




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

Building a 64-Node Hadoop Cluster

Last Updated: October 25, 2013



Building Architectures to Solve Business Problems



About the Authors



Raghunath Nambiar

Raghunath Nambiar, 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.



Ajay Singh

Ajay Singh, Hortonworks

Ajay Singh is Director, Technology Alliances at Hortonworks. Ajay is responsible for design & validation of ecosystem solutions to optimally integrate, deploy & operate Hortonworks Data Platform.



Manankumar Trivedi

Manankumar Trivedi, 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.



Karthik Kulkarni

Karthik Kulkarni, Cisco Systems

Karthik Kulkarni is a Technical Marketing Engineer at Cisco Data Center Business Group focusing on Big Data and Hadoop technologies.

Acknowledgment

The authors acknowledge contributions of Ashwin Manjunatha, and Sindhu Sudhir in developing the Cisco UCS Common Platform Architecture (CPA) for Big Data with Hortonworks Cisco Validated Design.



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

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

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 Hortonworks

Audience

This document describes the architecture and deployment procedures of Hortonworks Data Platform (HDP) on a 64 node cluster based Cisco UCS Common Platform Architecture (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 HDP on the Cisco UCS CPA for Big Data.

Introduction

Hadoop has become a strategic data platform embraced by mainstream enterprises as it offers the fastest path for businesses to unlock value in big data while maximizing existing investments. The Hortonworks Data Platform (HDP) is a 100% open source distribution of Apache Hadoop that is truly enterprise grade having been built, tested and hardened with enterprise rigor. The combination of HDP and Cisco UCS provides industry-leading platform for Hadoop based applications.

Cisco UCS Common Platform Architecture for Big Data

The Cisco UCS solution for HDP is based on [Cisco 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**—provide high-bandwidth, low-latency connectivity for servers, with integrated, unified management provided for all connected devices by Cisco UCS Manager. Deployed in redundant pairs, Cisco fabric interconnects offer the full active-active redundancy, performance, and exceptional scalability needed to support the large number of nodes that are typical in clusters serving Big Data applications. Cisco UCS Manager enables rapid and consistent server configuration using service profiles and automation of the ongoing system



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

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

maintenance activities such as firmware updates across the entire cluster as a single operation. Cisco UCS Manager also offers advanced monitoring with options to raise alarms and send notifications about the health of the entire cluster.

- **Cisco UCS 2200 Series Fabric Extenders**—extends the network into each rack, acting as remote line cards for fabric interconnects and providing highly scalable and extremely cost-effective connectivity for a large number of nodes.
- **Cisco UCS C-Series Rack-Mount Servers**—Cisco UCS C240M3 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)**—the unique Cisco UCS Virtual Interface Cards incorporate next-generation converged network adapter (CNA) technology from Cisco, and offer dual 10Gbps 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**—resides within the Cisco UCS 6200 Series Fabric Interconnects. It makes the system self-aware and self-integrating, managing 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.

Hortonworks Data Platform (HDP)

The Hortonworks Data Platform (HDP) is an enterprise-grade, hardened Apache Hadoop distribution that enables you to store, process, and manage large data sets.

Apache Hadoop is an open-source software framework that allows for the distributed processing of large data sets across clusters of computers using simple programming models. It is designed for high-availability and fault-tolerance, and can scale from a single server up to thousands of machines.

The Hortonworks Data Platform combines the most useful and stable versions of Apache Hadoop and its related projects into a single tested and certified package. Hortonworks offers the latest innovations from the open source community, along with the testing and quality you expect from enterprise-quality software.

The Hortonworks Data Platform is designed to integrate with and extend the capabilities of your existing investments in data applications, tools, and processes. With Hortonworks, you can refine, analyze, and gain business insights from both structured and unstructured data – quickly, easily, and economically.

Hortonworks - Key Features and Benefits

With the Hortonworks Data Platform, enterprises can retain and process more data, join new and existing data sets, and lower the cost of data analysis. Hortonworks enables enterprises to implement the following data management principles:

- **Retain as much data as possible**—Traditional data warehouses age, and over time will eventually store only summary data. Analyzing detailed records is often critical to uncovering useful business insights.
- **Join new and existing data sets**—Enterprises can build large-scale environments for transactional data with analytic databases, but these solutions are not always well suited to processing nontraditional data sets such as text, images, machine data, and online data. Hortonworks enables enterprises to incorporate both structured and unstructured data in one comprehensive data management system.
- **Archive data at low cost**—It is not always clear what portion of stored data will be of value for future analysis. Therefore, it can be difficult to justify expensive processes to capture, cleanse, and store that data. Hadoop scales easily, so you can store years of data without much incremental cost, and find deeper patterns that your competitors may miss.
- **Access all data efficiently**—Data needs to be readily accessible. Apache Hadoop clusters can provide a low-cost solution for storing massive data sets while still making the information readily available. Hadoop is designed to efficiently scan all of the data, which is complimentary to databases that are efficient at finding subsets of data.
- **Apply data cleansing and data cataloging**—Categorize and label all data in Hadoop with enough descriptive information (metadata) to make sense of it later, and to enable integration with transactional databases and analytic tools. This greatly reduces the time and effort of integrating with other data sets, and avoids a scenario in which valuable data is eventually rendered useless.
- **Integrate with existing platforms and applications**—There are many business intelligence (BI) and analytic tools available, but they may not be compatible with your particular data warehouse or DBMS. Hortonworks connects seamlessly with many leading analytic, data integration, and database management tools.

The Hortonworks Data Platform is the foundation for the next-generation enterprise data architecture – one that addresses both the volume and complexity of today’s data.

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 Cluster Configuration**—offers a balance of compute power with IO bandwidth optimized for price and performance. It is built using Cisco UCS C240M3 Rack-Mount Servers powered by two Intel Xeon E5-2665 processors (16 cores) with 256 GB of memory and 24 1TB SFF disk drives.
- **High Capacity Cluster Configuration**—optimized for low cost per terabyte, is built using Cisco UCS C240M3 Rack-Mount Servers powered by two Intel Xeon E5-2640 processors (12 cores) with 128GB memory and 12 3TB LFF disk drives.



Note

This CVD describes the installation process for a 64-node High Performance Cluster configuration.

The High Performance Cluster configuration consists of the following:

- Two Cisco UCS 6296UP Fabric Interconnects
- Eight Cisco Nexus 2232PP Fabric Extenders (two per rack)
- 64 Cisco UCS C240M3 Rack-Mount Servers (16 per rack)

- Four Cisco R42610 standard racks
- Eight vertical power distribution units (PDU) (country specific)

Rack and PDU Configuration

Each rack consists of two vertical PDU. The master rack consists of two Cisco UCS 6296UP Fabric Interconnects, two Cisco Nexus 2232PP Fabric Extenders and sixteen Cisco UCS C240M3 Servers, connected to each of the vertical PDUs for redundancy; thereby, ensuring availability during power source failure. The expansion racks also consists of 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, similar to master rack.



Note

Contact your Cisco representative for country specific information.

Table 1 and Table 2 describe the rack configurations of rack 1 (master rack) and racks 2-4 (expansion racks).

Table 1 Rack Configuration For The Master Rack (Rack-1)

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

Table 1 *Rack Configuration For The Master Rack (Rack-1)*

Cisco 42U Rack	Master Rack
8	Cisco UCS C240M3
7	
6	Cisco UCS C240M3
5	
4	Cisco UCS C240M3
3	
2	Cisco UCS C240M3
1	

Table 2 *Rack Configuration for the Expansion Rack (Racks 2-4)*

Cisco 42U Rack	Master Rack
42	Unused
41	Unused
40	Unused
39	Unused
38	Cisco Nexus FEX 2232PP
37	Cisco Nexus FEX 2232PP
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	

Cisco 42U Rack	Master Rack
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 Cisco UCS C240M3 Rack Server is equipped with Intel Xeon E5-2665 processors, 256 GB of memory, Cisco UCS Virtual Interface Card (VIC)1225, LSI MegaRAID SAS 9266-8i storage controller and 24 x 1TB 7.2K Serial Advance Technology Attachment (SATA) disk drives.

Figure 1 illustrates the ports on the Cisco Nexus 2232PP fabric extender connecting to the Cisco UCS C240M3 servers. Sixteen Cisco UCS C240M3 servers are used in the master rack configurations.

Figure 1 Fabric Topology





Figure 2 illustrates the port connectivity between the Cisco Nexus 2232PP FEX and the Cisco UCS C240M3 server.

Figure 2 Connectivity Diagram of Cisco Nexus 2232PP FEX and Cisco UCS C240M3 Servers



For more information on physical connectivity and 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

Figure 3 depicts a 64-node cluster, and each link represents 8 x 10 Gigabit link.

Figure 3 64 -Node Cluster Configuration



Software Distributions and Versions

Hortonworks Data Platform (HDP)

The Hortonworks Data Platform supported is HDP 1.3. For more information, see: <http://www.hortonworks.com>

RHEL

The Operating System supported is Red Hat Enterprise Linux Server 6.2. For more information on the Linux support, see: www.redhat.com.

Software Versions

Table 3 describes the software versions tested and validated in this document.

Table 3 Software Versions Summary

Layer	Component	Version or Release
Compute	Cisco UCS C240M3	1.4.7cc

Table 3 **Software Versions Summary**

Layer	Component	Version or Release
Network	Cisco UCS 6296UP	UCS 2.1(1e)
	Cisco UCS VIC1225 Firmware	2.1(1a)
	Cisco UCS VIC1225 Driver	2.1.1.41
	Cisco Nexus 2232PP	5.1(3)N2(2.11a)
Storage	LSI 9266-8i Firmware	23.7.0-0039
	LSI 9266-8i Driver	06.504.01.00
Software	Red Hat Enterprise Linux Server	6.2 (x86_64)
	Cisco UCS Manager	2.1(1e)
	Hortonworks Data Platform	1.3

**Note**

To download the latest drivers, see:

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

Fabric Configuration

This section provides details for configuring a fully redundant, highly available Cisco UCS 6296 Fabric Interconnect.

1. Initial setup of the Fabric Interconnect A and B.
2. Connect to IP address of Fabric Interconnect A using web browser.
3. Launch the Cisco UCS Manager.
4. Edit the chassis discovery policy.
5. Enable server and uplink ports.
6. Create pools and polices for service profile template.
7. Create Cisco Service Profile template and 64 service profiles.
8. Start discover process.
9. Associate to server.

Performing Initial Setup of Cisco UCS 6296 Fabric Interconnects

This section describes the steps to perform the initial setup of the Cisco UCS 6296 Fabric Interconnects A and B.

Configure Fabric Interconnect A

Follow these steps to configure the Fabric Interconnect A:

1. Connect to the console port on the first Cisco UCS 6296 Fabric Interconnect.

2. At the prompt to enter the configuration method, enter console to continue.
3. If asked to either perform 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.
12. Enter the Mgmt0 IPv4 netmask.
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 enter yes to save the configuration.
20. Wait for the login prompt to make sure the configuration has been saved.

Configure Fabric Interconnect B

Follow these steps to configure the Fabric Interconnect B:

1. Connect to the console port on the second Cisco UCS 6296 Fabric Interconnect.
2. When prompted to enter the configuration method, enter console to continue.
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 that was configured for the first Fabric Interconnect.
5. Enter the Mgmt0 IPv4 address.
6. Enter yes to save the configuration.
7. Wait for the login prompt to confirm that the configuration has been saved.

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 login to Cisco UCS Manager.

1. Open a Web browser and navigate to the Cisco UCS 6296 Fabric Interconnect cluster address.
2. Click the **Launch** link to download the Cisco UCS Manager software.

3. If prompted, accept the security certificates.
4. When prompted, enter the username as admin and the administrative password.
5. Click **Login**.

Upgrading UCSM Software to Version 2.1(1e)

This document assumes the uses of Cisco UCS 2.1(1e). Make sure that the Cisco UCS C-Series version 2.1(1e) software bundle is installed on the Cisco UCS Fabric Interconnects.

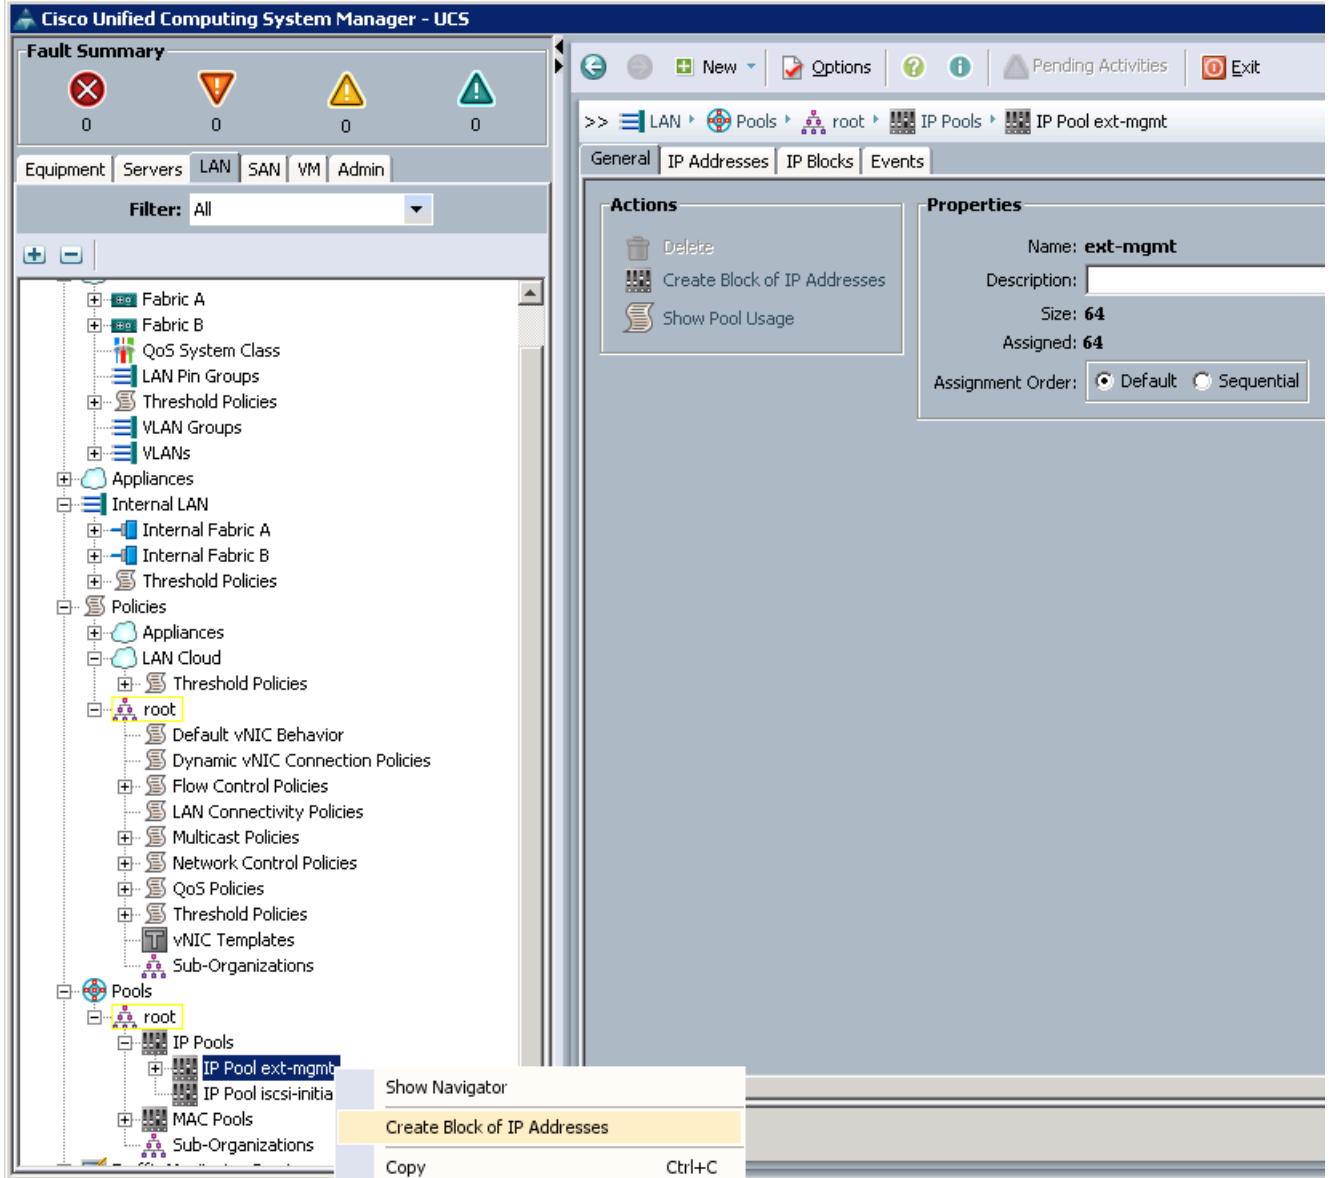
To upgrade the Cisco UCS Manager software and Cisco UCS 6296 Fabric Interconnect software to version 2.1(1e), see: [Upgrading Cisco UCS from Release 2.0 to Releases 2.1](#)

Adding Block of IP Addresses for KVM Access

Follow these steps to create a block of KVM IP addresses for the server access in Cisco UCS environment.

1. Click the **LAN** tab.
2. Select **Pools > IPPools > IP Pool ext-mgmt**.
3. Right-click **Management IP Pool**.
4. Select **Create Block of IP Addresses** as shown in [Figure 4](#)

Figure 4 Adding Block of IP Addresses for KVM Access Part 1



5. Enter the starting IP address of the block and number of IPs needed, the subnet and the gateway information as shown in Figure 5.

Figure 5 Adding Block of IP Addresses for KVM Access Part 2

Create Block of IP Addresses

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 Secondary DNS: 0.0.0.0

OK Cancel

- Click **OK** to create the IP Address block as shown in [Figure 6](#).

Figure 6 Adding Block of IP Addresses for KVM Access Part 3

Create Block of IP Addresses

From: 10.29.160.120 Size: 64

Subnet Mask: 255.255.255.0 Default Gateway: 10.29.160.1

Primary DNS: 0.0.0.0 Secondary DNS: 0.0.0.0

OK Cancel

- Click **OK**.

Editing The Chassis Discovery Policy

This section provides details for modifying the chassis discovery policy. Setting the discovery policy ensures easy addition of the Cisco UCS B-Series chassis or fabric extenders for the Cisco UCS C-Series servers in future.

1. Click the **Equipment** tab.
2. In the right pane, click the **Policies** tab.
3. Click the **Global Policies** tab.
4. In the Chassis/FEX Discovery Policy area, select 8-link from the drop-down list for Action field as shown in [Figure 7](#).

Figure 7 Changing The Chassis Discovery Policy

The screenshot displays the Cisco UCS Management Center interface. On the left, the 'Equipment' tree is expanded to show 'Fabric Interconnects'. The main pane shows the 'Chassis/FEX Discovery Policy' configuration. The 'Action' dropdown menu is set to '8 Link', which is highlighted with a red circle and the number '1'. Other policy settings include 'Link Grouping Preference' (None selected), 'Rack Server Discovery Policy' (Action: Immediate selected), 'Rack Management Connection Policy' (Action: Auto Acknowledged selected), 'Power Policy' (Redundancy: N+1 selected), 'MAC Address Table Aging' (Aging Time: Mode Default selected), and 'Global Power Allocation Policy' (Allocation Method: Policy Driven Chassis Group Cap selected).

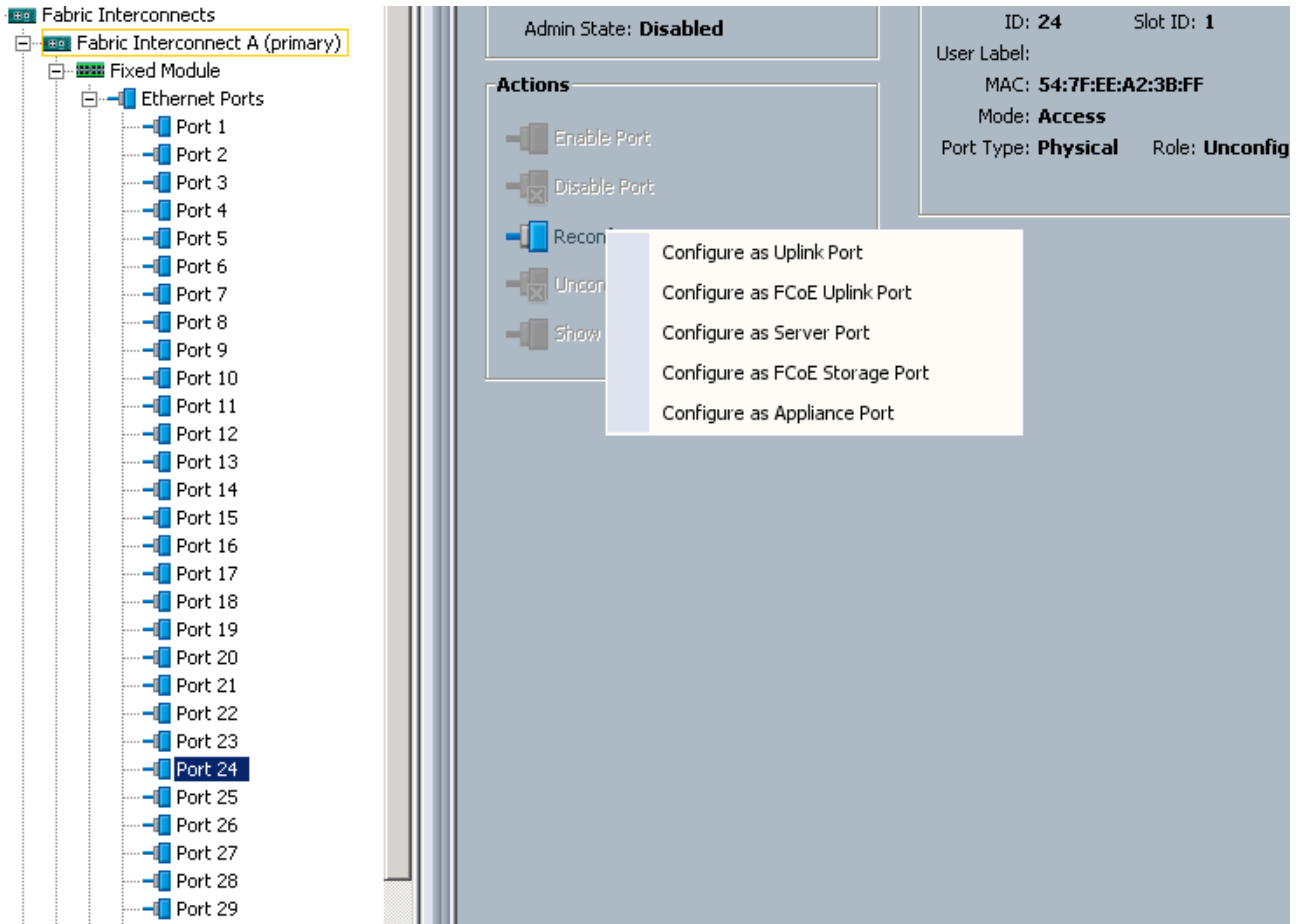
5. Click **Save Changes**.
6. Click **OK**.

Enabling The Server Ports and Uplink Ports

Follow these steps to enable the server and configure the uplink ports:

1. Click the **Equipment** tab.
2. Select **Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module**.
3. Expand the Unconfigured Ethernet Ports.
4. Select all the ports that are connected to the Cisco 2232PP FEX (eight per FEX), right-click and select **Reconfigure > Configure as a Server Port**.
5. Select port 1 that is connected to the uplink switch, right-click, then select **Reconfigure > Configure as Uplink Port**.
6. Select **Show Interface** and select 10GB for Uplink Connection.
7. Click **Yes** and then **OK** to continue.
8. Select **Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module**.
9. Expand the Unconfigured Ethernet Ports section.
10. Select all the ports that are connected to the Cisco 2232 Fabric Extenders (eight per Fex), right-click and select **Reconfigure > Configure as Server Port**.
11. Click **Yes** and then **OK** to continue.
12. Select port number 1, which is connected to the uplink switch, right-click and select **Reconfigure > Configure as Uplink Port**.

Figure 8 Enabling Server Ports



13. Select **Show Interface** and select 10GB for Uplink Connection.
14. Click **Yes** and then **OK** to continue.

Figure 9 shows all the configured uplink and Server ports.

Figure 9 Server and Uplink Ports Summary

Port	Primary ID	Secondary ID	MAC Address	Role
Port 1	1	1	54:7F:EE:A2:3B:E8	Network
Port 2	1	2	54:7F:EE:A2:3B:E9	Server
Port 3	1	3	54:7F:EE:A2:3B:EA	Server
Port 4	1	4	54:7F:EE:A2:3B:EB	Server
Port 5	1	5	54:7F:EE:A2:3B:EC	Server
Port 6	1	6	54:7F:EE:A2:3B:ED	Server
Port 7	1	7	54:7F:EE:A2:3B:EE	Server
Port 8	1	8	54:7F:EE:A2:3B:EF	Server
Port 9	1	9	54:7F:EE:A2:3B:F0	Server
Port 10	1	10	54:7F:EE:A2:3B:F1	Server
Port 11	1	11	54:7F:EE:A2:3B:F2	Server
Port 12	1	12	54:7F:EE:A2:3B:F3	Server
Port 13	1	13	54:7F:EE:A2:3B:F4	Server
Port 14	1	14	54:7F:EE:A2:3B:F5	Server
Port 15	1	15	54:7F:EE:A2:3B:F6	Server
Port 16	1	16	54:7F:EE:A2:3B:F7	Server
Port 17	1	17	54:7F:EE:A2:3B:F8	Server
Port 18	1	18	54:7F:EE:A2:3B:F9	Server
Port 19	1	19	54:7F:EE:A2:3B:FA	Server
Port 20	1	20	54:7F:EE:A2:3B:FB	Server
Port 21	1	21	54:7F:EE:A2:3B:FC	Server
Port 22	1	22	54:7F:EE:A2:3B:FD	Server
Port 23	1	23	54:7F:EE:A2:3B:FE	Server
Port 24	1	24	54:7F:EE:A2:3B:FF	Server
Port 25	1	25	54:7F:EE:A2:3C:00	Server
Port 26	1	26	54:7F:EE:A2:3C:01	Server
Port 27	1	27	54:7F:EE:A2:3C:02	Server
Port 28	1	28	54:7F:EE:A2:3C:03	Server
Port 29	1	29	54:7F:EE:A2:3C:04	Server
Port 30	1	30	54:7F:EE:A2:3C:05	Server
Port 31	1	31	54:7F:EE:A2:3C:06	Server
Port 32	1	32	54:7F:EE:A2:3C:07	Server

Creating Pools for Service Profile Templates

Creating an Organization

Organizations are used as a means to arrange and restrict access to various groups within the IT organization, and enable multi-tenancy of the compute resources. This document does not use organizations; however, the steps to create an organizations are given for future reference.

Follow these steps to configure an organization in the Cisco UCS Manager:

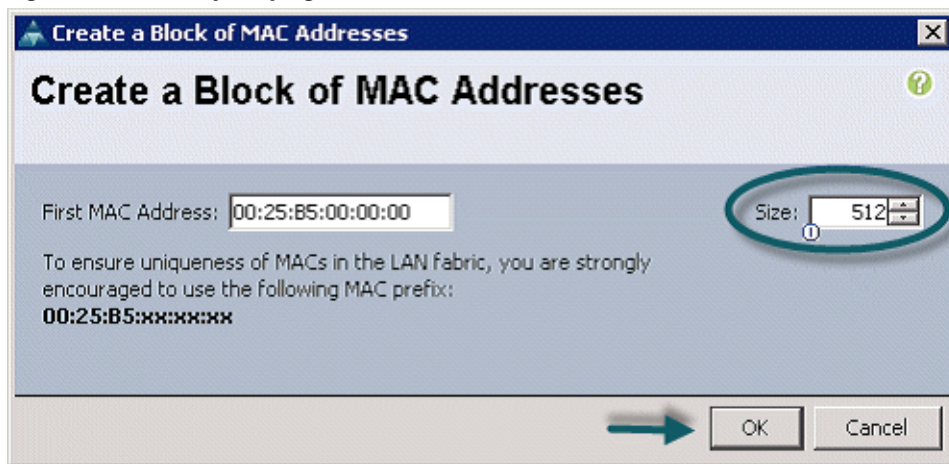
1. Click **New** in the left corner of the UCS Manager GUI.
2. Select **Create Organization** from the options.
3. Enter a name for the organization.
4. (Optional) Enter a description for the organization.
5. Click **OK**.

Creating MAC Address Pools

Follow these steps to create MAC address pools:

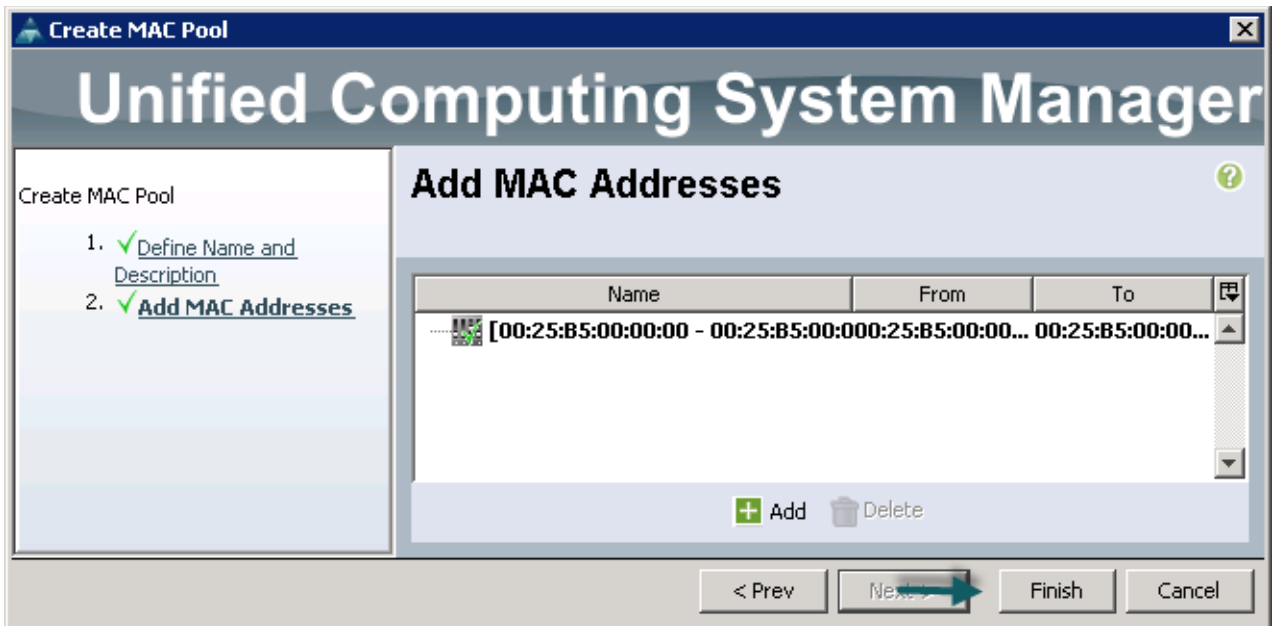
1. Click the **LAN** tab.
2. Select **Pools > root**.
3. Right-click **MAC Pools** under the root organization.
4. Select **Create MAC Pool** to create the MAC address pool. Enter ucs as the name of the MAC pool.
5. (Optional) Enter a description of the MAC pool.
6. Click **Next**.
7. Click **Add**.
8. Specify a starting MAC address.
9. Specify a size of the MAC address pool, which is sufficient to support the available server resources as shown in [Figure 10](#).
10. Click **OK**.

Figure 10 Specifying the First MAC Address and Size



11. Click **Finish** as shown in [Figure 11](#).

Figure 11 Adding MAC Addresses



12. Click **OK** to confirm the addition of the MAC addresses.

Configuring VLANs

Table 4 describes the VLANs that are configured in this design solution.

Table 4 VLAN Configurations

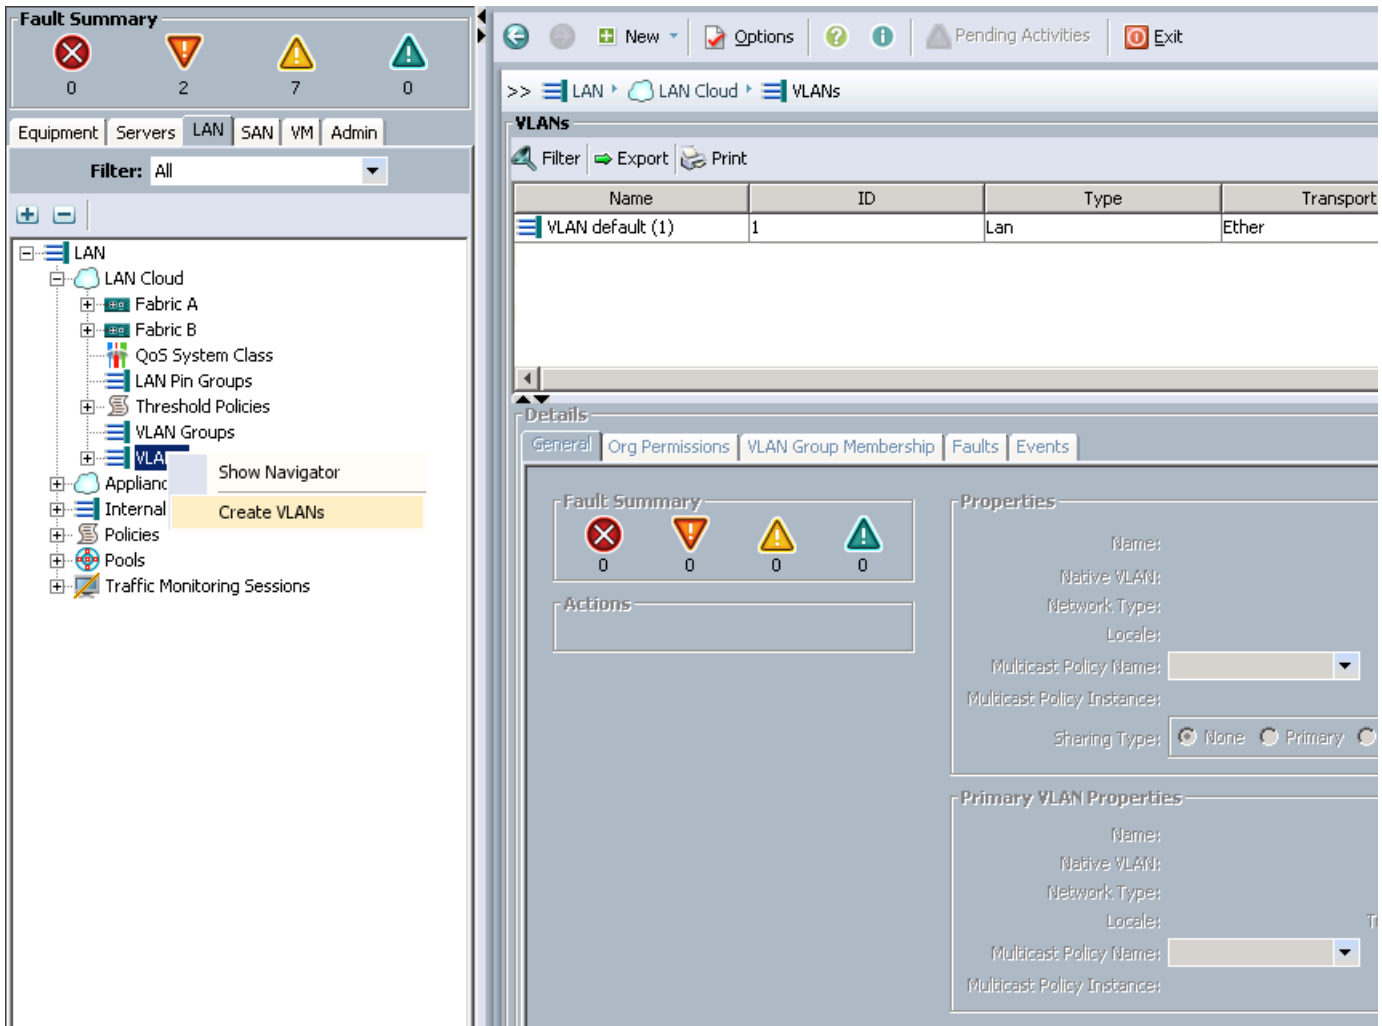
VLAN	Fabric	NIC Port	Function	Failover
vlan160_mgmt	A	eth0	Management, user connectivity	Fabric Failover B
vlan12_HDFC	B	eth1	Hadoop	Fabric Failover A
vlan11_DATA	A	eth2	Hadoop and/or SAN/NAS access, ETL	Fabric Failover B

All of the VLANs created should be trunked to the upstream distribution switch connecting the fabric interconnects. In this deployment, 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 secondary interconnect and/or SAN/NAS access, heavy ETL, and so on.

Follow these steps to configure VLANs in Cisco UCS Manager:

1. Click the **LAN** tab.
2. Select **LAN > VLANs**.
3. Right-click the VLANs under the root organization.
4. Select **Create VLANs** to create the VLAN as shown in Figure 12.

Figure 12 Creating VLAN



5. Enter vlan160_mgmt in the VLAN Name/Prefix text box as shown in Figure 13.
6. Click the **Common/Global** radio button.
7. Enter 160 in the VLAN IDs text box.
8. Click **OK** and then click **Finish**.
9. Click **OK**.

Figure 13 Creating Management VLAN

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")

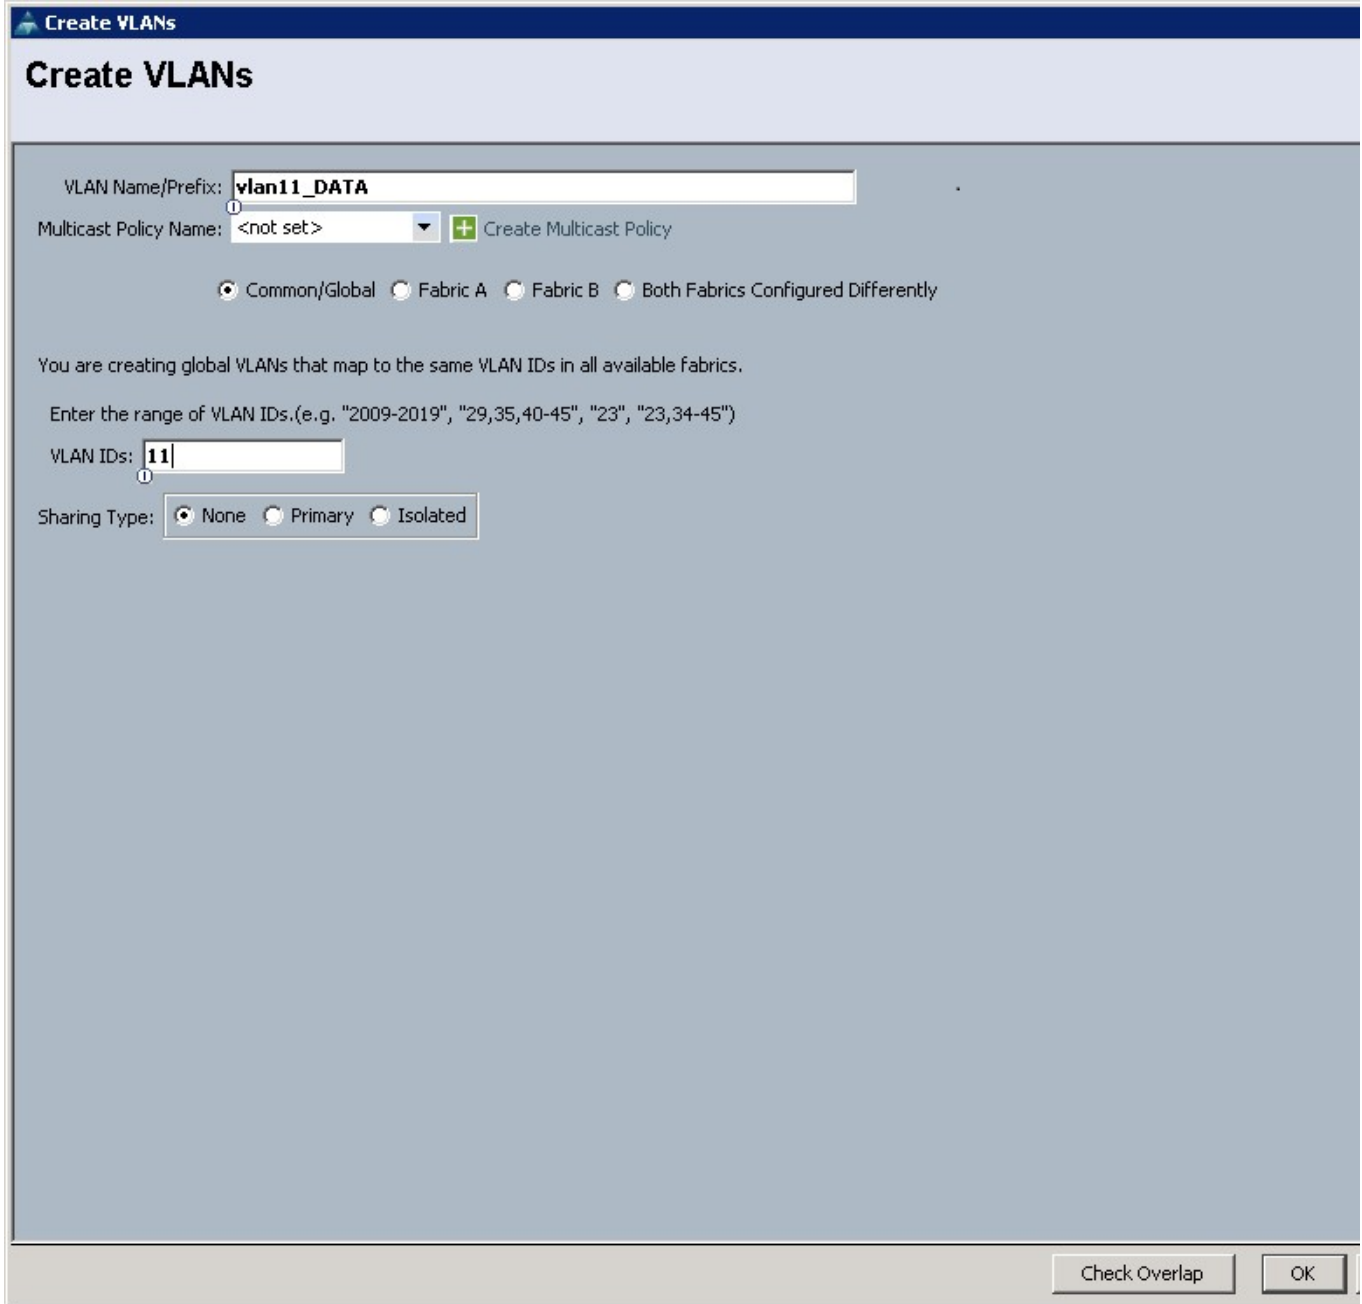
VLAN IDs:

Sharing Type:
 None
 Primary
 Isolated

10. Click the **LAN** tab.
11. Select **LAN > VLANs**.
12. Right-click the VLANs under the root organization.
13. Select **Create VLANs** to create the VLAN as shown in [Figure 14](#).
14. Enter `vlan11_DATA` in the VLAN Name/Prefix text box.
15. Click the **Common/Global** radio button.
16. Enter 11 in the VLAN IDs text box.

17. Click **OK** and then click **Finish**.
18. Click **OK**.

Figure 14 *Creating VLAN for Data*



19. Click the **LAN** tab.
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` in the VLAN Name/Prefix text box as shown in [Figure 15](#).

24. Click the **Common/Global** radio button.
25. Enter 12 in the VLAN IDs text box.
26. Click **OK** and then click **Finish**.

Figure 15 *Creating VLAN for Hadoop Data*

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

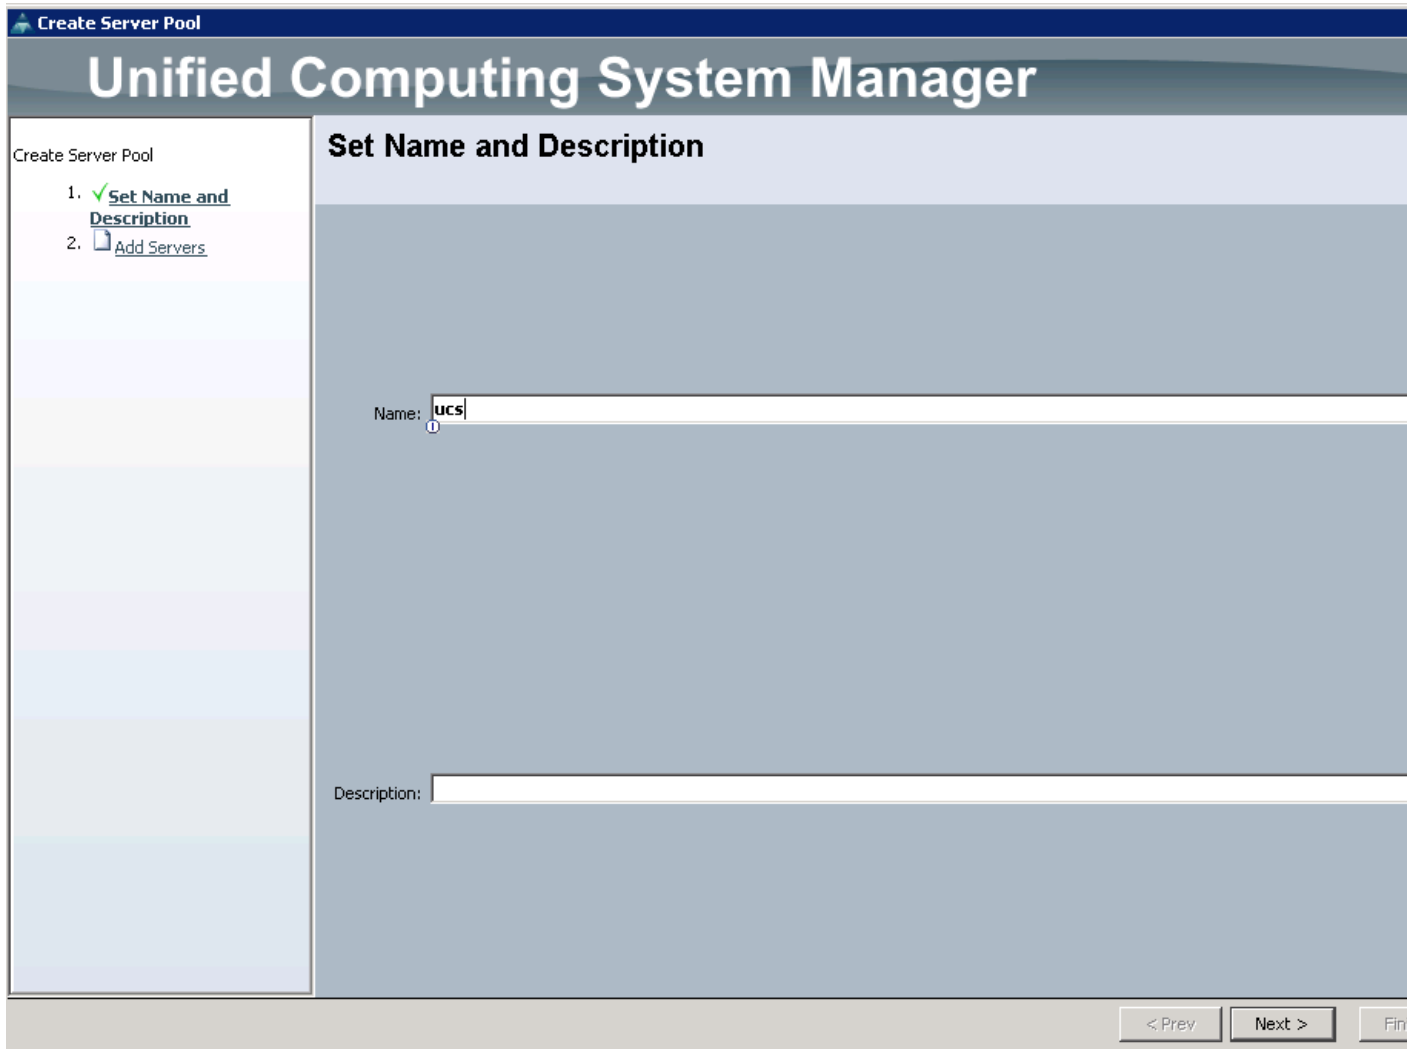
Creating a Server Pool

A server pool contains a set of servers. These servers typically share the same characteristics such as their location in the chassis, 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 the server pool policies and server pool policy qualifications to automate the server assignment.

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

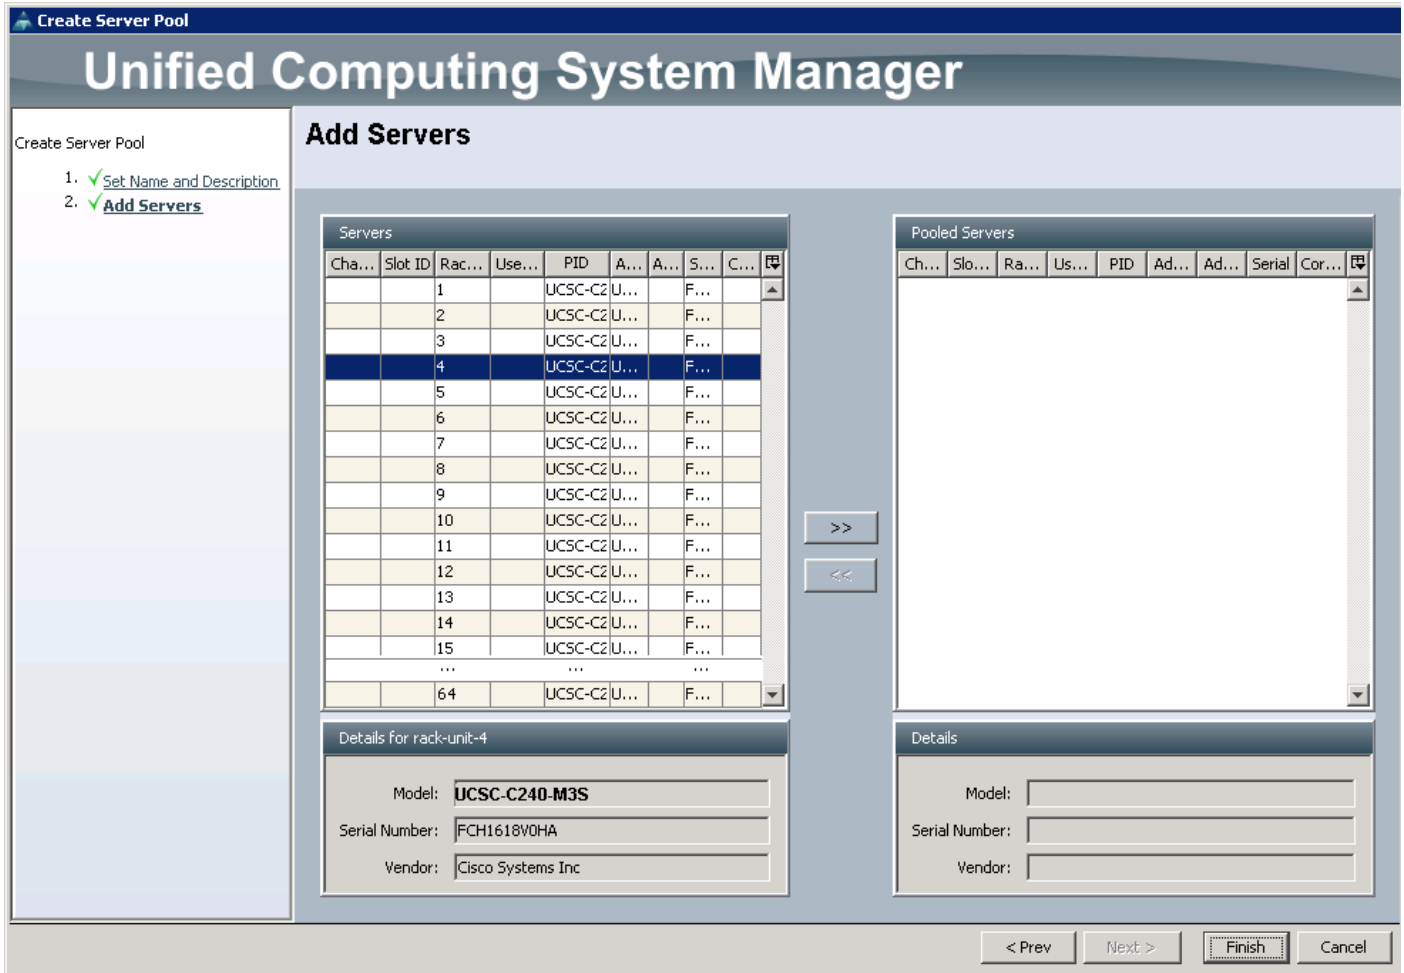
1. Click the **Servers** tab.
2. Select **Pools > root**.
3. Right-click the **Server Pools**.
4. Select **Create Server Pool**.
5. Enter the required name (ucs) for the server pool in the name text box as shown in [Figure 16](#).
6. (Optional) Enter a description for the organization.

Figure 16 *Setting Name and Description of the Server Pool*



7. Click **Next** to add the servers.
8. Select all the Cisco UCS C240M3 servers to be added to the server pool that were previously created (ucs), then **Click >>** to add them to the pool as shown in [Figure 17](#).
9. Click **OK**, and then click **Finish**.

Figure 17 Adding Servers to the Server Pool



Creating Policies for Service Profile Templates

This section provides you the procedure to create the following policies for the service profile template:

- [Creating a Host Firmware Package Policy, page 32](#)
- [Creating QoS Policies, page 32](#)
- [Creating a Local Disk Configuration Policy, page 36](#)
- [Creating a Server BIOS Policy, page 37](#)
- [Creating a Boot Policy, page 41](#)

Creating a Host Firmware Package Policy

Firmware management policies allow the administrator to select the corresponding firmware packages for a given server configuration. The components that can be configured include adapters, BIOS, board controllers, FC adapters, HBA options, ROM and storage controller.

Follow these steps to create a host firmware management policy for a given server configuration using the Cisco UCS Manager:

1. Click the **Servers** tab in the UCS Manager.
2. Select **Policies > root**.
3. Right-click **Host Firmware Packages**.
4. Select **Create Host Firmware Package**.
5. Enter the required host firmware package name (ucs) as shown in [Figure 18](#).
6. Click the **Simple** radio button to configure the host firmware package.
7. Select the appropriate Rack Package value.
8. Click **OK** to complete creating the management firmware package.
9. Click **OK**.

Figure 18 *Creating Host Firmware Package*

The screenshot shows the 'Create Host Firmware Package' dialog box. The title bar reads 'Create Host Firmware Package'. The main title is 'Create Host Firmware Package'. There are two input fields: 'Name:' with the value 'ucs' and 'Description:' which is empty. Below these is a question: 'How would you like to configure the Host Firmware Package?' with two radio buttons: 'Simple' (selected) and 'Advanced'. At the bottom, there are two dropdown menus: 'Blade Package:' set to '<not set>' and 'Rack Package:' set to '2.1(1e)'.

Creating QoS Policies

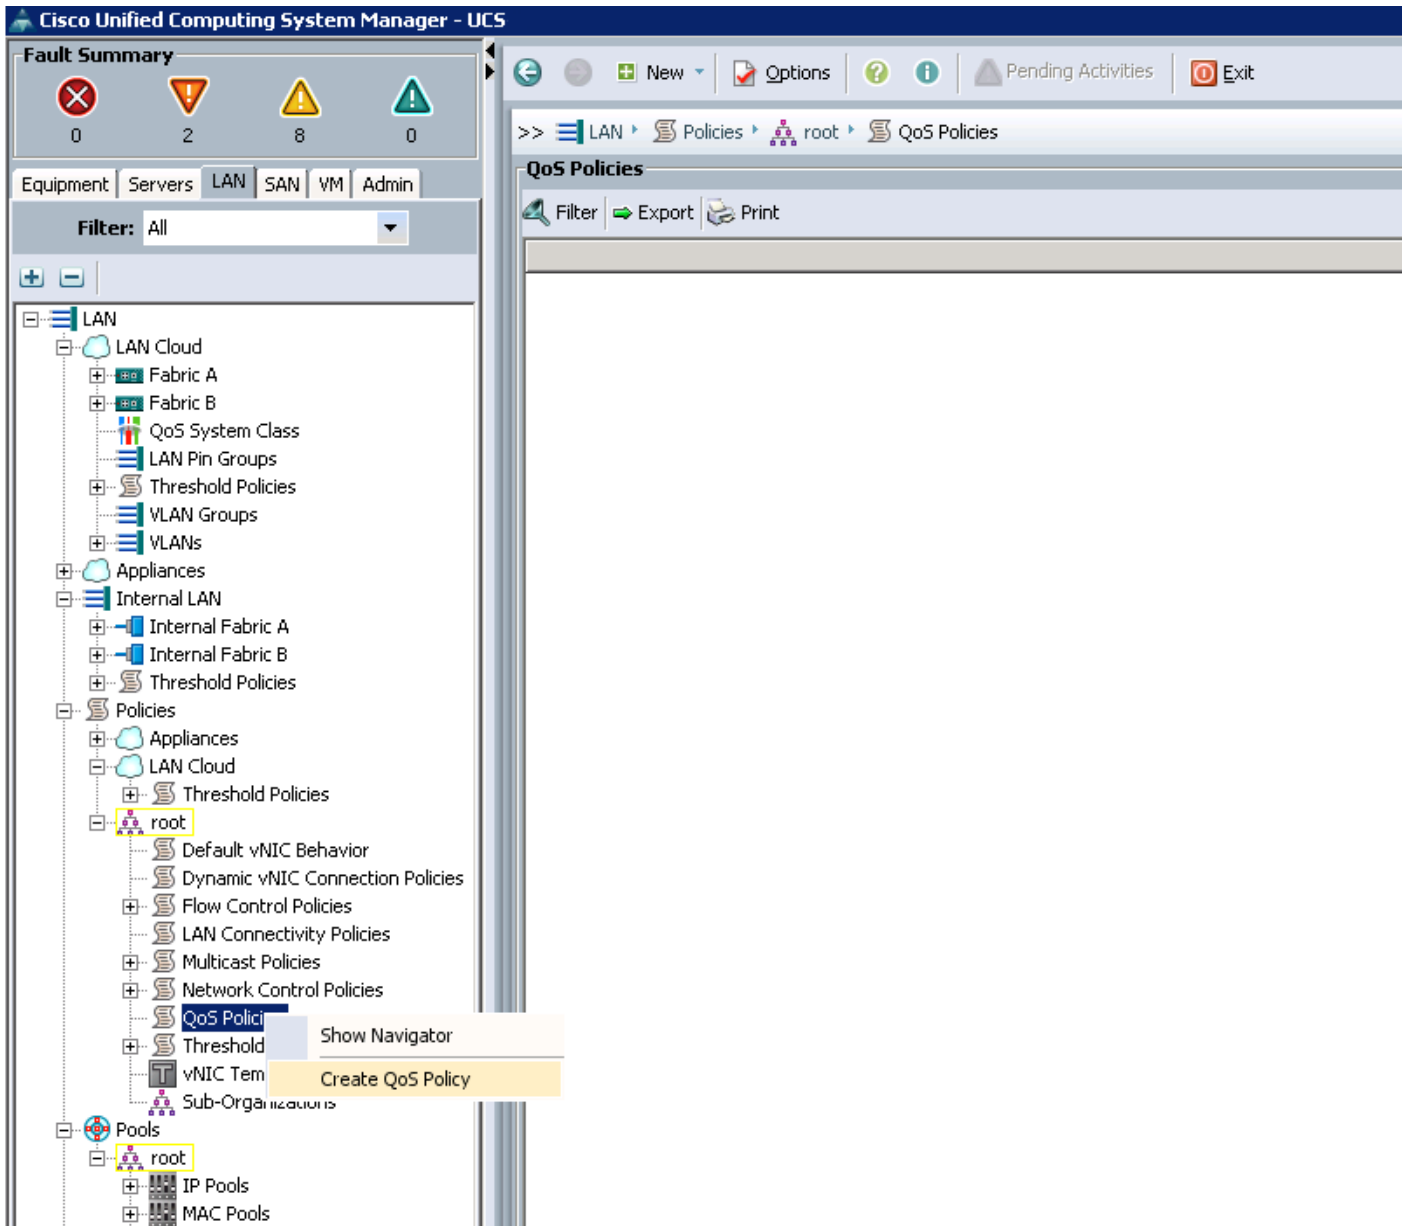
This section describes the procedure to create the Best Effort QoS Policy and Platinum QoS policy.

Creating the Best Effort Policy

Follow these steps to create the Best Effort Policy:

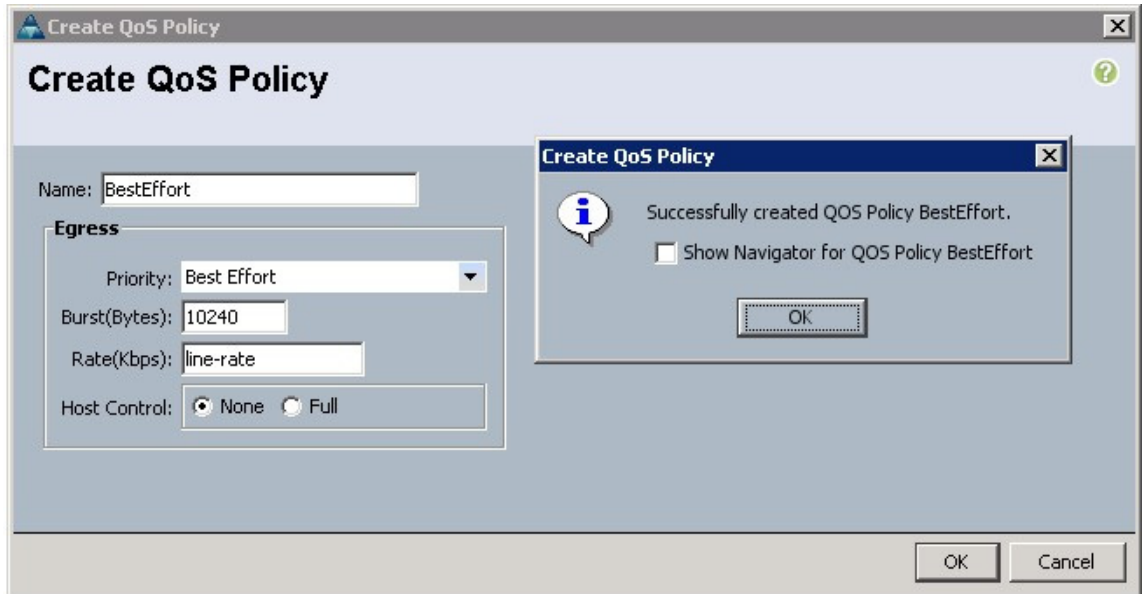
1. Click the **LAN** tab.
2. Select **Policies > root**.
3. Right-click **QoS Policies**.
4. Select **Create QoS Policy** as shown in [Figure 19](#).

Figure 19 Creating QoS Policy



5. Enter BestEffort as the name of the policy as shown in Figure 20.
6. Select **BestEffort** from the drop down menu.
7. Keep the Burst (Bytes) field as default (10240).
8. Keep the Rate (Kbps) field as default (line-rate).
9. Keep the **Host Control** radio button as default (none).
10. Click **OK** to complete creating the Policy.
11. Click **OK**.

Figure 20 Creating BestEffort QoS Policy

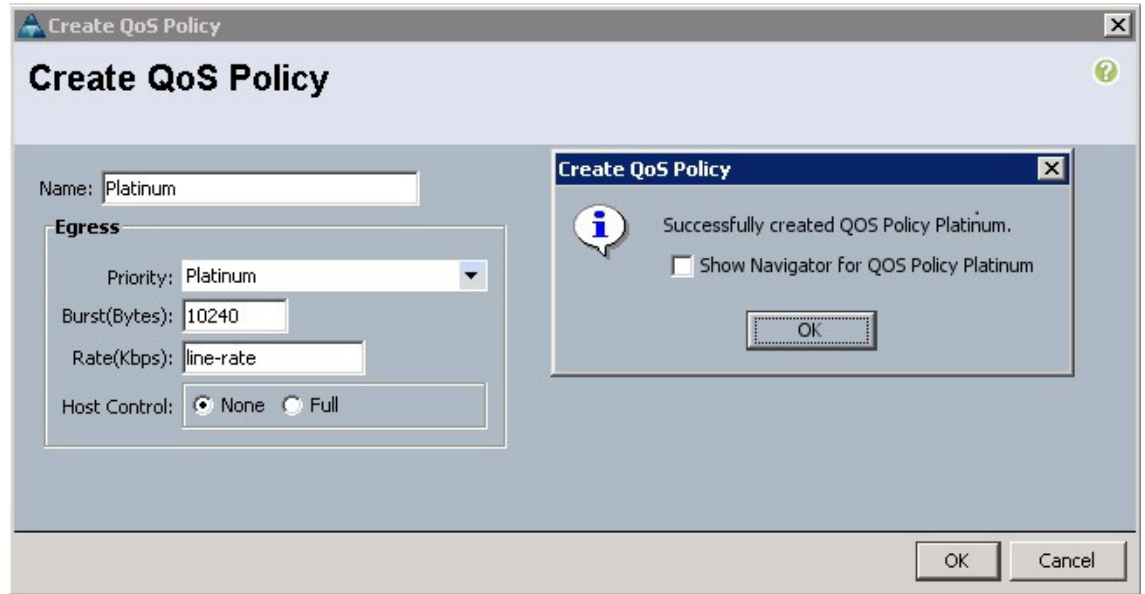


Creating a Platinum Policy

Follow these steps to create the Platinum QoS policy:

1. Click the **LAN** tab.
2. Select **Policies > root**.
3. Right-click **QoS Policies**.
4. Select **Create QoS Policy**.
5. Enter Platinum as the name of the policy as shown in [Figure 21](#).
6. Select Platinum from the drop down menu.
7. Keep the Burst (Bytes) field as default (10240).
8. Keep the Rate (Kbps) field as default (line-rate).
9. Keep the **Host Control** radio button as default (none).
10. Click **OK** to complete creating the Policy.
11. Click **OK**.

Figure 21 Creating Platinum QoS Policy

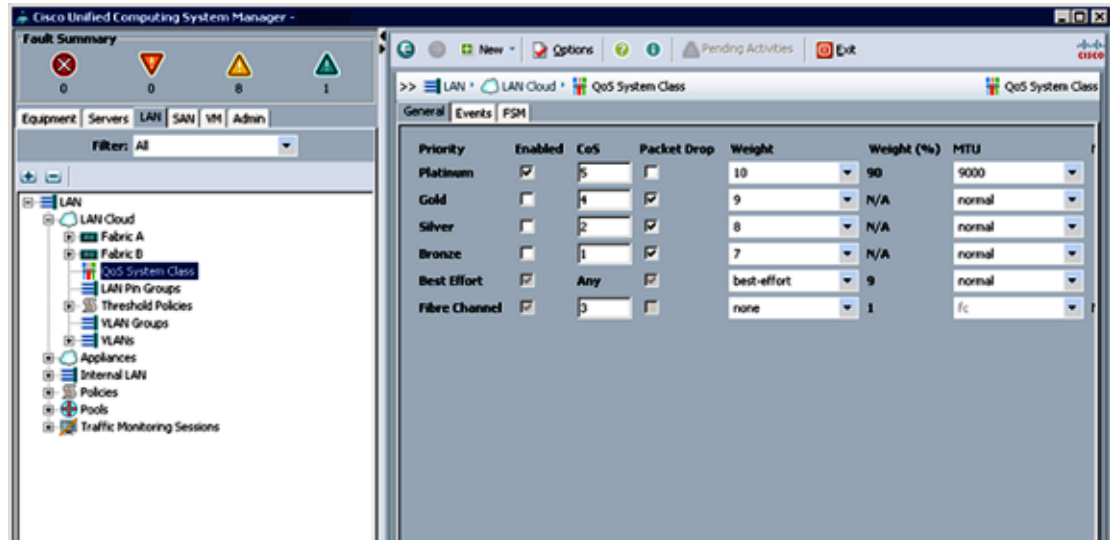


Setting Jumbo Frames

Follow these steps to set up Jumbo frames and enable the QoS:

1. Click the **LAN** tab in the Cisco UCS Manager.
2. Select **LAN Cloud > QoS System Class**.
3. In the right pane, click the **General** tab.
4. In the Platinum row, enter 9000 for MTU as shown in [Figure 22](#).
5. Check the **Enabled** check box.
6. Click **Save Changes**.
7. Click **OK**.

Figure 22 Setting Jumbo Frames



Creating a Local Disk Configuration Policy

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

1. Click the **Servers** tab.
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 as shown in [Figure 23](#).
6. Select **Any Configuration** from the drop-down list to set the Mode.
7. Uncheck the **Protect Configuration** check box.
8. Click **OK** to complete creating the Local Disk Configuration Policy.
9. Click **OK**.

Figure 23 Configuring Local Disk Policy

Create Local Disk Configuration Policy

Name: ←

Description:

Mode:

Protect Configuration:

If **Protect Configuration** is set, the local disk configuration is preserved if the service profile is disassociated with the server.
In that case, a configuration error will be raised when a new service profile is associated with that server if the local disk configuration in that profile is different.

Creating a Server BIOS Policy

The BIOS policy feature in Cisco UCS automates the BIOS configuration process. The traditional mode of setting the BIOS is manual and is often error-prone. By creating a BIOS policy and assigning the policy to a server or group of servers, you can enable transparency within the BIOS settings configuration.



Note

BIOS settings can have a significant performance impact, depending on the workload and the applications. The BIOS settings listed in this section is for configurations optimized for best performance which can be adjusted based on the application, performance and energy efficiency requirements.

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

1. Select the **Servers** tab.
2. Select **Policies > root**.
3. Right-click **BIOS Policies**.
4. Select **Create BIOS Policy**.
5. Enter the preferred BIOS policy name.
6. Change the BIOS settings as shown in [Figure 24](#).

Figure 24 Creating Server BIOS Policy

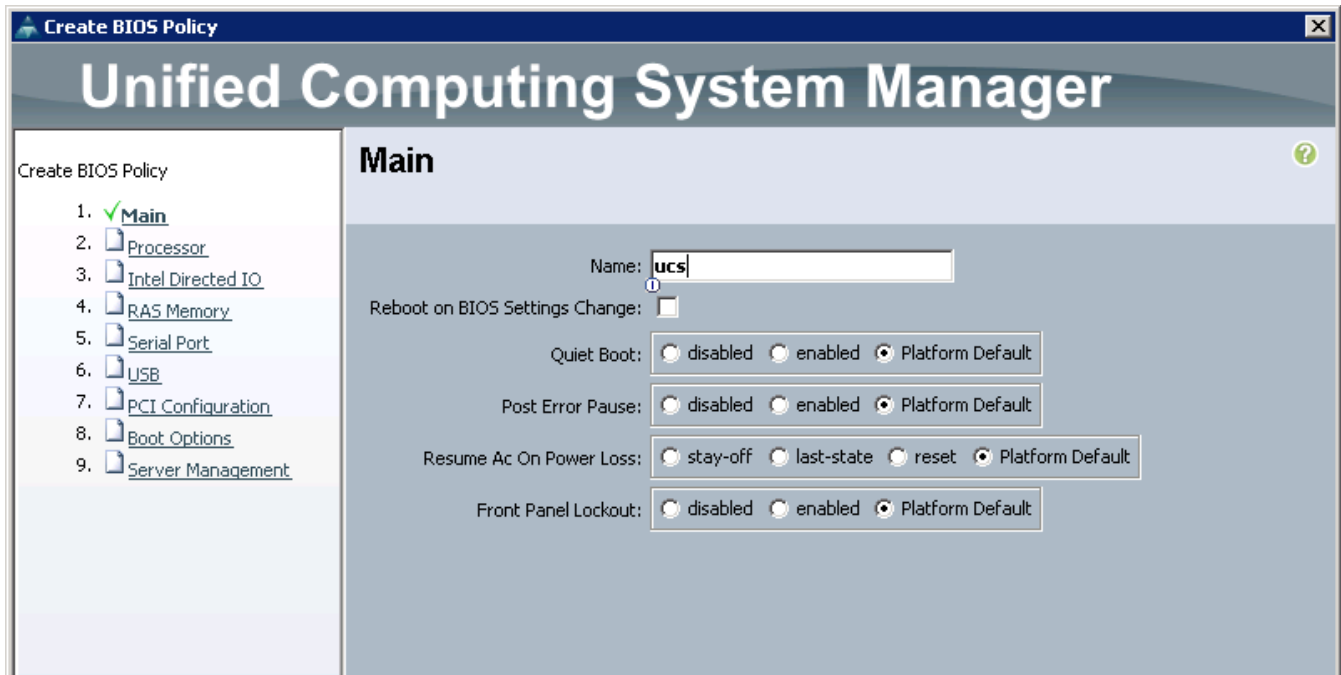


Figure 25 and Figure 26 show the Processor and Intel Directed IO properties settings in the BIOS Policy.

Figure 25 Creating Server BIOS Policy for Processor

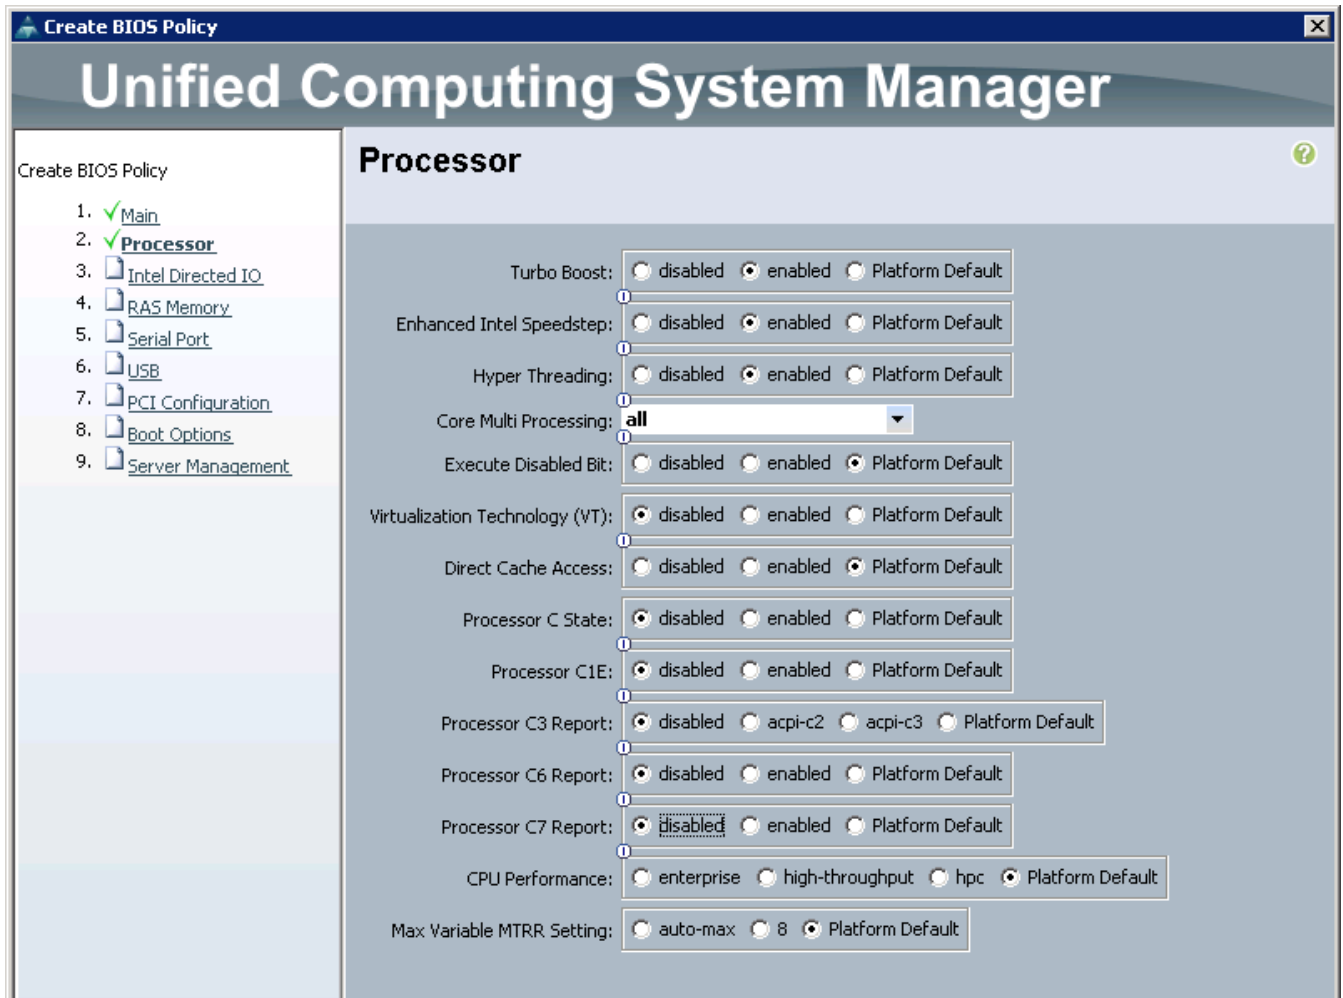
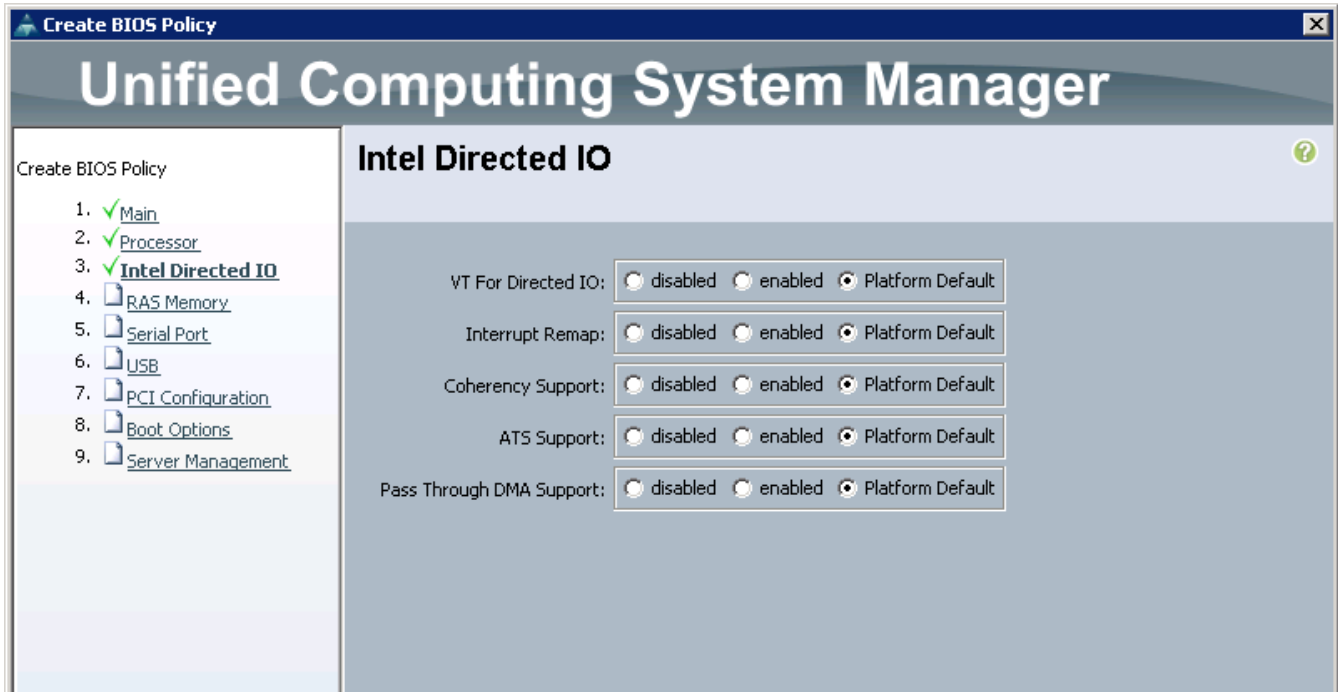
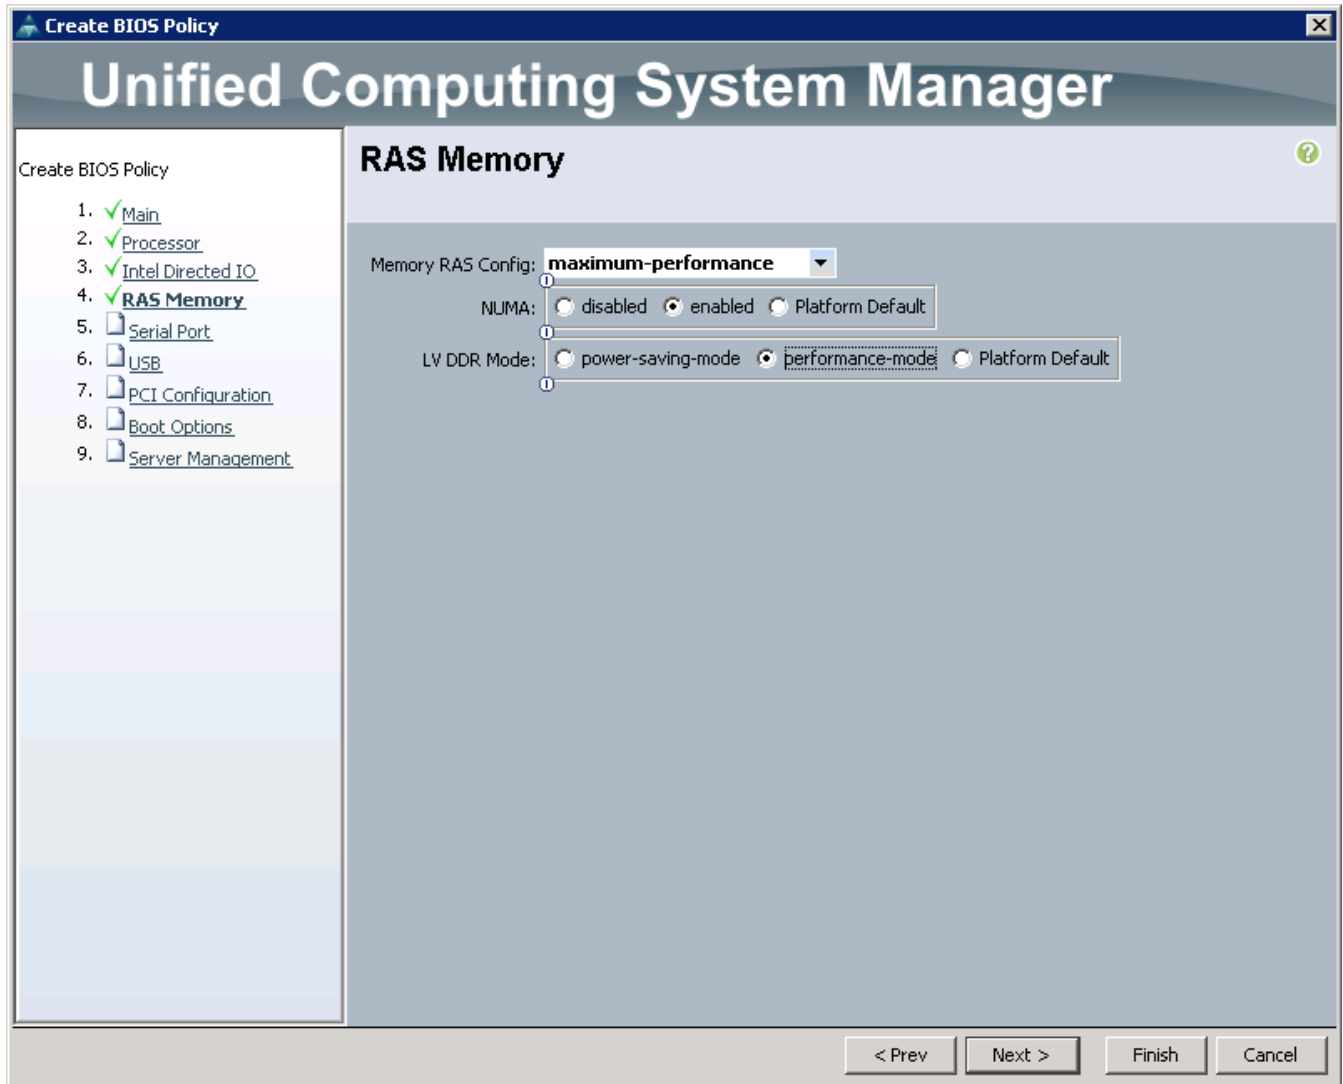


Figure 26 Creating Server BIOS Policy for Intel Directed IO



7. Set the RAS Memory settings and click **Next** as shown in Figure 27.

Figure 27 Creating Server BIOS Policy for Memory



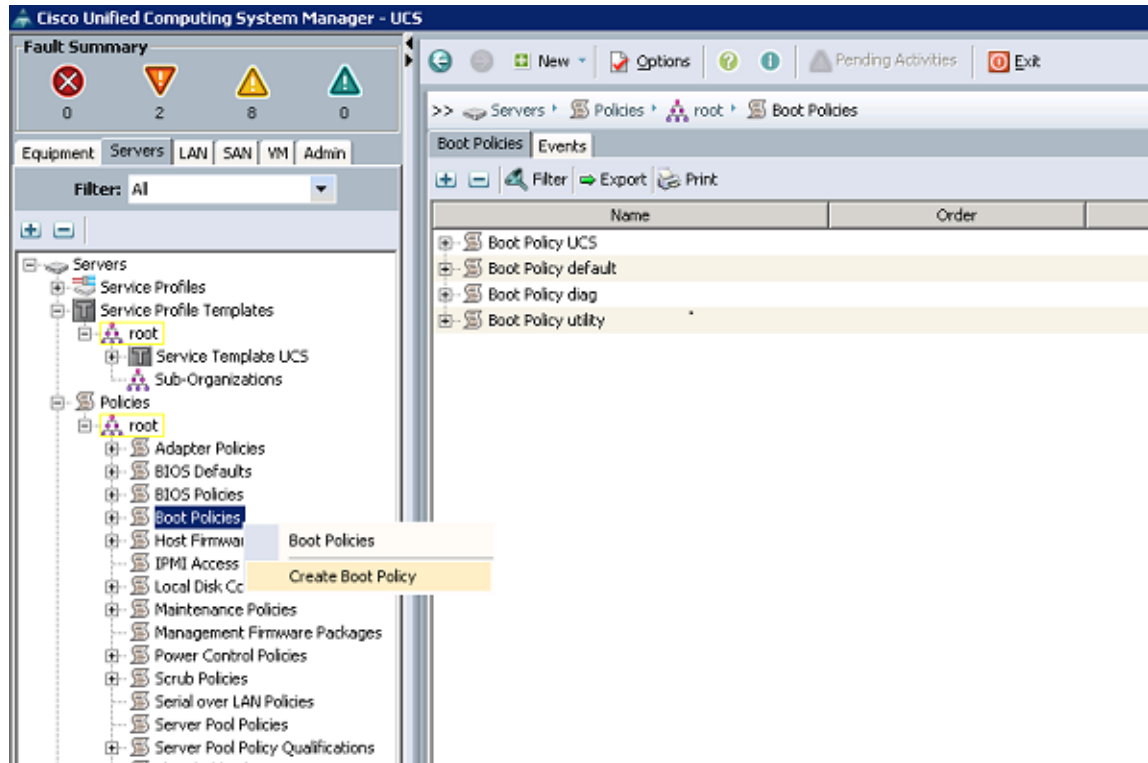
8. Click **Finish** to complete creating the BIOS Policy.
9. Click **OK**.

Creating a Boot Policy

Follow these steps to create a boot policy within Cisco UCS Manager:

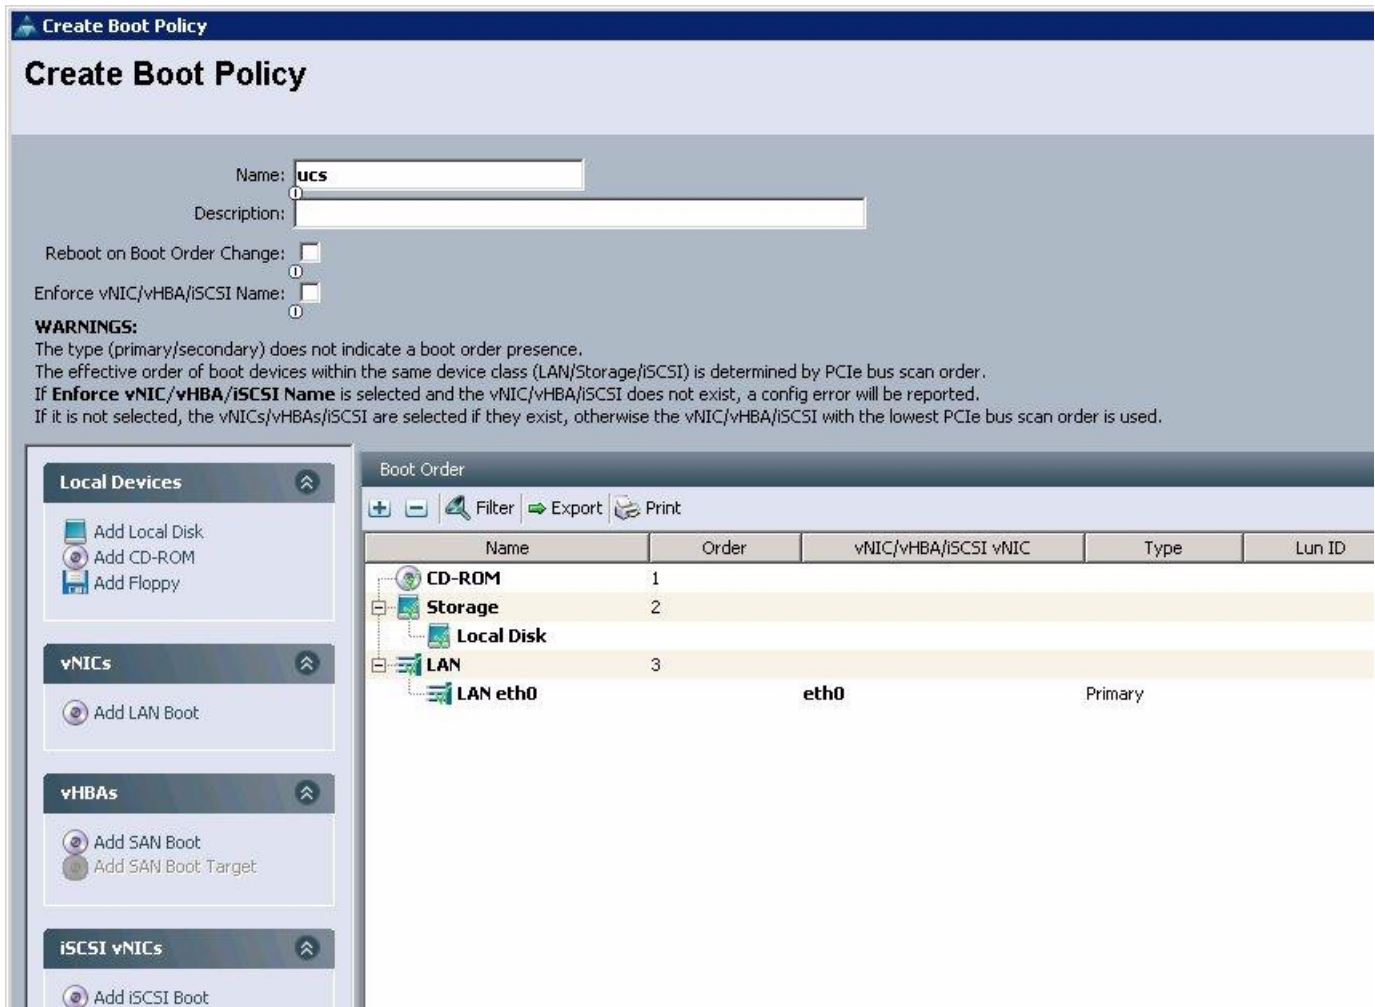
1. Select the **Servers** tab.
2. Select **Policies > root**.
3. Right-click the **Boot Policies**.
4. Select **Create Boot Policy** as shown in [Figure 28](#).

Figure 28 Creating Boot Policy Part 1



5. Enter **ucs** as the boot policy name as shown in Figure 29.
6. (Optional) Enter a description for the boot policy.
7. Keep the Reboot on Boot Order Change check box as default (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 Policy Part 2



Create Boot Policy

Name:

Description:

Reboot on Boot Order Change:

Enforce vNIC/vHBA/iSCSI Name:

WARNINGS:
 The type (primary/secondary) does not indicate a boot order presence.
 The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
 If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
 If it is not selected, the vNICs/vHBAs/iSCSI are selected if they exist, otherwise the vNIC/vHBA/iSCSI with the lowest PCIe bus scan order is used.

Local Devices

- Add Local Disk
- Add CD-ROM
- Add Floppy

vNICs

- Add LAN Boot

vHBAs

- Add SAN Boot
- Add SAN Boot Target

iSCSI vNICs

- Add iSCSI Boot

Boot Order

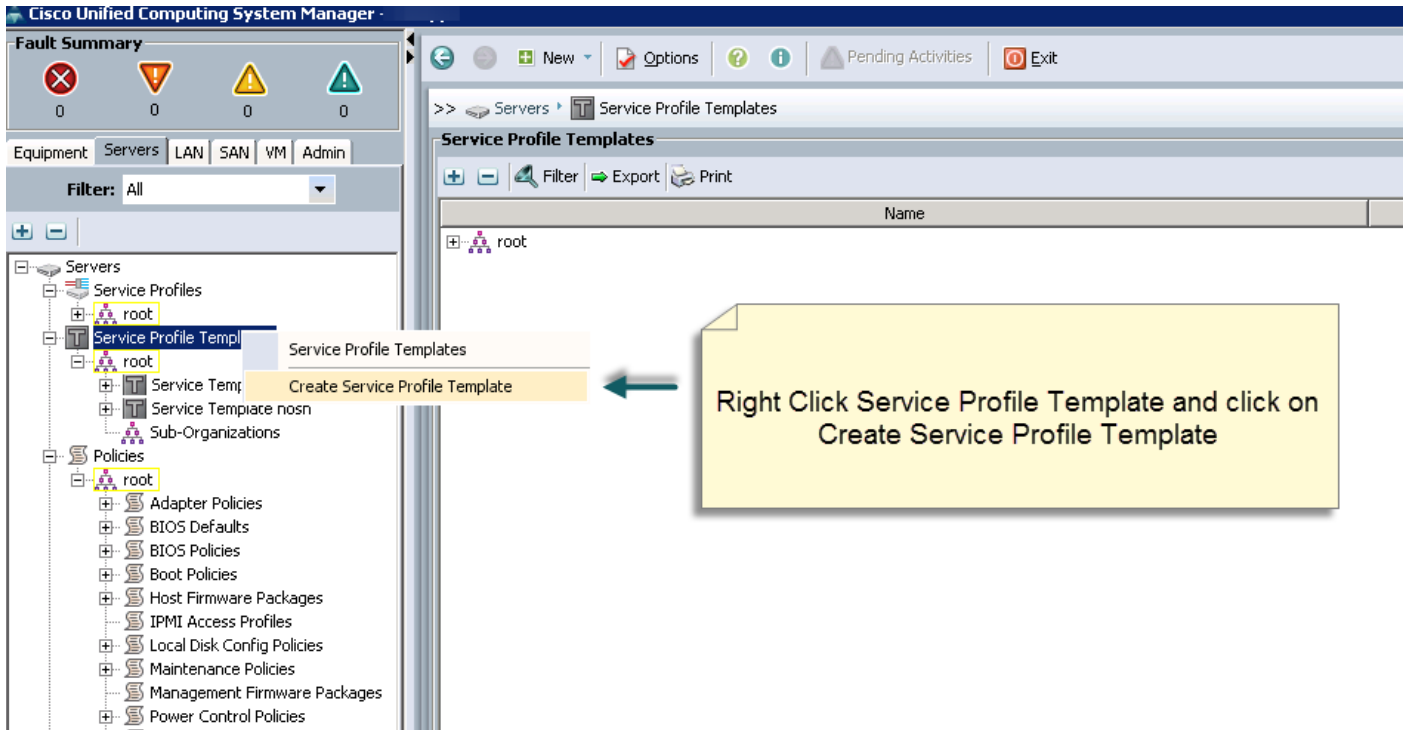
Name	Order	vNIC/vHBA/iSCSI vNIC	Type	Lun ID
CD-ROM	1			
Storage	2			
Local Disk	3			
LAN eth0	3	eth0	Primary	

Creating a Service Profile Template

Follow these steps to create a service profile template in Cisco UCS Manager:

1. Click the **Servers** tab.
2. Select **Policies > root**.
3. Right-click **root**.
4. Select **Create Service Profile Template** as shown in [Figure 30](#).

Figure 30 Creating Service Profile Template



5. The Create Service Profile Template window appears. Do the following (see [Figure 31](#)):
 - a. In the Identify Service Profile Template window, enter the name of the service profile template as ucs.
 - b. Click the **Updating Template** radio button.
 - c. In the UUID section, select **Hardware Default** as the UUID pool.
6. Click **Next** to continue.

Figure 31 Identify Service Profile Template

Create Service Profile Template

Unified Computing System Manager

Identify Service Profile Template

You must enter a name for the service profile template and specify the template type. You can also specify how a UUID will be assigned to this template and enter a description.

Name:

The template will be created in the following organization. Its name must be unique within this organization.

Where: **org-root**

The template will be created in the following organization. Its name must be unique within this organization.

Type: Initial Template Updating Template

Specify how the UUID will be assigned to the server associated with the service generated by this template.

UUID

UUID Assignment:

The UUID assigned by the manufacturer will be used.
Note: This UUID will not be migrated if the service profile is moved to a new server.

Optionally enter a description for the profile. The description can contain information about when and where the service profile should be used.

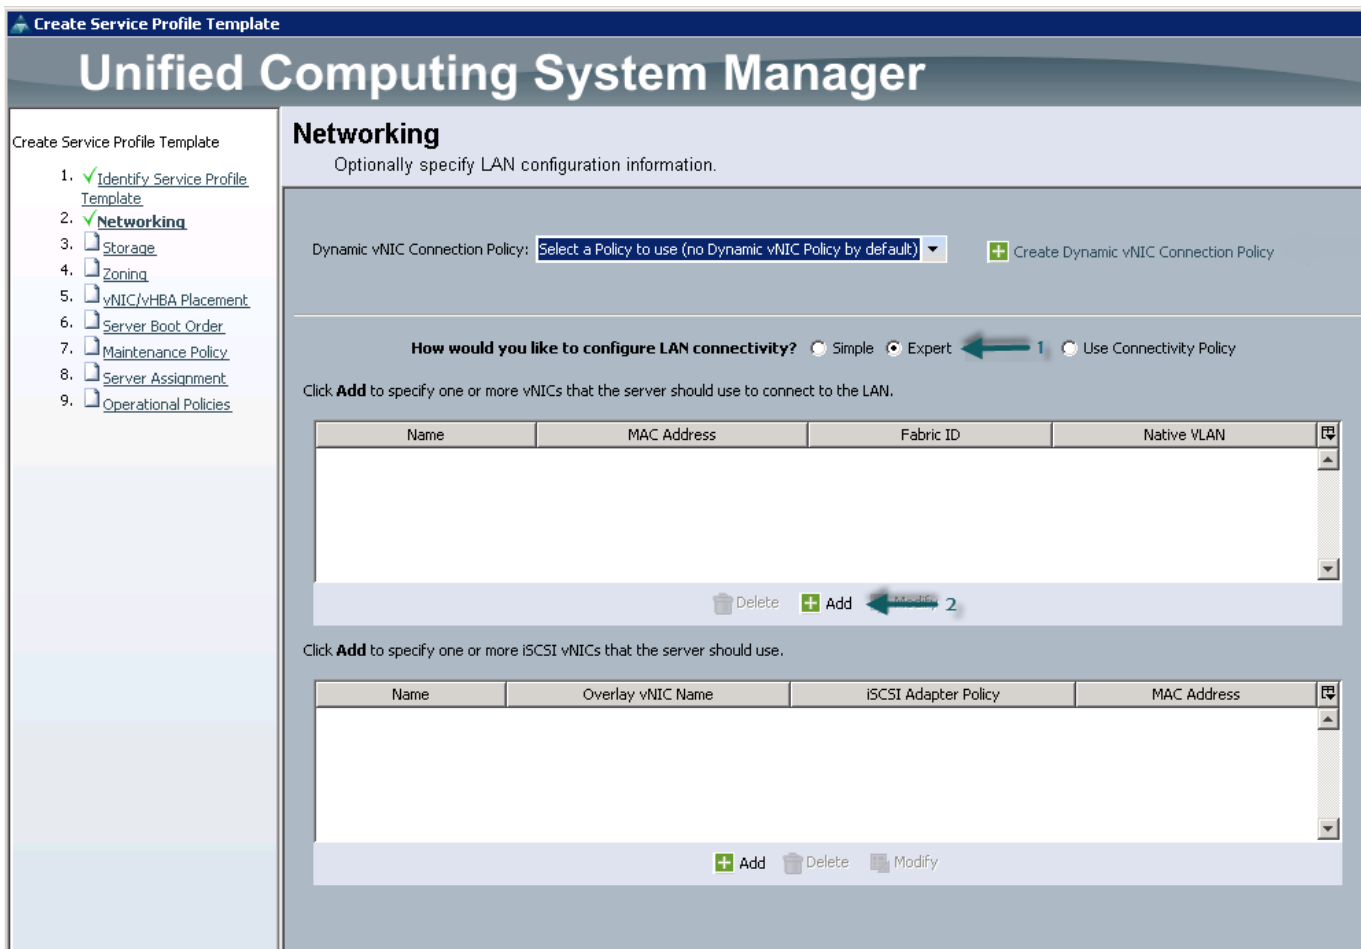
< Prev Next > Finish Cancel

Configuring the Network Settings for the Template

In the Networking window, follow these steps to configure the network settings in the Cisco UCS Manager:

1. Keep the Dynamic vNIC Connection Policy field at the default as shown in [Figure 32](#).
2. Click the **Expert** radio button to define How would you like to configure LAN connectivity?
3. Click **Add** to add a vNIC to the template. The Modify vNIC window appears.

Figure 32 Configuring Network Settings for the Template



4. In the Modify vNIC window, enter name for the vNIC as eth0 as shown in [Figure 33](#).
5. Select ucs in the MAC Address Assignment pool.
6. Click the **Fabric A** radio button and check the **Enable failover** check box for the Fabric ID.
7. Check the **vlan160_mgmt** check box for VLANs.
8. Click the **Native VLAN** radio button.
9. Select MTU size as 1500.
10. Select adapter policy as **Linux**.
11. Keep the Dynamic vNIC connection policy as <no set>.
12. Select **QoS Policy** as BestEffort.
13. Keep the Network Control Policy as default.
14. Click **OK**.

Figure 33 Configuring vNIC eth0

Modify vNIC

Name: **eth0**

Use vNIC Template:

Create vNIC Template

MAC Address

MAC Address Assignment: **ucs(512/512)**

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

Fabric ID: Fabric A Fabric B Enable Failover

VLANs

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

MTU:

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:

Operational Parameters

Adapter Performance Profile

Adapter Policy:

Dynamic vNIC Connection Policy:

QoS Policy:

Network Control Policy:

15. The Modify vNIC window appears. Enter the name of the vNIC as eth1 as shown in Figure 34.
16. Select ucs in the MAC Address Assignment pool.
17. Click 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.
19. Select MTU size as 9000.

20. Select adapter policy as Linux.
21. Keep the Dynamic vNIC Connection Policy as <no set>.
22. Select **QoS Policy** as Platinum.
23. Keep the Network Control Policy as default.
24. Click **OK**.

Figure 34 Configuring vNIC eth1

Modify vNIC

Name: **eth1**

Use vNIC Template:

Create vNIC Template

MAC Address

MAC Address Assignment: **ucs(512/512)**

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

Fabric ID: Fabric A Fabric B Enable Failover

VLANs

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

MTU:

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:

Operational Parameters

Adapter Performance Profile

Adapter Policy: **Linux**

Dynamic vNIC Connection Policy: **<not set>**

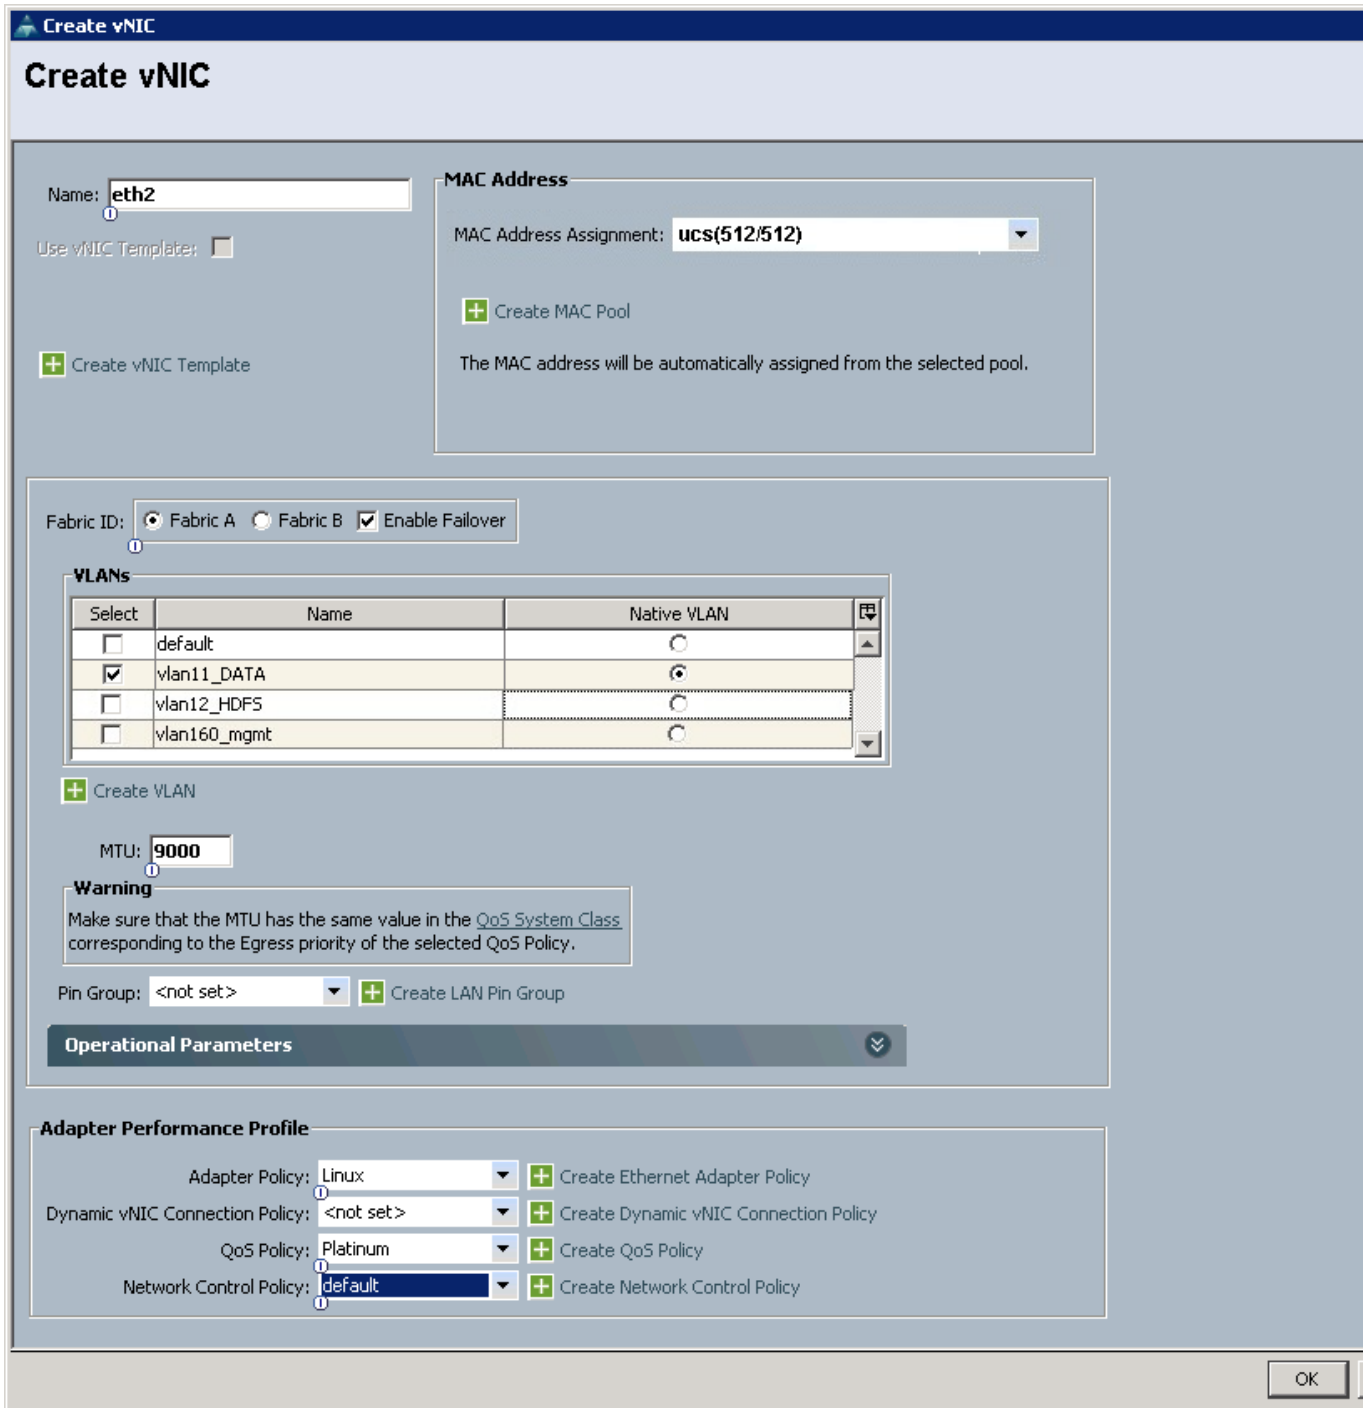
QoS Policy: **Platinum**

Network Control Policy: **default**

25. The Create vNIC window appears. Enter the name of the vNIC as eth2 as shown in Figure 35.
26. Select ucs in the MAC Address Assignment pool.
27. Click the **Fabric A** radio button and check the **Enable failover** check box for the Fabric ID.
28. Check the **vlan11_DATA** check box for VLANs and select the **Native VLAN** radio button.
29. Select MTU size as 9000.

30. Select adapter policy as **Linux**.
31. Keep the Dynamic vNIC Connection Policy as <no set>.
32. Select **QoS Policy** as **Platinum**.
33. Keep the Network Control Policy as default.
34. Click **OK**.
35. Click **Next** in the Networking window to continue.

Figure 35 Configuring vNIC eth2



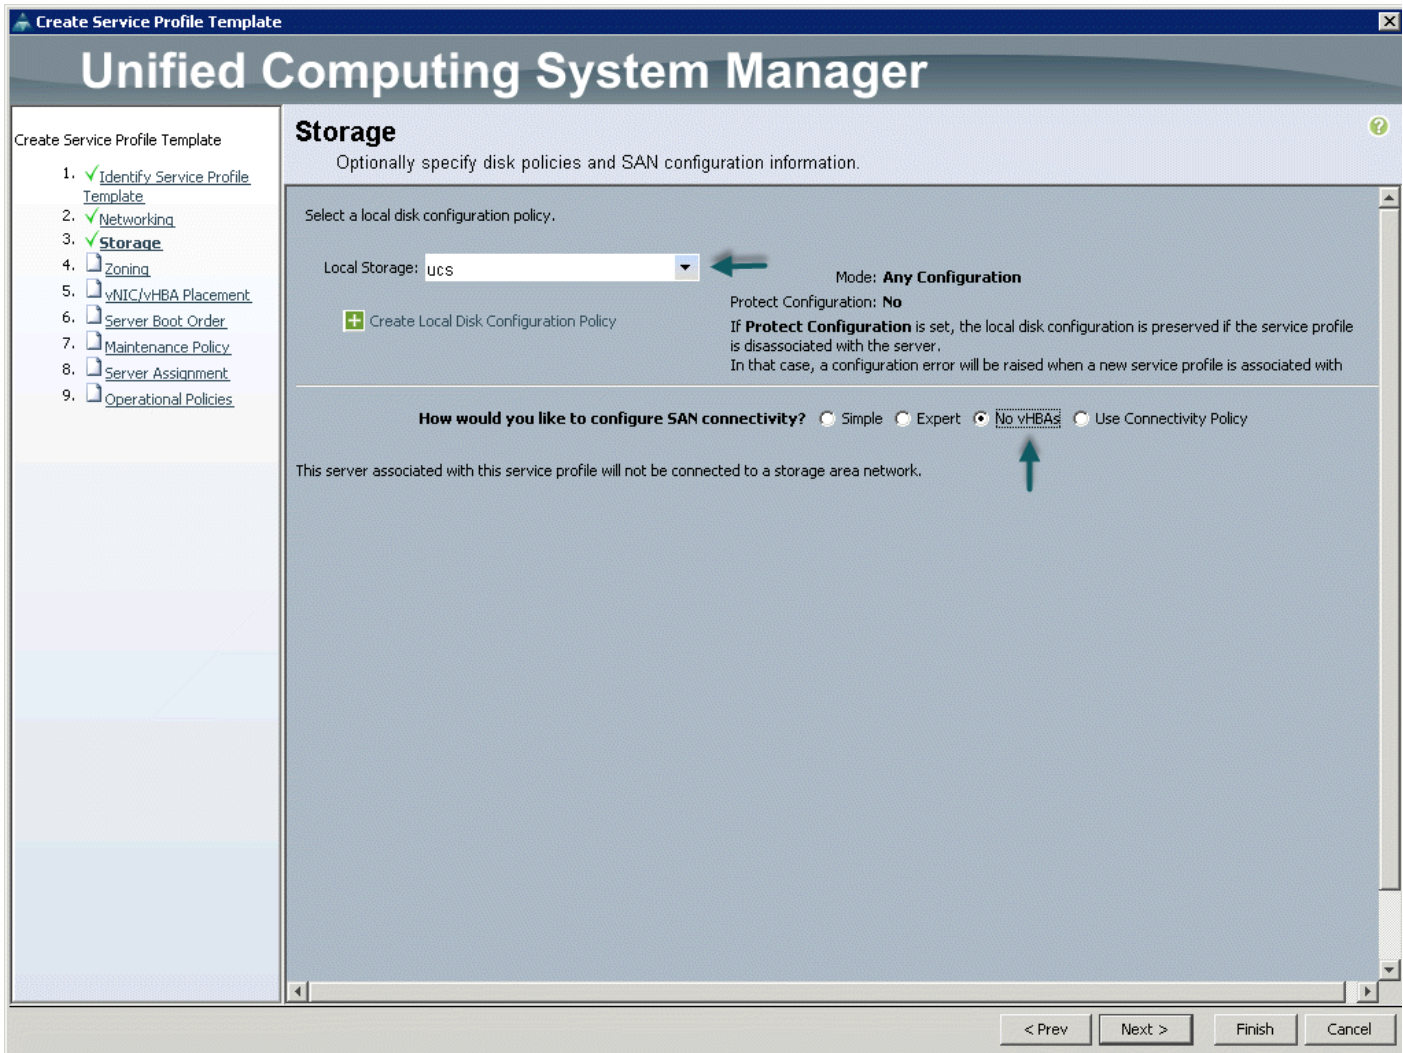
Configuring a Storage Policy for the Template

In the Storage window, follow these steps to configure a storage policy in Cisco UCS Manager:

1. Select ucs for the local disk configuration policy as shown in Figure 36.
2. Click the **No vHBAs** radio button to define How would you like to configure SAN connectivity?

3. Click **Next** to continue.

Figure 36 *Configuring Storage settings*



4. Click **Next** in the Zoning window to continue.

Configuring a vNIC/vHBA Placement for the Template

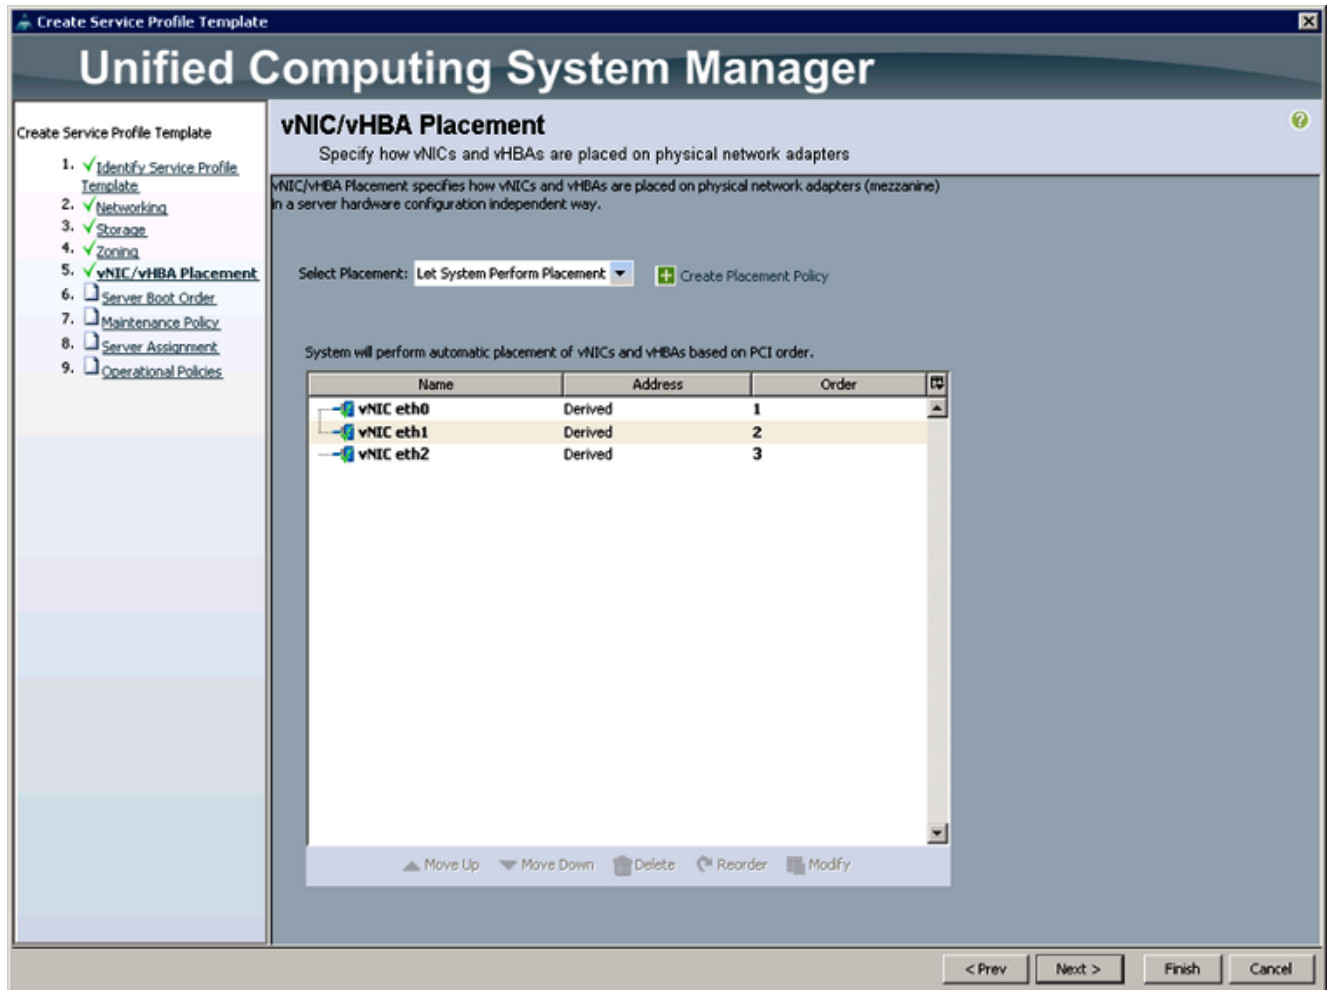
In the vNIC/vHBA window, follow these steps to configure a vNIC/vHBA placement policy in Cisco UCS Manager:

1. Select the Default Placement Policy option for the Select Placement field as shown in [Figure 37](#).
2. Select eth0, eth1 and eth2 assign the vNICs in the following order:
 - a. eth0
 - b. eth1
 - c. eth2

Review to make sure that all vNICs are assigned in the appropriate order.

3. Click **Next** to continue.

Figure 37 vNIC/vHBA Placement

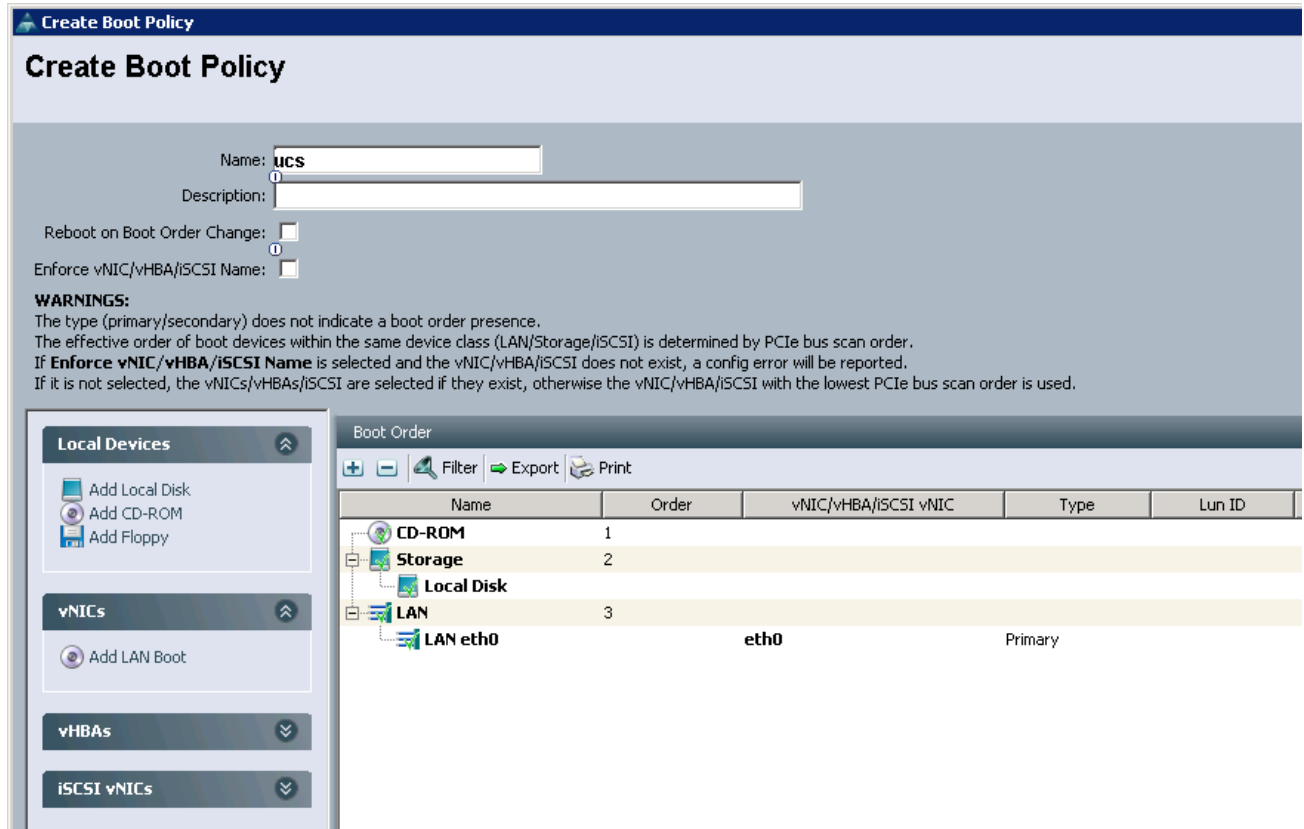


Configuring a Server Boot Order for the Template

In the Server Boot Order window, follow these steps to set the boot order for servers in Cisco UCS Manager:

1. Select **ucs** in the Boot Policy name field as shown in [Figure 38](#).
2. Check the **Enforce vNIC/vHBA/iSCSI Name** check box.
Review to make sure that all the boot devices are created and identified.
3. Verify that the boot devices are in the correct boot sequence.
4. Click **OK**.

Figure 38 Creating Boot Policy



5. Click **Next** to continue.

In the Maintenance Policy window, keep the default no policy as we have not created a policy. Click **Next** to continue to the next window.

Configuring Server Assignment for the Template

In the Server Assignment window, follow these steps to assign the servers to the pool in Cisco UCS Manager:

1. Select **ucs** for the **Pool Assignment** field as shown in [Figure 39](#).
2. Keep the Server Pool Qualification field as default.
3. Select **ucs** in **Host Firmware Package**.

Figure 39 Server Assignment

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

1. ✓ Identify Service Profile Template
2. ✓ Networking
3. ✓ Storage
4. ✓ Zoning
5. ✓ vNIC/vHBA Placement
6. ✓ Server Boot Order
7. ✓ Maintenance Policy
8. ✓ **Server Assignment**
9. Operational Policies

Server Assignment

Optionally specify a server pool for this service profile template.

Pool Assignment: **ucs** ← + Create Server Pool

Select the power state to be applied when this profile is associated with the server.

Up Down

The service profile template will be associated with one of the servers in the selected pool. If desired, you can specify an additional server pool policy qualification that the selected server must meet. To do so, select the qualification from the list.

Server Pool Qualification: **<not set>**

Restrict Migration:

Firmware Management (BIOS, Disk Controller, Adapter)

If you select a host firmware policy for this service profile, the profile will update the firmware on the server that it is associated with. Otherwise the system uses the firmware already installed on the associated server.

Host Firmware: **ucs** ← + Create Host Firmware Package

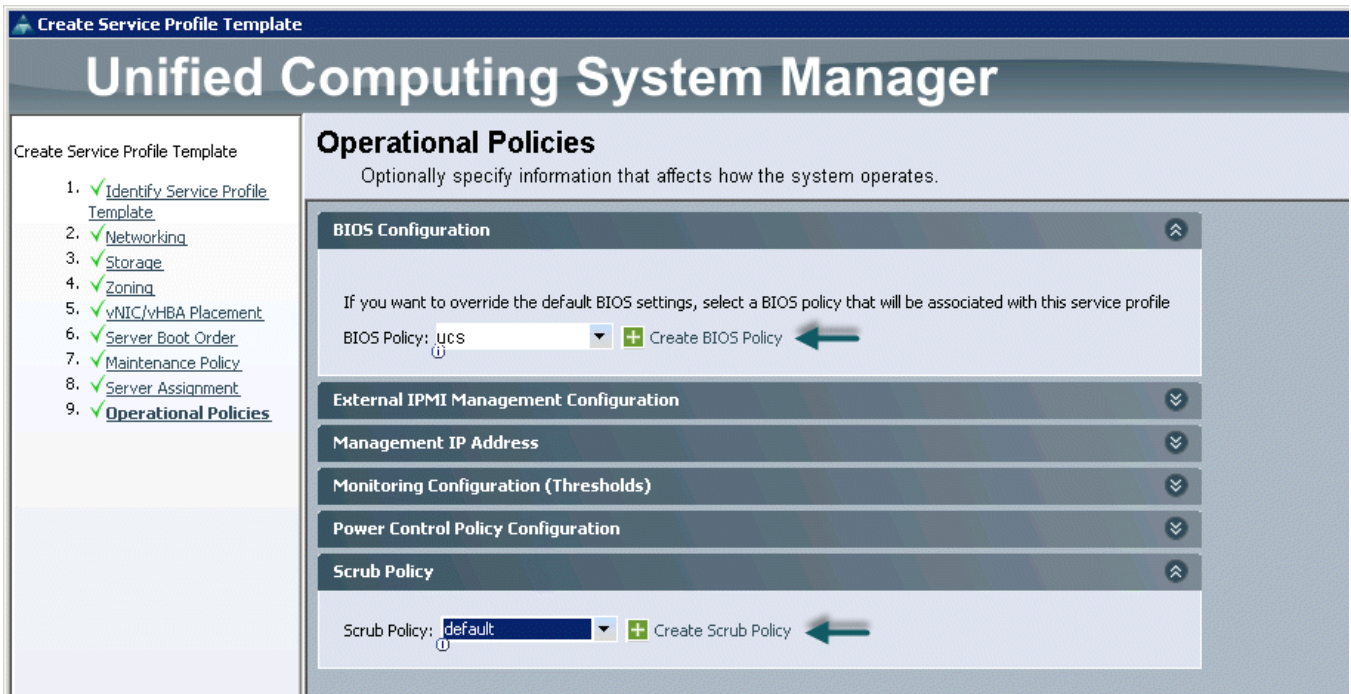
< Prev

Configuring Operational Policies for the Template

In the Operational Policies window, follow these steps:

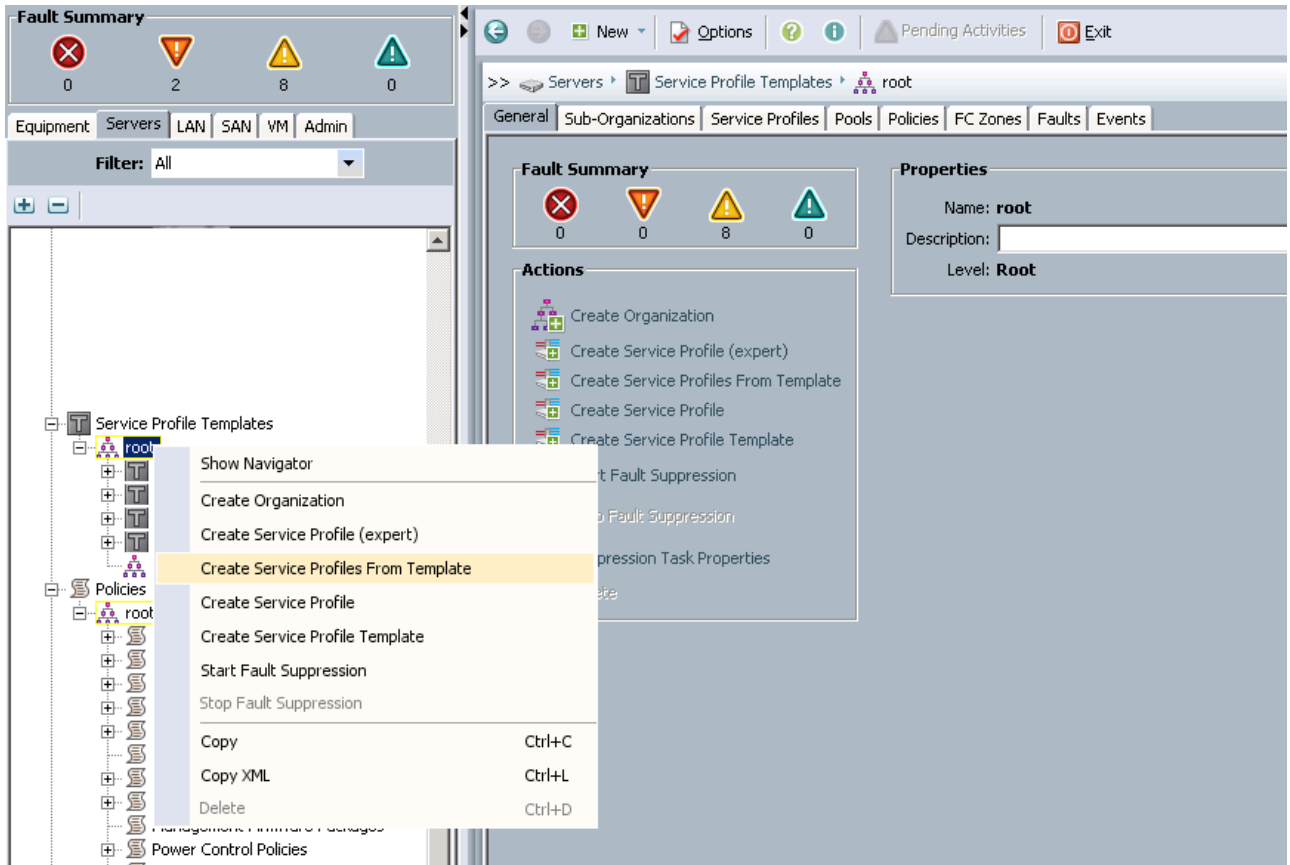
1. Select **ucs** in the **BIOS Policy** field as shown in [Figure 40](#).
2. Click **Finish** to create the Service Profile template.
3. Click **OK**.

Figure 40 Selecting BIOS Policy



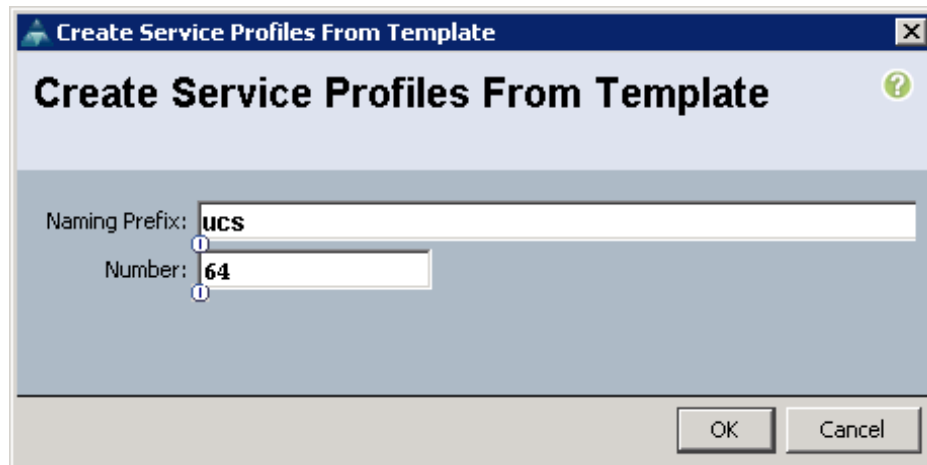
4. Click the **Servers** tab.
 - a. Select **Service Profile Templates > root**.
 - b. Right-click root and select **Create Service Profile Template** as shown in [Figure 41](#).

Figure 41 *Creating Service Profiles from Template*



- c. The Create Service Profile from Template window appears. Enter the name and number of nodes in the Name and Number fields as shown in Figure 42.

Figure 42 *Selecting Name and Total Number of Service Profiles*



The Cisco UCS Manager discovers the servers and automatically associate these servers with service profiles. Figure 43 illustrates the service profiles associated with all the 64-nodes.

Figure 43 Cisco UCS Manager showing 64 Nodes

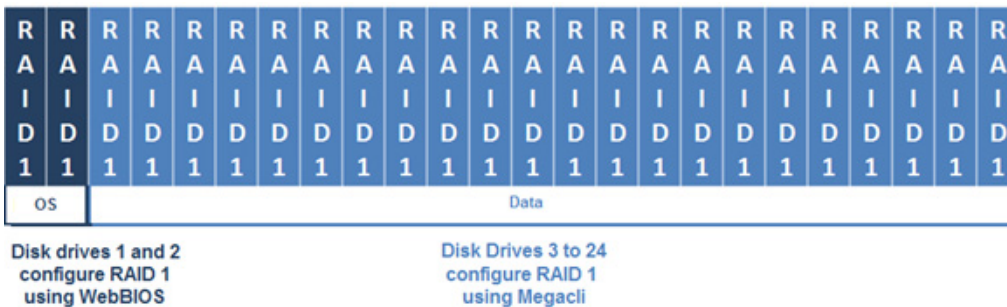
Name	Overall Status	PID	Model	User Label	Cores	Memory	Adapters	NICs	HBAs	Operability	Power State	Assoc State
Server 1	Ok	UCSC-C240-M35	Cisco UCS C240 M3		16	262144	1	2	0	Operable	On	Associated
Server 2	Ok	UCSC-C240-M35	Cisco UCS C240 M3		16	262144	1	2	0	Operable	On	Associated
Server 3	Ok	UCSC-C240-M35	Cisco UCS C240 M3		16	262144	1	2	0	Operable	On	Associated
Server 4	Ok	UCSC-C240-M35	Cisco UCS C240 M3		16	262144	1	2	0	Operable	On	Associated
Server 5	Ok	UCSC-C240-M35	Cisco UCS C240 M3		16	262144	1	2	0	Operable	On	Associated
Server 6	Ok	UCSC-C240-M35	Cisco UCS C240 M3		16	262144	1	2	0	Operable	On	Associated
Server 7	Ok	UCSC-C240-M35	Cisco UCS C240 M3		16	262144	1	2	0	Operable	On	Associated
Server 8	Ok	UCSC-C240-M35	Cisco UCS C240 M3		16	262144	1	2	0	Operable	On	Associated
Server 9	Ok	UCSC-C240-M35	Cisco UCS C240 M3		16	262144	1	2	0	Operable	On	Associated
Server 10	Ok	UCSC-C240-M35	Cisco UCS C240 M3		16	262144	1	2	0	Operable	On	Associated
Server 11	Ok	UCSC-C240-M35	Cisco UCS C240 M3		16	262144	1	2	0	Operable	On	Associated
Server 12	Ok	UCSC-C240-M35	Cisco UCS C240 M3		16	262144	1	2	0	Operable	On	Associated
Server 13	Ok	UCSC-C240-M35	Cisco UCS C240 M3		16	262144	1	2	0	Operable	On	Associated
Server 14	Ok	UCSC-C240-M35	Cisco UCS C240 M3		16	262144	1	2	0	Operable	On	Associated
Server 15	Ok	UCSC-C240-M35	Cisco UCS C240 M3		16	262144	1	2	0	Operable	On	Associated
Server 16	Ok	UCSC-C240-M35	Cisco UCS C240 M3		16	262144	1	2	0	Operable	On	Associated
Server 64	Ok	UCSC-C240-M35	Cisco UCS C240 M3		16	262144	1	2	0	Operable	On	Associated

Configuring Disk Drives for Operating System on NameNode

Namenode and Secondary Namenode have a different RAID configuration compared to Datanodes. This section details the configuration of disk drives for OS on these nodes (rhel1 and rhel2). The disk drives are configured as RAID1, read ahead cache and write cache are enabled when the battery is available. The first two disk drives are used for the Operating System and the remaining 22 disk drives are used for the HDFS as described in the following sections.

There are several ways to configure RAID such as using the 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 the LSI WebBIOS Configuration Utility and the remaining drives are configured using Linux-based MegaCLI commands after the completion of the Operating System installation.

Figure 44 RAID 1 Configured Using LSI WebBIOS Utility and MegaCLI




Follow these steps to create RAID1 on the first disk drive to install the Operating System:

1. Boot the server, and do the following:
 - a. Press <Ctrl><H> to launch the **WebBIOS**.
 - b. Press Ctrl+H immediately. The Adapter Selection window appears.
2. Click **Start** to continue as shown in [Figure 45](#).

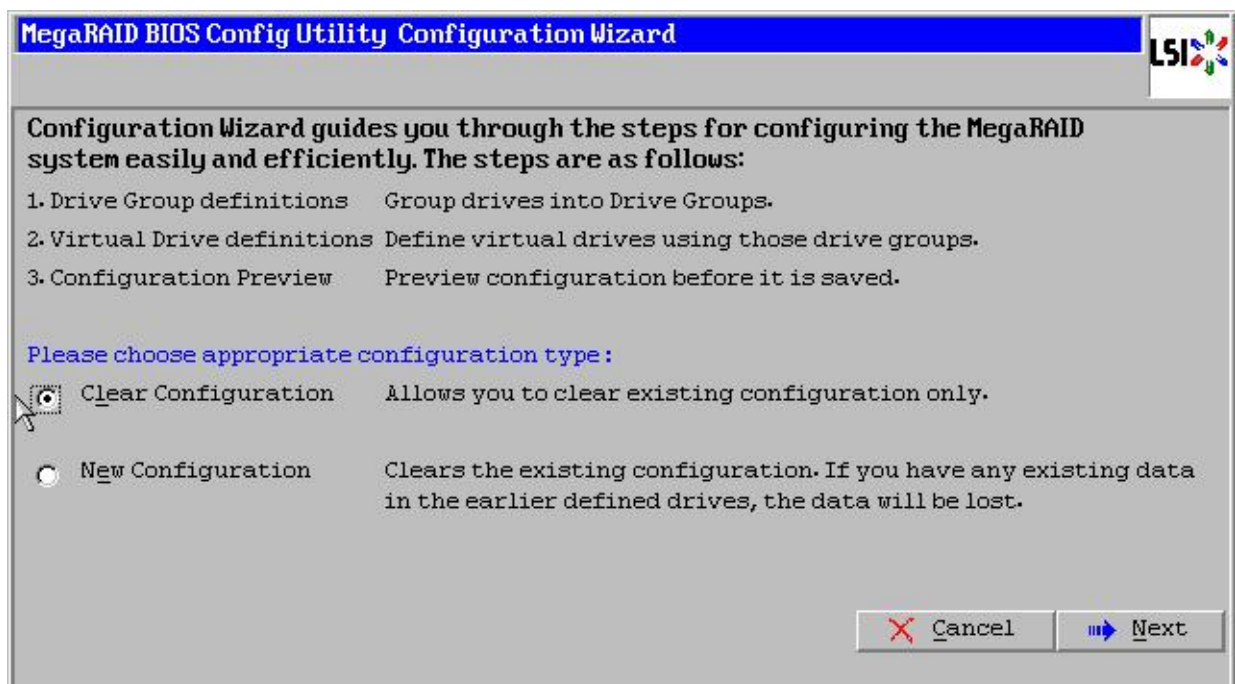
3. Click **Configuration Wizard**.

Figure 45 Adapter Selection for RAID Configuration



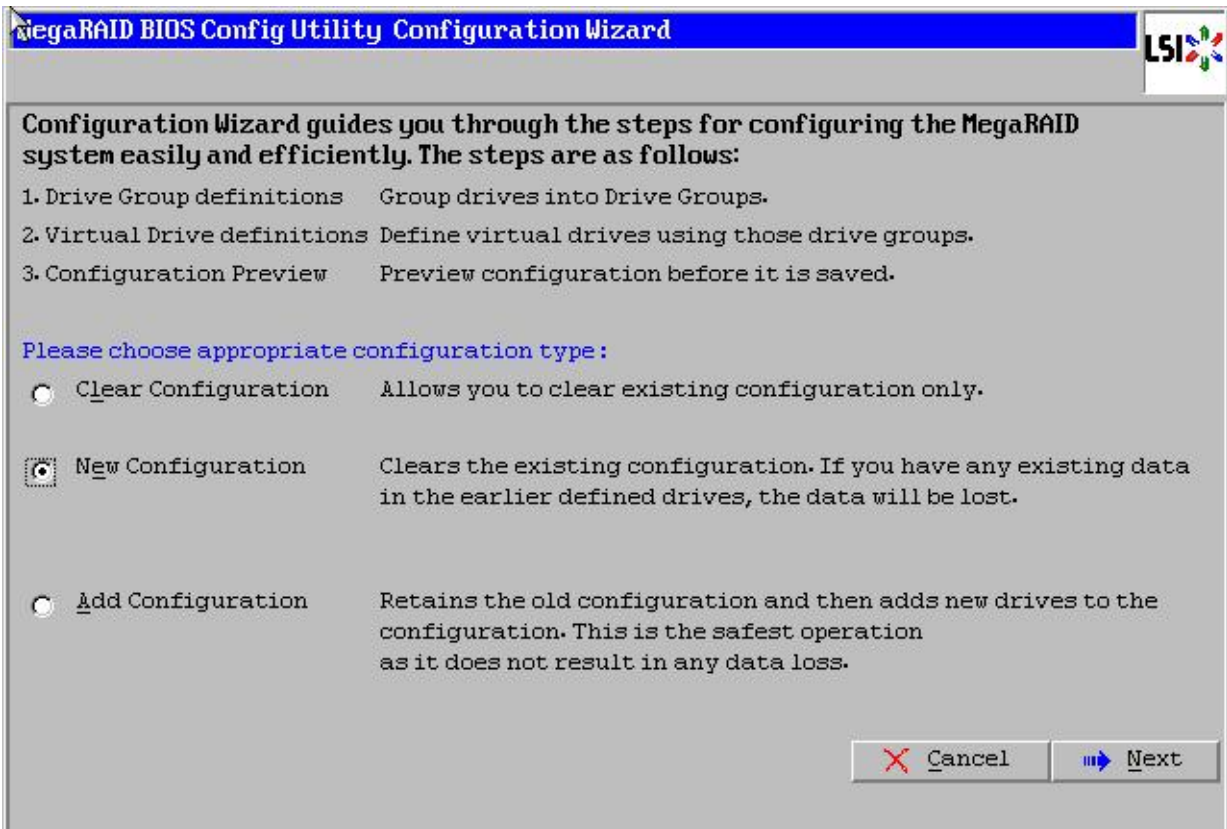
4. In the Configuration Wizard window, click the **Clear Configuration** radio button as shown in Figure 46.
5. Click **Next** to clear the existing configuration.

Figure 46 Clearing Current configuration on the controller



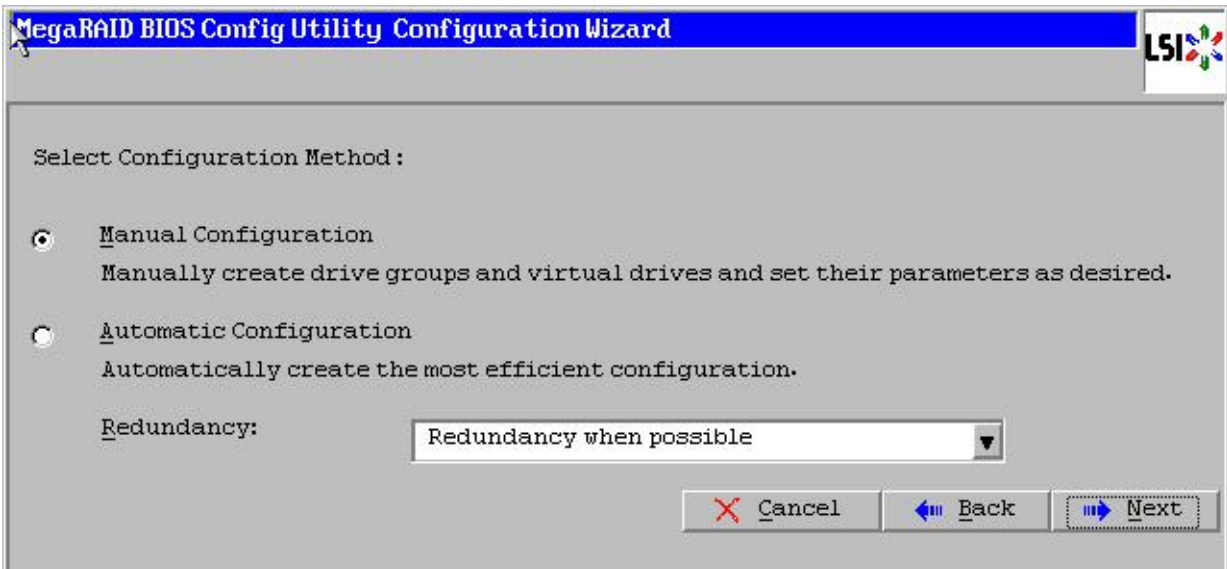
6. Click **Yes**.
7. In the Physical View, ensure that all the drives are Unconfigured Good.
8. In the Configuration Wizard window, click the **New Configuration** radio button as shown in Figure 47.
9. Click **Next**.

Figure 47 Choosing to create a New Configuration



10. Click the **Manual Configuration** radio button. This enables complete control over all attributes of the new storage configuration, such as, configuration of the drive groups, virtual drives and setting the parameters as shown in Figure 48.

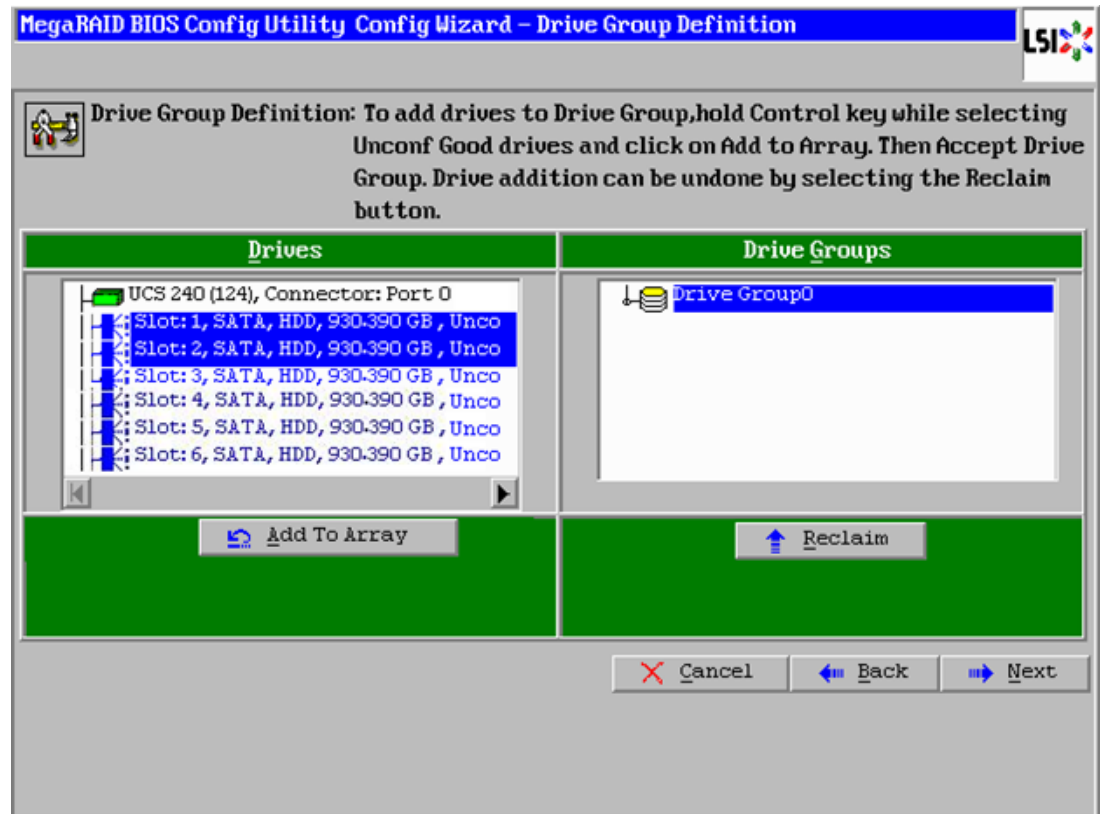
Figure 48 Choosing Manual Configuration Method



11. Click **Next**. The Drive Group Definition window appears.

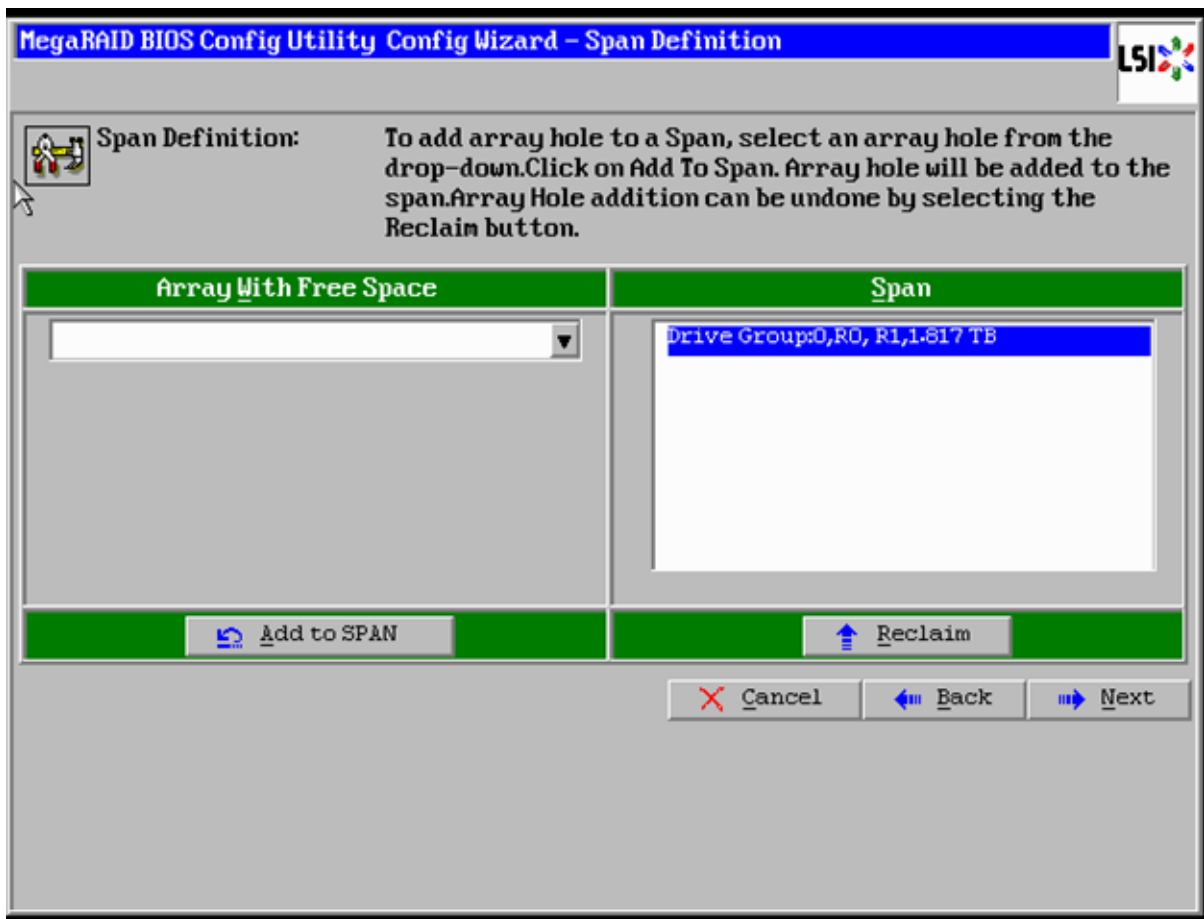
12. In the Drive Group Definition window, choose the first two drives to create drive groups as shown in Figure 49.
13. Click **Add to Array** to move the drives to a proposed drive group configuration in the Drive Groups pane.
14. Click **Accept DG** and click **Next**.

Figure 49 Selecting first drive and Adding to Drive Group



15. In the Span Definitions window, click **Add to SPAN** and click **Next** as shown in Figure 50.

Figure 50 Span Definition Window



16. In the Virtual Drive definitions window, do the following (see Figure 51):
 - a. Click on **Update Size**.
 - b. Change Strip Size to 1MB. A larger strip size ensures higher read performance.
 - c. From the Read Policy drop-down list, choose **Always Read Ahead**.
 - d. From the Write Policy drop-down list, choose **Write Back with BBU**.
 - e. Make sure RAID Level is set to RAID1.
 - f. Click **Accept** to accept the changes to the virtual drive definitions.
 - g. Click **Next**.



Note

Clicking on **Update Size** can change some of the settings in the window. Make sure all settings are correct before submitting the changes.

Figure 51 Defining Virtual Drive

MegaRAID BIOS Config Utility Config Wizard - Virtual Drive Definition

RAID Level: RAID 1

Strip Size: 1 MB

Access Policy: RW

Read Policy: Always Read Ahead

Write Policy: Write Back with BBU

IO Policy: Direct

Drive Cache: Unchanged

Disable BGI: No

Select Size: 930.390 GB

Update Size

Virtual Drives


Next LD, Possible RAID Levels
R1:930-390 GB

Accept Reclaim

Cancel Back Next

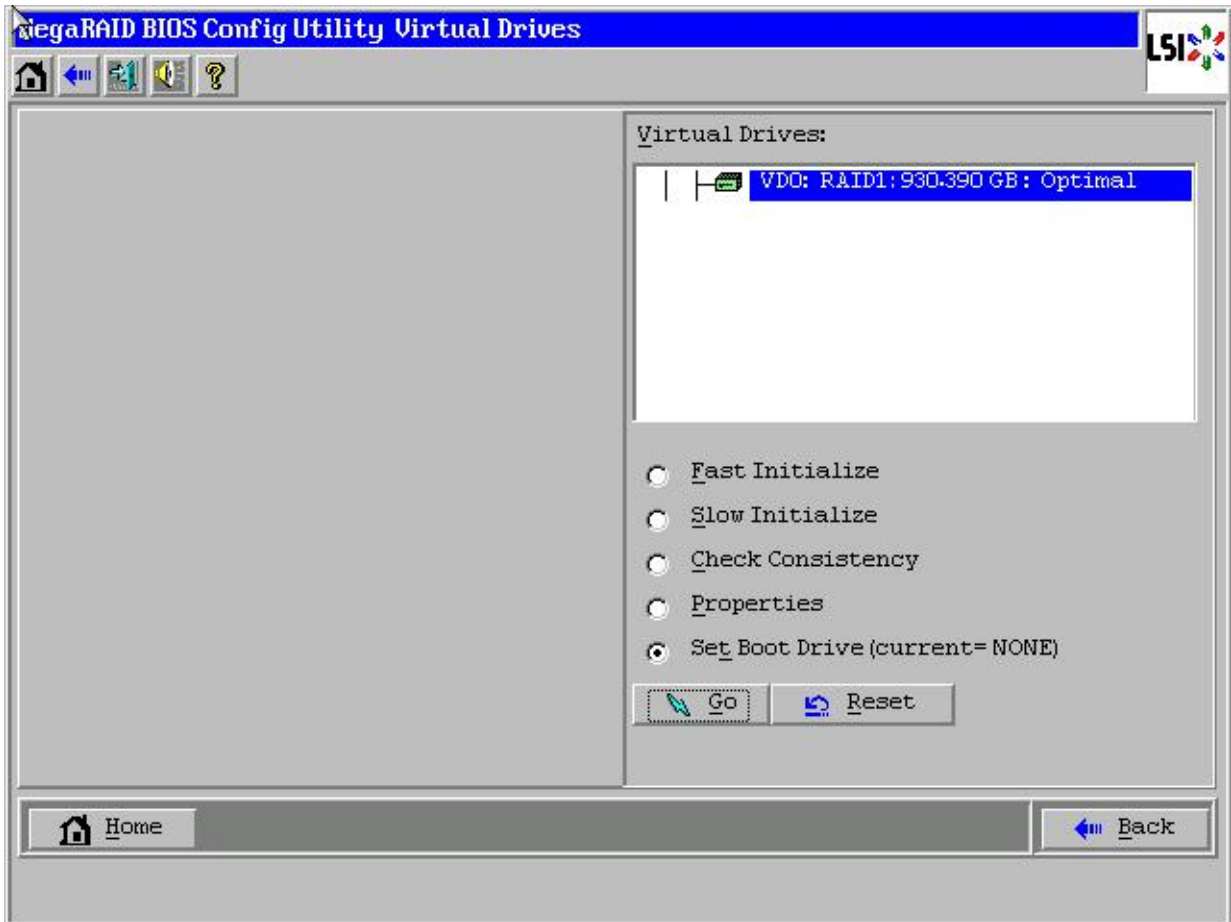
17. Click **Yes** to save the configuration.
18. In the Managing SSD Caching window, click **Cancel** as shown in [Figure 52](#).

Figure 52 SSD Caching Window



19. Click **Yes** in the confirmation page.
20. Set VD0 as the Boot Drive and click **Go** as shown in [Figure 53](#).

Figure 53 Setting Virtual Drive as Boot Drive



21. Click **Home**.
22. Review the configuration and click **Exit**.

Configuration of disks 2 to 24 are done using Linux based MegaCLI commands described in [“Configuring Data Drives on NameNode”](#) section on page 97.

Configuring Disk Drives for Operating System on DataNodes

Nodes 3 through 64 are configured as DataNodes. This section details the configuration of disk drives for OS on the data nodes. The focus of this CVD is on 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 and write cache are enabled when the battery is available. The first disk drive is used for the Operating System and the remaining 23 disk drives are used for the HDFS as described in the following sections.



Note


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

There are several ways to configure RAID. RAID can be configured using LSI WebBIOS Configuration Utility embedded in the MegaRAID BIOS, booting DOS and running MegaCLI commands, Linux based MegaCLI commands, or by third party tools having MegaCLI. For this deployment, the first disk drive is configured using LSI WebBIOS Configuration Utility and the rest of them are configured using a Linux based MegaCLI commands after the completion of OS installation.

Follow these steps to create RAID0 on the first disk drive to install the Operating System:

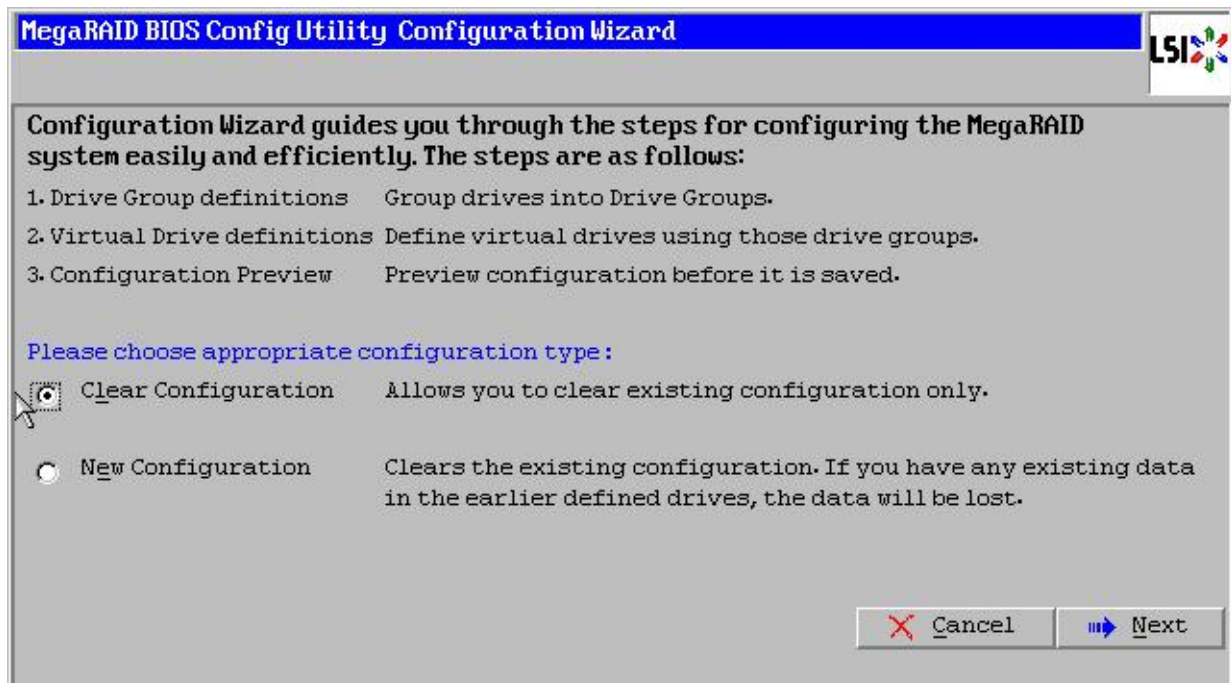
1. Boot the server, and do the following:
 - a. Press <Ctrl><H> to launch the **WebBIOS**.
 - b. Press Ctrl+H immediately. The Adapter Selection window appears.
2. Click **Start** to continue as shown in [Figure 45](#).
3. Click **Configuration Wizard**.

Figure 54 Adapter Selection for RAID Configuration



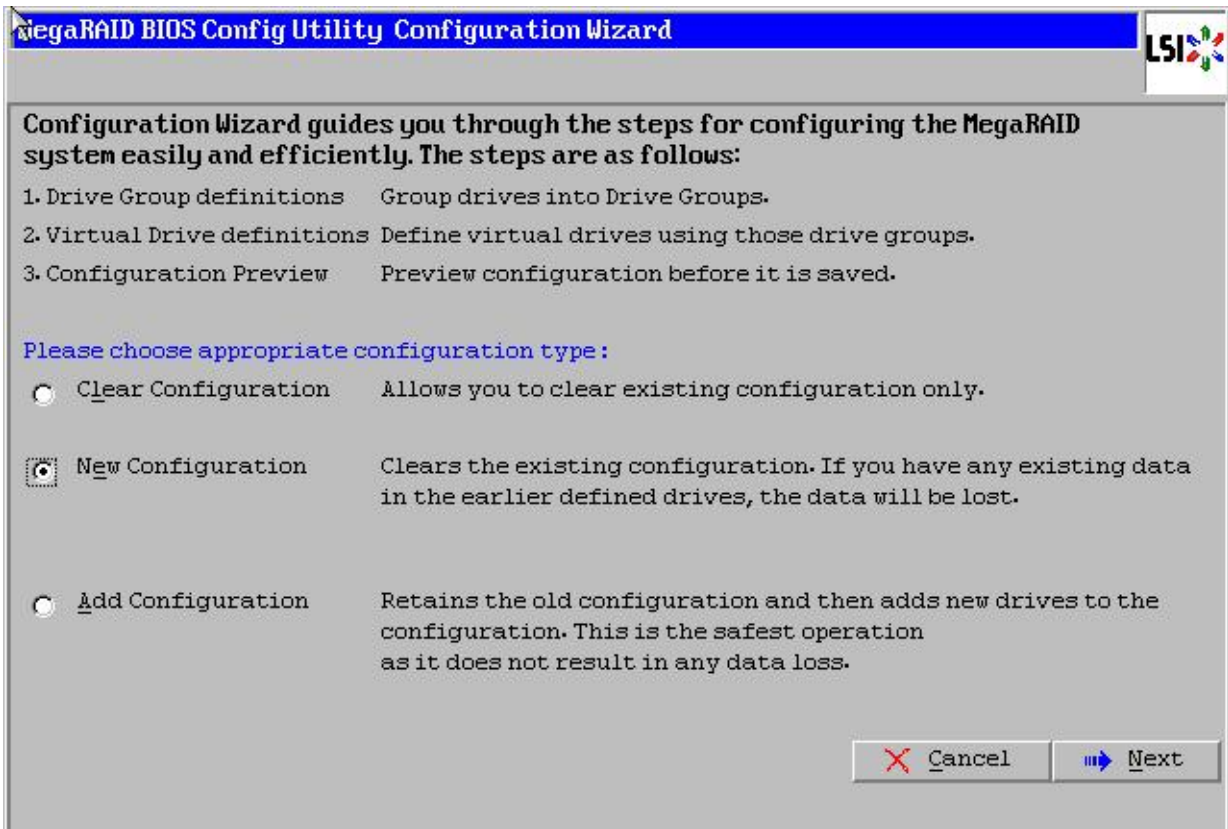
4. In the Configuration Wizard window, click the **Clear Configuration** radio button as shown in [Figure 46](#).
5. Click **Next** to clear the existing configuration.

Figure 55 Clearing Current configuration on the controller



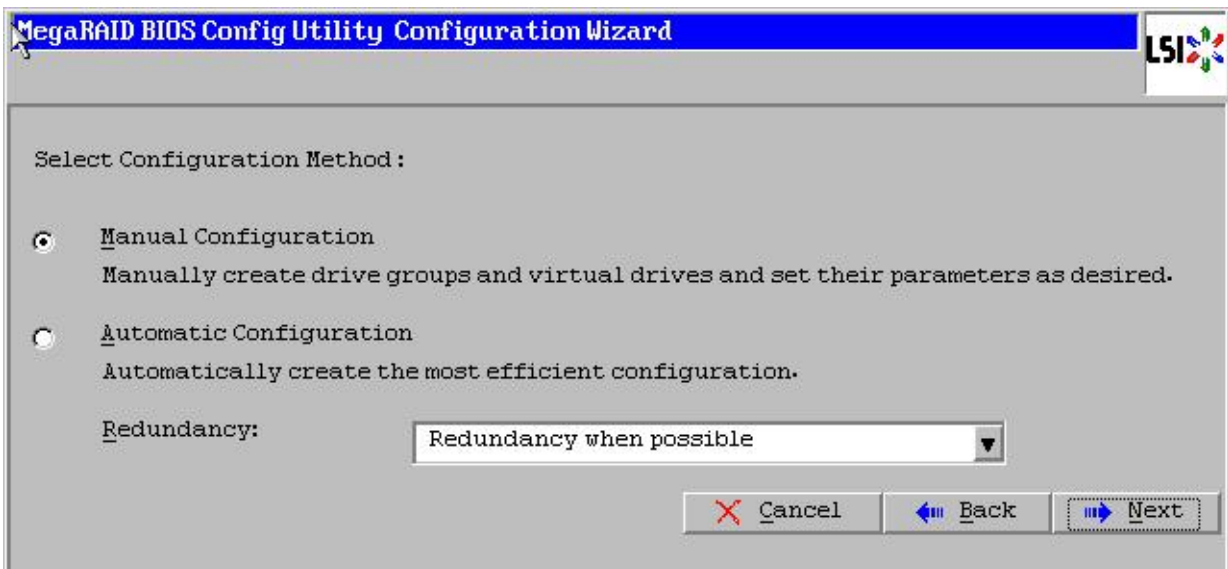
6. Click **Yes**.
7. In the Physical View, ensure that all the drives are Unconfigured Good.
8. In the Configuration Wizard window, click the **New Configuration** radio button as shown in [Figure 47](#).
9. Click **Next**.

Figure 56 Choosing to create a New Configuration



10. Click the **Manual Configuration** radio button. This enables complete control over all attributes of the new storage configuration, such as, configuration of the drive groups, virtual drives and setting the parameters as shown in Figure 48.

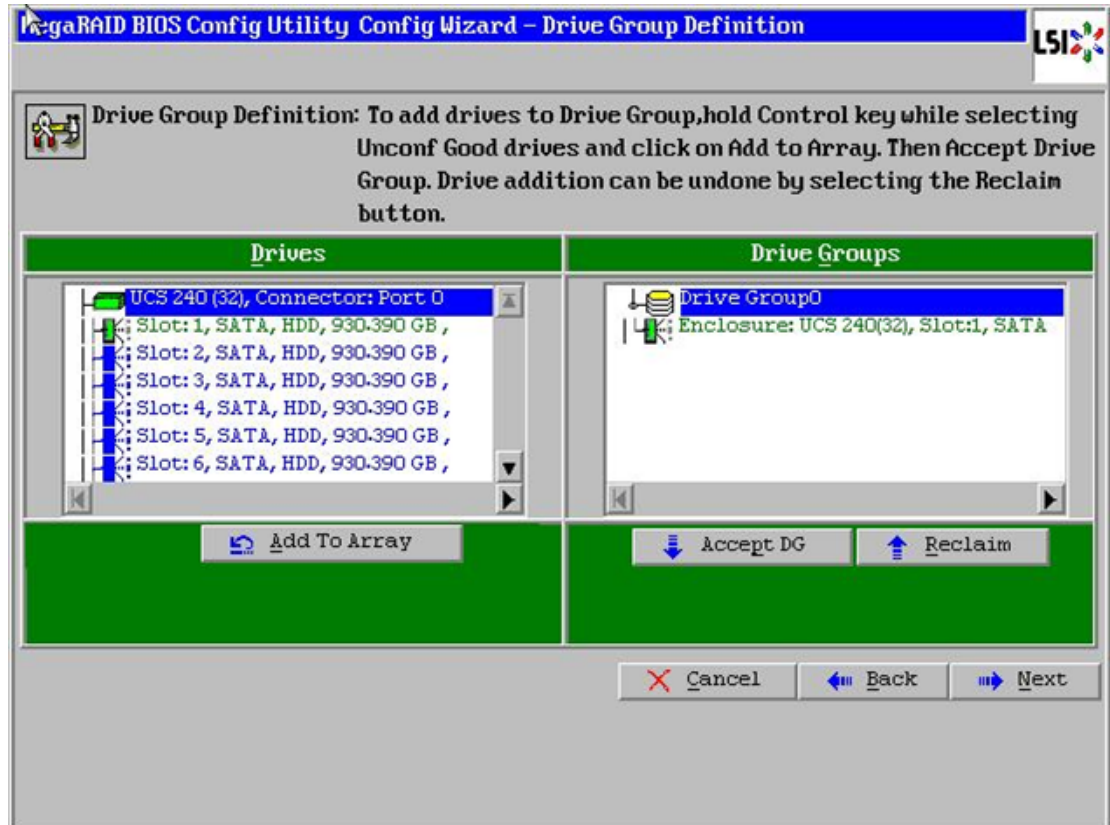
Figure 57 Choosing Manual Configuration Method



11. Click **Next**. The Drive Group Definition window appears.

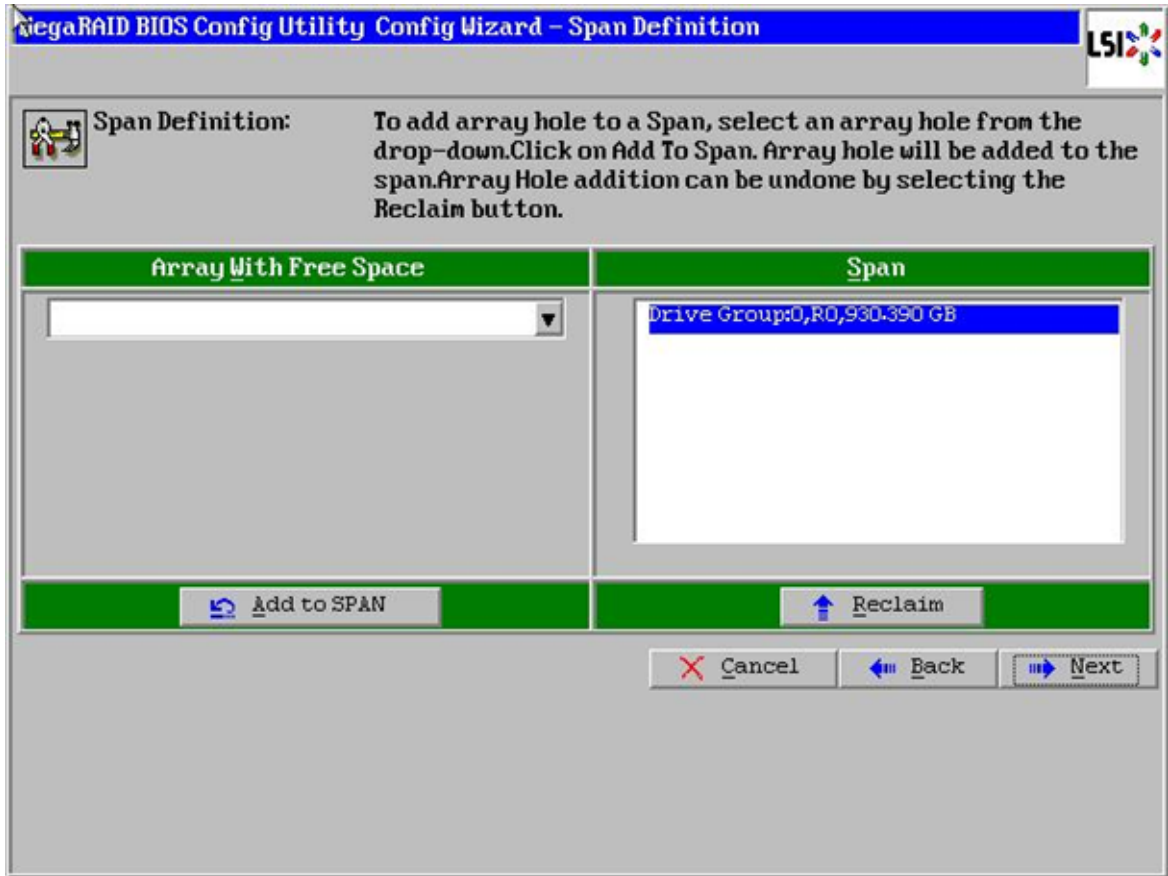
12. In the Drive Group Definition window, choose the first drive to create drive groups as shown in Figure 49.
13. Click **Add to Array** to move the drives to a proposed drive group configuration in the Drive Groups pane.
14. Click **Accept DG** and click **Next**.

Figure 58 Selecting first drive and Adding to Drive Group



15. In the Span Definitions window, click **Add to SPAN** and click **Next** as shown in Figure 50.

Figure 59 Span Definition Window



16. In the Virtual Drive definitions window, do the following (see [Figure 51](#)):
 - a. Click on **Update Size**.
 - b. Change Strip Size to 1MB. A larger strip size ensures higher read performance.
 - c. From the Read Policy drop-down list, choose **Always Read Ahead**.
 - d. From the Write Policy drop-down list, choose **Write Back with BBU**.
 - e. Make sure RAID Level is set to RAID0.
 - f. Click **Accept** to accept the changes to the virtual drive definitions.
 - g. Click **Next**.



Note

Clicking on **Update Size** can change some of the settings in the window. Make sure all settings are correct before submitting the changes.

Figure 60 Defining Virtual Drive

MegaRAID BIOS Config Utility Config Wizard - Virtual Drive Definition

RAID Level: RAID 0

Strip Size: 1 MB

Access Policy: RW

Read Policy: Always Read Ahead

Write Policy: Write Back with BBU

IO Policy: Direct

Drive Cache: Unchanged

Disable BGI: No

Select Size: 930.390 GB

Virtual Drives

Next LD, Possible RAID Levels
R0:930.390 GB

Accept Reclaim

Cancel Back Next

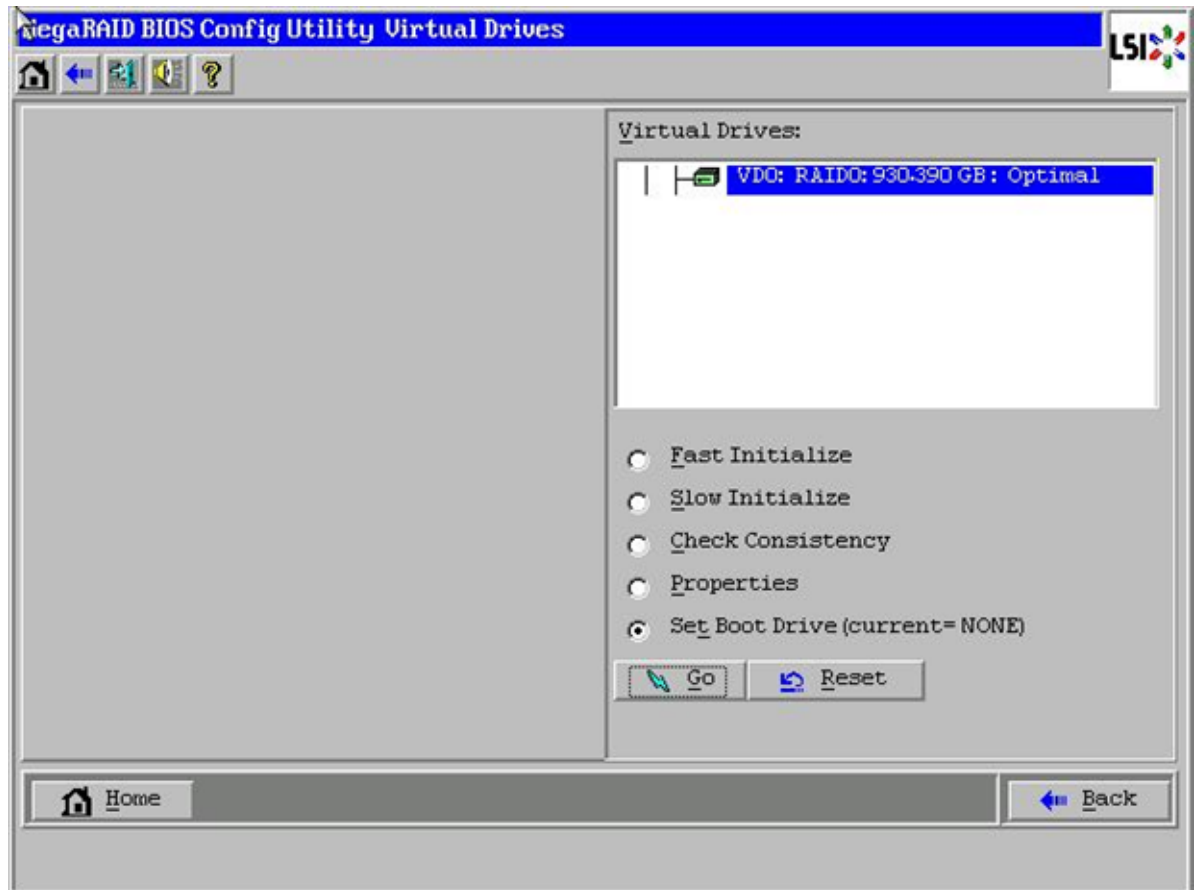
17. Click **Yes** to save the configuration.
18. In the Managing SSD Caching window, click **Cancel** as shown in [Figure 52](#).

Figure 61 SSD Caching Window



19. Click **Yes** in the confirmation page.
20. Set VD0 as the Boot Drive and click **Go** as shown in [Figure 53](#).

Figure 62 **Setting Virtual Drive as Boot Drive**



21. Click **Home**.
22. Review the configuration and click **Exit**.

Configuration of disks 3 to 24 are done using Linux based MegaCLI commands described in “[Configuring Data Drives on DataNodes](#)” section on page 100.

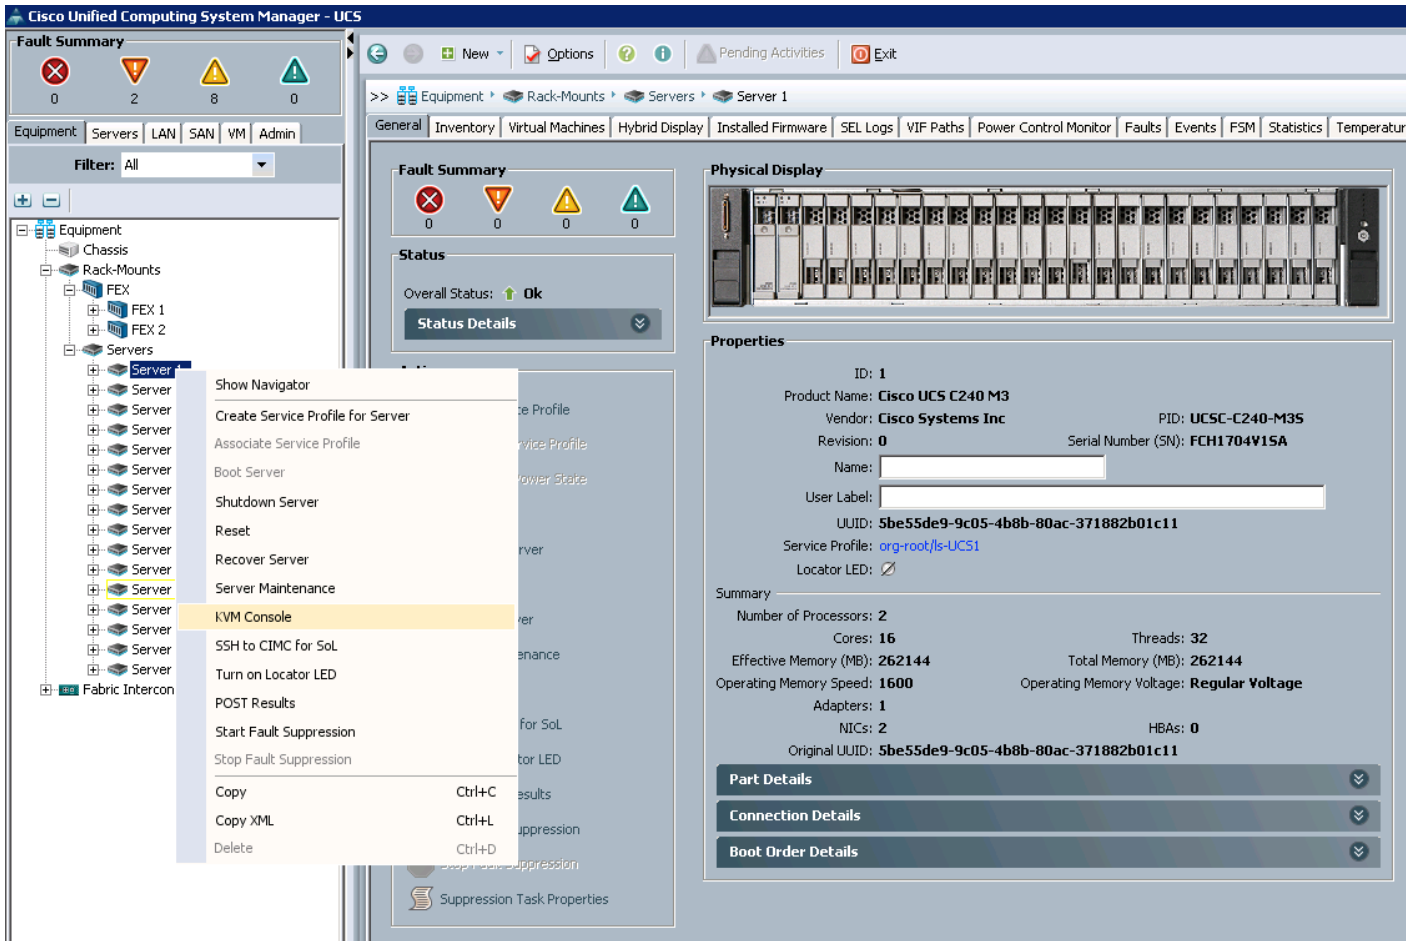
Installing Red Hat Linux 6.2 with KVM

The following section provides detailed procedures for installing Red Hat Linux 6.2.

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

1. Log in to the Cisco UCS 6296 Fabric Interconnect and launch the Cisco UCS Manager application.
2. Click the **Equipment** tab.
3. In the navigation pane expand **Rack-Mounts** and **Servers**.
4. Right-click on the **Server** and select **KVM Console** as shown in [Figure 63](#).

Figure 63 Selecting KVM Console Option



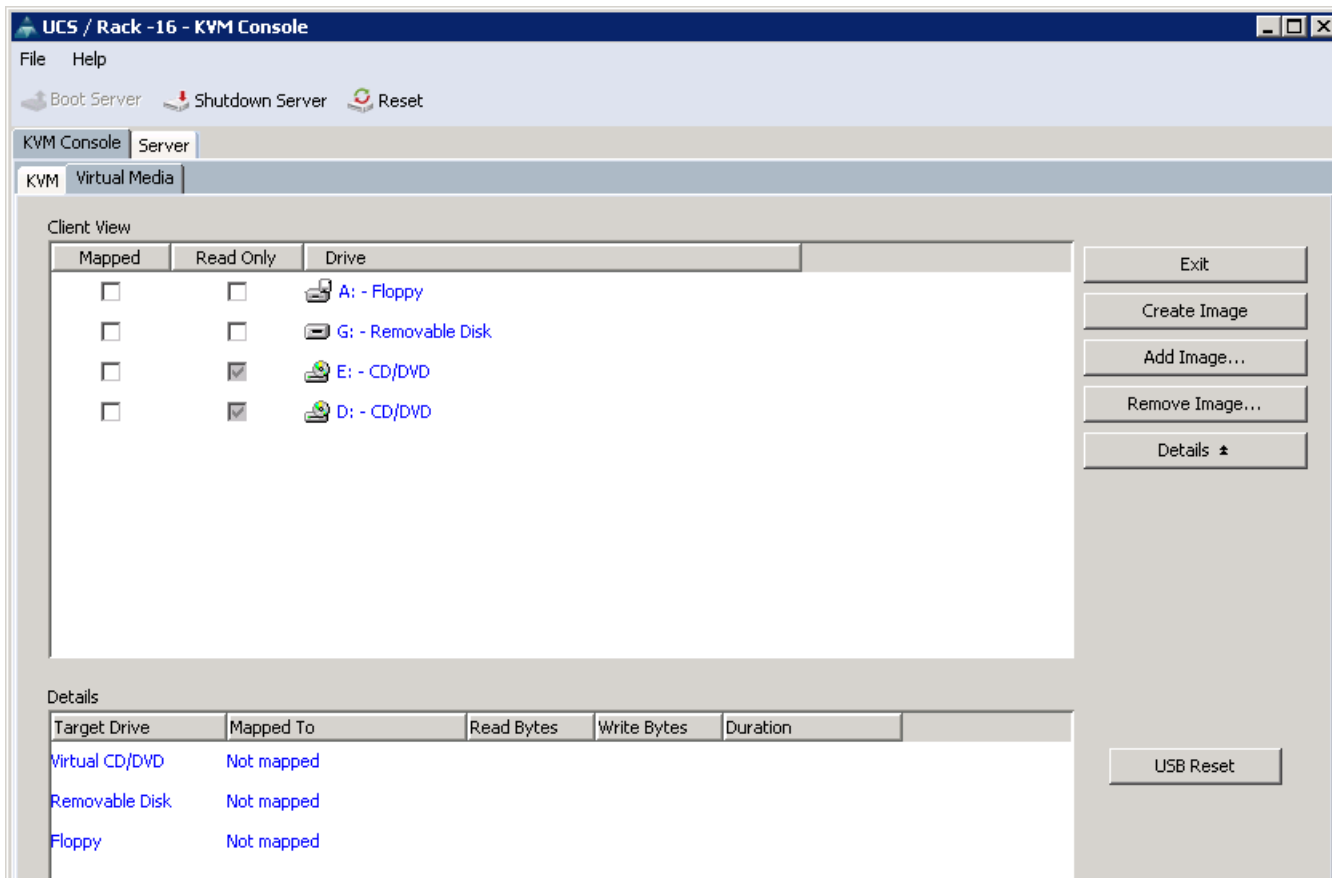
5. In the KVM window, select the **Virtual Media** tab as shown in [Figure 64](#).
6. Click **Add Image** button in the Client View selection window.
7. Browse to the Red Hat Enterprise Linux Server 6.2 installer ISO image file.



Note

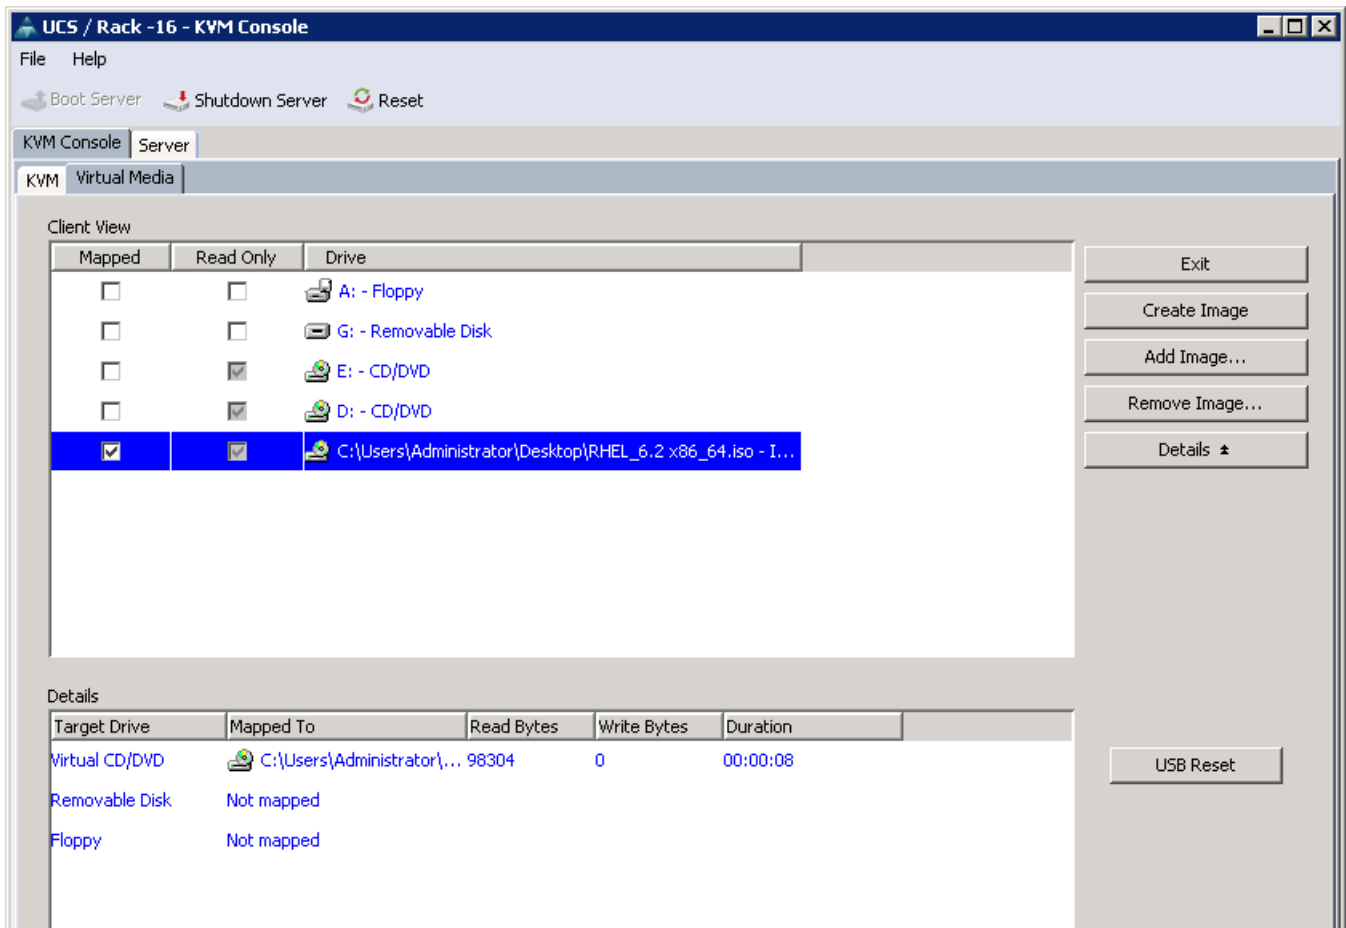
The Red Hat Enterprise Linux 6.2 DVD is assumed to be available on the client machine.

Figure 64 Adding an ISO Image



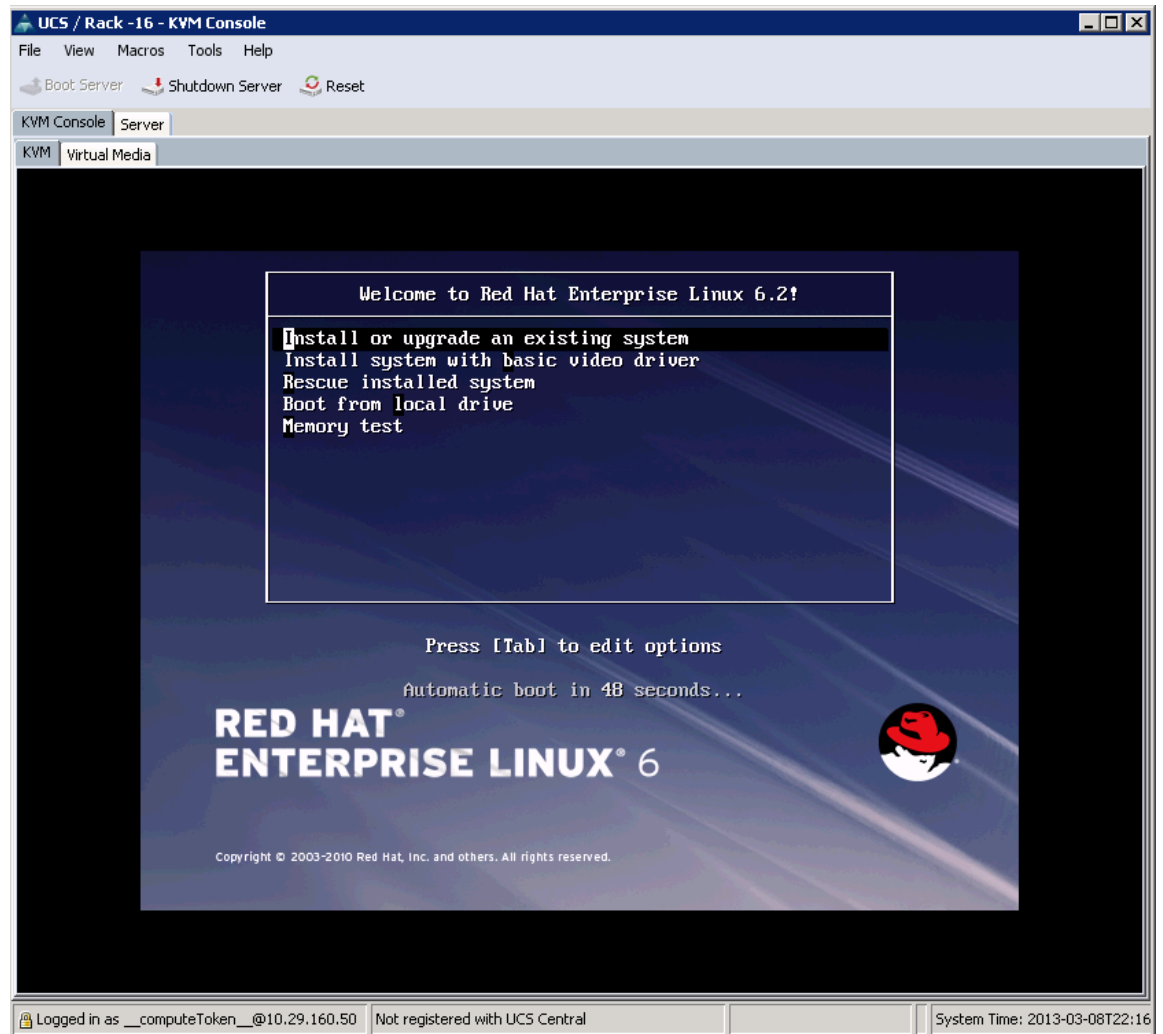
8. Click **Open** to add the image to the list of virtual media.
9. Check the **Mapped** check box for the image you just added as shown in [Figure 65](#).

Figure 65 Mapping ISO Image



10. In the **KVM** console, select the **KVM** tab to monitor the bootup.
11. In the **KVM** console, click **Boot Server**.
12. Click **OK**.
13. Click **OK** to reboot the system.
On reboot, the server detects the presence of the Red Hat Enterprise Linux Server 6.2 install media.
14. Select **Install or Upgrade an Existing System** option as shown in [Figure 66](#).

Figure 66 **Select Install Option**



15. Skip the Media test as the ISO Image is used for the installation.
16. Click **Next**. The Red Hat Linux Welcome Screen appears.
17. Select the Language for the installation.
18. Click the **Basic Storage Devices** radio button.
19. Click the **Fresh Installation** radio button.
20. Enter the host name of the server and click **Next**.
21. Click **Configure Network**. The Network Connections window appear.
22. In the Network Connections window, select the **Wired** tab.
23. Select the interface System eth0 and click **Edit**.
24. Editing System eth0 appears as shown in [Figure 67](#).
25. Check the **Connect automatically** check box.
26. Select Manual in the Method drop-down list.
27. Click **Add** and enter IP Address, the netmask and the gateway.

For this demonstration, the following values have been used:

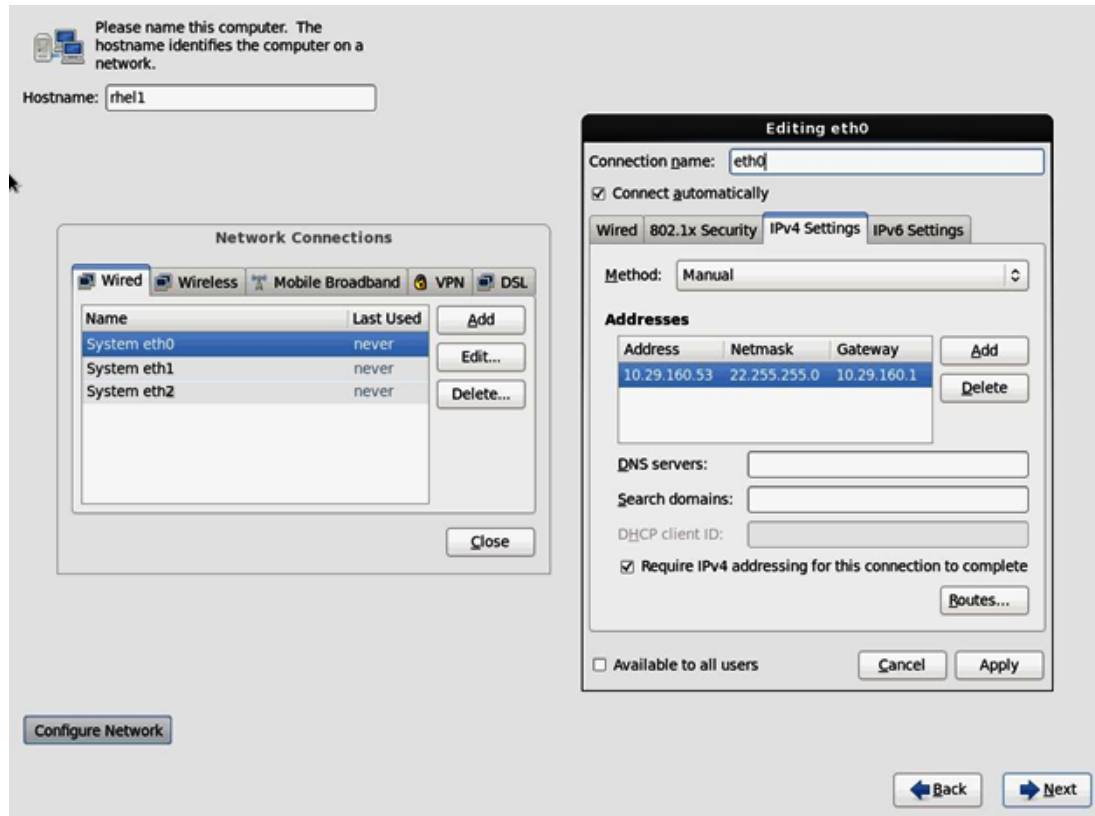
IP Address: 10.29.160.53

Netmask: 255.255.255.0

Gateway: 10.29.160.1

28. Add DNS servers (optional).
29. Click **Apply**.

Figure 67 *Configuring Network for eth0*

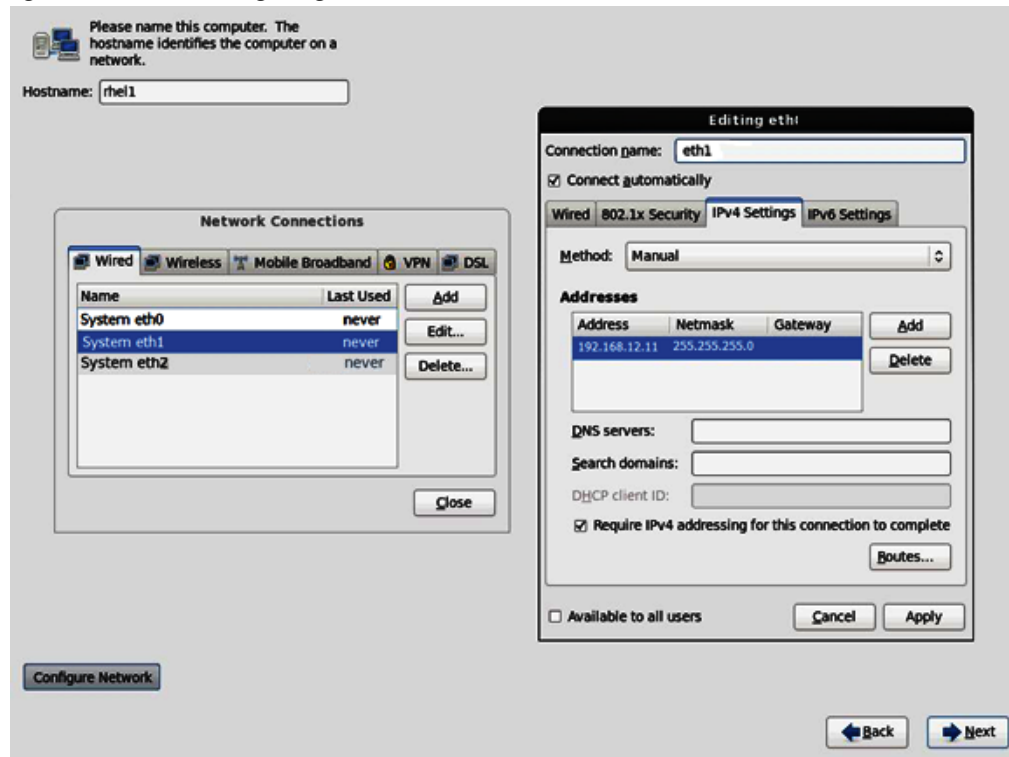


30. Repeat the steps 26 to steps 32 to configure the network for the System eth1. The following values have been used (see [Figure 68](#)):

IP Address: 192.168.12.11

Netmask: 255.255.255.0

Figure 68 **Configuring Network for eth1**



31. Repeat the steps 26 to steps 32 to configure the network for System eth2. The following values have been used:
 - IP Address: 192.168.11.11
 - Netmask: 255.255.255.0
32. Select the appropriate time zone.
33. Enter the root password and click **Next**.
34. Select **Use All Space** and click **Next** as shown in [Figure 69](#).
35. Choose an appropriate boot drive.

Figure 69 **Selecting Install Option**

Which type of installation would you like?

Use All Space
Removes all partitions on the selected device(s). This includes partitions created by other operating systems.
Tip: This option will remove data from the selected device(s). Make sure you have backups.

Replace Existing Linux System(s)
Removes only Linux partitions (created from a previous Linux installation). This does not remove other partitions you may have on your storage device(s) (such as VFAT or FAT32).
Tip: This option will remove data from the selected device(s). Make sure you have backups.

Shrink Current System
Shrinks existing partitions to create free space for the default layout.

Use Free Space
Retains your current data and partitions and uses only the unpartitioned space on the selected device(s), assuming you have enough free space available.

Create Custom Layout
Manually create your own custom layout on the selected device(s) using our partitioning tool.

Encrypt system

Review and modify partitioning layout

36. Click **Write changes to the disk** and click **Next**.

37. Select **Basic Server** and click **Next** as shown in [Figure 70](#).

Figure 70 **Selecting Type of Installation**

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

Resilient Storage

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

Customize later Customize now

38. After the installer has finished loading, it will continue with the installation.

39. Reboot the system after the installation is complete.

Repeat the above steps (1 to 39) to install the Red Hat Linux on servers 2 through 64.

**Note**

You can automate the OS installation and configuration of the nodes through the Preboot Execution Environment (PXE) boot or through third party tools.

[Table 5](#) describes the hostnames and their corresponding IP addresses.

Table 5 **Hostnames and IP Addresses**

Hostname	eth0	eth1	eth2
rhe11	10.29.160.53	192.168.12.11	192.168.11.11
rhe12	10.29.160.54	192.168.12.12	192.168.11.12
rhe13	10.29.160.55	192.168.12.13	192.168.11.13

Table 5 Hostnames and IP Addresses

Hostname	eth0	eth1	eth2
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
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
...
rhel64	10.29.160.116	192.168.12.74	192.168.11.74

Post OS Install Configuration

Choose one of the nodes of the cluster or a separate node as an Admin Node for management such as HDP installation, parallel shell, creating a local Red Hat repo and others. In this document, we have used rhel1 for management.

Setting Up Password-less Login

To manage all of the cluster 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 passwords.

Once Red Hat Linux is installed across all the nodes in the cluster, follow these steps 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.

```
[root@rhell ~]# 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@rhell
The key's randomart image is:
+--[ RSA 2048 ]-----+
| .=ooo. |
| ..+ |
| +. |
| +. |
| . +. S |
| .oo .o . |
| .o.o.o . |
| +. .o . |
| E.. .o |
+-----+
```

3. Run the following commands from the admin node to copy the public key `id_rsa.pub` to all the nodes of the cluster. The `.ssh-copy-id` command appends the keys to the remote-host.

`.ssh/authorized_key`.

```
for IP in {53..116}; do echo -n "$IP -> "; ssh-copy-id -i ~/.ssh/id_rsa.pub
10.29.160.$IP; done
```

4. Enter yes at the command prompt to continue connecting.
5. Enter the password of the remote host to login.

Installing and Configuring Parallel SSH

Installing Parallel-SSH

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

1. Download the pssh.

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

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

```
[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
```

2. Run the following command to 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
```

```
[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 and all the DataNodes in the cluster. This file is passed as a parameter to pssh to identify the nodes and run the commands on them.

```
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
...
10.29.160.116

vi /root/datanodes
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
...
10.29.160.116

```

Installing Cluster Shell

1. Download cluster shell (clush) and install it on rhel1.

Cluster shell is available from the Extra Packages for Enterprise Linux (EPEL) repository.

```

wget
http://dl.fedoraproject.org/pub/epel//6/x86_64/clustershell-1.6-1.el6.noarch.rpm
scp clustershell-1.6-1.el6.noarch.rpm rhel1:/root/

```

2. Login to rhel1 and install cluster shell.

```
yum install clustershell-1.6-1.el6.noarch.rpm
```

3. Edit /etc/clustershell/groups file to include hostnames for all the nodes of the cluster.

```
For 64 node cluster all: rhel[1-64]
```



Note

Configuring EPEL repository is discussed in detail in another section.

Configuring /etc/hosts

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

1. Run the following command to 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
192.168.12.11 rhel1
192.168.12.12 rhel2
192.168.12.13 rhel3
192.168.12.14 rhel4
192.168.12.15 rhel5
192.168.12.16 rhel6
192.168.12.17 rhel7
192.168.12.18 rhel8
192.168.12.19 rhel9
192.168.12.20 rhel10
192.168.12.21 rhel11
192.168.12.22 rhel12
192.168.12.23 rhel13
192.168.12.24 rhel14
192.168.12.25 rhel15
192.168.12.26 rhel16
...
192.168.12.74 rhel64
```

2. Run the following command to deploy /etc/hosts from the admin node (rhel1) to all the nodes:

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

```
[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
:
:
[64] 11:40:27 [SUCCESS] 10.29.160.116
```

Creating Red Hat Local Repository

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.

1. Login to rhel1 node, and run the following command to 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 rhel1.

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

Assuming the Red Hat ISO file is located in your working directory.

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

4. Copy the contents of the ISO to the **/var/www/html/rhelrepo** directory.

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

5. Run the following command on the rhel1 to create a .repo file that enables 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

The yum command based on the repo file requires httpd to be running on rhel1 so that the other nodes can access the repository.

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/
```

```
[root@rhel1 ~]# pscp -h /root/allnodes /var/www/html/rhelrepo/rheliso.repo
/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
⋮
⋮
⋮
[64] 15:00:09 [SUCCESS] 10.29.160.116
```

7. To use the repository files on rhel1 without httpd, edit the baseurl of the 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. Run `pssh -h /root/allnodes "yum clean all"` command:

```
[root@rhel1 ~]# 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
⋮
⋮
⋮
[64] 12:14:10 [SUCCESS] 10.29.160.116
```


Creating the Red Hat Repository Database

1. Install the createrepo package.
2. Use the createrepo package to regenerate the repository database(s) for the local copy of the RHEL DVD contents.
3. Purge the yum caches:

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

```
[root@rhel1 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 essential for performance and bug fixes.

To download the latest LSI drivers, see:

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

1. 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`

```
[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
```

2. Download and transfer `kmod-megaraid_sas-v06.504.01.00.rpm` driver to the admin node (rhel1).
3. Run the following commands to install the rpm on all nodes of the cluster:

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

```
[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
:
:
[64] 15:46:54 [SUCCESS] 10.29.160.116
```

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

```
[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
:
:
[64] 15:49:16 [SUCCESS] 10.29.160.116
```

4. Run the following command to verify the version of kmod-megaraid_sas driver is used on all the nodes (confirm all versions are same):

```
pssh -h /root/allnodes "modinfo megaraid_sas | head -5"
```

```

-----
filename:      /lib/modules/2.6.32-220.el6.x86_64/extra/megaraid_sas/megaraid_sas.ko
description:   LSI MegaRAID SAS Driver
author:        megaraidlinux@lsi.com
version:       06.504.01.00
license:       GPL

```

Installing httpd

1. Install httpd on the admin node to host repositories.

The Red Hat repository is hosted using http on the admin node, and this machine is accessible by all the hosts in the cluster.

```
yum -y install httpd
```

2. Add the ServerName as 10.29.160.53:80, and make the necessary changes to the server configuration file.

```
/etc/httpd/conf/httpd.conf
```

```

#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. Run the following command to make sure that the httpd is able to read the repofiles:

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

4. Run the following command to start httpd:

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

JDK Installation

Download Java SE 6 Development Kit (JDK)

Using a web browser, click on the following link:

<http://www.oracle.com/technetwork/java/index.html>

and download the latest Java™ SE 6 Development Kit (JDK™6).

Once the JDK6 package has been downloaded, place it in the /var/www/html/JDK/ directory.

Install JDK6 on All Node

Create the following script `install_jdk.sh` to install JDK:

Script `install_jdk.sh`

```
# Copy and install JDK
cd /tmp/
curl http://10.29.160.53/JDK/jdk-6u41-linux-x64.bin -O -L
sh ./jdk-6u41-linux-x64.bin -noregister
```

Copy script `disable_services.sh` to all nodes and run the script on all nodes:

```
pscp -h /root/pssh.hosts /root/install_jdk.sh /root/
pssh -h /root/pssh.hosts "/root/install_jdk.sh"
```

Extjs Installation

From a host connected to the Internet, download the Extjs and transfer it to `rhel1`.

```
wget
http://s3.amazonaws.com/public-repo-1.hortonworks.com/HDP-UTILS-1.1.0.15/repos/centos6
/extjs/extjs-2.2-1.noarch.rpm
```

Copy the `extjs` rpm to all nodes from the admin node.

```
pscp -h /root/allnodes /root/extjs-2.2-1.noarch.rpm /root/
```

Install `extjs` on all nodes.

```
pssh -h /root/allnodes "yum -y install /root/extjs-2.2-1.noarch.rpm"
```

NTP Configuration

The Network Time Protocol (NTP) is used to synchronize the time of all the nodes within the cluster. The Network Time Protocol daemon (`ntpd`) sets and maintains the system time of day in sync with the timeserver located in the admin node (`rhel1`). Configuring NTP is critical for any Hadoop Cluster. If server clocks in the cluster drift out of sync, serious problems can occur in the HBase and other services.



Note

Installing an internal NTP server keeps the cluster synchronized even when an outside NTP server is inaccessible.

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. Run the following commands to create `/root/ntp.conf` on the admin node and copy it to all the 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
```

```
[root@rhell ~]# for SERVER in {54..116}; 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
```

3. Run the following command in the admin node (rhell) to copy ntp.conf file from the admin node to /etc of all the nodes:

```
for SERVER in {54..116};
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.

4. Run the following commands to synchronize the time and restart NTP daemon on all the nodes:

```
pssh -h /root/allnodes "yum install -y ntpdate"
pssh -h /root/allnodes "service ntpd stop"
pssh -h /root/allnodes "ntpdate rhell"
pssh -h /root/allnodes "service ntpd start"
```

```
[root@rhell ~]# 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
:
:
:
[64] 13:38:55 [SUCCESS] 10.29.160.116
```

- Run the following command to restart the NTP daemon on all the reboots:

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

```
[root@rhell ~]# 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
:
:
:
[64] 13:52:55 [SUCCESS] 10.29.160.116
```

Enabling Syslog

Syslog must be enabled on each node to preserve logs regarding killed processes or failed jobs. Modern versions such as syslog-ng and rsyslog are possible, making it more difficult to ascertain if a syslog daemon is present.

Run any of the commands to confirm if the service is properly configured:

```
clush -B -a rsyslogd -v
clush -B -a service rsyslog status
```

Setting Ulimit

On each node, `ulimit -n` specifies the number of inodes that can be opened simultaneously. With the default value of 1024, the system appears to be out of disk space and shows no inodes available. This value should be set to 64000 on every node.

Higher values are unlikely to result in an appreciable performance gain.

1. For setting ulimit on Redhat, run the command **Edit /etc/security/limits.conf** and add the following lines:

```
root soft nofile 64000
root hard nofile 64000
```

2. To verify the ulimit setting, run the following command:

```
clush -B -a ulimit -n
```

The command should report 64000 as the ulimit.



Note

The ulimit values are applied on a new shell. Running the command on a node on an earlier instance of a shell shows old values.

Disabling SELinux

SELinux must be disabled during the HDP installation procedure and cluster setup. SELinux can be enabled after installation and while the cluster is running.

SELinux can be disabled by editing `/etc/selinux/config` and changing the SELINUX line to SELINUX=disabled.

1. Run the following command to disable SELINUX on all nodes:

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

```
[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
:      :      :
[64] 14:07:40 [SUCCESS] 10.29.160.116
```

```
pssh -h /root/allnodes "setenforce 0"
```

**Note**

This command fails if SELinux is already disabled.

Setting TCP Retries Parameter

Adjusting the `tcp_retries` parameter for the system network enables faster detection of failed nodes. Given the advanced networking features of UCS, this is a safe and recommended change (failures observed at the Operating System layer are mostly serious rather than transitory). On each node, set the number of TCP retries to 5 can help detect unreachable nodes with less latency.

1. Edit the file `/etc/sysctl.conf` and add the following line:

```
net.ipv4.tcp_retries2=5
```

2. Save the file and run the following command.

```
clush -B -a sysctl -p
```

Disabling the Linux Firewall

The default Linux firewall settings are far too restrictive for any Hadoop deployment. Since the Cisco UCS Big Data deployment is performed in the isolated network, there is no need to leave the iptables service running.

1. Run the following commands to disable the iptables:

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

```
[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
⋮
⋮
⋮
[64] 14:13:25 [SUCCESS] 10.29.160.116
```

2. Run the following command to check if the iptables are disabled:

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



```
[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
:
:
:
[64] 14:13:25 [SUCCESS] 10.29.160.116
```

Configuring Data Drives on NameNode

This section provides the steps to configure data drives on the NameNode.

The first two disk drives are configured for the Operating System on the nodes, rhell1 and rhell2, as shown in “[Configuring Disk Drives for Operating System on NameNode](#)” section on page 58. The remaining disk drives can be configured similarly or by using MegaCli.

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

```
scp /root/MegaCli64 rhell1:/root/
scp /root/Lib_Utills-1.00-08.noarch.rpm rhell1:/root/
scp /root/Lib_Utills2-1.00-01.noarch.rpm rhell1:/root/
```

2. To copy all the three files to all the nodes, run the following commands:

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

```
[root@rhell1 ~]# 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
⋮
[64] 13:00:40 [SUCCESS] 10.29.160.116
```

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

```
[root@rhell1 ~]# 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
⋮
[64] 13:01:26 [SUCCESS] 10.29.160.116
```

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

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

```

[root@rhell1 ~]# pssh -h /root/allnodes "rpm -ivh Lib_Utills*"
[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
:
:
:
[64] 13:02:05 [SUCCESS] 10.29.160.116

```

4. Run the following script as root user on NameNode and Secondary NameNode to create the virtual drives.

```

vi /root/raid1.sh
./MegaCli64 -cfdgdadd
r1[$1:3,$1:4,$1:5,$1:6,$1:7,$1:8,$1:9,$1:10,$1:11,$1:12,$1:13,$1:14,$1:15,$1:16,$1:17,$1:18,$1:19,$1:20,$1:21,$1:22,$1:23,$1:24] wb ra nocachedbadbbu strpsz1024 -a0

```

The above script requires enclosure ID as a parameter. Run the following command to get enclosure id.

```

./MegaCli64 pdlist -a0 | grep Enc | grep -v 252 | awk '{print $4}' | sort | uniq -c | awk '{print $2}'

```

```

chmod 755 raid1.sh

```

Run MegaCli script as follows

```

./raid1.sh <EnclosureID obtained by running the command above>

```

WB: Write back

RA: Read Ahead

NoCachedBadBBU: Do not write cache when the BBU is bad.

Strpsz1024: Strip Size of 1024K



Note

The command above will not override any existing configuration. To clear and reconfigure the existing configurations, see *Embedded MegaRAID Software Users Guide* available at: www.lsi.com.

Configuring the Filesystem for NameNodes

To Configure the filesystem for NameNodes, run the following script:

```

vi /root/driveconf.sh
#!/bin/bash
disks_count=`lsblk -id | grep sd | wc -l`
if [ $disks_count -eq 2 ]; then
    echo "Found 2 disks"
else
    echo "Found $disks_count disks. Expecting 2. Exiting.."
    exit 1
fi
[[ "-x" == "${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
        (( $? )) && continue
        /bin/mkdir -p /HDP/${Y}
        (( $? )) && continue
        /bin/mount -t xfs -o allocsize=128m,noatime,nobarrier,nodiratime ${X}1 /HDP/${Y}
        (( $? )) && continue
        echo "LABEL=${Y} /HDP/${Y} xfs allocsize=128m,noatime,nobarrier,nodiratime 0 0" >>
        /etc/fstab
    fi
done
    
```

Configuring Data Drives on DataNodes

This section provides the steps to configure data drives on DataNodes.

The first disk drive is configured for the Operating System on all the DataNodes, rhel3 to rhel64 as shown in [“Configuring Disk Drives for Operating System on DataNodes” section on page 65](#). The remaining disk drives can be configured similarly or by using MegaCli.

Run the following command from the admin node to create the virtual drives with RAID 0 configurations on all the DataNodes.

```

pssh -h /root/datanodes "./MegaCli64 -cfgeachdiskraid0 WB RA direct NoCachedBadBBU
strpsz1024 -a0"
    
```

WB: Write back

RA: Read Ahead

NoCachedBadBBU: Do not write cache when the BBU is bad

Strpsz1024: Strip Size of 1024K

**Note**

The above command will not override existing configurations. To clear and reconfigure the existing configurations, see *Embedded MegaRAID Software Users Guide* available at: www.lsi.com.

Configuring the Filesystem for DataNodes

This section describes the procedure to configure the filesystem for DataNodes.

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

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

```
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
        (( $? )) && continue
        /bin/mkdir -p /HDP/${Y}
        (( $? )) && continue
        /bin/mount -t xfs -o allocsize=128m,noatime,nobarrier,nodiratime ${X}1
        /HDP/${Y}
        (( $? )) && continue
        echo "LABEL=${Y} /HDP/${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 DataNodes.

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

3. Run the following command from the admin node to run the script across all DataNodes.

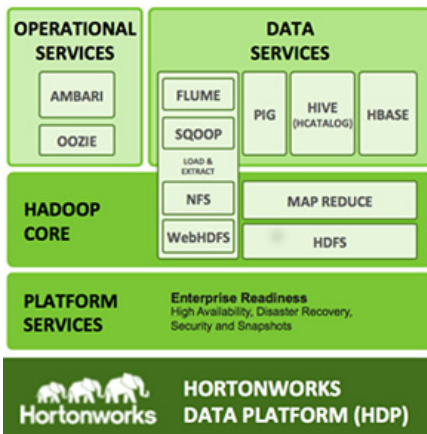
```
pssh -h /root/datanodes "./driveconf.sh"
```

```
[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
:
:
[64] 16:15:35 [SUCCESS] 10.29.160.116
```

Installing HDP

HDP is an enterprise grade, hardened Hadoop distribution. HDP combines Apache Hadoop and its related projects into a single tested and certified package. It offers the latest innovations from the open source community with the testing and quality you expect from the enterprise quality software. HDP components are depicted in [Figure 71](#).

Figure 71 HDP Components



Prerequisites for HDP Installation

This section details the prerequisites for HDP installation such as setting up of EPEL and HDP Repo.

Hortonworks and EPEL Repo

From a host connected to the Internet, download the EPEL and Hortonworks repositories as shown below and transfer it to rhel1.

1. Download EPEL Repository from the system connected to the Internet.

```
mkdir -p /tmp/Hortonworks
cd /tmp/Hortonworks
rpm -Uvh
http://download.fedoraproject.org/pub/epel/6/x86_64/epel-release-6-8.noarch.rpm
reposync -r epel
```

2. Download Hortonworks HDP Repo

```
wget
http://public-repo-1.hortonworks.com/HDP/centos6/HDP-1.3.0.0-centos6-rpm.tar.gz
```

3. Download Hortonworks HDP-Utils Repo

```
wget
http://public-repo-1.hortonworks.com/HDP-UTILS-1.1.0.15/repos/centos6/HDP-UTILS-1.1.0.15-centos6.tar.gz
```

4. Download Ambari Repo

```
wget
http://public-repo-1.hortonworks.com/ambari/centos6/ambari-1.2.3.7-centos6.tar.gz
```

5. Copy the repository directory to admin node

```
scp -r /tmp/Hortonworks/ rhel1:/var/www/html
```

6. Extract the files

```
login to rhel1
cd /var/www/html/Hortonworks
tar -zxvf HDP-1.3.0.0-centos6-rpm.tar.gz
tar -zxvf HDP-UTILS-1.1.0.15-centos6.tar.gz
tar -zxvf ambari-1.2.3.7-centos6.tar.gz
```

7. Create the hdp.repo file with following contents:

```
vi /etc/yum.repos.d/hdp.repo
[HDP-1.3.0.0]
name=Hortonworks Data Platform Version - HDP-1.3.0.0
baseurl=http://10.29.160.53/Hortonworks/HDP/centos6/1.x/GA/1.3.0.0/
gpgcheck=0
enabled=1
priority=1
```

8. Create the hdp-utils.repo file with following contents:

```
vi /etc/yum.repos.d/hdp-utils.repo
[HDP-UTILS-1.1.0.15]
name=Hortonworks Data Platform Version -HDP-UTILS-1.1.0.15
baseurl=http://10.29.160.53/Hortonworks/HDP-UTILS-1.1.0.15/repos/centos6
gpgcheck=0
enabled=1
priority=1
```

9. Create the Ambari repo file with following contents:

```
vi /etc/yum.repos.d/ambari.repo
[Updates-ambari-1.2.3.7]
name=ambari-1.2.3.7 - Updates
```

```
baseurl=http://rhell/Hortonworks/ambari/centos6/1.x/updates/1.2.3.7
gpgcheck=0
enabled=1
priority=1
```

10. Create epel.repo

```
cd /var/www/html/Hortonworks/epel
createrepo .
vi /etc/yum.repos.d/epel.repo
name=Extra Packages for Enterprise Linux 6 - $basearch
baseurl=http://rhell/Hortonworks/epel/
enabled=1
gpgcheck=0
priority=1
```

From the admin node copy the repo files to `/etc/yum.repos.d/` of all the nodes of the cluster.

```
pscp -h /root/allnodes /etc/yum.repos.d/hdp* /etc/yum.repos.d/
pscp -h /root/allnodes /etc/yum.repos.d/ambari.repo /etc/yum.repos.d/
pscp -h /root/allnodes /etc/yum.repos.d/epel.repo /etc/yum.repos.d/
```

HDP Installation

To install HDP, issue the following CLI commands:

Install and Setup Ambari Server on rhel1

```
yum install ambari-server
```

Setup Ambari Server

```
ambari-server setup -j $JAVA_HOME
```

```
[root@rhell ~]# ambari-server setup -j $JAVA_HOME
Using python /usr/bin/python2.6
Run postgresql initdb
Run postgresql start
Starting postgresql service: [ OK ]
Setup ambari-server
Checking SELinux...
SELinux status is 'disabled'
Checking iptables...
iptables is disabled now
Checking PostgreSQL...
Configuring database...
Configuring PostgreSQL...
Backup for pg_hba found, reconfiguration not required
Checking JDK...
WARNING: JAVA_HOME /usr/java/jdk1.6.0_37 must be valid on ALL hosts
Completing setup...
Ambari Server 'setup' finished successfully
```


Configure Ambari Server to use Local Repository

Edit `redhat6` and `centos6` sections of the Ambari `repoinfo.xml` to point to local repository.

```
vi /var/lib/ambari-server/resources/stacks/HDP/Local/1.3.0/repos/repoinfo.xml
```

Replace the xml element `<os type="redhat6"> .. </os>` with

```
<os type="redhat6">
  <repo>
    <baseurl>http://10.29.160.53/Hortonworks/HDP/centos6/1.x/GA/1.3.0.0</baseurl>
    <repoid>HDP-1.3.0</repoid>
    <reponame>HDP</reponame>
  </repo>
  <repo>
    <baseurl>http://10.29.160.53/Hortonworks/epel</baseurl>
    <repoid>HDP-epel</repoid>
    <reponame>HDP-epel</reponame>

  <mirrorslist><![CDATA[http://mirrors.fedoraproject.org/mirrorlist?repo=epel-6&arch=$basearch]]></mirrorslist>
</os>
```

Replace the xml element `<os type="centos6"> .. </os>` with

```
<os type="centos6">
  <repo>
    <baseurl>http://10.29.160.53/Hortonworks/HDP/centos6/1.x/GA/1.3.0.0</baseurl>
    <repoid>HDP-1.3.0</repoid>
    <reponame>HDP</reponame>
  </repo>
  <repo>
    <baseurl>http://10.29.160.53/Hortonworks/epel</baseurl>
    <repoid>HDP-epel</repoid>
    <reponame>HDP-epel</reponame>

  <mirrorslist><![CDATA[http://mirrors.fedoraproject.org/mirrorlist?repo=epel-6&arch=$basearch]]></mirrorslist>
</os>
```

Start Ambari Server

```
ambari-server start
```

Confirm Ambari Server Startup

```
ps -ef | grep ambari-server
```

Login to Ambari Server

Once the Ambari service has been started, access the Ambari Install Wizard through the browser.

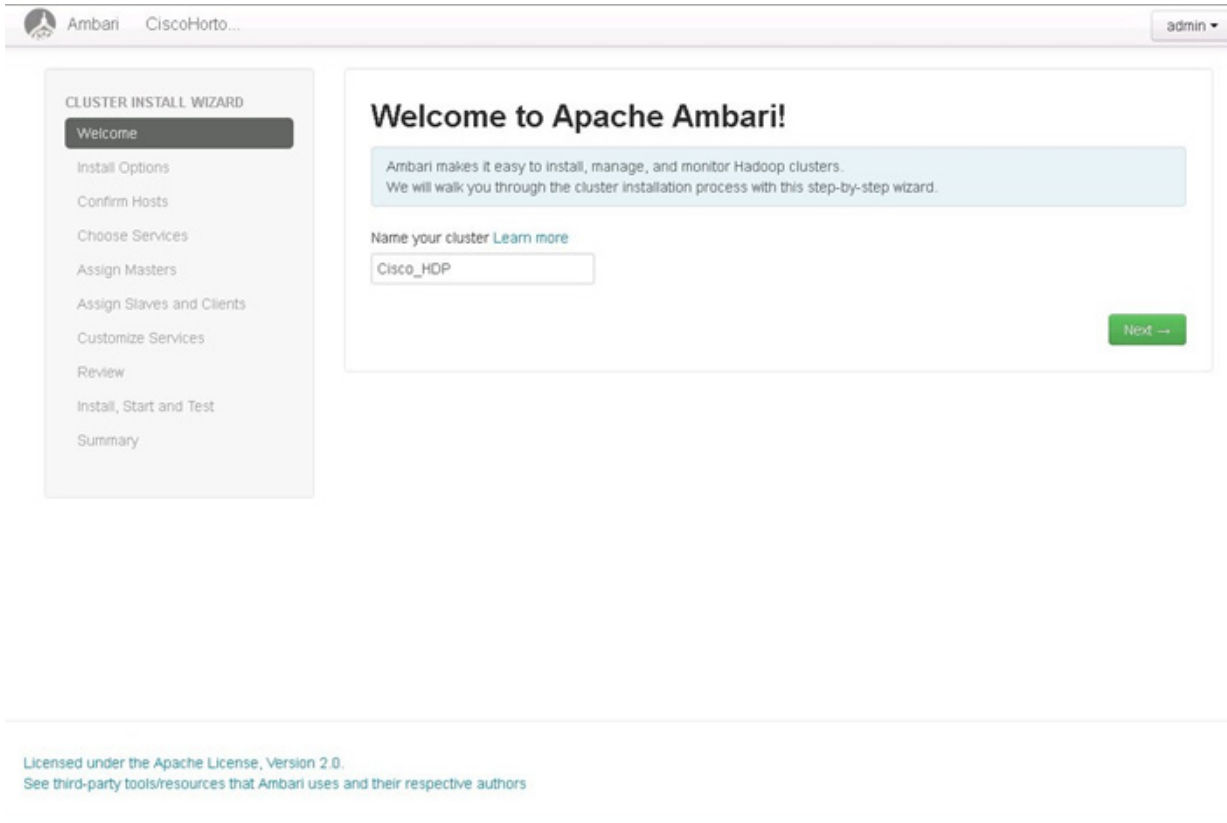
1. Go to `http://<ip address for rhel1>:8080`.
2. Login to the Ambari Server using the default username/password as `admin/admin`, which can be changed at a later period of time.

Create Cluster Name

Follow these steps to create the cluster name:

1. In the Ambari Server home page, enter Cisco_HDP in the name field.

Figure 72 Apache Ambari - Home Page



HDP Cluster Installation

In order to build a cluster, the install wizard needs to have a general information about how the cluster needs to be set up. For this, you need to provide the Fully Qualified Domain Name (FQDN) of each one of the hosts. The wizard also needs to access the private key file that was created in “[Setting Up Password-less Login](#)” section on page 82. It uses these to locate all the hosts in the system and to access and interact with them securely.

1. Use the Target Hosts text box to enter the list of host names, one by one. You can also give a range within the brackets to indicate larger sets of hosts.
2. Select the option Provide your SSH Private Key in the Ambari cluster install wizard.
 - a. Copy the contents of the file `/root/.ssh/id_rsa` on `rhel1` and paste it in the text area provided by the Ambari cluster install wizard.



Note

Make sure there is no extra white space after the text-----END RSA PRIVATE KEY-----

Figure 73 Copying Contents from `/root/.ssh/id_rsa`

```
[root@rhell1 ~]# cat /root/.ssh/id_rsa
-----BEGIN RSA PRIVATE KEY-----
MIIEoQIBAAKCAQEAYDOIRbk4mBZrirc0/gOM2iYT2h4vxkIXA/uvQVPthFreUdgT
Zehw/Qtdk7meeqhgqsHmb1CriF0m6SxvPEXW2cGoAx75hZwTuDIR3Qlvk6oYUmDW
BKq5TMfUMKfD7tknkGkg5N+YHsPCoNILLz/Wqc01hZZ0tiCmrxeRnPGS1JY74/Db
A0BewMuNajAoVppPD6cLGF6/NKORpEDUnCuwe5pCRV5tko+gzBeBF5oeCS6Ya6I7
ns0HplJXV0Mv23SNUwl3cswbqLdrr3atG6YrieVrmmr/PlrKmp192tzQ1mHZMBqG
w1RJTILjyqW0gp5g7NQBGeM7sX4V6Omzv4vmzwIBIwKCAQEAg4+UEI+o2PjKVCuX
2h+XEwMUXCJ3KoneYbPr2nj7KxckYas/8oLN6B1pYROUB3X2YZVc6hBwuLI+JDMk
hrGNMALqWdjtHU10yX/9HDlmlDyTo9k8LvPY2q8zqvHnJ+3Jisi92Dspc01xRRxQ
wnpofjAm1CDx5Wxp4MZYX9HynCcKmheFefobLys6gloxd84eHW1y6b0xU1dh7hsQ
pcK+xpDfW1sHYFbvckTuCHUAezF4+uBT5F0PMid7PwzrvbXKA65ABuezv9gg2/I1
PekIkRvbosniFbBUi2ZOS1uN/gsaZgmSQ9gTarJlV8zMy6K31LEtcock12LzHRX2
5sEx6wKBgQD9CiKc0HfiulrQWW5cLTDJU8wzTINK4M91Qb2LohfFuZfluiA13Ref
yiL9MjE3A5Mnn9pcRxMmXXPF4t9iuh3+3tCsr1TzPml4WT+Fipa9sh+3J22HKgm
pCquAEdoFRK4oP3/yYQg95gie2SC9sB0z6zVohdyNUvnkiMb9vwi3wKBgQDKiyTi
Yu421owsYKfZ7YjomjRKUFaH4CKtnyJy1SM3wFPRnZJd4BUaMq0DaTxr2tW4si+4
t88M8XS6FHGHymSqRtL0tYzMLmmwUtjCLN2QfqSeg1NovekXXL0iUzel8PL3Z0H
AeBj0/GLQ3SF/PGWMokCwNtaJoV/xldBdIsqEQKBgEERPBmx8UVF3NZ9ZYvtMYO
09KtsU3Ex52x0ad1Vpht5Tssmolkv06TEE+8cw4lfZx5j+vXwxh+bjozBj30/Dwc
GGGbrQbrkKscs5HLL3Z5+QqtWepB4hiQnUKvnVVHP1QMJA6S53YxCdz7KHlypnqq
bkWQfKhw2QEiUivDKuRlAoGASzr/EkIAtUfFb5GdbjOn4V3Y6Gb7ky3DvNS1BhSm
rk7ADAdTnzX5NZ3L08gAf9Tws+ppfx+zTfNIOMFmNY1Y9EpyJs0S/1adLEoroWu
sC8J8bu/5RNWk8z+z9s5zwUrd5tXT2cY1J8t1KQgtWyUPxoVoe/ccfENA5LP872S
xnsCgYAFRE4SbB416p9miir1+gNCiHM9N+FmHmMcP/y80QL/MoAYoHB1Tn8cwVu
l+sju4bWGUZvnGMWxwPEU5zVBra+yShh309IwjP/1kpCNWz7CX+/ui6FY+s1ZxTr
t5P/Avh0vUKMhRFjXFQoY5yqNUkasvIu6S8Q1unl8N2IhEgw1g==
-----END RSA PRIVATE KEY-----
```

3. Check the check boxes Use Local Software Repository and path to 64-bit JDK.
4. Provide a path to java install directory.
5. Click **Register and Confirm** to continue.

Figure 74 *Confirming the HDP Installation*

CLUSTER INSTALL WIZARD

- Welcome
- Install Options
- Confirm Hosts
- Choose Services
- Assign Masters
- Assign Slaves and Clients
- Customize Services
- Review
- Install, Start and Test
- Summary

Install Options

Enter the list of hosts to be included in the cluster and provide your SSH key.

Target Hosts

Enter a list of hosts using the Fully Qualified Domain Name (FQDN), one per line. Or use [Pattern Expressions](#)

```
rhe[1-64]
```

Host Registration Information

Provide your [SSH Private Key](#) (id_rsa for root) and use SSH to automatically register hosts

```
|+sju4bWGUZvnGMWxwpEU5zvBra+yShh309lwjP/1kpCNWz7CX+
/ul6FY+sjZxTt
t5P/AvhOvUKMhRFjxFQoY5yqNUkasvlu6S8Q1uni8N2ihEgw1g==
```

Perform [manual registration](#) on hosts and do not use SSH

Advanced Options

Use a [Local Software Repository](#) instead of downloading software packages from the internet

Path to 64-bit JDK [JAVA_HOME](#)

```
/usr/java/jdk1.6.0_37
```

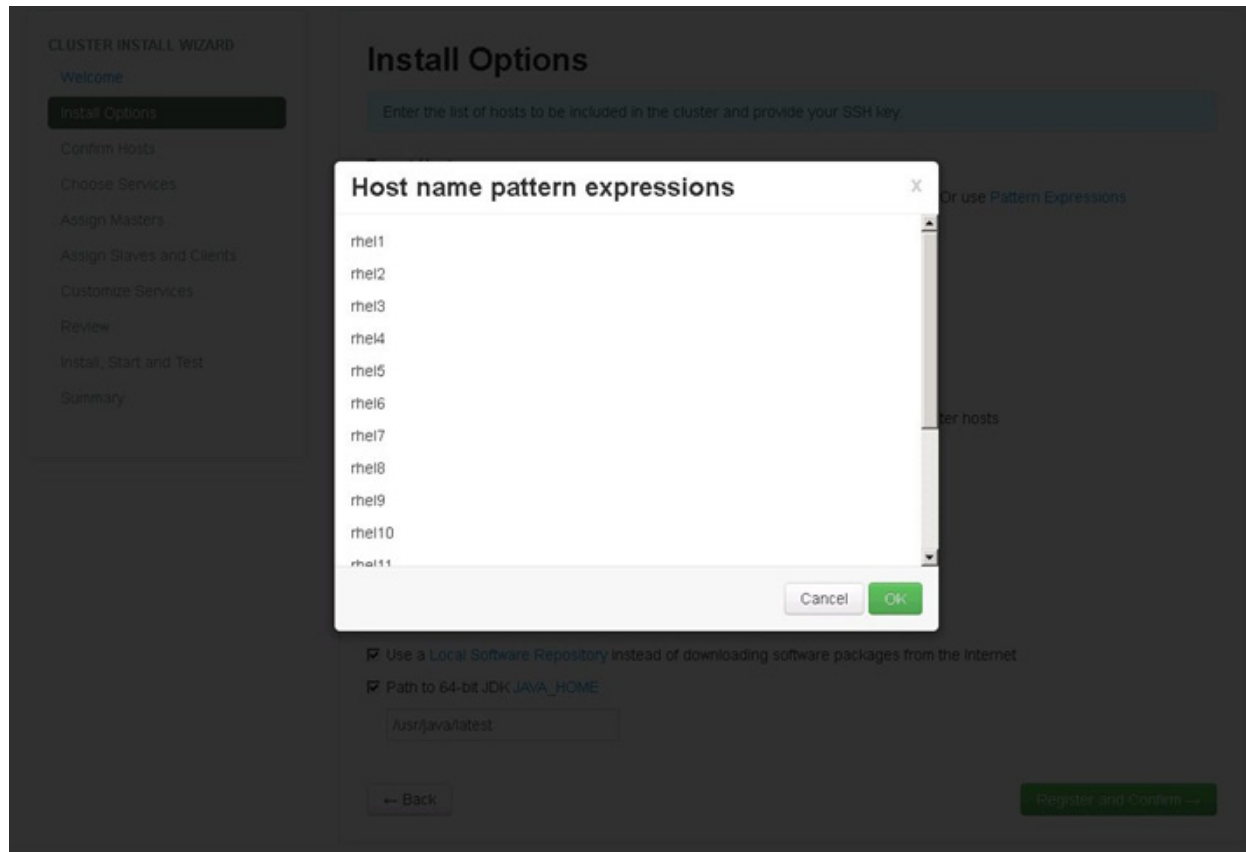
[← Back](#)

[Register and Confirm →](#)

Host Name Pattern Expressions

[Figure 75](#) shows a list of target host names using pattern expressions.

Figure 75 Host Name Pattern Expressions



Confirming Hosts

This screen allows you to make sure that the Ambari server has located all the required hosts for the cluster and to make sure that the hosts have correct directories, packages, and processes to continue the installation process.

You can remove the undesired hosts that were selected by the Ambari server. To remove all the undesired hosts, check the appropriate check boxes provided against each of the hosts and then click

. To remove a single host, click .

Figure 76 *Confirming Hosts to be Included in the Cluster*

Confirm Hosts

Registering your hosts.
Please confirm the host list and remove any hosts that you do not want to include in the cluster.

Remove Selected Show: All (64) | Installing (0) | Registering (0) | Success (60) | Fail (0)

<input type="checkbox"/>	Host	Progress	Status	Action
<input type="checkbox"/>	rhe11	<div style="width: 100%; background-color: green;"></div>	Success	<input type="button" value="Remove"/>
<input type="checkbox"/>	rhe12	<div style="width: 100%; background-color: green;"></div>	Success	<input type="button" value="Remove"/>
<input type="checkbox"/>	rhe13	<div style="width: 100%; background-color: green;"></div>	Success	<input type="button" value="Remove"/>
<input type="checkbox"/>	rhe14	<div style="width: 100%; background-color: green;"></div>	Success	<input type="button" value="Remove"/>
<input type="checkbox"/>	rhe15	<div style="width: 100%; background-color: green;"></div>	Success	<input type="button" value="Remove"/>
<input type="checkbox"/>	rhe16	<div style="width: 100%; background-color: green;"></div>	Success	<input type="button" value="Remove"/>
<input type="checkbox"/>	rhe17	<div style="width: 100%; background-color: green;"></div>	Success	<input type="button" value="Remove"/>
<input type="checkbox"/>	rhe18	<div style="width: 100%; background-color: green;"></div>	Success	<input type="button" value="Remove"/>
<input type="checkbox"/>	rhe19	<div style="width: 100%; background-color: green;"></div>	Success	<input type="button" value="Remove"/>
<input type="checkbox"/>	rhe110	<div style="width: 100%; background-color: green;"></div>	Success	<input type="button" value="Remove"/>

Some warnings were encountered while performing checks against the above hosts.
[Click here to see the warnings.](#)

Choose Services

HDP is made up of a number of components. See [Understand the Basics](#) for more information.

1. Select all to preselect all items.
2. When you have made your selections, click **Next**.

Figure 77 Choosing Services for the Cluster

CLUSTER INSTALL WIZARD

- Welcome
- Install Options
- Confirm Hosts
- Choose Services**
- Assign Masters
- Assign Slaves and Clients
- Customize Services
- Review
- Install, Start and Test
- Summary

Choose Services

Choose which services you want to install on your cluster.

Service	all minimum	Version	Description
<input checked="" type="checkbox"/> HDFS		1.2.0.1.3.0.0	Apache Hadoop Distributed File System
<input checked="" type="checkbox"/> MapReduce		1.2.0.1.3.0.0	Apache Hadoop Distributed Processing Framework
<input checked="" type="checkbox"/> Nagios		3.2.3	Nagios Monitoring and Alerting system
<input checked="" type="checkbox"/> Ganglia		3.2.0	Ganglia Metrics Collection system
<input checked="" type="checkbox"/> Hive + HCat + ZooKeeper		0.11.0.1.3.0.0	Data warehouse system for ad-hoc queries & analysis of large datasets and table & storage management service
<input checked="" type="checkbox"/> HBase + ZooKeeper		0.94.6.1.3.0.0	Non-relational distributed database and centralized service for configuration management & synchronization
<input checked="" type="checkbox"/> Pig		0.11.1.1.3.0.0	Scripting platform for analyzing large datasets
<input checked="" type="checkbox"/> Sqoop		1.4.3.1.3.0.0	Tool for transferring bulk data between Apache Hadoop and structured data stores such as relational databases
<input checked="" type="checkbox"/> Oozie		3.3.2.1.3.0.0	System for workflow coordination and execution of Apache Hadoop jobs

[← Back](#) [Next →](#)

Assign Masters

The Ambari install wizard attempts to assign the master nodes for various services that have been selected for the appropriate hosts in the cluster. [Figure 78](#) shows the current service assignments by the host, the hostname and its number of CPU cores and RAM size.

1. Reconfigure the service assignment to match the [Table 6](#):

Table 6 Service Assignment

Service Name	Host
NameNode	rhel1
SNameNode	rhel2
JobTracker	rhel2
Nagios Server	rhel1
Ganglia Collector	rhel1
Hive Server2	rhel2
HBase Master	rhel2
Oozie Server	rhel1
ZooKeeper	rhel1, rhel2, rhel3



Note On a small cluster (<16 nodes), consolidate all the master services to run on a single node.

2. Click **Next**.

Figure 78 Assigning Master Components

CLUSTER INSTALL WIZARD

- Welcome
- Install Options
- Confirm Hosts
- Choose Services
- Assign Masters**
- Assign Slaves and Clients
- Customize Services
- Review
- Install, Start and Test
- Summary

Assign Masters

Assign master components to hosts you want to run them on.
 • HiveServer2, Hive Metastore, and WebHCat Server will be hosted on the same server.

NameNode: rhe1 (252.3 GB, 32 cores)

SNameNode: rhe2 (252.3 GB, 32 cores)

JobTracker: rhe2 (252.3 GB, 32 cores)

Nagios Server: rhe1 (252.3 GB, 32 cores)

Ganglia Collector: rhe1 (252.3 GB, 32 cores)

HiveServer2: rhe2 (252.3 GB, 32 cores)

Hive Metastore: rhe2

WebHCat Server: rhe2

HBase Master: rhe2 (252.3 GB, 32 cores)

Oozie Server: rhe1 (252.3 GB, 32 cores)

ZooKeeper: rhe1 (252.3 GB, 32 cores)

ZooKeeper: rhe2 (252.3 GB, 32 cores)

ZooKeeper: rhe3 (252.3 GB, 32 cores)

Hosts and their assigned components:

- rhe1 (252.3 GB, 32 cores): NameNode, Nagios Server, Ganglia Collector, Oozie Server, ZooKeeper
- rhe2 (252.3 GB, 32 cores): SNameNode, JobTracker, HiveServer2, Hive Metastore, WebHCat Server, HBase Master, ZooKeeper
- rhe3 (252.3 GB, 32 cores): ZooKeeper

4 hosts not running master services

← Back Next →

Assign Slaves and Clients

The Ambari install wizard attempts to assign the slave components (DataNodes, TaskTrackers, and RegionServers) to appropriate hosts in the cluster. Reconfigure the service assignment to match [Figure 79](#):

1. Assign DataNode, TaskTracker, RegionServer nodes rhe13 to rhe164.
2. Assign Client to all nodes.
3. Click **Next**.

Figure 79 Assigning Slave and Client Components to Hosts

Assign slave and client components to hosts you want to run them on.
Hosts that are assigned master components are shown with •.
Client will install HDFS Client, MapReduce Client, Hive Client, HCat Client, HBase Client, Pig, Sqoop, Oozie Client and ZooKeeper Client.

Host	all none	all none	all none	all none
rhe1 •	<input type="checkbox"/> DataNode	<input type="checkbox"/> TaskTracker	<input type="checkbox"/> RegionServer	<input checked="" type="checkbox"/> Client
rhe2 •	<input type="checkbox"/> DataNode	<input type="checkbox"/> TaskTracker	<input type="checkbox"/> RegionServer	<input checked="" type="checkbox"/> Client
rhe3 •	<input checked="" type="checkbox"/> DataNode	<input checked="" type="checkbox"/> TaskTracker	<input checked="" type="checkbox"/> RegionServer	<input checked="" type="checkbox"/> Client
rhe4	<input checked="" type="checkbox"/> DataNode	<input checked="" type="checkbox"/> TaskTracker	<input checked="" type="checkbox"/> RegionServer	<input checked="" type="checkbox"/> Client
rhe5	<input checked="" type="checkbox"/> DataNode	<input checked="" type="checkbox"/> TaskTracker	<input checked="" type="checkbox"/> RegionServer	<input checked="" type="checkbox"/> Client
rhe6	<input checked="" type="checkbox"/> DataNode	<input checked="" type="checkbox"/> TaskTracker	<input checked="" type="checkbox"/> RegionServer	<input checked="" type="checkbox"/> Client
rhe7	<input checked="" type="checkbox"/> DataNode	<input checked="" type="checkbox"/> TaskTracker	<input checked="" type="checkbox"/> RegionServer	<input checked="" type="checkbox"/> Client
rhe8	<input checked="" type="checkbox"/> DataNode	<input checked="" type="checkbox"/> TaskTracker	<input checked="" type="checkbox"/> RegionServer	<input checked="" type="checkbox"/> Client
rhe9	<input checked="" type="checkbox"/> DataNode	<input checked="" type="checkbox"/> TaskTracker	<input checked="" type="checkbox"/> RegionServer	<input checked="" type="checkbox"/> Client
rhe10	<input checked="" type="checkbox"/> DataNode	<input checked="" type="checkbox"/> TaskTracker	<input checked="" type="checkbox"/> RegionServer	<input checked="" type="checkbox"/> Client

← Back Next →

Customize Services

Customize Services window in the cluster install wizard presents a number of configuration settings to manage Hadoop components. The configuration settings can be done based on your requirements under each of the tabs as shown in Figure 132. This window shows the default settings for each of the configuration options, but you can modify the settings to meet specific requirements.

Following are the configurations available in the cluster install wizard:

- [HDFS, page 113](#)
- [MapReduce, page 115](#)
- [Hive/HCat, page 117](#)
- [WebHCat, page 118](#)
- [HBase, page 119](#)
- [ZooKeeper, page 120](#)
- [Oozie, page 121](#)
- [Nagios, page 122](#)
- [Misc, page 123](#)

The following sections provide details of each of these configurations.

HDFS

Update the HDFS configurations as shown in [Table 7](#), and [Figure 80](#) and [Figure 81](#):

Table 7 HDFS Configurations

Property Name	Value
NameNode Java Heap Size	4GB
Reserved space for HDFS	4GB
DataNode Volumes Failure Toleration	5

Figure 80 Customize Services - HDFS Configuration Window Part 1

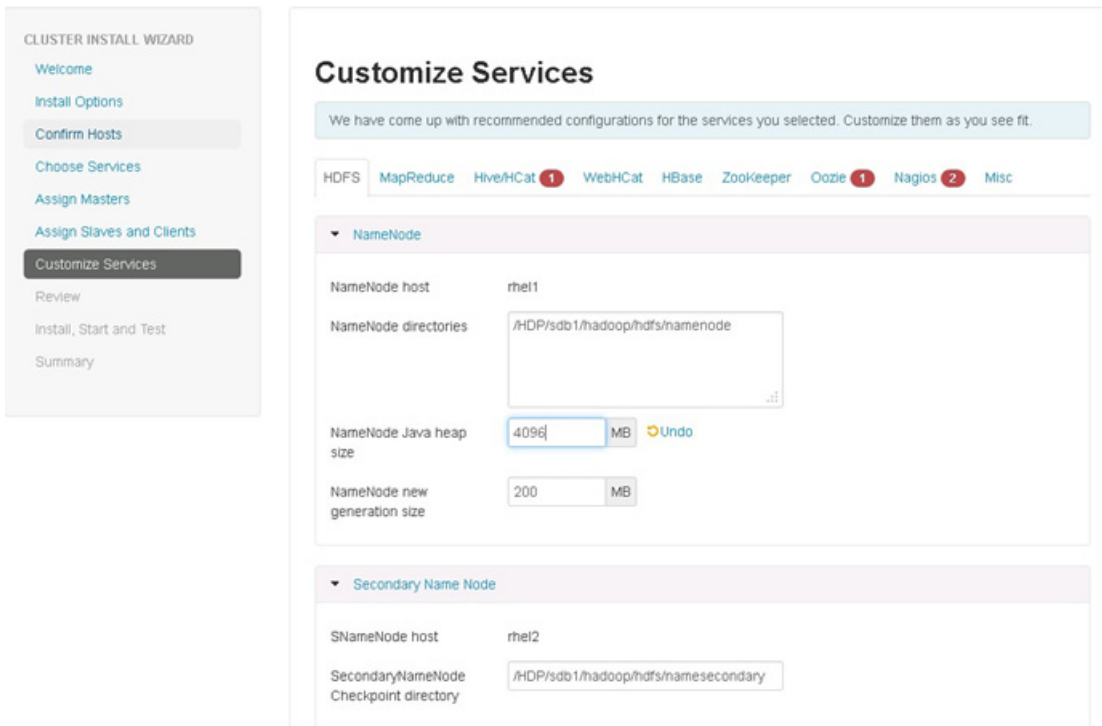


Figure 81 Customize Services - HDFS Configuration Window Part 2

▼ DataNode

DataNode hosts rhel3 and 61 others

DataNode directories

/HDP/sdb1/hadoop/hdfs/data
 /HDP/sdc1/hadoop/hdfs/data
 /HDP/sdd1/hadoop/hdfs/data
 /HDP/sde1/hadoop/hdfs/data
 /HDP/sdf1/hadoop/hdfs/data
⋮

DataNode maximum Java heap size MB

DataNode volumes failure toleration Undo

▼ General

WebHDFS enabled Undo

Hadoop maximum Java heap size MB

Reserved space for HDFS GB

HDFS Maximum Checkpoint Delay seconds

HDFS Maximum Edit Log Size for Checkpointing GB

MapReduce

Update the MapReduce configuration as shown in [Table 8](#), and [Figure 82](#) and [Figure 83](#):

Table 8 MapReduce Configurations

Property Name	Value
Job Tracker Maximum Java Heap Size	4GB
Number of Map Slots per Node	24
Number of Reduce Slots per Node	12
Java Options for MapReduce Tasks	4GB
Map-side sort buffer memory	1GB

Figure 82 Customize Services - MapReduce Configuration Window Part 1

HDFS
MapReduce
Hive/HCat 1
WebHCat
HBase
ZooKeeper
Oozie 1
Nagios 2
Misc

Choose Services

Assign Masters

Assign Slaves and Clients

Customize Services

Review

Install, Start and Test

Summary

▼ JobTracker

JobTracker host: rhel2

JobTracker new generation size: MB

JobTracker maximum new generation size: MB

JobTracker maximum Java heap size: MB Undo

▼ TaskTracker

TaskTracker hosts: rhel3 and 61 others

MapReduce local directories:

/HDP/sdb1/hadoop/mapred
 /HDP/sdc1/hadoop/mapred
 /HDP/sdd1/hadoop/mapred
 /HDP/sde1/hadoop/mapred
 /HDP/sdf1/hadoop/mapred
 ..:

Number of Map slots per node: Undo

Number of Reduce slots per node: Undo

Java options for MapReduce tasks: MB Undo

Figure 83 Customize Services - MapReduce Configuration Window Part 2

▼ [General](#)

MapReduce Capacity Scheduler	<input type="text" value="org.apache.hadoop.mapred.CapacityTaskS"/>	
Cluster's Map slot size (virtual memory)	<input type="text" value="-1"/>	<input type="button" value="MB"/>
Cluster's Reduce slot size (virtual memory)	<input type="text" value="-1"/>	<input type="button" value="MB"/>
Upper limit on virtual memory for single Map task	<input type="text" value="-1"/>	<input type="button" value="MB"/>
Upper limit on virtual memory for single Reduce task	<input type="text" value="-1"/>	<input type="button" value="MB"/>
Default virtual memory for a job's map-task	<input type="text" value="-1"/>	<input type="button" value="MB"/>
Default virtual memory for a job's reduce-task	<input type="text" value="-1"/>	<input type="button" value="MB"/>
Map-side sort buffer memory	<input style="border: 1px solid #add8e6;" type="text" value="1024"/>	<input type="button" value="MB"/> Undo
Limit on buffer	<input type="text" value="0.9"/>	
Job log retention (hours)	<input type="text" value="24"/>	<input type="button" value="hours"/>
Maximum number tasks for a Job	<input type="text" value="-1"/>	
Enable Job Diagnostics	<input checked="" type="checkbox"/>	

Hive/HCat

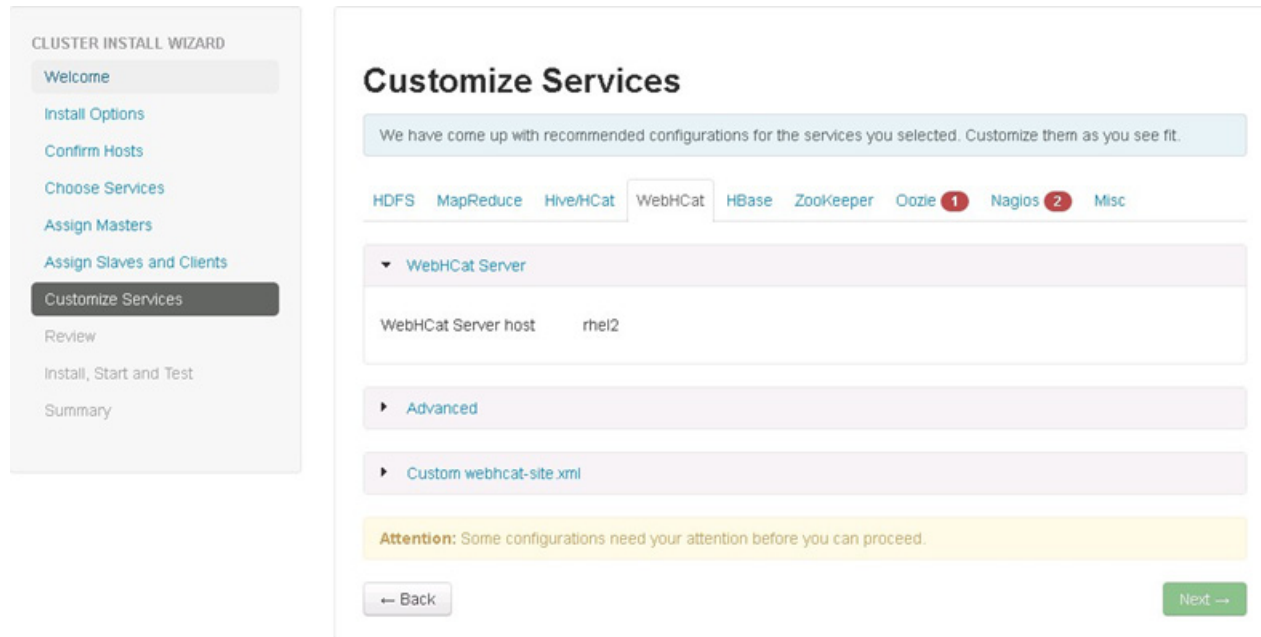
Enter the hive database password as per the organizational policy as shown in [Figure 84](#).

Figure 84 *Customize Services - Hive/HCat Window*

The screenshot displays the 'Customize Services' window for Hive/HCat configuration. On the left is a sidebar titled 'CLUSTER INSTALL WIZARD' with the following steps: Welcome, Install Options, Confirm Hosts, Choose Services, Assign Masters, Assign Slaves and Clients, **Customize Services** (highlighted), Review, Install, Start and Test, and Summary. The main content area is titled 'Customize Services' and includes a message: 'We have come up with recommended configurations for the services you selected. Customize them as you see fit.' Below this is a horizontal menu with tabs for HDFS, MapReduce, Hive/HCat (selected), WebHCat, HBase, ZooKeeper, Oozie (with a red '1' badge), Nagios (with a red '2' badge), and Misc. The 'Hive Metastore' section is expanded, showing the following configuration: Hive Metastore host: rhel2; Database Type: MySQL; Hive Database: New MySQL Database, Existing MySQL Database; Database Host: rhel2; Database Name: hive; Database Username: hive; Database Password: two masked password fields with an 'Undo' button. Below the configuration are sections for 'Advanced' and 'Custom hive-site.xml'. A yellow attention box states: 'Attention: Some configurations need your attention before you can proceed.' At the bottom, there are 'Back' and 'Next' buttons.

WebHCat

We can restore the default settings, no changes needed as shown in [Figure 85](#).

Figure 85 *Customize Services - WebHcat Configuration*

HBase

Update the HBase configurations as shown in [Table 9](#), and [Figure 86](#):

Table 9 *HBase Configurations*

Property Name	Value
HBase Master Maximum Java Heap Size	4GB
HBase RegionServers Maximum Java Heap Size	32GB

Figure 86 *Customize Services - HBase Configuration Window*

CLUSTER INSTALL WIZARD

- Welcome
- Install Options
- Confirm Hosts
- Choose Services
- Assign Masters
- Assign Slaves and Clients
- Customize Services
- Review
- Install, Start and Test
- Summary

Customize Services

We have come up with recommended configurations for the services you selected. Customize them as you see fit.

HDFS MapReduce Hive/HCat WebHCat HBase ZooKeeper Oozie 1 2 Nagios Misc

▼ HBase Master

HBase Master hosts rhel2

HBase Master Maximum Java heap size MB Undo

▼ RegionServer

RegionServer hosts rhel3 and 13 others

HBase RegionServers maximum Java heap size MB Undo

HBase RegionServer Handler

HBase Region Major Compaction ms

HBase Region Block Multiplier

HBase Region Memstore Flush Size bytes

ZooKeeper

We can restore the default settings in the ZooKeeper window, no changes needed as shown in [Figure 85](#).

Figure 87 *Customize Services - ZooKeeper Window*

The screenshot displays the 'Customize Services' window for the ZooKeeper Server. On the left is a 'CLUSTER INSTALL WIZARD' sidebar with steps: Welcome, Install Options, Confirm Hosts, Choose Services, Assign Masters, Assign Slaves and Clients, **Customize Services**, Review, Install, Start and Test, and Summary. The main window title is 'Customize Services' and contains a message: 'We have come up with recommended configurations for the services you selected. Customize them as you see fit.' Below this is a service selection bar with tabs for HDFS, MapReduce, Hive/HCat, WebHCat, HBase, ZooKeeper, Oozie (with a red '1' badge), Nagios (with a red '2' badge), and Misc. The 'ZooKeeper Server' section is expanded, showing the following configuration fields:

- ZooKeeper Server hosts: rhel1 and 2 others
- ZooKeeper directory: /HDP/sdb1/hadoop/zookeeper
- Length of single Tick: 2000 ms
- Ticks to allow for sync at Init: 10
- Ticks to allow for sync at Runtime: 5
- Port for running ZK Server: 2181

Below the configuration fields is an 'Advanced' section (collapsed) and a yellow 'Attention' box stating: 'Attention: Some configurations need your attention before you can proceed.' At the bottom, there are 'Back' and 'Next' buttons.

Oozie

Enter the Oozie database password as per the organizational policy as shown in [Figure 88](#).

Figure 88 *Customize Services - Oozie Window*

The screenshot shows the 'Customize Services' window for Oozie configuration. The sidebar on the left contains the following navigation options: Welcome, Install Options, Confirm Hosts, Choose Services (highlighted), Assign Masters, Assign Slaves and Clients, Customize Services (highlighted), Review, Install, Start and Test, and Summary. The main content area is titled 'Customize Services' and includes a message: 'We have come up with recommended configurations for the services you selected. Customize them as you see fit.' Below this is a navigation bar with tabs for HDFS, MapReduce, Hive/HCat, WebHCat, HBase, ZooKeeper, Oozie (selected), Nagios (with a red '2' notification), and Misc. The Oozie configuration section is expanded to show the following fields:

- Oozie Server host: rhel1
- Database Type: Derby
- Oozie Database: New Derby Database
- Database Name:
- Database Username:
- Database Password:
- Oozie Data Dir:

A tooltip for the Database Username field indicates: 'Database Username: oozie_metastore_user_name. Database user name to use to connect to the database.' Below the configuration fields are two expandable sections: 'Advanced' and 'Custom oozie-site.xml'. A yellow attention banner reads: 'Attention: Some configurations need your attention before you can proceed.' At the bottom, there are 'Back' and 'Next' buttons.

Nagios

Update the Nagios configuration as shown in [Figure 89](#).

Enter the following in the Nagios window:

- Nagios admin password as per organizational policy.
- Hadoop admin email.

Figure 89 *Customize Services - Nagios Window*

CLUSTER INSTALL WIZARD

- Welcome
- Install Options
- Confirm Hosts
- Choose Services
- Assign Masters
- Assign Slaves and Clients
- Customize Services**
- Review
- Install, Start and Test
- Summary

Customize Services

We have come up with recommended configurations for the services you selected. Customize them as you see fit.

HDFS MapReduce HiveHCat WebHCat HBase ZooKeeper Oozie **Nagios** Misc

▼ General

Nagios Admin username

Nagios Admin password Undo

Hadoop Admin email Undo

← Back Next →

Misc

We can restore the default settings in the Misc window, no changes needed as shown in [Figure 90](#).

Figure 90 *Customize Services - Misc Window*

CLUSTER INSTALL WIZARD

- Welcome
- Install Options
- Confirm Hosts
- Choose Services
- Assign Masters
- Assign Slaves and Clients
- Customize Services**
- Review
- Install, Start and Test
- Summary

Customize Services

We have come up with recommended configurations for the services you selected. Customize them as you see fit.

HDFS MapReduce HiveHCat WebHCat HBase ZooKeeper Oozie Nagios **Misc**

▼ General

Ganglia rrd cached base directory

▼ Users and Groups

Proxy group for Hive, WebHCat, and Oozie

HDFS User

MapReduce User

HBase User

Hive User

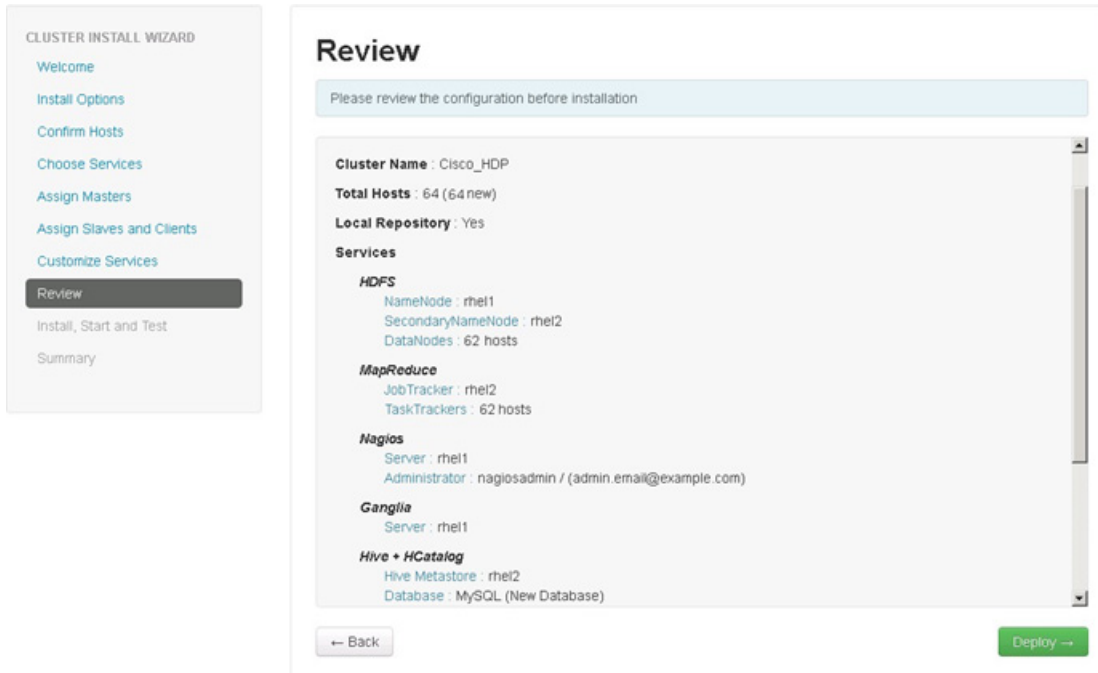
HCat User

WebHCat User

Review

Make sure the Review window shows all the configurations that you have done. Then click **Deploy** as shown in [Figure 91](#). If any changes are to be made, use the left navigation bar to return to the appropriate screen.

Figure 91 Review Window



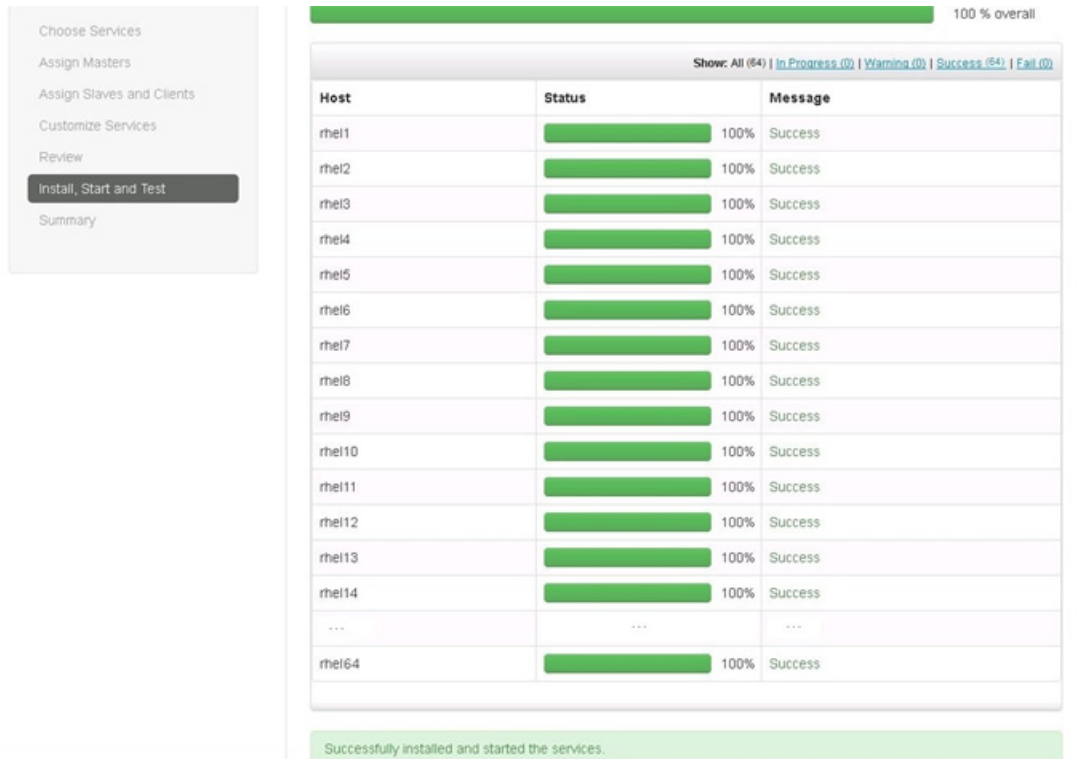
The installation process is shown by the progress indicator as shown in [Figure 92](#). Each component when installed, gets started with a simple test which is run on each of the components. The overall status of the installed components are shown by the progress bar besides every host.

To see the specific information on what tasks have been completed per host, click the link in the Message column for the appropriate host. In the Tasks pop-up, select individual task to see the related log files. Select filter conditions by using the drop-down list. To see a larger version of the log contents, click **Open**. And to copy the contents to the clipboard, click **Copy**.

Depending on the components being installed per host, the entire process may take 30 or more minutes.

Click **Next**, when successfully installed and started the services message at the bottom of window appears as shown in [Figure 92](#).

Figure 92 Cluster Install Wizard - Install, Start and Test Window



100 % overall

Show: All (64) | In Progress (0) | Warning (0) | Success (64) | Fail (0)

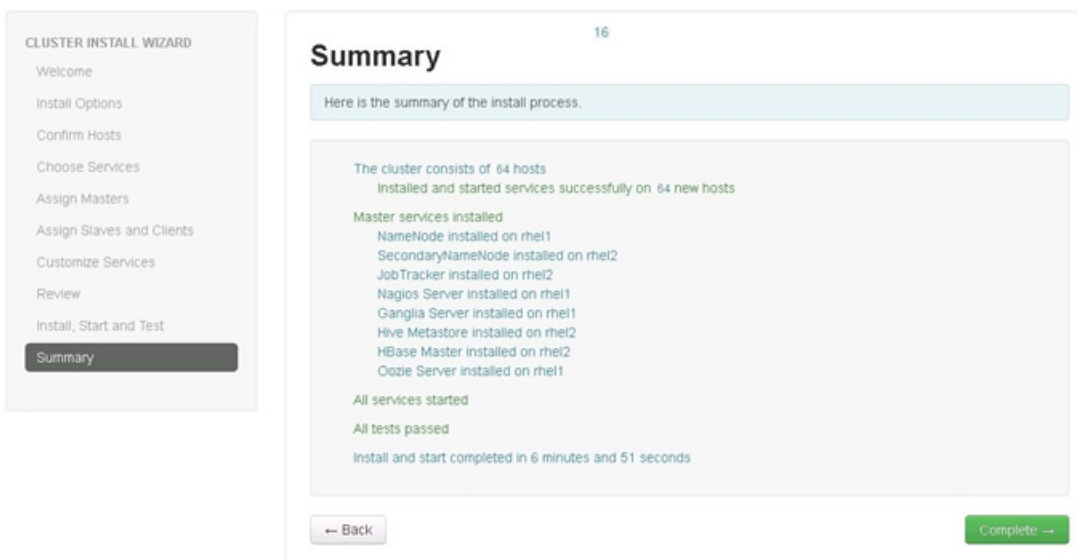
Host	Status	Message
rhel1	100%	Success
rhel2	100%	Success
rhel3	100%	Success
rhel4	100%	Success
rhel5	100%	Success
rhel6	100%	Success
rhel7	100%	Success
rhel8	100%	Success
rhel9	100%	Success
rhel10	100%	Success
rhel11	100%	Success
rhel12	100%	Success
rhel13	100%	Success
rhel14	100%	Success
...
rhel64	100%	Success

Successfully installed and started the services.

Summary of Installation Process

The summary page shows the accomplished tasks after the completion of cluster installation.

Figure 93 Cluster Install Wizard- Summary Window



16

Summary

Here is the summary of the install process.

The cluster consists of 64 hosts
Installed and started services successfully on 64 new hosts

Master services installed

- NameNode installed on rhel1
- SecondaryNameNode installed on rhel2
- JobTracker installed on rhel2
- Nagios Server installed on rhel1
- Ganglia Server installed on rhel1
- Hive Metastore installed on rhel2
- HBase Master installed on rhel2
- Oozie Server installed on rhel1

All services started

All tests passed

Install and start completed in 6 minutes and 51 seconds

← Back Complete →

Conclusion

Hadoop has become a popular data management application across all the verticals. The Cisco CPA for Big Data for HDP offers a dependable deployment model for enterprise Hadoop that offer a fast and predictable path for businesses to unlock the value in big data.

The configuration details provided in this document can be extended to clusters of various sizes depending on application demands. Up to 160 servers (10 racks) can be supported without an additional switching in a single UCS domain. Each additional rack requires two Cisco Nexus 2232PP 10GigE Fabric Extenders and 16 Cisco UCS C240M3 Rack-Mount Servers. Scaling beyond 10 racks (160 servers) can be implemented by interconnecting multiple UCS domains using Nexus 6000/7000 Series switches. The solution is scalable to thousands of servers and to hundreds of petabytes storage, and is managed from a single pane using [Cisco UCS Central](#).

Bill of Material

This section provides the hardware and software components used in the design setup for deploying the 64-node High Performance Cluster.

[Table 10](#) describes the BOM for the master rack; [Table 11](#) describes the BOM for expansion racks (rack 2 to 4); and [Table 12](#) and [Table 13](#) describe the BOM for the software components

Table 10 *Bill of Material for Base Rack*

Part Number	Description	Quantity
UCS-EZ-BD-HC	High Capacity Rack	1
UCS-EZ-INFRA-FI96	Cisco UCS 6296 FI w/ 18p LIC, Cables Bundle	2 (included)
N2K-UCS2232PF	Cisco Nexus 2232PP with 16 FET (2 AC PS, 1 FAN (Std Airflow)	2 (included)
UCS-EZ-C240-2640	Cisco UCS C240M3 LFF w/ 2640, 16x16GB, VIC 1225, 2PS (if High Capacity Config)	16 (included)
UCS-EZ-C240-2665	Cisco UCS C240 M3 SFF w/ 2665, 16x16GB, VIC 1225, 2PS (if High Performance Config)	16 (included)
SFP-H10GB-CU3M=	10GBASE-CU SFP+ Cable 3 Meter	28 (included)
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)	6

Table 11 Bill of Material for Expansion Racks

Part Number	Description	Quantity
N2K-UCS2232PF	Cisco Nexus 2232PP with 16 FET	6
CON-SNTP-UCS2232	SMARTNET 24X7X4 Cisco Nexus 2232PP	6
UCS-EZ-C240-2640	Cisco UCS C240M3 LFF w/ 2640, 16x16GB, VIC 1225, 2PS (if High Capacity Config)	48
UCS-EZ-C240-2665	Cisco UCS C240 M3 SFF w/ 2665, 16x16GB, VIC 1225, 2PS (if High Performance Config)	48
CON-SNTP-C24065EZ	SMARTNET 24X7X4 Cisco UCS C240M3 Server	48
SFP-H10GB-CU1M	10GBASE-CU SFP+ Cable 1 Meter (server 1-8 to FEX in rack 1-4, 8 x FEX 1 to FI A, 8 x FEX 2 to FI B in rack 1)	80
FP-H10GB-CU2M	10GBASE-CU SFP+ Cable 2 Meter (server 9-16 to FEX in rack 1-4)	64
SFP-H10GB-CU3M	10GBASE-CU SFP+ Cable 3 Meter (8 x FEX 1 to FI A, 8 x FEX 2 to FI B in rack 2 and 3)	32
SFP-H10GB-CU5M	10GBASE-CU SFP+ Cable 5 Meter (8 x FEX 1 to FI A, 8 x FEX 3 to FI B in rack 2 and 3)	16
RACK-UCS2	Cisco R42610 standard rack w/side panels	4
RP208-30-1P-U-2=	Cisco RP208-30-U-2 Single Phase PDU 20x C13 4x C19 (Country Specific)	8
CON-UCW3-RPDUX	UC PLUS 24X7X4 Cisco RP208-30-U-X Single Phase PDU 2x (Country Specific)	24

Table 12 RedHat Enterprise Linux License

Red Hat Enterprise Linux		
RHEL-2S-1G-3A	Red Hat Enterprise Linux	64
CON-ISV1-RH2S1G3A	3 year Support for Red Hat Enterprise Linux	64

Table 13 Hortonworks Software

HDP		
HDP	Hortonworks Data Platform	64