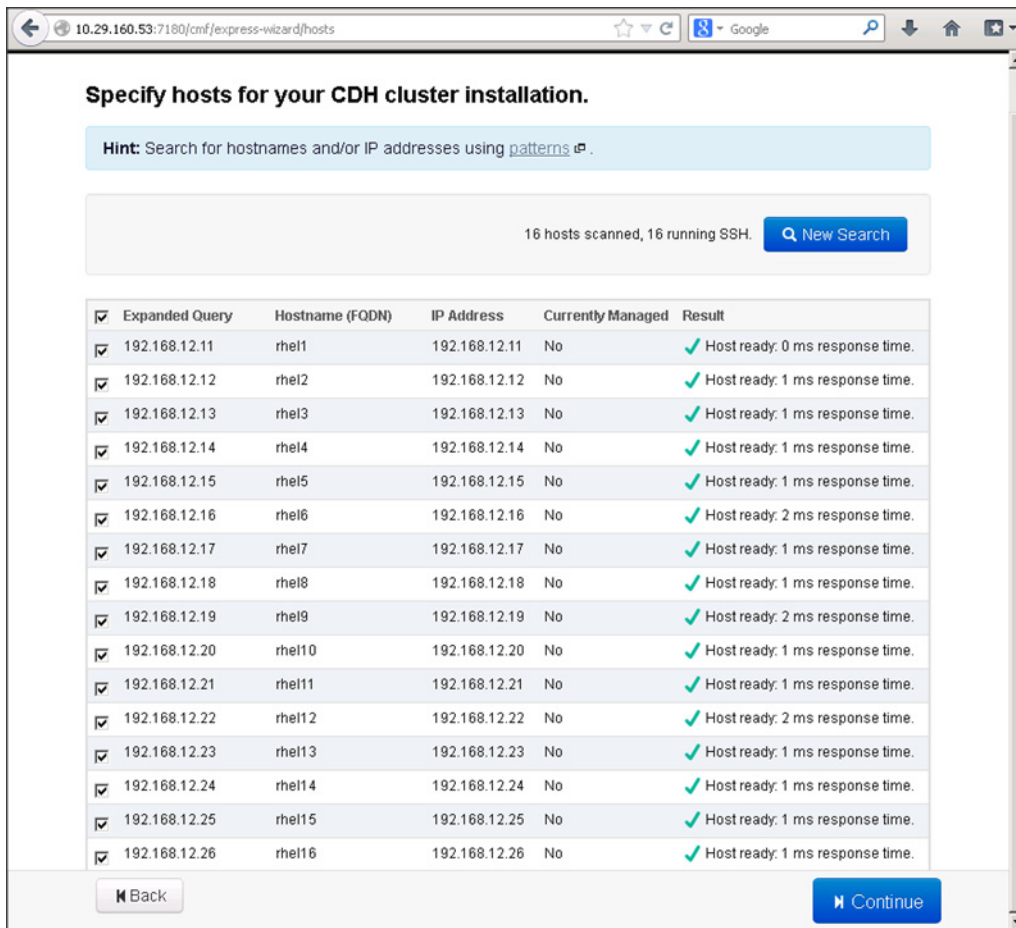
Specify hosts for your CDH cluster installation.



7. Cloudera Manager will “discover” the nodes in the cluster. Verify that all desired nodes have been found and selected for installation.

- Click **Install CDH On Selected Host**. CDH is Cloudera Distribution for Apache Hadoop.

Figure 112 Verifying and Selecting the Hosts



- For the method of installation, select the **Use Package** radio button.
- For the CDH version, select the **CDH4** radio button.
- For the specific release of CDH you want to install in your hosts, select Custom Repository radio button.
- Enter the following URL for the repository within the admin node.
`http://10.29.160.53/clouderarepo/cloudera-cdh4/`
- For the specific Impala release, select the **Custom Repository** radio button.
- Enter the following URL for the repository within the admin node.
`http://10.29.160.53/clouderarepo/cloudera-impala`
- For the specific release of Cloudera Manager, select the **Custom Repository** radio button.
- Enter the URL for the repository within the admin node.
`http://10.29.160.53/clouderarepo/cloudera-manager/`

Figure 113 **Selecting the CDH Version**

Choose Method:

- Use Packages
- Use Parcels (Recommended)

Select the version of CDH

- CDH4
- CDH3

Select the specific release of CDH you want to install on your hosts.

- Latest Release of CDH4
- CDH 4.2.1
- CDH 4.2.0
- CDH 4.1.4
- CDH 4.1.3
- CDH 4.1.2
- CDH 4.1.1
- CDH 4.1.0
- CDH 4.0.1
- CDH 4.0.0
- Custom Repository

Example for SLES, Redhat or other RPM based distributions:

http://archive.cloudera.com/cdh4/redhat/5/x86_64/cdh/4/

Example for Ubuntu or other Debian based distributions:

deb http://archive.cloudera.com/cdh4/ubuntu/lucid/amd64/cdh/ lucid-cdh contrib

Select the specific release of Impala you want to install on your hosts.

- Latest Release of Impala
- Custom Repository
- None

Note: Impala is supported only on CDH 4.1 or later deployments.

Select the specific release of the Cloudera Manager Agent you want to install on your hosts.

- Matched release for this Cloudera Manager server
- Custom Repository

Example for SLES, Redhat or other RPM based distributions:

http://archive.cloudera.com/cm4/redhat/5/x86_64/cm/4/

Example for Ubuntu or other Debian based distributions:

deb http://archive.cloudera.com/cm4/ubuntu/lucid/amd64/cm/ lucid-cm4 contrib

Enter a custom URL for the location of the GPG signing key (applies to all custom repositories and without Internet access).

Custom GPG Key URL:

Example for SLES, Redhat or other RPM based distributions:

http://archive.cloudera.com/redhat/cdh/RP#-GPG-KEY-cloudera

Example for Ubuntu or other Debian based distributions:

http://archive.cloudera.com/debian/archive.key

◀ Back

1 2 3 4 5

▶ Continue

17. Provide SSH login credentials for the cluster and click **Start Installation**.

Figure 114 Login Credentials to Start CDH Installation

Provide SSH login credentials.

Root access to your hosts is required to install the Cloudera packages. This installer will connect to your hosts via SSH and log in either directly as root or as another user with password-less sudo privileges to become root.

Login to all hosts as: root
 Another User:

You may connect via password or public-key authentication for the user selected above.

Authentication Method: All hosts accept same password
 All hosts accept same private key

Enter Password:

Confirm Password:

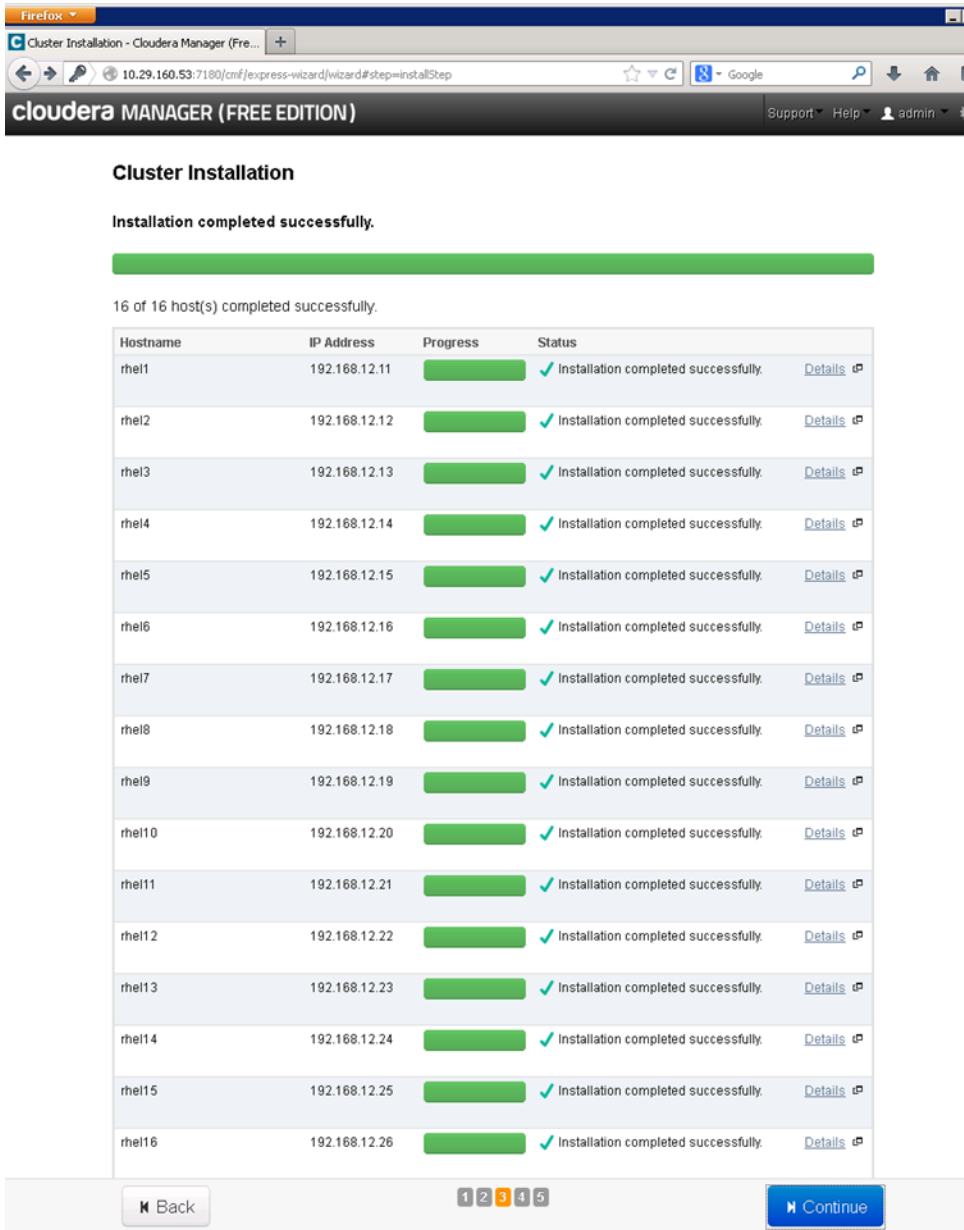
SSH Port:

Number of simultaneous installations:
(Running a large number of installations at once can consume large amounts of network bandwidth and other system resources)

[Back](#) [Start Installation](#)

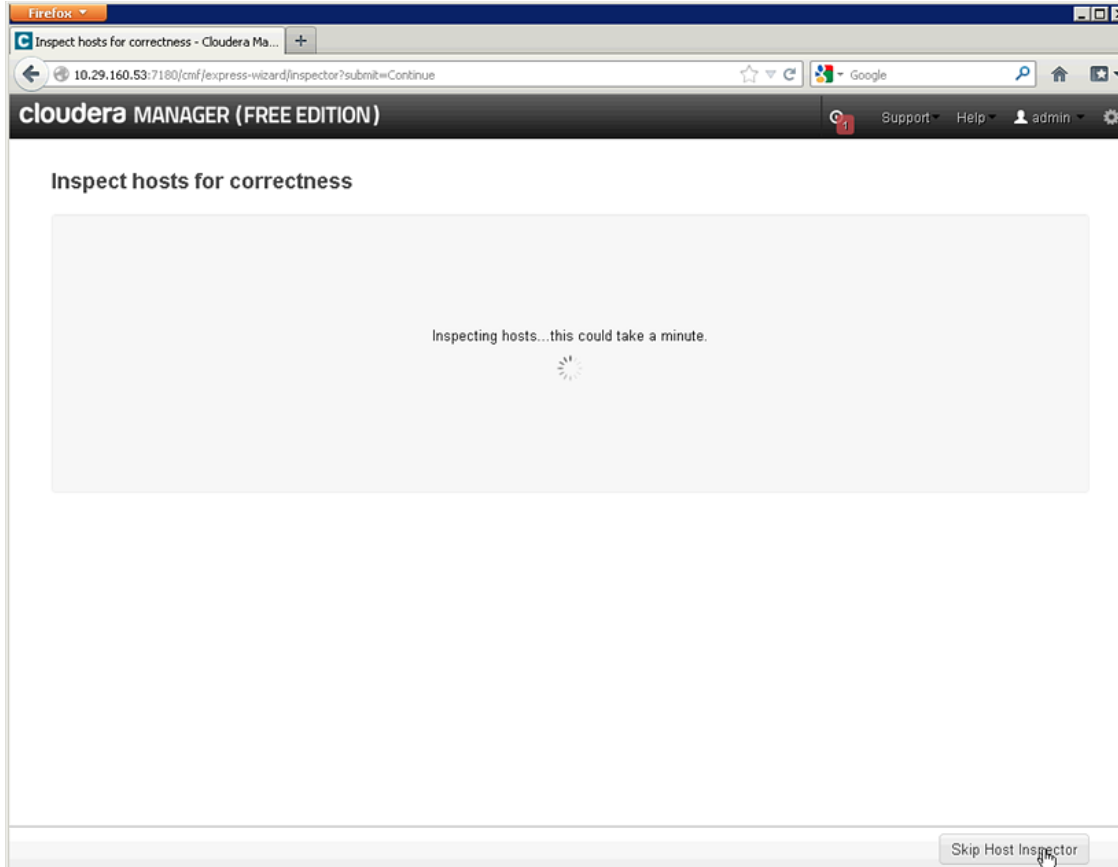
18. Make sure the installation across all the hosts is complete.
19. After the installation is complete, click **Continue**.

Figure 115 Installation in Progress



20. Wait for Cloudera Manager to inspect the hosts on which it has just performed the installation.

Figure 116 *Inspecting Hosts*



21. Review and verify the summary. Click **Continue**.

Figure 117 Inspecting Hosts for Correctness Part 1

The screenshot shows the Cloudera Manager (FREE EDITION) interface. At the top, the browser address bar shows the URL: 10.29.160.53:7180/cm/express-wizard/wizard#step=hostInspectorStep. Below the browser title bar, the page header reads "cloudera MANAGER (FREE EDITION)".

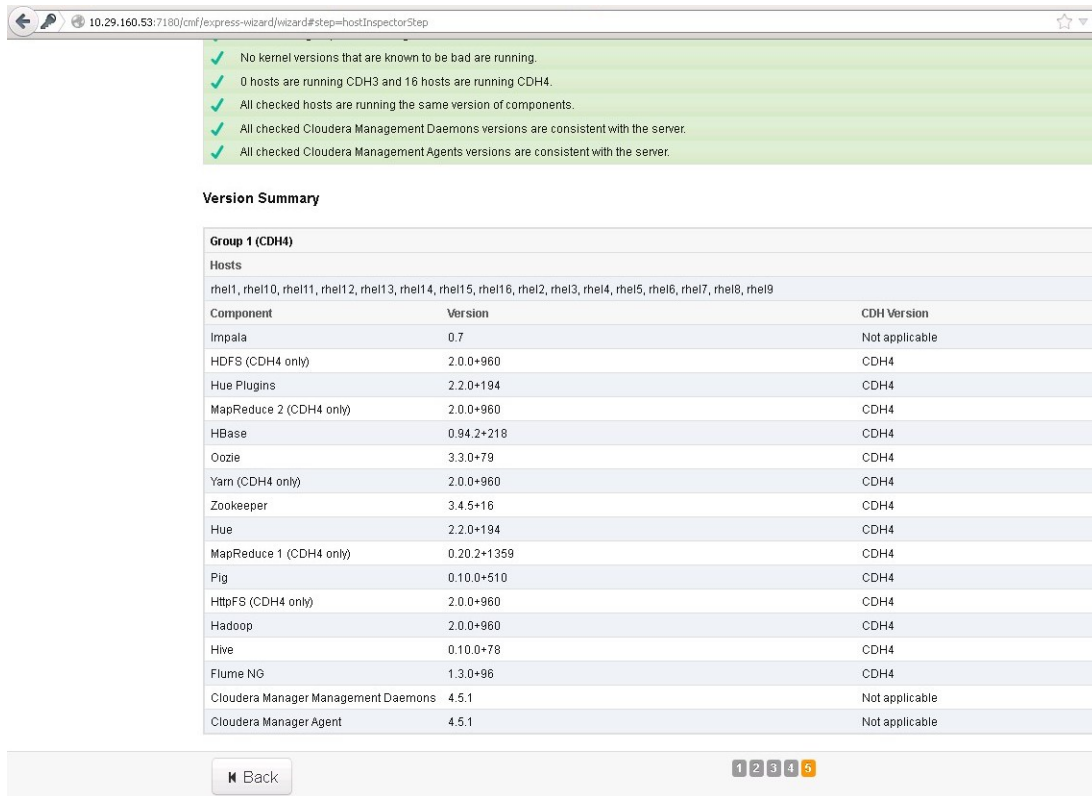
The main content area is titled "Cluster Installation". Underneath, there is a section "Inspect hosts for correctness" with a "Run Again" button.

Below this is a "Validations" section, which contains a list of 14 green checkmarks indicating successful checks:

- Inspector ran on all 16 hosts.
- Individual hosts resolved their own hostnames correctly.
- No errors were found while looking for conflicting init scripts.
- No errors were found while checking /etc/hosts.
- All hosts resolved localhost to 127.0.0.1.
- All hosts checked resolved each other's hostnames correctly.
- Host clocks are approximately in sync (within ten minutes).
- Host time zones are consistent across the cluster.
- No users or groups are missing.
- No kernel versions that are known to be bad are running.
- 0 hosts are running CDH3 and 16 hosts are running CDH4.
- All checked hosts are running the same version of components.
- All checked Cloudera Management Daemons versions are consistent with the server.
- All checked Cloudera Management Agents versions are consistent with the server.

Below the validations is a "Version Summary" section. It shows a table for "Group 1 (CDH4)".

Group 1 (CDH4)		
Hosts		
rhe11, rhe10, rhe11, rhe12, rhe13, rhe14, rhe15, rhe16, rhe2, rhe3, rhe4, rhe5, rhe6, rhe7, rhe8, rhe9		
Component	Version	CDH Version
<div style="display: flex; justify-content: space-between; align-items: center;"> ⏪ Back 1 2 3 4 5 </div>		

Figure 118 *Inspecting Hosts for Correctness Part2*


10.29.160.53:7180/cmf/express-wizard/wizard#step=hostInspectorStep

- ✓ No kernel versions that are known to be bad are running.
- ✓ 0 hosts are running CDH3 and 16 hosts are running CDH4.
- ✓ All checked hosts are running the same version of components.
- ✓ All checked Cloudera Management Daemons versions are consistent with the server.
- ✓ All checked Cloudera Management Agents versions are consistent with the server.

Version Summary

Group 1 (CDH4)

Hosts
rhe11, rhe110, rhe111, rhe112, rhe113, rhe114, rhe115, rhe116, rhe12, rhe13, rhe14, rhe15, rhe16, rhe17, rhe18, rhe19

Component	Version	CDH Version
Impala	0.7	Not applicable
HDFS (CDH4 only)	2.0.0+960	CDH4
Hue Plugins	2.2.0+194	CDH4
MapReduce 2 (CDH4 only)	2.0.0+960	CDH4
HBase	0.94.2+218	CDH4
Oozie	3.3.0+79	CDH4
Yarn (CDH4 only)	2.0.0+960	CDH4
Zookeeper	3.4.5+16	CDH4
Hue	2.2.0+194	CDH4
MapReduce 1 (CDH4 only)	0.20.2+1359	CDH4
Pig	0.10.0+510	CDH4
HttpFS (CDH4 only)	2.0.0+960	CDH4
Hadoop	2.0.0+960	CDH4
Hive	0.10.0+78	CDH4
Flume NG	1.3.0+96	CDH4
Cloudera Manager Management Daemons	4.5.1	Not applicable
Cloudera Manager Agent	4.5.1	Not applicable

Back 1 2 3 4 5

22. select services that need to be started on the cluster.

Figure 119 *Selecting CDH Version and Services*

Choose the CDH4 services that you want to install on your cluster.

Choose a combination of services to install.

- Core Hadoop**
HDFS, MapReduce, Oozie, Hive, and Hue
- Core with Real-Time Delivery**
HDFS, MapReduce, ZooKeeper, HBase, Oozie, Hive, and Hue
- Core with Real-Time Query (Beta)**
HDFS, MapReduce, Impala, Oozie, Hive, and Hue
- All Services**
HDFS, MapReduce, ZooKeeper, HBase, Impala (Beta), Oozie, Hive, and Hue
The current Impala release is beta software and not recommended for use in production.
- Custom Services**
Choose your own services. Services required by chosen services must also be selected. Note that Flume service can be added after your initial cluster has been set up.

[Inspect Role Assignments](#)[Continue](#)

23. This is one of the critical steps in the installation. Inspect and customize the role assignments of all the nodes based on your requirements and click **Continue**.

Figure 120 Reviewing the Role Assignments Part1

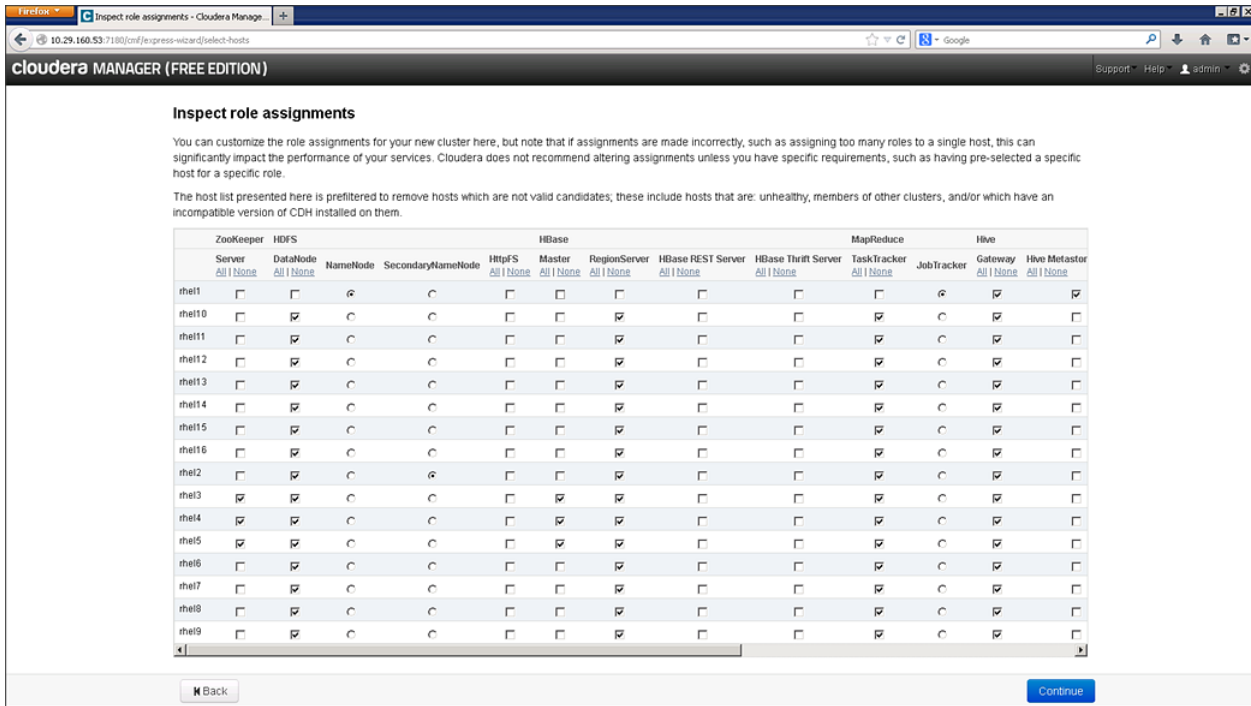
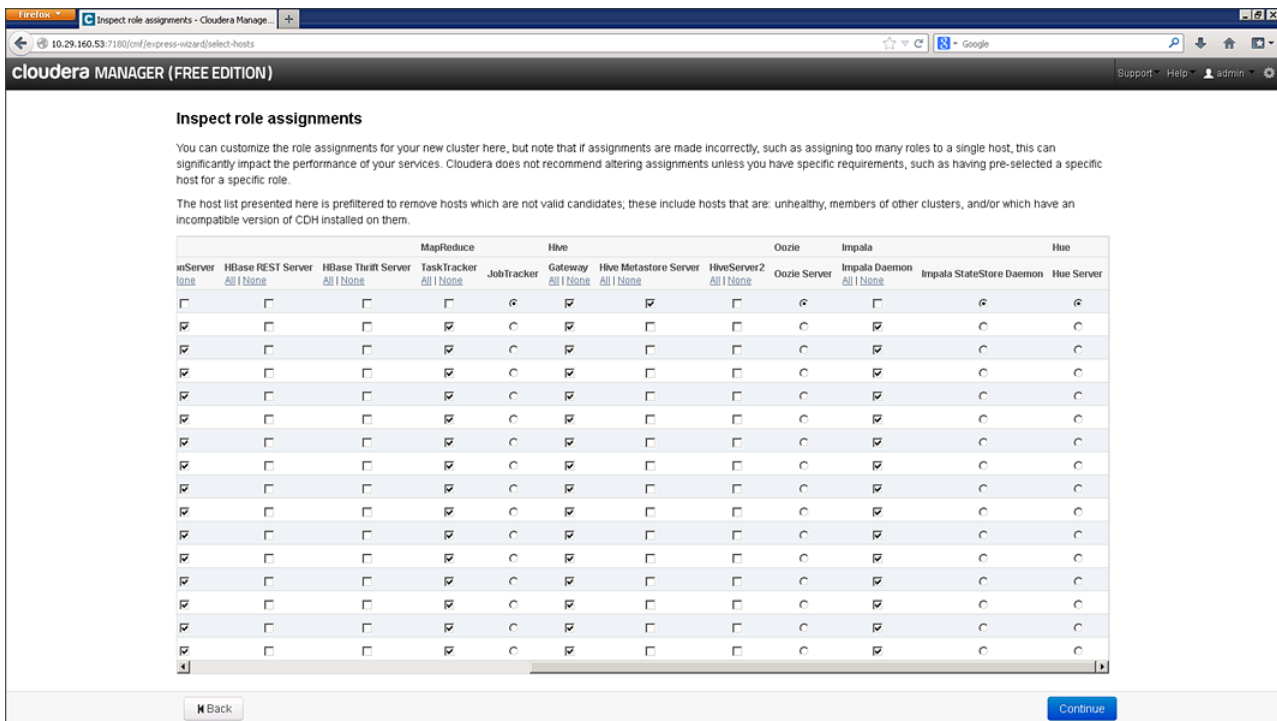


Figure 121 Reviewing the Role Assignments Part2



Scaling the Cluster

The role assignment recommendation above is for clusters of up to 16 servers. For clusters of 16 to 64 nodes the recommendation is to dedicate one server for name node and a second server for secondary name node and job tracker. For larger clusters larger than 64 nodes the recommendation is to dedicate one server each for name node, secondary name node and job tracker.

1. Select the **Use Embedded Database** radio button.
2. Click **Test Connection** and click **Continue**.

Figure 122 Database Setup

Database Setup

On this page you configure and test database connections. If using custom databases, create the databases first according to the **Installing and Configuring an External Database** section of the [Installation Guide](#).

When using the Embedded Database, passwords are auto generated. Please copy them down.

Use Embedded Database
 Use Custom Databases

Hive ✔ Skipped. Will create database in later step.

Database Host Name:	Database Type:	Database Name :	Username:	Password:
rhel1:7432	PostgreSQL	hive	hive	p09MaGRbeb

3. Review and customize the configuration changes based on your requirements

Figure 123 **Reviewing the Configuration Changes Part1**

Review configuration changes

Set the following configuration values for your new role(s). Required values are marked with *.

Group	Parameter	Recommended Value	Description
Service hbase1			
Service-Wide	HDFS Root Directory* hbase.rootdir	/hbase default value	The HDFS directory shared by HBase RegionServers.
Service hdfs1			
DataNode (Base) Show Members	DataNode Data Directory* dfs.datanode.data.dir	<input type="text" value="/CDH/sdb1/dfs/dn"/> <input type="text" value="/CDH/sdc1/dfs/dn"/> <input type="text" value="/CDH/sdd1/dfs/dn"/> <input type="text" value="/CDH/sde1/dfs/dn"/> <input type="text" value="/CDH/sdf1/dfs/dn"/> <input type="text" value="/CDH/sdg1/dfs/dn"/> <input type="text" value="/CDH/sdh1/dfs/dn"/> <input type="text" value="/CDH/sdi1/dfs/dn"/> <input type="text" value="/CDH/sdj1/dfs/dn"/> <input type="text" value="/CDH/sdk1/dfs/dn"/> <input type="text" value="/CDH/sdl1/dfs/dn"/> <input type="text" value="/CDH/sdm1/dfs/dn"/>	Comma-delimited list of directories on the local file system where the DataNode stores HDFS block data. Typical values are /dataN/dfs/dn for N= 1, 2, 3... These directories should be mounted using the noatime option and the disks should be configured using JBOD. RAID is not recommended.
		<input type="button" value="Back"/>	<input type="button" value="Continue"/>

Figure 124 **Reviewing the Configuration Changes Part2**

DataNode (Base) Show Members	DataNode Failed Volumes Tolerated dfs.datanode.failed.volumes.tolerated	<input type="text" value="11"/> Reset to the default value: 0	The number of volumes that are allowed to fail before a DataNode stops offering service. By default, any volume failure will cause a DataNode to shutdown.
NameNode (Base) Show Members	NameNode Data Directories* dfs.name.node.name.dir	<input type="text" value="/CDH/sdc1/dfs/nn"/> <input type="text" value="/CDH/sdd1/dfs/nn"/> Reset to empty default value	Determines where on the local file system the NameNode should store the name table (fsimage). For redundancy, enter a comma-delimited list of directories to replicate the name table in all of the directories. Typical values are /dataN/dfs/nn where N=1..3.
SecondaryNameNode (Base) Show Members	HDFS Checkpoint Directory* dfs.name.node.checkpoint.dir	<input type="text" value="/CDH/sdb1/dfs/snn"/> Reset to empty default value	Determines where on the local file system the DFS SecondaryNameNode should store the temporary images to merge. For redundancy, enter a comma-delimited list of directories to replicate the image in all of the directories. Typical values are /dataN/dfs/snn for N= 1, 2, 3...
Service hive1			
Service-Wide	Hive Warehouse Directory hive.metastore.warehouse.dir	/user/hive/warehouse default value	Hive warehouse directory is the location in HDFS where Hive's tables are stored. Note that Hive's default value for its warehouse directory is 'user/hive/warehouse'.
Hive Metastore Server (Base) Show Members	Hive Metastore Server Port hive.metastore.port	9083 default value	Port on which Hive Metastore Server will listen for connections.
Service mapreduce1			
JobTracker (Base) Show Members	JobTracker Local Data Directory* mapred.local.dir	<input type="text" value="/CDH/sdc1/mapred/jt"/> Reset to empty default value	Directory on the local filesystem where the JobTracker stores job configuration data. Directories that do not exist are ignored. A single directory is sufficient; a list of multiple directories will not cause problems.
TaskTracker (Base) Show Members	TaskTracker Local Data Directory List* mapred.local.dir	<input type="text" value="/CDH/sdb1/mapred/local"/> <input type="text" value="/CDH/sdc1/mapred/local"/> <input type="text" value="/CDH/sdd1/mapred/local"/> <input type="text" value="/CDH/sde1/mapred/local"/> <input type="text" value="/CDH/sdf1/mapred/local"/> <input type="text" value="/CDH/sdq1/mapred/local"/>	List of directories on the local filesystem where a TaskTracker stores intermediate data files. To spread disk I/O, enter a comma-separated list of directories on different devices. Directories that do not exist are ignored. Typical values are /dataN/mapred/local for N= 1, 2, 3...
		<input type="button" value="Back"/>	<input type="button" value="Continue"/>

Figure 125 *Reviewing the Configuration Changes Part3*

/CDH/sdk1/mapred/local	⊕ ⊖		
/CDH/sdl1/mapred/local	⊕ ⊖		
/CDH/sdm1/mapred/local	⊕ ⊖		
/CDH/sdn1/mapred/local	⊕ ⊖		
/CDH/sdo1/mapred/local	⊕ ⊖		
/CDH/sdp1/mapred/local	⊕ ⊖		
/CDH/sdq1/mapred/local	⊕ ⊖		
/CDH/sdr1/mapred/local	⊕ ⊖		
/CDH/sds1/mapred/local	⊕ ⊖		
/CDH/sdt1/mapred/local	⊕ ⊖		
/CDH/sdu1/mapred/local	⊕ ⊖		
/CDH/sdv1/mapred/local	⊕ ⊖		
/CDH/sdw1/mapred/local	⊕ ⊖		
/CDH/sdx1/mapred/local	⊕ ⊖		
Reset to empty default value ↶			
Service oozie1			
Oozie Server (Base) Show Members Ⓞ	Oozie Server Data Directory	<code>/var/lib/oozie/data</code> default value	Directory where the Oozie Server will place its data. Only applicable when using Derby as the database type.
Service zookeeper1			
Server (Base) Show Members Ⓞ	Data Directory dataDir	<code>/var/lib/zookeeper</code> default value	The disk location that ZooKeeper will use to store its database snapshots.
Server (Base) Show Members Ⓞ	Transaction Log Directory dataLogDir	<code>/var/lib/zookeeper</code> default value	The disk location that ZooKeeper will use to store its transaction logs.
⏪ Back		Continue	

4. Click **Continue** to start running the cluster services.

Figure 126 Starting the Cluster Services**Starting your cluster services.**

Completed 18 of 18 steps.

- ✓ Waiting for ZooKeeper Service to initialize
Finished waiting
- ✓ Starting ZooKeeper Service
Service started successfully.
- ✓ Checking if the name directories of the NameNode are empty. Formatting HDFS only if empty.
Successfully formatted NameNode.
- ✓ Starting HDFS Service
Service started successfully.
- ✓ Creating HDFS /tmp directory
Successfully created HDFS directory /tmp.
- ✓ Creating HBase root directory
Successfully created HBase root directory.
- ✓ Starting HBase Service
Service started, but only 17/18 roles started
- ✓ Starting MapReduce Service
Service started successfully.
- ✓ Creating Hive Metastore Database
Created Hive Metastore Database.
- ✓ Creating Hive Metastore Database Tables
Created Hive Metastore Database Tables successfully.
- ✓ Creating Hive warehouse directory
Successfully created Hive warehouse directory.
- ✓ Starting Hive Service
Service started successfully.
- ✓ Creating Oozie database

[Continue](#)

5. Hadoop services are installed, configured and now running on all the nodes of the cluster. Click **Continue** to complete the installation.

Figure 127 Installation Completion

10.29.160.53:7180/cm/express-wizard/finish

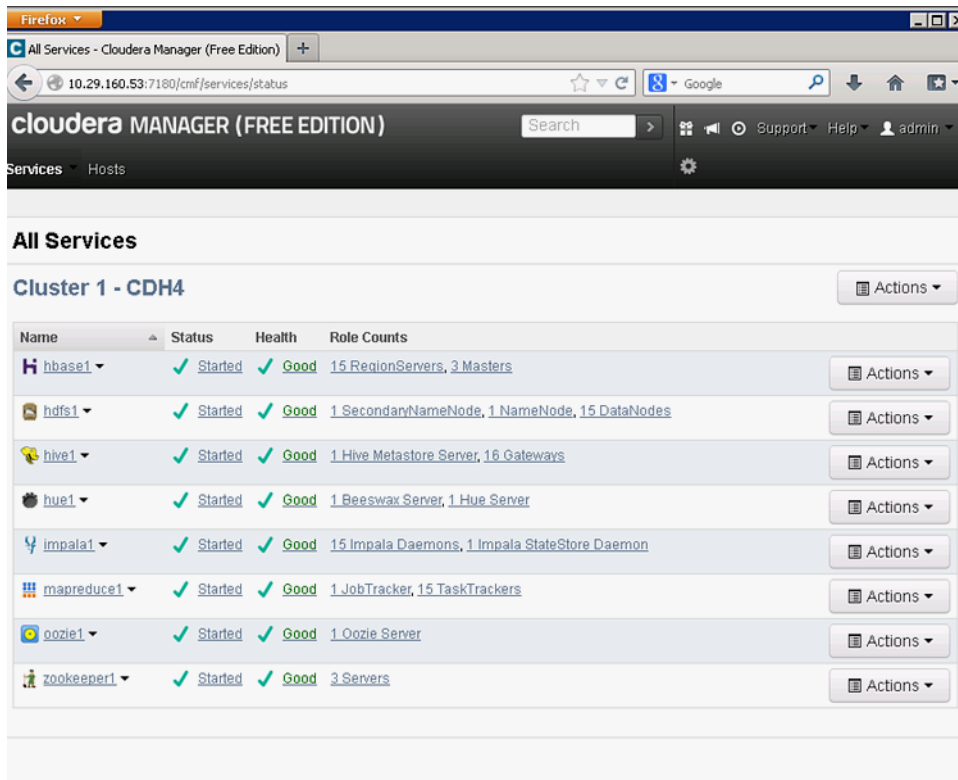
cloudera MANAGER (FREE EDITION)

Congratulations!

The Hadoop services are installed, configured, and running on your cluster.

6. Cloudera Manager will now show the status of all Hadoop services running on the cluster.

Figure 128 Service Status of the Cluster



Conclusion

Cisco UCS Common Platform Architecture for Big Data offers a comprehensive stack for enterprise Big Data deployments. Together, Cisco and Cloudera are well positioned to help organizations exploit the valuable business insights found in all their data, regardless of whether it’s structured, semi structured or unstructured. The solution offers industry-leading performance, scalability and advanced management capabilities while reducing the risks involved in Big Data deployments.

The rack level configuration detailed in the document can be extended to multiple rack scale. Up to 160 servers (10 racks) can be supported with no additional switching in a single UCS domain. Each additional rack requires two Cisco Nexus 2232PP 10GigE Fabric Extenders and 16 Cisco UCS C240 M3 Rack-Mount Servers. Scaling beyond 10 racks (160 servers) can be implemented by interconnecting multiple UCS domains using Nexus 6000/7000 Series switches, scalable to thousands of servers and to hundreds of petabytes storage, and managed from a single pane using [UCS Central](#).

Bill of Materials

See [Table 5](#) and [Table 6](#) provides BOM for the high performance rack and high capacity rack configuration respectively, [Table 7](#) provides BOM for rack and PDU. [Table 8](#) and [Table 9](#) provide BOM for RHEL and Cloudera software components used in the deployment model.

Table 5 *BOM for High Performance Rack*

Part Number	Description	Quantity
UCS-EZ-BD-HP	High Performance Rack	1
UCS-EZ-INFRA-FI96	UCS 6296 FI w/ 18p LIC, Cables Bundle	2 (included)
N2K-UCS2232PF	Nexus 2232PP with 16 FET (2 AC PS, 1 FAN (Std Airflow))	2 (included)
UCS-EZ-C240-2665	UCS C240 M3 SFF w/ 2665, 16x16GB, VIC 1225, 2PS	16 (included)
SFP-H10GB-CU3M=	10GBASE-CU SFP+ Cable 3 Meter	28 (included)

Table 6 *BOM for High Capacity Rack*

Part Number	Description	Quantity
UCS-EZ-BD-HC	High Capacity Rack	1
UCS-EZ-INFRA-FI96	UCS 6296 FI w/ 18p LIC, Cables Bundle	2 (included)
N2K-UCS2232PF	Nexus 2232PP with 16 FET (2 AC PS, 1 FAN (Std Airflow))	2 (included)
UCS-EZ-C240-2640	UCS C240 M3 LFF w/ 2640, 16x16GB, VIC 1225, 2PS	16 (included)
SFP-H10GB-CU3M=	10GBASE-CU SFP+ Cable 3 Meter	28 (included)

Table 7 *BOM for Rack and PDUs*

Part Number	Description	Quantity
RACK-UCS2	Cisco R42610 standard rack w/side panels	1
RP208-30-1P-U-2=	Cisco RP208-30-U-2 Single Phase PDU 20x C13 4x C19 (Country Specific)	2
CON-UCW3-RPDUX	UC PLUS 24X7X4 Cisco RP208-30-U-X Single Phase PDU 2x (Country Specific)	2

Table 8 *Red Hat Enterprise Linux License*

Red Hat Enterprise Linux	Description	Quantity
RHEL-2S-1G-3A	Red Hat Enterprise Linux	16
CON-ISV1-RH2S1G1A	Support for Red Hat Enterprise Linux	16

Table 9 *Cloudera Software License*

Software	Description	Quantity
CECO-2407	Cloudera Enterprise Core CDH	16