




Cisco UCS Integrated Infrastructure for Big Data with Hortonworks Data Platform With Optional Tiered Storage Extension

Last Updated: September 9, 2015



Building Architectures to Solve Business Problems



About the Authors

Raghunath Nambiar, Distinguished Engineer, Data Center Business Group (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.



Manankumar Trivedi

Manankumar Trivedi, Technical Marketing Engineer, Data Center Solutions Group (Cisco Systems)

Manankumar Trivedi is a Technical Marketing Engineer in the Data Center Solutions Group at Cisco Systems. He is part of solution engineering team focusing on big data infrastructure and performance.



Ali Bajwa

Ali Bajwa, Senior Partner Solutions Engineer, Technology Alliances Team, Hortonworks

Ali Bajwa is a Senior Partner Solutions Engineer at Hortonworks and works as part of the Technology Alliances team. His focus is to evangelize and assist partners integrate with Hortonworks Data Platform.

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.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <http://www.cisco.com/go/trademarks>. Third-party trademarks mentioned 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. (1110R).

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

© 2015 Cisco Systems, Inc. All rights reserved.

Acknowledgment

The authors acknowledge contributions of Karthik Kulkarni, Karthik Karupasamy, Ashwin Manjunatha and Sindhu Sudhir in developing this document.



Cisco UCS Integrated Infrastructure for Big Data with Hortonworks Data Platform

Audience

This document describes the architecture and deployment procedures for Hortonworks Data Platform (HDP 2.2) on a 64 Cisco UCS C240 M4 node cluster along with 4 archival node (Cisco UCS C3160) based on Cisco UCS Integrated Infrastructure 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 2.2 on Cisco UCS Integrated Infrastructure 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 Integrated Infrastructure for Big Data with Tiered Storage

The Cisco UCS solution for Hortonworks is based on [Cisco UCS Integrated Infrastructure 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:



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

Copyright © 2015 Cisco Systems, Inc. All rights reserved.

Cisco UCS 6200 Series Fabric Interconnects

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, automating ongoing system 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.

Figure 1 Cisco UCS 6296UP 96-Port Fabric Interconnect



Cisco UCS C-Series Rack Mount Servers

Cisco UCS C-Series Rack Mount C220 M4 High-Density Rack servers (Small Form Factor Disk Drive Model) and Cisco UCS C240 M4 High-Density Rack servers (Small Form Factor Disk Drive Model) are enterprise-class systems that support a wide range of computing, I/O, and storage-capacity demands in compact designs. Cisco UCS C-Series Rack-Mount Servers are based on Intel Xeon E5-2600 v3 product family and 12-Gbps SAS throughput, delivering significant performance and efficiency gains over the previous generation of servers. The servers use dual Intel Xeon processor E5-2600 v3 series CPUs and support up to 768 GB of main memory (128 or 256 GB is typical for big data applications) and a range of disk drive and SSD options. 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-optimized option, along with 4 Gigabit Ethernet LAN-on-motherboard (LOM) ports. Cisco UCS virtual interface cards 1227 (VICs) designed for the M4 generation of Cisco UCS C-Series Rack Servers are optimized for high-bandwidth and low-latency cluster connectivity, with support for up to 256 virtual devices that are configured on demand through Cisco UCS Manager.

Figure 2 Cisco UCS C240 M4 Rack Server



Cisco UCS C3160 Rack Server

Cisco UCS C3160 Rack Server is an advanced, modular rack server with extremely high storage density. Based on the Intel Xeon processor E5-2600 v2 series, it offers up to 360 TB of local storage in a compact 4-rack-unit (4RU) form factor. Because all its hard-disk drives are individually hot-swappable, and with its built-in enterprise-class Redundant Array of Independent Disks (RAID) redundancy, the Cisco UCS C3160 helps you achieve the highest levels of data availability. The Cisco UCS C3160 is ideal for Snapshots, active archiving, compliance, media storage, and distributed file systems for scenarios in which high storage capacity is important. Cisco UCS virtual interface cards 1227 (VICs) designed for the M4 generation of Cisco UCS C-Series Rack Servers and C3160 are optimized for high-bandwidth and low-latency cluster connectivity, with support for up to 256 virtual devices that are configured on demand through Cisco UCS Manager.

Figure 3 Cisco UCS C3160 Server



Cisco UCS Virtual Interface Cards (VICs)

Cisco UCS Virtual Interface Cards (VICs), unique to Cisco, Cisco UCS Virtual Interface Cards incorporate next-generation converged network adapter (CNA) technology from Cisco, and offer dual 10-Gbps ports designed for use with Cisco UCS C-Series Rack-Mount Servers. Optimized for virtualized networking, these cards deliver high performance and bandwidth utilization and support up to 256 virtual devices. The Cisco UCS Virtual Interface Card (VIC) 1227 is a dual-port, Enhanced Small Form-Factor Pluggable (SFP+), 10 GigabitEthernet Ethernet and Fiber Channel over Ethernet (FCoE)-capable, PCI Express (PCIe) modular LAN on motherboard (mLOM) adapter. It is designed exclusively for the M4 generation of Cisco UCS C-Series Rack Servers and the C3160 dense storage servers.

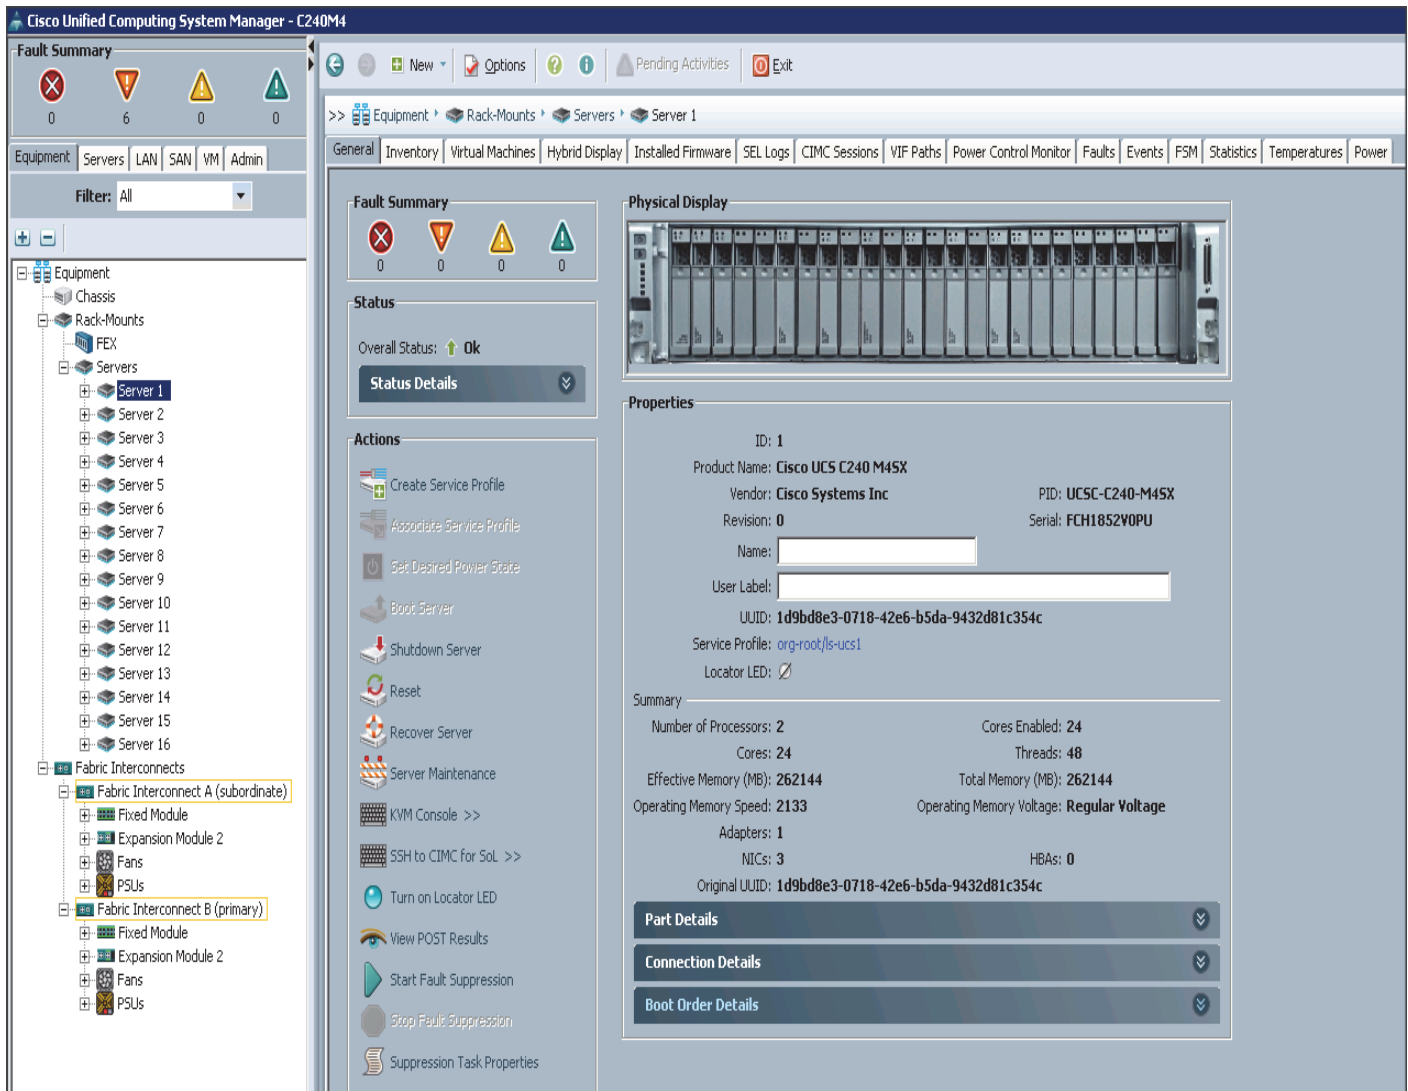
Figure 4 Cisco UCS VIC 1227



Cisco UCS Manager

Cisco UCS Manager resides within the Cisco UCS 6200 Series Fabric Interconnects. It makes the system self-aware and self-integrating, managing all of the system components as a single logical entity. Cisco UCS Manager can be accessed through an intuitive graphical user interface (GUI), a command-line interface (CLI), or an XML application-programming interface (API). Cisco UCS Manager uses service profiles to define the personality, configuration, and connectivity of all resources within Cisco UCS, radically simplifying provisioning of resources so that the process takes minutes instead of days. This simplification allows IT departments to shift their focus from constant maintenance to strategic business initiatives.

Figure 5 Cisco UCS Manager



Cisco UCS Director Express for Big Data

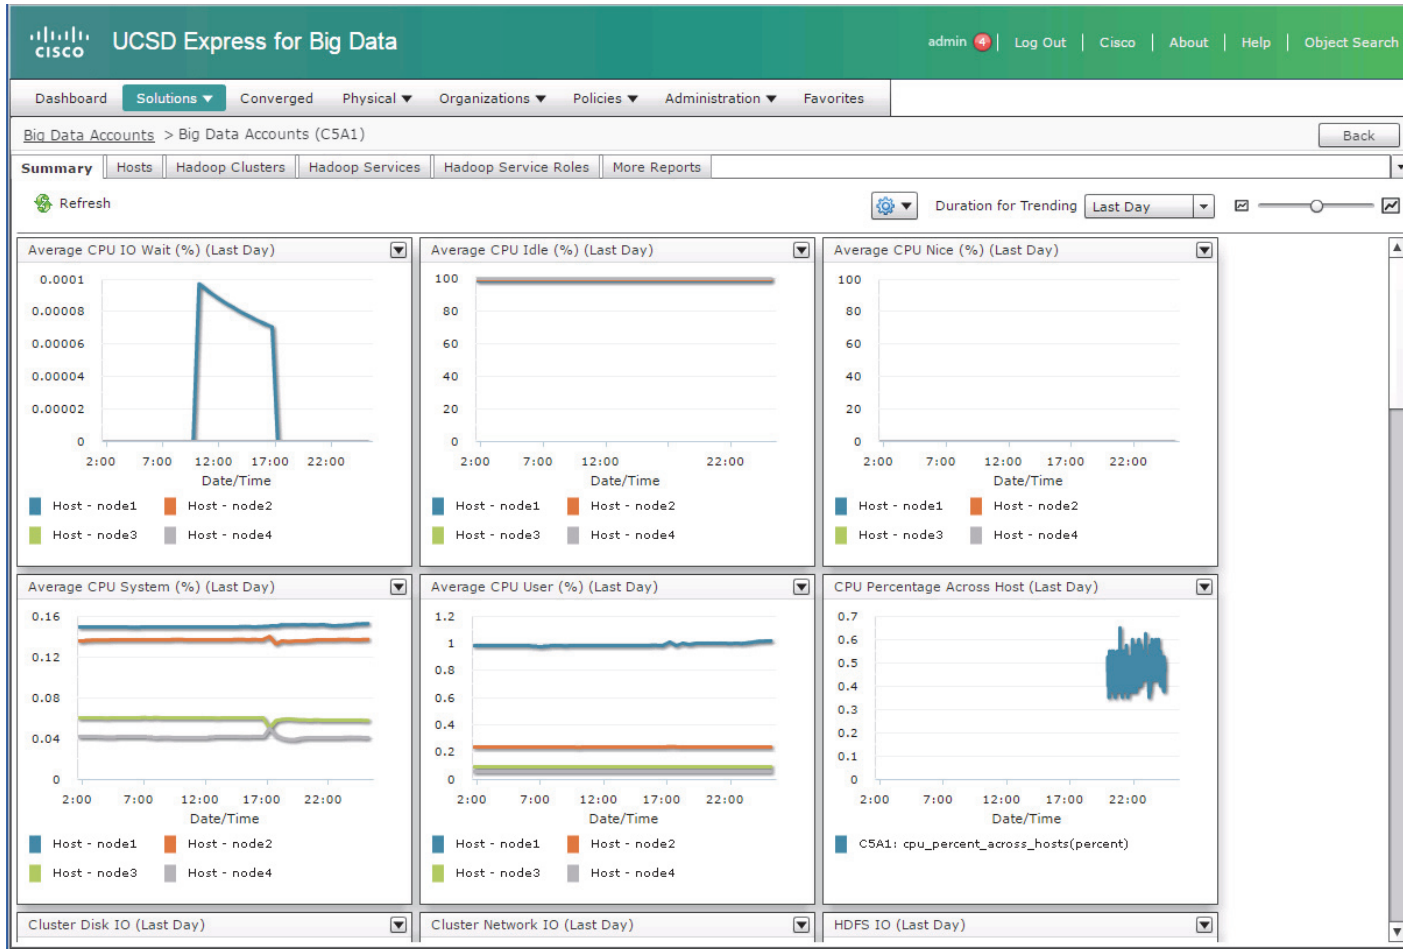
Cisco UCS Director Express for Big Data provides a single-touch solution that automates deployment of Hadoop Distributions on leading Cisco UCS Integrated Infrastructure for Big Data.

It also provides a single management pane across both physical infrastructure and Hadoop software. All elements of the infrastructure are handled automatically with little need for user input. Through this approach, configuration of physical computing, internal storage, and networking infrastructure is integrated with the deployment of operating systems, Java packages, and Hadoop along with the provisioning of Hadoop services. Cisco UCS Director Express for Big Data is integrated with major Hadoop distributions from Hortonworks, Cloudera, and MapR, providing single-pane management across the entire infrastructure.

It complements and communicates with Hadoop managers, providing a system wide perspective and enabling administrators to correlate Hadoop activity with network and computing activity on individual Hadoop nodes.

The appendix section describes on how to go about configuring Cisco UCS Director Express for Big Data and deploying popular Hadoop distributions such as Cloudera, MapR and Hortonworks on the Cisco UCS Integrated Infrastructure for Big Data cluster.

Figure 6 Cisco USCD Express for Big Data



Hortonworks Data Platform (HDP 2.2)

The Hortonworks Data Platform 2.2 (HDP 2.2) 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.

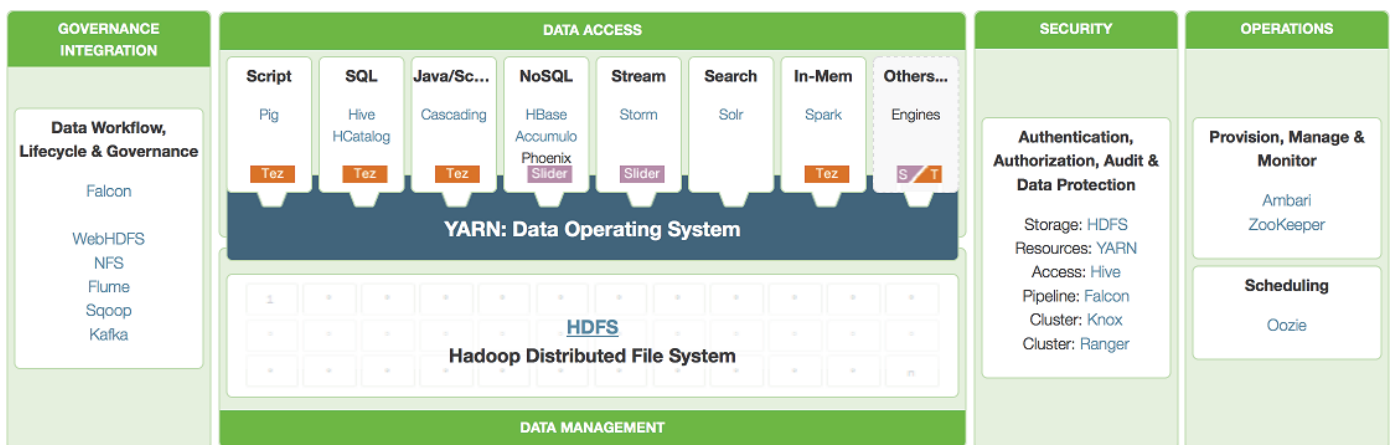
The Hortonworks Data Platform combines the most useful and stable versions of Apache Hadoop and its related projects into a single tested and certified package. Hortonworks offers the latest innovations from the open source community, along with the testing and quality you expect from enterprise-quality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

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

Key Features of HDP 2.2

Hortonworks Data Platform enables Enterprise Hadoop: the full suite of essential Hadoop capabilities that are required by the enterprise and that serve as the functional definition of any data platform technology. This comprehensive set of capabilities is aligned to the following functional areas: Data Management, Data Access, Data Governance and Integration, Security, and Operations.

Figure 7 Hortonworks Data Platform



HDP 2.2 incorporates many new innovations that have happened in Hadoop and its supporting ecosystem of projects. Some of the key projects are listed below.

Tiered Storage in HDFS

With HDP 2.2, HDFS provides the ability to utilize heterogeneous storage media within the HDFS cluster to enable the following tiered storage scenarios:

- **Hot Data Tier:** Provides a storage tier that consists of C240M4 servers to store datasets that require high speed storage access.
- **Archival Data Tier:** Provides storage dense tier that consists of C3160 server to store less frequently accessed datasets.

This is explained in detail in post HDP installation section.

Enterprise SQL at Scale in Hadoop

While YARN has allowed new engines to emerge for Hadoop, one of the popular integration point with Hadoop continues to be SQL and Apache Hive is still the defacto standard.

New capabilities in HDP 2.2 include:

- **Updated SQL Semantics for Hive Transactions for Update and Delete:** ACID transactions provide atomicity, consistency, isolation, and durability. This helps with streaming and baseline update scenarios for Hive such as modifying dimension tables or other fact tables.

- **Improved Performance of Hive with a Cost Based Optimizer:** The cost based optimizer for Hive, uses statistics to generate several execution plans and then chooses the most efficient path as it relates system resources required to complete the operation. This presents a major performance increase for Hive.

Apache Tez

Apache Tez is an extensible framework for building high performance batch and interactive data processing applications, coordinated by YARN in Apache Hadoop. Tez improves the MapReduce paradigm by dramatically improving its speed, while maintaining MapReduce's ability to scale to petabytes of data. Important Hadoop ecosystem projects like Apache Hive and Apache Pig use Apache Tez, as do a growing number of third party data access applications developed for the broader Hadoop ecosystem.

Hive with Tez

As the defacto standard for SQL-In-Hadoop, Apache Hive is optimal for both batch and interactive queries at petabyte scale. Hive embeds Tez so that it can translate complex SQL statements into highly optimized, purpose-built data processing graphs that strike the right balance between performance, throughput, and scalability. Apache Tez innovations drove many of the Hive performance improvements delivered by the Stinger Initiative, a broad community effort that included contributions from 145 engineers across 44 different organizations. Tez helps make Hive interactive.

Kafka for Processing the Internet of Things

[Apache Kafka](#) has quickly become the standard for high-scale, fault-tolerant, publish-subscribe messaging system for Hadoop. It is often used with Storm and Spark so as to stream events in to Hadoop in real time and its application within the “Internet of things” uses cases is tremendous.

Apache Flume

Flume is a distributed, reliable, and available service for efficiently collecting, aggregating, and moving large amounts of streaming data into the Hadoop Distributed File System (HDFS). It has a simple and flexible architecture based on streaming data flows, and is robust and fault tolerant with tunable reliability mechanisms for failover and recovery.

Apache Sqoop

Sqoop is a tool designed for efficiently transferring bulk data between Apache Hadoop and structured data stores such as relational databases. Sqoop imports data from external structured data stores into HDFS or related systems like Hive and HBase. Sqoop can also be used to extract data from Hadoop and export it to external structured data stores such as relational databases and enterprise data warehouses. Sqoop works with relational databases such as Teradata, Netezza, Oracle, MySQL, Postgres, and HSQLDB.

Apache Knox

Knox provides perimeter security so that the enterprise can confidently extend Hadoop access to more of those new users while also maintaining compliance with enterprise security policies. Knox also simplifies Hadoop security for users who access the cluster data and execute jobs. It integrates with prevalent identity management and SSO systems and allows identities from those enterprise systems to be used for seamless, secure access to Hadoop clusters.

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

This CVD describes architecture and deployment procedures for Hortonworks Data Platform (HDP 2.2) on a 64 Cisco UCS C240 M4 node cluster along with 4 archival node (Cisco UCS C3160) based on Cisco UCS Integrated Infrastructure for Big Data. This solution describes in detail the configuration of HDP 2.2 on Cisco UCS Integrated Infrastructure along with Archival nodes (UCS C3160) and defining storage policies for data placement.

The current version of the Cisco UCS Integrated Infrastructure for Big Data offers the following configuration depending on the compute and storage requirements:

Table 1 Cisco UCS Integrated Infrastructure for Big Data Configuration Details

Performance Optimized	Capacity Optimized	Extreme Capacity
16 Cisco UCS C240 M4 Rack Servers (SFF), each with: <ul style="list-style-type: none"> • 2 Intel Xeon processors E5-2680 v3 CPUs • 256 GB of memory • Cisco 12-Gbps SAS Modular Raid Controller with 2-GB flash-based write cache (FBWC) • 24 1.2-TB 10K SFF SAS drives (460 TB total) • 2 120-GB 6-Gbps 2.5-inch Enterprise Value SATA SSDs for Boot • Cisco UCS VIC 1227 (with 2 10 GE SFP+ ports) 	16 Cisco UCS C240 M4 Rack Servers (LFF), each with: <ul style="list-style-type: none"> • 2 Intel Xeon processors E5-2620 v3 CPU • 128 GB of memory • Cisco 12-Gbps SAS Modular Raid Controller with 2-GB FBWC • 12 4-TB 7.2K LFF SAS drives (768 TB total) • 2 120-GB 6-Gbps 2.5-inch Enterprise Value SATA SSDs for Boot • Cisco UCS VIC 1227 (with 2 10 GE SFP+ ports) 	2 Cisco UCS C3160 Rack Servers, each with: <ul style="list-style-type: none"> • 2 Intel Xeon processors E5-2695 v2 CPUs • 256 GB of memory • Cisco 12-Gbps SAS Modular Raid Controller with 4-GB FBWC • 60 4 TB (or 6TB) 7.2K LFF SAS drives (480 TB or 720 TB total) • 2 120-GB 6-Gbps 2.5-inch Enterprise Value SATA SSDs for Boot • 2 Cisco UCS VIC 1227 (each with 2 10 GE SFP+ ports) • 2 built-in 10 GE LOM ports



Note

This CVD describes the install process of HDP 2.2 for a 64 node (2 Master node + 62 Data node) of Performance Optimized Cluster configuration along with 4 Archival Nodes using Cisco UCS C3160 Servers.

The Performance cluster configuration consists of the following:

- Two Cisco UCS 6296UP Fabric Interconnects
- 64 UCS C240 M4 Rack-Mount servers (16 per rack)
- 4 UCS C3160 Rack Server (1 per rack)
- Four Cisco R42610 standard racks
- Eight vertical power distribution units (PDUs) (Country Specific)

Rack and PDU Configuration

Each rack consists of two vertical PDUs. The master rack consists of two Cisco UCS 6296UP Fabric Interconnects, sixteen Cisco UCS C240 M4 Servers and one Cisco UCS C3160 connected to each of the vertical PDUs for redundancy; thereby, ensuring availability during power source failure. The expansion racks consists of sixteen Cisco UCS C240 M4 Servers and one Cisco UCS C3160 connected to each of the vertical PDUs for redundancy; thereby, ensuring availability during power source failure, similar to the master rack.



Note

Please contact your Cisco representative for country specific information.

Table 2 and Table 3 describe the rack configurations of rack 1 (master rack) and racks 2-4 (expansion racks).

Table 2 *Rack 1 (Master Rack)*

Cisco 42URack	Master Rack
42	Cisco UCS FI 6296UP
41	
40	Cisco UCS FI 6296UP
39	
38	Unused
37	Unused
36	Cisco UCS C240 M4
35	
34	Cisco UCS C240 M4
33	
32	Cisco UCS C240 M4
31	
30	Cisco UCS C240 M4
29	
28	Cisco UCS C240 M4
27	

Table 2 **Rack 1 (Master Rack)**

Cisco 42URack	Master Rack
26	Cisco UCS C240 M4
25	
24	Cisco UCS C240 M4
23	
22	Cisco UCS C240 M4
21	
20	Cisco UCS C240 M4
19	
18	Cisco UCS C240 M4
17	
16	Cisco UCS C240 M4
15	
14	Cisco UCS C240 M4
13	
12	Cisco UCS C240 M4
11	
10	Cisco UCS C240 M4
9	
8	Cisco UCS C240 M4
7	
6	Cisco UCS C240 M4
5	
4	Cisco UCS C3160
3	
2	
1	

Table 3 **Rack 2-4 (Expansion Racks)**

Cisco 42URack	Expansion Rack
42	Unused
41	Unused
40	Unused
39	Unused
38	Unused
37	Unused

Table 3 Rack 2-4 (Expansion Racks)

Cisco 42URack	Expansion Rack
36	Cisco UCS C240 M4
35	
34	Cisco UCS C240 M4
33	
32	Cisco UCS C240 M4
31	
30	Cisco UCS C240 M4
29	
28	Cisco UCS C240 M4
27	
26	Cisco UCS C240 M4
25	
24	Cisco UCS C240 M4
23	
22	Cisco UCS C240 M4
21	
20	Cisco UCS C240 M4
19	
18	Cisco UCS C240 M4
17	
16	Cisco UCS C240 M4
15	
14	Cisco UCS C240 M4
13	
12	Cisco UCS C240 M4
11	
10	Cisco UCS C240 M4
9	
8	Cisco UCS C240 M4
7	
6	Cisco UCS C240 M4
5	
4	Cisco UCS C3160
3	
2	
1	

Port Configuration on Fabric Interconnects

Table 4 Port Types and Port Numbers

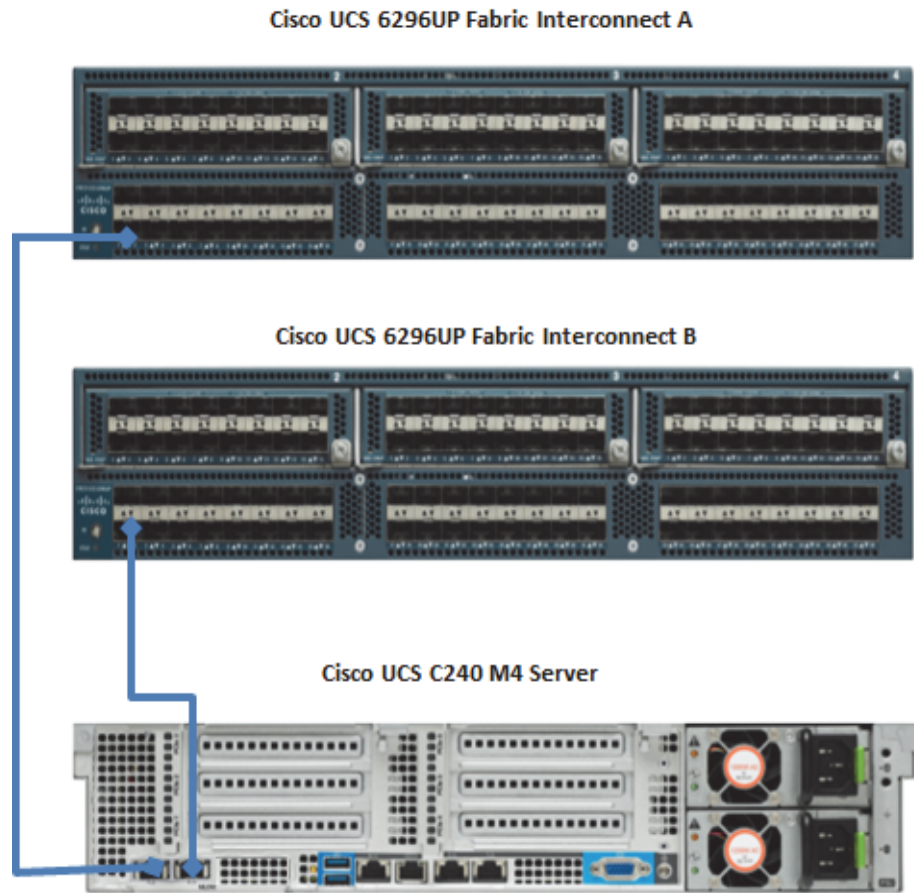
Port Type	Port Number
Network	1
Appliance	2 to 5
Server	6 to 69

Server Configuration and Cabling for C240M4

The C240 M4 rack server is equipped with Intel Xeon E5-2680 v3 processors, 256 GB of memory, Cisco UCS Virtual Interface Card 1227, Cisco 12-Gbps SAS Modular Raid Controller with 2-GB FBWC, 24 1.2-TB 10K SFF SAS drives, 2 120-GB SATA SSD for Boot.

Figure 8, illustrates the port connectivity between the Fabric Interconnect and Cisco UCS C240 M4 server. Sixteen Cisco UCS C240 M4 servers are used in Master rack configurations.

Figure 8 Fabric Topology for C240 M4

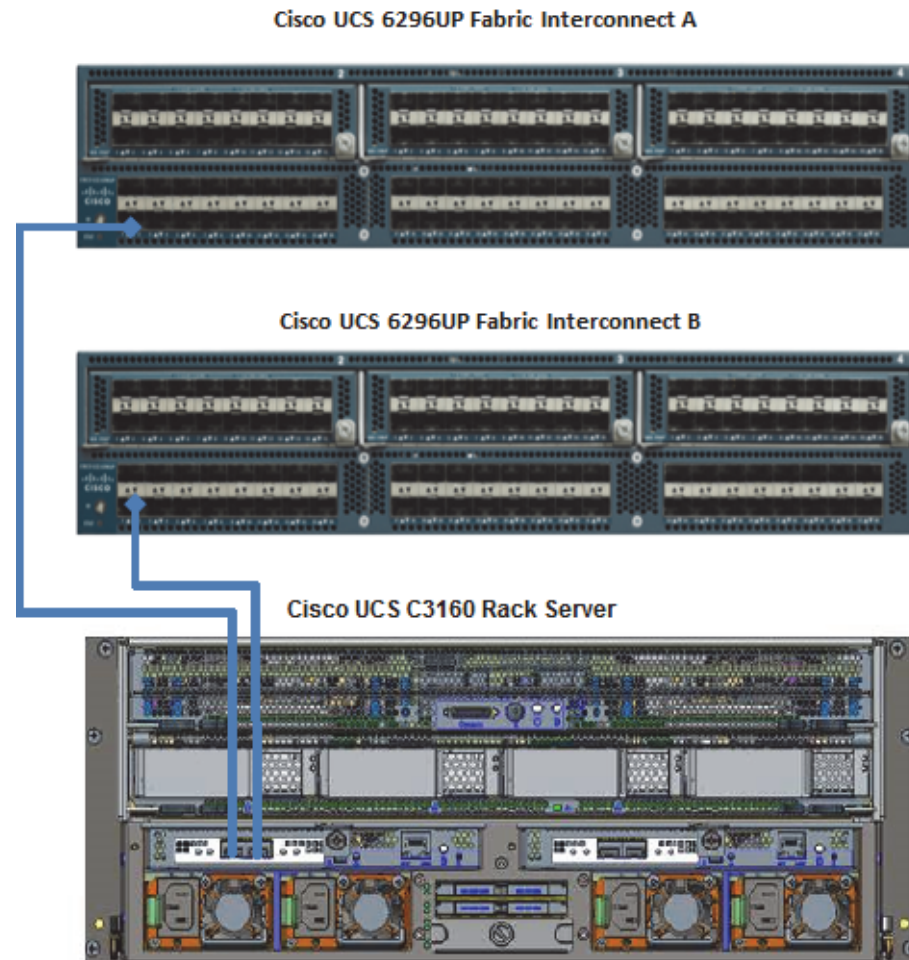


Server Configuration and Cabling for C3160

The C3160 rack server is equipped with Intel Xeon E5-2695 v2 processors, 256 GB of memory, 2 Cisco UCS Virtual Interface Card 1227, Cisco 12-Gbps SAS Modular Raid Controller with 4-GB FBWC, 60 4-TB 7.2K LFF SAS drives, 2 120-GB SATA SSD for Boot.

Figure 9, illustrates the port connectivity between the Fabric Interconnect and Cisco UCS C3160 server as an Appliance port. One Cisco UCS C3160 server is used in master rack configurations.

Figure 9 *Fabric Topology for C3160*




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 10 depicts a 64-node cluster along with 4 archival nodes. Every rack has 16 Cisco UCS C240 M4 servers along with 1 Cisco UCS C3160 as an archival server. Each link in the figure represents 16 x 10 Gigabit Ethernet link from each of the 16 servers connecting to a Cisco UCS Fabric Interconnect as a Direct Connect along with the Cisco UCS C3160 connected as an Appliance port to the Fabric Interconnect. Every server is connected to both the Cisco UCS Fabric Interconnects shown with dual link.

Figure 10 68 Nodes Cluster Configuration



Software Distributions and Versions

The software distributions required versions are listed below.

Hortonworks Data Platform (HDP 2.2)

The Hortonworks Data Platform supported is HDP 2.0. For more information visit <http://www.hortonworks.com>

Red Hat Enterprise Linux (RHEL)

The operating system supported is Red Hat Enterprise Linux 6.5. For more information visit <http://www.redhat.com>

Software Versions

The software versions tested and validated in this document are shown in table 5.

Table 5 *Software Versions*

Layer	Component	Version or Release
Compute	Cisco UCS C240-M4	C240M4.2.0.3d
	Cisco UCS C3160	C3160M3.2.0.2.*
Network	Cisco UCS 6296UP	UCS 2.2(3d)A
	Cisco UCS VIC1227 Firmware	4.0(1d)
	Cisco UCS VIC1227 Driver	2.1.1.66
Storage	LSI SAS 3108	24.5.0-0020
Software	Red Hat Enterprise Linux Server	6.5 (x86_64)
	Cisco UCS Manager	2.2(3d)
	HDP	2.2

**Note**

- The latest drivers can be downloaded from the link below:
<https://software.cisco.com/download/release.html?mdfid=283862063&flowid=25886&softwareid=283853158&release=1.5.7d&reind=AVAILABLE&rellifecycle=&reltype=latest>
- The latest supported RAID controller driver is already included with the RHEL 6.5 operating system.

Fabric Configuration

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

1. Initial setup of the Fabric Interconnect A and B.
2. Connect to UCS Manager using virtual IP address or using the web browser.
3. Launch UCS Manager.
4. Enable server, uplink and appliance ports.
5. Start discovery process.
6. Create pools and policies for Service profile template.
7. Create Service Profile template and 64 Service profiles.
8. Associate Service Profiles to servers.

Performing Initial Setup of Cisco UCS 6296 Fabric Interconnects

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

Configure 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, answer `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, answer `y`.
16. Enter the DNS IPv4 address.
17. Answer `y` to set up the default domain name.
18. Enter the default domain name.
19. Review the settings that were printed to the console, and if they are correct, answer `yes` to save the configuration.
20. Wait for the login prompt to make sure the configuration has been saved.

Configure 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. Answer `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 to accept security certificates, accept as necessary.
4. When prompted, enter **admin** for the user-name and enter the administrative password.

5. Click **Login** to log in to the Cisco UCS Manager.

Upgrading Cisco UCS Manager Software to Version 2.2(3d)

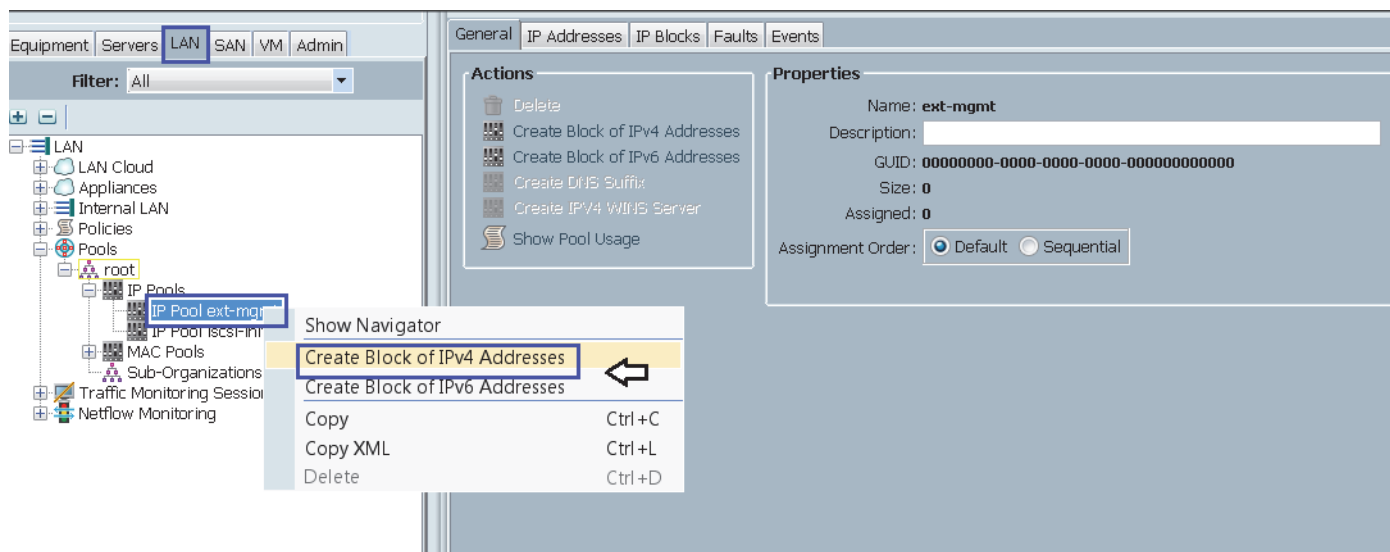
This document assumes the use of UCS 2.2(3d). Refer to [Upgrading between Cisco UCS 2.0 Releases](#) to upgrade the Cisco UCS Manager software and UCS 6296 Fabric Interconnect software to version 2.2(3d). Also, make sure the UCS C-Series version 2.2(3d) software bundles is installed on the Fabric Interconnects.

Adding Block of IP Addresses for KVM Access

These steps provide details for creating a block of KVM IP addresses for server access in the Cisco UCS environment.

1. Select the **LAN** tab at the top of the left window.
2. Select **Pools > IP Pools > IP Pool ext-mgmt**.
3. Right-click **IP Pool ext-mgmt**
4. Select **Create Block of IPv4 Addresses**.

Figure 11 Adding Block of IPv4 Addresses for KVM Access Part 1



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

Figure 12 Adding Block of IPv4 Addresses for KVM Access Part 2

The screenshot shows a dialog box titled "Create Block of IPv4 Addresses". The fields are as follows:

- 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

The "OK" button is highlighted in red.

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

Figure 13 Adding Block of IPv4 Addresses for KVM Access Part 3

The screenshot shows a dialog box titled "Create Block of IPv4 Addresses". The fields are as follows:

- From: 10.29.160.30
- 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

The "OK" button is highlighted in blue.

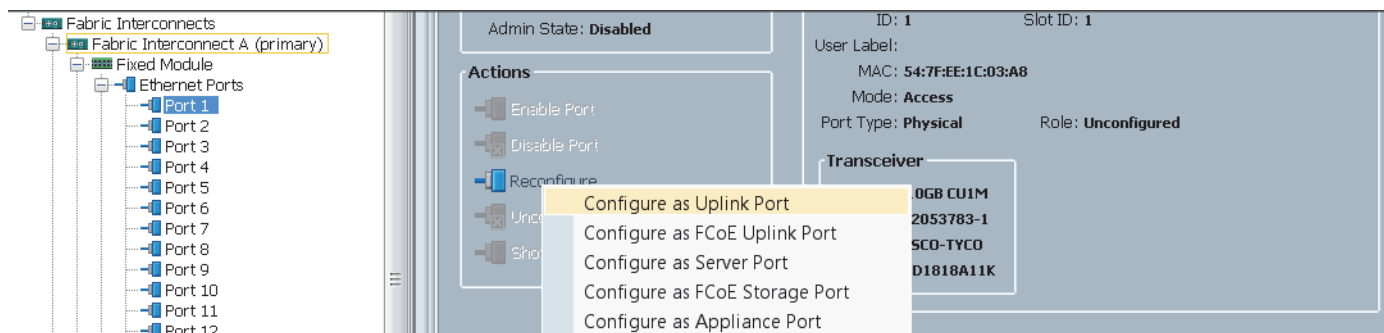
Enabling Uplink Port

These steps provide details for enabling uplinks ports.

1. Select the Equipment tab on the top left of the window.
2. Select **Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module**.
3. Expand the Unconfigured Ethernet Ports section.
4. Select **port 1**, that is connected to the uplink switch, right-click, then select **Reconfigure > Configure as Uplink Port**.
5. Select **Show Interface** and select 10GB for Uplink Connection.
6. A pop-up window appears to confirm your selection. Click **Yes**, then click **OK** to continue.
7. Select **Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module**.
8. Expand the Unconfigured Ethernet Ports section.

9. Select **port 1**, that is connected to the uplink switch, right-click, then select **Reconfigure > Configure as Uplink Port**.
10. Select **Show Interface** and select 10GB for Uplink Connection.
11. A pop-up window appears to confirm your selection. Click **Yes**, then click **OK** to continue.

Figure 14 Enabling Uplink Ports



Configuring VLANs

VLANs are configured as in shown in table 6.

Table 6 VLAN Configurations

VLAN	Fabric	NIC Port	Function	Failover
default(VLAN1)	A	eth0	Management, User connectivity	Fabric Failover to B
vlan11_DATA1	B	eth1	Hadoop	Fabric Failover to A
vlan12_DATA2	A	eth2	Hadoop with multiple NICs support	Fabric Failover to B

All of the VLANs created need to be trunked to the upstream distribution switch connecting the fabric interconnects. For this deployment default VLAN1 is configured for management access (Installing and configuring OS, clustershell commands, setup NTP, user connectivity, etc) and vlan11_DATA1 is configured for Hadoop Data traffic.

With some Hadoop distributions supporting multiple NICs, where Hadoop uses multiple IP subnets for its data traffic, vlan12_DATA2 can be configured to carry Hadoop Data traffic allowing use of both the Fabrics (10 GigE on each Fabric allowing 20Gbps active-active connectivity).

Further, if there are other distributed applications co-existing in the same Hadoop cluster, then these applications could use vlan12_DATA2 providing full 10GigE connectivity to this application on a different fabric without affecting Hadoop Data traffic (here Hadoop is not enabled for multi-NIC).



Note

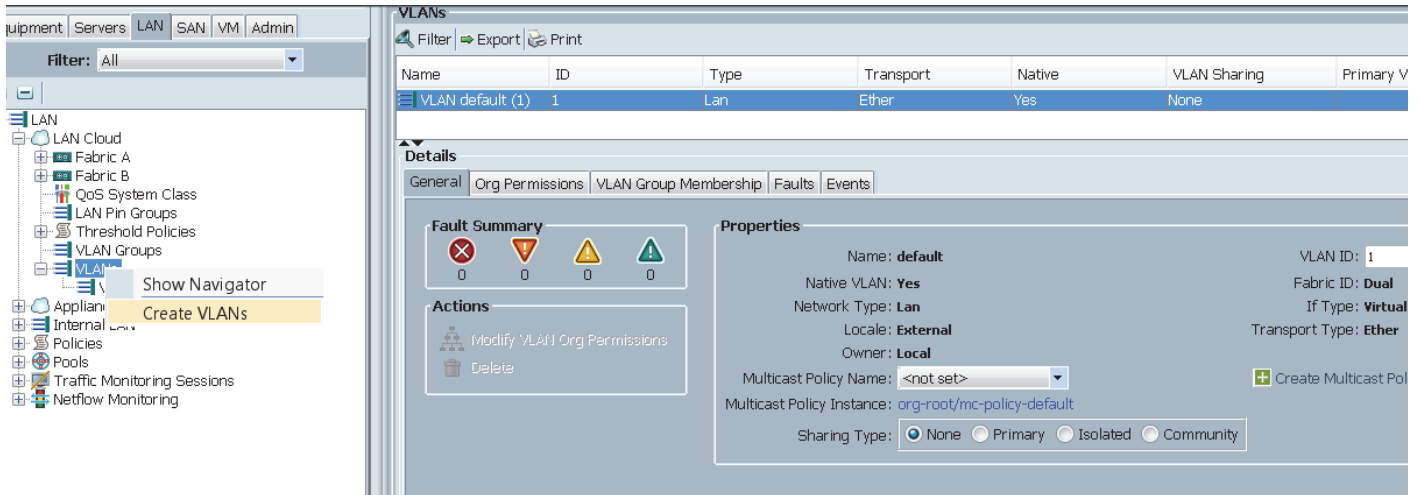
- All applications talking to Hadoop should be able to reach Hadoop VLAN. That is, all applications should be able to access all the Hadoop nodes.
- We are using default VLAN1 for management traffic.

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

1. Select the **LAN** tab in the left pane in the UCS Manager GUI.


2. Select **LAN > VLANs**.
3. Right-click the **VLANs** under the root organization.
4. Select **Create VLANs** to create the VLAN.

Figure 15 *Creating VLAN*




5. Enter **vlan11_DATA1** for the VLAN Name.
6. Click the **Common/Global** radio button for the **vlan11_DATA1**.
7. Enter **11** on VLAN IDs of the Create VLAN IDs.
8. Click **OK** and then, click **Finish**.
9. Click **OK** in the success message box.

Figure 16 *Creating VLAN for Data*



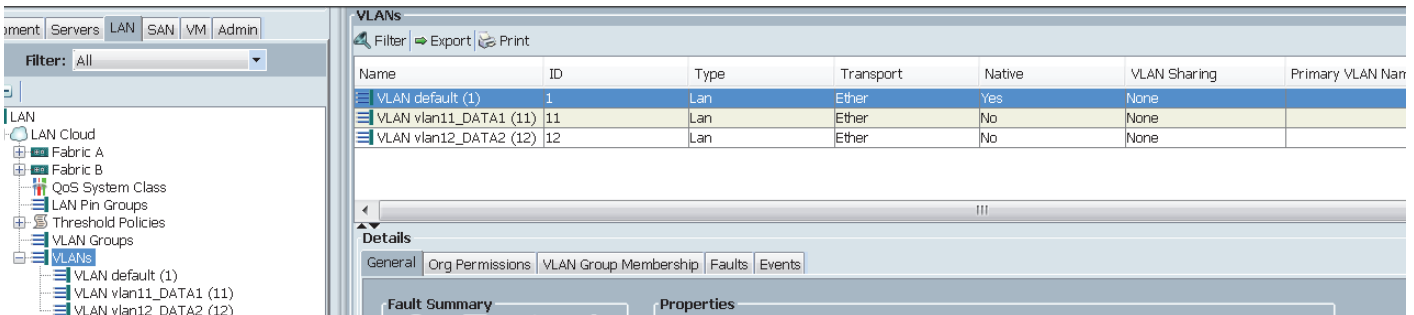
10. Select the **LAN** tab in the left pane again
11. Select **LAN > VLANs**.
12. Right-click the **VLANs** under the root organization.
13. Select **Create VLANs** to create the VLAN.
14. Enter `vlan12_DATA2` for the VLAN Name.
15. Click the **Common/Global** radio button for the `vlan12_DATA2`.
16. Enter 12 on VLAN IDs of the Create VLAN IDs.
17. Click **OK** and then, click **Finish**.

Figure 17 Creating VLAN for Hadoop Data



18. The below screenshot shows the created VLANs.

Figure 18 List of VLANs created for Hadoop Data



Create VLAN for Appliance Port

These steps provide details for creating VLAN for Appliance port configuration.

1. Select the **LAN** tab in the left pane in the UCS Manager GUI.
2. Select **LAN > Appliances > VLANs**.
3. Right-click **VLANs** under the root organization.
4. Select **Create VLANs** to create the VLAN.
5. Enter `vlan11_Appliance` for the VLAN Name.
6. Click the **Common/Global** radio button.
7. Enter 11 for VLAN ID.

Figure 19 Creating VLAN for Appliance Port 1

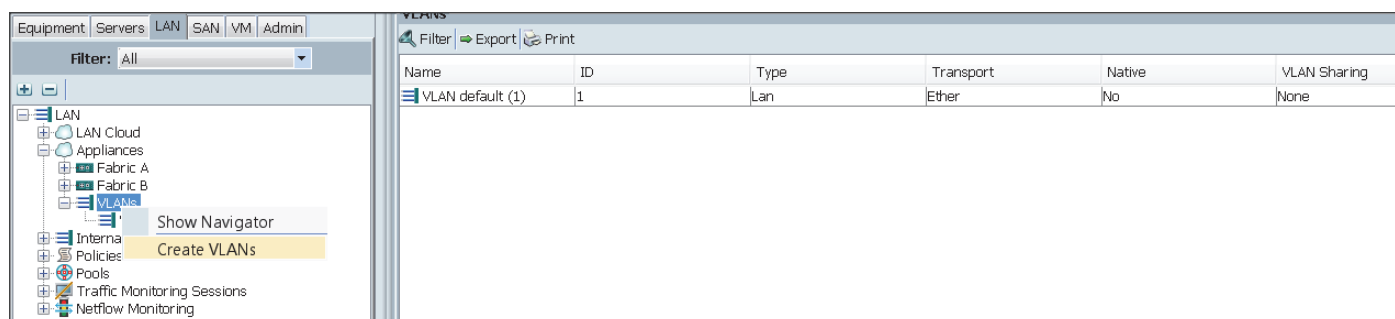



Figure 20 Creating VLAN for Appliance Port 2

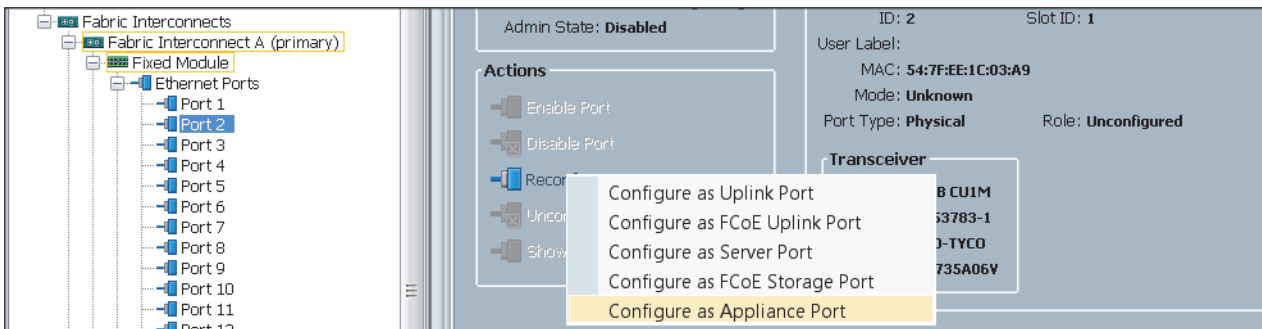


Configuring Appliance Port

These steps provide details for configuring Appliance ports.

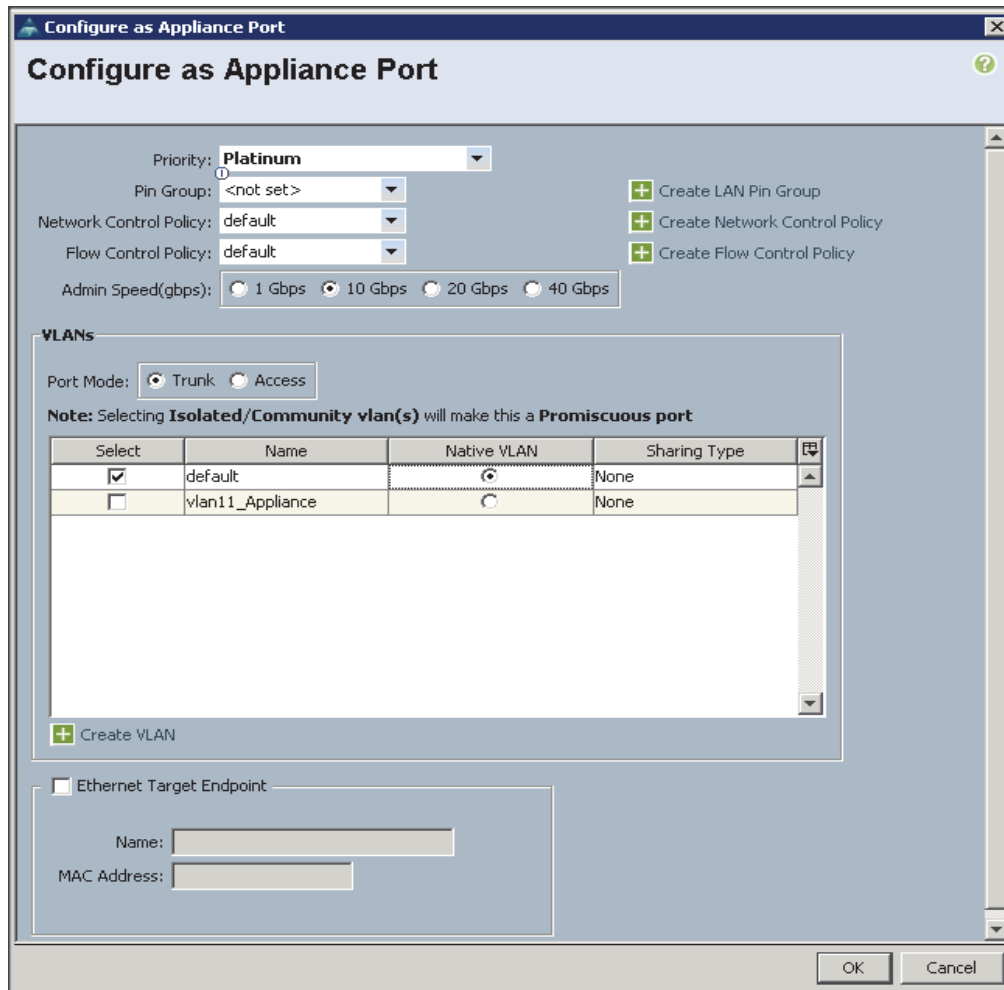
1. Select the **Equipment** tab on the top left of the window.
2. Select **Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module**.
3. Expand the **Unconfigured Ethernet Ports** section.
4. Select port 2, right-click the port, and select **Reconfigure > Configure as an Appliance Port**.

Figure 21 *Configure as Appliance Port 1*



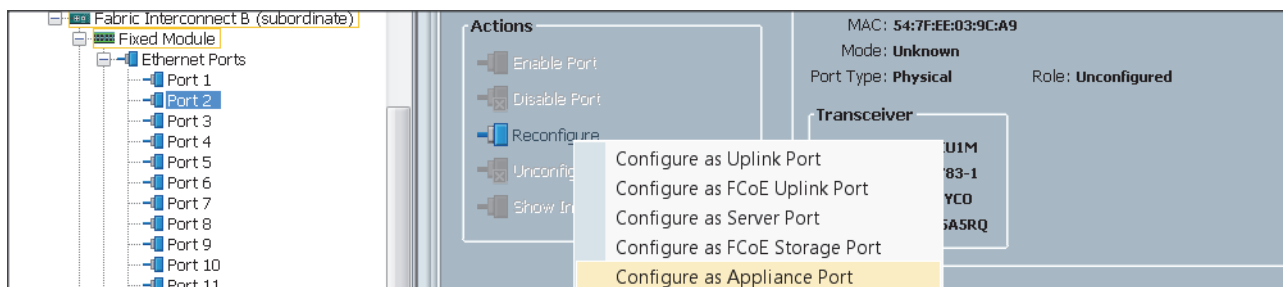
5. A confirmation message box appears. Click **Yes**, then **OK** to continue.
6. Select **Platinum** for the Priority.
7. Keep the Pin Group as <not set>.
8. Keep the Network Control Policy as Default.
9. Keep the Flow Control Policy as Default.
10. Click the **10Gbps** radio button for the Admin Speed.
11. Click the **Trunk** radio button for the Port Mode.
12. Select **Default VLAN**, and click the **Native VLAN** radio button.

Figure 22 Configure as Appliance Port 2



13. Select the **Equipment** tab on the top left of the window.
14. Select **Equipment > Fabric Interconnects > Fabric Interconnect B (Subordinate) > Fixed Module**.
15. Expand the Unconfigured Ethernet Ports section.
16. Select port 2, right-click the port, and select **Reconfigure > Configure as an Appliance Port**.

Figure 23 Configure as Appliance Port 3



17. A confirmation message box appears. Click **Yes**, then **OK** to continue.
18. Select **Platinum** for the Priority.
19. Keep the Pin Group as <not set>.
20. Keep the Network Control Policy as Default.
21. Keep the Flow Control Policy as Default.
22. Click the **10Gbps** radio button for the Admin Speed.
23. Click the **Trunk** radio button for the Port Mode.
24. Select **vlan11_Appliance**, and click the **Native VLAN** radio button.

Figure 24 Configure as Appliance Port 4

Configure as Appliance Port

Priority: **Platinum**

Pin Group: <not set>

Network Control Policy: default

Flow Control Policy: default

Admin Speed(gbps): 1 Gbps 10 Gbps 20 Gbps 40 Gbps

VLANs

Port Mode: Trunk Access

Note: Selecting **Isolated/Community vlan(s)** will make this a **Promiscuous port**

Select	Name	Native VLAN	Sharing Type
<input type="checkbox"/>	default	<input type="radio"/>	None
<input checked="" type="checkbox"/>	vlan11_Appliance	<input checked="" type="radio"/>	None

Ethernet Target Endpoint

Name:

MAC Address:

OK Cancel

25. Repeat steps 1 through 24 for configuring appliance port on port 3, 4, and 5 for configuring Cisco UCS C3160 on the expansion racks.

Enabling Server Ports

These steps provide details for enabling server ports.

26. Select the **Equipment** tab on the top left of the window.
27. Select **Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module**.

28. Expand the Unconfigured Ethernet Ports section.
29. Select all the ports that are connected to the Servers right-click them, and select **Reconfigure > Configure as a Server Port**.
30. A pop-up window appears to confirm your selection. Click **Yes** then **OK** to continue.
31. Select **Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module**.
32. Expand the Unconfigured Ethernet Ports section.
33. Select all the ports that are connected to the Servers right-click them, and select **Reconfigure > Configure as a Server Port**.
34. A pop-up window appears to confirm your selection. Click **Yes** then **OK** to continue.

Figure 25 Enabling Server Ports

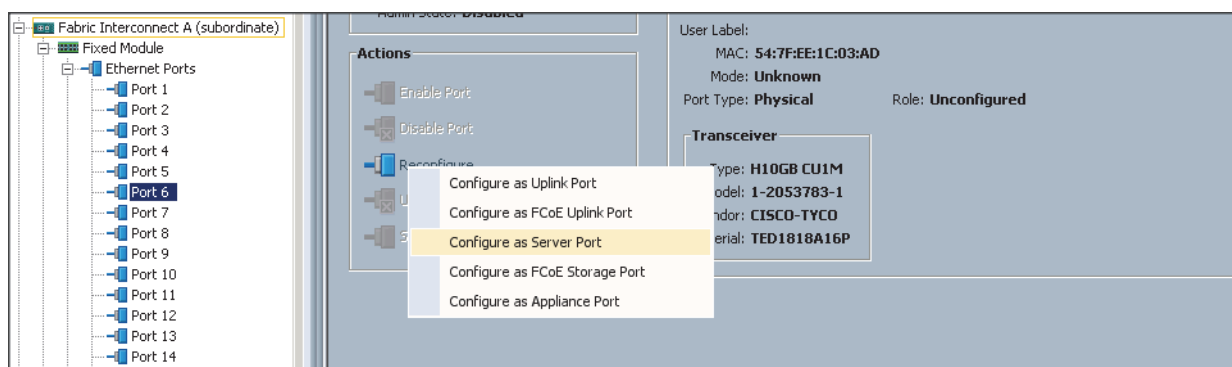


Figure 26 Showing Servers, Appliance and Uplink Ports

Slot	Port ID	MAC	If Role	If Type	Overall Status	Administrative State
1	1	54:7F:EE:1C:03:A8	Network	Physical	↑ Up	↑ Enabled
1	2	54:7F:EE:1C:03:A9	Appliance Storage	Physical	↑ Up	↑ Enabled
1	3	54:7F:EE:1C:03:AA	Appliance Storage	Physical	↑ Up	↑ Enabled
1	4	54:7F:EE:1C:03:AB	Appliance Storage	Physical	↑ Up	↑ Enabled
1	5	54:7F:EE:1C:03:AC	Appliance Storage	Physical	↑ Up	↑ Enabled
1	6	54:7F:EE:1C:03:AD	Server	Physical	↑ Up	↑ Enabled
1	7	54:7F:EE:1C:03:AE	Server	Physical	↑ Up	↑ Enabled
1	8	54:7F:EE:1C:03:AF	Server	Physical	↑ Up	↑ Enabled
1	9	54:7F:EE:1C:03:B0	Server	Physical	↑ Up	↑ Enabled
1	10	54:7F:EE:1C:03:B1	Server	Physical	↑ Up	↑ Enabled
1	11	54:7F:EE:1C:03:B2	Server	Physical	↑ Up	↑ Enabled
1	12	54:7F:EE:1C:03:B3	Server	Physical	↑ Up	↑ Enabled
1	13	54:7F:EE:1C:03:B4	Server	Physical	↑ Up	↑ Enabled
1	14	54:7F:EE:1C:03:B5	Server	Physical	↑ Up	↑ Enabled
1	15	54:7F:EE:1C:03:B6	Server	Physical	↑ Up	↑ Enabled
1	16	54:7F:EE:1C:03:B7	Server	Physical	↑ Up	↑ Enabled
1	17	54:7F:EE:1C:03:B8	Server	Physical	↑ Up	↑ Enabled
1	18	54:7F:EE:1C:03:B9	Server	Physical	↑ Up	↑ Enabled
1	19	54:7F:EE:1C:03:BA	Server	Physical	↑ Up	↑ Enabled
1	20	54:7F:EE:1C:03:BB	Server	Physical	↑ Up	↑ Enabled
1	21	54:7F:EE:1C:03:BC	Server	Physical	↑ Up	↑ Enabled
1	22	54:7F:EE:1C:03:BD	Server	Physical	↑ Up	↑ Enabled
1	23	54:7F:EE:1C:03:BE	Server	Physical	↑ Up	↑ Enabled
1	24	54:7F:EE:1C:03:BF	Server	Physical	↑ Up	↑ Enabled
1	25	54:7F:EE:1C:03:C0	Server	Physical	↑ Up	↑ Enabled
1	26	54:7F:EE:1C:03:C1	Server	Physical	↑ Up	↑ Enabled
1	27	54:7F:EE:1C:03:C2	Server	Physical	↑ Up	↑ Enabled
1	28	54:7F:EE:1C:03:C3	Server	Physical	↑ Up	↑ Enabled
1	29	54:7F:EE:1C:03:C4	Server	Physical	↑ Up	↑ Enabled
1	30	54:7F:EE:1C:03:C5	Server	Physical	↑ Up	↑ Enabled
1	31	54:7F:EE:1C:03:C6	Server	Physical	↑ Up	↑ Enabled
1	32	54:7F:EE:1C:03:C7	Server	Physical	↑ Up	↑ Enabled

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, thereby enabling multi-tenancy of the compute resources. This document does not assume the use of Organizations; however the necessary steps are provided for future reference.

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

1. Click **New** on the top left corner in the right pane in 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**.
6. Click **OK** in the success message box.

Creating MAC Address Pools

Follow these steps to create MAC address pools:

1. Select the **LAN** tab on the left of the window.
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 for the name of the MAC pool.
5. (Optional) Enter a description of the MAC pool.
6. Select Assignment Order Sequential.
7. Click **Next**.
8. Click **Add**.
9. Specify a starting MAC address.
10. Specify a size of the MAC address pool, which is sufficient to support the available server resources.
11. Click **OK**.

Figure 27 Creating MAC Pool Window

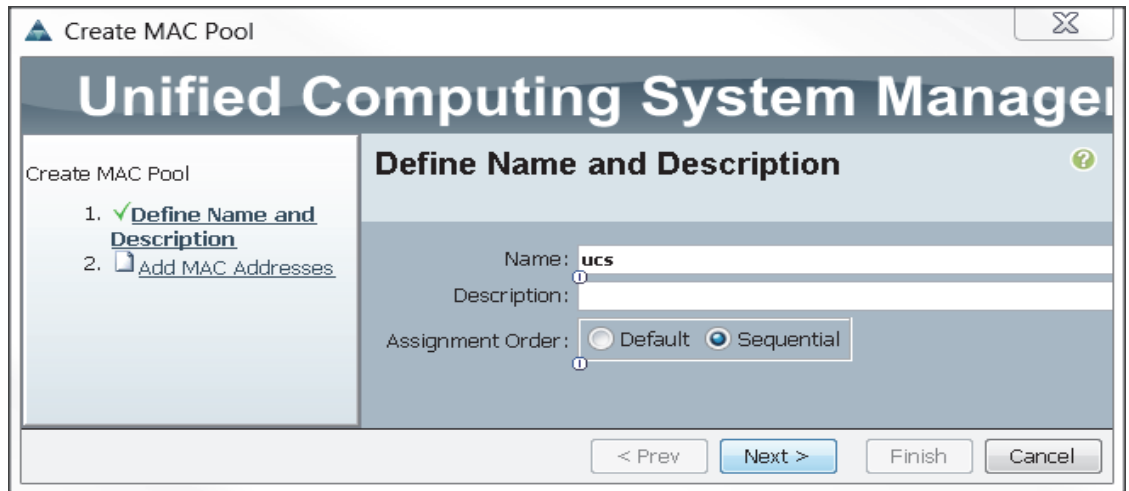
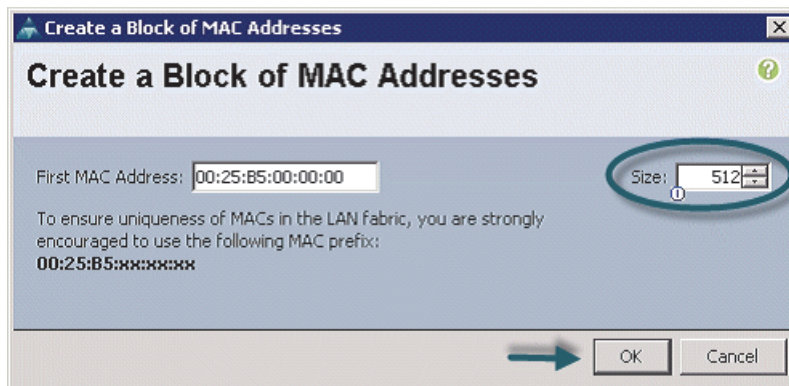
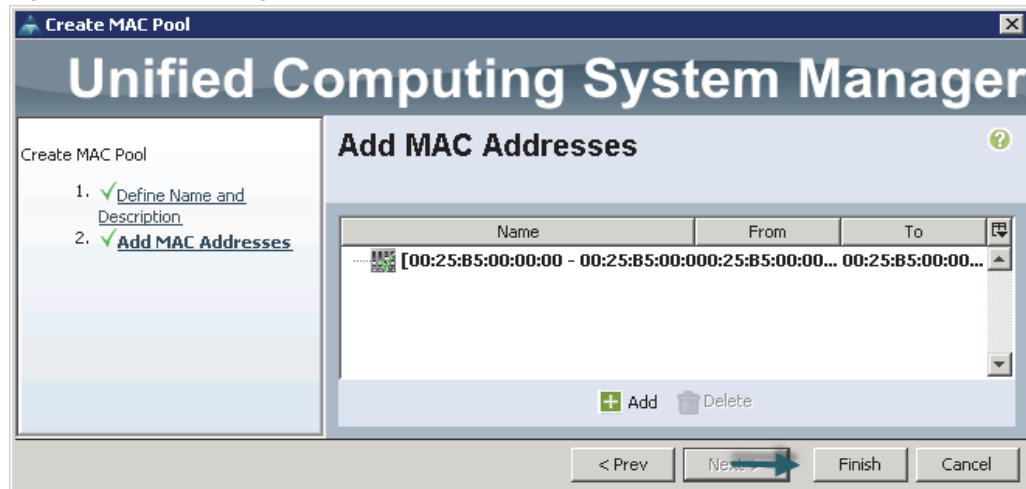


Figure 28 Specifying First MAC Address and Size



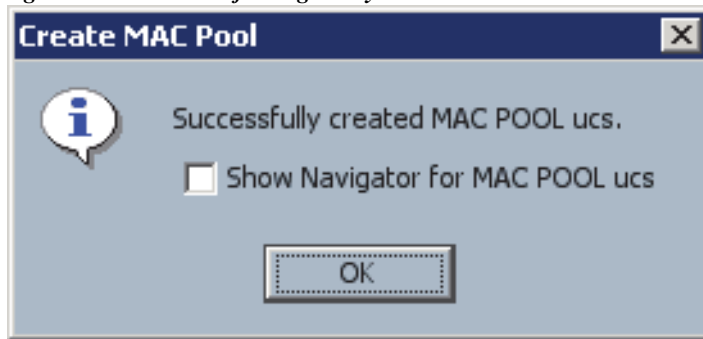
12. Click **Finish**.

Figure 29 Adding MAC Addresses



13. When the message box displays, click **OK**.

Figure 30 *Confirming Newly Added MAC Pool*



Creating Server Pools

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

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

1. Select the **Servers** tab in the left pane in the UCS Manager GUI.
2. Select **Pools > root**.
3. Right-click the **Server Pools**.
4. Select **Create Server Pool**.
5. Enter your required name (ucs) for the Server Pool in the name text box.
6. (Optional) enter a description for the organization
7. Click **Next** to add the servers.

Figure 31 Setting Name and Description of Server Pool

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

8. Select all the Cisco UCS C240M4SX servers to be added to the server pool you previously created (ucs), then Click >> to add them to the pool.
9. Click **Finish**.
10. Click **OK**, and then click **Finish**.

Figure 32 Adding Servers to the Server Pool



Creating Policies for Service Profile Templates

Creating Host Firmware Package Policy

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

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

1. Select the **Servers** tab in the left pane in the UCS Manager GUI.
2. Select **Policies > root**.
3. Right-click **Host Firmware Packages**.
4. Select **Create Host Firmware Package**.
5. Enter your required Host Firmware package name (ucs).

6. Click the **Simple** radio button to configure the Host Firmware package.
7. Select the appropriate Rack package that you have.
8. Click **OK** to complete creating the management firmware package.
9. Click **OK**.

Figure 33 *Creating Host Firmware Package*

Create Host Firmware Package

Name:

Description:

How would you like to configure the Host Firmware Package? Simple Advanced

Blade Package:

Rack Package:

OK Cancel

Creating QoS Policies

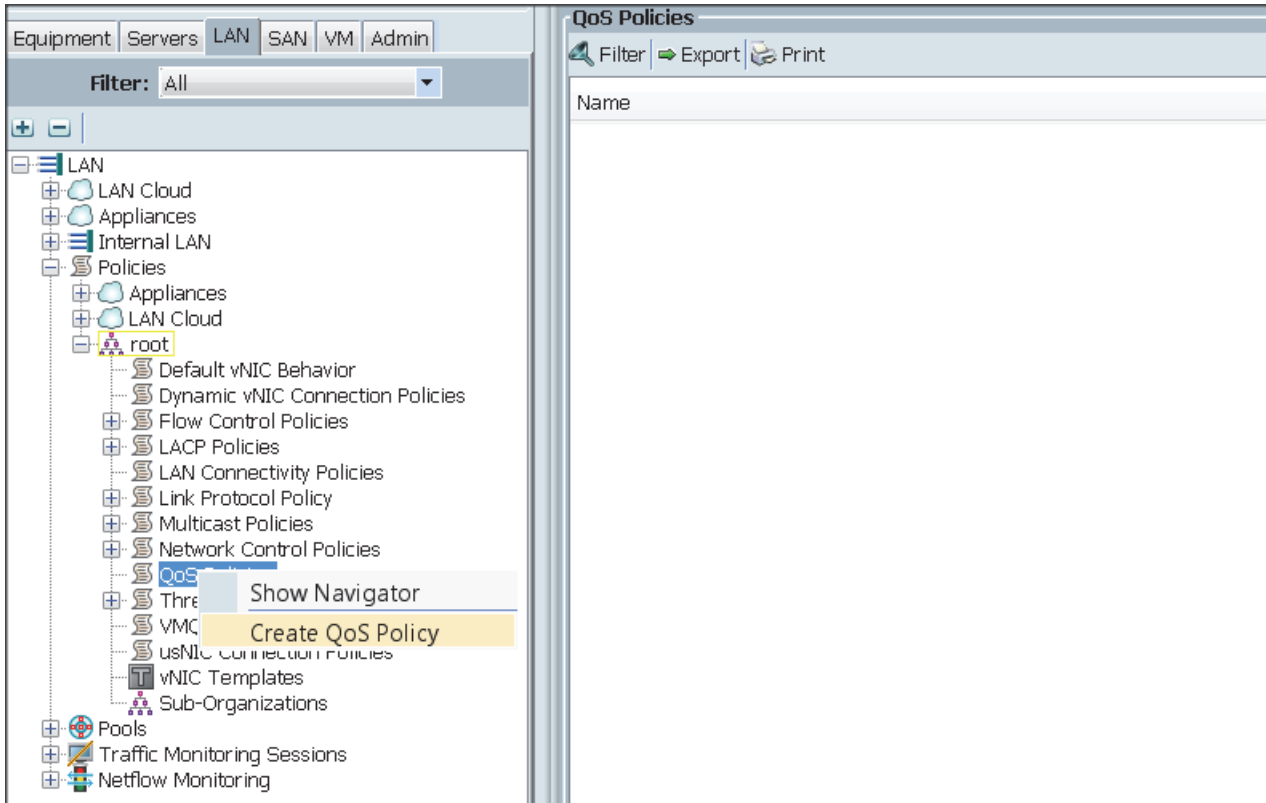
Follow these steps to create the QoS policy for a given server configuration using the Cisco UCS Manager GUI:

Best Effort Policy

1. Select the **LAN** tab in the left pane in the UCS Manager GUI.
2. Select **Policies > root**.

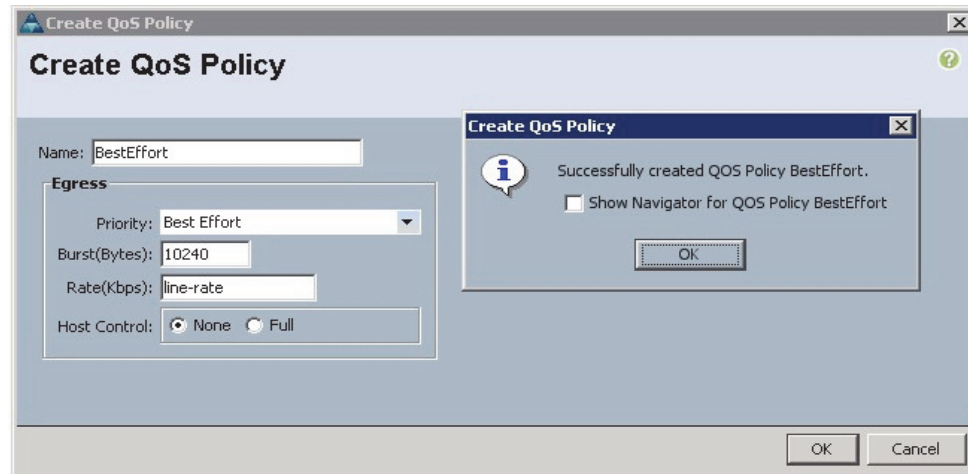
3. Right-click **QoS Policies**.
4. Select **Create QoS Policy**.

Figure 34 *Creating QoS Policy*



5. Enter BestEffort as the name of the policy.
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 Host Control radio button as default (none).
10. Once the pop-up window appears, click **OK** to complete the creation of the Policy.

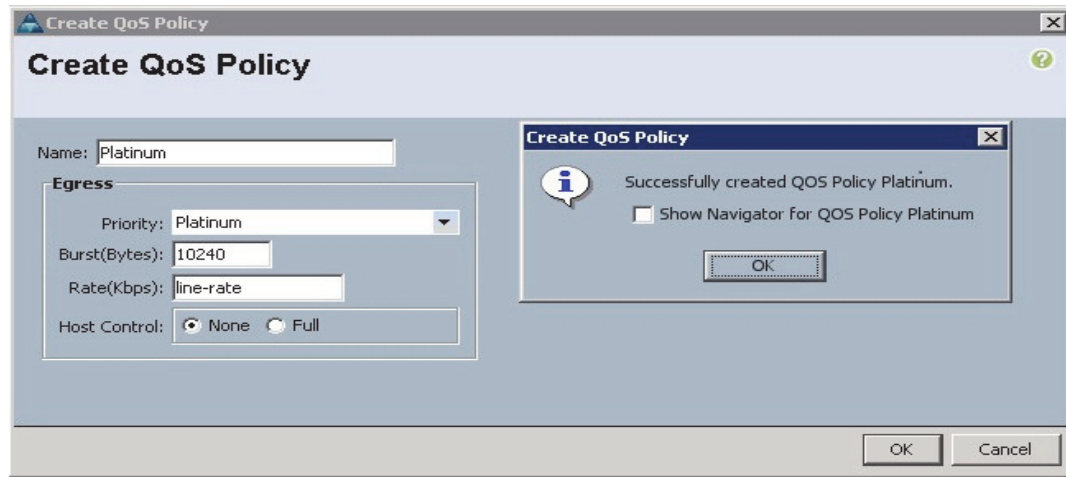
Figure 35 Creating BestEffort QoS Policy



Platinum Policy

1. Select the **LAN** tab in the left pane in the UCS Manager GUI.
2. Select **Policies > root**.
3. Right-click **QoS Policies**.
4. Select **Create QoS Policy**.
5. Enter Platinum as the name of the policy.
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 Host Control radio button as default (none).
10. Once the pop-up window appears, click **OK** to complete the creation of the Policy.

Figure 36 Creating Platinum QoS Policy

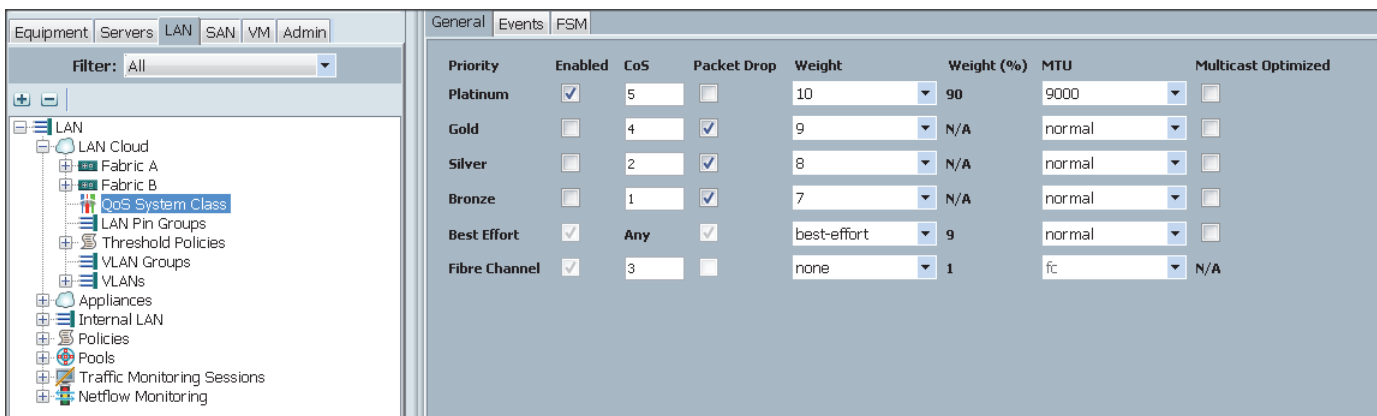


Setting Jumbo Frames

Follow these steps for setting up the Jumbo frames and enabling QoS:

1. Select the **LAN** tab in the left pane in the UCS Manager GUI.
2. Select **LAN Cloud > QoS System Class**.
3. In the right pane, select the **General** tab
4. In the Platinum row, enter 9000 for MTU.
5. Check the **Enabled** Check box next to Platinum.
6. In the Best Effort row, select best-effort for weight.
7. In the Fiber Channel row, select none for weight.
8. Click **Save Changes**.
9. Click **OK**.

Figure 37 Setting Jumbo Frames

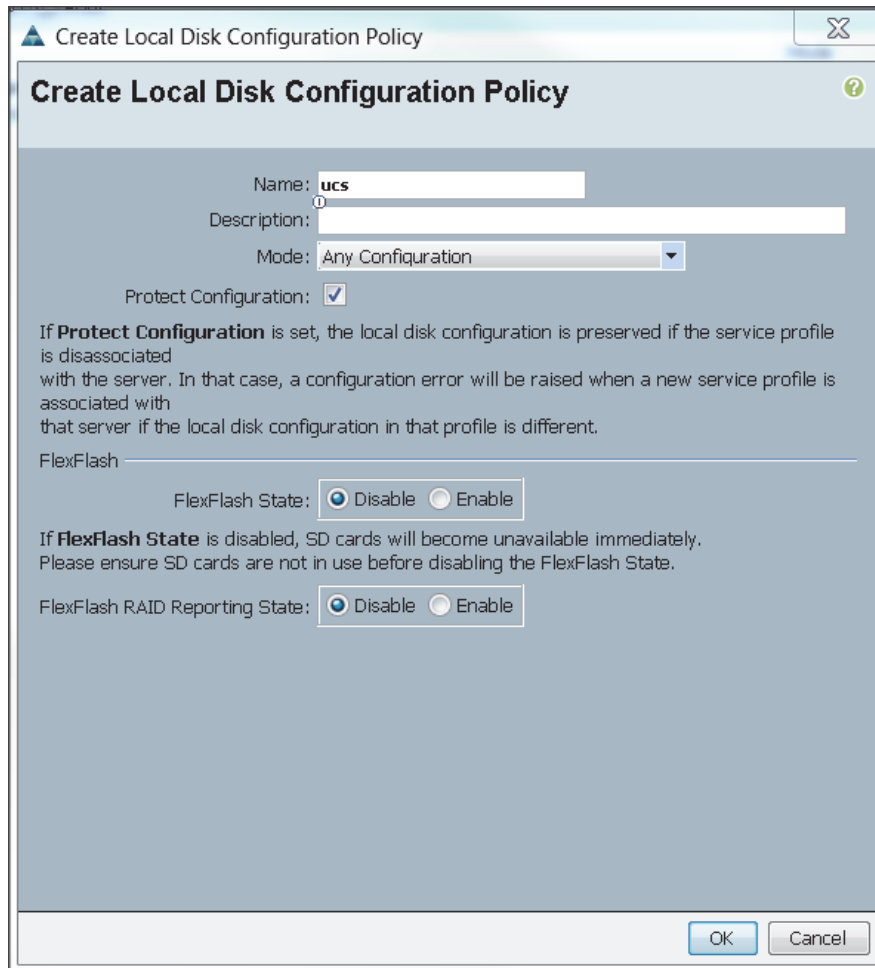


Creating Local Disk Configuration Policy

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

1. Select the **Servers** tab on the left pane in the UCS Manager GUI.
2. Go to **Policies > root**.
3. Right-click Local Disk Config Policies.
4. Select Create Local Disk Configuration Policy.
5. Enter **ucs** as the local disk configuration policy name.
6. Change the Mode to Any Configuration. Check the **Protect Configuration** box.
7. Keep the FlexFlash State field as default (Disable).
8. Keep the FlexFlash RAID Reporting State field as default (Disable).
9. Click **OK** to complete the creation of the Local Disk Configuration Policy.
10. Click **OK**.

Figure 38 Configuring Local Disk Policy



Creating Server BIOS Policy

The BIOS policy feature in Cisco UCS automates the BIOS configuration process. The traditional method of setting the BIOS is done manually 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 GUI:

1. Select the **Servers** tab in the left pane in the UCS Manager GUI.
2. Select **Policies > root**.
3. Right-click **BIOS Policies**.

4. Select Create BIOS Policy.
5. Enter your preferred BIOS policy name (ucs).
6. Change the BIOS settings as per the following figures:

Figure 39 Creating Server BIOS Policy

Create BIOS Policy

Unified Computing System Manager

Create BIOS Policy

1. **Main**
2. Processor
3. Intel Directed IO
4. RAS Memory
5. Serial Port
6. USB
7. PCI
8. QPI
9. LOM and PCIe Slots
10. Boot Options
11. Server Management

Main

Name:

Description:

Reboot on BIOS Settings Change:

Quiet Boot: disabled enabled Platform Default

Post Error Pause: disabled enabled Platform Default

Resume Ac On Power Loss: stay-off last-state reset Platform Default

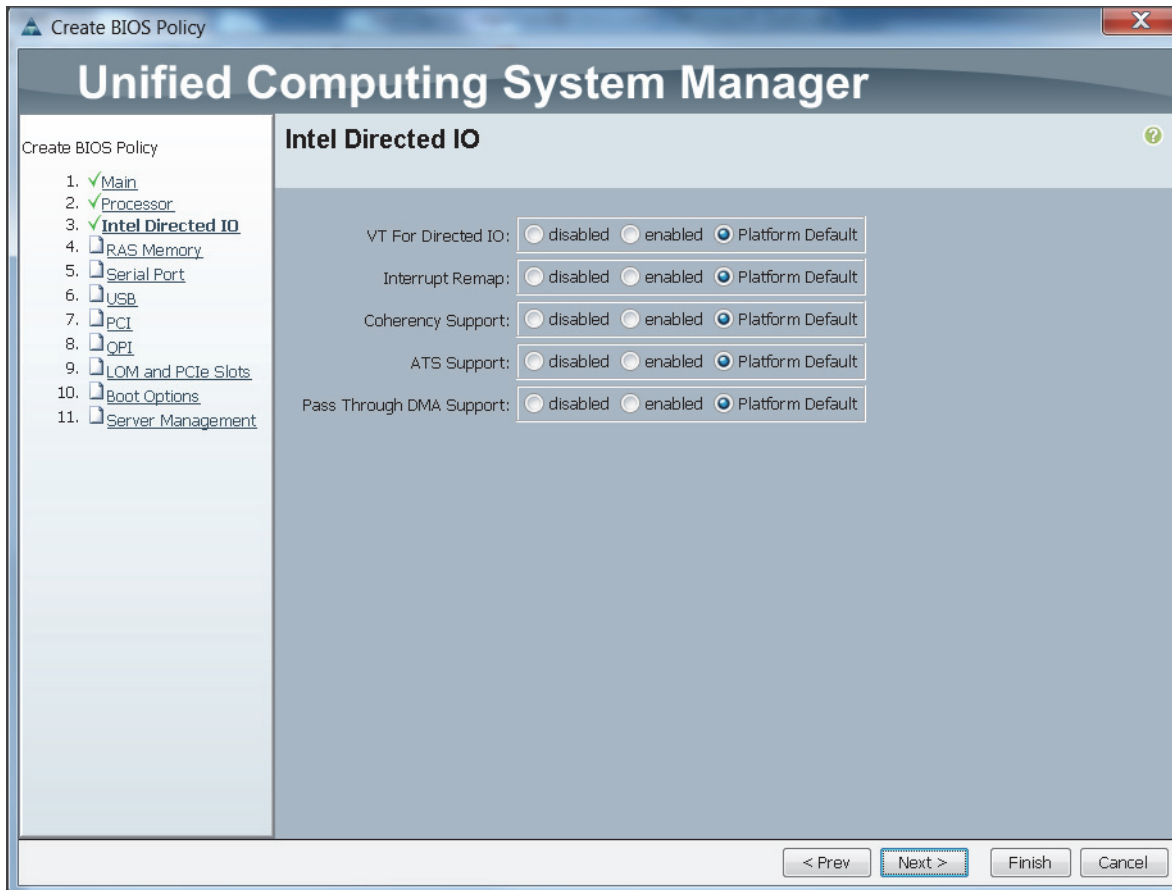
Front Panel Lockout: disabled enabled Platform Default

< Prev Next > Finish Cancel

Figure 40 Creating Server BIOS Policy for Processor

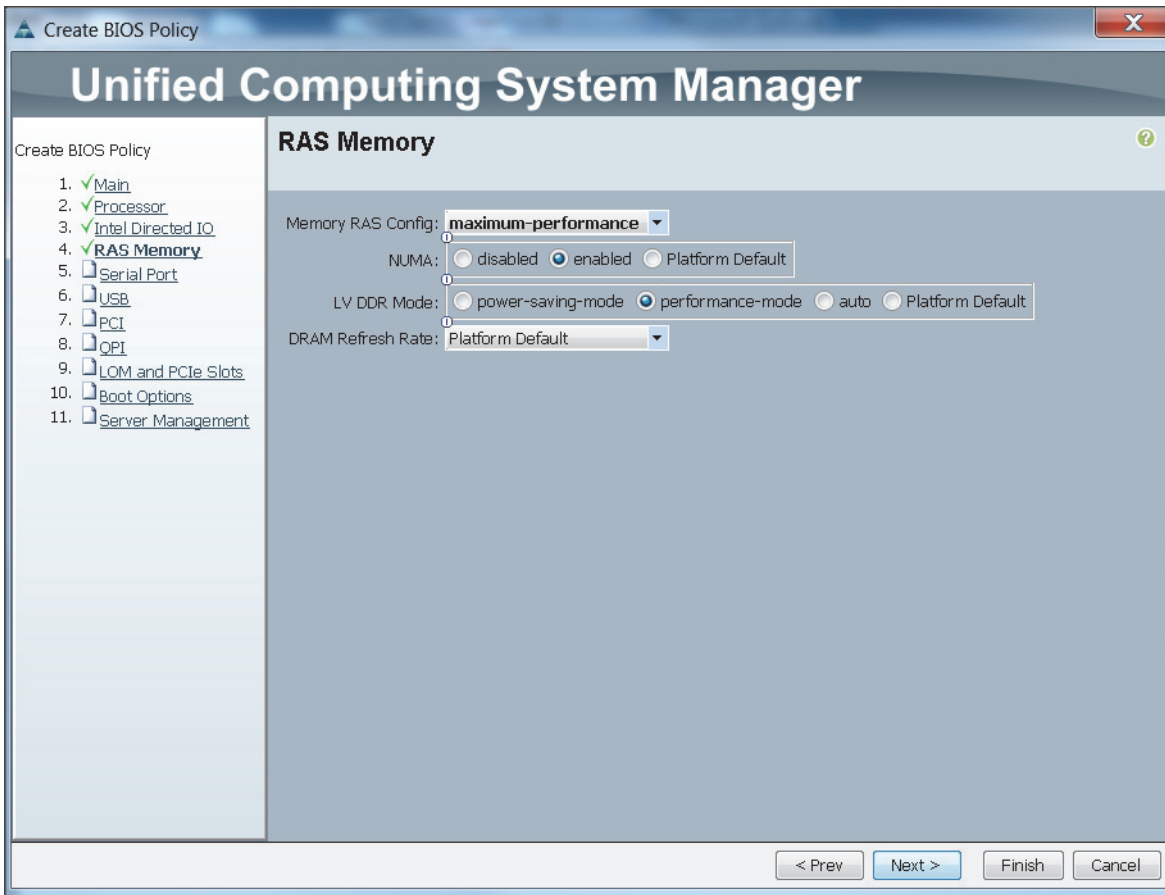


Figure 41 Creating Server BIOS Policy for Intel Directed IO



7. Click **Finish** to complete creating the BIOS policy.
8. Click **OK**.

Figure 42 Creating Server BIOS Policy for Memory

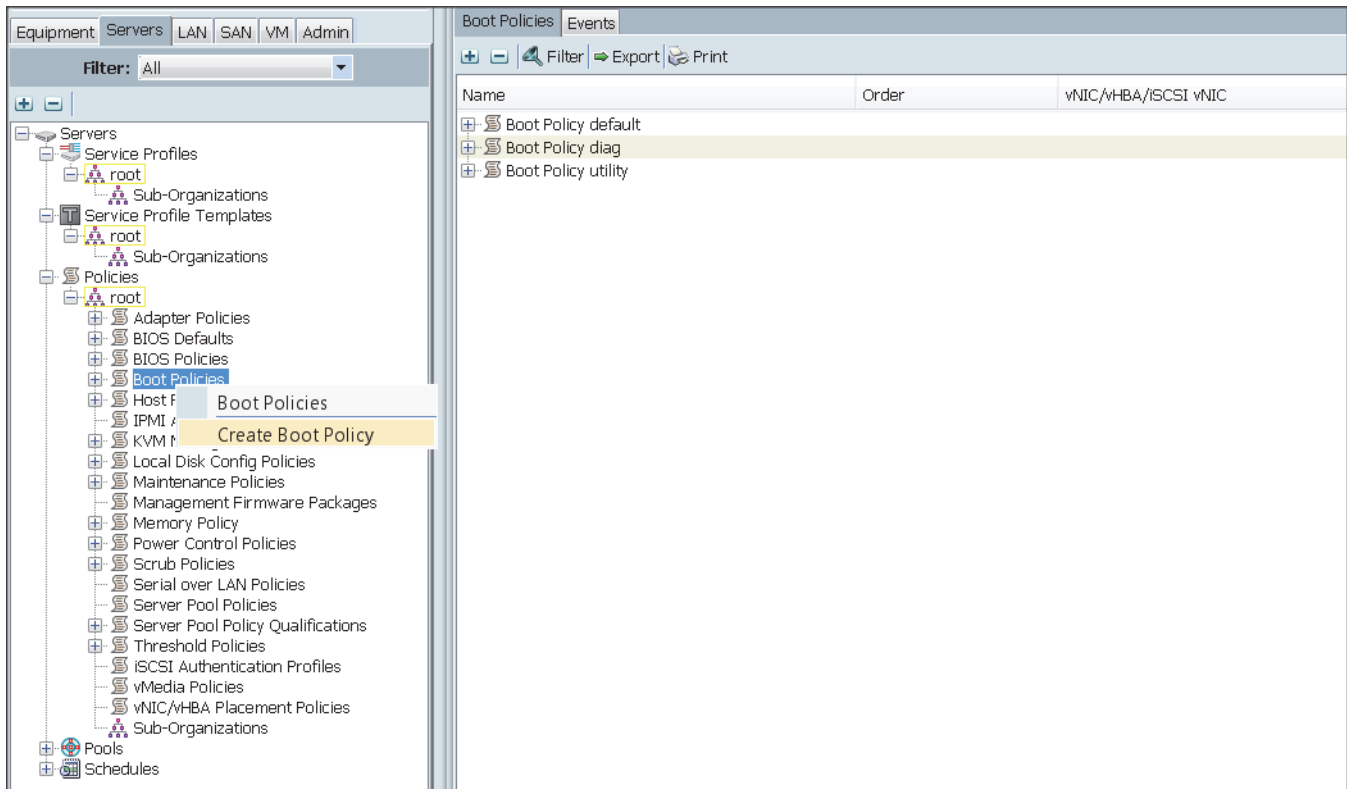


Creating Boot Policy

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

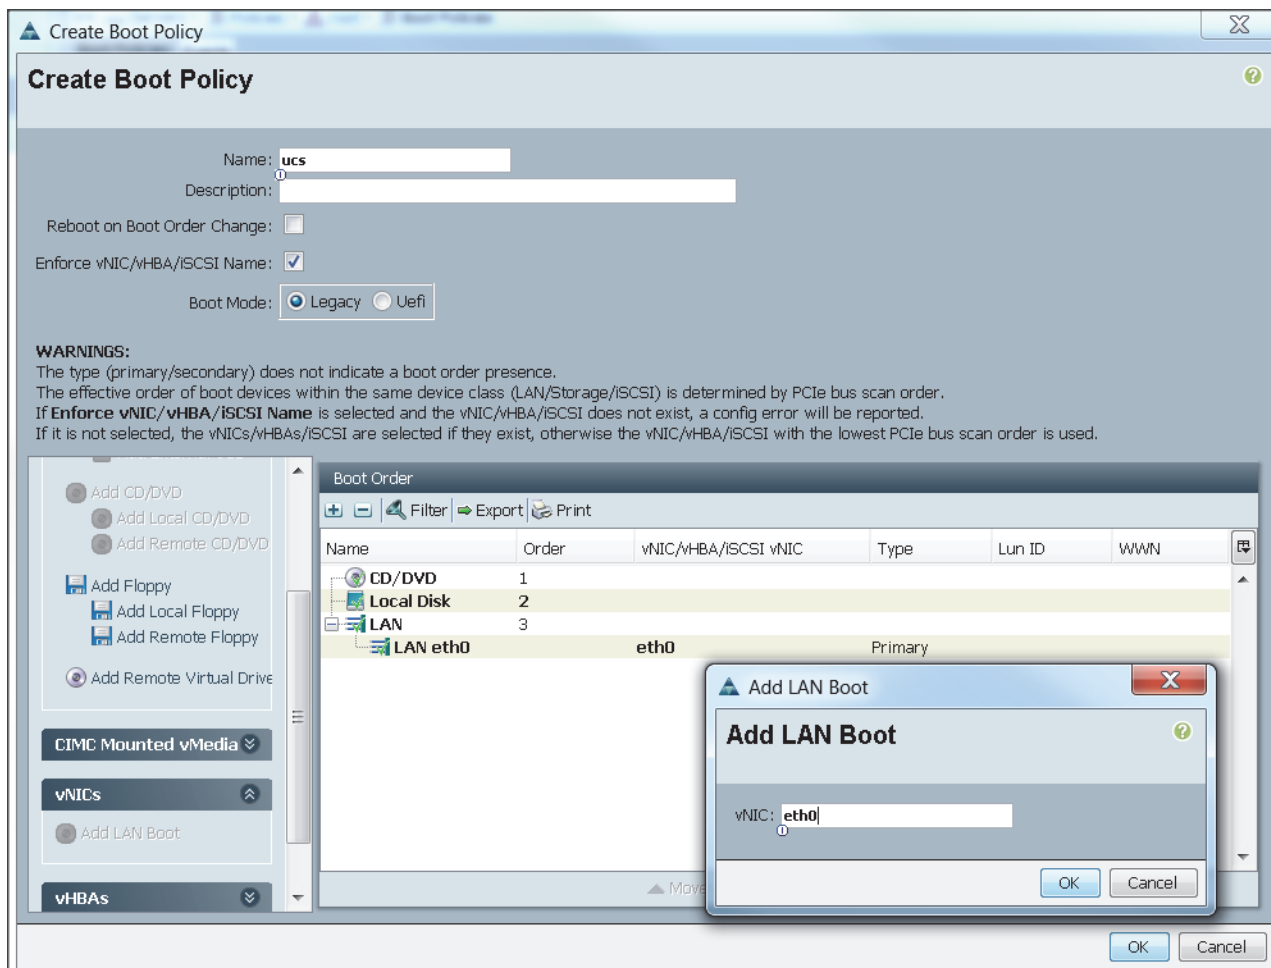
1. Select the **Servers** tab in the left pane in the UCS Manager GUI.
2. Select **Policies > root**.
3. Right-click the **Boot Policies**.
4. Select **Create Boot Policy**.

Figure 43 Creating Boot Policy Part 1



5. Enter ucs as the boot policy name.
6. (Optional) enter a description for the boot policy.
7. Keep the Reboot on Boot Order Change check box unchecked.
8. Keep Enforce vNIC/vHBA/iSCSI Name check box checked.
9. Keep Boot Mode Default (Legacy).
10. Expand **Local Devices > Add CD/DVD** and select **Add Local CD/DVD**.
11. Expand Local Devices and select **Add Local Disk**.
12. Expand vNICs and select Add LAN Boot and enter eth0.
13. Click **OK** to add the Boot Policy.
14. Click **OK**.

Figure 44 Creating Boot Policy Part 2

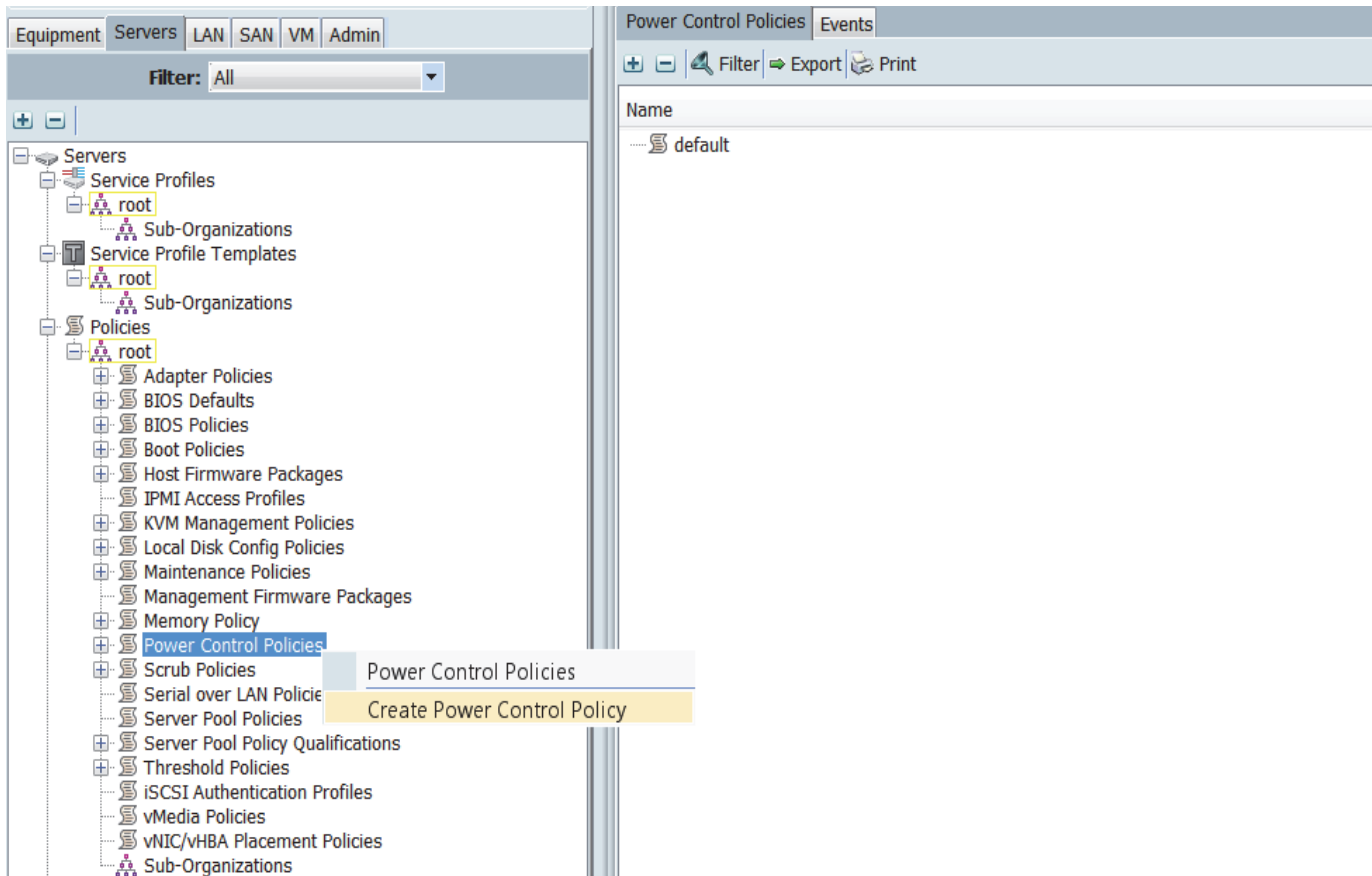


Creating Power Control Policy

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

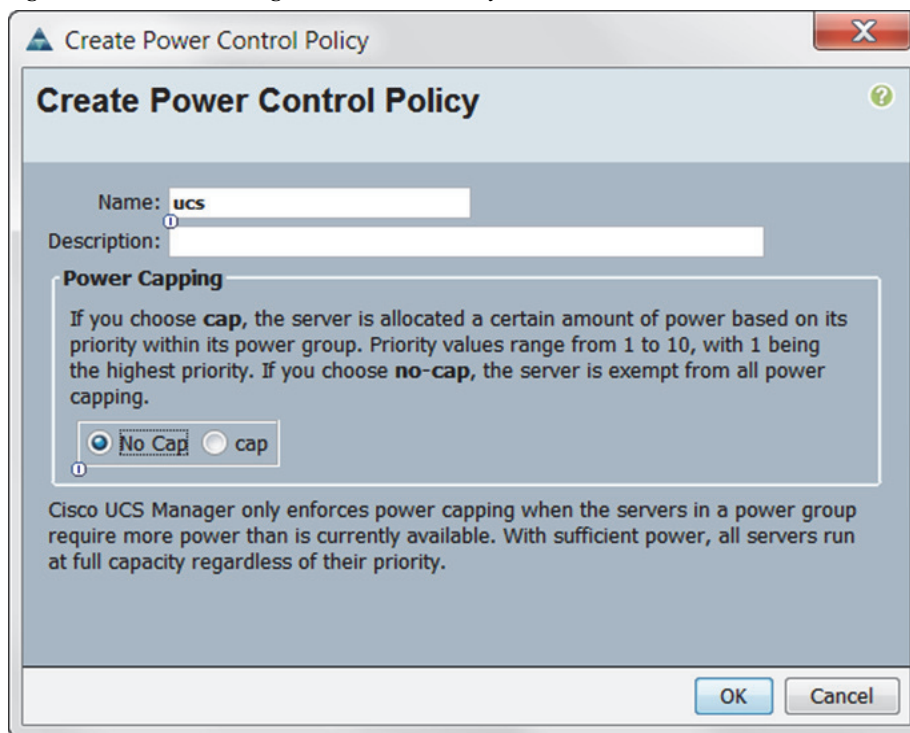
15. Select the **Servers** tab in the left pane in the UCS Manager GUI.
16. Select **Policies > root**.
17. Right-click the **Power Control Policies**.
18. Select **Create Power Control Policy**.

Figure 45 Creating Power Control Policy Part 1



19. Enter ucs as the Power Control policy name.
20. (Optional) enter a description for the boot policy.
21. Select **No cap** for Power Capping selection.
22. Click **OK** to the Power Control Policy.
23. Click **OK**.

Figure 46 Creating Power Control Policy Part 2

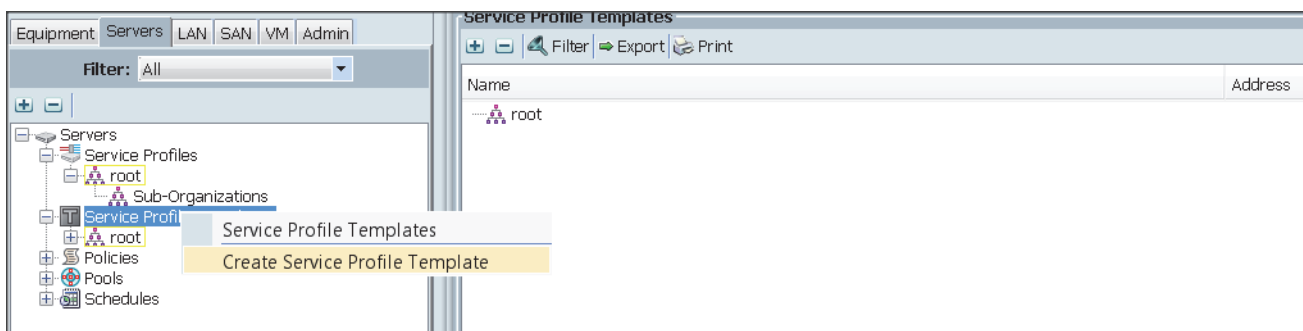


Creating Service Profile Template

To create a service profile template, follow these steps:

1. Select the **Servers** tab in the left pane in the UCS Manager GUI.
2. Right-click **Service Profile Templates**.
3. Select **Create Service Profile Template**.

Figure 47 Creating Service Profile Template



4. The Create Service Profile Template window appears.

These steps below provide a detailed configuration procedure to identify the service profile template:

- a. Name the service profile template as **ucs**. Click the **Updating Template** radio button.

- b. In the UUID section, select **Hardware Default** as the UUID pool.
- c. Click **Next** to continue to the next section.

Figure 48 Identify Service Profile Template

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.

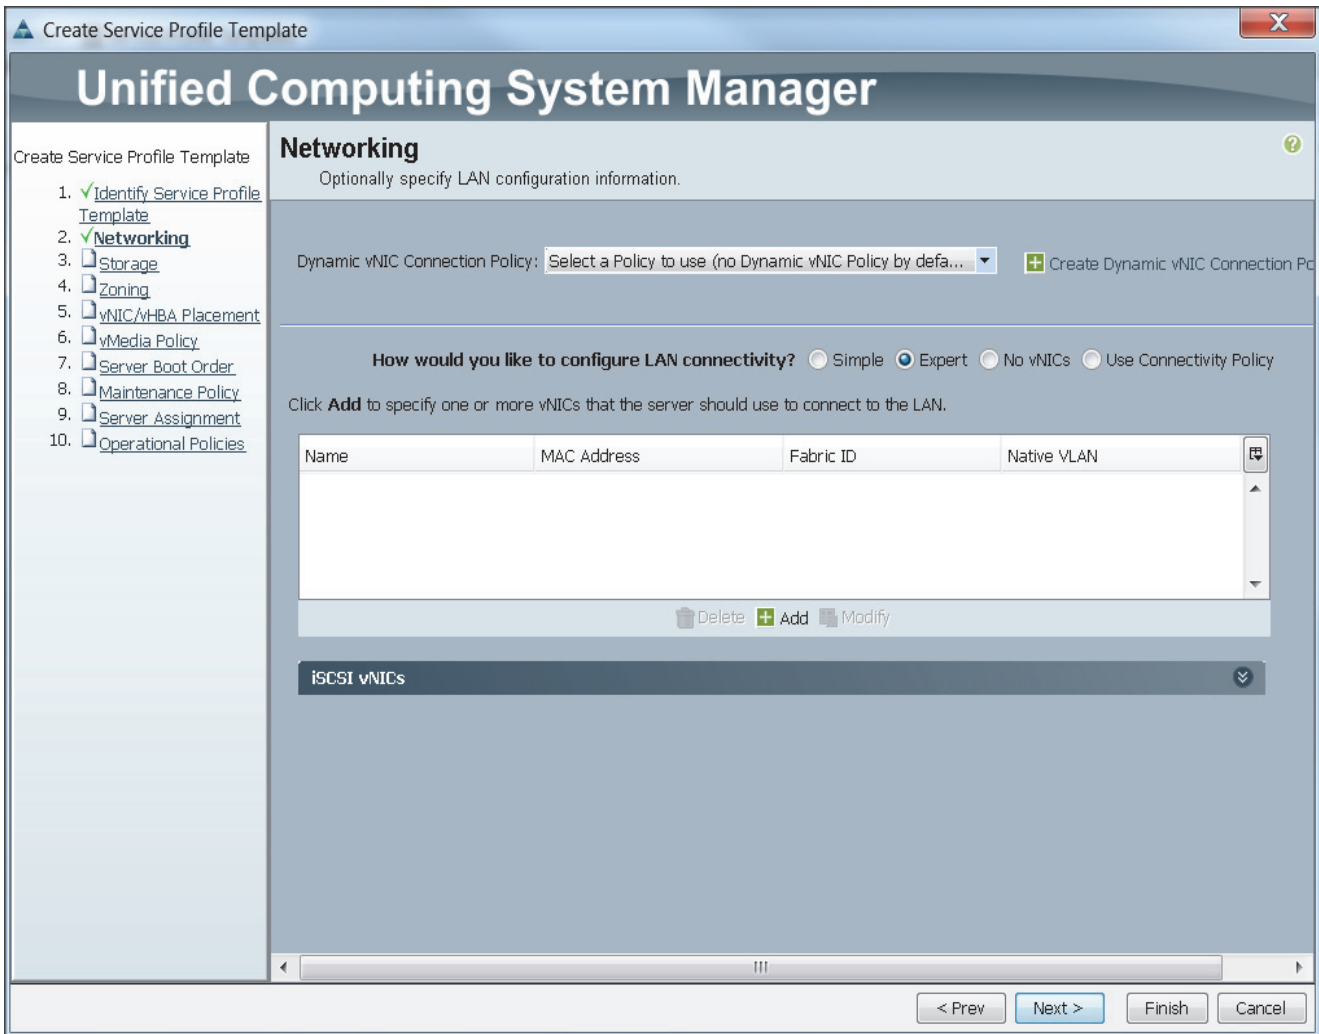
Optionally enter a description for the profile. The description can contain information about when and where the service profile should be used.

< Prev Next > Finish Cancel

Configuring Network Settings for the Template

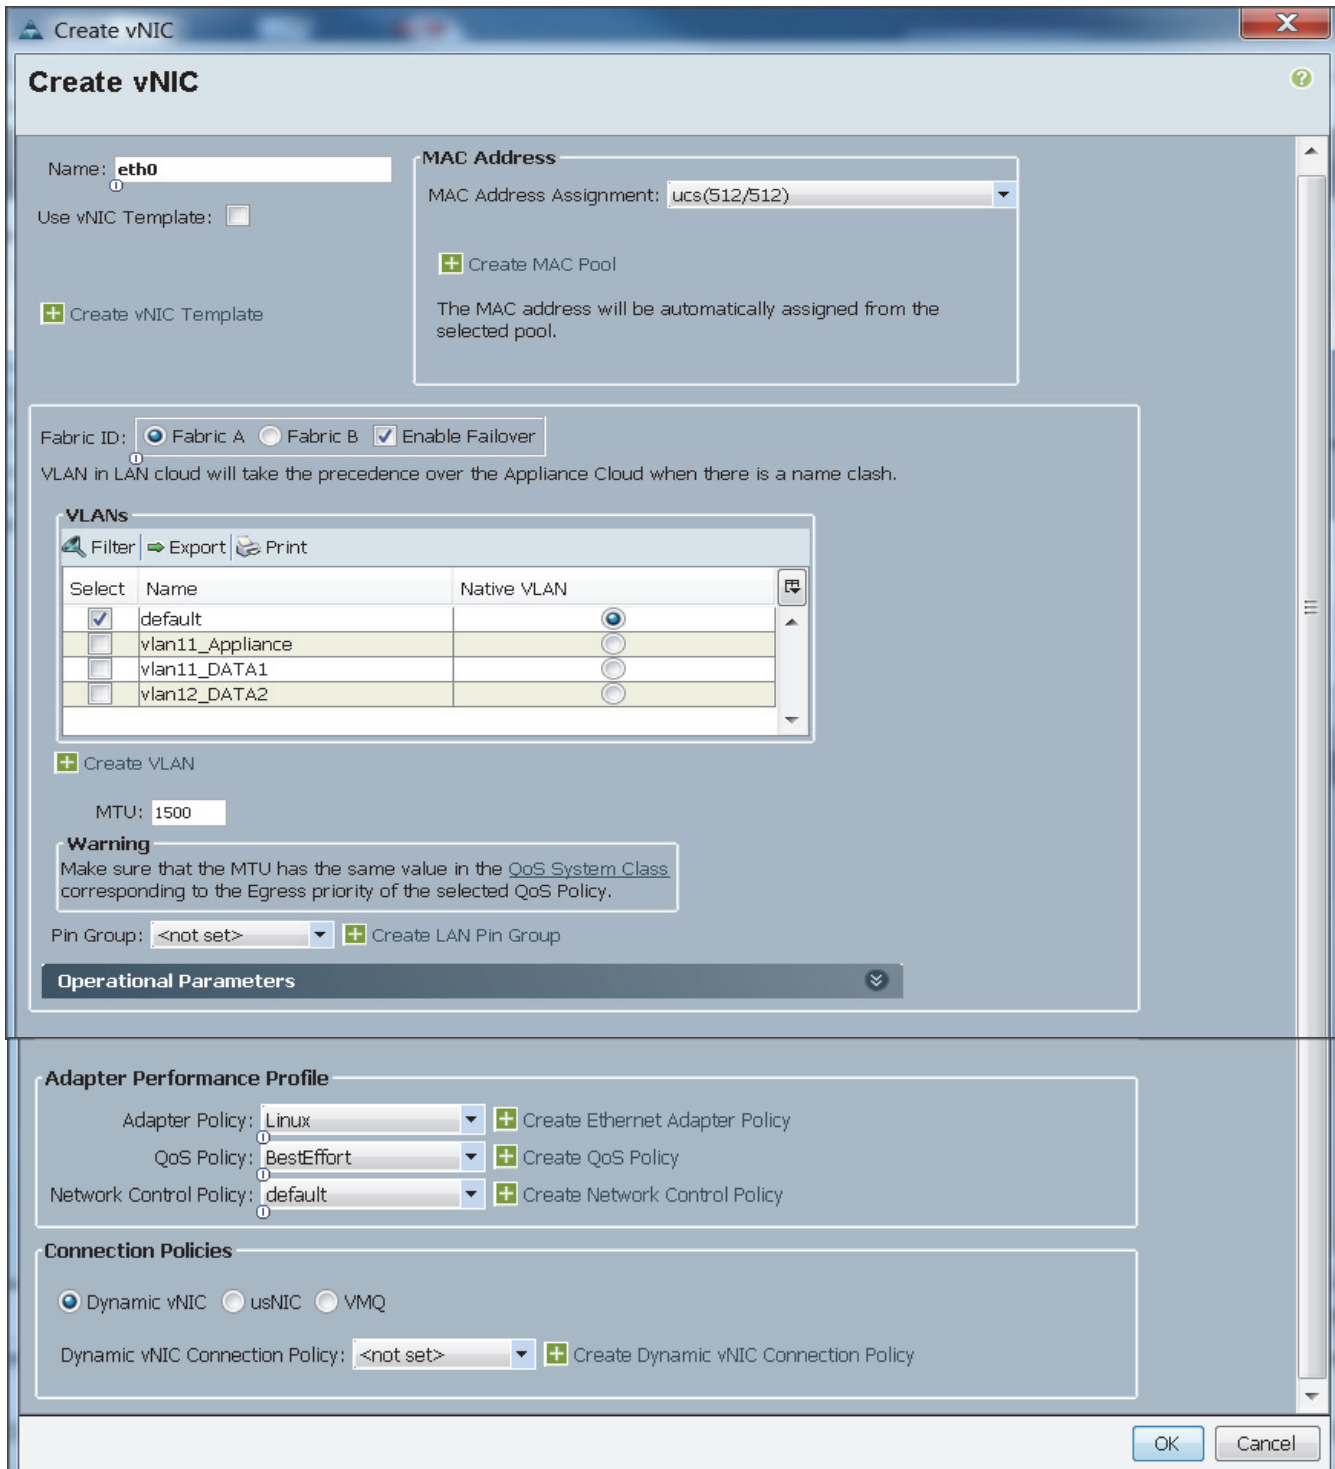
1. Keep the Dynamic vNIC Connection Policy field at the default.
2. Click the **Expert** radio button for the option, **how would you like to configure LAN connectivity?**
3. Click **Add** to add a vNIC to the template.

Figure 49 Configuring Network Settings for the Template



4. The Create vNIC window displays. Name the vNIC as eth0.
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 **default** check box for VLANs and click the **Native VLAN** radio button.
8. Select MTU size as 1500
9. Select adapter policy as Linux
10. Select QoS Policy as BestEffort.
11. Keep the Network Control Policy as Default.
12. Keep the Connection Policies as Dynamic vNIC.
13. Keep the Dynamic vNIC Connection Policy as <not set>.
14. Click **OK**.

Figure 50 Configuring vNIC eth0



15. Click **Add** to add a vNIC to the template.
16. The Create vNIC window appears. Name the vNIC eth1.
17. Select ucs in the Mac Address Assignment pool.

18. Click the **Fabric B** radio button and Check the **Enable failover** check box for the Fabric ID.
19. Check the **vlan11_DATA1** check box for VLANs, and click the **Native VLAN** radio button
20. Select MTU size as 9000
21. Select adapter policy as Linux
22. Select QoS Policy as Platinum.
23. Keep the Network Control Policy as Default.
24. Keep the Connection Policies as Dynamic vNIC.
25. Keep the Dynamic vNIC Connection Policy as <not set>.
26. Click **OK**.

Figure 51 Configuring vNIC eth1

Create vNIC

Name:

Use vNIC Template:

MAC Address

MAC Address Assignment:

+ Create MAC Pool

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

Fabric ID: Fabric A Fabric B Enable Failover

VLAN in LAN cloud will take the precedence over the Appliance Cloud when there is a name clash.

VLANs

Filter Export Print

Select	Name	Native VLAN
<input type="checkbox"/>	default	<input type="radio"/>
<input type="checkbox"/>	vlan11_Appliance	<input type="radio"/>
<input checked="" type="checkbox"/>	vlan11_DATA1	<input checked="" type="radio"/>
<input type="checkbox"/>	vlan12_DATA2	<input type="radio"/>

+ Create VLAN

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: **+** Create LAN Pin Group

Operational Parameters

Adapter Performance Profile

Adapter Policy: **+** Create Ethernet Adapter Policy

QoS Policy: **+** Create QoS Policy

Network Control Policy: **+** Create Network Control Policy

Connection Policies

Dynamic vNIC usNIC VMQ

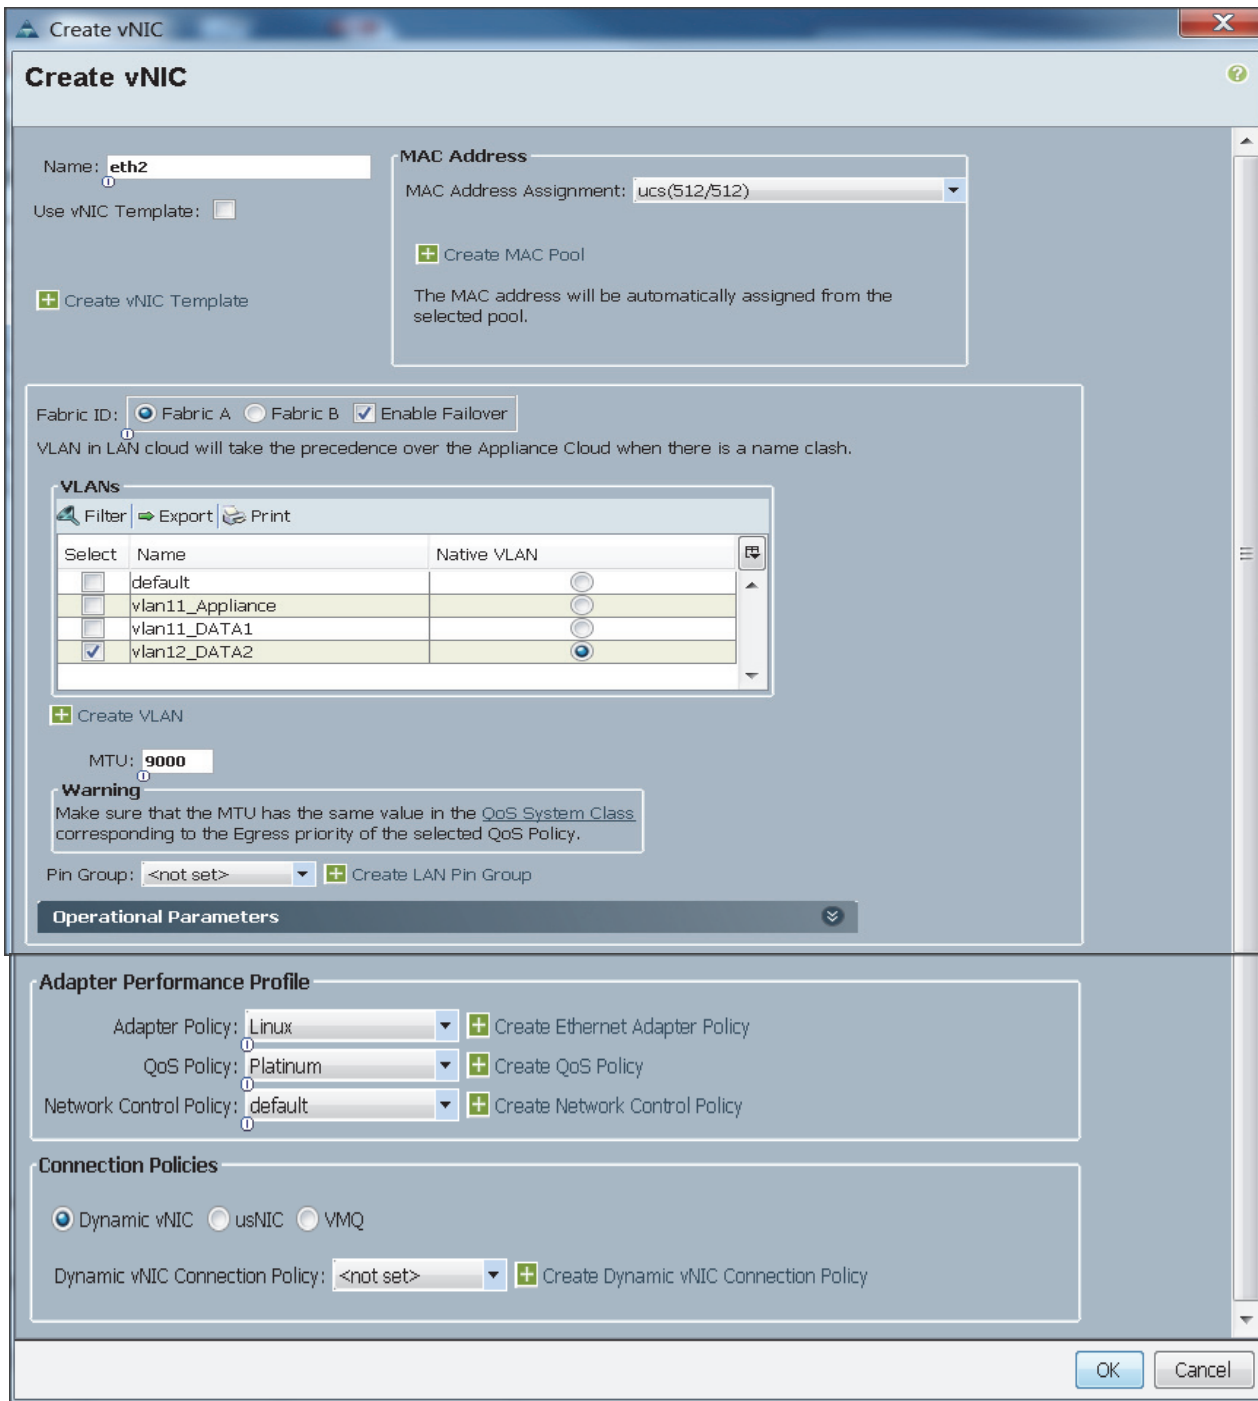
Dynamic vNIC Connection Policy: **+** Create Dynamic vNIC Connection Policy

OK Cancel

27. Click **Add** to add a vNIC to the template.
28. The Create vNIC window appears. Name the vNIC eth2.
29. Select ucs in the Mac Address Assignment pool.
30. Click the **Fabric A** radio button, and then Check the **Enable failover** check box for the Fabric ID.

31. Check the **vlan12_DATA2** check box for VLANs, and then click the **Native VLAN** radio button.
32. Select MTU size as 9000.
33. Select adapter policy as Linux.
34. Select QoS Policy as Platinum.
35. Keep the Network Control Policy as Default.
36. Keep the Connection Policies as Dynamic vNIC.
37. Keep the Dynamic vNIC Connection Policy as <not set>.
38. Click **OK**.

Figure 52 Configuring vNIC eth2



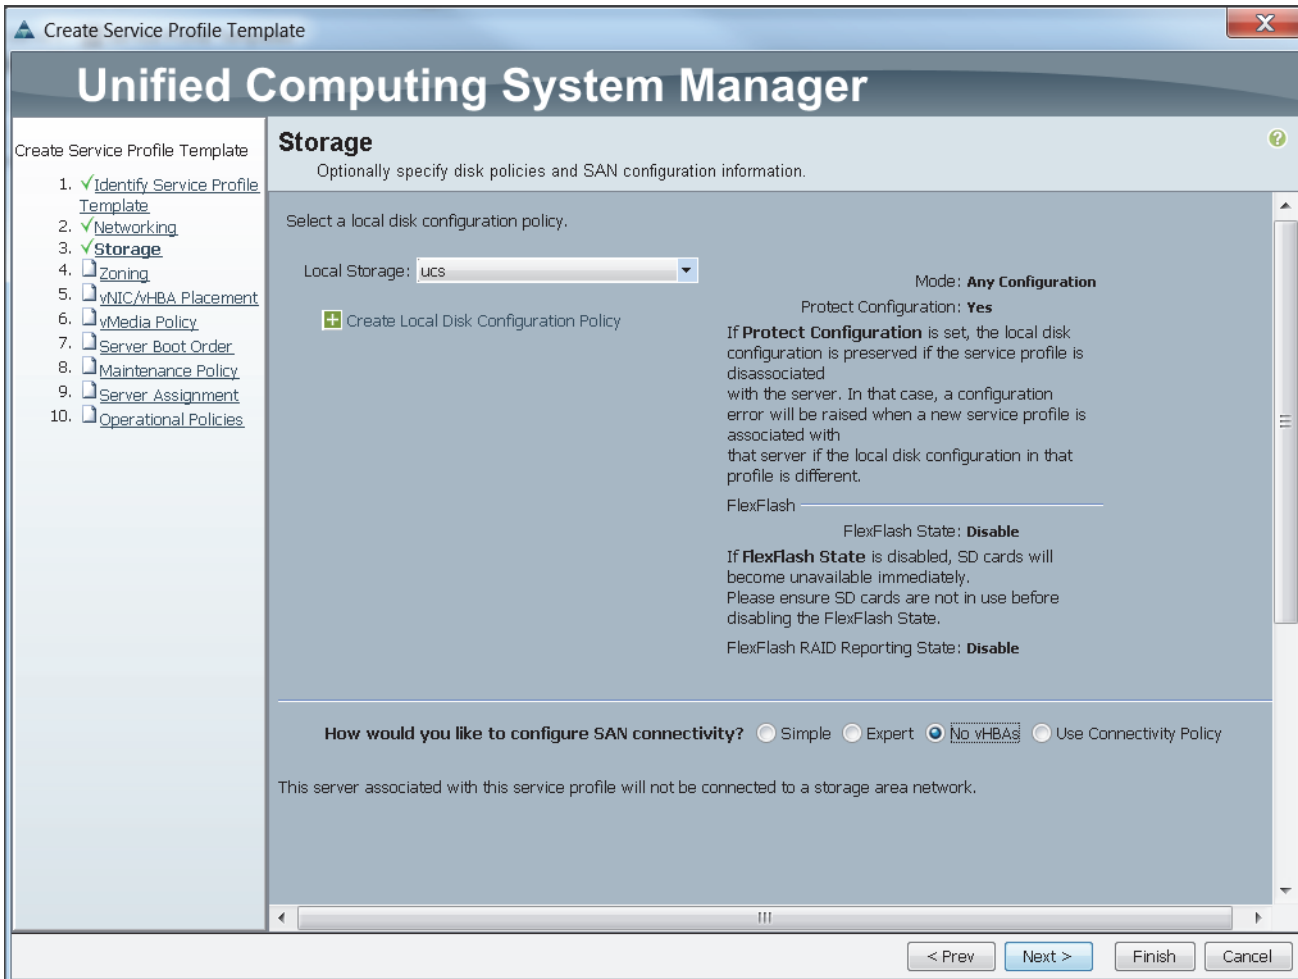
Configuring Storage Policy for the Template

Follow these steps to configure storage policies:

1. Select ucs for the local disk configuration policy.

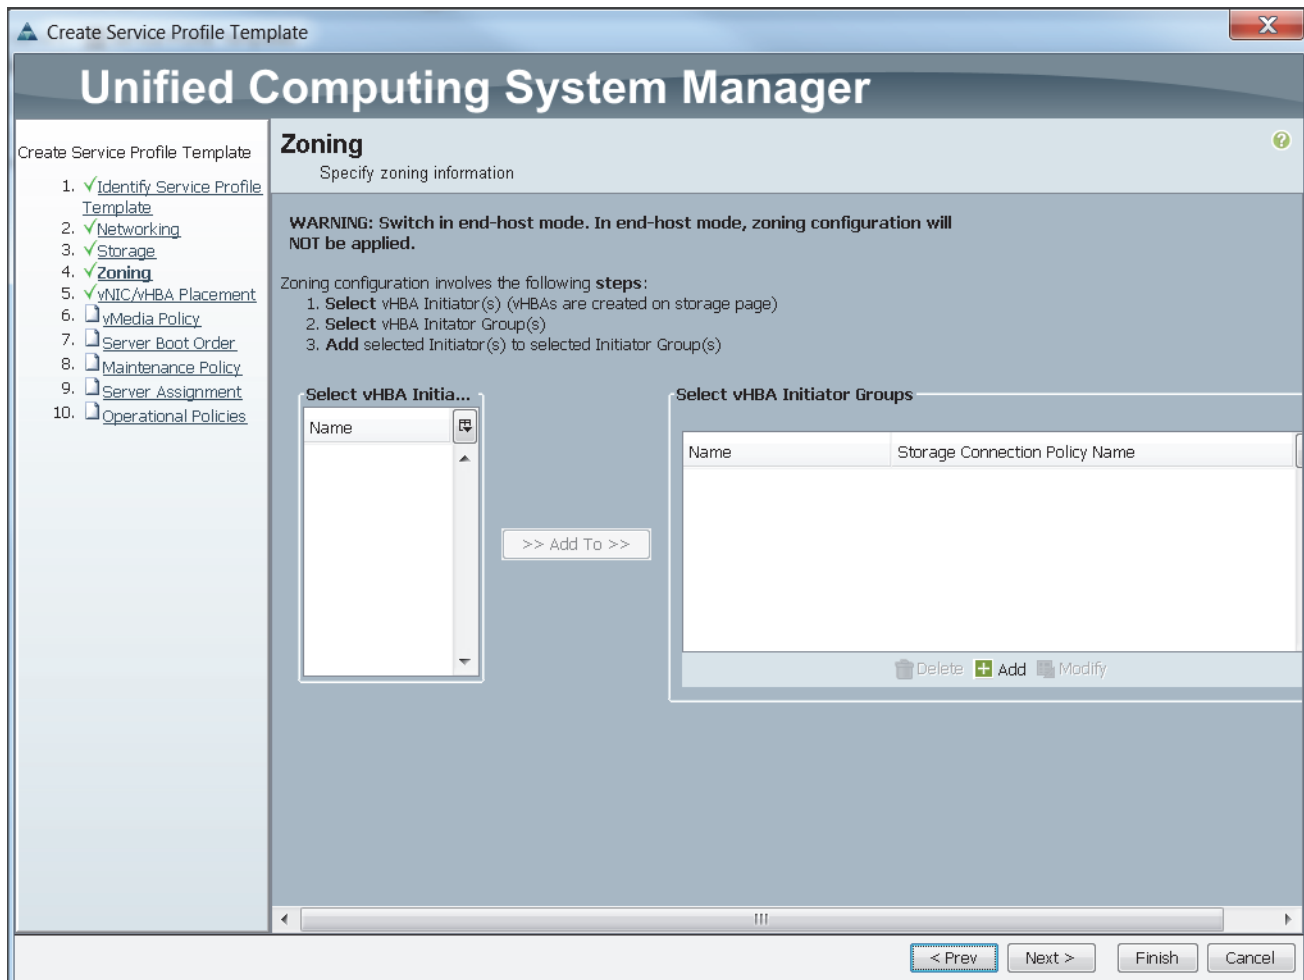
2. Click the **No vHBAs** radio button for the option, **How would you like to configure SAN connectivity?**
3. Click **Next** to continue to the next section.

Figure 53 Configuring Storage Settings



4. Click **Next** once the zoning window appears to go to the next section.

Figure 54 Configure Zoning

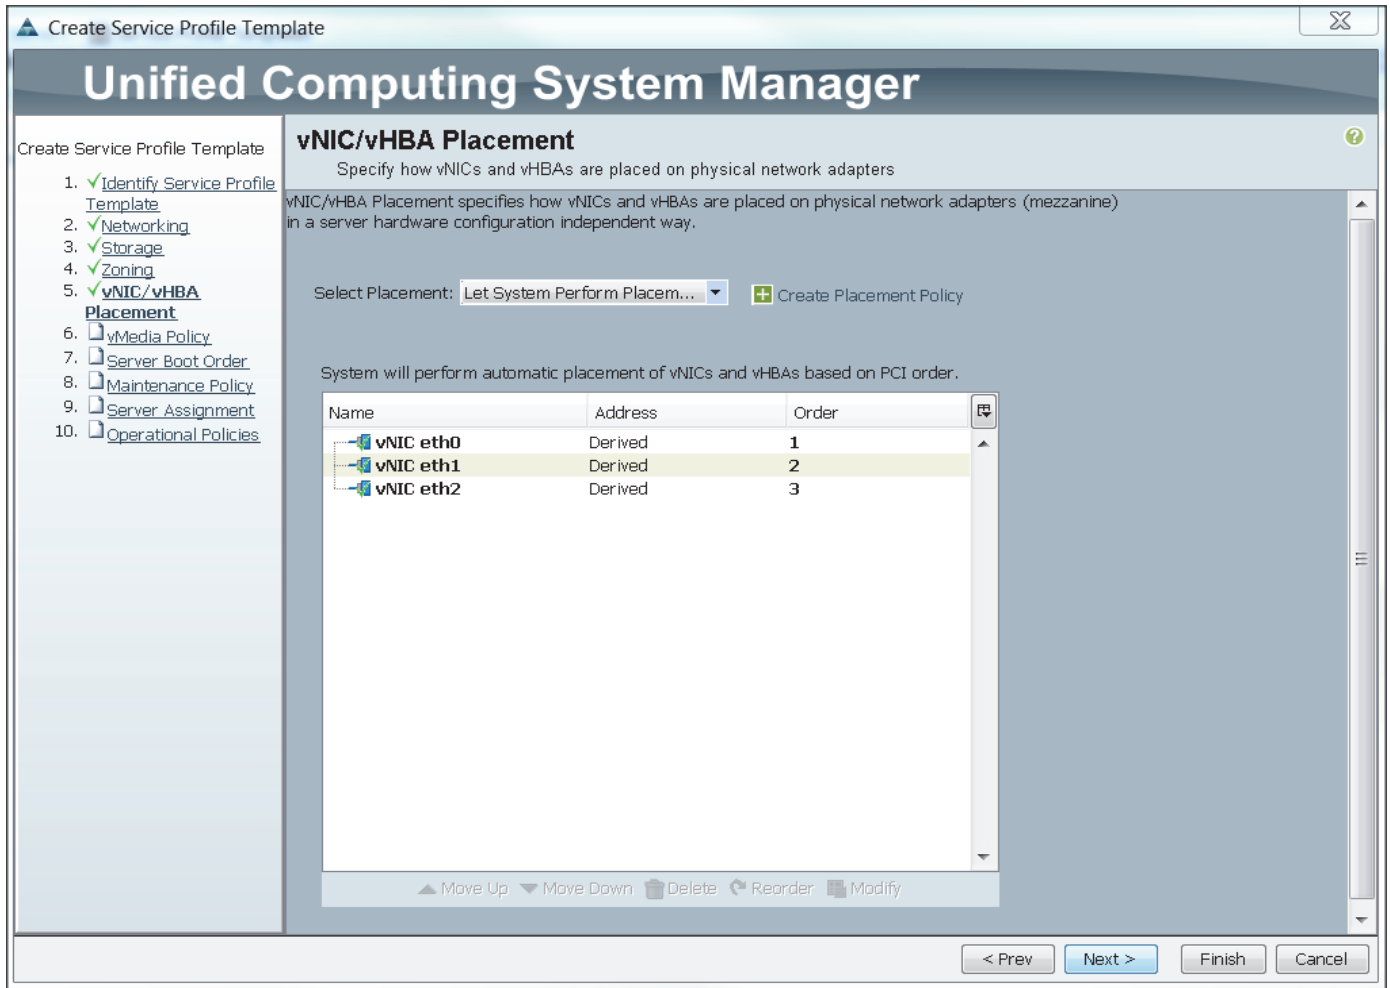


Configuring vNIC/vHBA Placement for the Template

Follow these steps to configure vNIC/vHBA placement policy:

1. Select the Default Placement Policy option for the Select Placement field.
2. Select eth0, eth1 and eth2 assign the vNICs in the following order:
 - a. eth0
 - b. eth1
 - c. eth2
3. Review to make sure that all of the vNICs were assigned in the appropriate order.
4. Click **Next** to continue to the next section.

Figure 55 vNIC/vHBA Placement



Configuring vMedia Policy for the Template

1. Click Next once the vMedia Policy window appears to go to the next section.

Figure 56 UCSM vMedia Policy Window

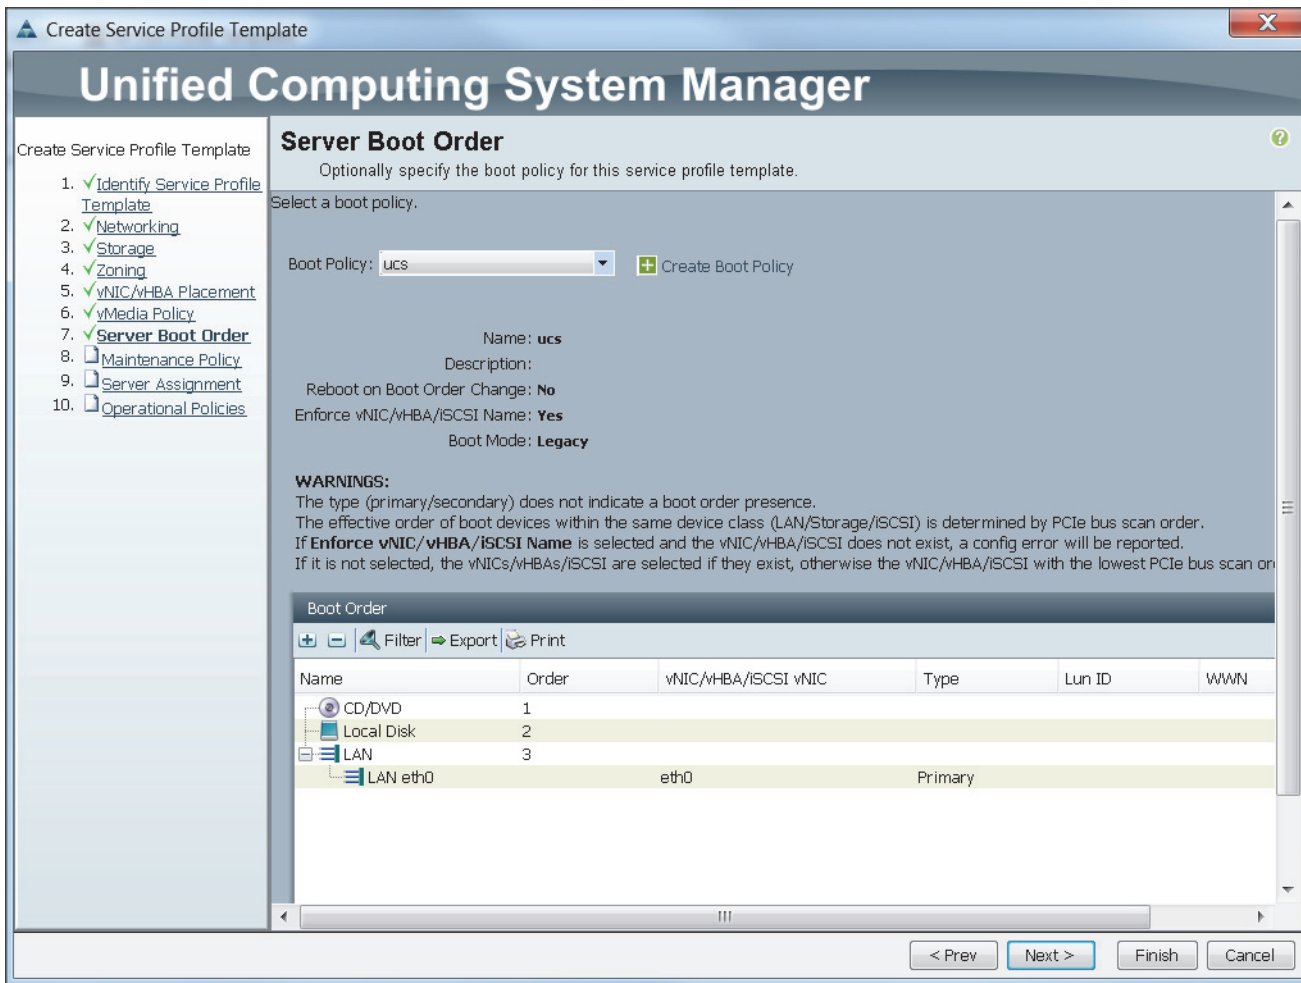


Configuring Server Boot Order for the Template

Follow these steps to set the boot order for servers:

1. Select ucs in the Boot Policy name field.
2. Review to make sure that all of the boot devices were created and identified.
3. Verify that the boot devices are in the correct boot sequence.
4. Click **OK**.
5. Click **Next** to continue to the next section.

Figure 57 Creating Boot Policy



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

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

Configuring Server Assignment for the Template

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

3. Select **ucs** for the Pool Assignment field.
4. Keep the Server Pool Qualification field at default.
5. Select **ucs** in Host Firmware Package.

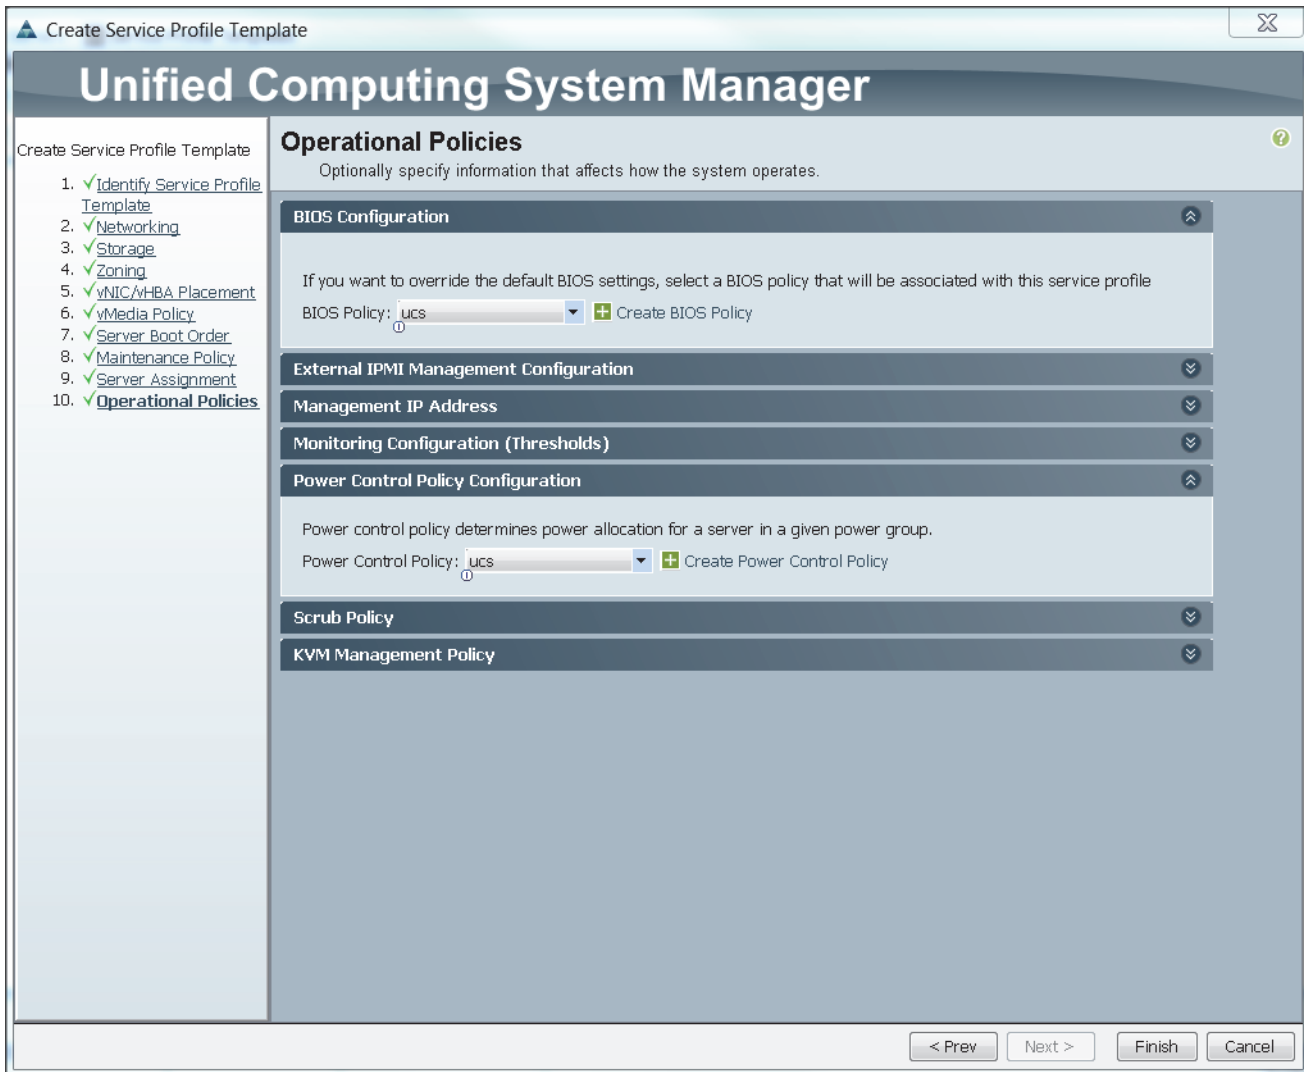
Figure 58 Server Assignment

Configuring Operational Policies for the Template

In the Operational Policies Window, follow these steps:

6. Select **ucs** in the BIOS Policy field.
7. Select **ucs** in the Power Control Policy field.
8. Click **Finish** to create the Service Profile template.
9. Click **OK** in the pop-up window to proceed.

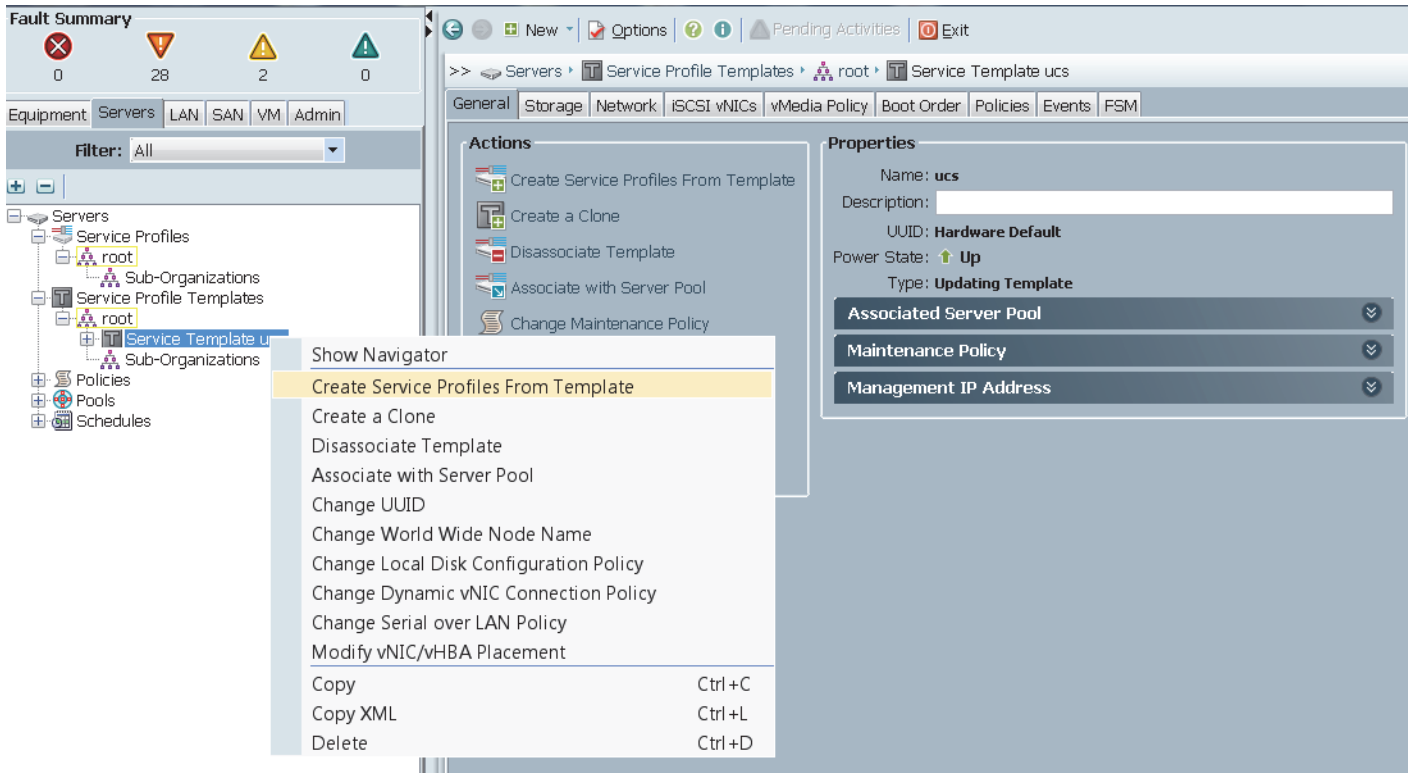
Figure 59 Selecting BIOS and Power Control Policy



Select the **Servers** tab in the left pane of the UCS Manager GUI.

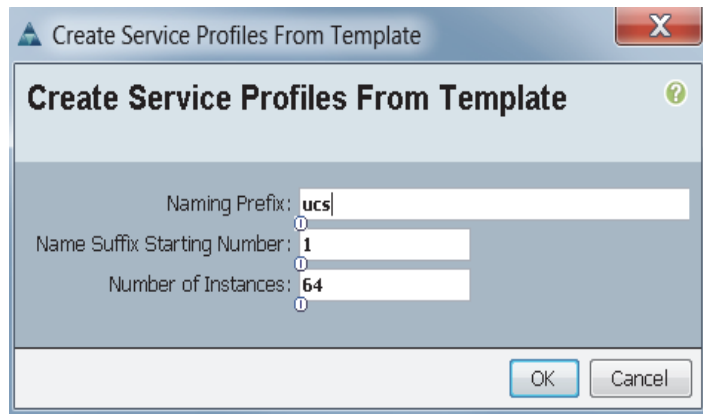
1. Go to Service Profile **Templates > root**.
2. Right-click Service Profile Templates ucs.
3. Select Create Service Profiles From Template.

Figure 60 *Creating Service Profiles from Template*



4. The Create Service Profile from Template window appears.

Figure 61 *Selecting Name and Total number of Service Profiles*



Association of the Service Profiles will take place automatically.
The Final Cisco UCS Manager window is shown in Figure 46.

Figure 62 UCS Manager showing all Nodes

Name	Overall Status	PID	Model	Serial	User Label	Cores	Memory	Adapters	NICs	HBAs	Operability	Power State	Assoc State	Profile	Fault
Server 1	Ok	UCSC-C240-...	Cisco UCS C...	FCH1852V0PU		24	262144	1	3	0	Operable	On	Associated	org-root/ls-...	N/A
Server 2	Ok	UCSC-C240-...	Cisco UCS C...	FCH1850V36U		24	262144	1	3	0	Operable	On	Associated	org-root/ls-...	N/A
Server 3	Ok	UCSC-C240-...	Cisco UCS C...	FCH1844V0QK		24	262144	1	3	0	Operable	On	Associated	org-root/ls-...	N/A
Server 4	Ok	UCSC-C240-...	Cisco UCS C...	FCH1852V0PY		24	262144	1	3	0	Operable	On	Associated	org-root/ls-...	N/A
Server 5	Ok	UCSC-C240-...	Cisco UCS C...	FCH1851V1Z2		24	262144	1	3	0	Operable	On	Associated	org-root/ls-...	N/A
Server 6	Ok	UCSC-C240-...	Cisco UCS C...	FCH1852V0L4		24	262144	1	3	0	Operable	On	Associated	org-root/ls-...	N/A
Server 7	Ok	UCSC-C240-...	Cisco UCS C...	FCH1852V0Q3		24	262144	1	3	0	Operable	On	Associated	org-root/ls-...	N/A
Server 8	Ok	UCSC-C240-...	Cisco UCS C...	FCH1851V23J		24	262144	1	3	0	Operable	On	Associated	org-root/ls-...	N/A
Server 9	Ok	UCSC-C240-...	Cisco UCS C...	FCH1852V0MF		24	262144	1	3	0	Operable	On	Associated	org-root/ls-...	N/A
Server 10	Ok	UCSC-C240-...	Cisco UCS C...	FCH1852V0PP		24	262144	1	3	0	Operable	On	Associated	org-root/ls-...	N/A
Server 11	Ok	UCSC-C240-...	Cisco UCS C...	FCH1851V213		24	262144	1	3	0	Operable	On	Associated	org-root/ls-...	N/A
Server 12	Ok	UCSC-C240-...	Cisco UCS C...	FCH1852V0QF		24	262144	1	3	0	Operable	On	Associated	org-root/ls-...	N/A
Server 13	Ok	UCSC-C240-...	Cisco UCS C...	FCH1851V243		24	262144	1	3	0	Operable	On	Associated	org-root/ls-...	N/A
Server 14	Ok	UCSC-C240-...	Cisco UCS C...	FCH1851V216		24	262144	1	3	0	Operable	On	Associated	org-root/ls-...	N/A
Server 15	Ok	UCSC-C240-...	Cisco UCS C...	FCH1852V0NA		24	262144	1	3	0	Operable	On	Associated	org-root/ls-...	N/A
Server 16	Ok	UCSC-C240-...	Cisco UCS C...	FCH1852V0NA		24	262144	1	3	0	Operable	On	Associated	org-root/ls-...	N/A

Configuring CIMC Access Using the CIMC Configuration Utility on C3160

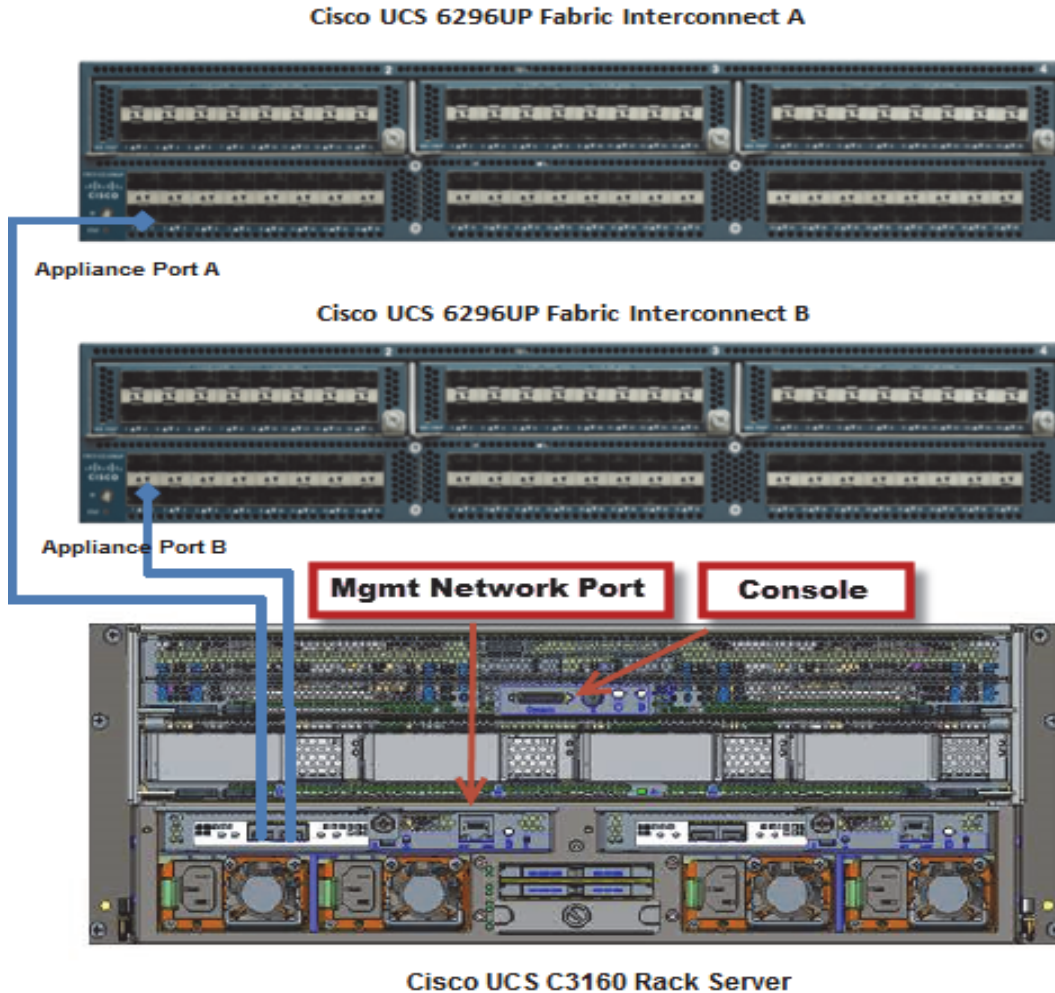
Introduction

The following sections provide an overview of how to setup the CIMC network connectivity for UCS C3160.

1. Cisco C-Series Rack Servers provides a physical local video and two USB connections for a USB keyboard, video monitor, and USB mouse connection through the front and back panel of the rack server using the Cisco provided dongle connector.
2. All rack servers can have up to 4 active KVM over IP sessions in addition to the local connection at front or rear panels. all active sessions have full control of the console.
3. KVM over IP supports text and graphics modes of the graphics controller and needs no manual setting to view data.

Cable Connectivity

Figure 63 Cable connections for C3160 servers:



Connection for C3160 Server:

1. Connect video monitor, USB keyboard and USB mouse to Cisco C3160 rack servers through the back panel using the Cisco provided dongle connector.
2. Connect the network port of the C3160 Server to Management port of the Management switch.

Power up the KVM

Complete the steps below in order to power up the server:

1. Plug the power cord into the chassis.



Note

- CIMC initializes system standby (Power Off mode).



Note

- CIMC is active and can be controlled through GUI or CLI, if you know the IP address.
2. Depress Front Panel Power:
 - The Fans begin to spin up.
 - Then POST sequence begins.
 - At this point you can either boot up or begin the installation process.
 - Note that for large memory, models can display a blank screen for several minutes.

F8 to Configure and View CIMC IP

While in BIOS you can press F8 for the CIMC IP address configuration and password reset.

1. Set **NIC** mode to Dedicated.
2. Set **NIC** redundancy to None
3. Choose **IPv4** for Static configuration.
4. Enter the CIMC IP, subnet and gateway information.
5. After entering IP information, press **F5** in order to display the configured IP.

```

Cisco IMC Configuration Utility Version 2.0 Cisco Systems, Inc.
*****
NIC Properties
NIC mode                               NIC redundancy
Dedicated:      [X]                    None:           [X]
Shared LOM:     [ ]                    Active-standby: [ ]
Cisco Card:     [ ]                    Active-active:  [ ]
SIOC Slot:     1
IP (Basic)
IPV4:           [X]                    IPV6:          [ ]
DHCP enabled   [ ]
CIMC IP:       10.29.160.230
Prefix/Subnet: 255.255.255.0
Gateway:       10.29.160.1
Pref DNS Server: 0.0.0.0
VLAN (Advanced)
VLAN enabled:  [ ]
VLAN ID:      1
Priority:      0
*****
<Up/Down>Selection <F10>Save <Space>Enable/Disable <F5>Refresh <ESC>Exit
<F1>Additional settings
    
```

6. Press **F1** and enter Additional Settings (optional).


```

Cisco IMC Configuration Utility Version 2.0 Cisco Systems, Inc.
*****
Common Properties
Hostname:      C3160-FCH1834J73U
Dynamic DNS:   [X]
DDNS Domain:
FactoryDefaults
Factory Default:      [ ]
Default User(Basic)
Default password:
Reenter password:
Port Profiles
Reset:              [ ]
Name:
*****
<Up/Down>Selection  <F10>Save  <Space>Enable/Disable  <F5>Refresh  <ESC>Exit
<F2>PreviousPage

```

7. Press **F10** in order to save the configuration.
8. Press **ESC** to exit.

```


Press <F2> Setup, <F6> Boot Menu, <F7> Diagnostics, <F8>Cisco IMC Configuration,
<F12> Network Boot

Bios Version : C3160M3.2.0.2a.0.090920140606
Platform ID  : C3160M3

Cisco IMC IPv4 Address : 10.29.160.230
Cisco IMC MAC Address  : FC:5B:39:A0:0A:E4

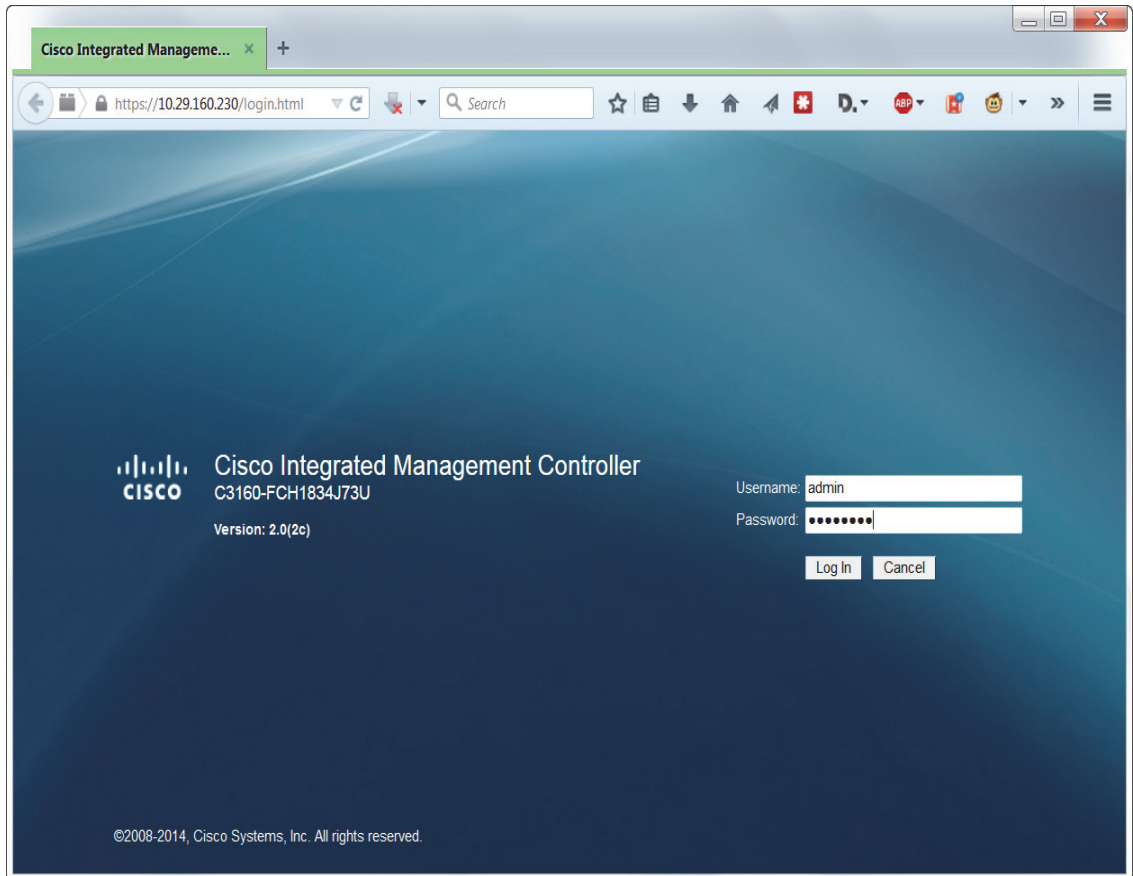
Processor(s) Intel(R) Xeon(R) CPU E5-2695 v2 @ 2.40GHz
Total Memory = 256 GB Effective Memory = 256 GB
Memory Operating Speed 1866 Mhz

```

Access CIMC

1. Then point a Web browser to the configured CIMC IP address `http://10.29.160.230`
 - Default username: admin
 - Default password: password

Figure 64 Cisco Integrated Management Window



2. Once logged in successfully. The server can be controlled using CIMC

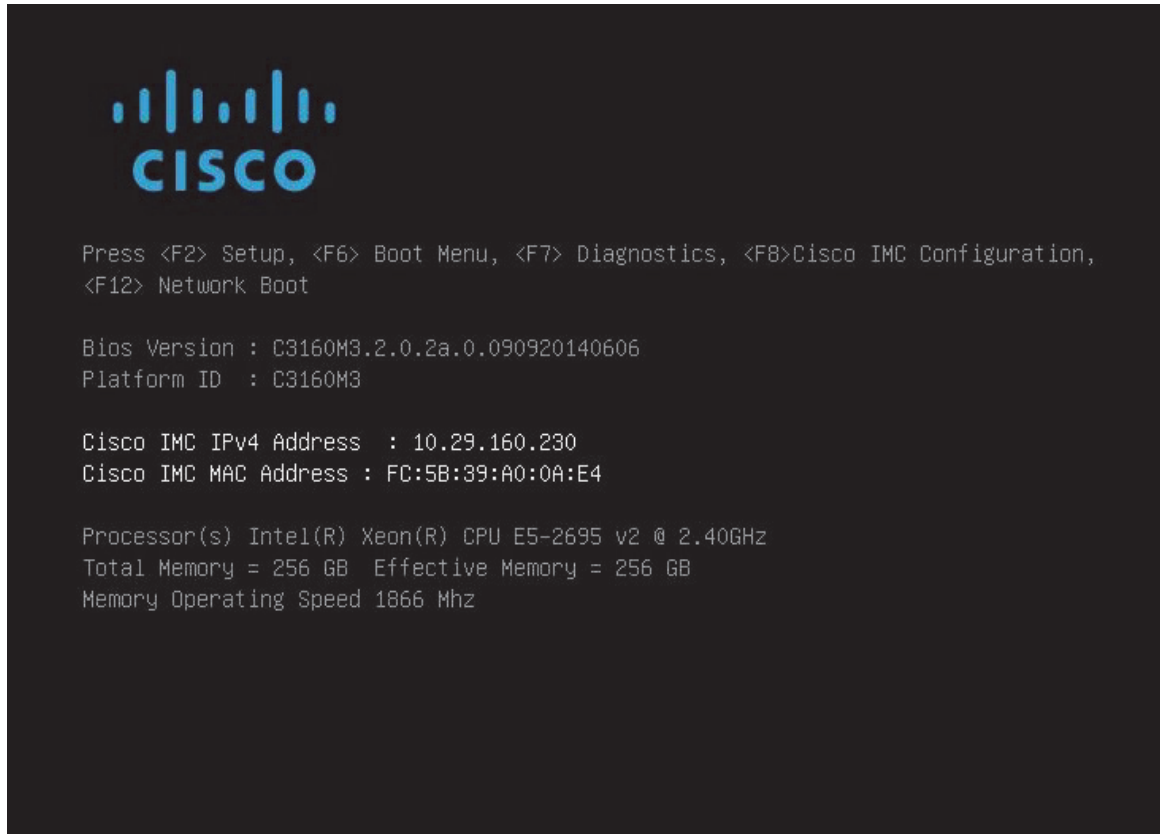
Figure 65 Cisco Integrated Management Controller

The screenshot displays the Cisco Integrated Management Controller (CIMC) web interface for a UCS C3160 server. The interface is organized into several sections:

- Overall Server Status:** Shows a green checkmark and the word "Good".
- Navigation:** Includes tabs for "Server", "Admin", and "Storage". A left sidebar lists menu items: Summary, Inventory, Sensors, Remote Presence, BIOS, Power Policies, Faults and Logs, and Troubleshooting.
- Server Summary:**
 - Actions:** A list of server management actions: Power On Server, Power Off Server, Shut Down Server, Power Cycle Server, Hard Reset Server, Launch KVM Console, Turn On Locator LED, and Turn On Front Locator LED.
 - Server Properties:**
 - Product Name: UCS C3160
 - Serial Number: FCH18457MGK
 - PID: UCSC-C3X60-SVRNB
 - UUID: 2233FC9A-F7AD-4A5F-90CD-76DD7DA85679
 - BIOS Version: C3160M3.2.0.2a.0.090920140606
 - Description: (empty field)
 - Server Status:**
 - Power State: On
 - Overall Server Status: Good
 - Temperature: Good
 - Overall DIMM Status: Good
 - Power Supplies: Good
 - Fans: Good
 - Locator LED: Off
 - Front Locator LED: Off
 - Overall Storage Status: Good
 - Cisco Integrated Management Controller (Cisco IMC) Information:**
 - Hostname: C3160-FCH18457MGK
 - IP Address: 10.29.160.238
 - MAC Address: 74:A0:2F:42:36:24
 - Firmware Version: 2.0(2c)
 - Current Time (UTC): Wed Mar 18 21:46:49 2015
 - Local Time: Wed Mar 18 21:46:49 2015 UTC
 - Timezone: UTC (Select Timezone)

3. Click **Launch KVM Console**.

- Restart the server by using KVM Console, **Macros > Static Macros > Ctrl-Alt-Del.**



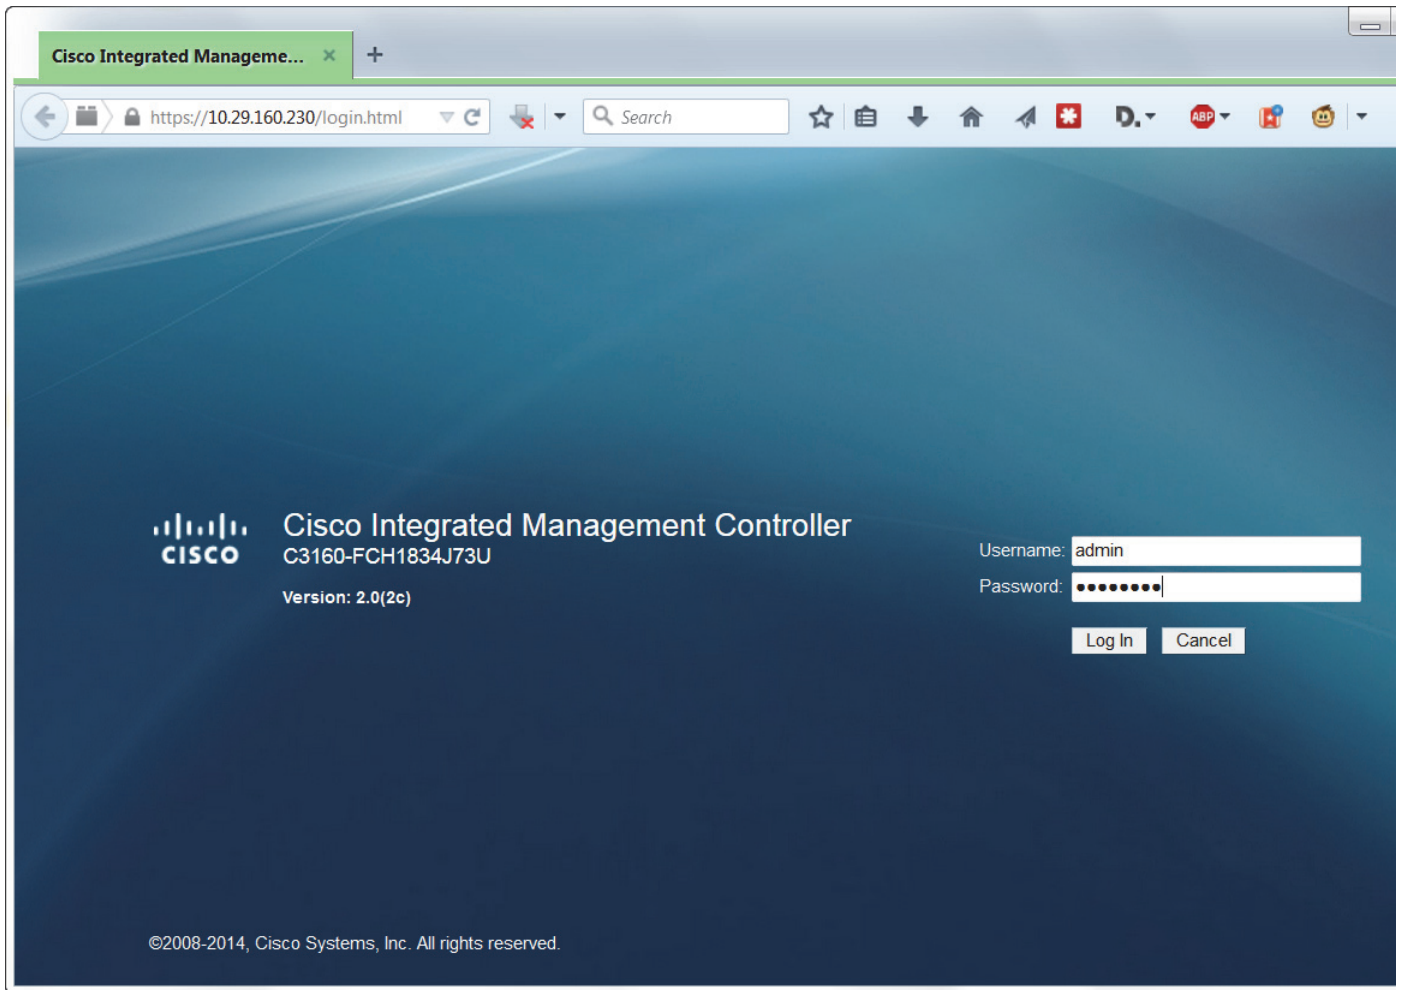
Installing Redhat Enterprise Linux 6.5 software Raid (OS based Mirroring) on C3160 System using CIMC

The following section provides detailed procedures for installing Red Hat Linux 6.5.

Access CIMC

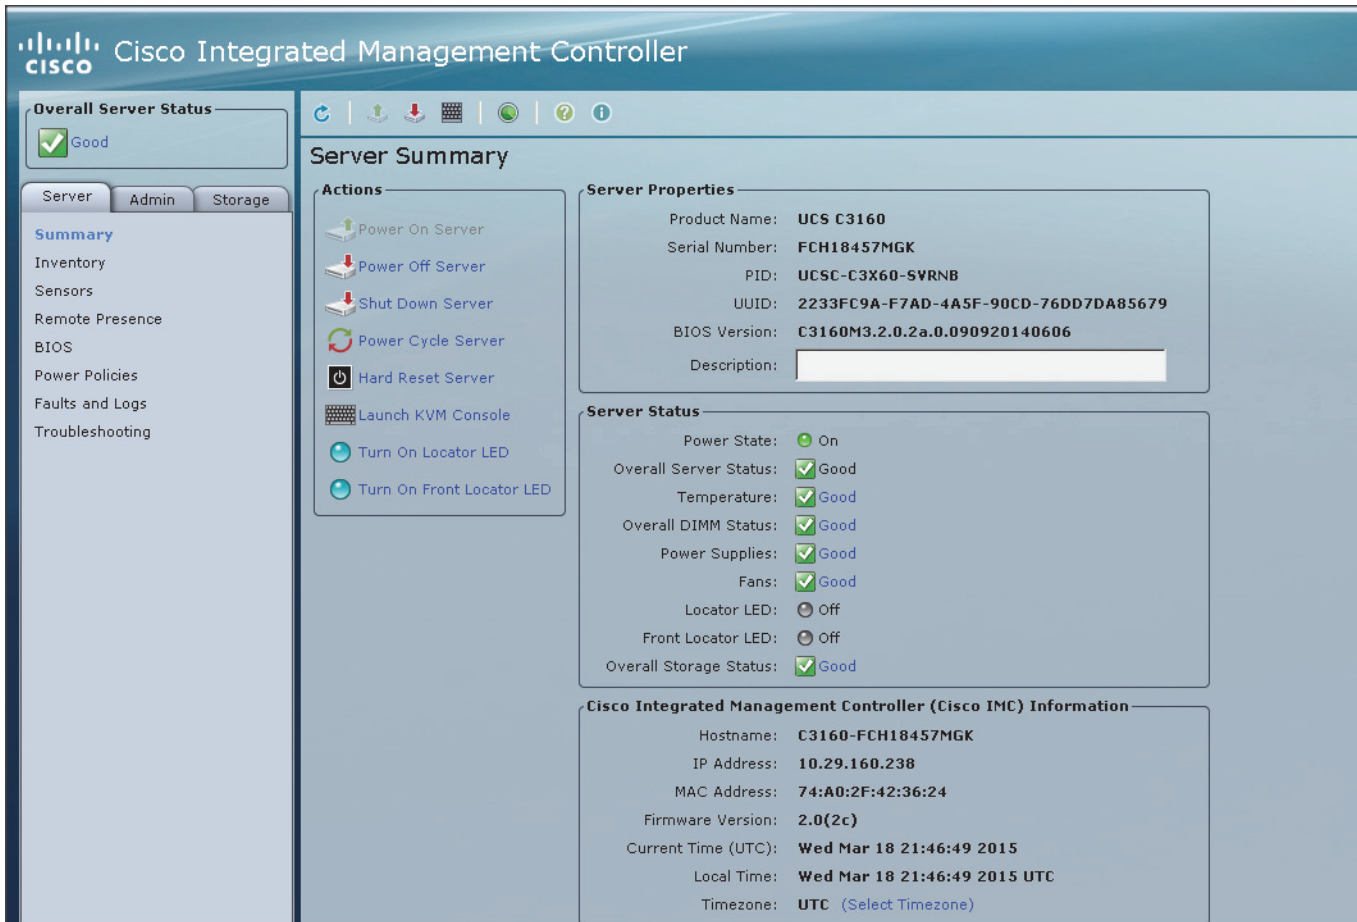
- Then point a Web browser to the configured CIMC IP address <http://10.29.160.230>
 - Default username: admin
 - Default password: password

Figure 66 CIMC Log in Page



2. Once logged in successfully. The server can be controlled using CIMC.

Figure 67 CIMC: Sever Summary Page




3. Click **Launch KVM Console**.


Restart the server by using KVM Console, **Macros > Static Macros > Ctrl-Alt-Del**

1. In the KVM window, select the **Virtual Media** tab.

2. Click the **Activate Virtual Devices** found under **Virtual Media** tab.



3. In the KVM window, select the Virtual Media tab and Click the **Map CD/DVD**.



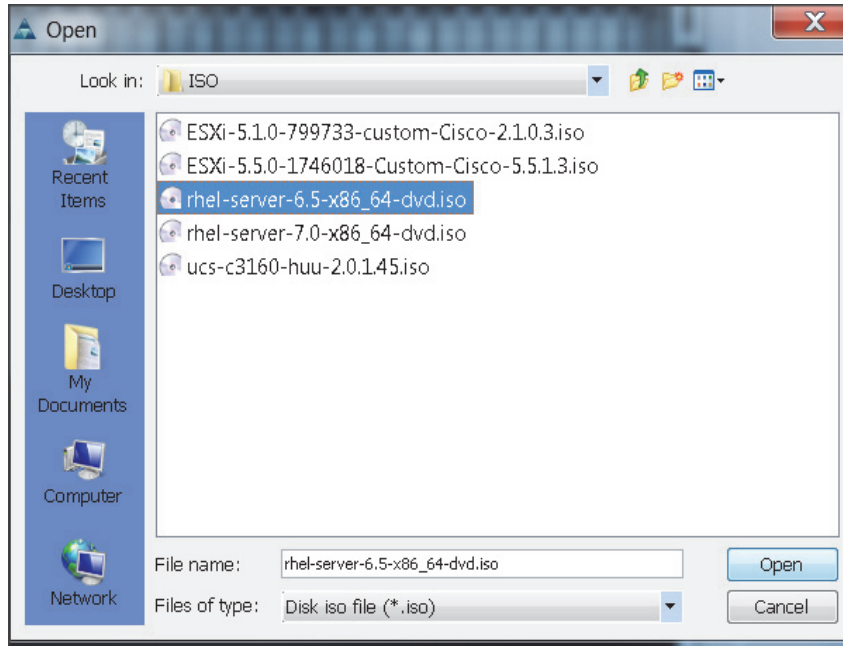
4. Browse to the Red Hat Enterprise Linux Server 6.5 installer ISO image file.

**Note**

The Red Hat Enterprise Linux 6.5 DVD is assumed to be on the client machine.

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

Figure 68 Browse to Red Hat Enterprise Linux ISO Image



6. In the KVM window, select the **KVM** tab to monitor during boot.
7. In the KVM window, select the **Macros > Static Macros > Ctrl-Alt-Del** button in the upper left corner.
8. Click **OK**.
9. Click **OK** to reboot the system.
10. On reboot, the machine detects the presence of the Red Hat Enterprise Linux Server 6.5 install media.
11. Select the **Install or upgrade an existing system**.

Figure 69 Red Hat Enterprise Linux Server 6.5 Install Media



12. Skip the Media test and start the installation

Figure 70 RHEL: Media Test and Start of Installation



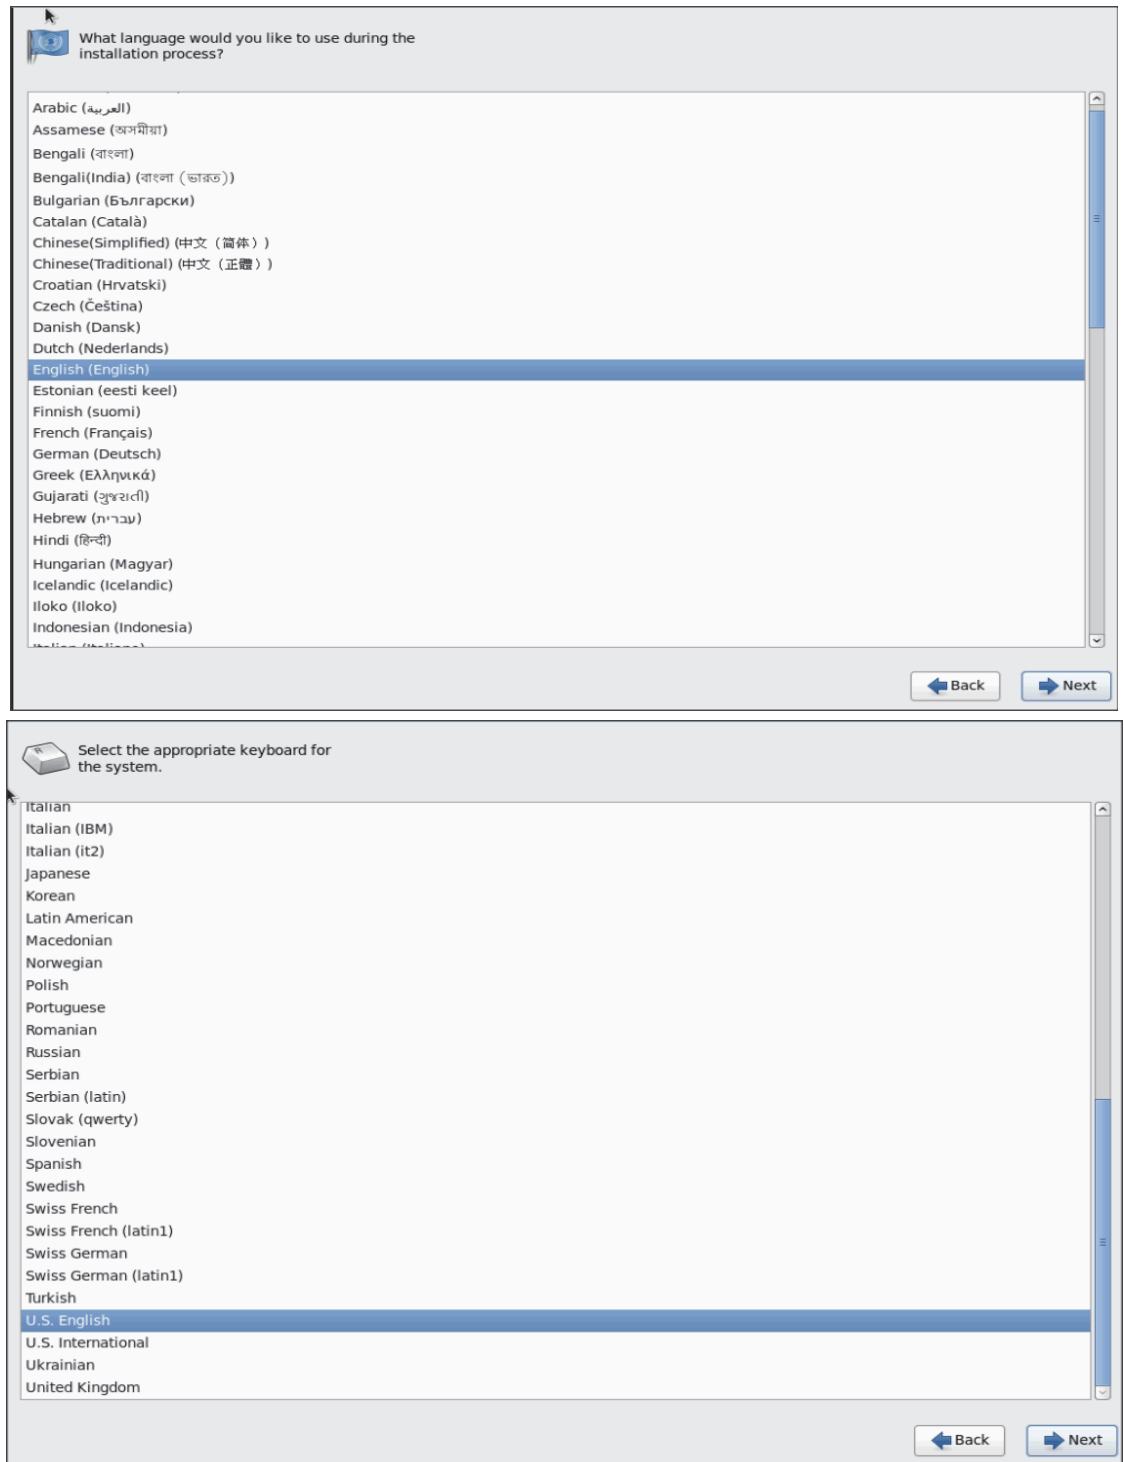
13. Click Next

Figure 71 *Red Hat Enterprise Linux Server 6.5 Install Media*



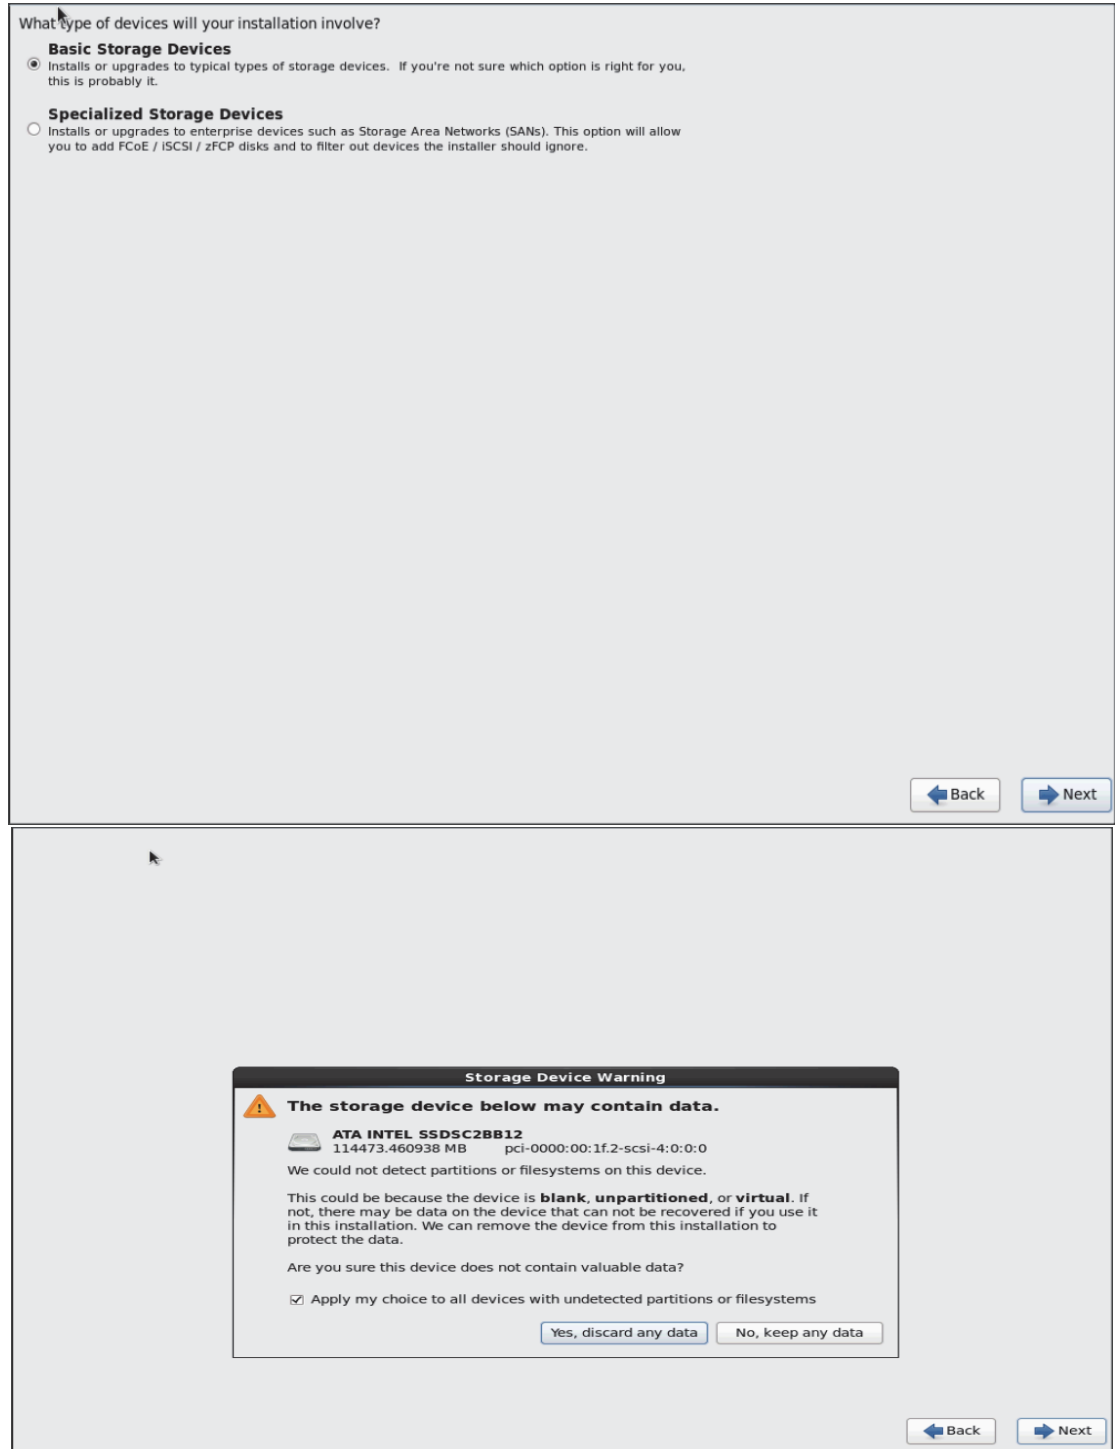
14. Select language of installation, and then Click **Next**

Figure 72 RHEL Installation: Language and Keyboard Selection



15. Select Basic Storage Devices and Click Next.

Figure 73 RHEL Installation: Storage Devices Selection



16. Provide hostname and configure Network for the host.

Figure 74 RHEL Installation: Specify Hostname

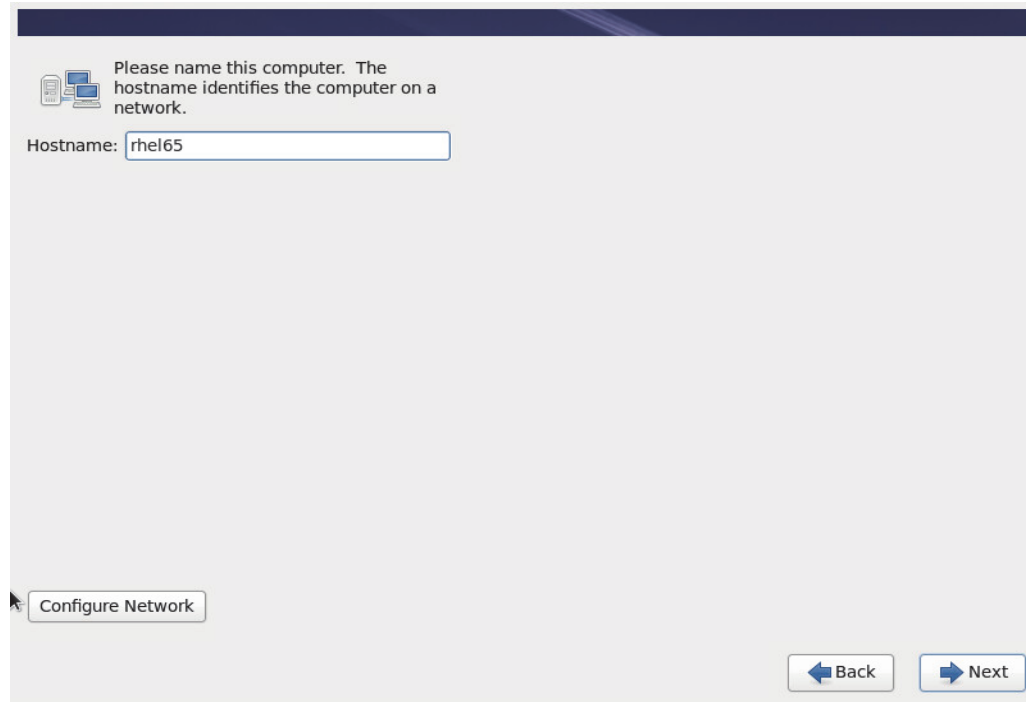


Figure 75 RHEL Installation: IPv4 Settings for eth0

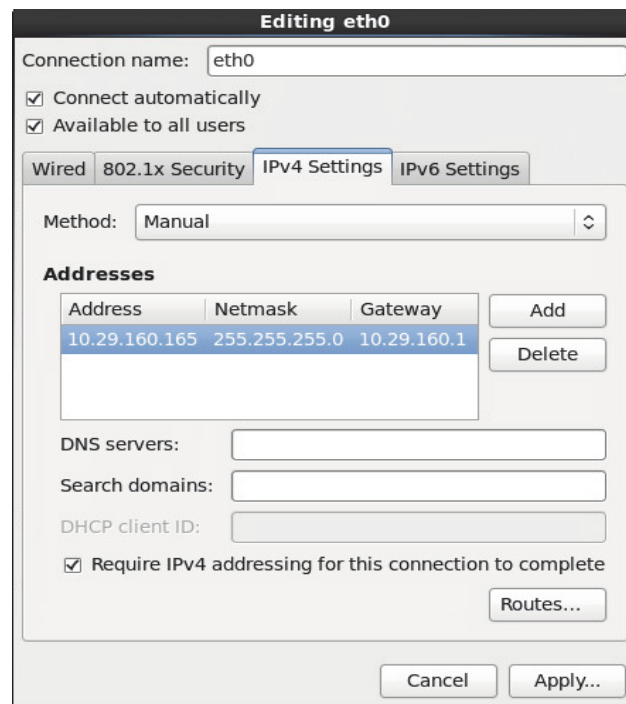


Figure 76 RHEL Installation: IPv4 Settings for eth1

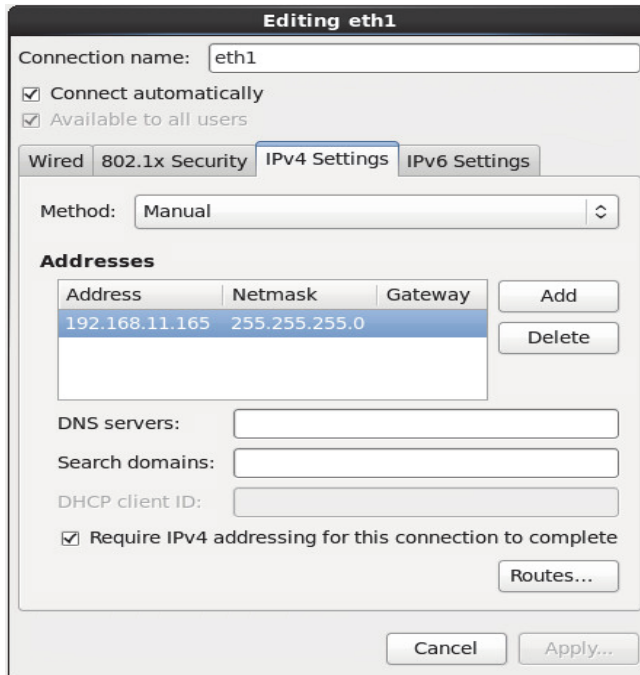


Figure 77 RHEL Installation: Location Selection

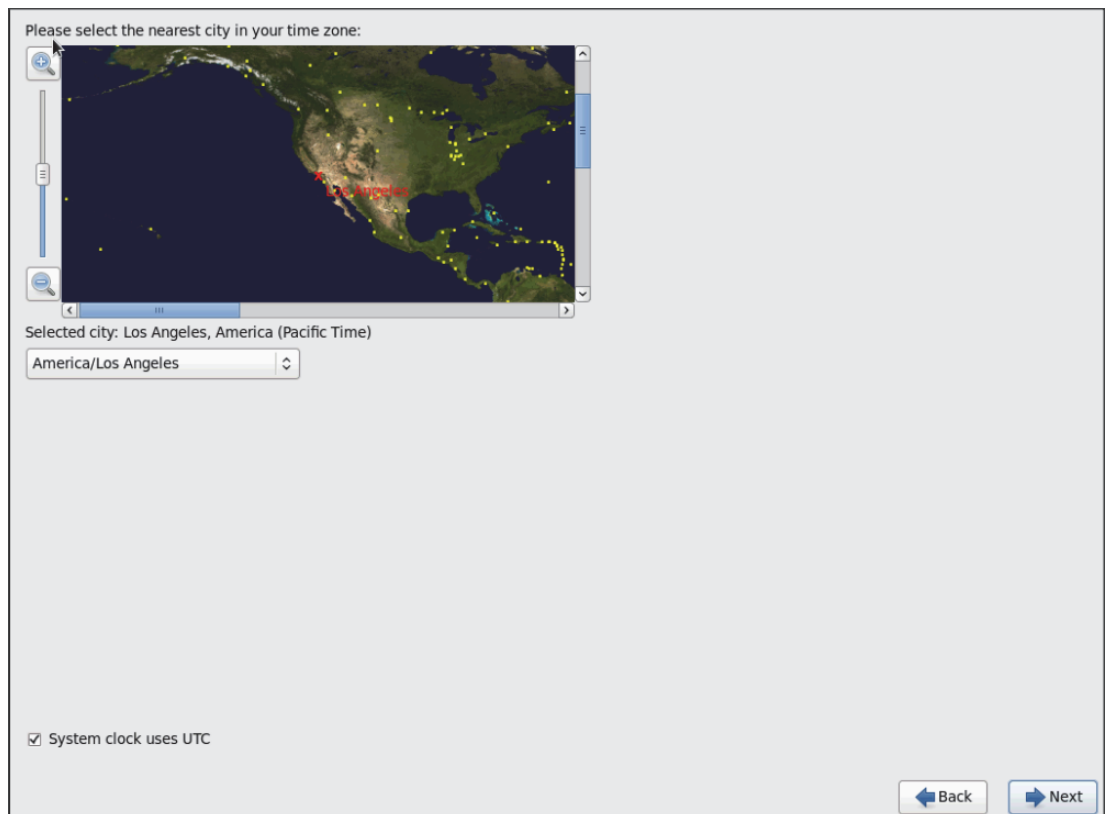
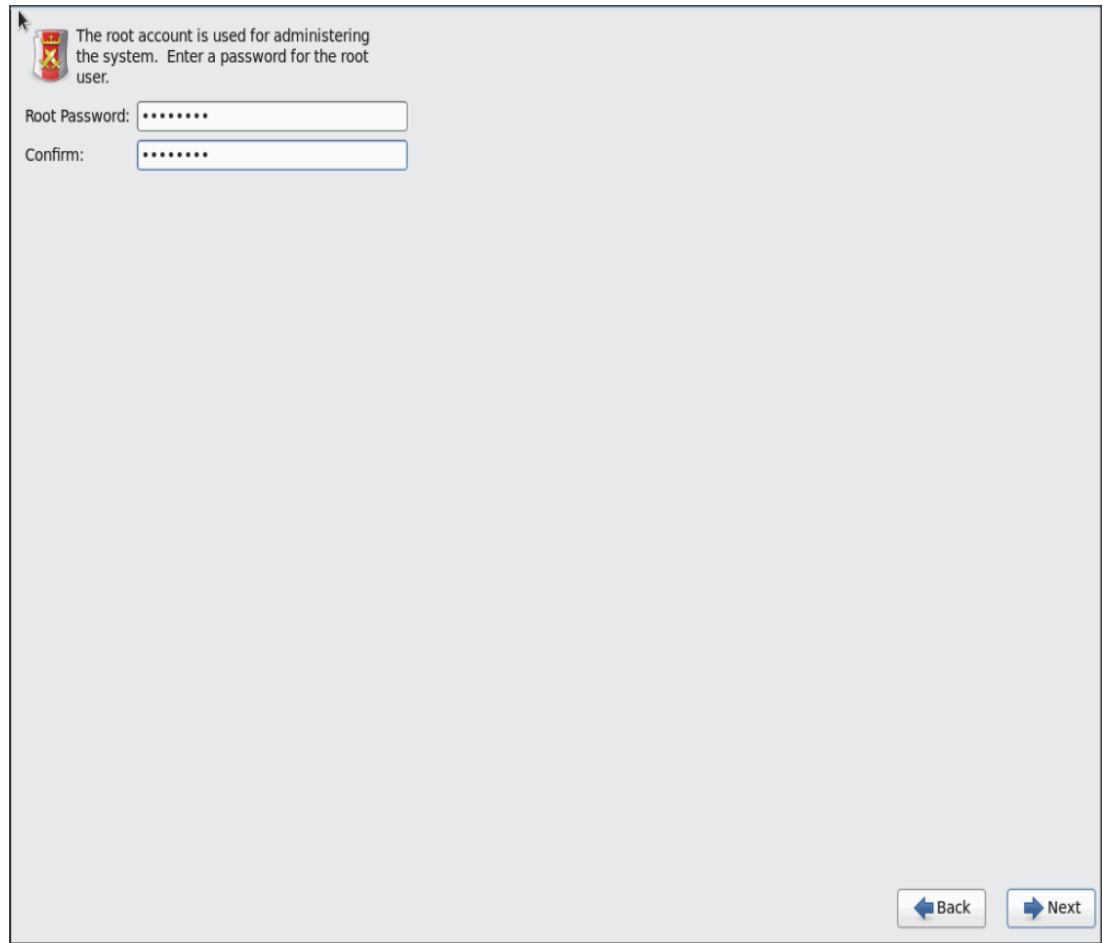


Figure 78 *RHEL Installation: Enter Root Credentials*



17. Choose **Create Custom Layout** for Installation type.

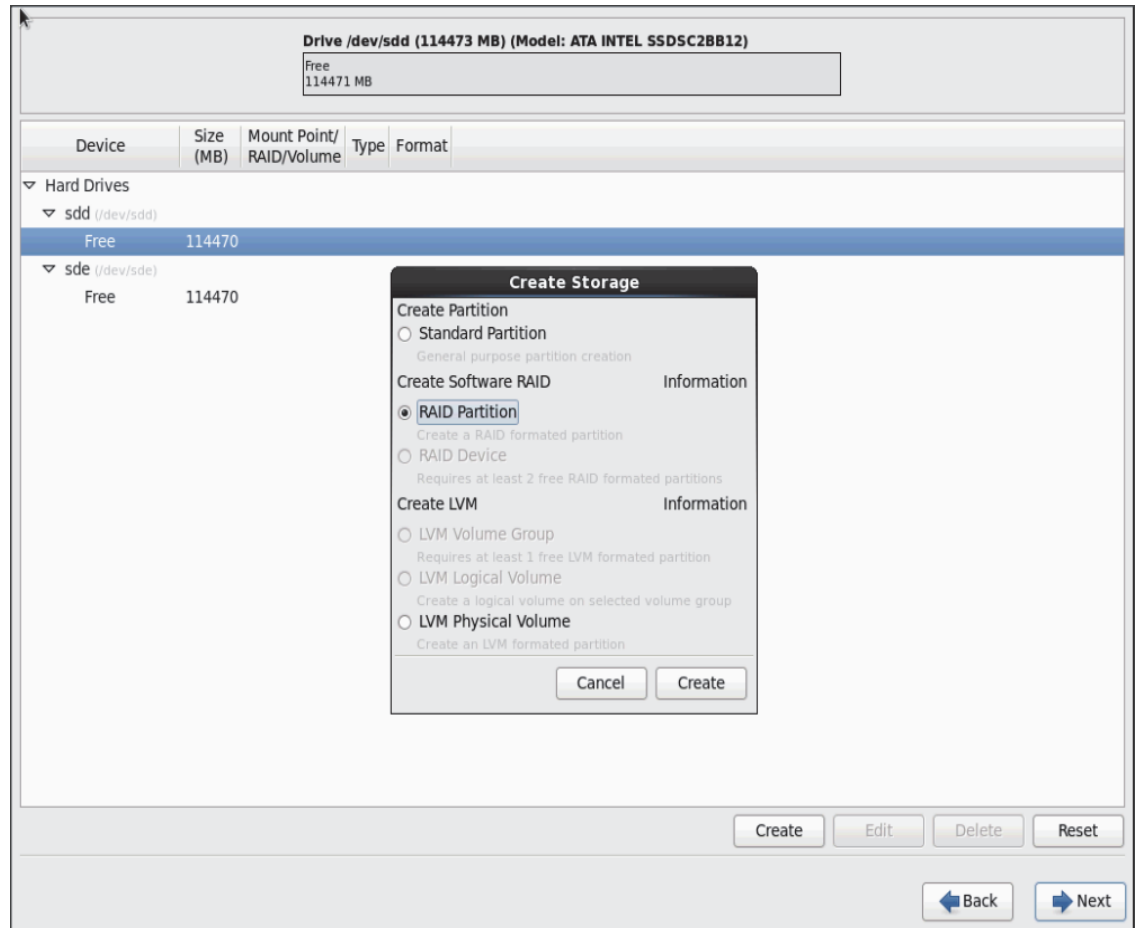
Figure 79 RHEL Installation: Custom Layout Creation



18. Following steps can be used to create two software RAID 1 partitions for boot and, or (root) partitions.

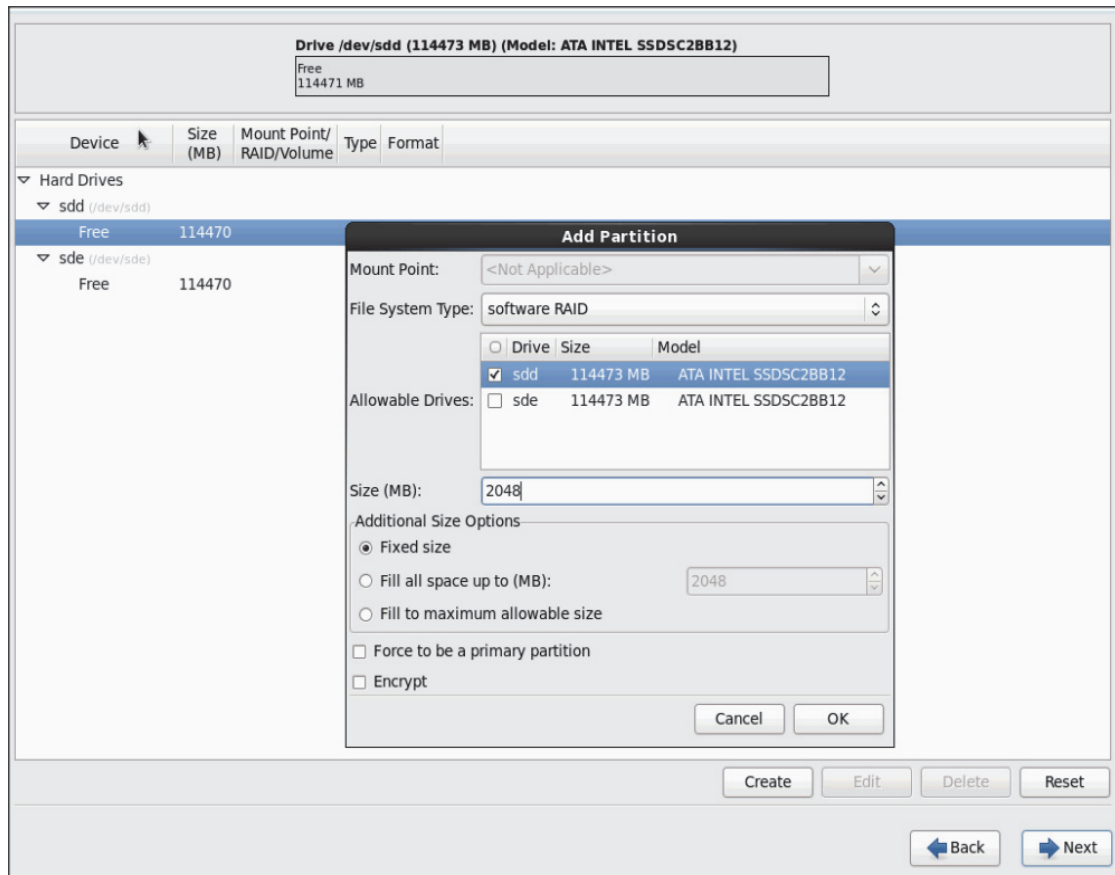
- a. Choose free volume and click on **Create** and choose **RAID Partition**.

Figure 80 *RHEL Installation: Create RAID Partition*



- b. Choose “Software RAID” for File system Type and set size for Boot volume

Figure 81 RHEL Installation: Add RAID Partition



19. Similarly, do the RAID configuration for the other free volume.

Figure 82 RHEL Installation: Create RAID Partition

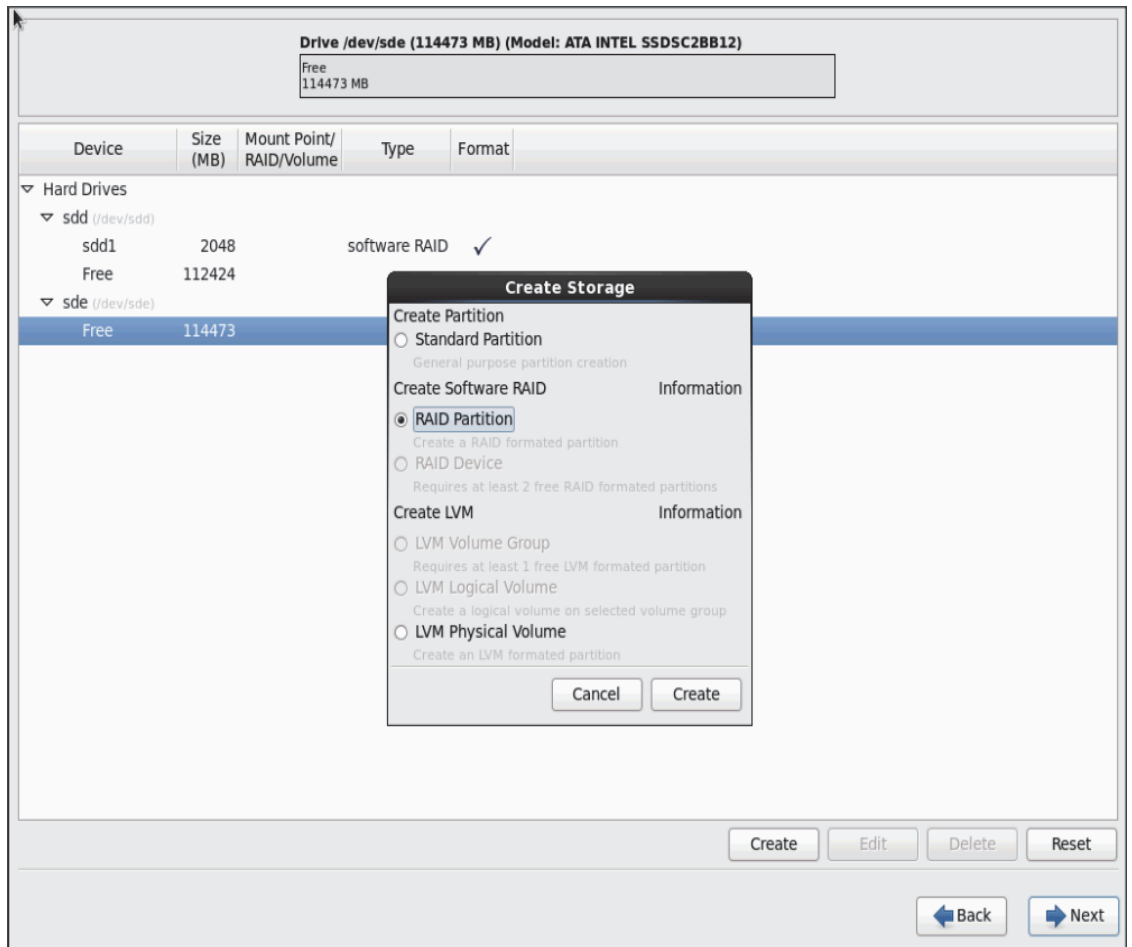
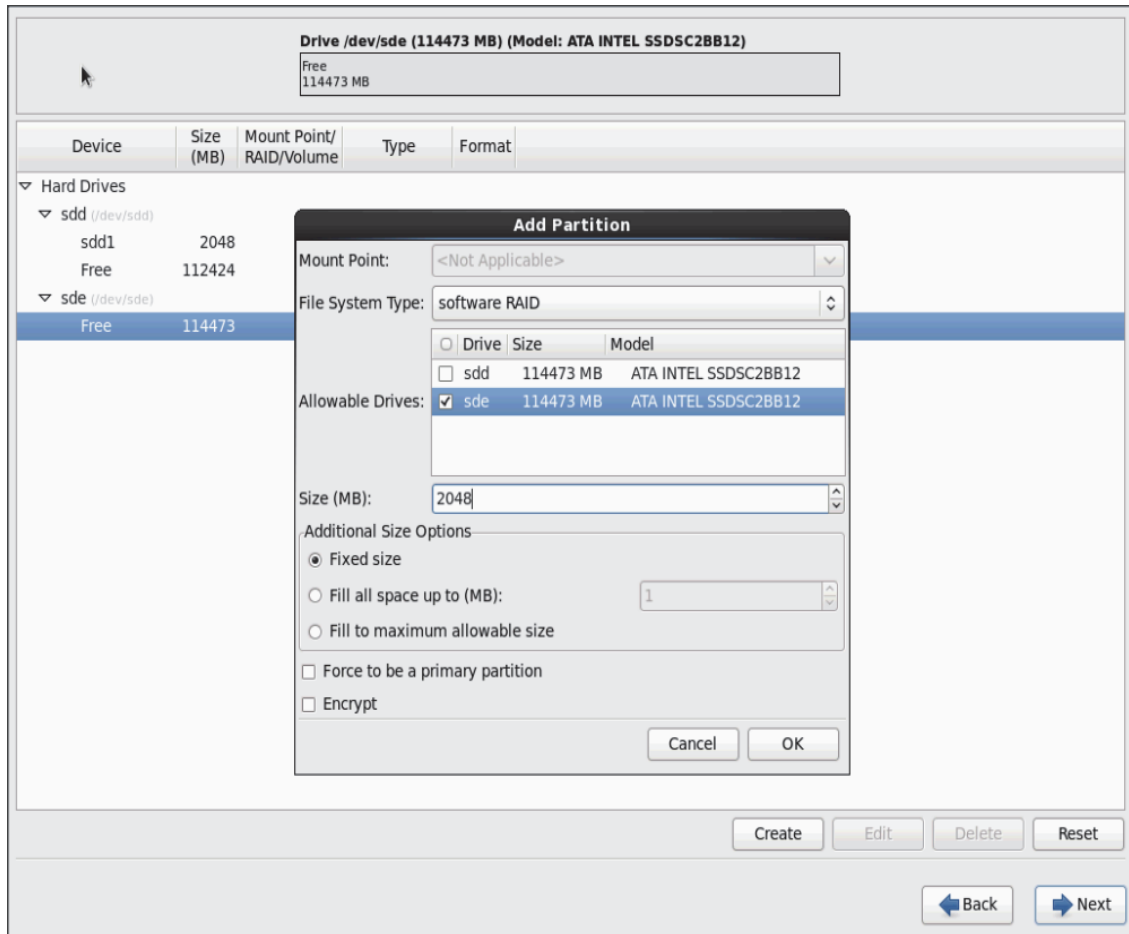


Figure 83 RHEL Installation: Add RAID Partition



- Now similarly create RAID partitions for root (/) partition on both the devices and use rest of the available space

Figure 84 RHEL Installation: Create RAID Partition

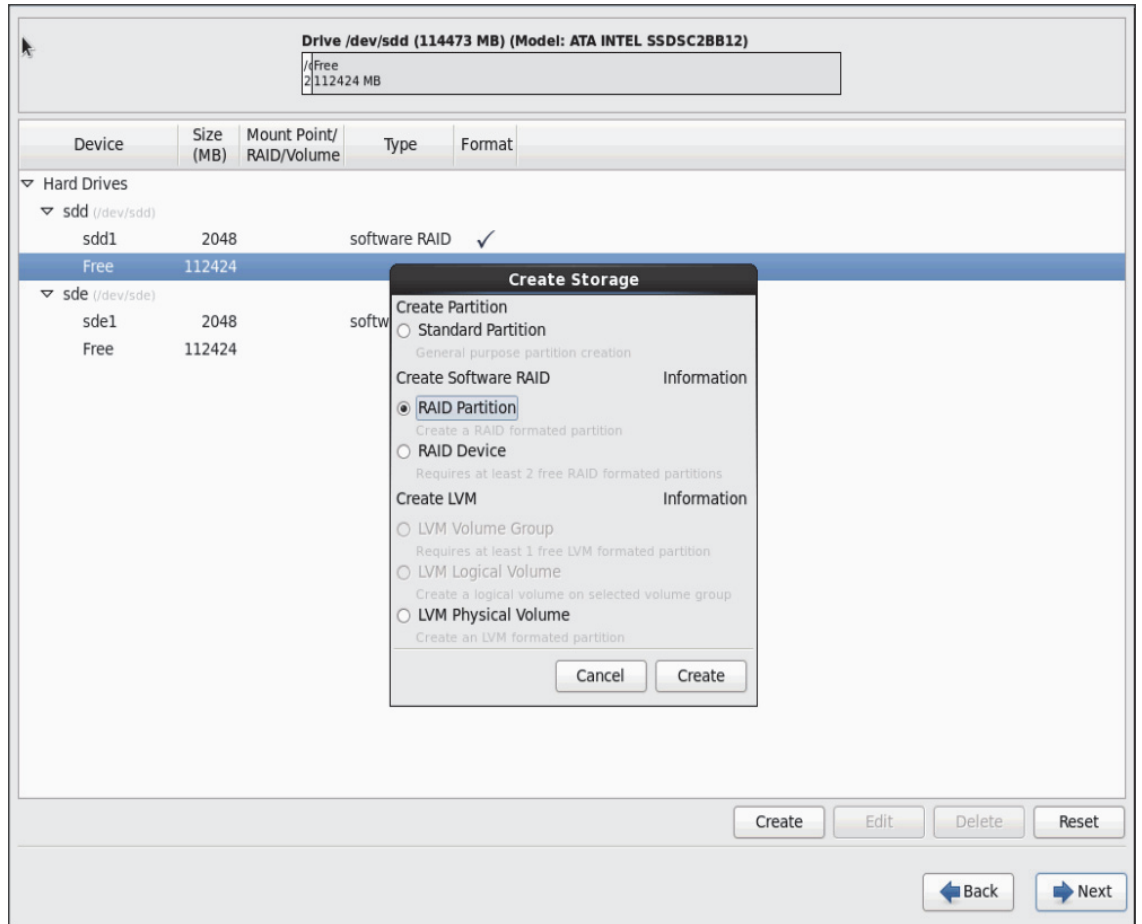


Figure 85 RHEL Installation: Add RAID Partition

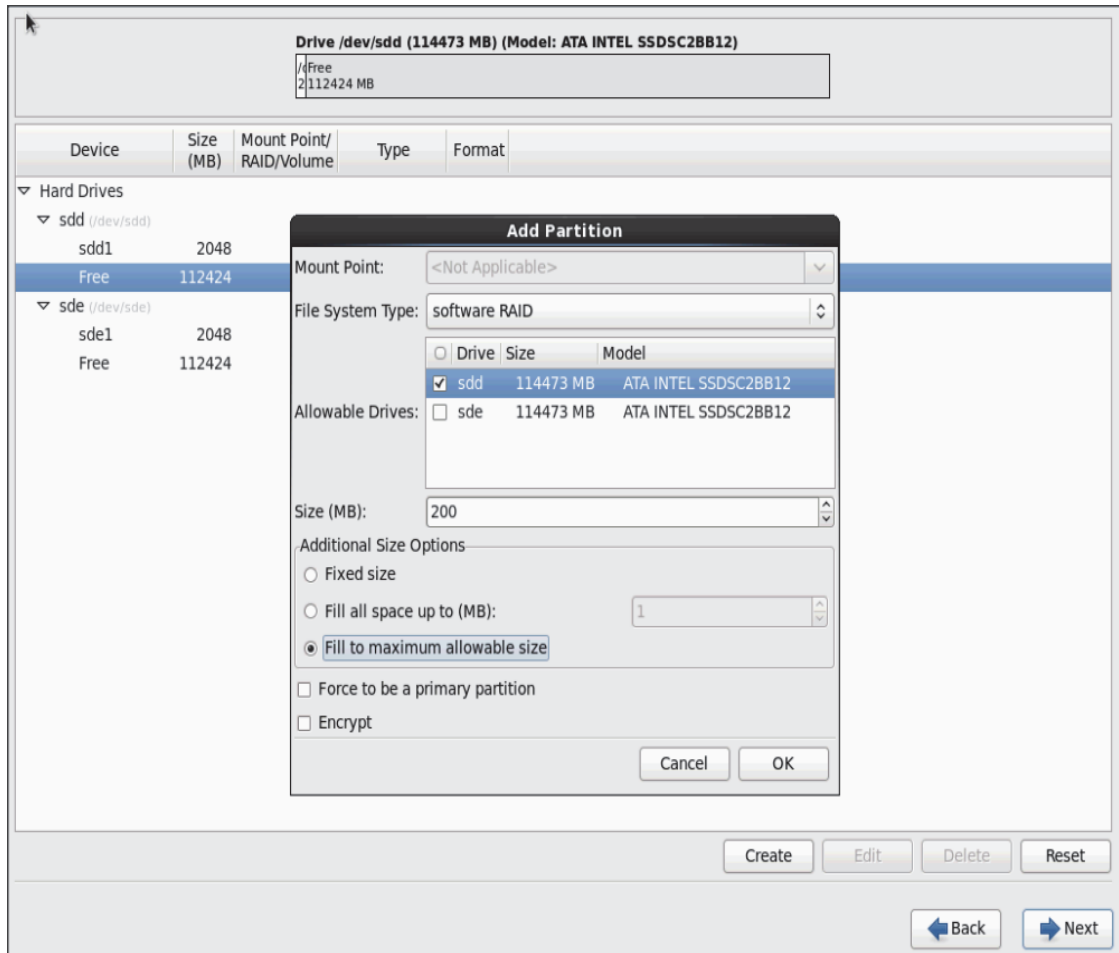


Figure 86 RHEL Installation: Create RAID Partition

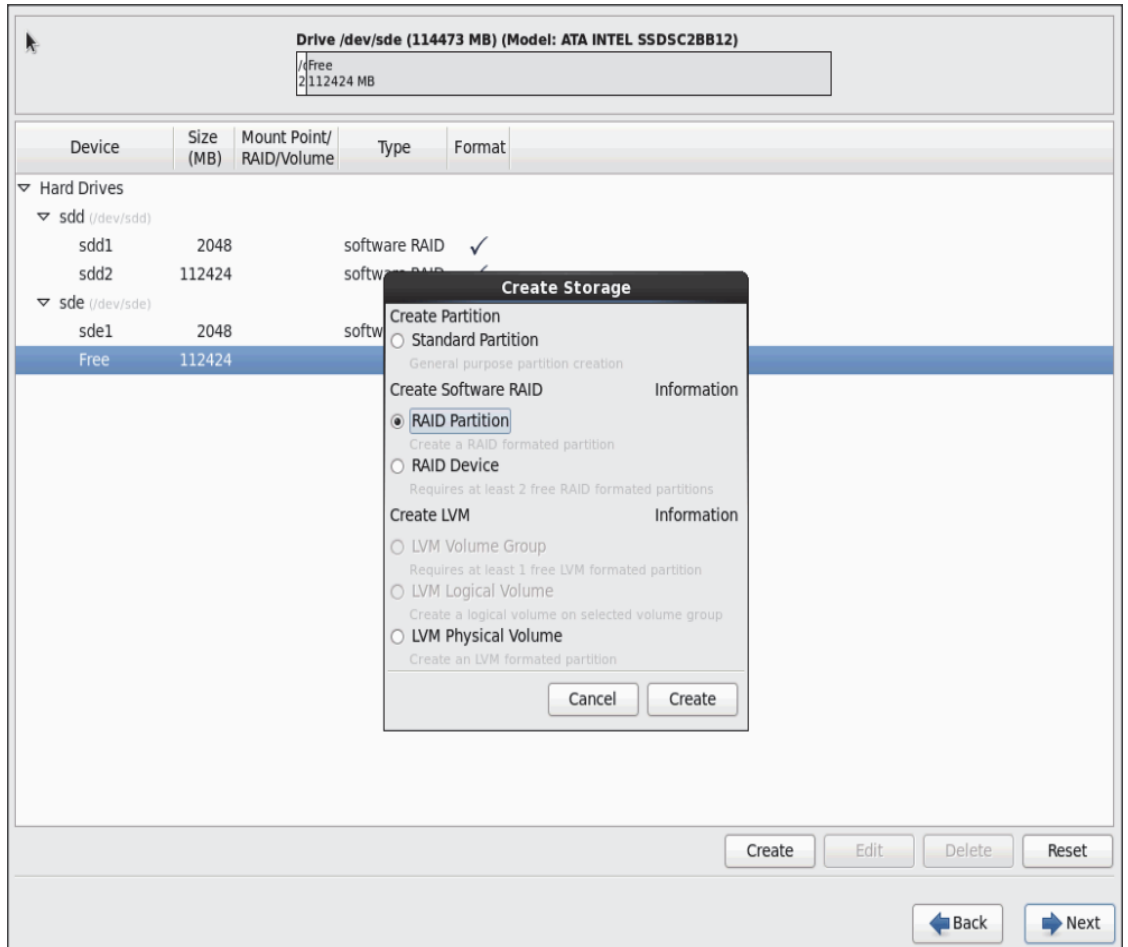
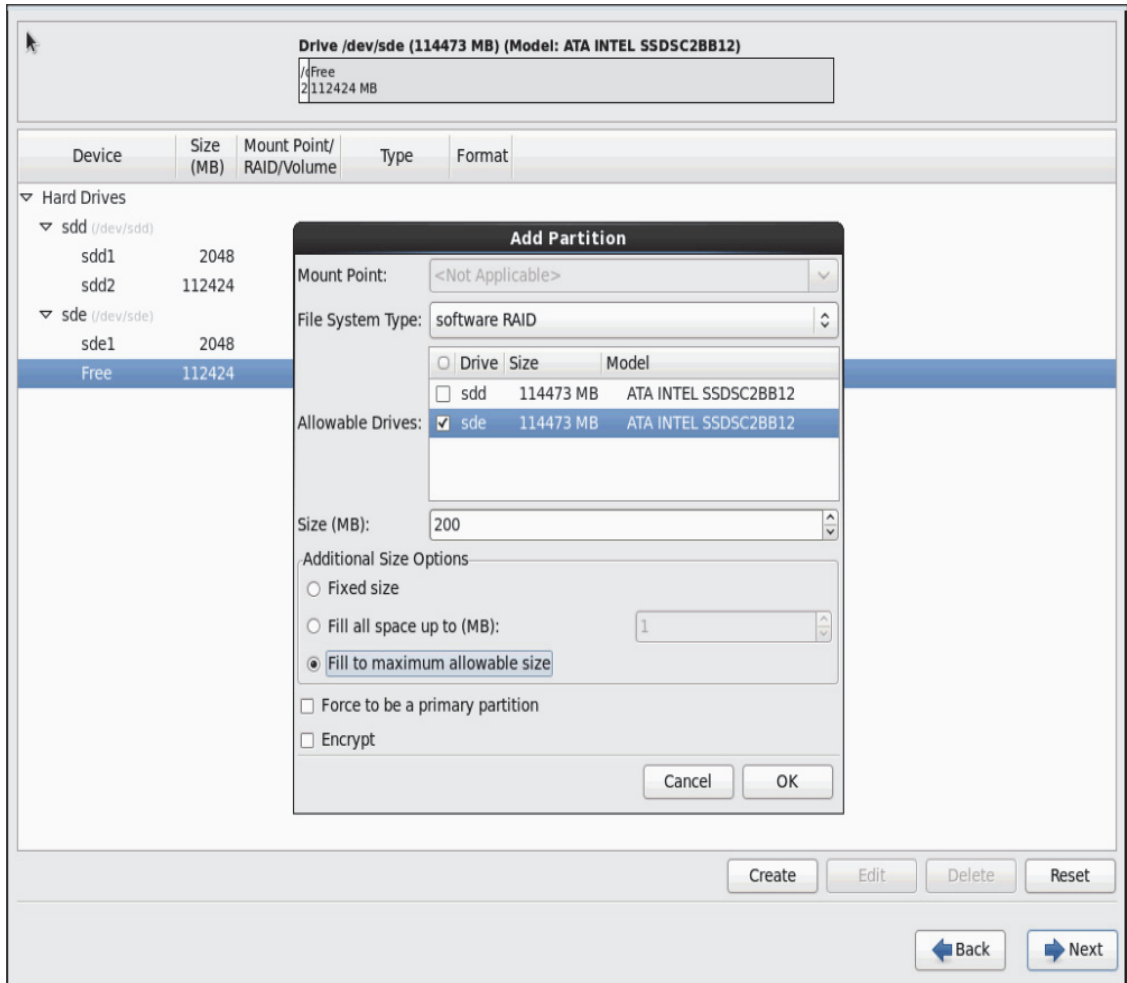
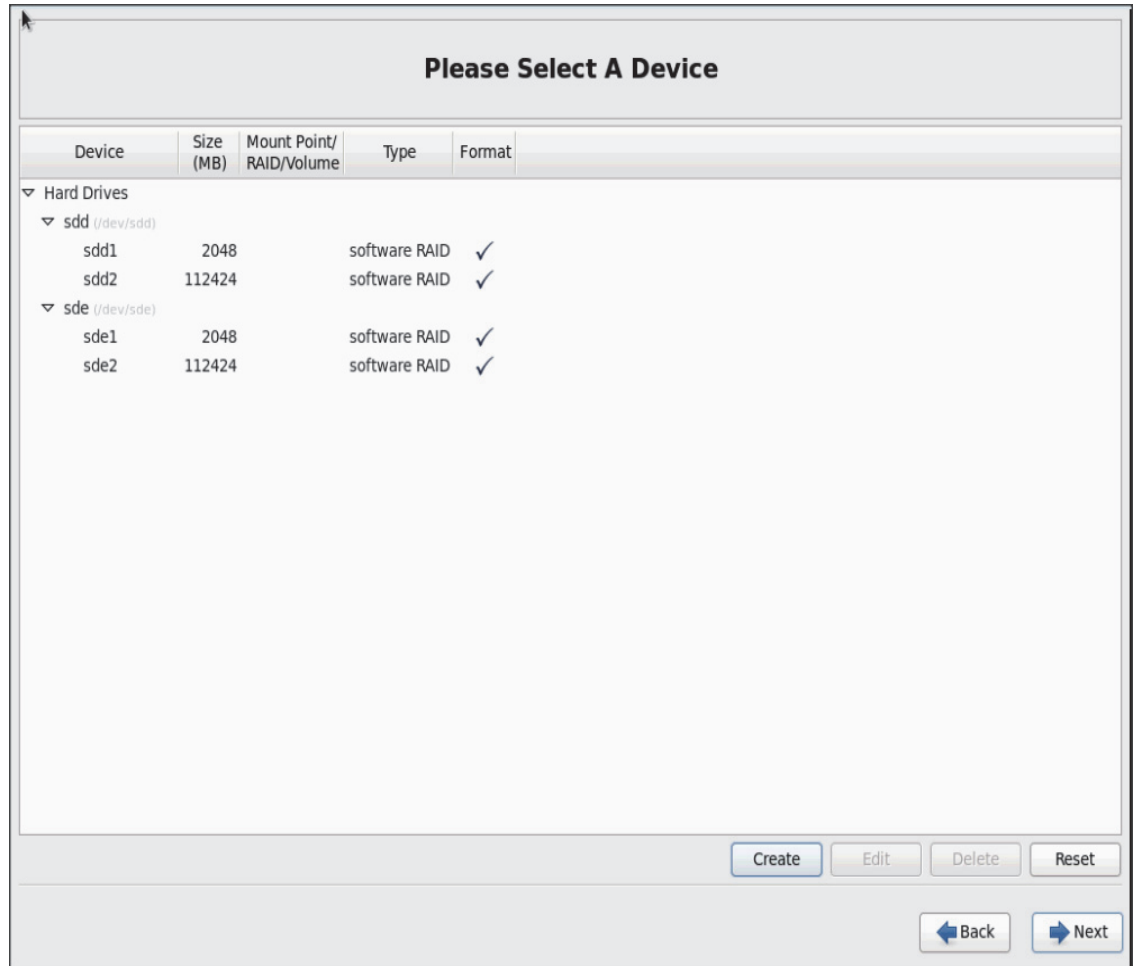


Figure 87 RHEL Installation: Add RAID Partition



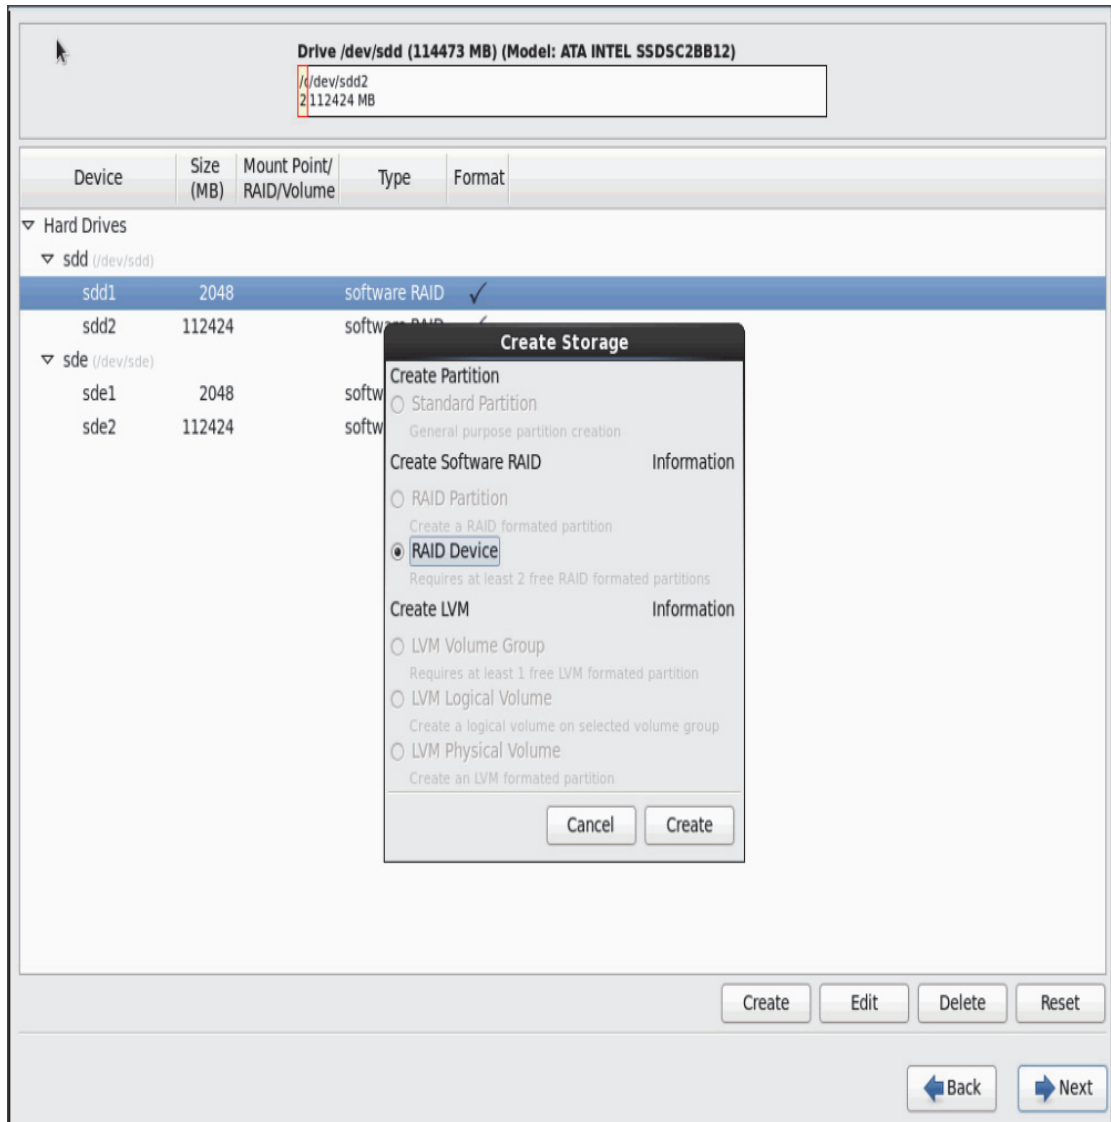
21. The above steps created 2 boot and 2 root (/) partitions. Following steps will RAID1 Devices

Figure 88 RHEL Installation: Selected RAID Devices



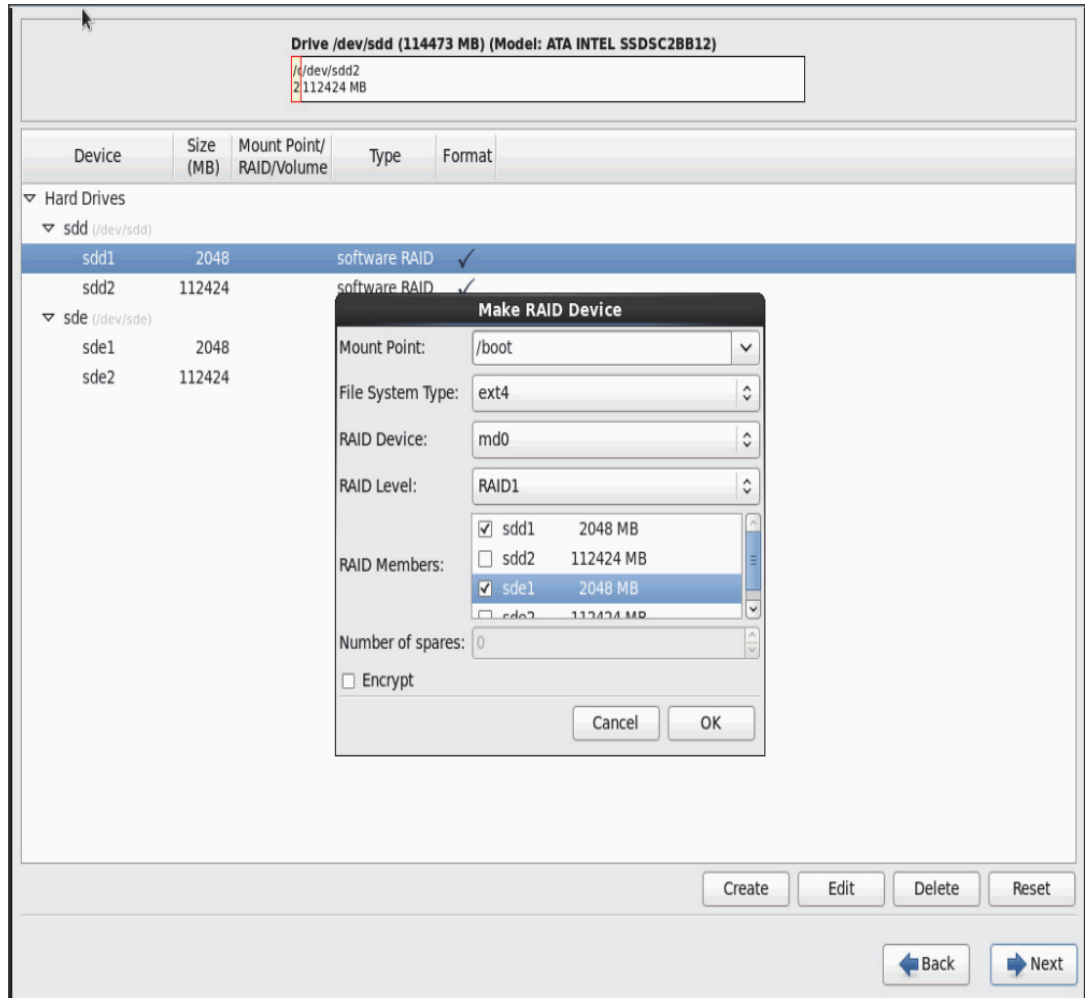
22. Choose one of the boot partitions and click on **Create > RAID Device**.

Figure 89 RHEL Installation: Select RAID Device



23. Choose this as /boot (boot device) and in RAID members, choose all the boot partitions created above in order to create a software RAID 1 for boot.

Figure 90 RHEL Installation: Make RAID Device



24. Similarly repeat for / partitions created above choosing both members with mount point as “/”.

Figure 91 RHEL Installation: Select RAID Device

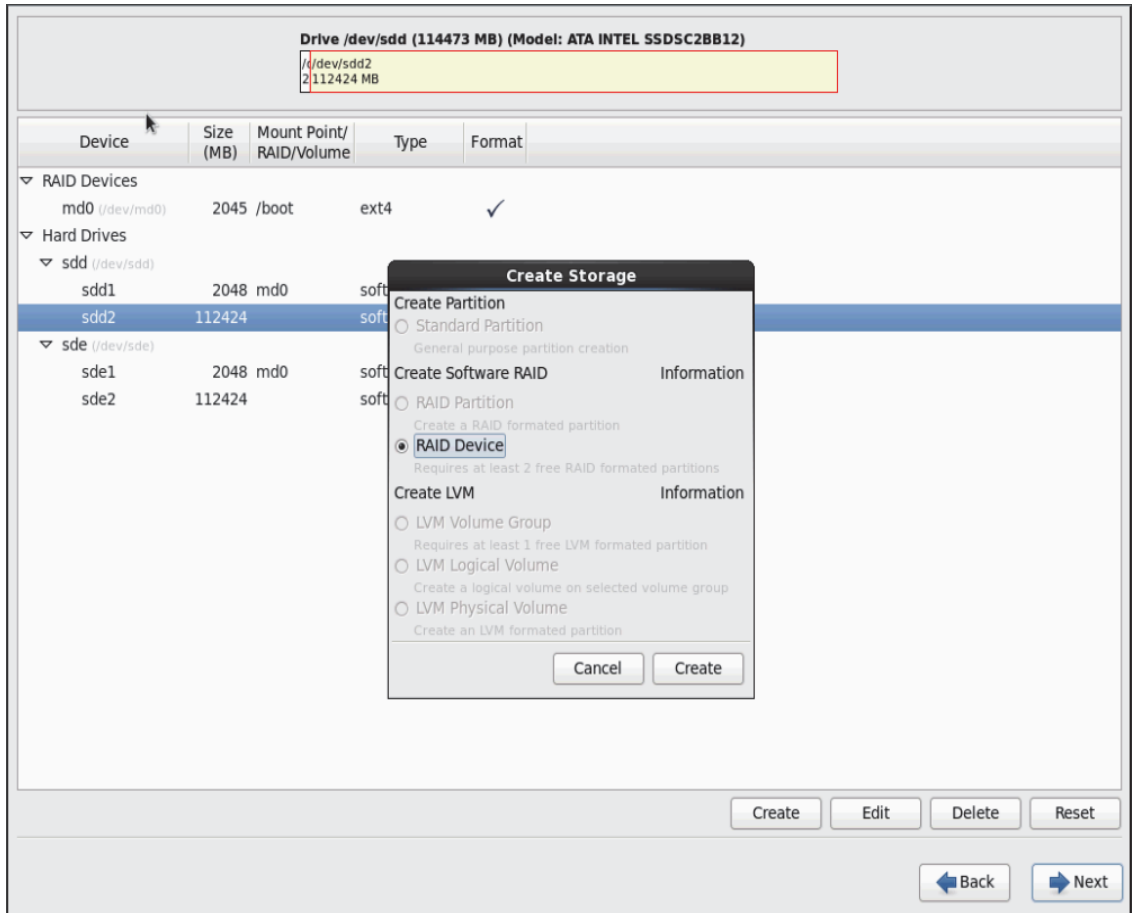


Figure 92 RHEL Installation: Make RAID Device

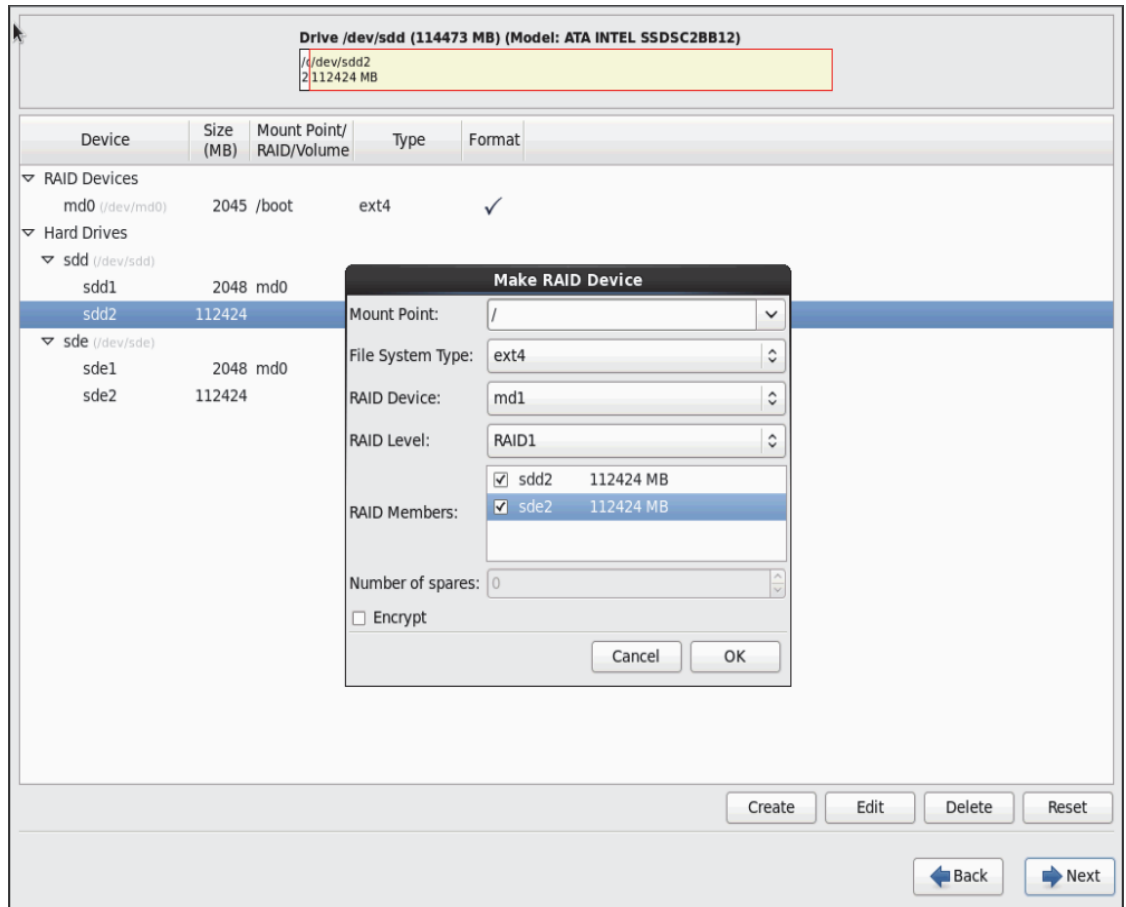
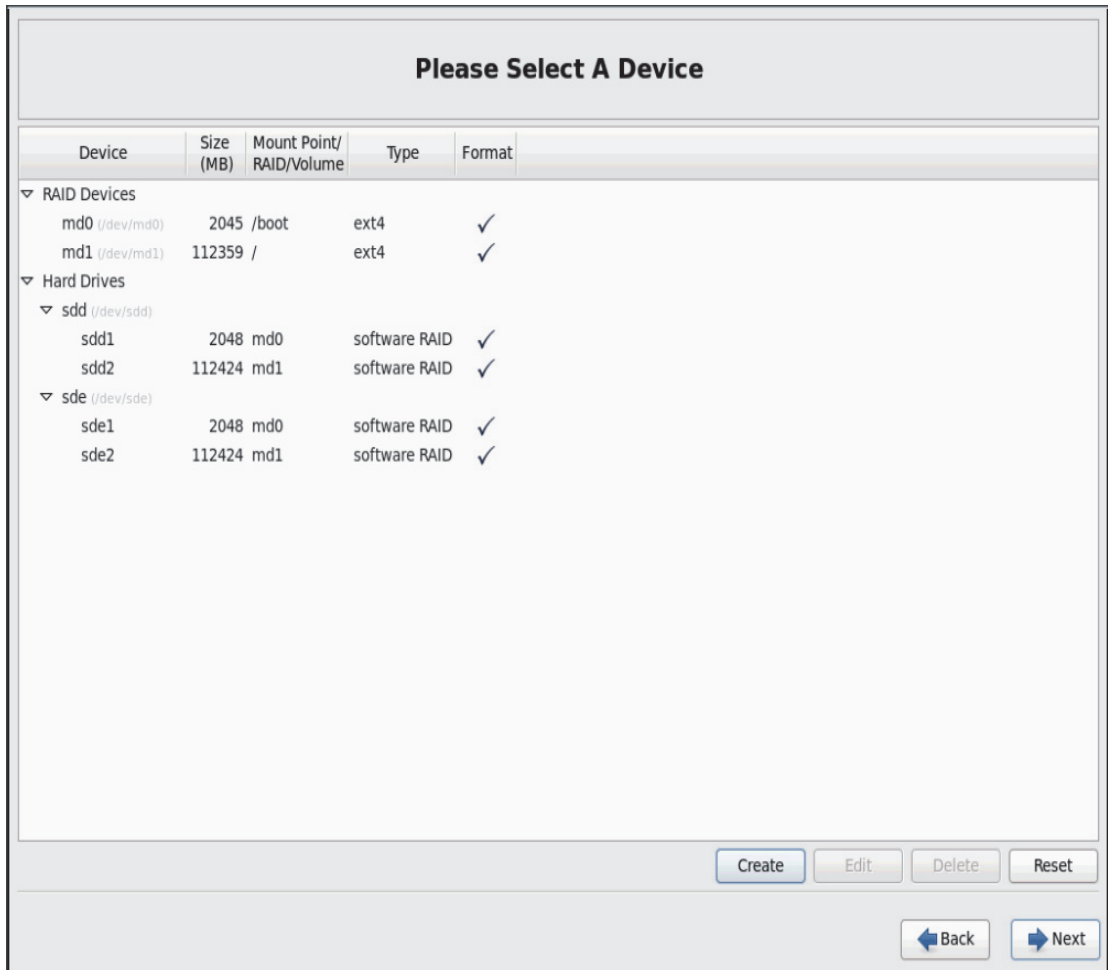
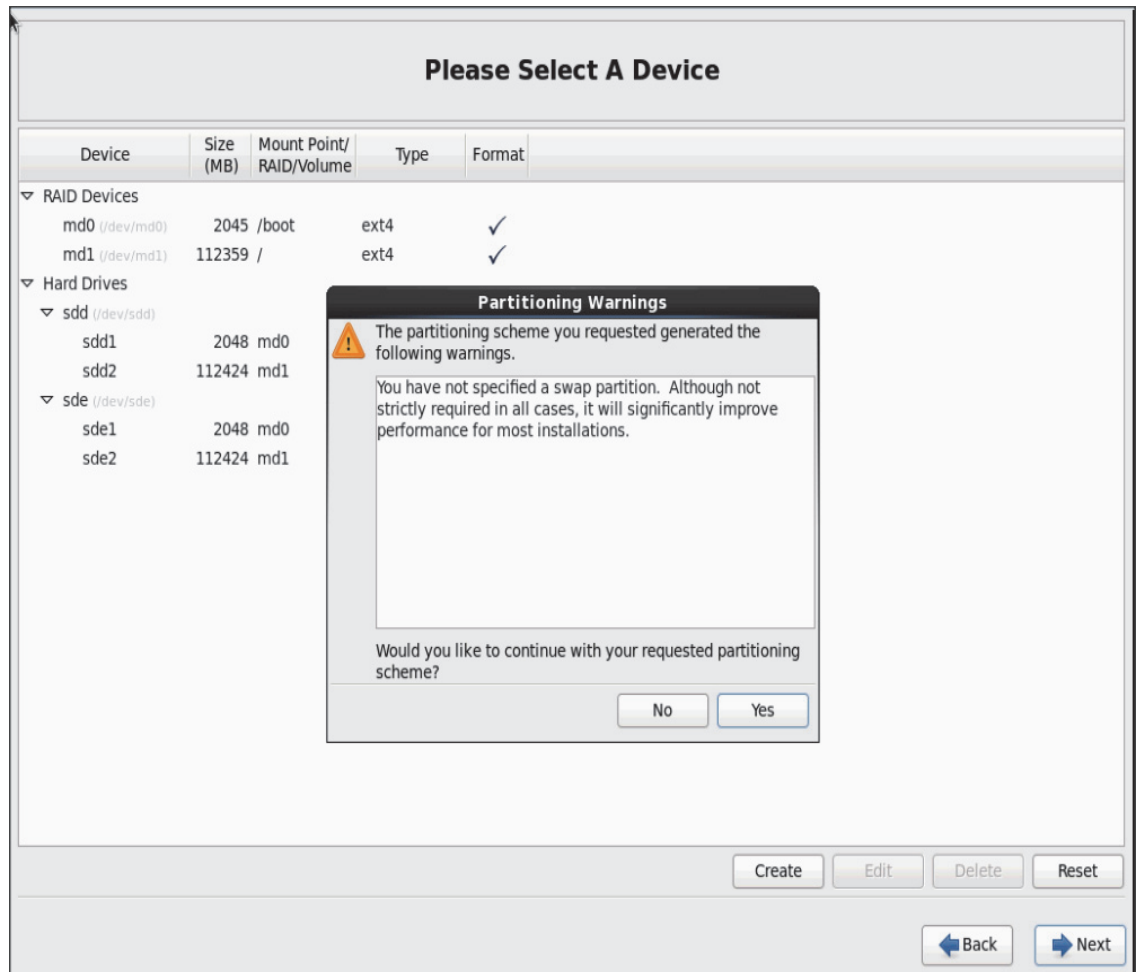


Figure 93 RHEL Installation: All the Selected Devices



25. Click on Next.

Figure 94 RHEL Installation: Warning before RAID Partitioning

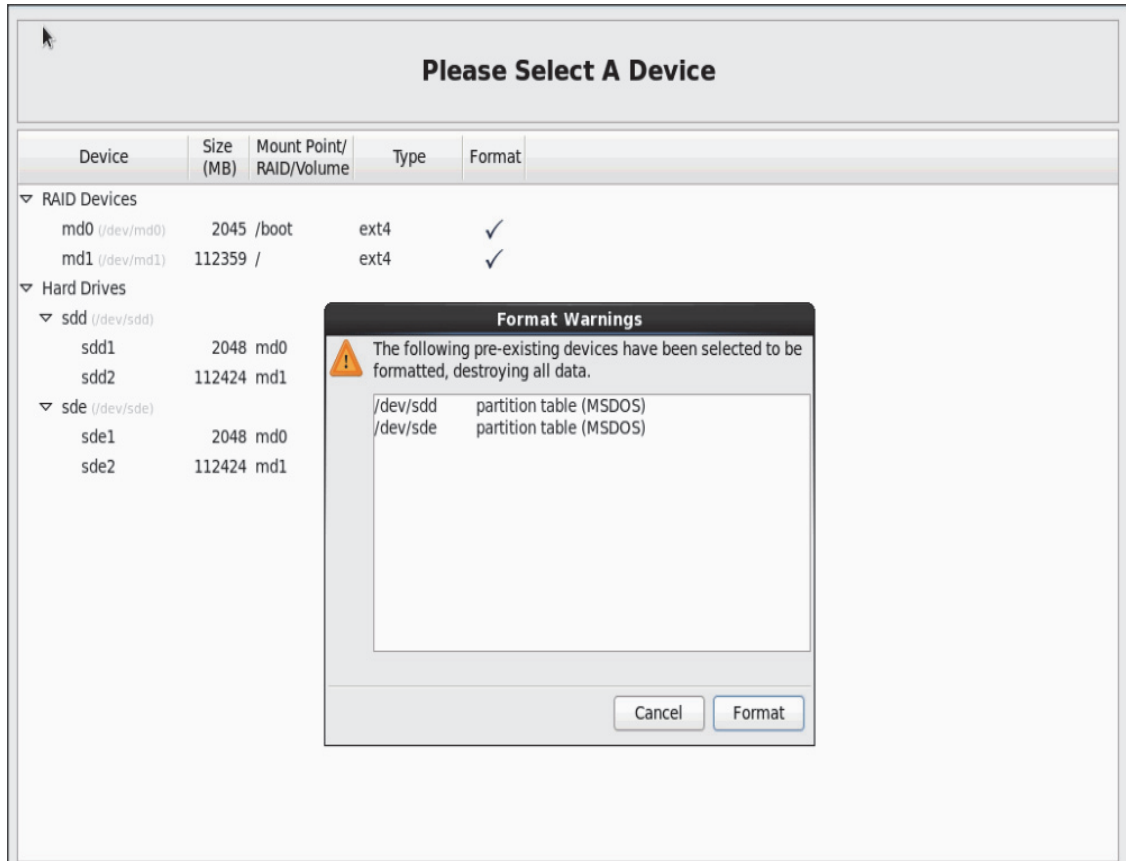


Note

Swap partition can be created using the similar steps, however, since these systems are high in memory, this step is skipped (click **Yes**).

26. Click **Next**, and then click **Format**.

Figure 95 RHEL Installation: Destroy Old Devices



27. Select default settings and click Next.

Figure 96 *RHEL Installation: Installing Boot Loader*

The screenshot shows the 'Installing Boot Loader' step in the RHEL installation process. It features two checkboxes: 'Install boot loader on /dev/sdd.' (checked) and 'Use a boot loader password' (unchecked). Below these are 'Change device' and 'Change password' buttons. A table titled 'Boot loader operating system list' contains one entry: 'Red Hat Enterprise Linux' on device '/dev/md1'. To the right of the table are 'Add', 'Edit', and 'Delete' buttons. At the bottom right are 'Back' and 'Next' navigation buttons.

Install boot loader on /dev/sdd.

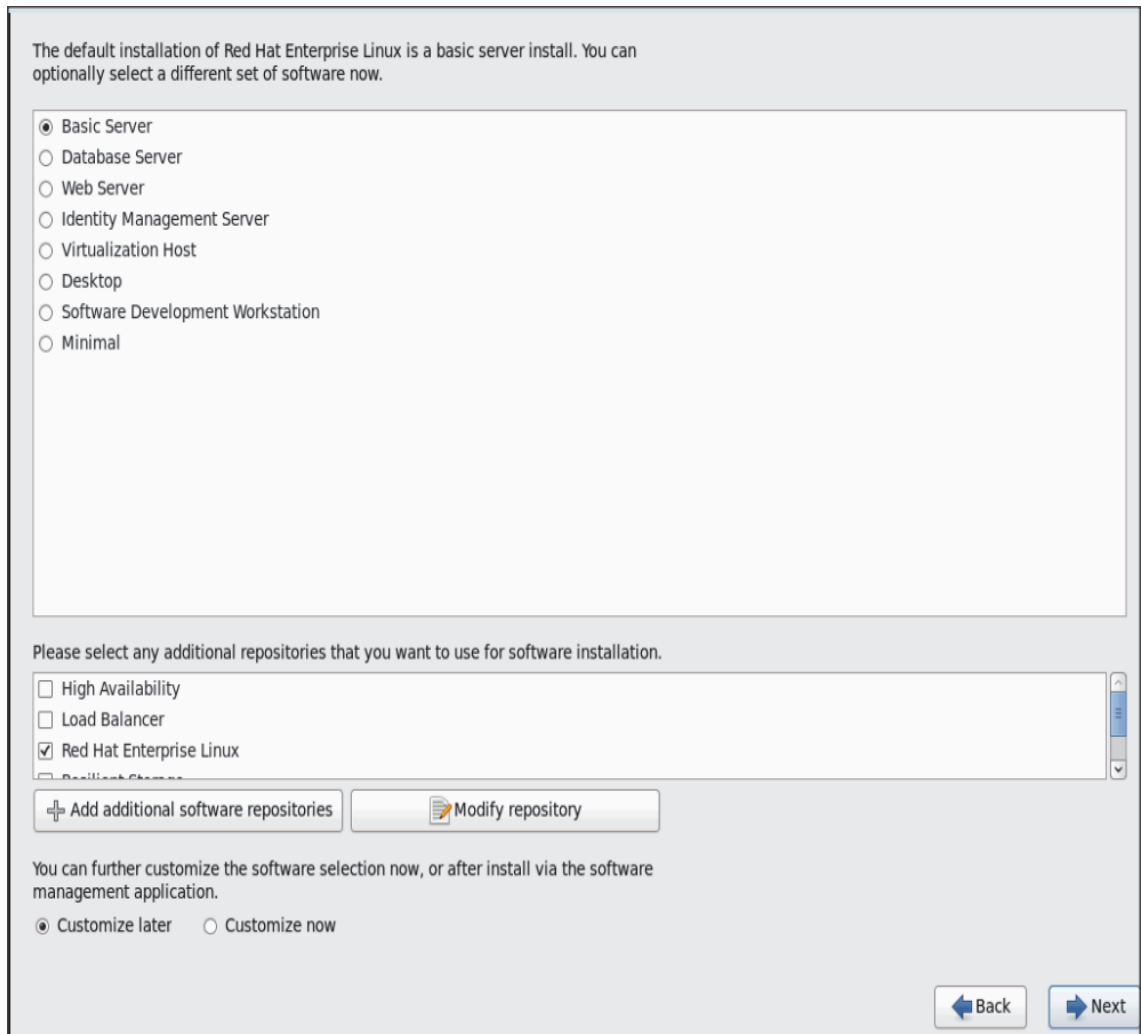
Use a boot loader password

Boot loader operating system list

Default	Label	Device
<input checked="" type="radio"/>	Red Hat Enterprise Linux	/dev/md1

28. Continue with RHEL Installation as shown below.

Figure 97 RHEL Installation: Keep the Default Installation Option



29. Once the installation is complete reboot the system.
Repeat the steps 1 through 29 to install Red Hat Linux 6.5 on Servers 66 through 68.



Note

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

Installing Red Hat Enterprise Linux 6.5 using software RAID on C240 M4 Systems

The following section provides detailed procedures for installing Red Hat Enterprise Linux 6.5 using Software RAID (OS based Mirroring) on Cisco UCS C240 M4 servers.

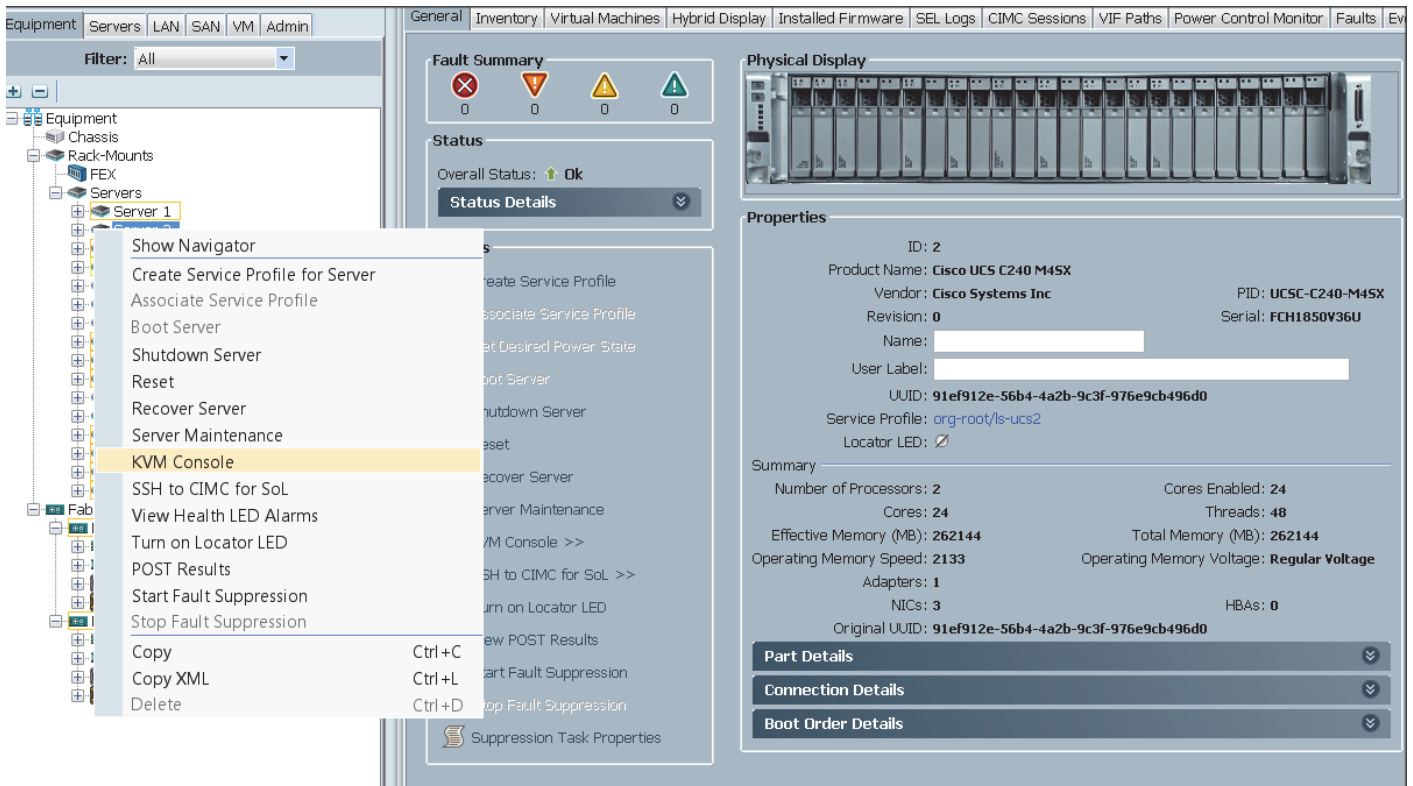
There are multiple methods to install Red Hat Linux operating system. The installation procedure described in this deployment guide uses KVM console and virtual media from Cisco UCS Manager.



Note This requires RHEL 6.5 DVD/ISO for the installation.


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

Figure 98 *Selecting KVM Console Option*




5. In the KVM window, select the **Virtual Media** tab.
6. Click the **Activate Virtual Devices** from the **Virtual Media** tab.

Figure 99 KVM Console



7. In the KVM window, select the **Virtual Media** tab and Click the **Map CD/DVD**.

Figure 100 KVM Console



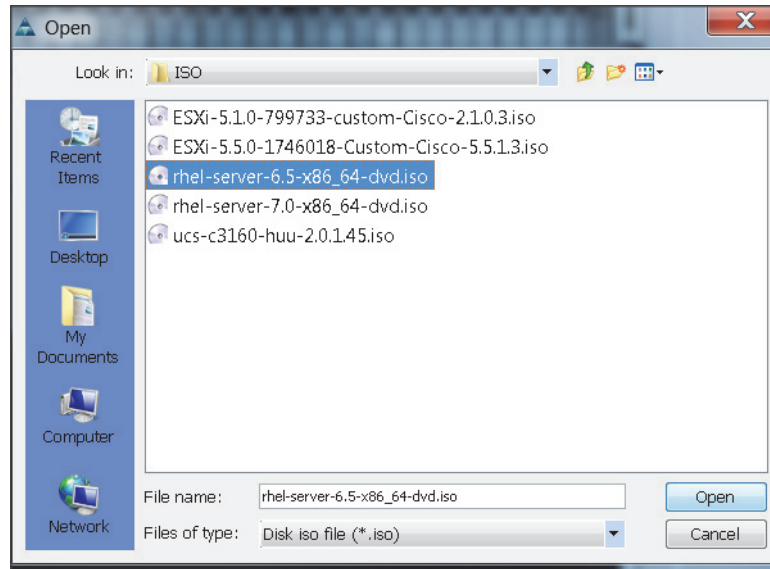
8. Browse to the Red Hat Enterprise Linux Server 6.5 installer ISO image file.

**Note**

The Red Hat Enterprise Linux 6.5 DVD is assumed to be on the client machine.

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

Figure 101 Browse to Red Hat Enterprise Linux ISO Image



10. In the KVM window, select the **KVM** tab to monitor during boot.
11. In the KVM window, select the **Macros > Static Macros > Ctrl-Alt-Del** button in the upper left corner.
12. Click **OK**.
13. Click **OK** to reboot the system.
14. On reboot, the machine detects the presence of the Red Hat Enterprise Linux Server 6.5 install media.
15. Select the **Install or Upgrade an Existing System**

Figure 102 RHEL Installation



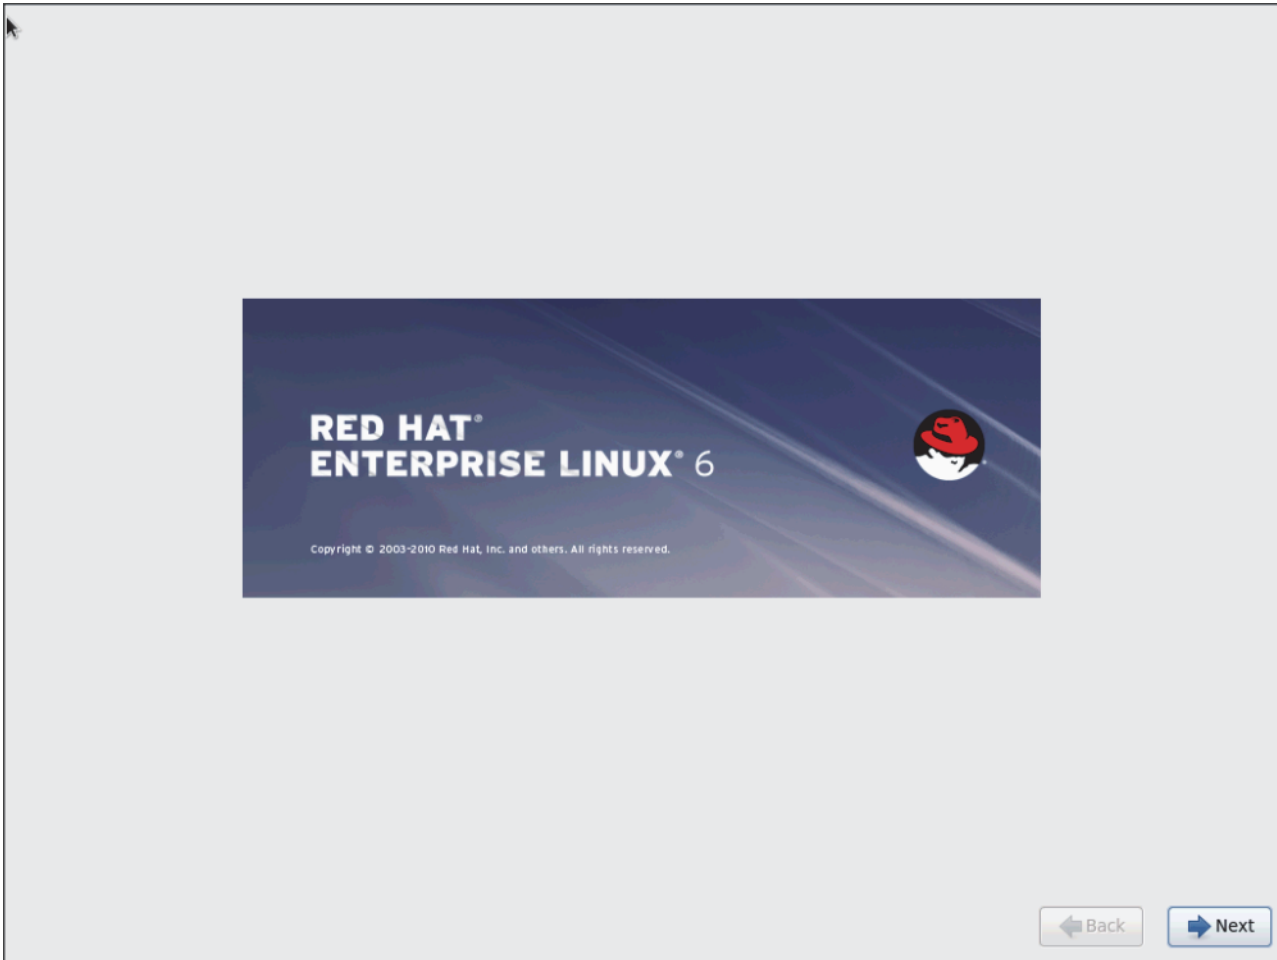
16. Skip the Media test and start the installation.

Figure 103 RHEL Installation: Media Test



17. Click Next

Figure 104 RHEL Installation: Installation Wizard



18. Select language of installation, and then Click **Next**

Figure 105 RHEL Installation: Language Selection

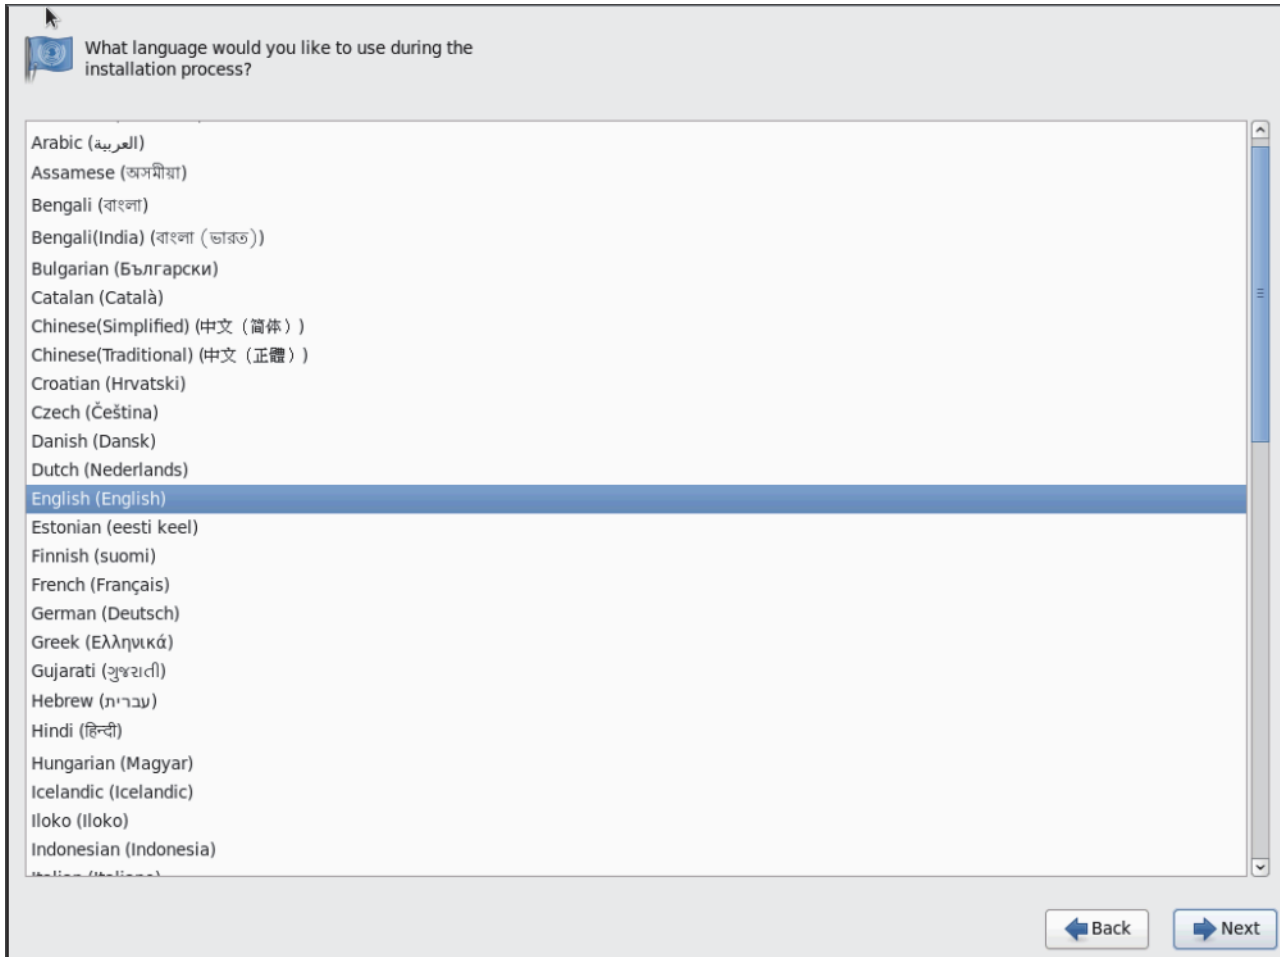
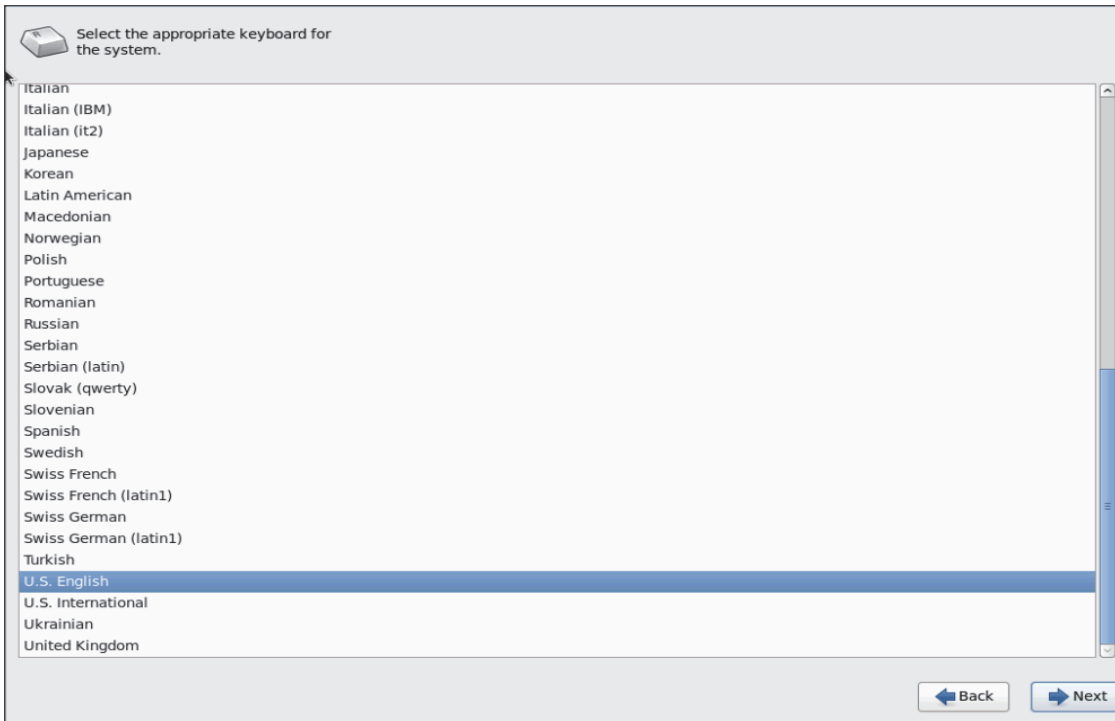
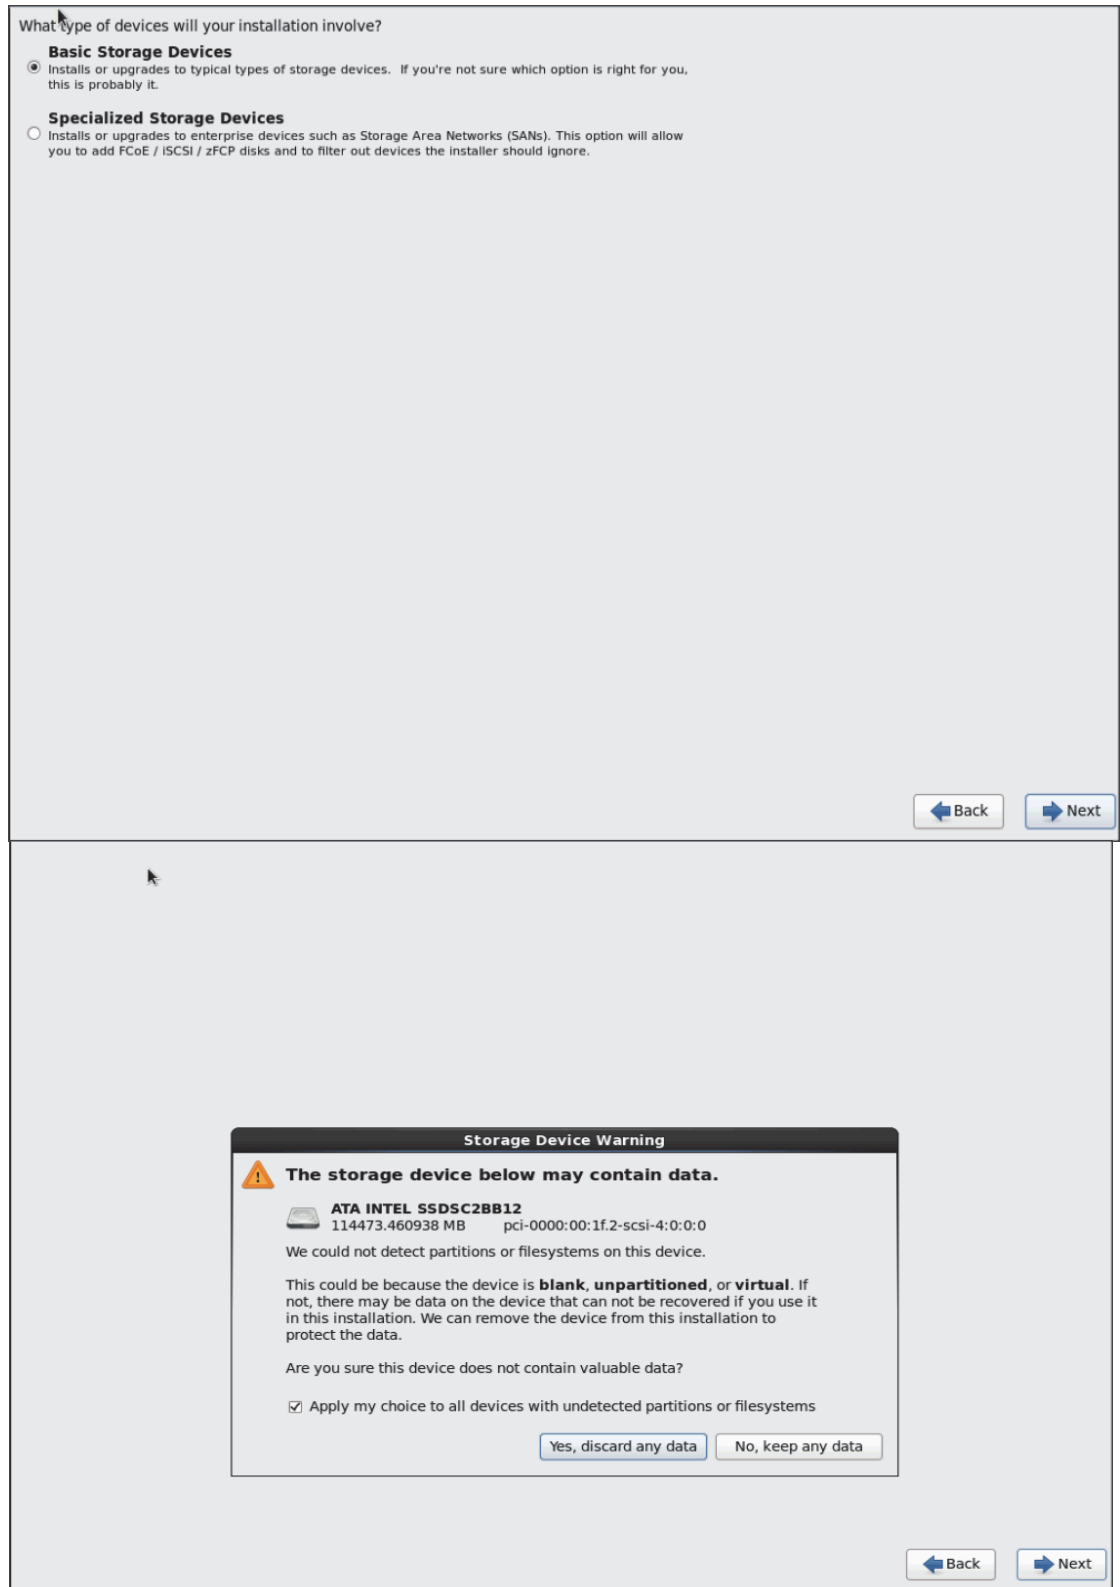


Figure 106 RHEL Installation: Language Selection



19. Select Basic Storage Devices and Click Next

Figure 107 RHEL Installation: Installation Type



20. Provide hostname and configure Network for the host.

Figure 108 RHEL Installation: Provide Host Name

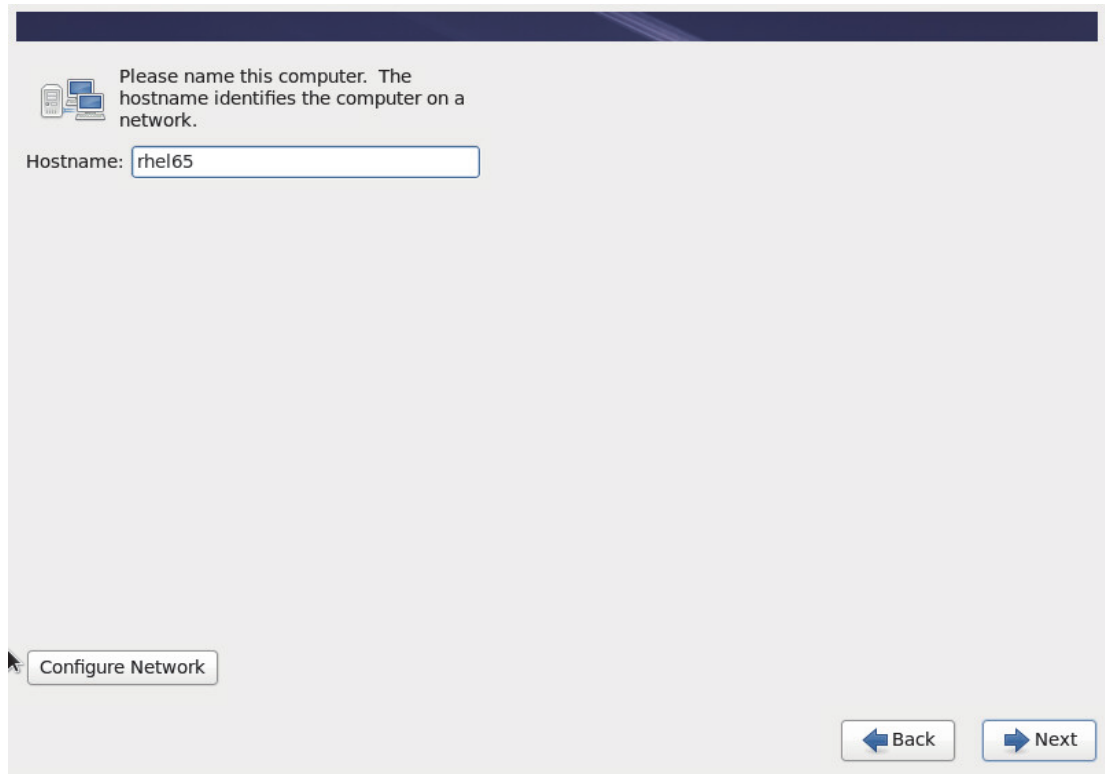


Figure 109 RHEL Installation: IPV4 Setting for eth0

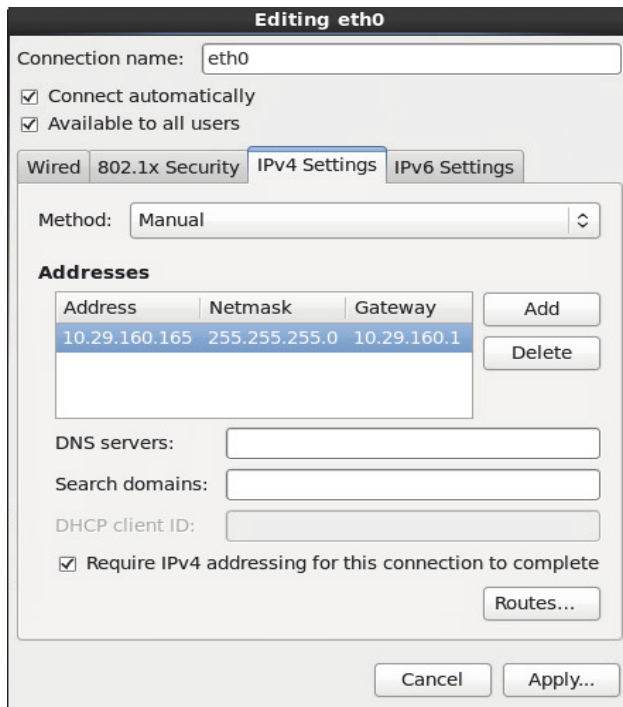


Figure 110 RHEL Installation: IPV4 Setting for eth1

Editing eth1

Connection name:

Connect automatically
 Available to all users

Wired | 802.1x Security | **IPv4 Settings** | IPv6 Settings

Method:

Addresses

Address	Netmask	Gateway
192.168.11.165	255.255.255.0	

DNS servers:

Search domains:

DHCP client ID:

Require IPv4 addressing for this connection to complete

Figure 111 *RHEL Installation: Selecting Location*

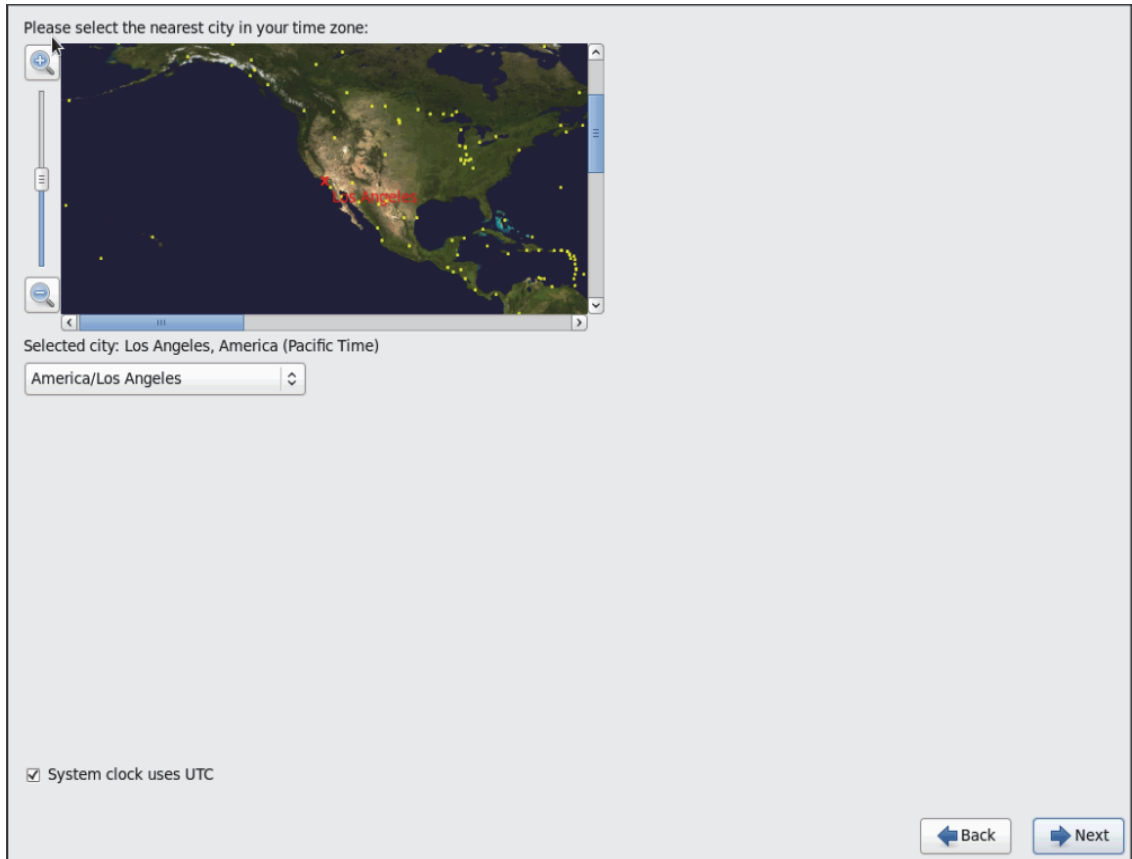
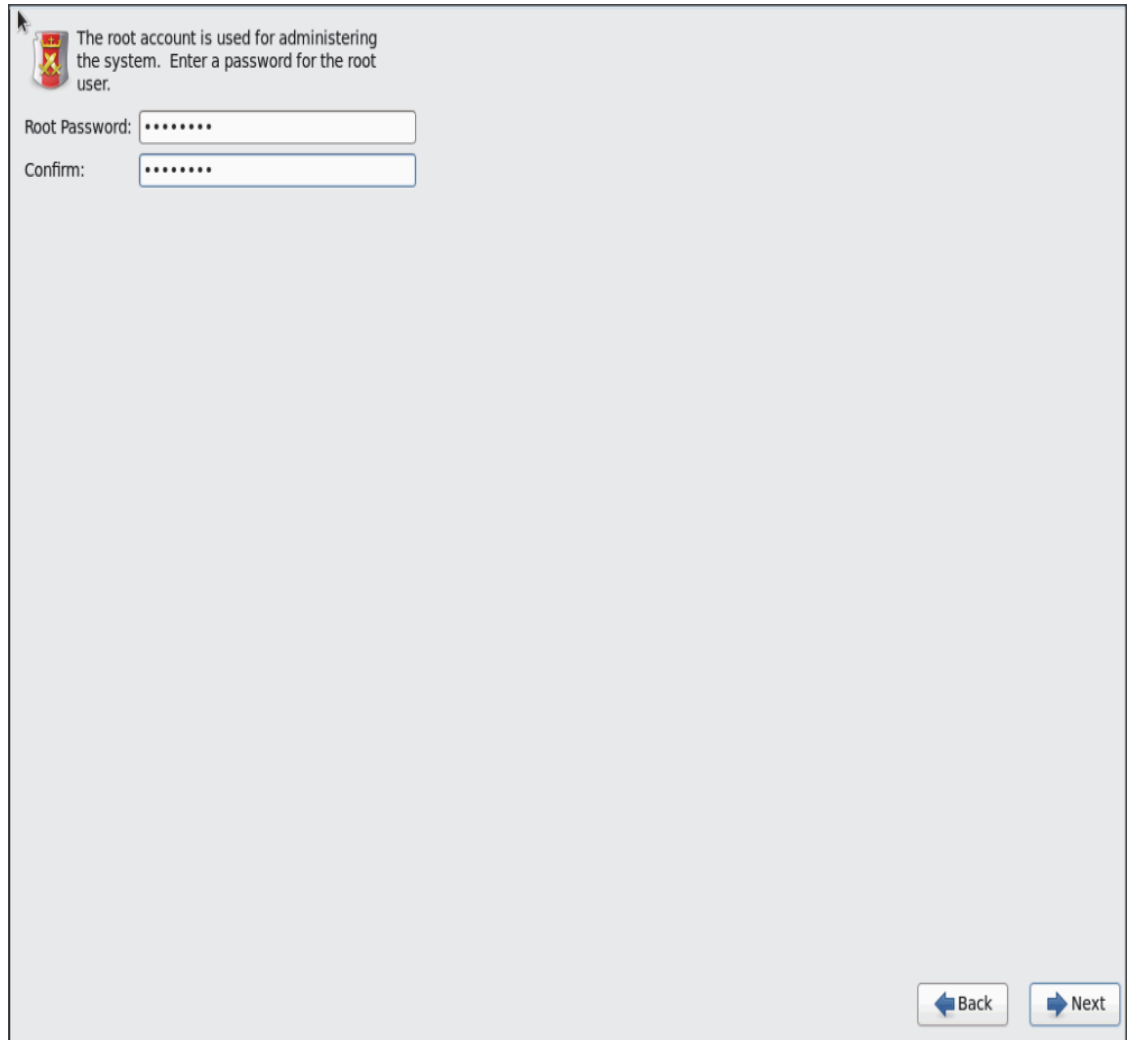


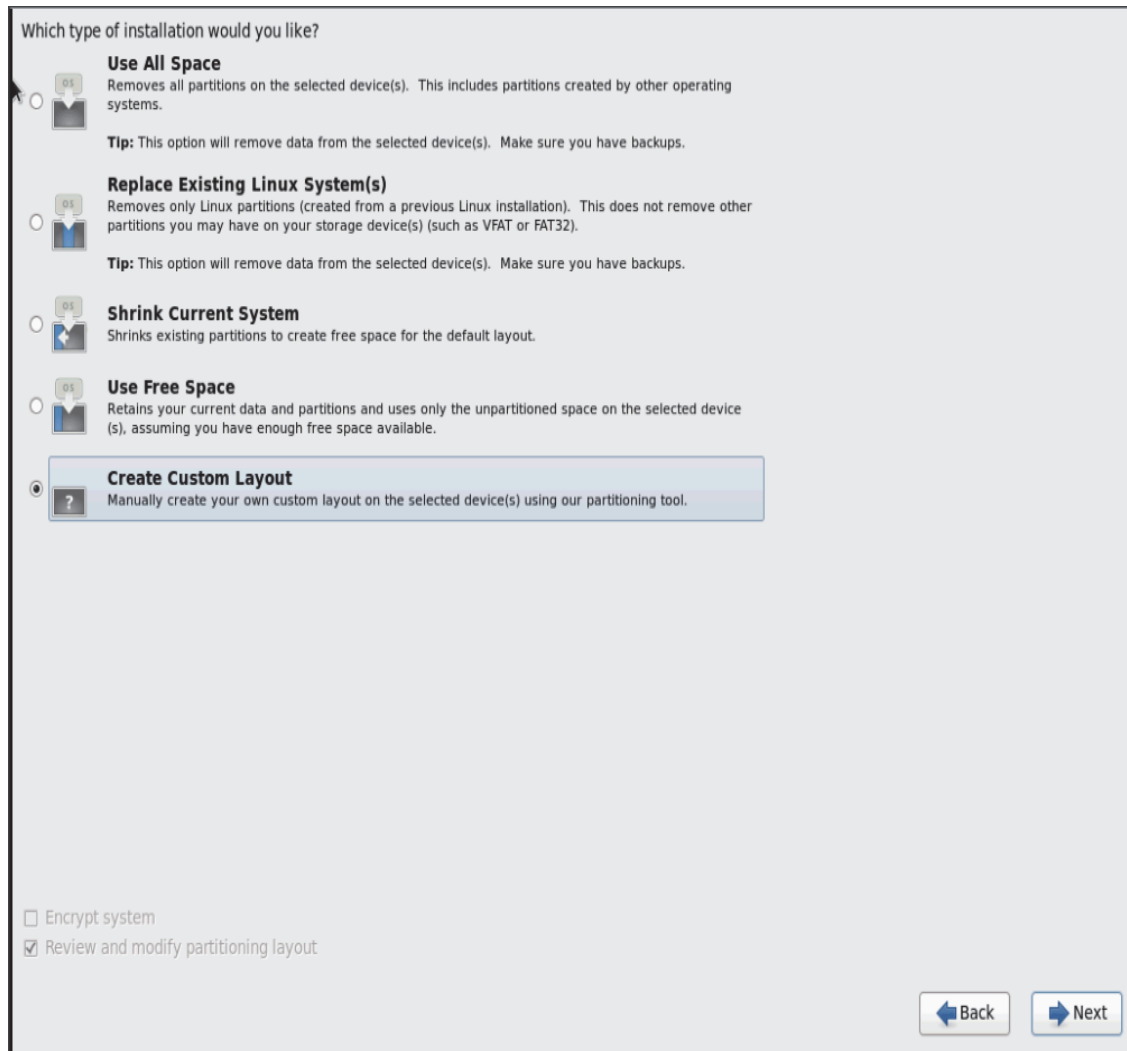
Figure 112 *RHEL Installation: Enter Root Credentials*



The screenshot shows a window titled "Enter Root Credentials" with a mouse cursor in the top-left corner. The text inside the window reads: "The root account is used for administering the system. Enter a password for the root user." Below this text are two input fields: "Root Password:" followed by a text box containing seven asterisks, and "Confirm:" followed by another text box containing seven asterisks. In the bottom-right corner of the window, there are two buttons: "Back" with a left-pointing arrow and "Next" with a right-pointing arrow.

21. Choose **Create custom layout** for Installation type.

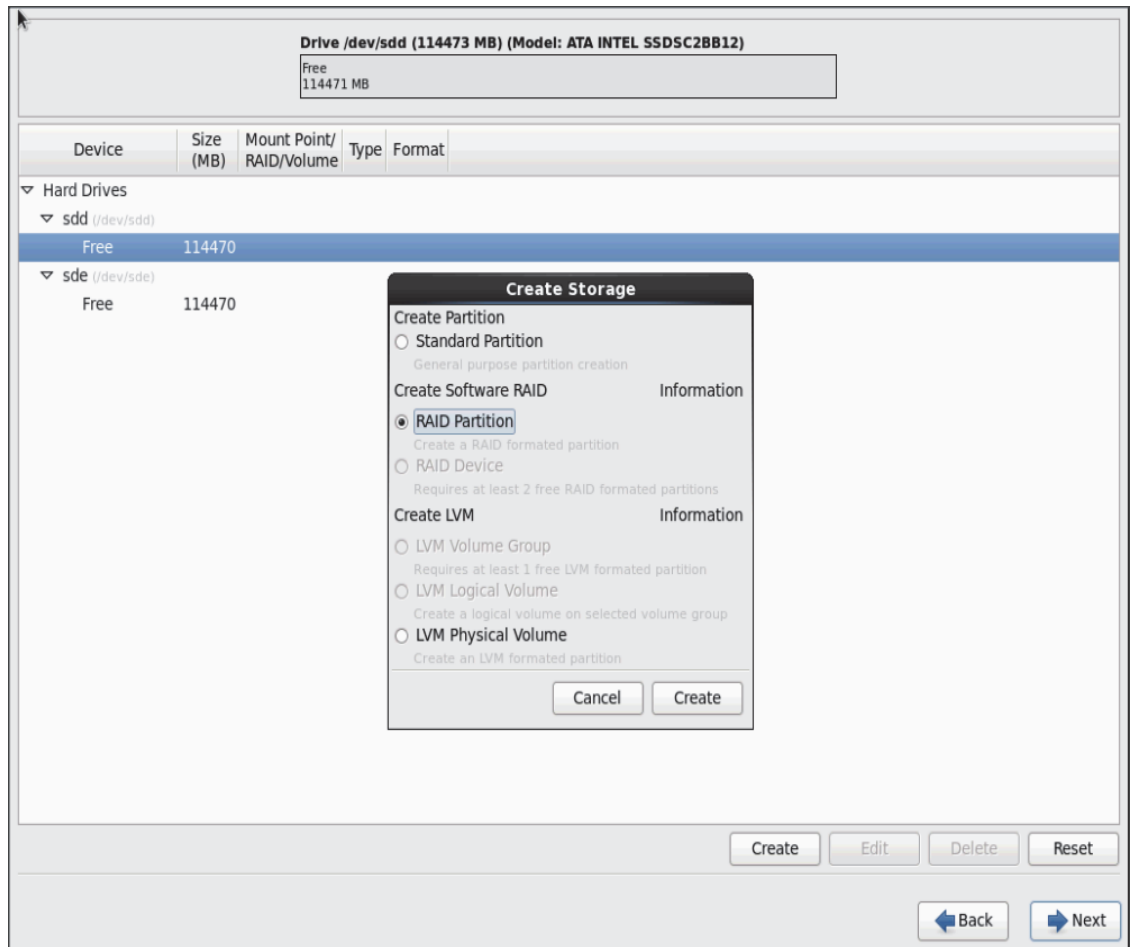
Figure 113 RHEL Installation: Create Custom Layout



Following steps can be used to create two software RAID 1 partitions for boot and / (root) partitions.

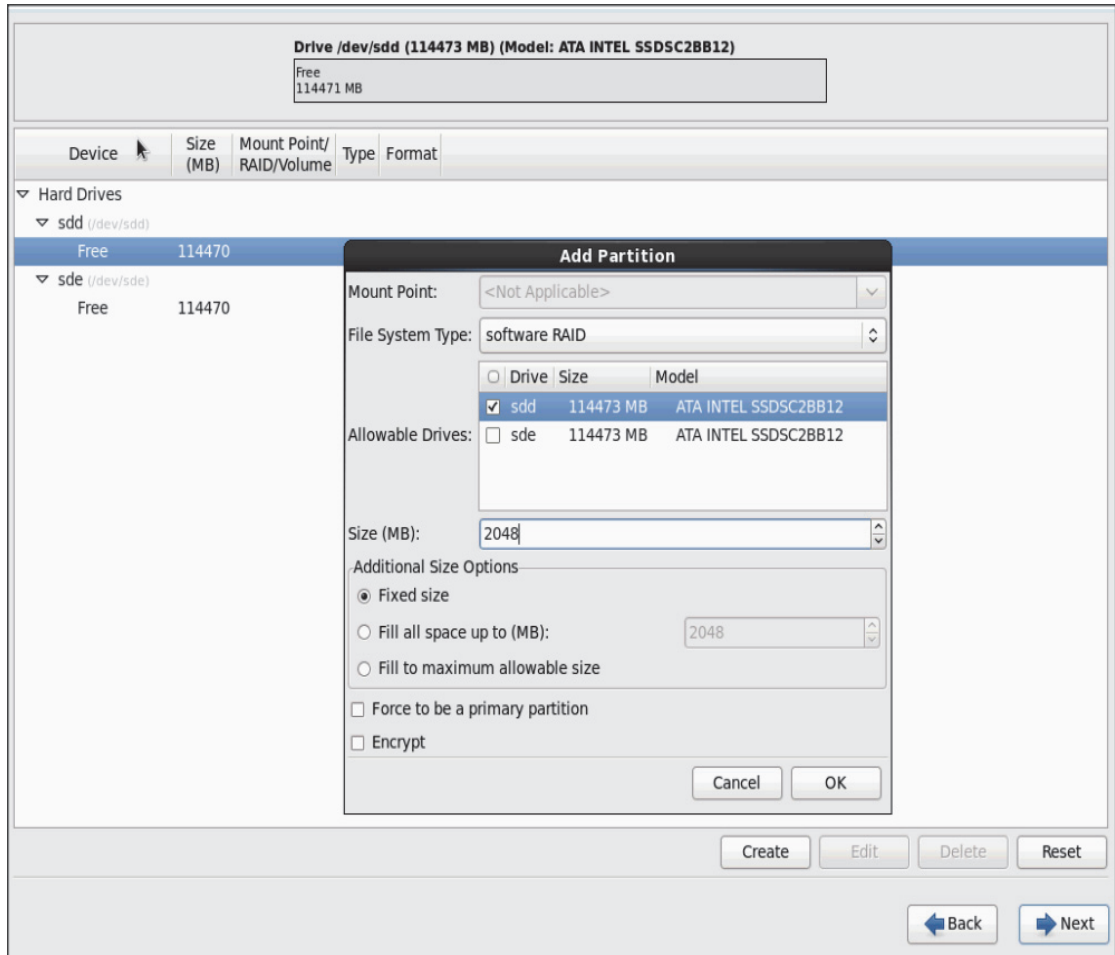
22. Choose free volume and click on **Create** and choose **RAID Partition**.

Figure 114 RHEL Installation: Create RAID Partition



23. Choose “Software RAID” for File system Type and set size for Boot volume.

Figure 115 RHEL Installation: Add Partition



24. Similarly, do the RAID partitioning for the other free volume.

Figure 116 RHEL Installation: Create RAID Partition

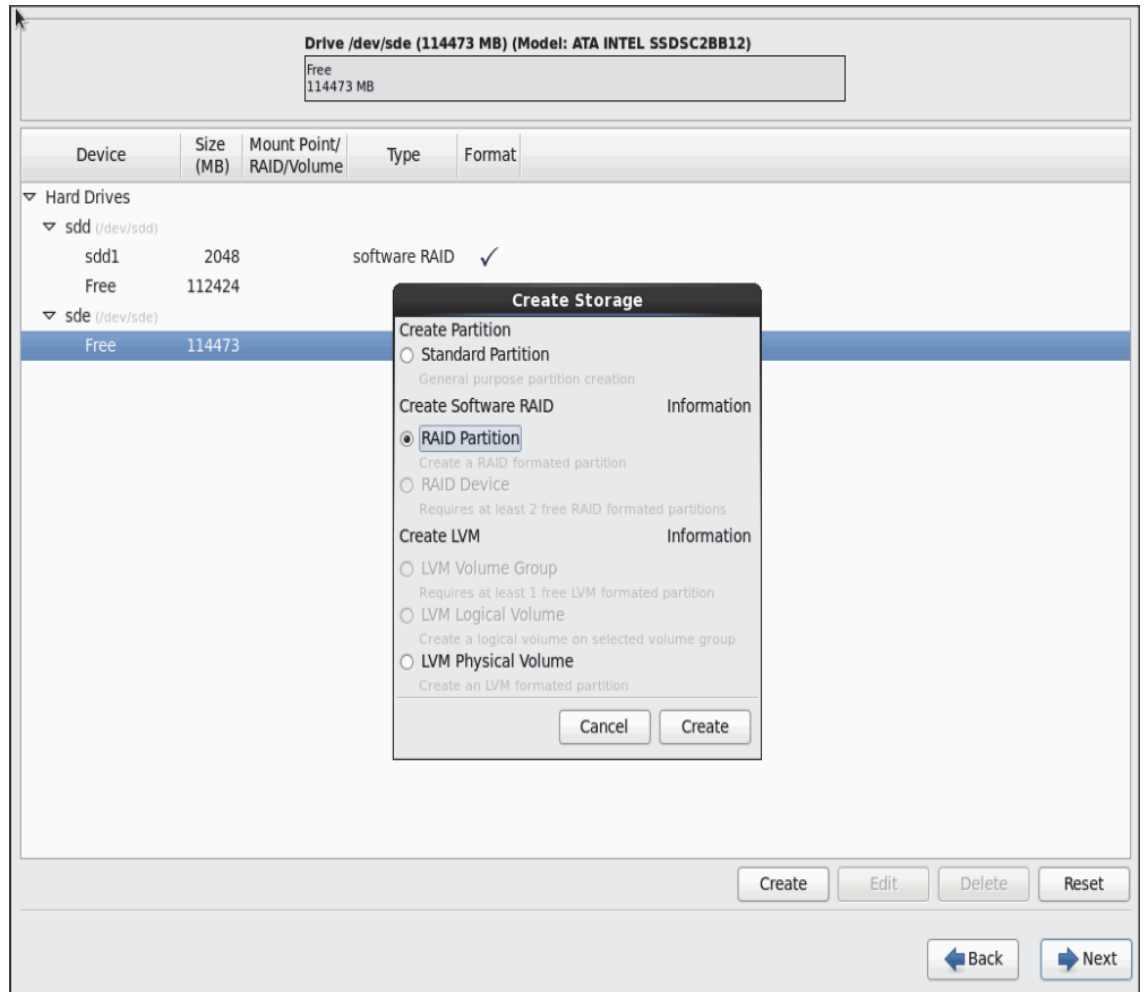
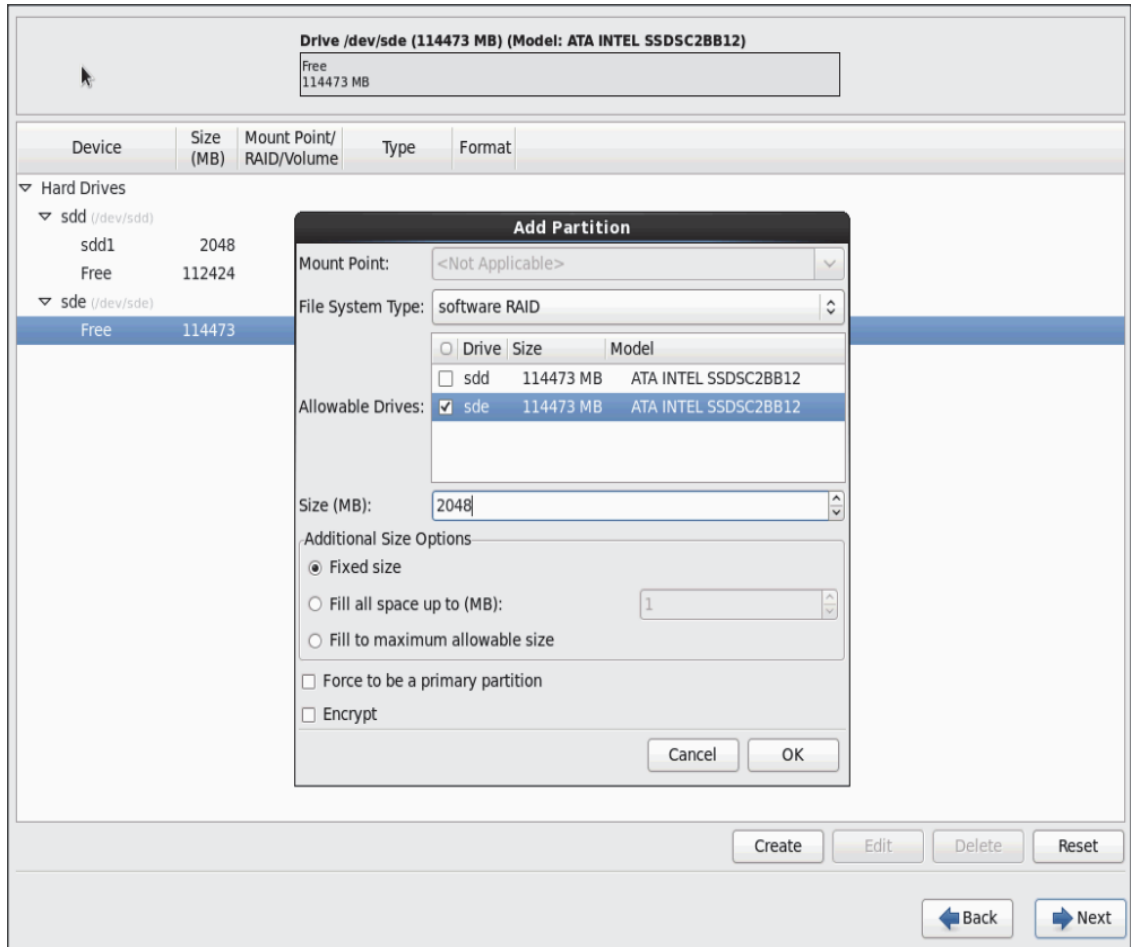


Figure 117 RHEL Installation: Add Partition



- Now similarly create RAID partitions for root (/) partition on both the devices and use rest of the available space.

Figure 118 RHEL Installation: Create RAID Partition

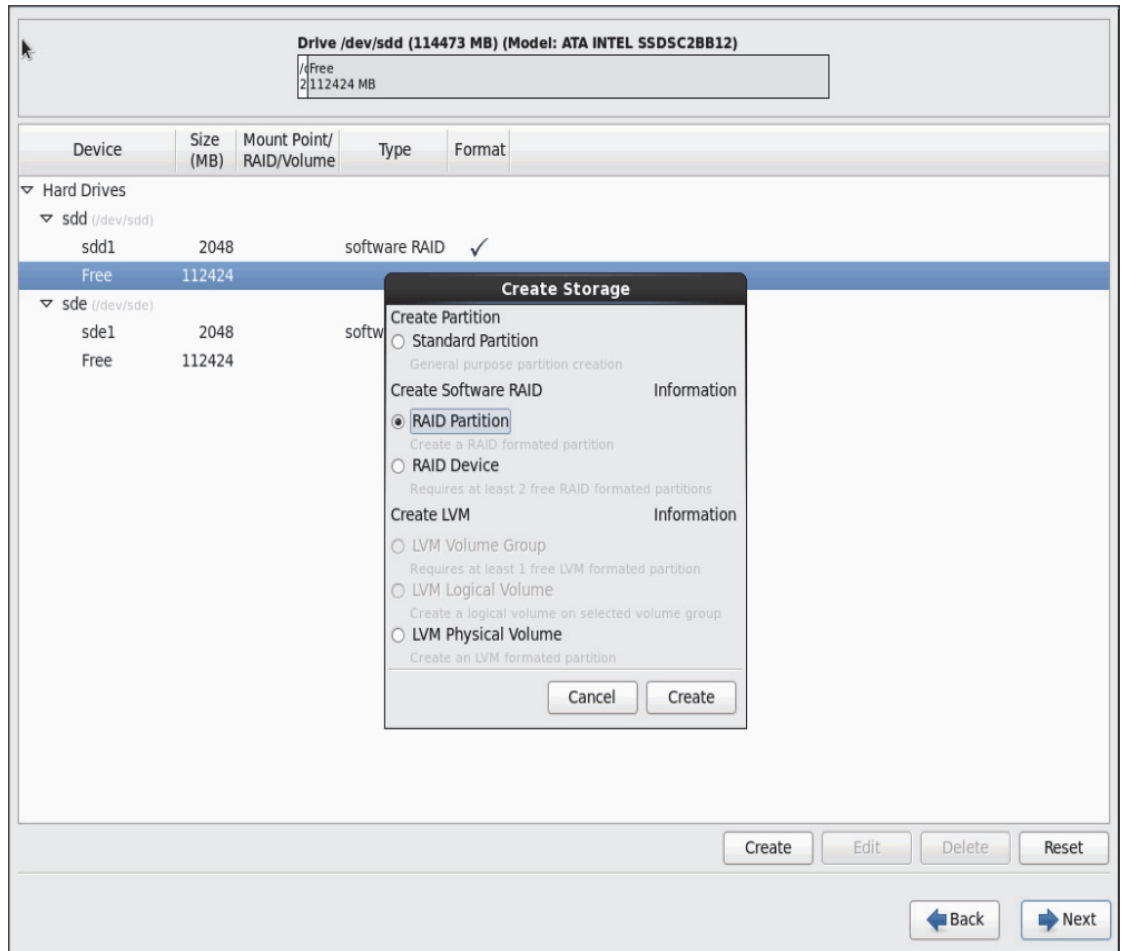


Figure 119 RHEL Installation: Add Partition

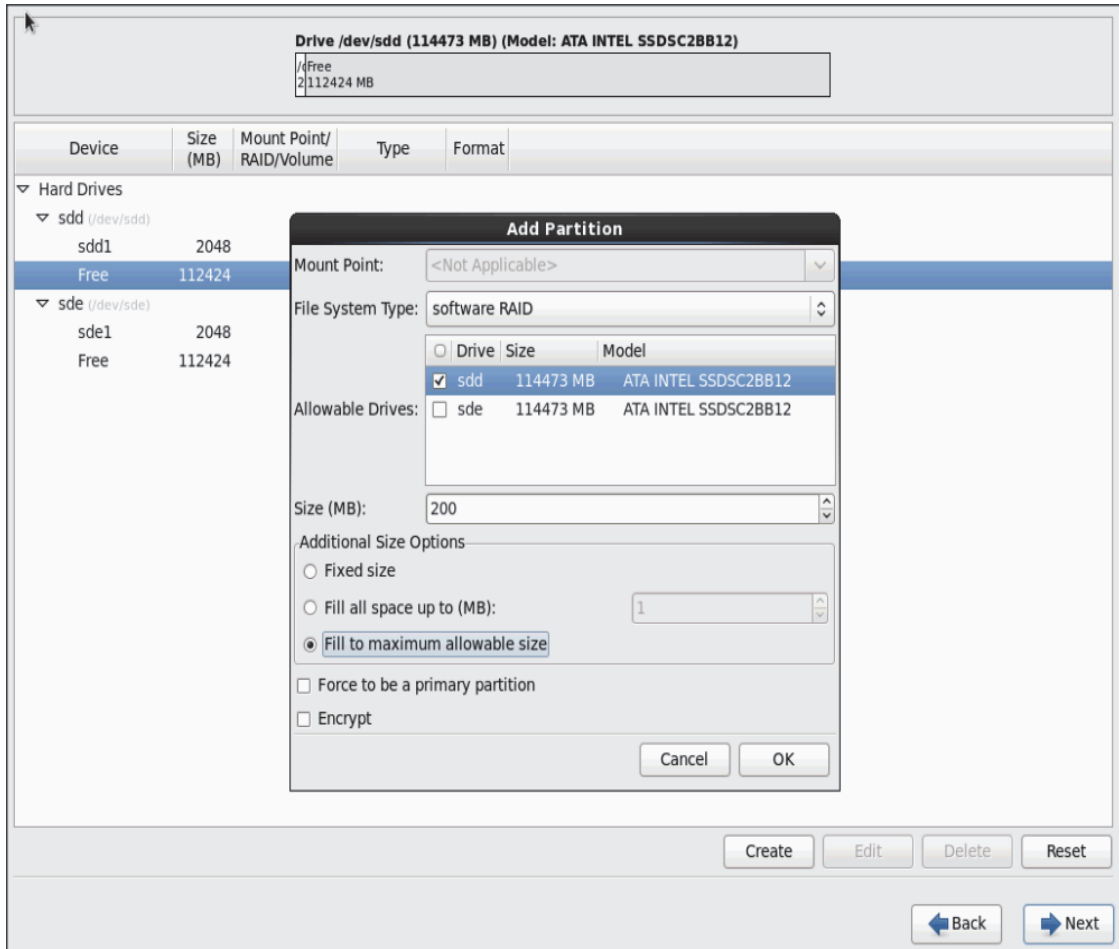


Figure 120 RHEL Installation: Create RAID Partition

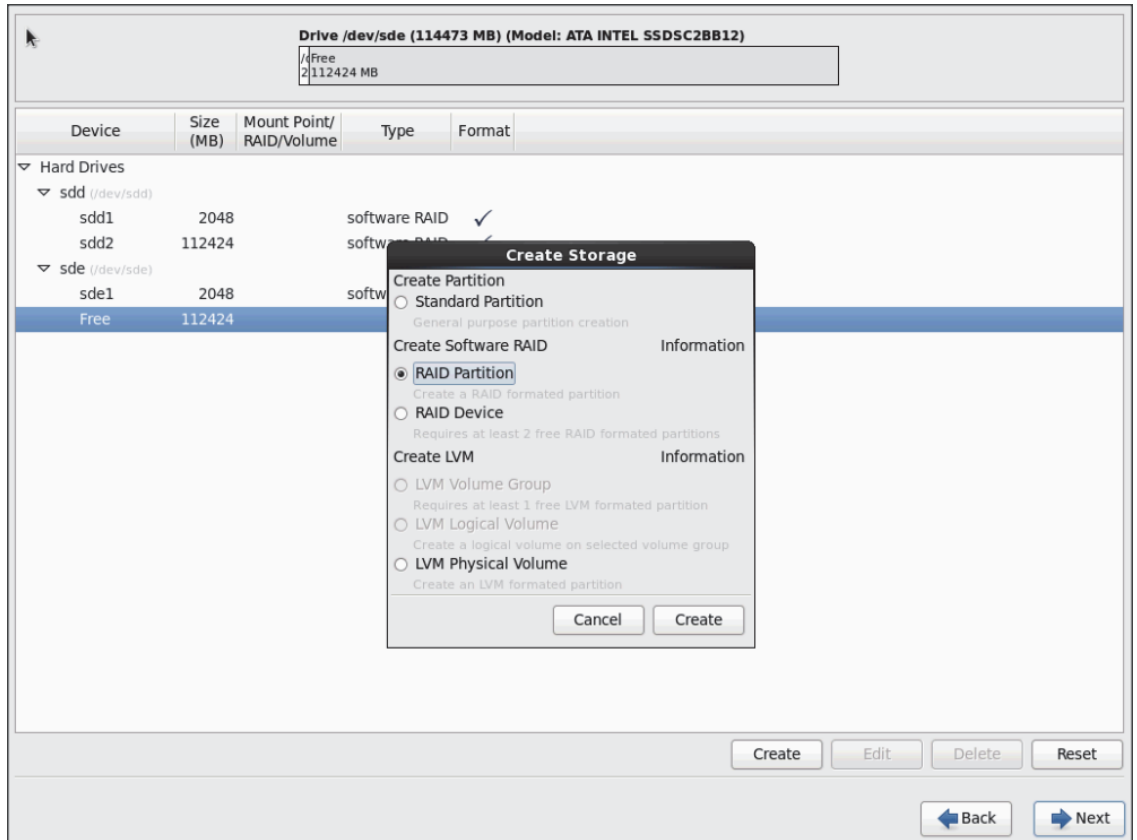
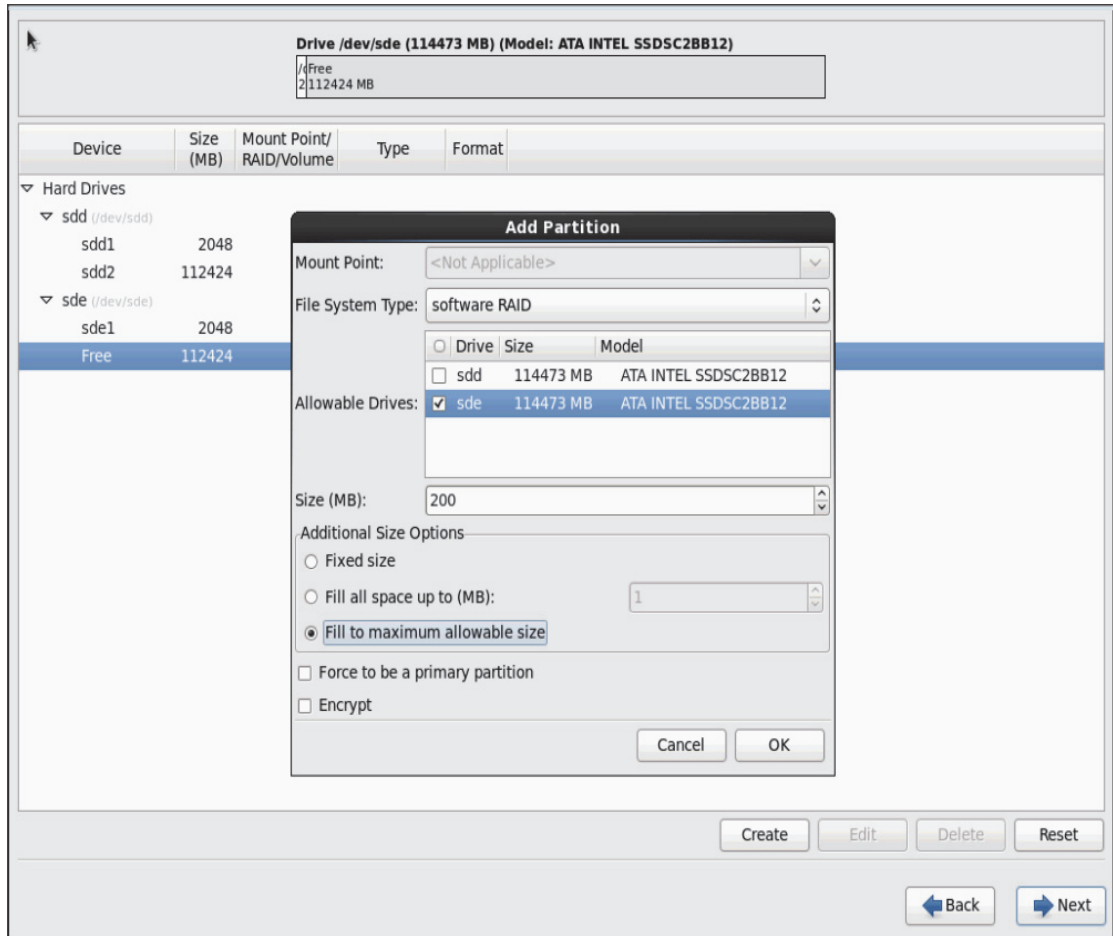
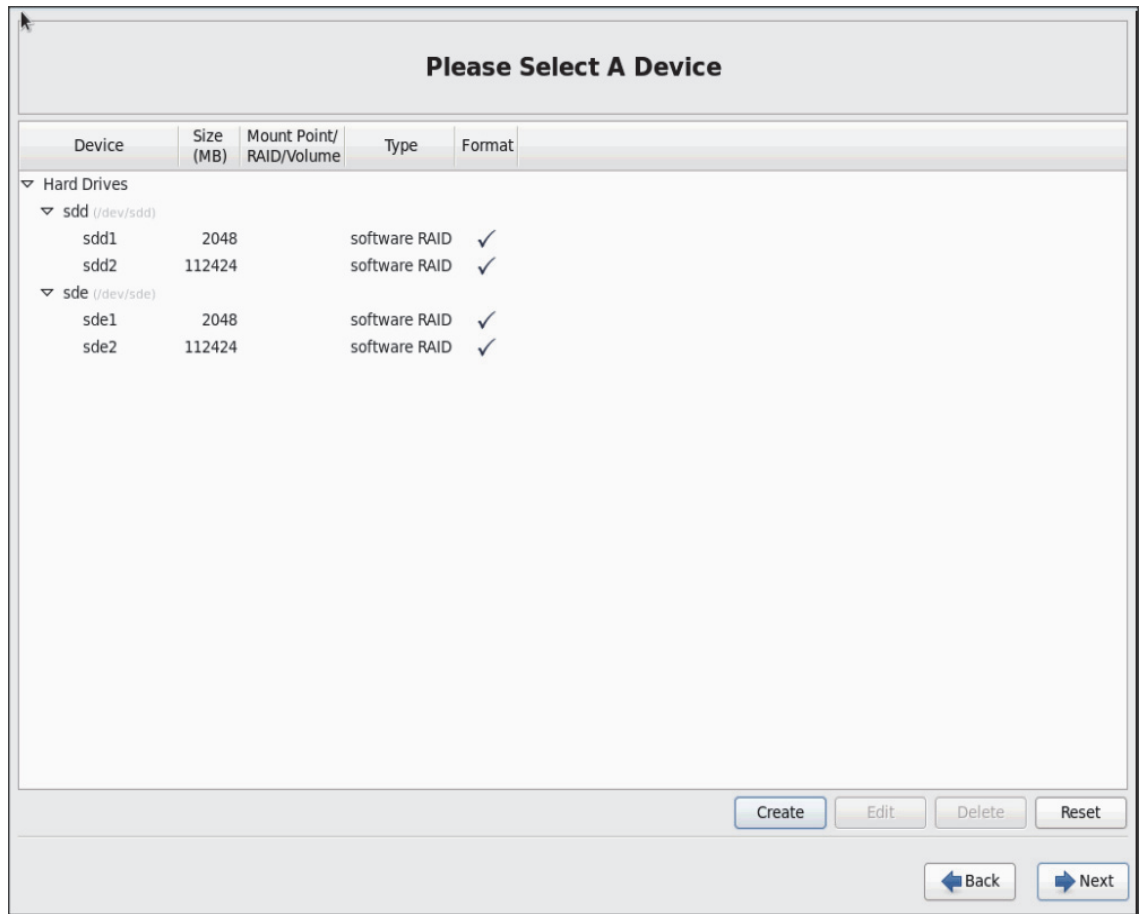


Figure 121 RHEL Installation: Add Partition

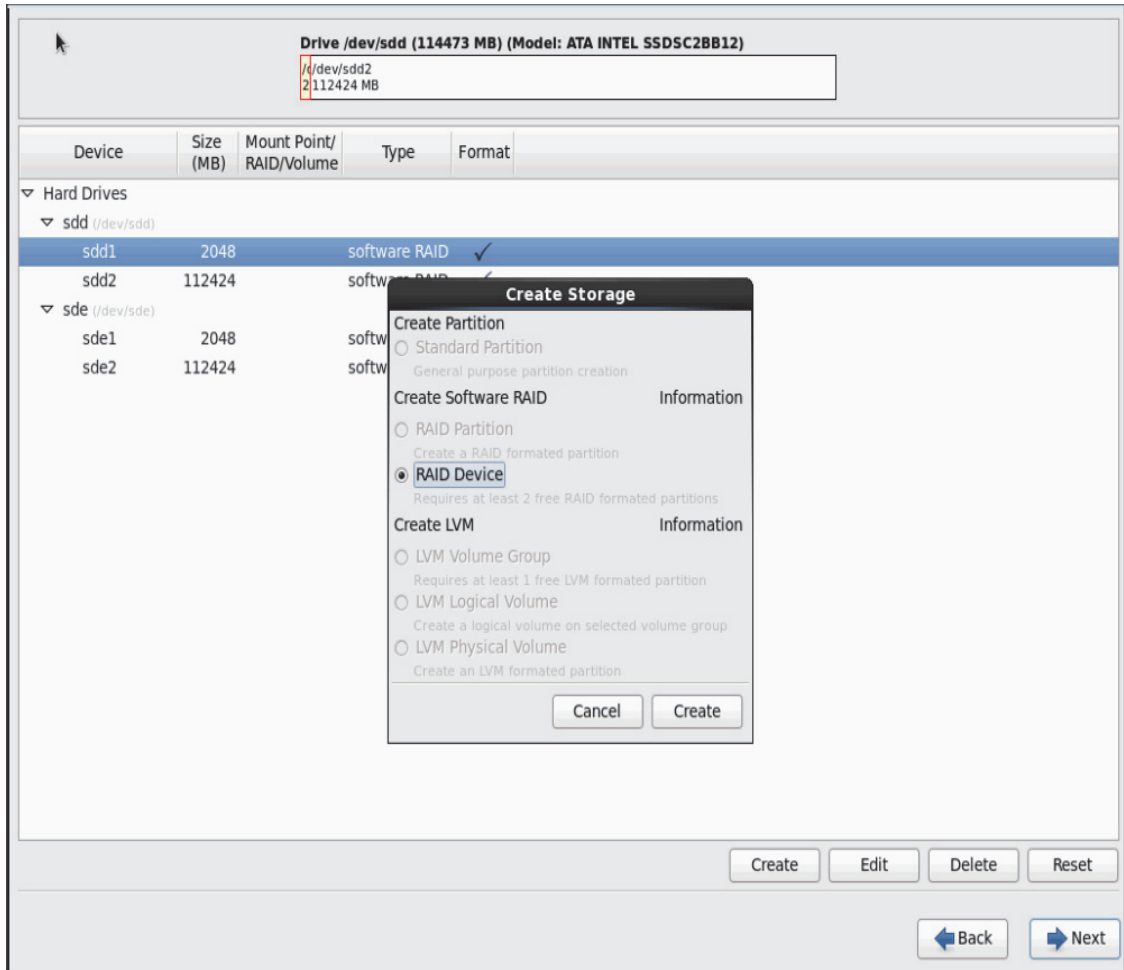


26. The above steps created 2 boot and 2 root (/) partitions. Following steps will RAID1 devices.

Figure 122 *RHEL Installation: Selected RAID Devices*

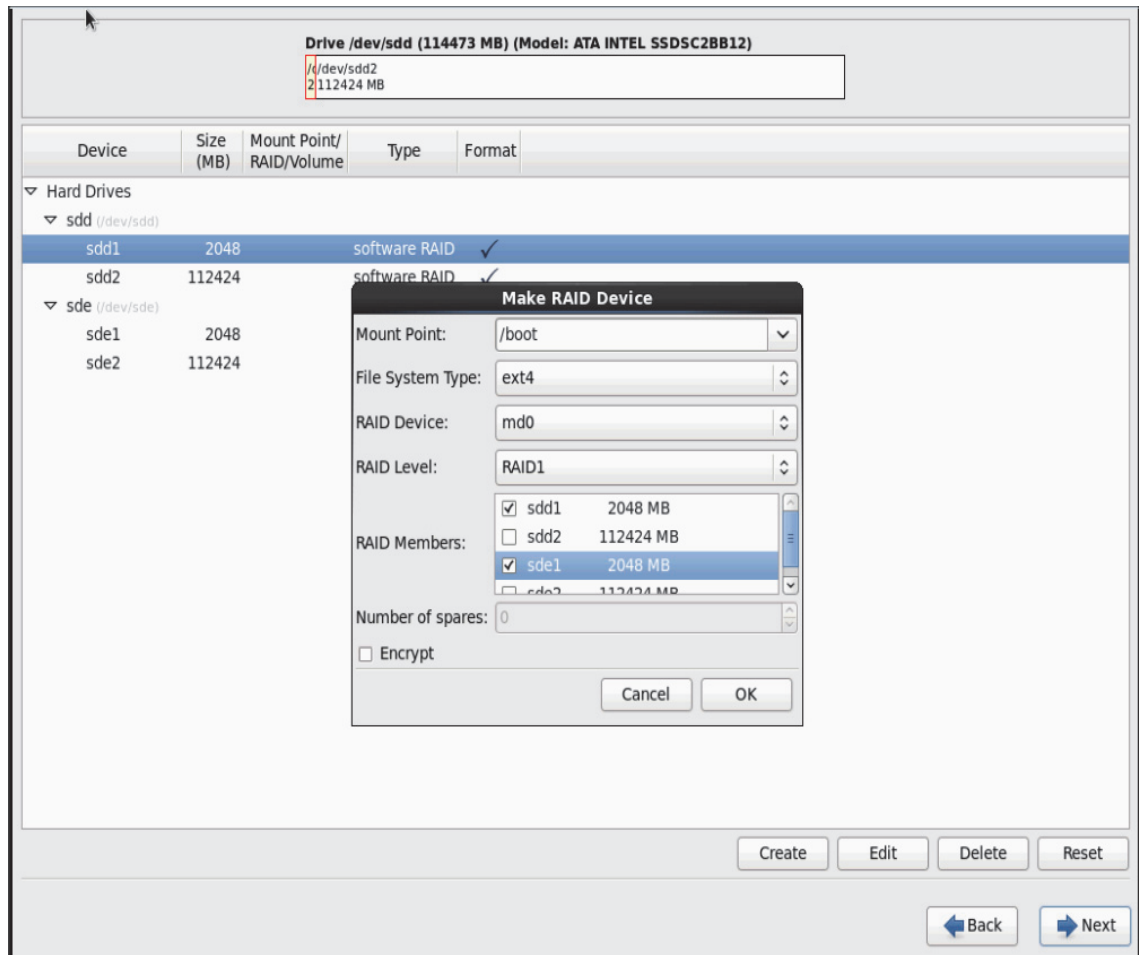
27. Choose one of the boot partitions and click on **Create > RAID Device**

Figure 123 RHEL Installation: Create RAID Device



28. Choose this as /boot (boot device) and in RAID members, choose all the boot partitions created above in order to create a software RAID 1 for boot

Figure 124 RHEL Installation: Make RAID Device



29. Similarly repeat for / partitions created above choosing both members with mount point as “/”.

Figure 125 RHEL Installation: Create RAID Device

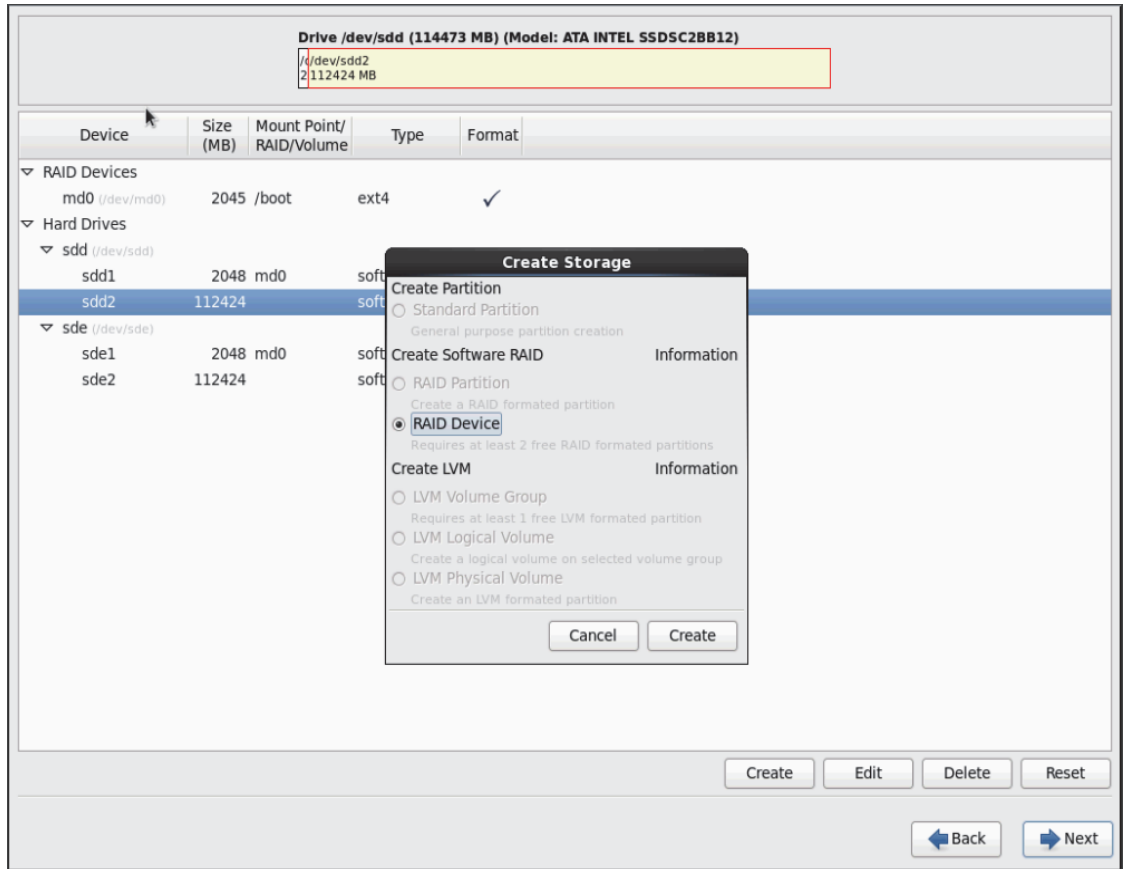


Figure 126 RHEL Installation: Make RAID Device

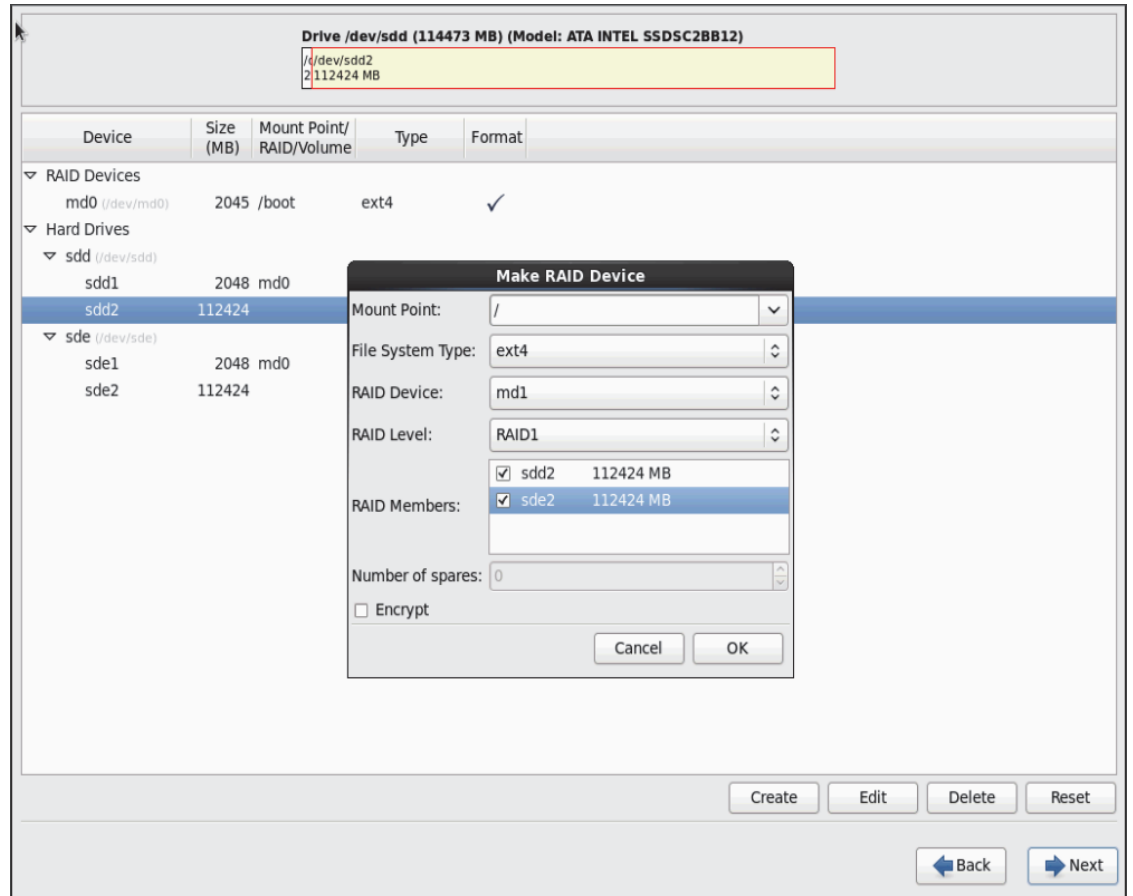
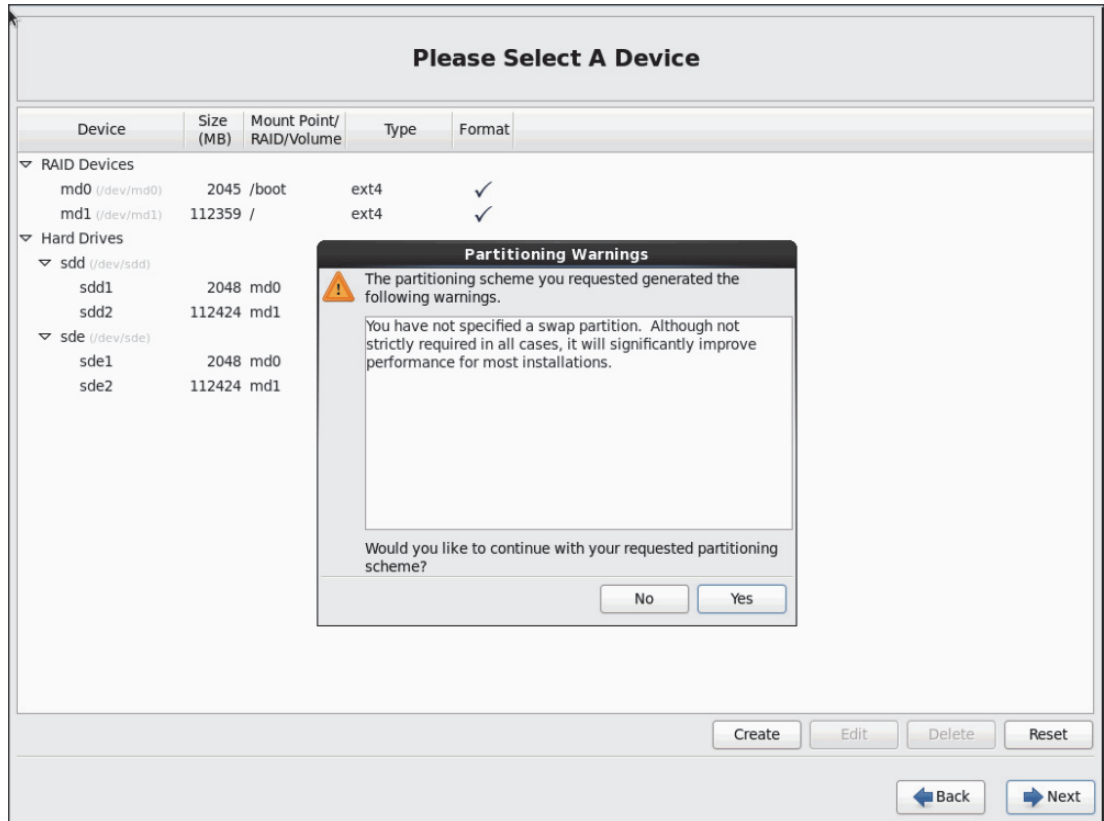


Figure 127 RHEL Installation: Selected RAID Devices



30. Click on Next.

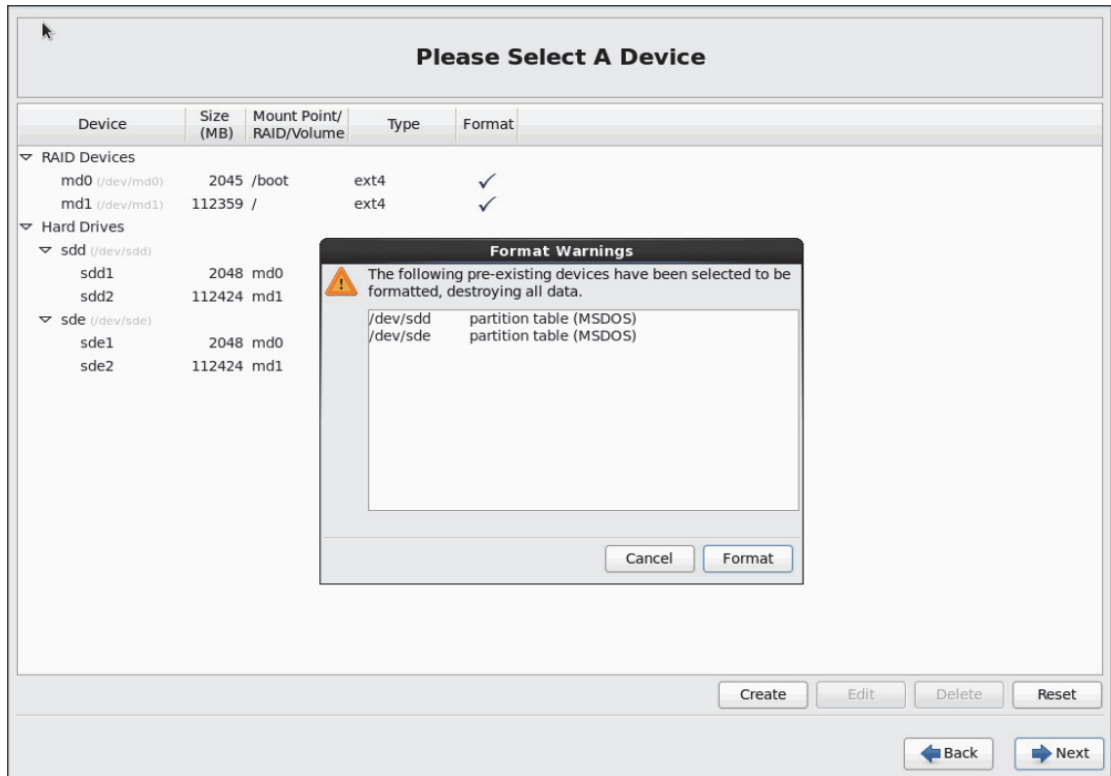
Figure 128 RHEL Installation: Partitioning Warning

**Note**

Swap partition can be created using the similar steps, however, since these systems are high in memory, this step is skipped (click **Yes**)

31. Click **Next**, and **Format**.

Figure 129 RHEL Installation: Format Warning



32. Select default settings and click Next.

Figure 130 RHEL Installation: Install Boot Loader

Install boot loader on /dev/sdd.

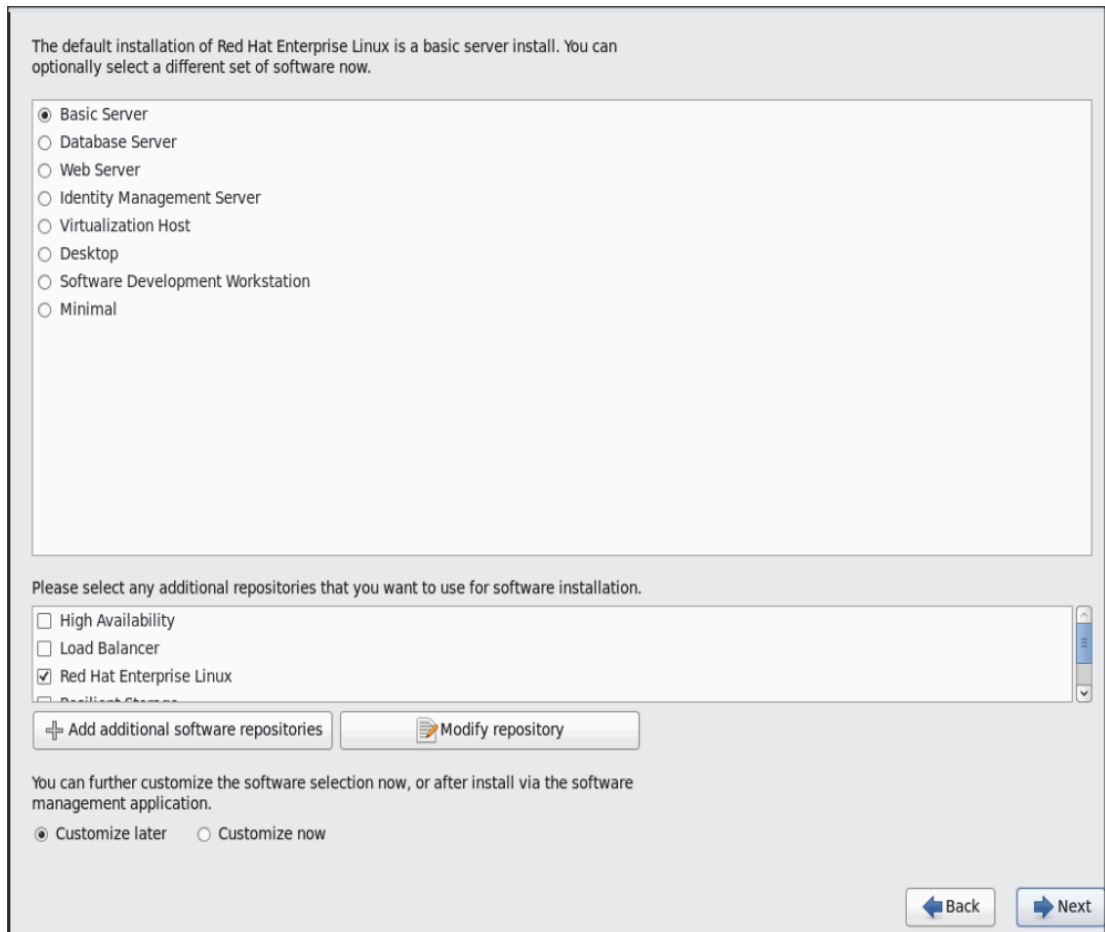
Use a boot loader password

Boot loader operating system list

Default	Label	Device
<input checked="" type="radio"/>	Red Hat Enterprise Linux	/dev/md1

33. Continue with RHEL Installation as shown below.

Figure 131 *RHEL Installation: Keep the Default Installation*



34. Once the installation is complete reboot the system.

Repeat the steps 1 through 34, to install Red Hat Enterprise Linux 6.5 on Servers 2 through 64.



Note

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

The host-names and their corresponding IP addresses are shown in Table 7.

Table 7 *Host-names and IP Addresses*

Hostname	eth0	eth1	eth2
rhel1	10.29.160.101	192.168.11.101	192.168.12.101
rhel2	10.29.160.102	192.168.11.102	192.168.12.102
rhel3	10.29.160.103	192.168.11.103	192.168.12.103
rhel4	10.29.160.104	192.168.11.104	192.168.12.104
rhel5	10.29.160.105	192.168.11.105	192.168.12.105
rhel6	10.29.160.106	192.168.11.106	192.168.12.106

Table 7 *Host-names and IP Addresses*

Hostname	eth0	eth1	eth2
rhel7	10.29.160.107	192.168.11.107	192.168.12.107
rhel8	10.29.160.108	192.168.11.108	192.168.12.108
rhel9	10.29.160.109	192.168.11.109	192.168.12.109
rhel10	10.29.160.110	192.168.11.110	192.168.12.110
rhel11	10.29.160.111	192.168.11.111	192.168.12.111
rhel12	10.29.160.112	192.168.11.112	192.168.12.112
rhel13	10.29.160.113	192.168.11.113	192.168.12.113
rhel14	10.29.160.114	192.168.11.114	192.168.12.114
rhel15	10.29.160.115	192.168.11.115	192.168.12.115
rhel16	10.29.160.116	192.168.11.116	192.168.12.116
rhel64	10.29.160.164	192.168.11.164	192.168.12.164
rhel65	10.29.160.165	192.168.11.165	NA
rhel66	10.29.160.166	192.168.11.166	NA
rhel67	10.29.160.167	192.168.11.167	NA
rhel68	10.29.160.168	192.168.11.168	NA

Post OS Install Configuration

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

Setting Up Password-less Login

To manage all of the clusters nodes from the admin node we need to setup password-less login. It assists in automating common tasks with cluster-shell (clush, a cluster wide parallel shell), and shell-scripts without having to use 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 10.29.160.101
```

2. Run the ssh-keygen command to create both public and private keys on the admin node.

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

```

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

```
for IP in {101..168}; do echo -n "$IP -> "; ssh-copy-id -i ~/.ssh/id_rsa.pub
10.29.160.$IP; done
```

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

Enter the password of the remote host.

Configuring /etc/hosts

Setup /etc/hosts on the Admin node and other nodes as follows; this is a pre-configuration to setup DNS as shown in the further section.

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

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

On Admin Node (rhel1)

```
vi /etc/hosts
127.0.0.1 local host localhost.localdomain localhost4 localhost4.localdomain4
::1 localhost localhost.localdomain localhost6 localhost6.localdomain6

10.29.160.101 rhel1.mgmt
10.29.160.102 rhel2.mgmt
10.29.160.103 rhel3.mgmt
10.29.160.104 rhel4.mgmt
10.29.160.105 rhel5.mgmt
10.29.160.106 rhel6.mgmt
10.29.160.107 rhel7.mgmt
10.29.160.108 rhel8.mgmt
10.29.160.109 rhel9.mgmt
10.29.160.110 rhel10.mgmt
10.29.160.111 rhel11.mgmt
```

```

10.29.160.112 rhel12.mgmt
10.29.160.113 rhel13.mgmt
10.29.160.114 rhel14.mgmt
10.29.160.115 rhel15.mgmt
10.29.160.116 rhel16.mgmt
...
10.29.160.168 rhel68.mgmt

192.168.11.101 rhel1
192.168.11.102 rhel2
192.168.11.103 rhel3
192.168.11.104 rhel4
192.168.11.105 rhel5
192.168.11.106 rhel6
192.168.11.107 rhel7
192.168.11.108 rhel8
192.168.11.109 rhel9
192.168.11.110 rhel10
192.168.11.111 rhel11
192.168.11.112 rhel12
192.168.11.113 rhel13
192.168.11.114 rhel14
192.168.11.115 rhel15
192.168.11.116 rhel16
...
192.168.11.168 rhel68

```

Setup ClusterShell

ClusterShell (or clush) is cluster wide shell to run commands on several hosts in parallel.

From the system connected to the Internet download Cluster shell (clush) and install it on rhel1. Cluster shell is available from EPEL (Extra Packages for Enterprise Linux) 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/
```

Login to rhel1 and install cluster shell

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

Edit `/etc/clustershell/groups` file to include host-names for all the nodes of the cluster. These set of hosts are taken when running clush with '-a' option

For 68 node cluster as in our CVD, set groups file as follows,

```
vi /etc/clustershell/groups
all: rhel[1-68]
```

```
[root@rhel1 ~]# vi /etc/clustershell/groups
[root@rhel1 ~]# cat /etc/clustershell/groups
all: rhel[1-68]
```



Note

For more information and documentation on ClusterShell, visit <https://github.com/cea-hpc/clustershell/wiki/UserAndProgrammingGuide>

**Note**

Clustershell will not work if not ssh to the machine earlier (as it requires to be in known_hosts file), for instance, as in the case below.

```
[root@Redhat-JB-R1 ~]# ssh rhel2
The authenticity of host 'rhel2 (10.0.127.52)' can't be established.
RSA key fingerprint is f2:0c:db:50:64:f1:9e:a6:7a:9d:c6:d4:8d:9e:e5:37.
Are you sure you want to continue connecting (yes/no)? █
```

```
[root@rhel1 ~]# ssh rhel5.mgmt
The authenticity of host 'rhel5.mgmt (10.29.160.105)' can't be established.
RSA key fingerprint is 7a:98:75:9a:6a:1a:80:a4:97:43:6c:8a:12:57:db:74.
Are you sure you want to continue connecting (yes/no)? █
```

Creating Red Hat Enterprise Linux (RHEL) 6.5 Local Repo

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. Log on to rhel1. Create a directory that would contain the repository.

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

2. Copy the contents of the Red Hat DVD to **/var/www/html/rhelrepo** directory.

3. Alternatively, if you have access to a Red Hat ISO Image, Copy the ISO file to rhel1.

```
scp rhel-server-6.5-x86_64-dvd.iso rhel1:/root/
```

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

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

```
mount -t iso9660 -o loop /root/rhel-server-6.5-x86_64-dvd.iso /mnt/rheliso/
```

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

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

```
[root@rhel1 ~]# mkdir -p /var/www/html/rhelrepo
[root@rhel1 ~]# mkdir -p /mnt/rheliso
[root@rhel1 ~]#
[root@rhel1 ~]# mount -t iso9660 -o loop /root/rhel-server-6.5-x86_64-dvd.iso /mnt/rheliso/
[root@rhel1 ~]# cp -r /mnt/rheliso/* /var/www/html/rhelrepo/
```

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

```
vi /var/www/html/rhelrepo/rheliso.repo
[rhel6.5]
name=Red Hat Enterprise Linux 6.5
baseurl=http://10.29.160.101/rhelrepo
gpgcheck=0
enabled=1
```

6. Now copy rheliso.repo file from **/var/www/html/rhelrepo** to **/etc/yum.repos.d** on rhel1

```
cp /var/www/html/rhelrepo/rheliso.repo /etc/yum.repos.d/
```

**Note**

Based on this repo file yum requires httpd to be running on rhel1 for other nodes to access the repository.

- Copy the **rheliso.repo** to all the nodes of the cluster.

```
clush -a -b -c /etc/yum.repos.d/rheliso.repo --dest=/etc/yum.repos.d/
```

```
[root@rhel1 ~]# clush -a -b -c /etc/yum.repos.d/rheliso.repo --dest=/etc/yum.repos.d/
```

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

**Note**

This step is needed to install software on Admin Node (rhel1) using the repo (such as httpd, createrepo, etc).

```
vi /etc/yum.repos.d/rheliso.repo
[rhel6.5]
name=Red Hat Enterprise Linux 6.5
baseurl=file:///var/www/html/rhelrepo
gpgcheck=0
enabled=1
```

- Creating the Red Hat Repository Database.

Install the createrepo package on admin node (rhel1). Use it to regenerate the repository database(s) for the local copy of the RHEL DVD contents.

```
yum -y install createrepo
```

```
[root@rhel1 ~]# yum -y install createrepo
Loaded plugins: product-id, refresh-packagekit, security, subscription-manager
This system is not registered to Red Hat Subscription Management. You can use subscription-manager to register.
rhel6.5 | 3.9 kB 00:00
rhel6.5/primary_db | 3.1 MB 00:00
Setting up Install Process
Resolving Dependencies
--> Running transaction check
--> Package createrepo.noarch 0:0.9.9-18.el6 will be installed
--> Processing Dependency: python-deltarpm for package: createrepo-0.9.9-18.el6.noarch
--> Running transaction check
--> Package python-deltarpm.x86_64 0:3.5-0.5.20090913git.el6 will be installed
--> Processing Dependency: deltarpm = 3.5-0.5.20090913git.el6 for package: python-deltarpm-3.5-0.5.20090913git.el6.x86_64
--> Running transaction check
```

- Run createrepo on the RHEL repository to create the repo database on admin node

```
cd /var/www/html/rhelrepo
createrepo .
```

```
[root@rhel1 rhelrepo]# createrepo .
Spawning worker 0 with 3763 pkgs
Workers Finished
Gathering worker results

Saving Primary metadata
Saving file lists metadata
Saving other metadata
Generating sqlite DBs
Sqlite DBs complete
```

11. Finally, purge the yum caches after httpd is installed (steps in section “Install Httpd”).

Configuring DNS

This section details setting up DNS using dnsmasq as an example based on the /etc/hosts configuration setup in the earlier section.

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

1. Disable Network manager on all nodes

```
clush -a -b service NetworkManager stop
clush -a -b chkconfig NetworkManager off
```

2. Update /etc/resolv.conf file to point to Admin Node

```
vi /etc/resolv.conf
nameserver 192.168.11.101
```



Note This step is needed if setting up dnsmasq on Admin node. Else this file should be updated with the correct nameserver.

3. Install and Start dnsmasq on Admin node

```
yum -y install dnsmasq
service dnsmasq start
chkconfig dnsmasq on
```

4. Deploy /etc/resolv.conf from the admin node (rhel1) to all the nodes via the following clush command:

```
clush -a -B -c /etc/resolv.conf
```



Note A clush copy without - --dest copies to the same directory location as the source-file directory.

5. Ensure DNS is working fine by running the following command on Admin node and any datanode

```
[root@rhel2 ~]# nslookup rhel1
Server:192.168.11.101
Address:192.168.11.101#53

Name: rhel1
Address: 192.168.11.101 •

[root@rhel2 ~]# nslookup rhel1.mgmt
Server: 192.168.11.101
```



```

Address: 192.168.11.101#53

Name: rhell.mgmt
Address: 10.29.160.101 •

[root@rhel2 ~]# nslookup 10.29.160.101
Server: 192.168.11.101
Address: 192.168.11.101#53

101.160.29.10.in-addr.arpa name = rhell.mgmt. •

```

Installing httpd

Setting up RHEL repo on the admin node requires httpd. This section describes the process of setting up one

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

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

```
yum -y install httpd
```

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

```
vi /etc/httpd/conf/httpd.conf
ServerName 10.29.160.101:80
```

```

[root@rhell ~]# vi /etc/httpd/conf/httpd.conf
[root@rhell ~]# cat /etc/httpd/conf/httpd.conf | grep ServerName
# ServerName gives the name and port that the server uses to identify itself.
#ServerName www.example.com:80
ServerName 10.29.160.101:80
# ServerName directive.
#   ServerName dummy-host.example.com

```

3. Start httpd

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

4. Purge the yum caches after httpd is installed (step followed from section Setup Red Hat Repo)

```
clush -a -B yum clean all
clush -a -B yum repolist
```

```

[root@rhell ~]# clush -a -B yum clean all
-----
rhel[1-17] (17)
-----
Loaded plugins: product-id, refresh-packagekit, security, subscription-manager
This system is not registered to Red Hat Subscription Management. You can use subscription-manager to register.
Cleaning repos: rhel6.5
Cleaning up Everything

```



Note

While suggested configuration is to disable SELinux as shown below, if for any reason SELinux needs to be enabled on the cluster, then ensure to run the following to make sure that the httpd is able to read the Yum repofiles `chcon -R -t httpd_sys_content_t /var/www/html/`

Upgrading Cisco Network driver for VIC1227

The latest Cisco Network driver is required for performance and updates. The latest drivers can be downloaded from the link below:

<https://software.cisco.com/download/release.html?mdfid=283862063&flowid=25886&softwareid=283853158&release=1.5.7d&relind=AVAILABLE&rellifecycle=&reltype=latest>

In the ISO image, the required driver `kmod-enic-2.1.1.66-rhel6u5.el6.x86_64.rpm` can be located at `\Linux\Network\Cisco\12x5x\RHEL\RHEL6.5`

From a node connected to the Internet, download, extract and transfer `kmod-enic-2.1.1.66-rhel6u5.el6.x86_64.rpm` to `rhell` (admin node).

Install the rpm on all nodes of the cluster using the following clush commands. For this example the rpm is assumed to be in present working directory of `rhell`.

```
[root@rhell ~]# clush -a -b -c kmod-enic-2.1.1.66-rhel6u5.el6.x86_64.rpm
[root@rhell ~]# clush -a -b "rpm -ivh kmod-enic-2.1.1.66-rhel6u5.el6.x86_64.rpm "
```

Ensure that the above installed version of `kmod-enic` driver is being used on all nodes by running the command “`modinfo enic`” on all nodes

```
[root@rhell ~]# clush -a -B "modinfo enic | head -5"
```

```
filename:      /lib/modules/2.6.32-431.el6.x86_64/extra/enic/enic.ko
version:      2.1.1.66
license:      GPL v2
author:       Scott Feldman <scofeldm@cisco.com>
description:  Cisco VIC Ethernet NIC Driver
```

Installing xfsprogs

From the admin node `rhell` run the command below to Install **xfsprogs** on all the nodes for xfs filesystem.

```
clush -a -B yum -y install xfsprogs
```

```
[root@rhell ~]# clush -a -B yum -y install xfsprogs
```

```

-----
rhel[1-17] (17)
-----
Loaded plugins: product-id, refresh-packagekit, security, subscription-manager
Setting up Install Process
Resolving Dependencies
--> Running transaction check
---> Package xfsprogs.x86_64 0:3.1.1-14.el6 will be installed
--> Finished Dependency Resolution

Dependencies Resolved

=====
Package                Arch          Version        Repository      Size
=====
Installing:
  xfsprogs             x86_64        3.1.1-14.el6   rhel6.5         724 k
=====
Transaction Summary
=====
Install      1 Package(s)

Total download size: 724 k
Installed size: 3.2 M
Downloading Packages:
Running rpm_check_debug
Running Transaction Test
Transaction Test Succeeded
Running Transaction
  Installing : xfsprogs-3.1.1-14.el6.x86_64                1/1
  Verifying  : xfsprogs-3.1.1-14.el6.x86_64                1/1

Installed:
  xfsprogs.x86_64 0:3.1.1-14.el6

Complete!

```

Setting up JAVA

HDP 2.2 requires JAVA 7, download `jdk-7u75-linux-x64.rpm` from [oracle.com](http://www.oracle.com/technetwork/java/javase/downloads/jdk7-downloads-1880260.html) (<http://www.oracle.com/technetwork/java/javase/downloads/jdk7-downloads-1880260.html>) to admin node (rhell).

Create the following files `java-set-alternatives.sh` and `java-home.sh` on admin node (rhell)

vi `java-set-alternatives.sh`

```

#!/bin/bash
for item in java javac javaws jar jps javah javap jcontrol jconsole jdb; do
  rm -f /var/lib/alternatives/$item
  alternatives --install /usr/bin/$item $item /usr/java/jdk1.7.0_75/bin/$item 9
  alternatives --set $item /usr/java/jdk1.7.0_75/bin/$item
done

```

vi `java-home.sh`

```

export JAVA_HOME=/usr/java/jdk1.7.0_75

```

Run the following commands on admin node (rhel1) to install and setup java on all nodes

1. Copying JDK rpm to all nodes

```
clush -b -a -c /root/jdk-7u75-linux-x64.rpm --dest=/root/
```

2. Make the two java scripts created above executable

```
chmod 755 ./java-set-alternatives.sh ./java-home.sh
```

3. Copying java-set-alternatives.sh to all nodes

```
clush -b -a -c ./java-set-alternatives.sh --dest=/root/
```

4. Extract and Install JDK on all nodes

```
clush -a -b rpm -ivh /root/jdk-7u75-linux-x64.rpm
```

5. Setup Java Alternatives

```
clush -b -a ./java-set-alternatives.sh
```

6. Ensure correct java is setup on all nodes (should point to newly installed java path)

```
clush -b -a "alternatives --display java | head -2"
```

7. Setup JAVA_HOME on all nodes

```
clush -b -a -c ./java-home.sh --dest=/etc/profile.d
```

8. Display JAVA_HOME on all nodes

```
clush -a -b "echo \$JAVA_HOME"
```

9. Display current java -version

```
clush -B -a java -version
```

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 synchronism 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 will occur with HBase and other services.

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

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
```

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

```
vi /root/ntp.conf
server 10.29.160.101
driftfile /var/lib/ntp/drift
restrict 127.0.0.1
restrict -6 ::1
includefile /etc/ntp/crypto/pw
keys /etc/ntp/keys
```

Copy `ntp.conf` file from the admin node to `/etc` of all the nodes by executing the following command in the admin node (rhel1)

```
for SERVER in {102..168}; do scp /root/ntp.conf
10.29.160.$SERVER:/etc/ntp.conf; done
```

```
[root@rhell ~]# for SERVER in {102..168}; do scp /root/ntp.conf 10.29.160.$SERVER:/etc/ntp.conf; done
ntp.conf 100% 136 0.1KB/s 00:00
ntp.conf 100% 136 0.1KB/s 00:00
ntp.conf 100% 136 0.1KB/s 00:00
ntp.conf 100% 136 0.1KB/s 00:00
ntp.conf 100% 136 0.1KB/s 00:00
ntp.conf 100% 136 0.1KB/s 00:00
ntp.conf 100% 136 0.1KB/s 00:00
ntp.conf 100% 136 0.1KB/s 00:00
ntp.conf 100% 136 0.1KB/s 00:00
ntp.conf 100% 136 0.1KB/s 00:00
ntp.conf 100% 136 0.1KB/s 00:00
ntp.conf 100% 136 0.1KB/s 00:00
ntp.conf 100% 136 0.1KB/s 00:00
ntp.conf 100% 136 0.1KB/s 00:00
ntp.conf 100% 136 0.1KB/s 00:00
ntp.conf 100% 136 0.1KB/s 00:00
ntp.conf 100% 136 0.1KB/s 00:00
ntp.conf 100% 136 0.1KB/s 00:00
ntp.conf 100% 136 0.1KB/s 00:00
ntp.conf 100% 136 0.1KB/s 00:00
```

**Note**

Instead of the above for loop, this could be run as a clush command with “-w” option.

```
clush -w rhel[2-68] -b -c /root/ntp.conf --dest=/etc
```

Do not use `clush -a -b -c /root/ntp.conf --dest=/etc` command as it overwrites `/etc/ntp.conf` on the admin node.

Run the following to synchronize the time and restart NTP daemon on all nodes

```
clush -a -B "yum install -y ntpdate"
clush -a -b "service ntpd stop"
clush -a -b "ntpdate rhell"
clush -a -b "service ntpd start"
```

Ensure restart of NTP daemon across reboots

```
clush -a -b "chkconfig ntpd on"
```

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 be sure that a syslog daemon is present. One of the following commands should suffice to confirm that the service is properly configured:

```
clush -B -a rsyslogd -v
```

```
[root@rhell1 ~]# clush -B -a rsyslogd -v
-----
rhel[1-17] (17)
-----
rsyslogd 5.8.10, compiled with:
  FEATURE_REGEX:                Yes
  FEATURE_LARGEFILE:             No
  GSSAPI Kerberos 5 support:     Yes
  FEATURE_DEBUG (debug build, slow code): No
  32bit Atomic operations supported: Yes
  64bit Atomic operations supported: Yes
  Runtime Instrumentation (slow code): No

See http://www.rsyslog.com for more information.
```

```
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.

For setting ulimit on Redhat, edit `/etc/security/limits.conf` on admin node rhell1 and add the following lines:

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

```
[root@rhell1 ~]# cat /etc/security/limits.conf | grep 64000
root soft nofile 64000
root hard nofile 64000
```

Copy the `/etc/security/limits.conf` file from admin node (rhell1) to all the nodes using the following command.

```
clush -a -b -c /etc/security/limits.conf --dest=/etc/security/
```

```
[root@rhell1 ~]# clush -a -b -c /etc/security/limits.conf --dest=/etc/security/
```

Verify the **ulimit** setting with the following steps:



Note

Ulimit values are applied on a new shell, running the command on a node on an earlier instance of a shell will show old values

Run the following command at a command line. The command should report 64000.

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

Disabling SELinux

SELinux must be disabled during the install 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. The following command will disable SELINUX on all nodes.

```
clush -a -b "sed -i 's/SELINUX=enforcing/SELINUX=disabled/g' /etc/selinux/config "
clush -a -b "setenforce 0"
```

```
[root@rhel1 ~]# clush -a -b "sed -i 's/SELINUX=enforcing/SELINUX=disabled/g' /etc/selinux/config "
```



Note

The above command may fail if SELinux is already disabled.

Set TCP Retries

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 most likely 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 on admin node `rhel1` and add the following lines:

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

Copy the `/etc/sysctl.conf` file from admin node (`rhel1`) to all the nodes using the following command.

```
clush -a -b -c /etc/sysctl.conf --dest=/etc/
```

2. Load the settings from default `sysctl` file `/etc/sysctl.conf` by running the below 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 UCS Big Data deployment will be in its own isolated network, there's no need to leave the IP tables service running.

```
clush -a -b "service iptables stop"
clush -a -b "chkconfig iptables off"
```

```
[root@rhel1 ~]# clush -a -b "service iptables stop"
[root@rhel1 ~]# clush -a -b "chkconfig iptables off"
```

Disable Swapping

In order to reduce Swapping, run the following on all nodes. Variable `vm.swappiness` defines how often swap should be used. 0 is No Swapping, 60 is the default value.

```
clush -a -b " echo '\vm.swappiness=0\' >> /etc/sysctl.conf"
```

Load the settings from default `sysctl` file `/etc/sysctl.conf`

```
clush -a -b "sysctl -p"
```

Disable Transparent Huge Pages

Disabling Transparent Huge Pages (THP) reduces elevated CPU usage caused by THP. From the admin node, run the following commands

```
clush -a -b "echo never >
/sys/kernel/mm/redhat_transparent_hugepage/enabled"
clush -a -b "echo never >
/sys/kernel/mm/redhat_transparent_hugepage/defrag"
```

The above command needs to be run for every reboot, hence, copy this command to `/etc/rc.local` so they are executed automatically for every reboot.

On Admin node, run the following commands

```
rm -f /root/thp_disable
echo "echo never > /sys/kernel/mm/redhat_transparent_hugepage/enabled" >>
/root/thp_disable
echo "echo never > /sys/kernel/mm/redhat_transparent_hugepage/defrag" >>
/root/thp_disable
```

Copy file to each node

```
clush -a -b -c /root/thp_disable
```

Append the content of file `thp_disable` to `/etc/rc.local`

```
clush -a -b "cat /root/thp_disable >> /etc/rc.local"
```

Install Openssl

Install Openssl and Openssl-devel version 1.0.1e-30 and above for RHEL6.5. This is a requirement for HDP 2.2 on all nodes. If openssl is already installed (generally the case), use the following command to upgrade openssl

```
clush -a -b -c /root/openssl-*
clush -a -b rpm -Uvh openssl-1.0.1e-*.rpm openssl-devel-1.0.1e-*.rpm
```

```
[root@rhel1 ~]# rpm -Uvh openssl-1.0.1e-30.el6_6.5.x86_64.rpm openssl-devel-1.0.1e-30.el6_6.5.x86_64.rpm
warning: openssl-1.0.1e-30.el6_6.5.x86_64.rpm: Header V3 RSA/SHA1 Signature, key ID c105b9de: NOKEY
Preparing... ##### [100%]
 1:openssl ##### [ 50%]
 2:openssl-devel ##### [100%]
```

(RPMs are available at:

http://mirror.centos.org/centos/6/updates/x86_64/Packages/openssl-1.0.1e-30.el6_6.5.x86_64.rpm and

http://mirror.centos.org/centos/6/updates/x86_64/Packages/openssl-devel-1.0.1e-30.el6_6.5.x86_64.rpm)



Note

This requires `krb5-devel` and `zlib-devel` as dependencies. If not installed, install it as follows on the nodes throwing error “`yum -y install krb5-devel zlib-devel`”

Disable IPv6 Defaults

Disable IPv6 as the addresses used are IPv4.

```
clush -a -b "echo '\net.ipv6.conf.all.disable_ipv6 = 1\' >> /etc/sysctl.conf"
```



```

clush -a -b "echo \'net.ipv6.conf.default.disable_ipv6 = 1\' >> /etc/sysctl.conf"
clush -a -b "echo \'net.ipv6.conf.lo.disable_ipv6 = 1\' >> /etc/sysctl.conf"
Load the settings from default sysctl file /etc/sysctl.conf

clush -a -b "sysctl -p"

```

Configuring Data Drives on Name Node

This section describes steps to configure non-OS disk drives as RAID1 using StorCli command as described below. All the drives are going to be part of a single RAID1 volume. This volume can be used for Staging any client data to be loaded to HDFS. This volume won't be used for HDFS data.

From the website download storcli:

http://www.lsi.com/downloads/Public/RAID%20Controllers/RAID%20Controllers%20Common%20Files/1.14.12_StorCLI.zip

Extract the zip file and copy storcli-1.14.12-1.noarch.rpm from the linux directory.

1. Download storcli and its dependencies and transfer to Admin node.

```
scp storcli-1.14.12-1.noarch.rpm rhell:/root/
```

2. Copy storcli rpm to all the nodes using the following commands:

```
clush -a -b -c /root/storcli-1.14.12-1.noarch.rpm --dest=/root/
```

3. Run the below command to install storcli on all the nodes

```
clush -a -b rpm -ivh storcli-1.14.12-1.noarch.rpm
```

4. Run the below command to copy storcli64 to root directory.

```
cd /opt/MegaRAID/storcli/
cp storcli64 /root/
```

```

[root@rhell ~]# cd /opt/MegaRAID/storcli/
[root@rhell storcli]# ls
install.log  libstorelibir-2.so  libstorelibir-2.so.14.07-0  storcli64
[root@rhell storcli]# cp storcli64 /root/

```

5. Copy storcli64 to all the nodes using the following commands:

```
clush -a -b -c /root/storcli64 --dest=/root/
```

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

```

vi /root/raid1.sh
./storcli64 -cfgldadd
r1[$1:1,$1:2,$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.
./storcli64 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 existing configurations refer to Embedded MegaRAID Software Users Guide available at www.lsi.com

Configuring Data Drives on Data Nodes

This section describes steps to configure non-OS disk drives as individual RAID0 volumes using StorCli command as described below. These volumes are going to be used for HDFS Data.

Issue the following command from the admin node to create the virtual drives with individual RAID 0 configurations on all the datanodes.

```
clush -w rhel[3-64] -B ./storcli64 -cfgeachdskraid0 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 command above will not override existing configurations. To clear and reconfigure existing configurations refer to Embedded MegaRAID Software Users Guide available at www.lsi.com

Configuring Data Drives on Archival Nodes

This section describes steps to configure non-OS disk drives as 4 RAID5 volumes using StorCli command as described below. These volumes are going to be used for HDFS Archival data.

1. Run the following script as root user on Archival Nodes to create the virtual drives.

```
vi /root/raid5.sh
./storcli64 /c0 add vd type=raid5 drives=$1:1-15 WB ra direct Strip=1024
./storcli64 /c0 add vd type=raid5 drives=$1:16-30 WB ra direct Strip=1024
./storcli64 /c0 add vd type=raid5 drives=$1:31-45 WB ra direct Strip=1024
./storcli64 /c0 add vd type=raid5 drives=$1:46-60 WB ra direct Strip=1024
```

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

```
./storcli64 pdlist -a0 | grep Enc | grep -v 252 | awk '{print $4}' | sort | uniq
-c | awk '{print $2}'
chmod 755 raid5.sh
```

- Run MegaCli script as follows

```
./raid5.sh <EnclosureID> obtained by running the command above
```

WB: Write back

RA: Read Ahead

Strpsz1024: Strip Size of 1024K

**Note**

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

Configuring the Filesystem for NameNodes, DataNodes and Archival nodes

The following script will format and mount the available volumes on each node whether it is namenode, Data node or Archival node. OS boot partition is going to be skipped. All drives are going to be mounted based on their UUID as /data/disk1, /data/disk2, and so on.

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

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

**Note**

The script assumes there are no partitions already existing on the data volumes. If there are partitions, then they have to be deleted first before running this script. This process is documented in the “Note” section at the end of the section

```
vi /root/driveconf.sh
#!/bin/bash
#Commented because the script intermittently fails on some occasions
[[ "-x" == "${1}" ]] && set -x && set -v && shift 1
count=1
for X in $(ls /dev/disk/by-id/scsi-*)
do
echo "$X considered"
D=${X##*/}
Y=${D:5}
if [[ -b ${X} && `sbin/parted -s ${X} print quit|bin/grep -c boot` -ne 0 ]]
then
echo "$X bootable - skipping."
continue
elif [[ ${Y} =~ SATA_INTEL_SSD* ]]
then
echo "$X bootable partition skipping"
else
echo "$X for formating"
/sbin/parted -s ${X} mklabel gpt quit -s
/sbin/parted -s ${X} mkpart 1 6144s 100% quit
#Identify drive mapping in /dev/sd*
drive=`ls -l ${X} | cut -d " " -f11 | cut -d "/" -f3`
drive_map="/dev/${drive}"
/sbin/mkfs.xfs -f -q -l size=65536b,lazy-count=1,su=256k -d sunit=1024,swidth=6144 -r
extsize=256k -L ${drive}1 ${drive_map}1
(( $? )) && continue
#Identify UUID
UUID=`blkid ${drive_map}1 | cut -d " " -f3 | cut -d "=" -f2 | sed 's/"//g'`
echo "UUID of ${drive_map}1 = ${UUID}"
/bin/mkdir -p /data/disk${count}
(( $? )) && continue
/bin/mount -t xfs -o allocsize=128m,noatime,nobarrier,nodiratime -U ${UUID}
/data/disk${count}
```

```
(( $? )) && continue
echo "UUID=${UUID} /data/disk${count} xfs allocsize=128m,noatime,nobarrier,nodiratime
0 0" >> /etc/fstab
((count++))
fi
done
```

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

```
chmod 755 /root/driveconf.sh
clush -a -B -c /root/driveconf.sh
```

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

```
clush -a -B /root/driveconf.sh
```

4. Run the following from the admin node to list the partitions and mount points

```
clush -a -B df -h
clush -a -B mount
clush -a -B cat /etc/fstab
```



Note

In-case there is need to delete any partitions, it can be done so using the following. Run command 'mount' to identify which drive is mounted to which device /dev/sd<?> umount the drive for which partition is to be deleted and run fdisk to delete as shown below.

Care to be taken **not to delete OS partition** as this will wipe out OS

```
mount
umount /data/disk1 # <-- disk1 shown as example
(echo d; echo w;) | sudo fdisk /dev/sd<?>
```

Cluster Verification

The section describes the steps to create the script cluster_verification.sh that helps to verify CPU, memory, NIC, storage adapter settings across the cluster on all nodes. This script also checks additional prerequisites such as NTP status, SELinux status, ulimit settings, JAVA_HOME settings and JDK version, IP address and hostname resolution, Linux version and firewall settings.

Create script cluster_verification.sh as follows on the Admin node (rhel1).

```
vi cluster_verification.sh
#!/bin/bash
shopt -s expand_aliases
# Setting Color codes
green='\e[0;32m'
red='\e[0;31m'
NC='\e[0m' # No Color
echo -e "${green} === Cisco UCS Integrated Infrastructure for Big Data \ Cluster
Verification === ${NC}"
echo ""
echo ""
echo -e "${green} ==== System Information ==== ${NC}"
echo ""
echo ""
echo -e "${green}System ${NC}"
clush -a -B " `which dmidecode` |grep -A2 '^System Information'"
echo ""
echo ""
echo -e "${green}BIOS ${NC}"
clush -a -B " `which dmidecode` | grep -A3 '^BIOS I'"
echo ""
```

```

echo ""
echo -e "${green}Memory ${NC}"
clush -a -B "cat /proc/meminfo | grep -i ^memt | uniq"
echo ""
echo ""
echo -e "${green}Number of Dimms ${NC}"
clush -a -B "echo -n 'DIMM slots: '; `which dmidecode` |grep -c \
'^[[:space:]]*Locator:'"
clush -a -B "echo -n 'DIMM count is: '; `which dmidecode` | grep \ "Size"| grep -c
"MB""
clush -a -B " `which dmidecode` | awk '/Memory Device$/|^$/ {print}' |\ grep -e
'^Mem' -e Size: -e Speed: -e Part | sort -u | grep -v -e 'NO \ DIMM' -e 'No Module
Installed' -e Unknown"
echo ""
echo ""
# probe for cpu info #
echo -e "${green}CPU ${NC}"
clush -a -B "grep '^model name' /proc/cpuinfo | sort -u"
echo ""
clush -a -B "`which lscpu` | grep -v -e op-mode -e ^Vendor -e family -e \ Model: -e
Stepping: -e BogoMIPS -e Virtual -e ^Byte -e ^NUMA node(s)'"
echo ""
echo ""
# probe for nic info #
echo -e "${green}NIC ${NC}"
clush -a -B "`which ifconfig` | egrep '(^|p)' | awk '{print \$1}' | \ xargs -l
`which ethtool` | grep -e ^Settings -e Speed"
echo ""
clush -a -B "`which lspci` | grep -i ether"
echo ""
echo ""
# probe for disk info #
echo -e "${green}Storage ${NC}"
clush -a -B "echo 'Storage Controller: '; `which lspci` | grep -i -e \ raid -e storage
-e lsi"
echo ""
clush -a -B "dmesg | grep -i raid | grep -i scsi"
echo ""
clush -a -B "lsblk -id | awk '{print \$1,\$4}'|sort | nl"
echo ""
echo ""
echo -e "${green} ===== Software ===== ${NC}"
echo ""
echo ""
echo -e "${green}Linux Release ${NC}"
clush -a -B "cat /etc/*release | uniq"
echo ""
echo ""
echo -e "${green}Linux Version ${NC}"
clush -a -B "uname -srvm | fmt"
echo ""
echo ""
echo -e "${green}Date ${NC}"
clush -a -B date
echo ""
echo ""
echo -e "${green}NTP Status ${NC}"
clush -a -B "ntpstat 2>&1 | head -1"
echo ""
echo ""
echo -e "${green}SELINUX ${NC}"
clush -a -B "echo -n 'SElinux status: '; grep ^SELINUX= \ /etc/selinux/config 2>&1"
echo ""
echo ""

```

```

echo -e "${green}IPTables ${NC}"
clush -a -B "`which chkconfig` --list iptables 2>&1"
echo ""
clush -a -B "`which service` iptables status 2>&1 | head -10"
echo ""
echo ""
echo -e "${green}Transparent Huge Pages ${NC}"
clush -a -B " cat /sys/kernel/mm/*transparent_hugepage/enabled"
echo ""
echo ""
echo -e "${green}CPU Speed${NC}"
clush -a -B "echo -n 'CPUspeed Service: '; `which service` cpuspeed \ status 2>&1"
clush -a -B "echo -n 'CPUspeed Service: '; `which chkconfig` --list \ cpuspeed 2>&1"
echo ""
echo ""
echo -e "${green}Java Version${NC}"
clush -a -B 'java -version 2>&1; echo JAVA_HOME is ${JAVA_HOME:-Not \ Defined!}'
echo ""
echo ""
echo -e "${green}Hostname Lookup${NC}"
clush -a -B " ip addr show"
echo ""
echo ""
echo -e "${green}Open File Limit${NC}"
clush -a -B 'echo -n "Open file limit(should be >32K): "; ulimit -n'

```

Change permissions to executable

```
chmod 755 cluster_verification.sh
```


Run the Cluster Verification tool from the admin node. This can be run before starting HDP 2.2 to identify any discrepancies in Post OS Configuration between the servers or during troubleshooting of any cluster / Hadoop issues.

```
./cluster_verification.sh
```

Installing HDP 2.2

HDP is an enterprise grade, hardened Hadoop distribution. HDP combines Apache Hadoop and its related projects into a single tested and certified package. HDP 2.2 includes more than a hundred new features and closes thousands of issues across [Apache Hadoop](#) and its related projects with the testing and quality expected from enterprise quality software. HDP 2.2 components are depicted in figure below. The following sections go in detail on how to install HDP 2.2 on the cluster configured as shown in the earlier sections.

Figure 132 HDP Components



Pre-Requisites for HDP Installation

This section details the pre-requisites for HDP Installation such as setting up of HDP Repositories.

Hortonworks Repo

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

```
mkdir -p /tmp/Hortonworks
cd /tmp/Hortonworks/
```

1. Download Hortonworks HDP Repo

```
wget http://public-repo-1.hortonworks.com/HDP/centos6/HDP-2.2.0.0-centos6-rpm.tar.gz
```

```
[root@Srv1 Hortonworks]# wget http://public-repo-1.hortonworks.com/HDP/centos6/HDP-2.2.0.0-centos6-rpm.tar.gz
--2015-03-06 17:02:05-- http://public-repo-1.hortonworks.com/HDP/centos6/HDP-2.2.0.0-centos6-rpm.tar.gz
Resolving public-repo-1.hortonworks.com... 54.192.118.226, 54.230.118.137, 54.230.116.98, ...
Connecting to public-repo-1.hortonworks.com|54.192.118.226|:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 3260497490 (3.0G) [application/x-tar]
Saving to: âHDP-2.2.0.0-centos6-rpm.tar.gzâ

7% [====>] 252,892,909 11.2M/s
```

2. Download Hortonworks HDP-Utils Repo

```
wget
http://public-repo-1.hortonworks.com/HDP-UTILS-1.1.0.20/repos/centos6/HDP-UTILS-1.1.0.
20-centos6.tar.gz
```

```
[root@Srv1 Hortonworks]# wget http://public-repo-1.hortonworks.com/HDP-UTILS-1.1.0.20/repos/centos6/HDP-UTILS-1.1.0.20-centos6.tar.gz
--2015-03-06 17:04:09-- http://public-repo-1.hortonworks.com/HDP-UTILS-1.1.0.20/repos/centos6/HDP-UTILS-1.1.0.20-centos6.tar.gz
Resolving public-repo-1.hortonworks.com... 54.230.119.106, 54.239.132.164, 54.239.132.162, ...
Connecting to public-repo-1.hortonworks.com|54.230.119.106|:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 14076549 (13M) [application/x-tar]
Saving to: âHDP-UTILS-1.1.0.20-centos6.tar.gzâ

85% [=====>] 11,981,127 3.05M/s eta 1s
```

3. Download Ambari Repo

```
wget http://public-repo-1.hortonworks.com/ambari/centos6/ambari-1.7.0-centos6.tar.gz
```

```
[root@Srv1 Hortonworks]# wget http://public-repo-1.hortonworks.com/ambari/centos6/ambari-1.7.0-centos6.tar.gz
--2015-03-06 17:05:23-- http://public-repo-1.hortonworks.com/ambari/centos6/ambari-1.7.0-centos6.tar.gz
Resolving public-repo-1.hortonworks.com... 54.192.118.219, 54.192.118.224, 54.230.118.187, ...
Connecting to public-repo-1.hortonworks.com|54.192.118.219|:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 103219329 (98M) [application/x-tar]
Saving to: âambari-1.7.0-centos6.tar.gzâ

8% [====>] 9,123,154 2.18M/s eta 58s
```

4. Copy the repository directory to the admin node

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

5. Extract the files

login to rhell

```
cd /var/www/html/Hortonworks
tar -zxvf HDP-2.2.0.0-centos6-rpm.tar.gz
tar -zxvf HDP-UTILS-1.1.0.20-centos6.tar.gz
tar -zxvf ambari-1.7.0-centos6.tar.gz
```

6. Create the hdp.repo file with following contents

```
vi /etc/yum.repos.d/hdp.repo
[HDP-2.2.0.0]
name=Hortonworks Data Platform Version - HDP-2.2.0.0
baseurl=http://rhell1/Hortonworks/HDP/centos6/2.x/GA/2.2.0.0
gpgcheck=0
enabled=1
priority=1
[HDP-UTILS-1.1.0.20]
name=Hortonworks Data Platform Utils Version - HDP-UTILS-1.1.0.20
baseurl=http://rhell1/Hortonworks/HDP-UTILS-1.1.0.20/repos/centos6
gpgcheck=0
enabled=1
priority=1
```



```
[root@rhell ~]# vi /etc/yum.repos.d/hdp.repo
[root@rhell ~]# cat /etc/yum.repos.d/hdp.repo
[HDP-2.2.0.0]
name=Hortonworks Data Platform Version - HDP-2.2.0.0
baseurl=http://rhell/Hortonworks/HDP/centos6/2.x/GA/2.2.0.0
gpgcheck=0
enabled=1
priority=1

[HDP-UTILS-1.1.0.20]
name=Hortonworks Data Platform Utils Version - HDP-UTILS-1.1.0.20
baseurl=http://rhell/Hortonworks/HDP-UTILS-1.1.0.20/repos/centos6
gpgcheck=0
enabled=1
priority=1
```

7. Create the Ambari repo file with following contents

```
vi /etc/yum.repos.d/ambari.repo
[Updates-ambari-1.7.0]
name=ambari-1.7.0 - Updates
baseurl=http://rhell/Hortonworks/ambari/centos6/1.x/updates/1.7.0
gpgcheck=0
enabled=1
priority=1
```

```
[root@rhell ~]# vi /etc/yum.repos.d/ambari.repo
[root@rhell ~]# cat /etc/yum.repos.d/ambari.repo
[Updates-ambari-1.7.0]
name=ambari-1.7.0 - Updates
baseurl=http://rhell/Hortonworks/ambari/centos6/1.x/updates/1.7.0
gpgcheck=0
enabled=1
priority=1
```

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

```
clush -a -b -c /etc/yum.repos.d/hdp.repo --dest=/etc/yum.repos.d/
clush -a -b -c /etc/yum.repos.d/ambari.repo --dest=/etc/yum.repos.d/
```

HDP Installation

Follow these steps to install HDP.

Install and Setup Ambari Server on rhell

```
yum -y install ambari-server
```

```
[root@rhell ~]# yum -y install ambari-server
Loaded plugins: product-id, refresh-packagekit, security, subscription-manager
This system is not registered to Red Hat Subscription Management. You can use subscription-manager to register.
Setting up Install Process
Resolving Dependencies
--> Running transaction check
---> Package ambari-server.noarch 0:1.7.0-169 will be installed
--> Finished Dependency Resolution

Dependencies Resolved

=====
Package                Arch             Version          Repository
=====
Installing:
ambari-server          noarch           1.7.0-169       Updates-ambari-1.7.0
Transaction Summary
-----
Install      1 Package(s)

Total download size: 96 M
Installed size: 123 M
Downloading Packages:
ambari-server-1.7.0-169.noarch.rpm                                | 96 MB
Running rpm_check_debug
Running Transaction Test
Transaction Test Succeeded
Running Transaction
Warning: RPMDB altered outside of yum.
Installing : ambari-server-1.7.0-169.noarch [#####]
```

Setup Ambari Server

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

```
[root@rhell ~]# ambari-server setup -j $JAVA_HOME -s
Using python /usr/bin/python2.6
Setup ambari-server
Checking SELinux...
SELinux status is 'disabled'
Ambari-server daemon is configured to run under user 'root'. Change this setting [y/n] (n)?
Adjusting ambari-server permissions and ownership...
Checking firewall...
Checking JDK...
WARNING: JAVA_HOME /usr/java/jdk1.7.0_75 must be valid on ALL hosts
WARNING: JCE Policy files are required for configuring Kerberos security. If you plan to use Kerberos
urisdiction Policy Files are valid on all hosts.
Completing setup...
Configuring database...
Enter advanced database configuration [y/n] (n)?
Default properties detected. Using built-in database.
Checking PostgreSQL...
Running initdb: This may take upto a minute.
Initializing database: [ OK ]

About to start PostgreSQL
Configuring local database...
Connecting to local database...done.
Configuring PostgreSQL...
Restarting PostgreSQL
Extracting system views...
..ambari-admin-1.7.0.169.jar

Adjusting ambari-server permissions and ownership...
Ambari Server 'setup' completed successfully.
```

Start Ambari Server

```
ambari-server start
```

Confirm Ambari Server Startup

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

```
[root@rhell ~]# ambari-server start
Using python /usr/bin/python2.6
Starting ambari-server
Ambari Server running with 'root' privileges.
Organizing resource files at /var/lib/ambari-server/resources...
Server PID at: /var/run/ambari-server/ambari-server.pid
Server out at: /var/log/ambari-server/ambari-server.out
Server log at: /var/log/ambari-server/ambari-server.log
Waiting for server start.....
Ambari Server 'start' completed successfully.
[root@rhell ~]# ps -ef | grep ambari-server
root      6069      1 59 01:31 pts/0    00:00:15 /usr/java/jdk1.7.0_75/bin/java -server -XX:NewRatio=3
dLimit -XX:CMSInitiatingOccupancyFraction=60 -Xms512m -Xmx2048m -Djava.security.auth.login.config=/et
java.security.krb5.conf=/etc/krb5.conf -Djavax.security.auth.useSubjectCredsOnly=false -cp /etc/ambar
r/lib64/qt-3.3/bin:/usr/local/sbin:/usr/local/bin:/sbin:/bin:/usr/sbin:/usr/bin:/root/bin:/sbin:/usr
ambari.server.controller.AmbariServer
root      6186    4896  0 01:31 pts/0    00:00:00 grep ambari-server
```

Log into Ambari Server

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

1. Point the browser to <http://<IP address for rhell>:8080>
2. Log in to the Ambari Server using the default username/password: **admin/admin**. This can be changed at a later period of time.

Figure 133 Ambari Server Login



Sign in

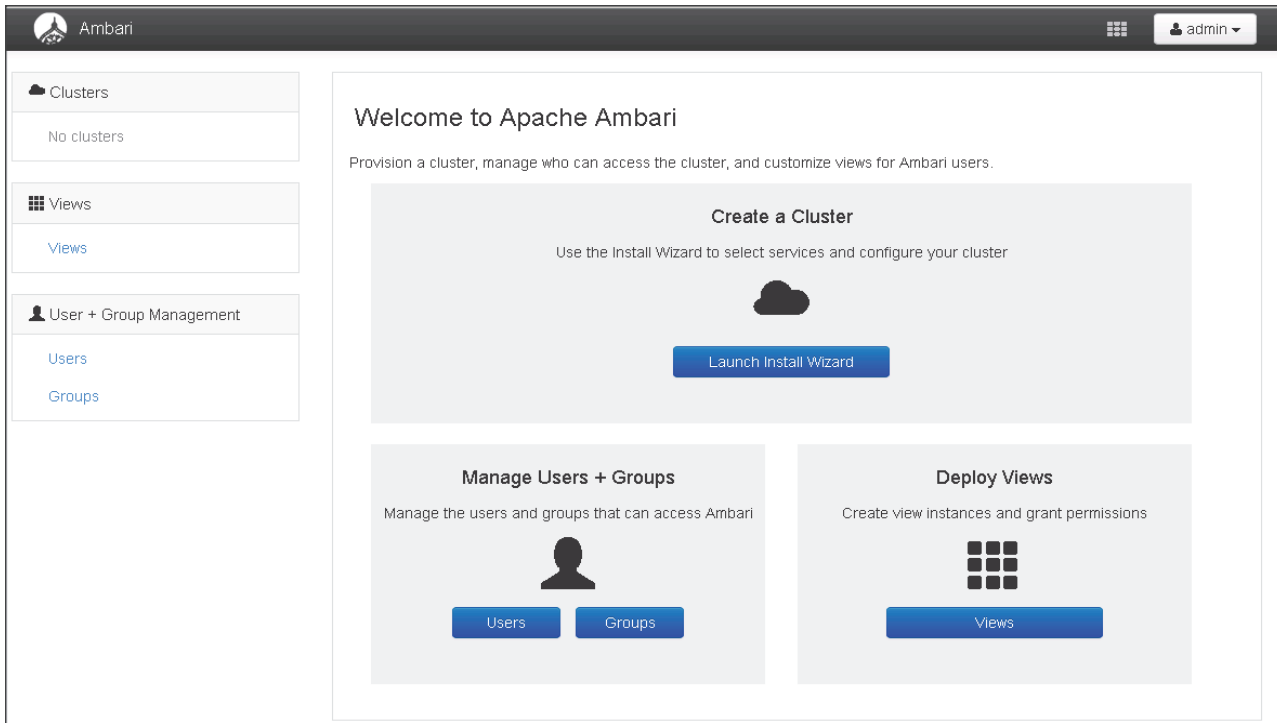
Username

Password

[Sign in](#)

Once logged in the **Welcome to Apache Ambari** window appears.

Figure 134 Apache Ambari

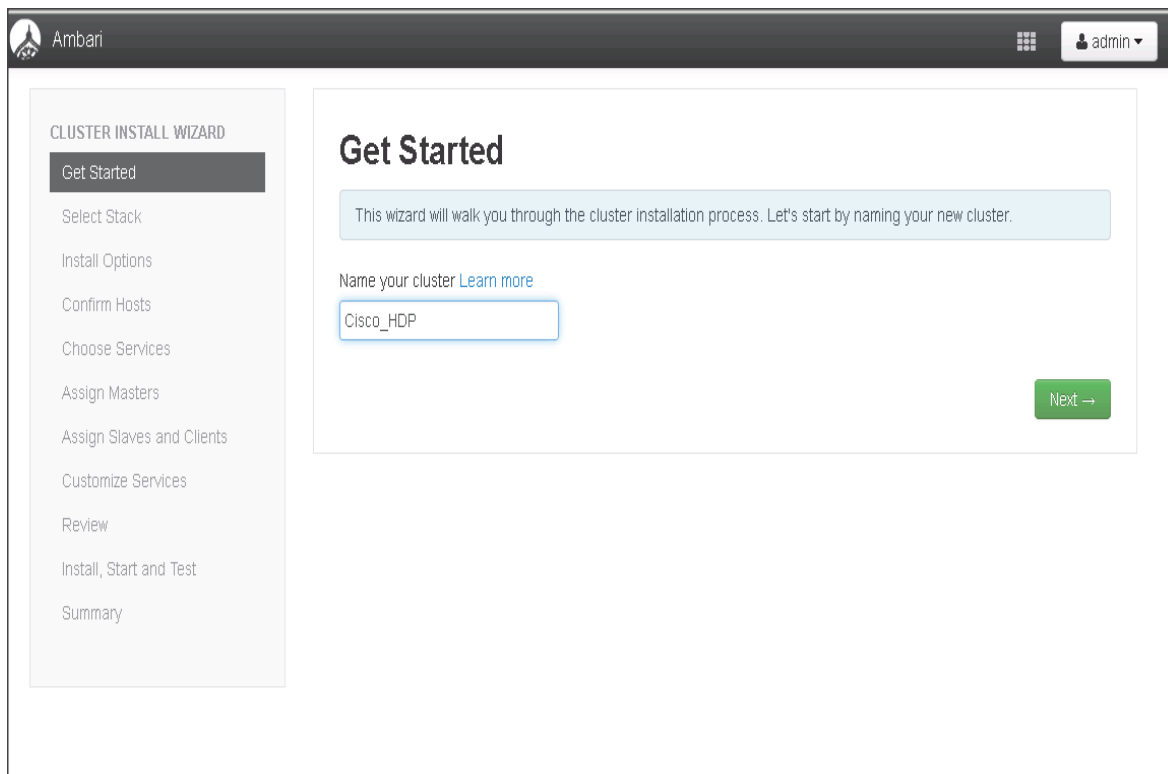


Create a Cluster

Use the following steps to create a cluster.

1. Click Launch install wizard button.
2. At the Get started page type “Cisco_HDP” as name for the cluster in the text box.
3. Click the **Next** button.

Figure 135 Ambari Getting Started



Select Stack

In the following screen

- Select HDP 2.2 stack.
- Expand “Advanced Repository Options”. Under the advanced repository option.
 - Select RedHat 6 checkbox.
 - Uncheck rest of the checkbox.
 - Update the Redhat 6 HDP-2.2 URL to <http://rhel1/Hortonworks/HDP/centos6/2.x/GA/2.2.0.0>
 - Update the Redhat 6 HDP-UTILS-1.1.0.20 URL to <http://rhel1/Hortonworks/HDP-UTILS-1.1.0.20/repos/centos6>

Figure 136 Ambari Screen

Ambari

admin

CLUSTER INSTALL WIZARD

- Get Started
- Select Stack**
- Install Options
- Confirm Hosts
- Choose Services
- Assign Masters
- Assign Slaves and Clients
- Customize Services
- Review
- Install, Start and Test
- Summary

Select Stack

Please select the service stack that you want to use to install your Hadoop cluster.

Stacks

- HDP 2.2
- HDP 2.1
- HDP 2.0
- HDP 1.3

Advanced Repository Options

Customize the repository Base URLs for downloading the Stack software packages. If your hosts do not have access to the internet, you will have to create a local mirror of the Stack repository that is accessible by all hosts and use those Base URLs here.

Important: When using local mirror repositories, you only need to provide Base URLs for the Operating System you are installing for your Stack. Uncheck all other repositories.

OS	Name	Base URL
<input type="checkbox"/> redhat5	HDP-2.2	<input type="text" value="http://public-repo-1.hortonworks.com/HDP/centos5/2.x/GA/2."/>
	HDP-UTILS-1.1.0.20	<input type="text" value="http://public-repo-1.hortonworks.com/HDP-UTILS-1.1.0.20/re"/>
<input checked="" type="checkbox"/> redhat6	HDP-2.2	<input type="text" value="http://rhel1/Hortonworks/HDP/centos6/2.x/GA/2.2.0.0"/>
	HDP-UTILS-1.1.0.20	<input type="text" value="http://rhel1/Hortonworks/HDP-UTILS-1.1.0.20/repos/centos6"/>
<input type="checkbox"/> suse11	HDP-2.2	<input type="text" value="http://public-repo-1.hortonworks.com/HDP/suse11sp3/2.x/GA"/>

HDP Installation

In order to build up the cluster, the install wizard needs to know general information about how the cluster has to be set up. This requires providing the Fully Qualified Domain Name (FQDN) of each of the host. The wizard also needs to access the private key file that was created in **Set Up Password-less SSH**. 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 per line. One can also use ranges inside 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.

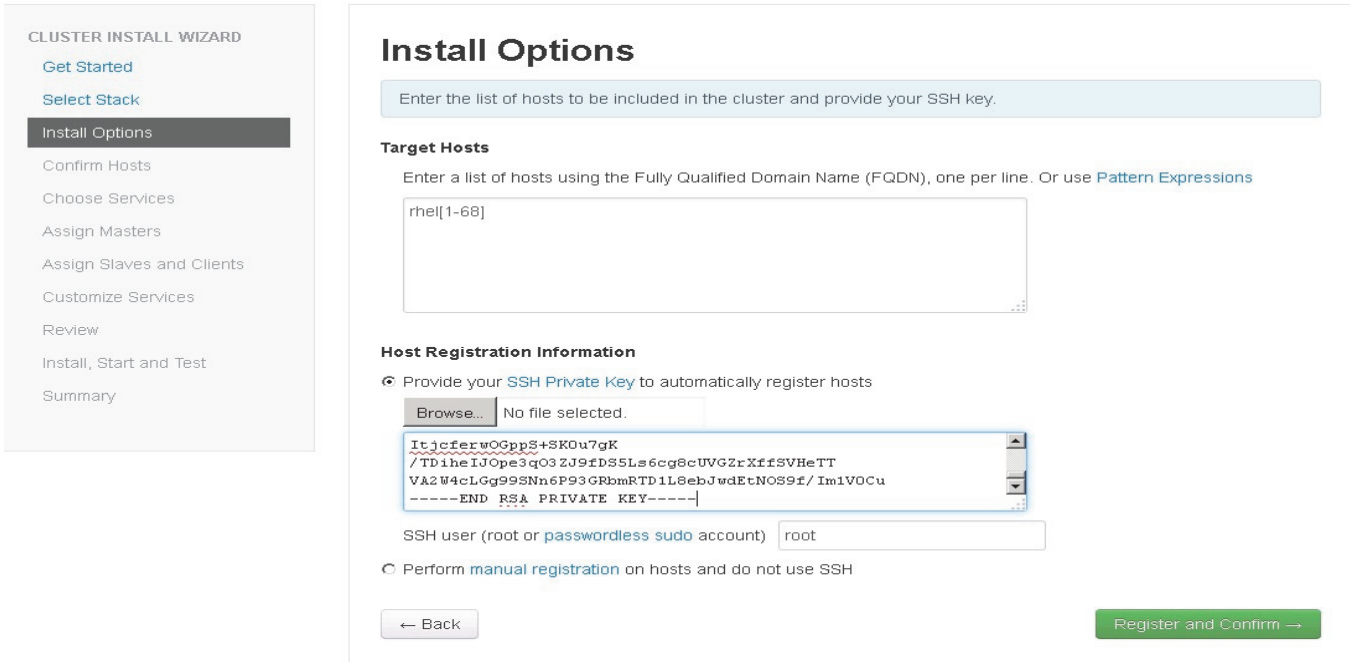


Note Make sure there is no extra white space after the text-----END RSA PRIVATE KEY-----

```
[root@rhell ~]# cat /root/.ssh/id_rsa
-----BEGIN RSA PRIVATE KEY-----
MIIEoQIBAAKCAQEAYDOIRbk4mBZrizec0/gOM2iYT2h4vxkIxA/uvQVPthFreUdgT
Zehw/Qtdk7meeqhgqsHmb1Crif0m6SxvPEXW2cGoAx75hZwTuDIR3Qlvk6oYUmDW
BKq5TMfUMKfd7tknkGkg5N+YHsPCoNILLz/Wqc01hZZ0tiCmrxeRnPGS1JY74/Db
A0BewMuNajAoVppPD6cLGF6/NKORpEDUnCuwe5pCRV5tko+gzBeBF5oeCS6Ya6I7
ns0HplJXV0Mv23SNUwl3cswbqLdrr3atG6YrieVrmmr/PlrKmp192tzQ1mHZMBqG
w1RJTILjygWogp5g7NQBGE7sX4V6Omzv4vmzwIBIwKCAQEAg4+UEI+o2PjKVCuX
2h+XEwMUXCJ3KoNEyBpr2nj7KxckYas/8oLN6B1pYR0UB3X2YZVc6hBwuLI+JDMk
hrGNMALqwdjthU10yX/9HDlmlDyTo9k8LvPY2q8zqvHnJ+3Jisi92Dspc01xRRxQ
wnpofjAmlCDx5Wxp4MZyX9HynCcKmheFefobLys6glox84eHWly6b0xU1dh7hsQ
pcK+xpdfWlshYFbvckTuCHUAezF4+uBT5F0PMiD7PwzrvbXKA65ABuezv9gg2/I1
PekIkRybosiFbBUi2ZOS1un/gsaZgmsQ9gTarJlV8zMy6K31LETcOckl2LZHRX2
5sEx6wKBGQD9CiKc0HfiulrQWW5cLTDJU8wzTiNK4M9lQb2LohffuzfluiAl3Ref
yil9MjE3A5Mnn9pcR:MmXXPF4t9iuLh3+3tCsriTzPml4WT+Fipa9sh+3JZ2HKgm
pCquAEdoFRK4oP3/yYQg95gie2SC9sB0z6zVohdyNUvnkiMb9vwi3wKBgQDKiyTi
Yu421owsYKfz7YjomjRKUFaH4CKtnyJy1SM3wFPRnzJd4BUaMq0DaTxr2tw4si+4
t88M8XS6FHGHymSqrTl0tYzmlmmwUtjCLNZQfqgSeg1NovekXxXL0iUzel8PL3ZOH
AeBj0/GLQ3SF/PGWMokCwnTaJoV/xldBdIsqEQKBgEERPBMx8UVF3Nz9ZYVqtMYO
O9KtsU3Ex52x0adlVpht5TsSmolkv06TEE+8cw4lfzX5j+vXwxh+bjozBj30/Dwc
GGGbrQbrkKscs5HLL3Z5+QqtWepB4hiQnUKvnVVHP1QMJA6S53YxCdz7KHlypnqq
bkWQFKhW2QEiUivDKuRlAoGASzr/EkIAtUfFb5Gdbj0n4V3Y6Gb7kY3DvNS1BhSm
rk7ADAdTnzX5Nz3L08gAf9Tws+ppfx+zTfNiNoMFmNYlY9EpyJs0S/1adLE0roWu
sC8J8bu/5RNWk8z+z9s5zwUrd5txT2cYlJ8t1KQGtWYUPxoVoe/ccfENA5LP872S
xnsCgYAFRE4SbB416p9miir1+gNCiihM9N+FmHmMcP/y80QL/MoAYoHB1Tn8cwVu
l+sju4bWGUVzvnGMWXwpEU5zVBra+yShh309IwjP/1kpCNWz7CX+/uI6FY+slZxTr
t5P/AvhOvUKMhRFjXFQoY5yqNUkasvIu6S8Q1unl8N2IhEgw1g==
-----END RSA PRIVATE KEY-----
```

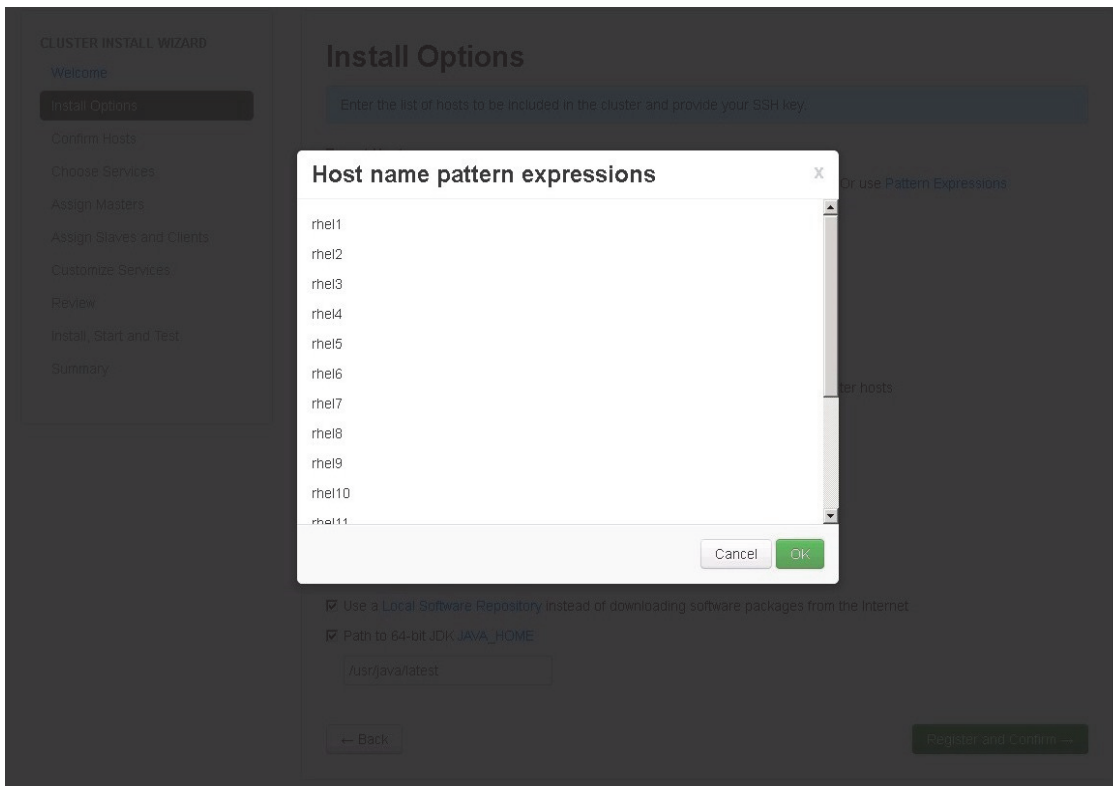
3. Click the **Register** and **Confirm** button to continue.

Figure 137 HDP Install Options



Hostname Pattern Expressions

Figure 138 Hostname Pattern Expressions



Confirm Hosts

This screen allows to ensure that Ambari has located the correct hosts for the cluster and to check those hosts to make sure they have the correct directories, packages, and processes to continue the install.

If any hosts were selected in error, it can be removed by selecting the appropriate **checkboxes** and clicking the grey **Remove Selected** button. To remove a single host, click the small white **Remove** button in the Action column.

When the lists of hosts are confirmed, click **Next**.

Ambari

admin

CLUSTER INSTALL WIZARD

- Get Started
- Select Stack
- Install Options
- Confirm Hosts**
- Choose Services
- Assign Masters
- Assign Slaves and Clients
- Customize Services
- Review
- Install, Start and Test
- Summary

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 (10) | Installing (0) | Registering (0) | Success (10) | Fail (0)

<input type="checkbox"/>	Host	Progress	Status	Action
<input type="checkbox"/>	rhel1	<div style="width: 100%; background-color: green;"></div>	Success	<input type="button" value="Remove"/>
<input type="checkbox"/>	rhel2	<div style="width: 100%; background-color: green;"></div>	Success	<input type="button" value="Remove"/>
<input type="checkbox"/>	rhel3	<div style="width: 100%; background-color: green;"></div>	Success	<input type="button" value="Remove"/>
<input type="checkbox"/>	rhel4	<div style="width: 100%; background-color: green;"></div>	Success	<input type="button" value="Remove"/>
<input type="checkbox"/>	rhel5	<div style="width: 100%; background-color: green;"></div>	Success	<input type="button" value="Remove"/>
<input type="checkbox"/>	rhel6	<div style="width: 100%; background-color: green;"></div>	Success	<input type="button" value="Remove"/>
<input type="checkbox"/>	rhel7	<div style="width: 100%; background-color: green;"></div>	Success	<input type="button" value="Remove"/>
<input type="checkbox"/>	rhel8	<div style="width: 100%; background-color: green;"></div>	Success	<input type="button" value="Remove"/>
<input type="checkbox"/>	rhel9	<div style="width: 100%; background-color: green;"></div>	Success	<input type="button" value="Remove"/>
<input type="checkbox"/>	rhel10	<div style="width: 100%; background-color: green;"></div>	Success	<input type="button" value="Remove"/>

Show: 25

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 139 Ambari Service Selection

Ambari admin

CLUSTER INSTALL WIZARD

- Get Started
- Select Stack
- 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 none	Version	Description
<input checked="" type="checkbox"/>		2.6.0.2.2.0.0	Apache Hadoop Distributed File System
<input checked="" type="checkbox"/>		2.6.0.2.2.0.0	Apache Hadoop NextGen MapReduce (YARN)
<input checked="" type="checkbox"/>		0.5.2.2.2.0.0	Tez is the next generation Hadoop Query Processing framework written on top of YARN.
<input checked="" type="checkbox"/>		3.5.0	Nagios Monitoring and Alerting system
<input checked="" type="checkbox"/>		3.5.0	Ganglia Metrics Collection system (RRDTool will be installed too)
<input checked="" type="checkbox"/>		0.14.0.2.2.0.0	Data warehouse system for ad-hoc queries & analysis of large datasets and table & storage management service
<input checked="" type="checkbox"/>		0.98.4.2.2.0.0	Non-relational distributed database and centralized service for configuration management & synchronization
<input checked="" type="checkbox"/>		0.14.0.2.2.0.0	Scripting platform for analyzing large datasets
<input checked="" type="checkbox"/>		1.4.5.2.2.0.0	Tool for transferring bulk data between Apache Hadoop and structured data stores such as relational databases
<input checked="" type="checkbox"/>		4.1.0.2.2.0.0	System for workflow coordination and execution of Apache Hadoop jobs. This also includes the installation of the optional Oozie Web Console which relies on and will install the ExtJS Library.
<input checked="" type="checkbox"/>		3.4.6.2.2.0.0	Centralized service which provides highly reliable distributed coordination
<input checked="" type="checkbox"/>		0.6.0.2.2.0.0	Data management and processing platform
<input checked="" type="checkbox"/>		0.9.3.2.2.0.0	Apache Hadoop Stream processing framework
<input checked="" type="checkbox"/>		1.5.2.2.2.0.0	A distributed service for collecting, aggregating, and moving large amounts of streaming data into HDFS
<input checked="" type="checkbox"/>		0.8.1.2.2.0.0	A high-throughput distributed messaging system
<input checked="" type="checkbox"/>		0.5.0.2.2.0.0	Provides a single point of authentication and access for Apache Hadoop services in a cluster
<input checked="" type="checkbox"/>		0.60.0.2.2.0.0	A framework for deploying, managing and monitoring existing distributed applications on YARN.

← Back Next →

Assign Masters

The Ambari install wizard attempts to assign the master nodes for various services that have been selected to appropriate hosts in the cluster. The right column shows the current service assignments by host, with the hostname and its number of CPU cores and amount of RAM indicated.

1. Reconfigure the service assignment to match the table shown below:

Table 8 *Reconfiguring the Service Assignment*

Service Name	Host
NameNode	rhel1
SNameNode	rhel2
HistoryServer	rhel2
ResouceManager	rhel2
Nagios Server	rhel1
Ganglia Collector	rhel1
HiveServer2	rhel2
HBase Master	rhel2
Oozie Server	rhel1
Zookeeper	rhel1, rhel2, rhel3
Kafka Broker	rhel1
Knox Gateway	rhel1
App Timeline Server	rhel2
Hive Metastore	rhel2
WebHCat Server	rhel2
Falcon Server	rhel2
DRPC Server	rhel2
Nimbus	rhel2
Storm UI Server	rhel2

Figure 140 Ambari Assign Masters

The screenshot displays the 'Assign Masters' configuration page in Ambari. On the left, a sidebar lists the 'CLUSTER INSTALL WIZARD' steps, with 'Assign Masters' selected. The main area is titled 'Assign Masters' and contains a list of services with dropdown menus for host selection. A note states: 'Assign master components to hosts you want to run them on. * HiveServer2, Hive Metastore and WebHCat Server will be hosted on the same host.'

Services and their assigned hosts:

- SNameNode: rhel2 (252.2 GB, 48 cores)
- NameNode: rhel1 (252.2 GB, 48 cores)
- History Server: rhel2 (252.2 GB, 48 cores)
- App Timeline Server: rhel2 (252.2 GB, 48 cores)
- ResourceManager: rhel2 (252.2 GB, 48 cores)
- Nagios Server: rhel1 (252.2 GB, 48 cores)
- Ganglia Server: rhel1 (252.2 GB, 48 cores)
- Hive Metastore: rhel2*
- WebHCat Server: rhel2*
- HiveServer2: rhel2 (252.2 GB, 48 cores)
- HBase Master: rhel2 (252.2 GB, 48 cores)
- Oozie Server: rhel1 (252.2 GB, 48 cores)
- ZooKeeper Server: rhel2 (252.2 GB, 48 cores)
- ZooKeeper Server: rhel3 (252.2 GB, 48 cores)
- ZooKeeper Server: rhel1 (252.2 GB, 48 cores)
- Falcon Server: rhel2 (252.2 GB, 48 cores)
- DRPC Server: rhel2 (252.2 GB, 48 cores)
- Nimbus: rhel2 (252.2 GB, 48 cores)
- Storm UI Server: rhel2 (252.2 GB, 48 cores)
- Kafka Broker: rhel1 (252.2 GB, 48 cores)
- Knox Gateway: rhel1 (252.2 GB, 48 cores)

Host assignments shown on the right:

- rhel1 (252.2 GB, 48 cores):** NameNode, Nagios Server, Ganglia Server, Oozie Server, ZooKeeper Server, Kafka Broker, Knox Gateway.
- rhel2 (252.2 GB, 48 cores):** SNameNode, History Server, App Timeline Server, ResourceManager, Hive Metastore, WebHCat Server, HiveServer2, HBase Master, ZooKeeper Server, Falcon Server, DRPC Server, Nimbus, Storm UI Server.
- rhel3 (252.2 GB, 48 cores):** ZooKeeper Server.

Navigation buttons: '← Back' and 'Next →'.



Note On a small cluster (<16 nodes), consolidate all master services to run on a single node. For large clusters (> 64 nodes), deploy master services across 3 nodes.

2. Click the **Next** button.

Assign Slaves and Clients

The Ambari install wizard attempts to assign the slave components (DataNodes, NodeManagers, RegionServers, Supervisor and Flume) to appropriate hosts in the cluster. Reconfigure the service assignment to match below:

1. Assign DataNode, NodeManager, RegionServer, Supervisor and Flume on nodes rhel3- rhel68
2. Assign Client to all nodes
3. Click the **Next** button.

Table 9 Client Service Name

Client Service Name	Host
DataNode	rhel3 to rhel68
NodeManager	rhel3 to rhel68
RegionServer	rhel3 to rhel68
Client	All nodes, rhel1-rhel68
Supervisor	rhel3-rhel68
Flume	rhel3-rhel68
Archival Nodes	rhel65 to rhel68

Figure 141 Ambari Slave and Client Assignment

Assign Slaves and Clients

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, MapReduce2 Client, YARN Client, Tez Client, HCat Client, Hive Client, HBase Client, Pig, Sqoop, Oozie Client, ZooKeeper Client, Falcon Client and Slider.

Host	all	none	all	none	all	none	all	none	all	none	all	none
rhel1 *	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
rhel2 *	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
rhel3 *	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
rhel4	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
rhel5	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
rhel6	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
rhel7	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
rhel8	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
rhel9	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
rhel10	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>

Show: 25 1 - 17 of 17

← Back Next →

Customize Services

This section presents with a set of tabs that manage configuration settings for Hadoop components. The wizard attempts to set reasonable defaults for each of the options here, but this can be modified to meet specific requirements. Following sections provide configuration guidance that should be refined to meet specific use case requirements.

Following are the changes made

- Memory and service level setting for each component and service level tuning.
- Create a separate configuration for archival nodes to account for them having different number of data drives (4 Volumes of RAID5 each with 14+1 drives) compared to data nodes (24 volumes). We will then point Ambari to these archive nodes and override the default configuration for these nodes.
- Customize the log locations of all the components to ensure growing logs do not cause the SSDs to run out of space.

HDFS

In Ambari, choose **HDFS Service** from the left tab and go to **Configs** tab.

Use the “Search” box on top to filter for the properties mentioned in the tab to update the values

HDFS JVM Settings

Update the following HDFS configurations in Ambari:

Table 10 HDFS Configurations

Property Name	Value
NameNode Java Heap Size	4096
Hadoop maximum Java heap size	4096
DataNode maximum Java heap size	4096
DataNode Volumes Failure Toleration	3

Manage Config Groups for Archival Nodes

Ambari initially assigns all hosts in your cluster to one default configuration group for each service installed. Since there are different number of data drives (volumes) on archival nodes (4) as compared to other data nodes (24), create a new group for archive nodes and override the default data node directories that were picked up.

1. Under the **HDFS** tab, click **Manage Config Groups**

Figure 142 *HDFS Manage Config Groups*

The screenshot shows the 'Customize Services' screen. On the left, a sidebar lists the installation steps: Get Started, Select Stack, Install Options, Confirm Hosts, Choose Services, Assign Masters, Assign Slaves and Clients, **Customize Services**, Review, Install, Start and Test, and Summary. The main content area is titled '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 horizontal list of services: HDFS, MapReduce2, YARN, Tez, Nagios (2), Ganglia, Hive (1), HBase, Pig, Sqoop, and Oozie (1). A second row shows ZooKeeper, Falcon, Storm, Flume, Kafka, Knox (1), Slider, and Misc. A 'Group' dropdown is set to 'HDFS Default (68)' and a 'Filter...' dropdown is empty. The 'NameNode' section is expanded, showing:

- NameNode hosts: rhel1
- NameNode directories: /data/disk1/hadoop/hdfs/namenode
- NameNode Java heap size: 1024 MB
- NameNode new generation size: 200 MB

2. Create an Archive group by clicking the + button on the left and entering “Archive” for the name.

Figure 143 *Manage HDFS Configuration Groups*

The screenshot shows the 'Manage HDFS Configuration Groups' dialog box. At the top, it says: 'You can apply different sets of HDFS configurations to groups of hosts by managing HDFS Configuration Groups and their host membership. Hosts belonging to a HDFS Configuration Group have the same set of configurations for HDFS. Each host belongs to one HDFS Configuration Group.' Below this is a list of groups, with 'HDFS Default (68)' selected. To the right, a list of hosts is shown: rhel1, rhel2, rhel3, rhel4, rhel5, rhel6, rhel7, rhel8, rhel9, rhel10, rhel11, rhel12, rhel13, rhel14, and rhel15. At the bottom, there are buttons for '+', '-', and a gear icon. Below these buttons, it says 'Overrides: 0 properties' and 'Description: Default cluster level HDFS configuration'. At the bottom right, there are 'Cancel' and 'Save' buttons.

Figure 144 Create New Configuration Group

Create New Configuration Group X

Name:

Description:

Now the Archive group should appear in the Configuration Groups.

Figure 145 Manage HDFS Configuration Groups

Manage HDFS Configuration Groups X

You can apply different sets of HDFS configurations to groups of hosts by managing HDFS Configuration Groups and their host membership. Hosts belonging to a HDFS Configuration Group have the same set of configurations for HDFS. Each host belongs to one HDFS Configuration Group.

HDFS Default (64)

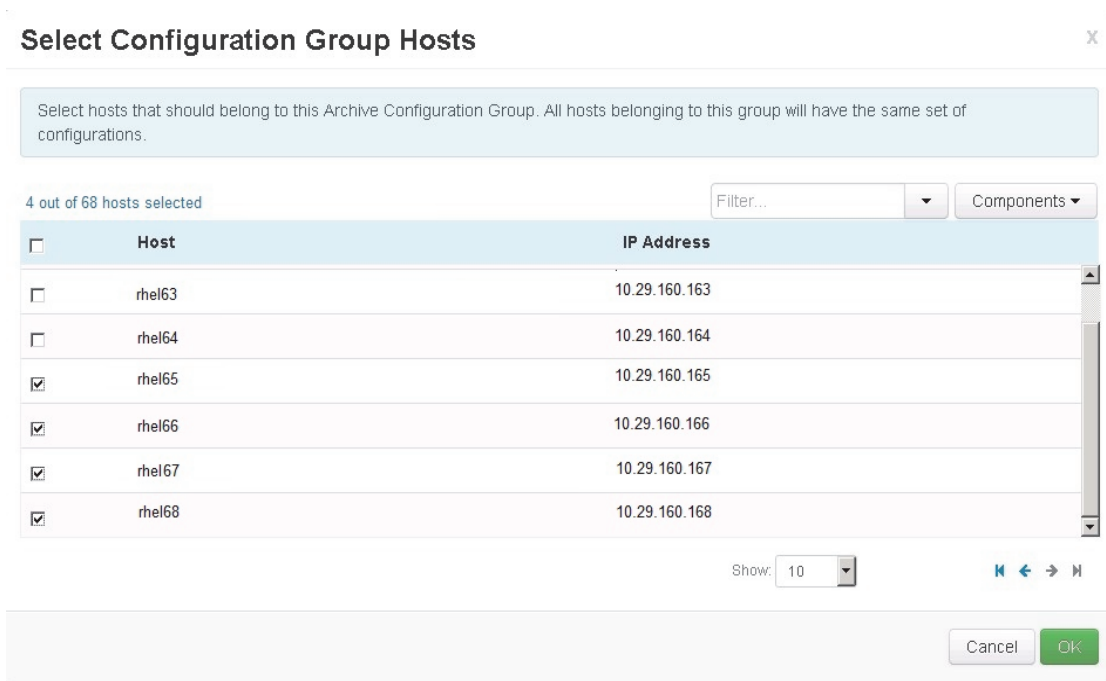
Archive (0)

Overrides: 0 properties

Description: 60 disks

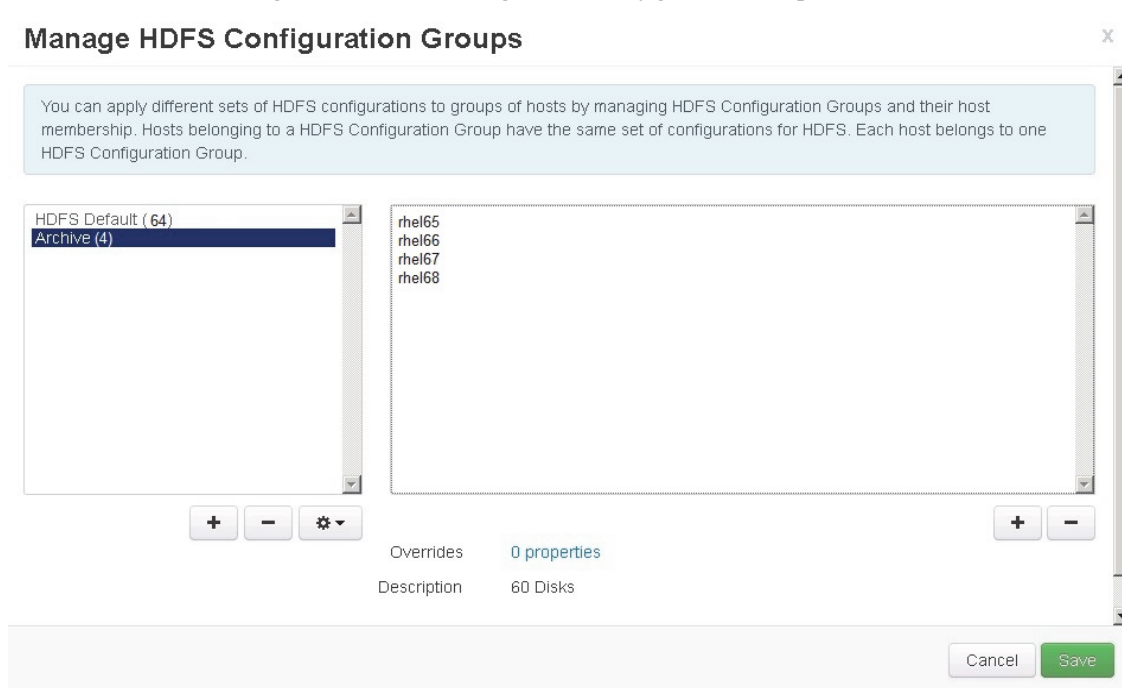
3. Add archive nodes to the archive group, by clicking on the **+** sign on the right and selecting the archive nodes, and click **OK** when done.

Figure 146 *Select Configuration Group Hosts*



Now you should see the archive nodes appear under the Archive configuration group. Click **Save** to commit the new configuration group you created.

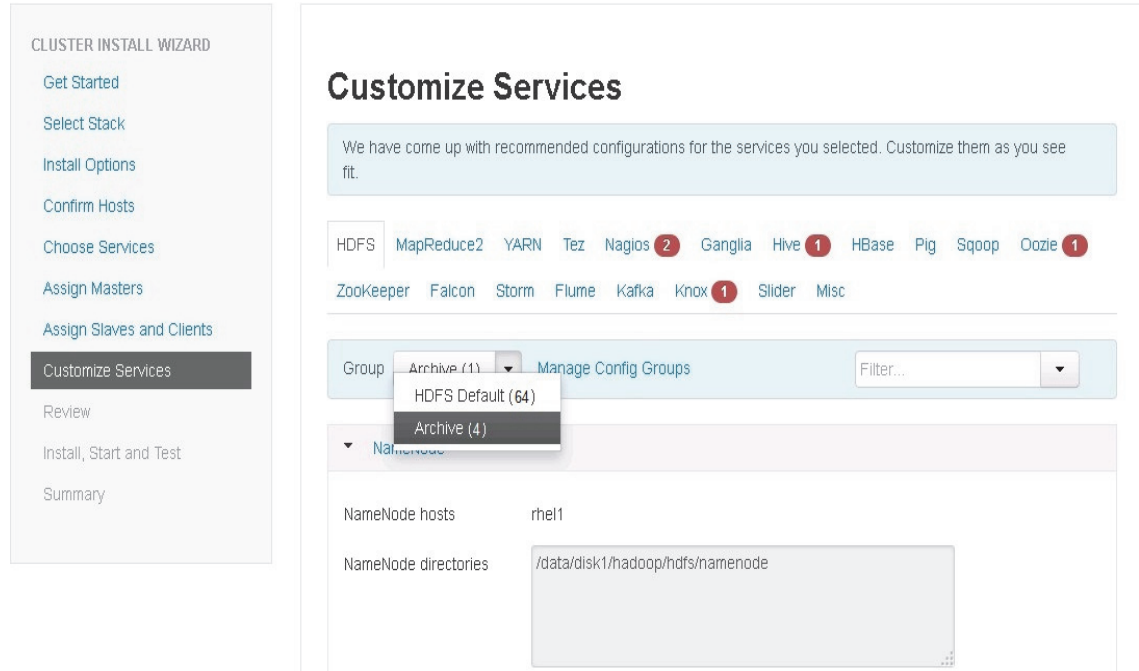
Figure 147 *Manage HDFS Configuration Groups*



Next we will configure the archive nodes to use different data directories than the data nodes by overriding this value by following the below steps.

4. Select HDFS Default group, scroll down to “Data Node directories” property and copy first four lines from the Data Node directories.
5. And then select **Archive config group**

Figure 148 *Manage HDFS: Customize Services Part1*



6. Scroll down to the “Data Node directories” property and click override button (to the right of the textbox) and paste those four lines from Data Node directories here.

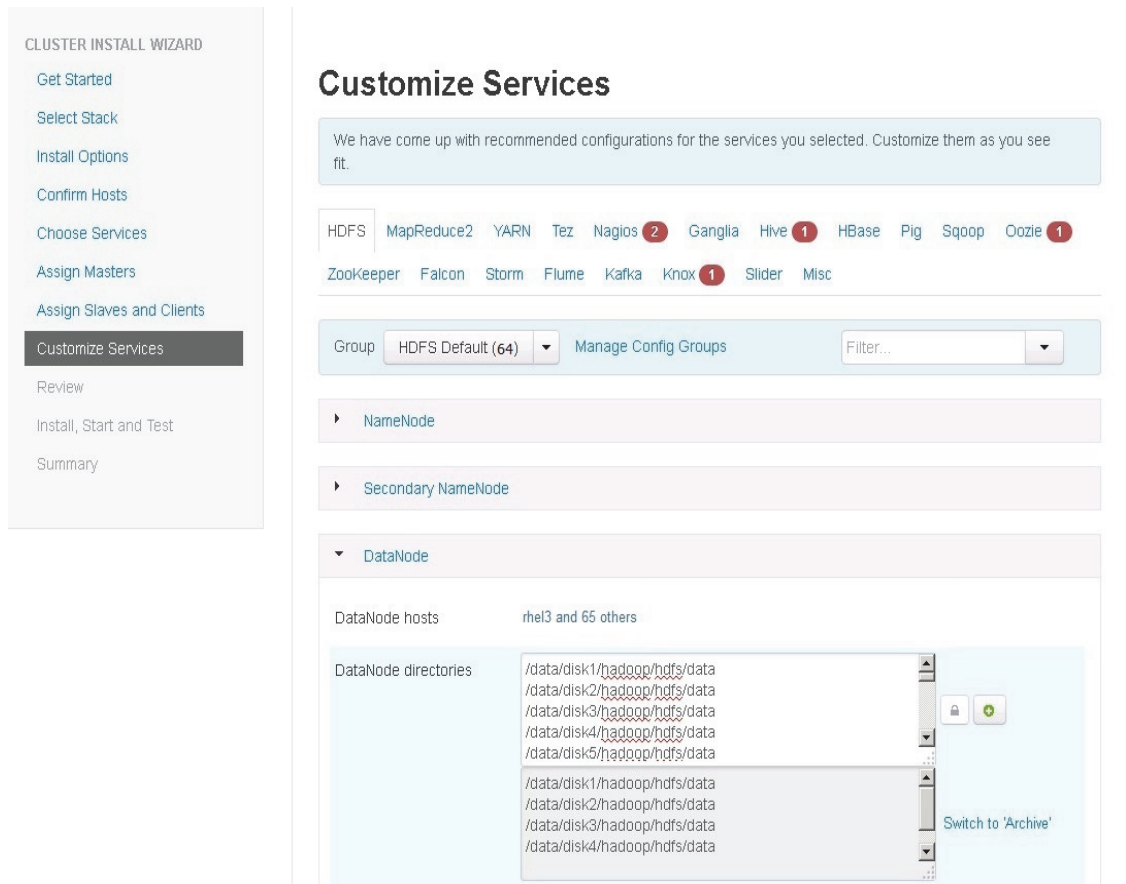
```
/data/disk1/hadoop/hdfs/data
/data/disk2/hadoop/hdfs/data
/data/disk3/hadoop/hdfs/data
/data/disk4/hadoop/hdfs/data
```

Figure 149 *Manage HDFS: Customize Services Part2*

The screenshot displays the 'Customize Services' interface. On the left, a sidebar titled 'CLUSTER INSTALL WIZARD' lists steps: Get Started, Select Stack, Install Options, Confirm Hosts, Choose Services, Assign Masters, Assign Slaves and Clients, **Customize Services**, 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, a horizontal list of services is shown: HDFS, MapReduce2, YARN, Tez, Nagios (2), Ganglia, Hive (1), HBase, Pig, Sqoop, and Oozie (1). A second row shows Zookeeper, Falcon, Storm, Flume, Kafka, Knox (1), Slider, and Misc. A 'Group' dropdown is set to 'Archive (4)'. The 'DataNode' section is expanded, showing 'DataNode hosts' as 'rhel3 and 65 others' and a list of 'DataNode directories' including paths like '/data/disk1/hadoop/hdfs/data'.

7. We are done customizing the Archive Config Group. For all remaining steps, select back the HDFS default group.

Figure 150 *Manage HDFS: Customize Services Part3*



Update Log Directory

Change the default log location by finding the Log Dir property and modifying the **/var** prefix to **/data/disk1**

Figure 151 Manage HDFS: Customize Services Part4

Customize Services

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

HDFS MapReduce2 YARN Tez Nagios 2 Ganglia Hive 1 HBase Pig Sqoop Oozie 1
ZooKeeper Falcon Storm Flume Kafka Knox 1 Slider Misc

Group HDFS Default (64) Manage Config Groups /var

Advanced hadoop-env

Hadoop Log Dir Prefix /data/disk1/log/hadoop

Hadoop PID Dir Prefix /var/run/hadoop

Hadoop Log Dir Prefix
hdfs_log_dir_prefix
Hadoop Log Dir Prefix

MapReduce2

In Ambari, choose MapReduce Service from the left tab and go to “Configs” tab.

Use the “Search” box on top to filter for the properties mentioned in the tab to update the values

Update the following MapReduce configurations.

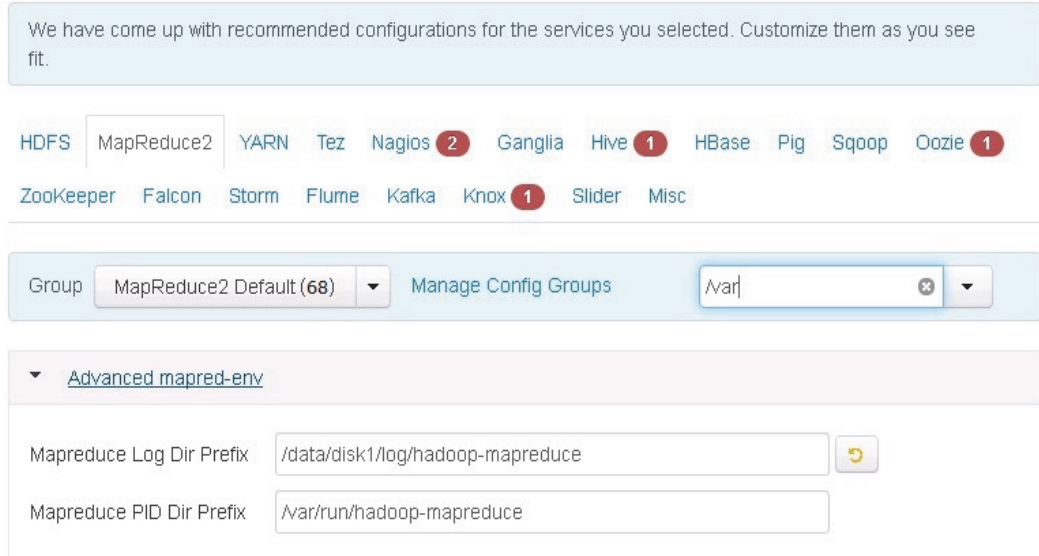
Table 11 MapReduce Property Name and Value

Property Name	Value
Default virtual memory for a job's map-task	4096
Default virtual memory for a job's reduce-task	8192
Map-side sort buffer memory	1638
yarn.app.mapreduce.am.resource.mb	4096
mapreduce.map.java.opts	-Xmx3276m
mapreduce.reduce.java.opts	-Xmx6552m
yarn.app.mapreduce.am.command-opts	-Xmx6552m

Similarly under MapReduce2 tab, change the default log location by finding the Log Dir property and modifying the /var prefix to /data/disk1.

Figure 152 MapReduce: Customize Services

Customize Services



YARN

In Ambari, choose YARN Service from the left tab and go to “Configs” tab.

Use the “Search” box on top to filter for the properties mentioned in the tab to update the values

Update the following YARN configurations

Table 12 Yarn Property Name and Value

Property Name	Value
ResourceManager Java heap size	4096
NodeManager Java heap size	2048
yarn.nodemanager.resource.memory-mb	184320
YARN Java heap size	4096
yarn.scheduler.minimum-allocation-mb	4096
yarn.scheduler.maximum-allocation-mb	184320

Similarly under YARN tab, change the default log location by finding the Log Dir property and modifying the **/var prefix** to **/data/disk1**

Figure 153 Yarn: Customize Services

Customize Services

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

HDFS MapReduce2 **YARN** Tez Nagios **2** Ganglia Hive **1** HBase Pig Sqoop Oozie **1**
 ZooKeeper Falcon Storm Flume Kafka Knox **1** Slider Misc

Group **YARN Default (68)** Manage Config Groups

Advanced yarn-env

YARN Log Dir Prefix 
 YARN PID Dir Prefix

Tez

No changes required.

Figure 154 Tez: Customize Services

Customize Services

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

HDFS MapReduce2 **YARN** **Tez** Nagios **2** Ganglia Hive **1** HBase Pig Sqoop Oozie **1**
 ZooKeeper Falcon Storm Flume Kafka Knox **1** Slider Misc

Group **Tez Default (68)** Manage Config Groups

▶ General

▶ Advanced tez-env

Nagios

On the **Nagios** tab, it is required to provide:

- Nagios admin password (as per organizational policy standards)
- Hadoop admin email

Figure 155 Nagios: Customize Services

Customize Services

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

HDFS MapReduce2 YARN Tez Nagios Ganglia Hive **1** HBase Pig Sqoop Oozie **1**

ZooKeeper Falcon Storm Flume Kafka Knox **1** Slider Misc

Group Nagios Default (68) Manage Config Groups Filter...

General

Nagios Admin username

Nagios Admin password

Hadoop Admin email

Ganglia

No changes are required.

Figure 156 Ganglia: Customize Services

Customize Services

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

HDFS MapReduce2 YARN Tez Nagios Ganglia Hive **1** HBase Pig Sqoop Oozie **1**

ZooKeeper Falcon Storm Flume Kafka Knox **1** Slider Misc

Group Ganglia Default (68) Manage Config Groups Filter...

General

Ganglia rrdcached base directory

Hive

In Ambari, choose HIVE Service from the left tab and go to **Configs** tab.

On the **Hive** tab, enter:



Note

Install Hive as is with just the following Log Dir changes, this uses MapReduce as engine. The Section on “Configuring Hive to use Apache Tez”, goes into details on tuning Hive for Performance by using Apache Tez as the execution engine and using Cost Based Optimizer to run queries.

Following are the changes for configuring Log Dir on Hive:

- Enter the Hive database password as per organizational policy.
- Change hive log directory by finding the Log Dir property and modifying the /var prefix to /data/disk1.
- Change the WebHCat log directory by finding the Log Dir property and modifying the /var prefix to /data/disk1.

Figure 157 Hive: Customize Services Part1

Customize Services

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

HDFS MapReduce2 YARN Tez Nagios Ganglia **Hive** HBase Pig Sqoop Oozie 1
 ZooKeeper Falcon Storm Flume Kafka Knox 1 Slider Misc

Group **Hive Default (68)** Manage Config Groups Filter...

Hive Metastore

Hive Metastore host rhel2

Database Type MySQL

Hive Database New MySQL Database
 Existing MySQL Database
 Existing PostgreSQL Database
 Existing Oracle Database

Database Host rhel2

Database Name

Database Username

Database Password

Figure 158 Hive: Customize Services Part2

Customize Services

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

HDFS **MapReduce2** YARN Tez Nagios Ganglia **Hive** HBase Pig Sqoop Oozie 1
 ZooKeeper Falcon Storm Flume Kafka Knox 1 Slider Misc

Group **Hive Default (68)** Manage Config Groups /var

Advanced hive-env

Hive Log Dir

Hive PID Dir

Advanced webhcat-env

WebHCat Log Dir

WebHCat PID Dir

HBase

In Ambari, choose HBASE Service from the left tab and go to **Configs** tab.

Use the “Search” box on top to filter for the properties mentioned in the tab to update the values

Update the following **HBASE configurations**:

Table 13 HBASE Configuration

Property Name	Value
HBase Master Maximum Java Heap Size	4096
HBase RegionServers Maximum Java Heap Size	16384



Note If you are not running HBase, keep the default value of 1024 for Java Heap size for HBase RegionServers and HBase Master.

Similarly under HBase tab, change the default log location by finding the Log Dir property and modifying the /var prefix to /data/disk1

Figure 159 HBase: Customize Services

Customize Services

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

HDFS MapReduce2 YARN Tez Nagios Ganglia Hive HBase Pig Sqoop Oozie 1
 ZooKeeper Falcon Storm Flume Kafka Knox 1 Slider Misc

Group HBase Default (68) Manage Config Groups /var

Advanced hbase-env

HBase Log Dir /data/disk1/log/hbase
 HBase PID Dir /var/run/hbase

Pig

No changes are required.

Customize Services

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

HDFS MapReduce2 YARN Tez Nagios Ganglia Hive HBase Pig Sqoop Oozie 1
 ZooKeeper Falcon Storm Flume Kafka Knox 1 Slider Misc

Group Pig Default (68) Manage Config Groups Filter...

Advanced pig-env

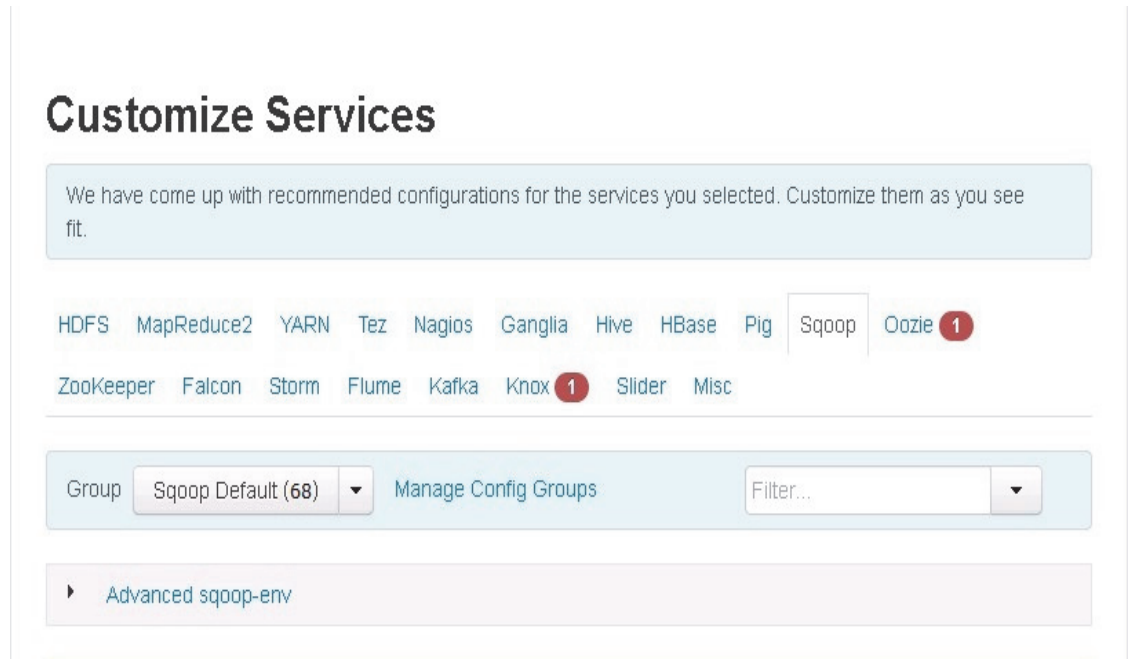
Advanced pig-log4j

Advanced pig-properties

Sqoop

No changes are required.

Figure 160 *Sqoop Customize Services*



Oozie

Similarly under **Oozie** tab, change the default log location by finding the Log Dir property and modifying the /var prefix to /data/disk1.

Also enter the oozie database password as per organizational policy.

Figure 161 Oozie: Customize Services Part1

Customize Services

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

HDFS MapReduce2 YARN Tez Nagios Ganglia Hive HBase Pig Sqoop **Oozie** ZooKeeper
 Falcon Storm Flume Kafka Knox **1** Slider Misc

Group **Oozie Default (68)** Manage Config Groups Filter...

Oozie Server

Oozie Server host rhel2

Database Type Derby

Oozie Database New Derby Database
 Existing MySQL Database
 Existing PostgreSQL Database
 Existing Oracle Database

Database Name

Database Username

Database Password

Figure 162 Oozie: Customize Services Part2

Customize Services

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

HDFS MapReduce2 YARN Tez Nagios Ganglia Hive HBase Pig Sqoop **Oozie** ZooKeeper
 Falcon Storm Flume Kafka Knox **1** Slider Misc

Group **Oozie Default (68)** Manage Config Groups /var

Advanced oozie-env

Oozie Log Dir

Oozie PID Dir

Zookeeper

Similarly under Zookeeper tab, change the default log location by finding the Log Dir property and modifying the /var prefix to /data/disk1

Figure 163 *Zookeeper: Customize Services*

Customize Services

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

HDFS MapReduce2 YARN Tez Nagios Ganglia Hive HBase Pig Sqoop Oozie **ZooKeeper**

Falcon Storm Flume Kafka Knox **1** Slider Misc

Group: ZooKeeper Default (68) Manage Config Groups /var

Advanced zookeeper-env

ZooKeeper Log Dir: /data/disk1/log/zookeeper

ZooKeeper PID Dir: /var/run/zookeeper

Falcon

Similarly under Falcon tab, change the default log location by finding the Log Dir property and modifying the /var prefix to /data/disk1.

Figure 164 Falcon: Customize Services

Customize Services

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

[HDFS](#)
[MapReduce2](#)
[YARN](#)
[Tez](#)
[Nagios](#)
[Ganglia](#)
[Hive](#)
[HBase](#)
[Pig](#)
[Sqoop](#)
[Oozie](#)
[ZooKeeper](#)

[Falcon](#)
[Storm](#)
[Flume](#)
[Kafka](#)
[Knox](#)
1
[Slider](#)
[Misc](#)

Group: Falcon Default (68) Manage Config Groups: /var

Advanced falcon-env

Falcon Log Dir:

Falcon PID Dir:

Storm

Similarly under Storm tab, change the default log location by finding the Log Dir property and modifying the /var prefix to /data/disk1

Figure 165 Storm: Customize Services

Customize Services

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

[HDFS](#)
[MapReduce2](#)
[YARN](#)
[Tez](#)
[Nagios](#)
[Ganglia](#)
[Hive](#)
[HBase](#)
[Pig](#)
[Sqoop](#)
[Oozie](#)
[ZooKeeper](#)

[Falcon](#)
[Storm](#)
[Flume](#)
[Kafka](#)
[Knox](#)
1
[Slider](#)
[Misc](#)

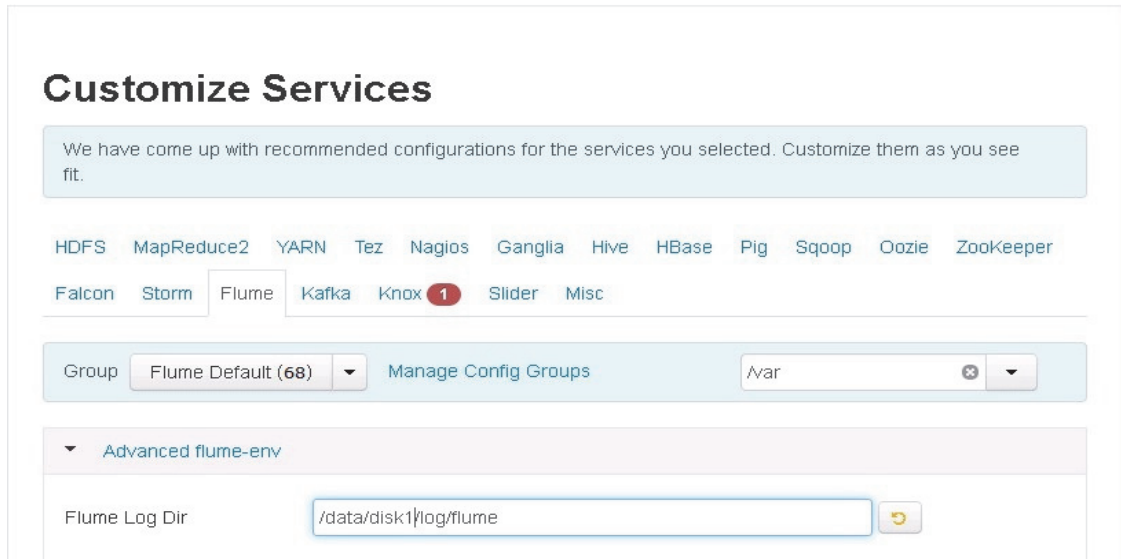
Group: Storm Default (68) Manage Config Groups /var

▾ Advanced storm-env

storm_log_dir	<input type="text" value="/data/disk1/log/storm"/>	+ ↺
storm_pid_dir	<input type="text" value="/var/run/storm"/>	+

Flume

Similarly under Flume tab, change the default log location by finding the Log Dir property and modifying the /var prefix to /data/disk1

Figure 166 *Flume: Customize Services*

Kafka

Similarly under Knox tab, change the default log location by finding the Log Dir property and modifying the /var prefix to /data/disk1

Figure 167 *Kafka: Customize Services Part1*

Customize Services

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

HDFS MapReduce2 YARN Tez Nagios Ganglia Hive HBase Pig Sqoop Oozie ZooKeeper
Falcon Storm Flume **Kafka** **Knox** ¹ Slider Misc

Group **Kafka Default (68)** Manage Config Groups /var

Advanced kafka-env

Kafka PID dir /var/run/kafka

kafka_log_dir /data/disk1/log/kafka

Figure 168 *Kafka: Customize Services Part2*

Customize Services

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

HDFS MapReduce2 YARN Tez Nagios Ganglia Hive HBase Pig Sqoop Oozie ZooKeeper
Falcon Storm Flume **Kafka** **Knox** Slider Misc

Group **Knox Default (68)** Manage Config Groups Filter...

Knox Gateway

Knox Gateway host rhel1

Knox Master Secret

Knox

For Knox, change the gateway port to 8444 to ensure no conflicts with local HTTP server.

Figure 169 Knox: Customize Services

Customize Services

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

[HDFS](#) [MapReduce2](#) [YARN](#) [Tez](#) [Nagios](#) [Ganglia](#) [Hive](#) [HBase](#) [Pig](#) [Sqoop](#) [Oozie](#) [ZooKeeper](#)
[Falcon](#) [Storm](#) [Flume](#) [Kafka](#) **Knox** [Slider](#) [Misc](#)

Group: Knox Default (68) [Manage Config Groups](#) gateway

Knox Gateway
 Knox Gateway host: rhel1

Advanced gateway-log4j

Advanced gateway-site

gateway.gateway.conf.dir	<input type="text" value="deployments"/>	<input type="button" value="⊕"/>
gateway.hadoop.kerberos.secured	<input type="text" value="false"/>	<input type="button" value="⊕"/>
gateway.path	<input type="text" value="gateway"/>	<input type="button" value="⊕"/>
gateway.port	<input type="text" value="8444"/>	<input type="button" value="⊕"/> <input type="button" value="↺"/>

Slider

No changes are required.

Figure 170 Slider: Customize Services

Customize Services

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

HDFS MapReduce2 YARN Tez Nagios Ganglia Hive HBase Pig Sqoop Oozie ZooKeeper
Falcon Storm Flume Kafka Knox Slider Misc

Group Slider Default (68) Manage Config Groups Filter...

Advanced slider-env

Advanced slider-log4j

Misc

No changes are required.

Figure 171 Misc: Customize Services

The screenshot shows the 'Customize Services' screen within the 'CLUSTER INSTALL WIZARD'. On the left is a navigation sidebar with the following items: Get Started, Select Stack, 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 contains the same text as Figure 170. The 'Misc' tab is selected in the service list. Below the service list, there is a section for 'Users and Groups' with a right-pointing arrow. At the bottom of the main content area, there are 'Back' and 'Next' buttons.

Review

The assignments that have been made are displayed. Check to ensure everything is correct before clicking on **Deploy** button. If any changes are to be made, use the left navigation bar to return to the appropriate screen.

Deploy

Once review is complete, click the **Deploy** button.

Figure 172 *Review the Assignments*

The screenshot shows the 'Review' step of the HDP installation wizard. On the left, a sidebar lists the steps: Get Started, Select Stack, Install Options, Confirm Hosts, Choose Services, Assign Masters, Assign Slaves and Clients, Customize Services, **Review**, Install, Start and Test, and Summary. The main content area is titled 'Review' and contains the following information:

- Admin Name :** admin
- Cluster Name :** Cisco_HDP
- Total Hosts :** 68 (68 new)
- Repositories:**
 - redhat6 (HDP-2.2): <http://rhel1/Hortonworks/HDP/centos6/2.x/GA/2.2.0.0>
 - redhat6 (HDP-UTILS-1.1.0.20): <http://rhel1/Hortonworks/HDP-UTILS-1.1.0.20/repos/centos6>
- Services:**
 - HDFS**
 - DataNode : 66 hosts
 - NameNode : rhel1
 - SNameNode : rhel2
 - YARN + MapReduce2**
 - App Timeline Server : rhel2
 - NodeManager : 66 hosts
 - ResourceManager : rhel2
 - Tez**
 - Client : 17 hosts

At the bottom of the main content area, there are '← Back' and 'Deploy →' buttons. A 'Print' button is also visible in the top right corner of the main content area.

The progress of the install is shown on the screen. Each component is installed and started and a simple test is run on the component. The next screen displays the overall status of the install in the progress bar at the top of the screen and a host-by-host status in the main section.

To see 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**, click the individual task to see the related log files. Select filter conditions by using the **Show** drop-down list. To see a larger version of the log contents, click the **Open** icon or to copy the contents to the clipboard, use the **Copy** icon.

Depending on which components are installing, the entire process may take 10 or more minutes.

When **successfully installed and started the services** appears, click **Next**.

Figure 173 *Installation and Test in Progress Page*

CLUSTER INSTALL WIZARD

- Get Started
- Select Stack
- Install Options
- Confirm Hosts
- Choose Services
- Assign Masters
- Assign Slaves and Clients
- Customize Services
- Review
- Install, Start and Test
- Summary

Install, Start and Test

Please wait while the selected services are installed and started.

100 % overall

Show: **All (68)** | [In Progress \(0\)](#) | [Warning \(0\)](#) | [Success \(68\)](#) | [Fail \(0\)](#)

Host	Status	Message
rhel1	<div style="background-color: #4CAF50; width: 100%; height: 10px; display: inline-block;"></div> 100%	Success
rhel2	<div style="background-color: #4CAF50; width: 100%; height: 10px; display: inline-block;"></div> 100%	Success
rhel3	<div style="background-color: #4CAF50; width: 100%; height: 10px; display: inline-block;"></div> 100%	Success
rhel4	<div style="background-color: #4CAF50; width: 100%; height: 10px; display: inline-block;"></div> 100%	Success
rhel5	<div style="background-color: #4CAF50; width: 100%; height: 10px; display: inline-block;"></div> 100%	Success
rhel7	<div style="background-color: #4CAF50; width: 100%; height: 10px; display: inline-block;"></div> 100%	Success
rhel8	<div style="background-color: #4CAF50; width: 100%; height: 10px; display: inline-block;"></div> 100%	Success
rhel9	<div style="background-color: #4CAF50; width: 100%; height: 10px; display: inline-block;"></div> 100%	Success
rhel10	<div style="background-color: #4CAF50; width: 100%; height: 10px; display: inline-block;"></div> 100%	Success
rhel11	<div style="background-color: #4CAF50; width: 100%; height: 10px; display: inline-block;"></div> 100%	Success
rhel12	<div style="background-color: #4CAF50; width: 100%; height: 10px; display: inline-block;"></div> 100%	Success
rhel13	<div style="background-color: #4CAF50; width: 100%; height: 10px; display: inline-block;"></div> 100%	Success
rhel14	<div style="background-color: #4CAF50; width: 100%; height: 10px; display: inline-block;"></div> 100%	Success
rhel15	<div style="background-color: #4CAF50; width: 100%; height: 10px; display: inline-block;"></div> 100%	Success

Summary of Install Process

The Summary page gives a summary of the accomplished tasks. Click **Complete**.

Figure 174 Summary

CLUSTER INSTALL WIZARD

- Get Started
- Select Stack
- Install Options
- Confirm Hosts
- Choose Services
- Assign Masters
- Assign Slaves and Clients
- Customize Services
- Review
- Install, Start and Test
- Summary

Summary

Here is the summary of the install process.

The cluster consists of (66) hosts
 Installed and started services successfully on (66) new hosts

Master services installed

- SNameNode installed on rhel2
- NameNode installed on rhel1
- History Server installed on rhel2
- ResourceManager installed on rhel2
- Nagios Server installed on rhel1
- Ganglia Server installed on rhel1
- HiveServer2 installed on rhel2
- HBase Master installed on rhel2
- Oozie Server installed on rhel1



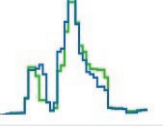
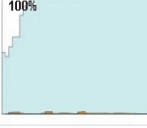
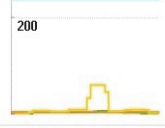





All services started
 All tests passed
 Install and start completed in 13 minutes and 48 seconds

Complete →

- HDFS
- MapReduce2
- YARN
- Tez
- Nagios
- Ganglia
- Hive
- HBase
- Pig
- Sqoop
- Oozie
- ZooKeeper
- Falcon
- Storm
- Flume
- Kafka
- Knox
- Slider

Actions ▾

Metrics ▾
Heatmaps
Config History

HDFS Disk Usage 	DataNodes Live 66/66	HDFS Links NameNode Secondary NameNode 66 DataNodes <small>More... ▾</small>	Memory Usage 	Network Usage 
CPU Usage 	Cluster Load 	NameNode Heap 	NameNode RPC 0.31 ms	NameNode CPU WIO 
NameNode Uptime 615.9 s	HBase Master Heap 	HBase Links HBase Master 66 RegionServers Master Web UI <small>More... ▾</small>	HBase Ave Load 1	HBase Master Uptime 560.7 s
ResourceManager Heap 	ResourceManager Uptime 502.9 s	NodeManagers Live 66/66	YARN Memory 	Supervisors Live 66/66

Archival Storage in HDFS

HDFS along with Apache Hadoop version 2.6.0 (included in HDP 2.2) provides the ability to utilize tiered (heterogeneous) storage media within the HDFS cluster to enable Archival storage to store less frequently accessed datasets and Flash based fast storage to store datasets that require lower latencies for read/write workloads.

In this section, we will go into details on using Cisco UCS C3160 as the Archival Storage.

Over time, frequency of reads on a dataset decreases (recently written data is more frequently accessed while aged data is less frequently accessed). This old dataset is deemed as “cold.” As the amount of data under storage grows, there is a need to optimize storage of such ‘cold’ datasets. An Archival storage tier, consisting of nodes with slow spinning high density storage drives and low compute power, provides cost efficiency for storing these “cold” datasets.

HDP 2.2 introduces an ‘ARCHIVE’ Storage Type and related Storage Policies – ‘Hot’, ‘Warm’, ‘Cold’.

The placement of data blocks is tied to the temperature of the data. The default storage type is DISK, which is on Data Nodes; by default every data is considered Hot and placed in Hot data-tier storage (UCS C240 M4).

Figure 175 HDP 2.2 Archive Storage Type



For WARM option, 2 replicas are placed on the drives of Data Nodes and one replica in Archival Storage.

In COLD option, all replicas are placed in Archival storage.

A directory within HDFS is then assigned a storage policy. Depending on the policy, data copied into that directory is then stored in the appropriately tagged data nodes.

For example – assign directory “/cold-data” to COLD storage policy

```
hdfs dfsadmin -setStoragePolicy /cold-data COLD
```

All disk volumes in the Archival storage tier nodes are tagged with the ‘ARCHIVE’ storage type. Administrators can then apply the ‘Cold’ Storage Policy to datasets that need to be stored on the Archival storage tier nodes. Since the ‘Cold’ Storage Policy is applied after the dataset has been created, the policy will be enforced when the HDFS Mover tool is run.

Sections below will go in detail on configuring Archival Storage.



Note

If HDP 2.2 is deployed on only a single Rack with 16 Cisco UCS C240 M4 servers along with only one Cisco UCS C3160 as the archival node, then as a best practice, use only Hot and Warm policies to ensure data redundancy. This is because, using COLD won’t provide 3 way replication as only one copy will be stored in Archival storage (which provides data redundancy through RAID5 but is not equivalent to 3 way replication). If there are two Archival nodes, this would still provide strong data redundancy, as there are two copies, one on each Archival node, which is further protected through RAID5 protection.

Steps:

1. Shutdown all data nodes.

Go to **Hosts > Actions > Filtered Hosts > DataNodes > Stop**

Figure 176 Ambari GUI: Stop the DataNodes

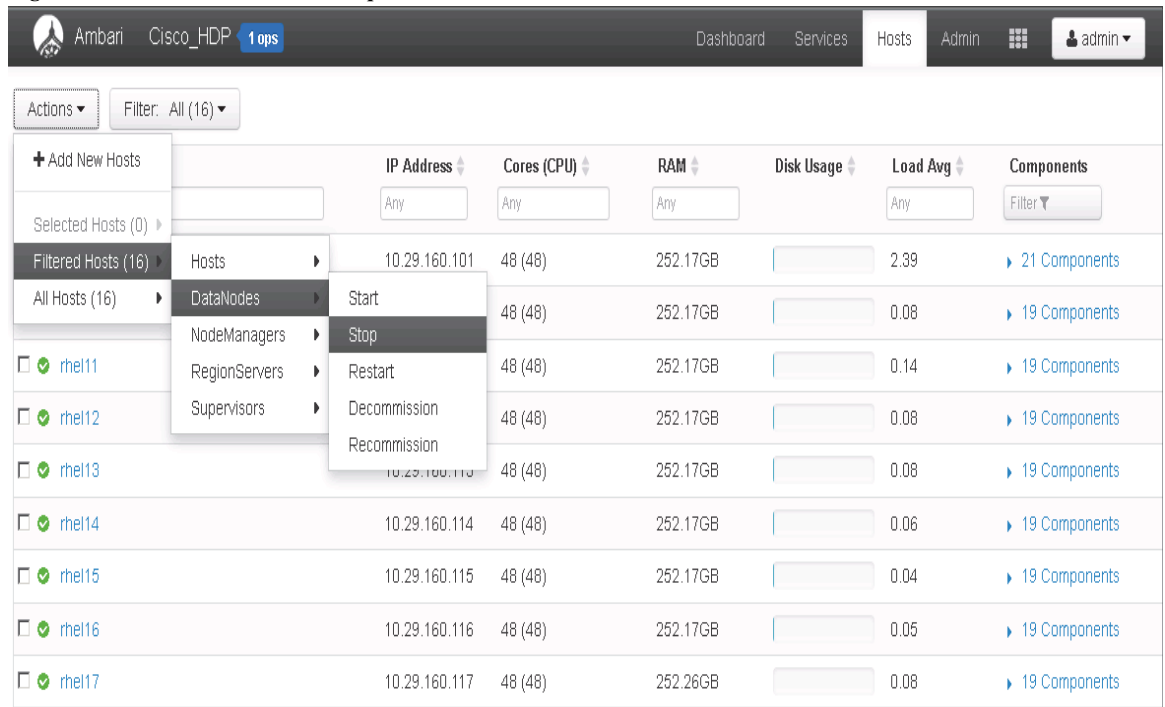



Figure 177 DataNodes Successfully Stopped



2. Assign the Archive Storage Type to the DataNode
Go to **HDFS > Config > Select Archive** in the Config groups

Edit all archive directories and add [ARCHIVE] as a prefix in front of each entry

```
[ARCHIVE] /data/disk1/hadoop/hdfs/data
[ARCHIVE] /data/disk2/hadoop/hdfs/data
[ARCHIVE] /data/disk3/hadoop/hdfs/data
[ARCHIVE] /data/disk4/hadoop/hdfs/data
```

Figure 178 HDFS Configuration

Save the configuration

3. Create cold and warm archive HDFS Directories.

```
sudo -u hdfs hadoop dfs -mkdir /archive-cold
sudo -u hdfs hadoop dfs -mkdir /archive-warm
```

```
[root@rhell ~]# sudo -u hdfs hadoop dfs -mkdir /archive-cold
DEPRECATED: Use of this script to execute hdfs command is deprecated.
Instead use the hdfs command for it.

[root@rhell ~]# sudo -u hdfs hadoop dfs -mkdir /archive-warm
DEPRECATED: Use of this script to execute hdfs command is deprecated.
Instead use the hdfs command for it.
```

4. Set Storage Policies on HDFS Directory and then confirm these

```
sudo -u hdfs hdfs dfsadmin -setStoragePolicy /archive-cold COLD
sudo -u hdfs hdfs dfsadmin -setStoragePolicy /archive-warm WARM
sudo -u hdfs hdfs dfsadmin -getStoragePolicy /archive-cold
sudo -u hdfs hdfs dfsadmin -getStoragePolicy /archive-warm
```

```
[root@rhel1 ~]# sudo -u hdfs hdfs dfsadmin -setStoragePolicy /archive-cold COLD
Set storage policy COLD on /archive-cold
[root@rhel1 ~]# sudo -u hdfs hdfs dfsadmin -setStoragePolicy /archive-warm WARM
Set storage policy WARM on /archive-warm
[root@rhel1 ~]# sudo -u hdfs hdfs dfsadmin -getStoragePolicy /archive-cold
The storage policy of /archive-cold:
BlockStoragePolicy{COLD:2, storageTypes=[ARCHIVE], creationFallbacks=[], replicationFallbacks=[]}
[root@rhel1 ~]# sudo -u hdfs hdfs dfsadmin -getStoragePolicy /archive-warm
The storage policy of /archive-warm:
BlockStoragePolicy{WARM:5, storageTypes=[DISK, ARCHIVE], creationFallbacks=[DISK, ARCHIVE], replicationFallbacks=[DISK, ARCHIVE]}
```

5. Start all the DataNodes.

Click **HDFS > Restart all > Confirm restart all**

Figure 179 Ambari GUI: Restart HDFS DataNodes

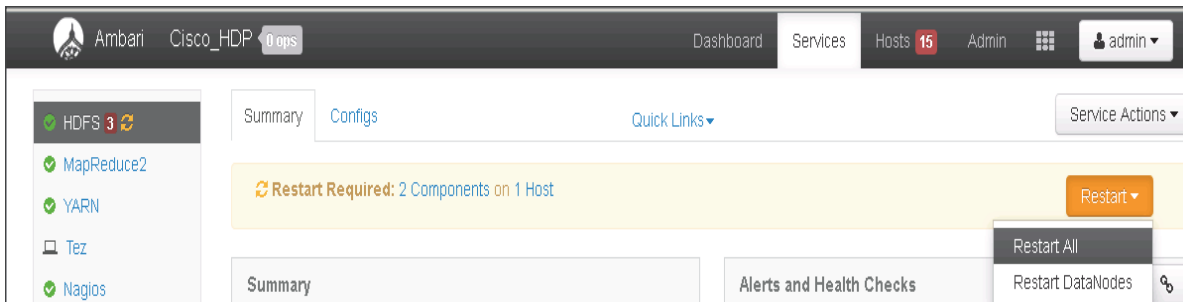
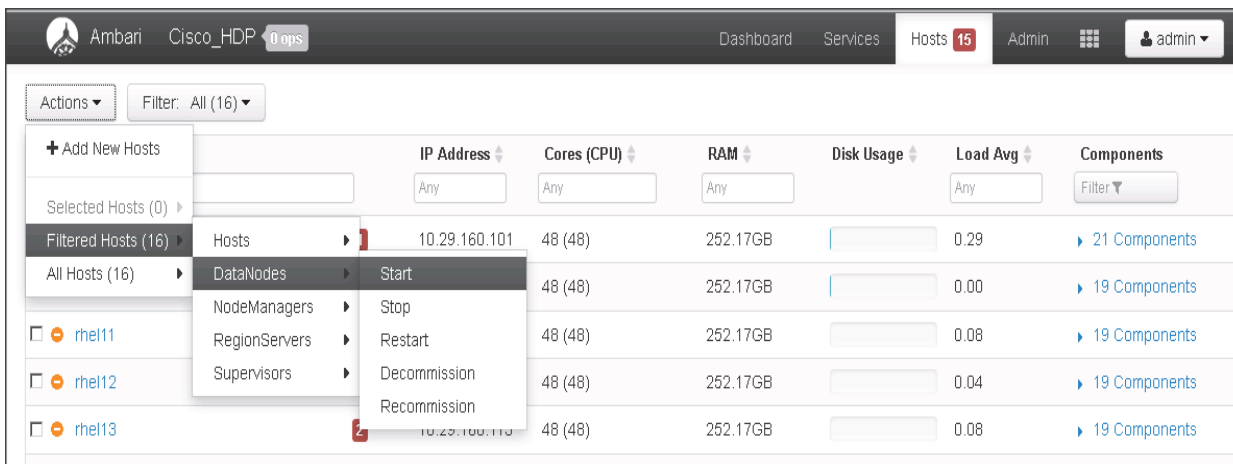


Figure 180 Ambari: Restart DataNodes



6. Copy test files into Warm and Cold Storage HDFS Directories.

Create your test file and copy into warm and cold HDFS directories.

```
echo "This is a tiered storage test file" > /tmp/storagetest.txt
sudo -u hdfs hadoop fs -put /tmp/storagetest.txt /archive-cold
sudo -u hdfs hadoop fs -put /tmp/storagetest.txt /archive-warm
sudo -u hdfs hadoop fs -cat /archive-cold/storagetest.txt
sudo -u hdfs hadoop fs -cat /archive-warm/storagetest.txt
```

```
[root@rhell1 tmp]# echo "This is a tiered storage test file" > /tmp/storagetest.txt
[root@rhell1 tmp]# sudo -u hdfs hadoop fs -put /tmp/storagetest.txt /archive-cold
[root@rhell1 tmp]# sudo -u hdfs hadoop fs -put /tmp/storagetest.txt /archive-warm
[root@rhell1 tmp]# sudo -u hdfs hadoop fs -cat /archive-cold/storagetest.txt
This is a tiered storage test file
[root@rhell1 tmp]# sudo -u hdfs hadoop fs -cat /archive-warm/storagetest.txt
This is a tiered storage test file
```

7. Use mover script to apply Storage Policies

```
sudo -u hdfs hdfs mover -p /archive-warm
sudo -u hdfs hdfs mover -p /archive-cold
```

```
[root@rhell1 ~]# sudo -u hdfs hdfs mover -p /archive-warm
15/03/18 03:25:06 INFO mover.Mover: namenodes = {hdfs://rhell1:8020=[/archive-warm]}
15/03/18 03:25:07 INFO balancer.KeyManager: Block token params received from NN: update interval=10hrs, 0sec, token lifetime=10hrs,
15/03/18 03:25:07 INFO block.BlockTokenSecretManager: Setting block keys
15/03/18 03:25:07 INFO balancer.KeyManager: Update block keys every 2hrs, 30mins, 0sec
15/03/18 03:25:07 INFO block.BlockTokenSecretManager: Setting block keys
15/03/18 03:25:07 INFO net.NetworkTopology: Adding a new node: /default-rack/192.168.11.113:50010
15/03/18 03:25:07 INFO net.NetworkTopology: Adding a new node: /default-rack/192.168.11.103:50010
15/03/18 03:25:07 INFO net.NetworkTopology: Adding a new node: /default-rack/192.168.11.105:50010
15/03/18 03:25:07 INFO net.NetworkTopology: Adding a new node: /default-rack/192.168.11.112:50010
15/03/18 03:25:07 INFO net.NetworkTopology: Adding a new node: /default-rack/192.168.11.107:50010
```

```
[root@rhell1 ~]# sudo -u hdfs hdfs mover -p /archive-cold
15/03/18 03:25:44 INFO mover.Mover: namenodes = {hdfs://rhell1:8020=[/archive-cold]}
15/03/18 03:25:45 INFO balancer.KeyManager: Block token params received from NN: update interval=10hrs, 0sec, token lifetime=10hrs,
15/03/18 03:25:45 INFO block.BlockTokenSecretManager: Setting block keys
15/03/18 03:25:45 INFO balancer.KeyManager: Update block keys every 2hrs, 30mins, 0sec
15/03/18 03:25:45 INFO block.BlockTokenSecretManager: Setting block keys
15/03/18 03:25:45 INFO net.NetworkTopology: Adding a new node: /default-rack/192.168.11.109:50010
15/03/18 03:25:45 INFO net.NetworkTopology: Adding a new node: /default-rack/192.168.11.112:50010
15/03/18 03:25:45 INFO net.NetworkTopology: Adding a new node: /default-rack/192.168.11.116:50010
15/03/18 03:25:45 INFO net.NetworkTopology: Adding a new node: /default-rack/192.168.11.105:50010
```

8. Run fsck to check the number of replicas and locations of the blocks

```
sudo -u hdfs hadoop fsck -racks -locations -blocks -files \
/archive-cold/storagetest.txt
sudo -u hdfs hadoop fsck -racks -locations -blocks -files \
/archive-warm/storagetest.txt
```

```

[root@rhell1 ~]# sudo -u hdfs hadoop fsck -racks -locations -blocks -files /archive-warm/storagetest.txt
DEPRECATED: Use of this script to execute hdfs command is deprecated.
Instead use the hdfs command for it.

Connecting to namenode via http://rhell1:50070
FSCK started by hdfs (auth:SIMPLE) from /192.168.11.101 for path /archive-warm/storagetest.txt at Wed Mar 18 03:29:47 EDT 2015
/archive-warm/storagetest.txt 35 bytes, 1 block(s): OK
0. BR-857809477-192.168.11.101-1426658417364:blk_1073742388_1637 len=35 repl=3 [/default-rack/192.168.11.107:50010, /default-rack/192.168.11.109:50010, /default-rack/192.168.11.109:50010]

Status: HEALTHY
Total size: 35 B
Total dirs: 0
Total files: 1
Total symlinks: 0
Total blocks (validated): 1 (avg. block size 35 B)
Minimally replicated blocks: 1 (100.0 %)
Over-replicated blocks: 0 (0.0 %)
Under-replicated blocks: 0 (0.0 %)
Mis-replicated blocks: 0 (0.0 %)
Default replication factor: 3
Average block replication: 3.0
Corrupt blocks: 0
Missing replicas: 0 (0.0 %)
Number of data-nodes: 14
Number of racks: 1
FSCK ended at Wed Mar 18 03:29:47 EDT 2015 in 1 milliseconds

The filesystem under path '/archive-warm/storagetest.txt' is HEALTHY

```

Under the WARM storage policy, two replicas are placed on data nodes while 1 replica is placed in Archival node. As shown in the screenshot, two replicas are stored in data nodes and one in Archival node.

Configuring Hive to use Apache Tez

Apache Hive was originally built for large-scale operational batch processing and it is very effective with reporting, data mining and data preparation use cases. These usage patterns remain very important but with widespread adoption of Hadoop, the enterprise requirement for Hadoop to become more real time or interactive has increased in importance as well.

With the Stinger initiative, Hive query time has improved dramatically, enabling Hive to support both batch and interactive workloads at speed and at scale.

Stinger Initiative was designed to enable Hive to answer human-time use cases (i.e. queries in the 5-30 second range) such as big data exploration, visualization, and parameterized reports through faster performance improvement to hive. One of the main change with Stinger Initiative was to run Hive queries with Apache Tez execution engine instead of the Map-reduce engine. Apache Tez innovations drove many of the Hive performance improvements delivered by the Stinger Initiative

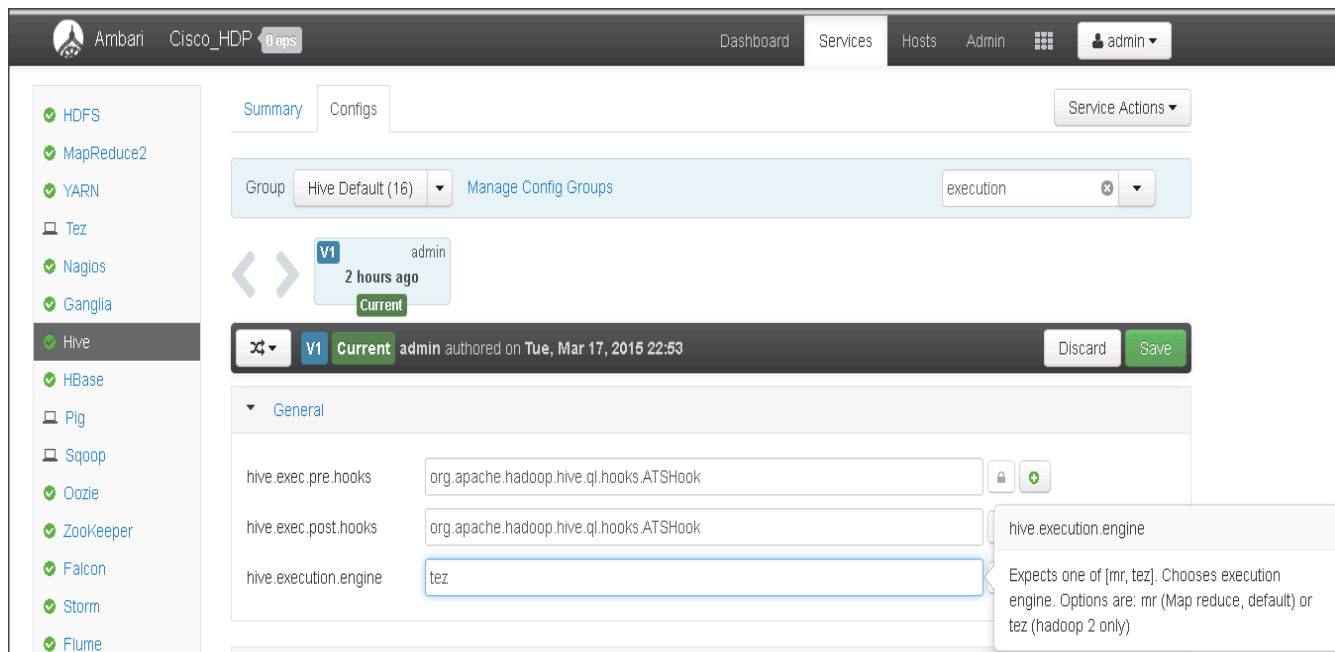
Following are the configurations to enable Hive for faster query (for more details such as using ORCFile, refer <http://hortonworks.com/blog/5-ways-make-hive-queries-run-faster/>)

Set Hive to use Apache Tez

In Ambari, first select **Hive** and navigate to **Configs** tab. Then, in the filter text box type “**execution**” to find the “hive.execution.engine” property. Change its value from **mr** to **tez** as shown below:

hive.execution.engine=tez

Figure 181 Ambari: Configuring Hive to use Apache Tez



Enable Vectorization

Vectorized query in execution improves performance of operations like scans, aggregations, filters and joins, by performing them in batches of 1024 rows at once instead of single row each time.

Instructions to enable vectorization: In Ambari, first select Hive and navigate to Configs tab. Then, in the filter text box type “execution” to find the “hive.vectorized.execution.enabled” and “hive.vectorized.execution.reduce.enabled” properties. Ensure both their values are set to “true”

```
hive.vectorized.execution.enabled = true
hive.vectorized.execution.reduce.enabled = true
```

Figure 182 Ambari: Set Preferences - Vectorization for Hive

The screenshot shows the Ambari configuration page for Hive. The left sidebar lists various services, with 'Hive' selected. The main content area is divided into 'Summary' and 'Configs' tabs. The 'Configs' tab is active, showing a configuration group 'Hive Default (16)' and a filter 'execution'. The configuration version is 'V1' and it was last updated '2 hours ago' by 'admin'. The configuration is currently 'Current' and was authored on 'Tue, Mar 17, 2015 22:53'. The 'General' section contains the following properties:

- hive.exec.pre.hooks: org.apache.hadoop.hive.ql.hooks.ATSHook
- hive.exec.post.hooks: org.apache.hadoop.hive.ql.hooks.ATSHook
- hive.execution.engine: tez

The 'Performance' section contains the following properties:

- hive.vectorized.execution.enabled: true
- hive.vectorized.execution.reduce.enabled: true

A tooltip for the 'hive.vectorized.execution.reduce.enabled' property states: "This flag should be set to true to enable vectorized mode of the reduce-side of query execution. The default value is true."

Enable cost based query optimization

Background: Hive optimizes each query's logical and physical execution plan before submitting for final execution. A recent addition to Hive, Cost-based optimization, performs further optimizations based on query cost, resulting in potentially different decisions: how to order joins, which type of join to perform, degree of parallelism and others.

Enable cost-based optimization (also known as CBO): In Ambari, first select Hive and navigate to Configs tab. Then, in the filter text box type "stats" to find the "hive.compute.query.using.stats", "hive.stats.fetch.column.stats" and "hive.stats.fetch.partition.stats" properties. Ensure all their values are set to "true". Then, in the filter text box type "cbo" and ensure that "hive.compute.query.using.stats" property is also set to true

```
hive.cbo.enable=true
hive.compute.query.using.stats=true
hive.stats.fetch.column.stats=true
hive.stats.fetch.partition.stats=true
```


Figure 183 Ambari: Set Preferences - Query Optimization for Hive

The screenshot displays the Ambari configuration interface for Hive. On the left, a sidebar lists various services, with 'Hive' selected. The main content area shows the 'Configs' tab for the 'Hive Default (16)' group. A version control bar indicates the current configuration (V1) was updated by 'admin' 2 hours ago. Below this, the 'Performance' section is expanded, showing the following configurations:

Property Name	Value
hive.stats.autogather	true
hive.stats.fetch.column.stats	true
hive.compute.query.using.stats	true

A tooltip for the 'hive.stats.fetch.column.stats' property provides the following description: "Annotation of operator tree with statistics information requires column statistics. Column statistics are fetched from metastore. Fetching column statistics for each needed column can be expensive when the number of columns is high. This flag can be used to disable fetching of column statistics from metastore."

Conclusion

Hadoop has evolved into a leading data management platform across all verticals. The Cisco UCS Integrated Infrastructure for Big Data with Hortonworks (HDP 2.2) with Tiered Storage offers a dependable deployment model for enterprise Hadoop that offers a fast and predictable path for businesses to unlock value in Big Data while providing customer with storage archival to ensure faster drives are used for Hot and Warm data.

The configuration detailed in the document can be extended to clusters of various sizes depending on what application demands. Up to 80 servers (5 racks) can be supported with no additional switching in a single UCS domain with no network over-subscription. Scaling beyond 5 racks (80 servers) can be implemented by interconnecting multiple UCS domains using Nexus 6000/7000 Series switches, scalable to thousands of servers and to hundreds of petabytes storage, and managed from a single pane using [UCS Central](#).

Bill of Materials

This section provides the BOM for 64 nodes Performance Optimized Cluster with 4 nodes Cisco UCS C3160 for Archival nodes. See Table 14 and Table 15 for BOM for the master rack, Table 16 and Table 17 for BOM for expansion racks (rack 2 to 4), Table 18 and 19 for software components.

Table 14 Bill of Materials for C240M4SX Base rack

Part Number	Description	Quantity
UCS-SL-CPA3-P	Performance Optimized Cluster	1
UCSC-C240-M4SX	UCS C240 M4 SFF 24 HD w/o CPU, mem, HD, PCIe, PS, railkt w/expndr	16
UCSC-MRAID12G	Cisco 12G SAS Modular Raid Controller	16
UCSC-MRAID12G-2GB	Cisco 12Gbps SAS 2GB FBWC Cache module (Raid 0/1/5/6)	16
UCSC-MLOM-CSC-02	Cisco UCS VIC1227 VIC MLOM - Dual Port 10Gb SFP+	16
CAB-9K12A-NA	Power Cord 125VAC 13A NEMA 5-15 Plug North America	32
UCSC-PSU2V2-1200W	1200W V2 AC Power Supply for 2U C-Series Servers	32
UCSC-RAILB-M4	Ball Bearing Rail Kit for C220 M4 and C240 M4 rack servers	16
UCSC-HS-C240M4	Heat Sink for UCS C240 M4 Rack Server	32
UCSC-SCCBL240	Supercap cable 250mm	16
UCS-CPU-E52680D	2.50 GHz E5-2680 v3/120W 12C/30MB Cache/DDR4 2133MHz	32
UCS-MR-1X162RU-A	16GB DDR4-2133-MHz RDIMM/PC4-17000/dual rank/x4/1.2v	256
UCS-HD12T10KS2-E	1.2 TB 6G SAS 10K rpm SFF HDD	384
UCS-SD120G0KSB-EV	120 GB 2.5 inch Enterprise Value 6G SATA SSD (BOOT)	32
UCSC-PCI-1C-240M4	Right PCI Riser Bd (Riser 1) 2onbd SATA bootdrvs+ 2PCI slts	16
UCS-FI-6296UP-UPG	UCS 6296UP 2RU Fabric Int/No PSU/48 UP/ 18p LIC	2
CON-SNTP-C240M4SX	SMARTNET 24X7X4 UCS C240 M4 SFF 24 HD w/o CPU, mem	16
CON-SNTP-FI6296UP	SMARTNET 24X7X4 UCS 6296UP 2RU Fabric Int/2 PSU/4 Fans	2
SFP-H10GB-CU3M	10GBASE-CU SFP+ Cable 3 Meter	32
UCS-ACC-6296UP	UCS 6296UP Chassis Accessory Kit	2
UCS-PSU-6296UP-AC	UCS 6296UP Power Supply/100-240VAC	4
N10-MGT012	UCS Manager v2.2	2
UCS-L-6200-10G-C	2rd Gen FI License to connect C-direct only	70
UCS-BLKE-6200	UCS 6200 Series Expansion Module Blank	6
UCS 6296UP Fan Module	UCS 6296UP Fan Module	8

Table 14 Bill of Materials for C240M4SX Base rack

Part Number	Description	Quantity
CAB-9K12A-NA	Power Cord 125VAC 13A NEMA 5-15 Plug North America	4
UCS-FI-E16UP	UCS 6200 16-port Expansion module/16 UP/ 8p LIC	4
RACK-UCS2	Cisco R42610 standard rack w/side panels	1
CON-OS-R42610	ONSITE 8X5XNBD Cisco R42610 expansion rack no side panel	1
RP208-30-1P-U-2=	Cisco RP208-30-U-2 Single Phase PDU 20x C13 4x C19 (Country Specific)	2
CON-OS-RPDUX	ONSITE 8X5XNBD Cisco RP208-30-U-X Single Phase PDU 2x	2

Table 15 Bill of Materials for C3160 Base Rack

Part Number	Description	Quantity
UCS-SA-C3160-D	Extreme Capacity	1
UCSC-C3160	Cisco UCS C3160 Base Chassis w/ 4x PSU, 2x120GB SSD RailKit	1
CAB-9K12A-NA	Replace with Power Cord 125VAC 13A NEMA 5-15 Plug North America	4
UCSC-C3X60-SBLKP	UCS C3x60 SIOC blanking plate	1
UCSC-PSU1-1050W	UCS C3X60 1050W Power Supply Unit	4
UCSC-C3X60-12SSD	Cisco UCS C3X60 2x120GB SATA Enterprise Value SSD	2
UCSC-C3X60-RAIL	UCS C3X60 Rack Rails Kit	1
UCSC-C3X60-SVRN4	Cisco C3X60 Server Node E5-2695 v2 CPU 256GB 4GB RAID cache	1
UCSC-HS-C3X60	Cisco UCS C3X60 Server Node CPU Heatsink	2
UCS-CPU-E52695B	2.40 GHz E5-2695 v2/115W 12C/30MB Cache/DDR3 1866MHz	2
UCS-MR-1X162RZ-A	16GB DDR3-1866-MHz RDIMM/PC3-14900/dual rank/x4/1.5v	16
UCSC-C3X60-R4GB	UCS C3X60 12G SAS RAID Controller with 4GB cache	1
UCSC-C3160-SIOC	Cisco UCS C3160 System IO Controller with mLOM mez adapter	1
UCSC-C3X60-56HD4	Cisco UCS C3X60 Four rows 56x 4TB Drives	1
UCSC-C3X60-HD4TB	UCS C3X60 4TB NL-SAS 7.2K HDD including C3X60 HDD carrier Top-load	56
UCSC-C3X60-EX16T	Cisco UCS C3X60 Disk Exp Tray w/ 4x 4TB	1

Part Number	Description	Quantity
UCS-HD4T7KS3-E	4TB SAS 7.2K RPM 3.5 inch HDD/hot plug/drive sled mounted	4
UCSC-MLOM-CSC-02	Cisco UCS VIC1227 VIC MLOM - Dual Port 10Gb SFP+	1
CON-SNTP-C3160VD1	UCS C3160 BD D Server	1
SFP-H10GB-CU3M	10GBASE-CU SFP+ Cable 3 Meter	2



Note If using 6 TB drives for C3160, use the following PID instead of 4TB drives.

UCSC-C3X60-56HD6	Cisco UCS C3X60 Four row of drives containing 56 x 6TB (Total)	1
UCSC-C3X60-HD6TB	UCS C3X60 6TB 12Gbps NL-SAS 7200RPM HDD w carrier- Top-load	56
UCSC-C3X60-EX24T	Cisco UCS C3160 Expander with 4x 6TB 7200RPM NL-SAS Drives	1
UCSC-C3X60-6TBRR	UCS C3X60 6TB 12Gbps NL-SAS 7200RPM HDD w carrier- Rear-load	4

Table 16 Bill of Materials for Expansion Racks

Part Number	Description	Quantity
UCSC-C240-M4SX	UCS C240 M4 SFF 24 HD w/o CPU, mem, HD, PCIe, PS, railkt w/expndr	48
UCSC-MRAID12G	Cisco 12G SAS Modular Raid Controller	48
UCSC-MRAID12G-2GB	Cisco 12Gbps SAS 2GB FBWC Cache module (Raid 0/1/5/6)	48
UCSC-MLOM-CSC-02	Cisco UCS VIC1227 VIC MLOM - Dual Port 10Gb SFP+	48
CAB-9K12A-NA	Power Cord 125VAC 13A NEMA 5-15 Plug North America	96
UCSC-PSU2V2-1200W	1200W V2 AC Power Supply for 2U C-Series Servers	96
UCSC-RAILB-M4	Ball Bearing Rail Kit for C220 M4 and C240 M4 rack servers	48
UCSC-HS-C240M4	Heat Sink for UCS C240 M4 Rack Server	48
UCSC-SCCBL240	Supercap cable 250mm	48
UCS-CPU-E52680D	2.50 GHz E5-2680 v3/120W 12C/30MB Cache/DDR4 2133MHz	96

Part Number	Description	Quantity
UCS-MR-1X162RU-A	16GB DDR4-2133-MHz RDIMM/PC4-17000/dual rank/x4/1.2v	768
UCS-HD12T10KS2-E	1.2 TB 6G SAS 10K rpm SFF HDD	1152
UCS-SD120G0KSB-EV	120 GB 2.5 inch Enterprise Value 6G SATA SSD (BOOT)	96
UCSC-PCI-1C-240M4	Right PCI Riser Bd (Riser 1) 2onbd SATA bootdrvs+ 2PCI slts	48
SFP-H10GB-CU3M=	10GBASE-CU SFP+ Cable 3 Meter	96
CON-SNTP-C240M4SX	SMARTNET 24X7X4 UCS C240 M4 SFF 24 HD w/o CPU, mem	48
RACK-UCS2	Cisco R42610 standard rack w/side panels	3
CON-OS-R42610	ONSITE 8X5XNBD Cisco R42610 expansion rack no side panel	3
RP208-30-1P-U-2=	Cisco RP208-30-U-2 Single Phase PDU 20x C13 4x C19 (Country Specific)	6
CON-OS-RPDUX	ONSITE 8X5XNBD Cisco RP208-30-U-X Single Phase PDU 2x	3

Table 17 *Bill of Materials for C3160 Expansion Rack*

Part Number	Description	Quantity
UCSC-C3160	Cisco UCS C3160 Base Chassis w/ 4x PSU, 2x120GB SSD RailKit	3
CAB-9K12A-NA	Replace with Power Cord 125VAC 13A NEMA 5-15 Plug North America	12
UCSC-C3X60-SBLKP	UCS C3x60 SIOC blanking plate	3
UCSC-PSU1-1050W	UCS C3X60 1050W Power Supply Unit	12
UCSC-C3X60-12SSD	Cisco UCS C3X60 2x120GB SATA Enterprise Value SSD	6
UCSC-C3X60-RAIL	UCS C3X60 Rack Rails Kit	3
UCSC-C3X60-SVRN4	Cisco C3X60 Server Node E5-2695 v2 CPU 256GB 4GB RAID cache	3
UCSC-HS-C3X60	Cisco UCS C3X60 Server Node CPU Heatsink	6
UCS-CPU-E52695B	2.40 GHz E5-2695 v2/115W 12C/30MB Cache/DDR3 1866MHz	6
UCS-MR-1X162RZ-A	16GB DDR3-1866-MHz RDIMM/PC3-14900/dual rank/x4/1.5v	48
UCSC-C3X60-R4GB	UCS C3X60 12G SAS RAID Controller with 4GB cache	3
UCSC-C3160-SIOC	Cisco UCS C3160 System IO Controller with mLOM mez adapter	3
UCSC-C3X60-56HD4	Cisco UCS C3X60 Four rows 56x 4TB Drives	3

Part Number	Description	Quantity
UCSC-C3X60-HD4TB	UCS C3X60 4TB NL-SAS 7.2K HDD including C3X60 HDD carrier Top-load	168
UCSC-C3X60-EX16T	Cisco UCS C3X60 Disk Exp Tray w/ 4x 4TB	3
UCS-HD4T7KS3-E	4TB SAS 7.2K RPM 3.5 inch HDD/hot plug/drive sled mounted	12
UCSC-MLOM-CSC-02	Cisco UCS VIC1227 VIC MLOM - Dual Port 10Gb SFP+	3
SFP-H10GB-CU3M	10GBASE-CU SFP+ Cable 3 Meter	6
CON-SNTP-C3160VD1	UCS C3160 BD D Server	3

Table 18 Red Hat Enterprise Linux License

Red Hat Enterprise Linux		
Part Number	Description	Quantity
RHEL-2S-1G-3A	Red Hat Enterprise Linux	68
CON-ISV1-RH2S1G3A	3 year Support for Red Hat Enterprise Linux	68

Table 19 Bill of Materials for Hortonworks



Note Choose one of the part numbers.

Part Number	Description	Quantity
UCS-BD-HDP-ENT=	HORTONWORKS ENTERPRISE EDITION	68
UCS-BD-HDP-EPL=	HORTONWORKS ENTERPRISE PLUS EDITION	68

Appendix

Cisco UCS Director Express 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.

As you consider Hadoop to meet your growing data and business needs, operational challenges often emerge. Despite its compelling advantages, Hadoop clusters can be difficult, complex, and time consuming to deploy. Moreover, with so much data increasing so quickly, there is a need to find ways to consistently deploy Hadoop clusters and manage them efficiently.

**Note**

The UCSD Express appliances (UCSD Express VM and Baremetal Agent VM) can also be installed on an existing VMware ESXi server with proper network connectivity (See Figure 174) to the UCS domain that manages the Hadoop servers. In such a case, skip the sections until Downloading the UCS Director Express software components.

UCS Director Express for Big Data

Cisco UCS Director Express for Big Data provides a single-touch solution that automates deployment of Hadoop on Cisco UCS Common Platform Architecture (CPA) for Big Data infrastructure. It also provides a single management pane across both Cisco UCS integrated infrastructure and Hadoop software. All elements of the infrastructure are handled automatically with little need for user input. Through this approach, configuration of physical computing, internal storage, and networking infrastructure is integrated with the deployment of operating systems, Java packages, and Hadoop along with the provisioning of Hadoop services. Cisco UCS Director Express for Big Data is integrated with major Hadoop distributions from Cloudera, MapR, and Hortonworks, providing single-pane management across the entire infrastructure. It complements and communicates with Hadoop managers, providing a system wide perspective and enabling administrators to correlate Hadoop activity with network and computing activity on individual Hadoop nodes.

Key features of UCS Director (UCSD) Express for Big Data

- **Faster and Easier Big Data Infrastructure Deployment:** Cisco UCS Director Express for Big Data extends the Cisco UCS Integrated Infrastructure for Big Data with one-click provisioning, installation, and configuration, delivering a consistent, repeatable, flexible, and reliable end-to-end Hadoop deployment.
- **Massive Scalability and Performance:** Cisco's unique approach provides appliance-like capabilities for Hadoop with flexibility that helps ensure that resources are deployed right the first time and can scale without arbitrary limitations.
- **Centralized Visibility:** Cisco UCS Director Express for Big Data provides centralized visibility into the complete infrastructure to identify potential failures and latent threats before they affect application and business performance.
- **Open and Powerful:** Provides open interfaces that allows further integration into third-party tools and services while allowing flexibility for your own add-on services.

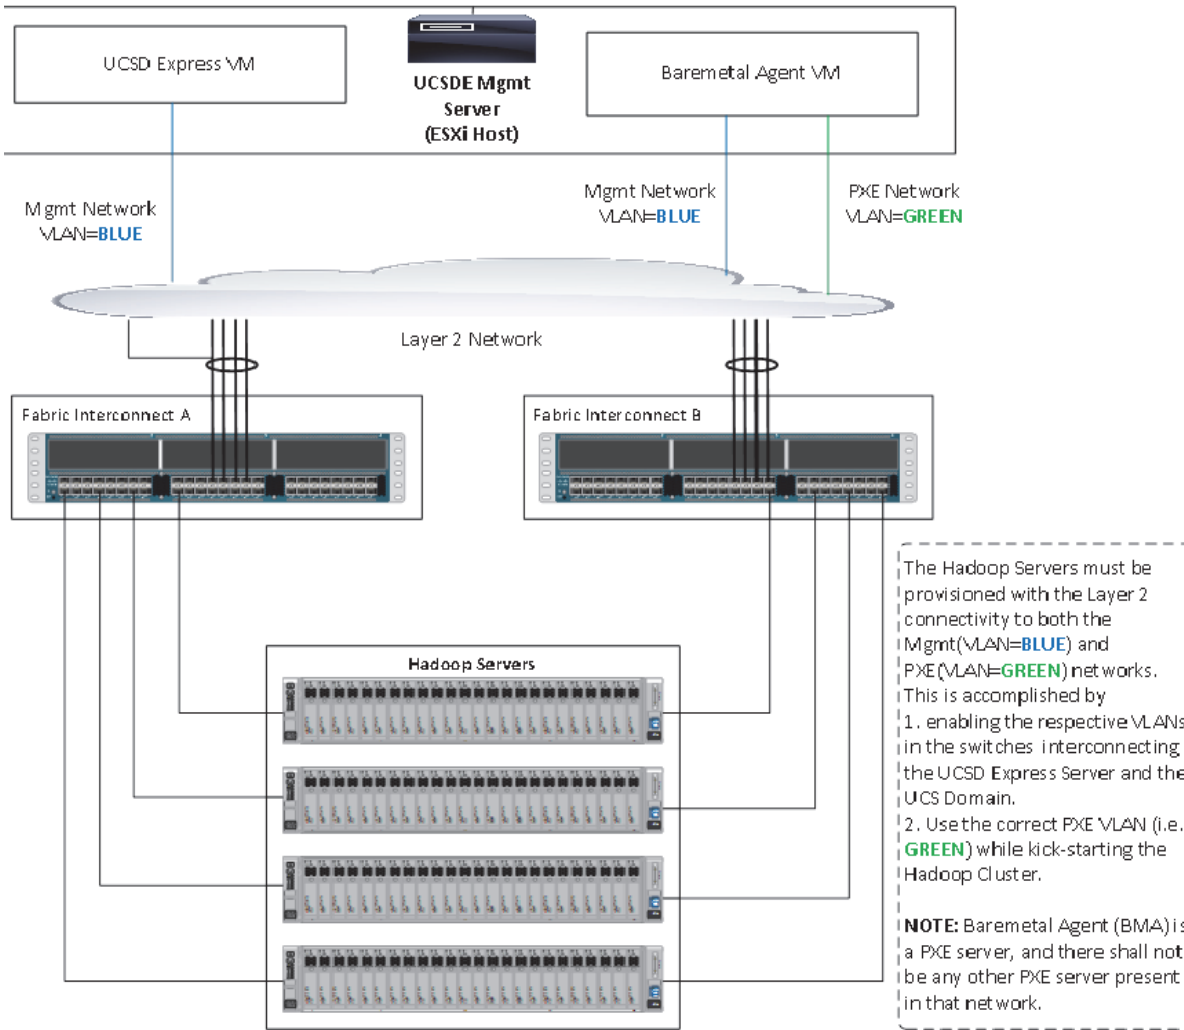
UCSD Express Management Server Configuration

The basic requirement for deploying and executing the UCSD Express software is a server with VMware ESXi based virtualization environment. Such a physical server machine with ESXi must be connected to the target Hadoop servers in the UCS domain by means of the management network and a dedicated PXE network.

The following are the potential network topologies:

1. The UCSD Express Management server is outside of the UCS Domain containing the C-Series servers that would be used to form the Hadoop cluster. For example, a standalone (CIMC managed) C220 M4 rack server provisioned with UCSD Express VMs is connected to the UCS Domain

Figure 184 UCSD Express Management Server that lives outside the UCS Domain




2. The UCSD Express Management server is hosted on a C220 M4 rack server that is connected to and managed by the same UCS Domain. This is the method used in this document.

Figure 185 UCS Express Management Server that is being managed as part of the same UCS Domain

The BMA VM is hosted on the UCSDE Mgmt server located within the UCS Domain. The BMA-VM's PXE interface (eth1) should be provisioned on the Fabric Interconnect B to avoid the PXE traffic leaving the UCS Domain.

NOTE: Baremetal Agent (BMA) is a PXE server, and there shall not be any other PXE server present in that network.

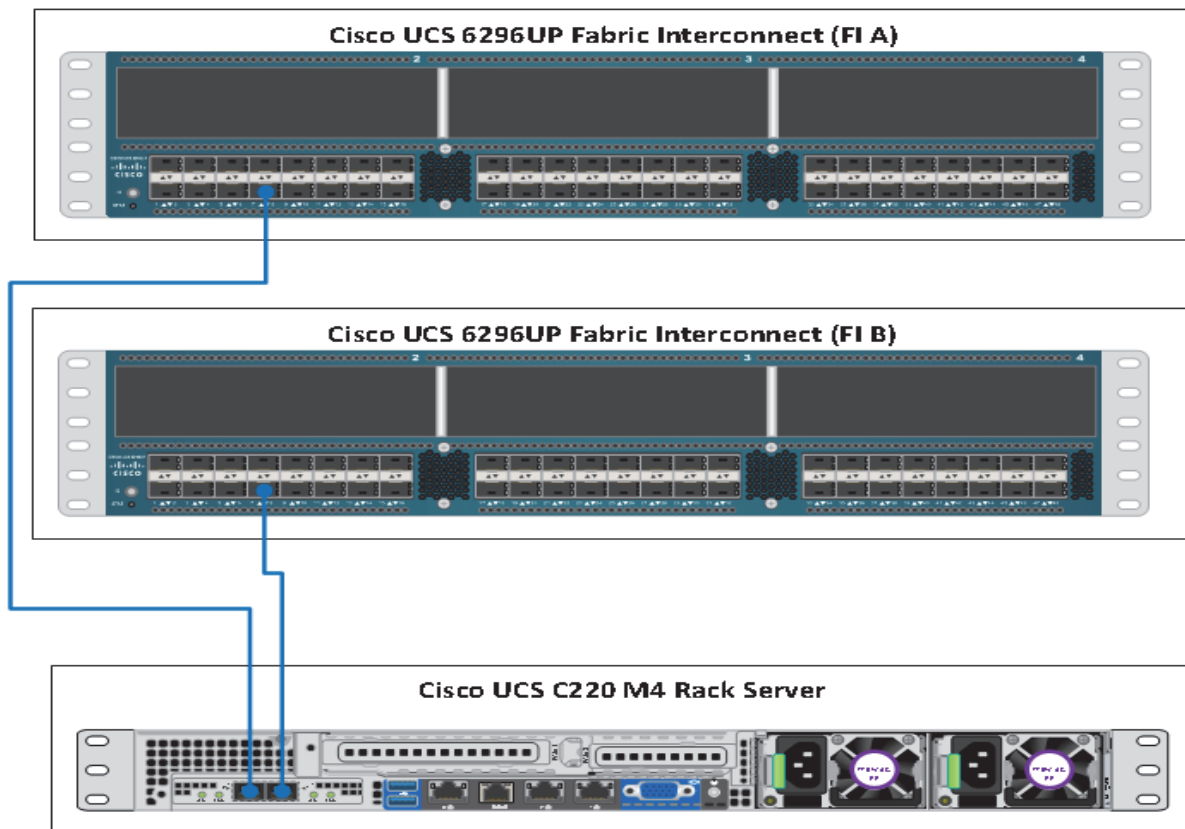


UCSD Management Server Cabling

For this deployment a C220 M4 rack server equipped with Intel Xeon E5-2620 v3 processors, 128 GB of memory, Cisco UCS Virtual Interface Card 1227, Cisco 12-Gbps SAS Modular Raid Controller with 512-MB FBWC, 4 X 600 GB 10K SFF SAS drives is used (any other Cisco UCS server can also be used for this purpose).

The C220 M4 server shall be connected to the UCS Fabric Interconnects as shown in Figure 188. The ports on the on the Fabric Interconnects must be configured as server ports.

Figure 186 Fabric Topology for C220 M4



Software Versions

The UCSD management server is a C220 M4 server that is managed by the UCS Manager. Refer to the software information section in the main part of this Cisco UCS Integrated Infrastructure for Big Data with Hortonworks. See Software Distributions and Versions. In addition, the following software distributions are necessary.

UCS Director Express for Big Data (1.1)

For more information visit
<http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-director-express-big-data-1-1/model.html>

VMware vSphere 5.5

UCS Director express requires the VMware vSphere 5.5 hypervisor. For more information see
<http://www.vmware.com>

Fabric Configuration

The UCSD management server is a C220 M4 server that is managed by the UCS Manager. Refer to the Fabric Configuration section in the main part of this document for more details.

Configuring VLANs

UCSD Express management server requires two network interfaces. It's service profile need to be

- Management Network – default (VLAN 1)
- PXE Network

Table 20 UCSD Express Management Server vNIC configurations

VLAN	Fabric	NIC Port	Function	Failover
default(VLAN1)	A	eth0	Management, User connectivity	Fabric Failover to B
vlan85_PXE	B	eth1	PXE	Fabric Failover to A

PXE VLAN dedicated for PXE booting purpose. Follow these steps in Configuring VLANs to create a dedicated VLAN for PXE. The management network shall continue to be on the default VLAN.

Other UCS configurations

Perform all other UCS configurations such as QOS policy, necessary policies and service profile template by following the documentation above. See the section Creating Pools for Service Profile Templates onwards in this Cisco UCS Integrated Infrastructure for Big Data with Hortonworks cisco validated design.



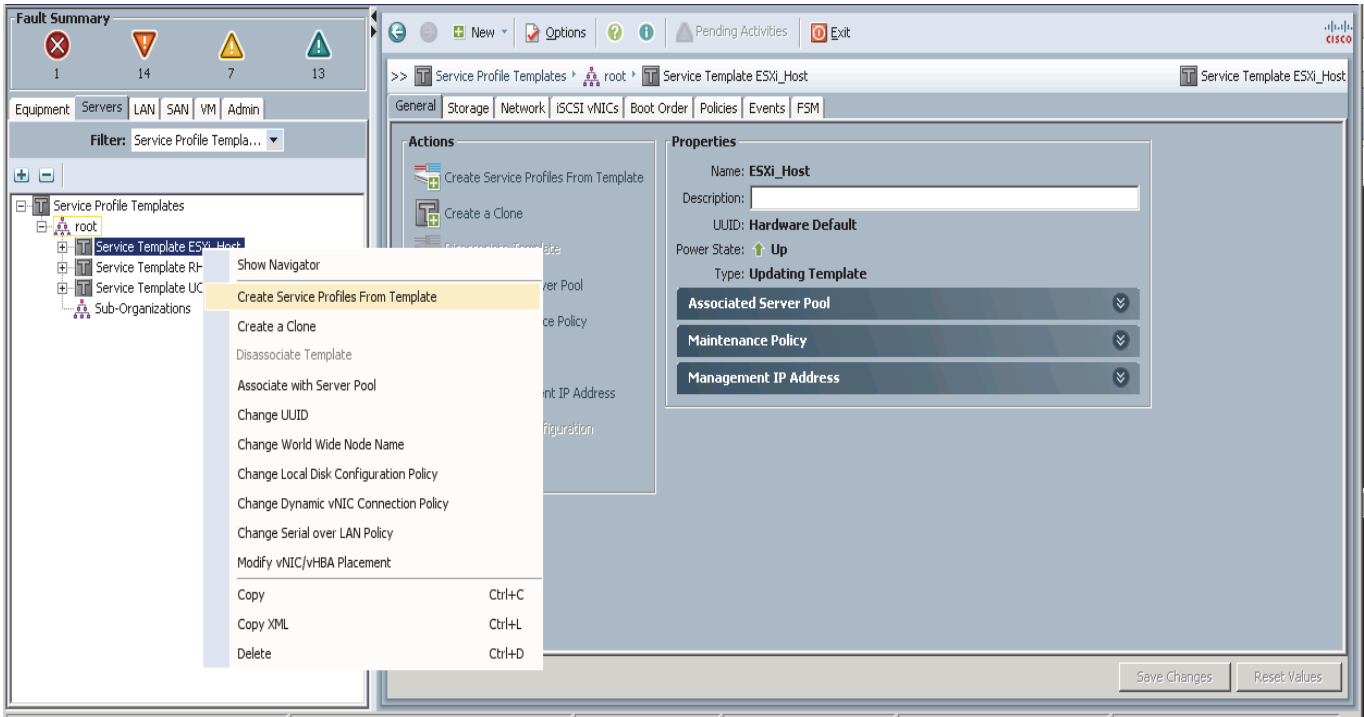
Note Create the service profile template named as ESXi_Host with two vNICs as shown in the above table. For vNIC eth0, select default VLAN as the native VLAN, and for vNIC eth1, select PXE VLAN (vlan85_PXE) as the native VLAN.

Creating Service Profile from the Template

Select the Servers tab in the left pane of the UCS Manager GUI.

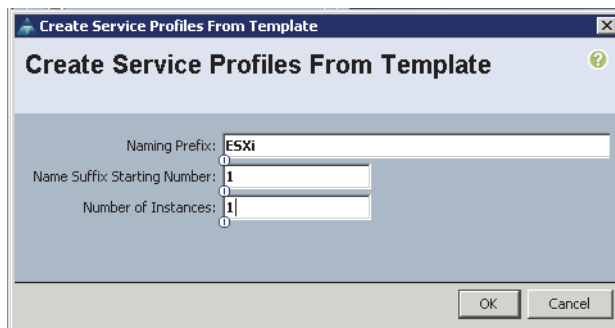
1. Go to Service Profile **Templates > root**.
2. Right-click **Service Profile Templates ESXi_Host**.
3. Select **Create Service Profiles From Template**.

Figure 187 Creating Service Profiles from Template



4. The Create Service Profile from Template window appears.

Figure 188 Selecting Name and Total number of Service Profiles



Association of the Service Profiles will take place automatically.

Installing VMware vSphere ESXi 5.5

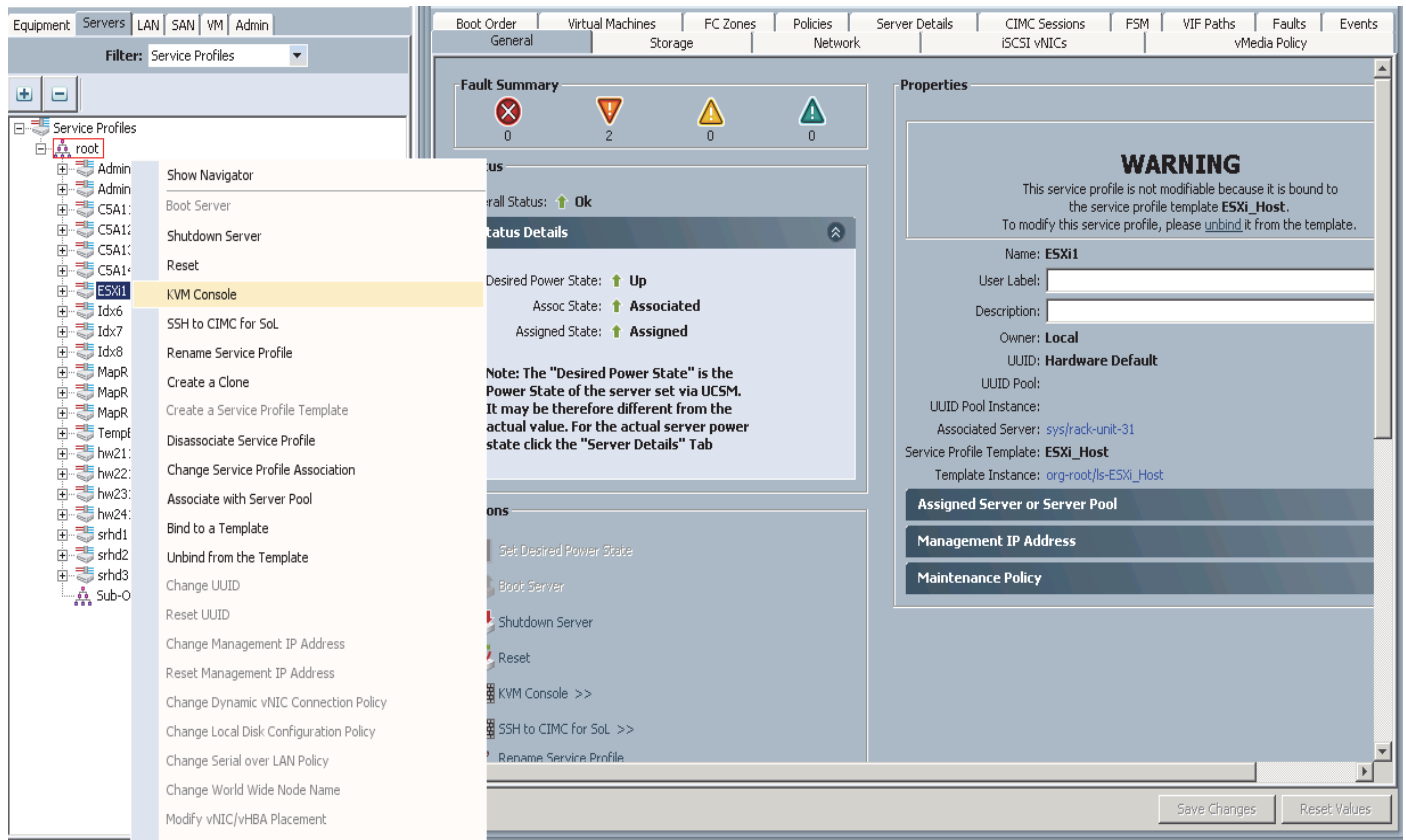
The following section provides detailed procedures for installing VMware vSphere ESXi 5.5.

There are multiple methods to install VMware vSphere ESXi 5.5. The installation procedure described in this deployment 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.

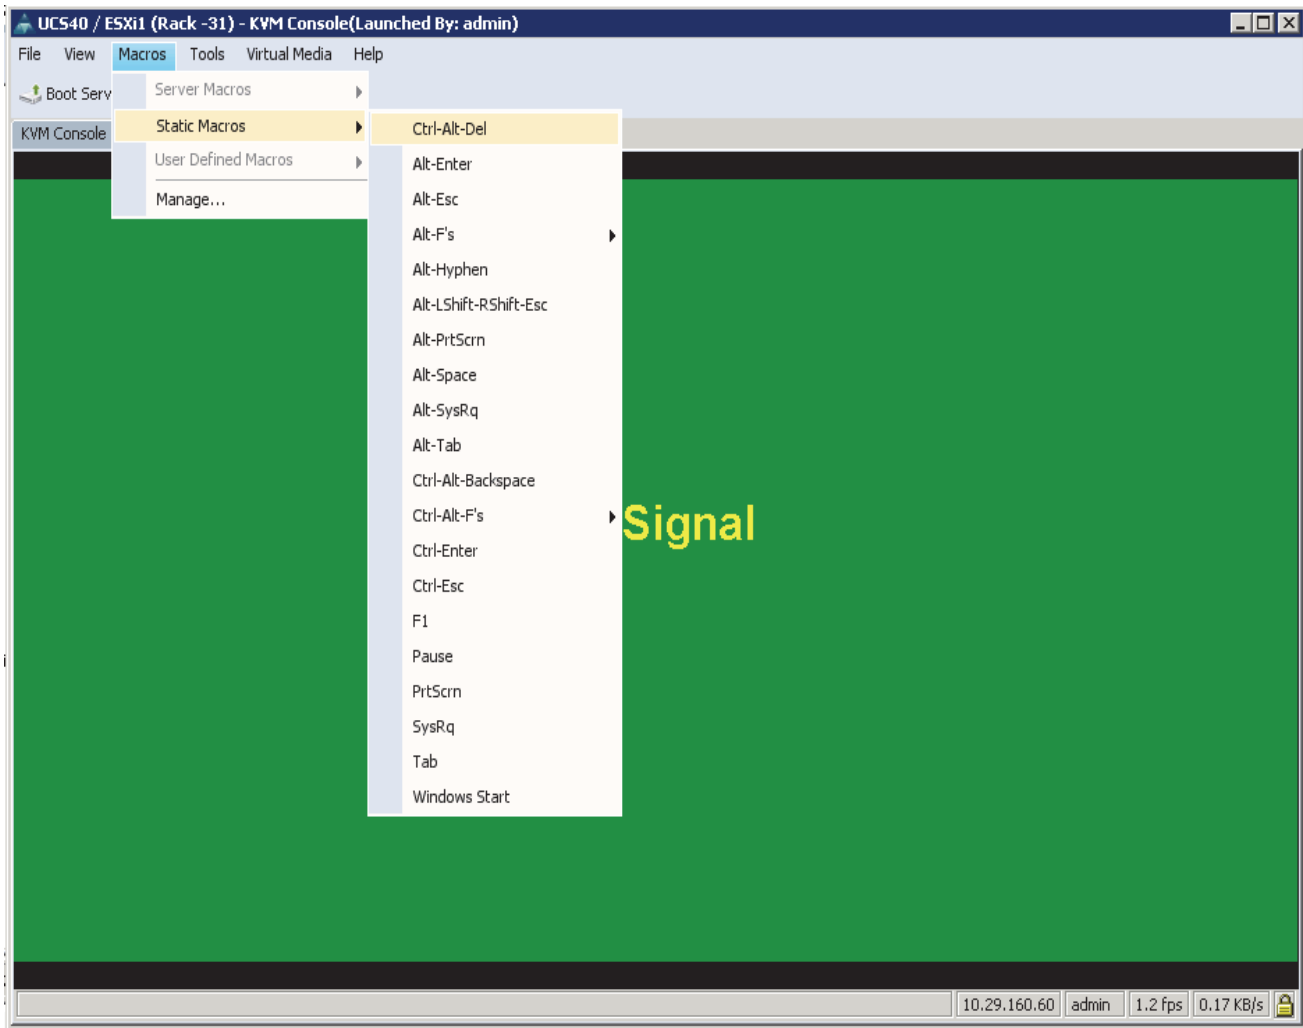
2. Select the **Servers** tab.
3. In the navigation pane expand Service Profiles.
4. Right click on the newly created service profile ESXi1 and select KVM Console.

Figure 189 *Selecting KVM Console*



5. In the KVM window, force a reboot by executing the **Ctrl-Alt-Del** macro.

Figure 190 Sending Ctrl-Alt-Del to Reset the Server



6. As the server goes through a reboot, monitor the progress via the KVM window. When the LSI MegaRAID SAS-MFI BIOS screen appears, press **Ctrl-R** to Enter the Cisco 12G SAS Modular Raid Controller BIOS Configuration Utility.

Figure 191 KVM Window displaying the LSI MegaRAID SAS-MFI BIOS screen

```

UCS40 / ESXi1 (Rack -31) - KVM Console(Launched By: admin)
File View Macros Tools Virtual Media Help
Boot Server Shutdown Server Reset
KVM Console Properties

LSI MegaRAID SAS-MFI BIOS
Version 6.19.05.0 (Build May 07, 2014)
Copyright(c) 2014 LSI Corporation

HA -0 (Bus 9 Dev 0) Cisco 12G SAS Modular Raid Controller
Battery Status: Fully charged
PCI Slot Number: 0

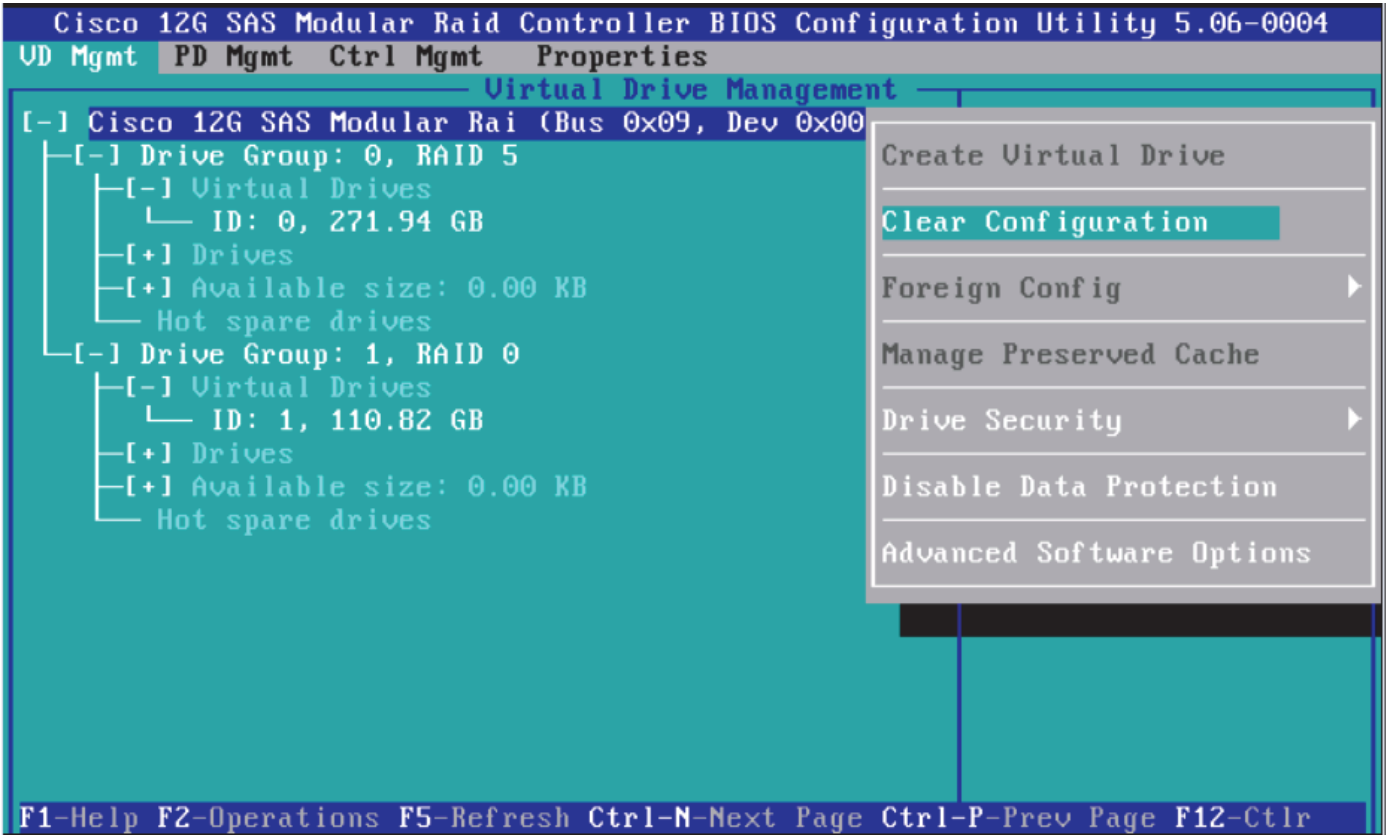
ID LUN VENDOR      PRODUCT                    REVISION      CAPACITY
-- -- -
      LSI          Cisco 12G SAS Modular Raid 4.250.00-3632 2048MB
20 0 SEAGATE      ST9146803SS                0004          140014MB
21 0 SEAGATE      ST9146803SS                0004          140014MB
22 0 TOSHIBA     MBF2300RC                  5704          286102MB
23 0 ATA        Micron_M500_MTFD           MU03          114473MB

0 JBOD(s) found on the host adapter
0 JBOD(s) handled by BIOS

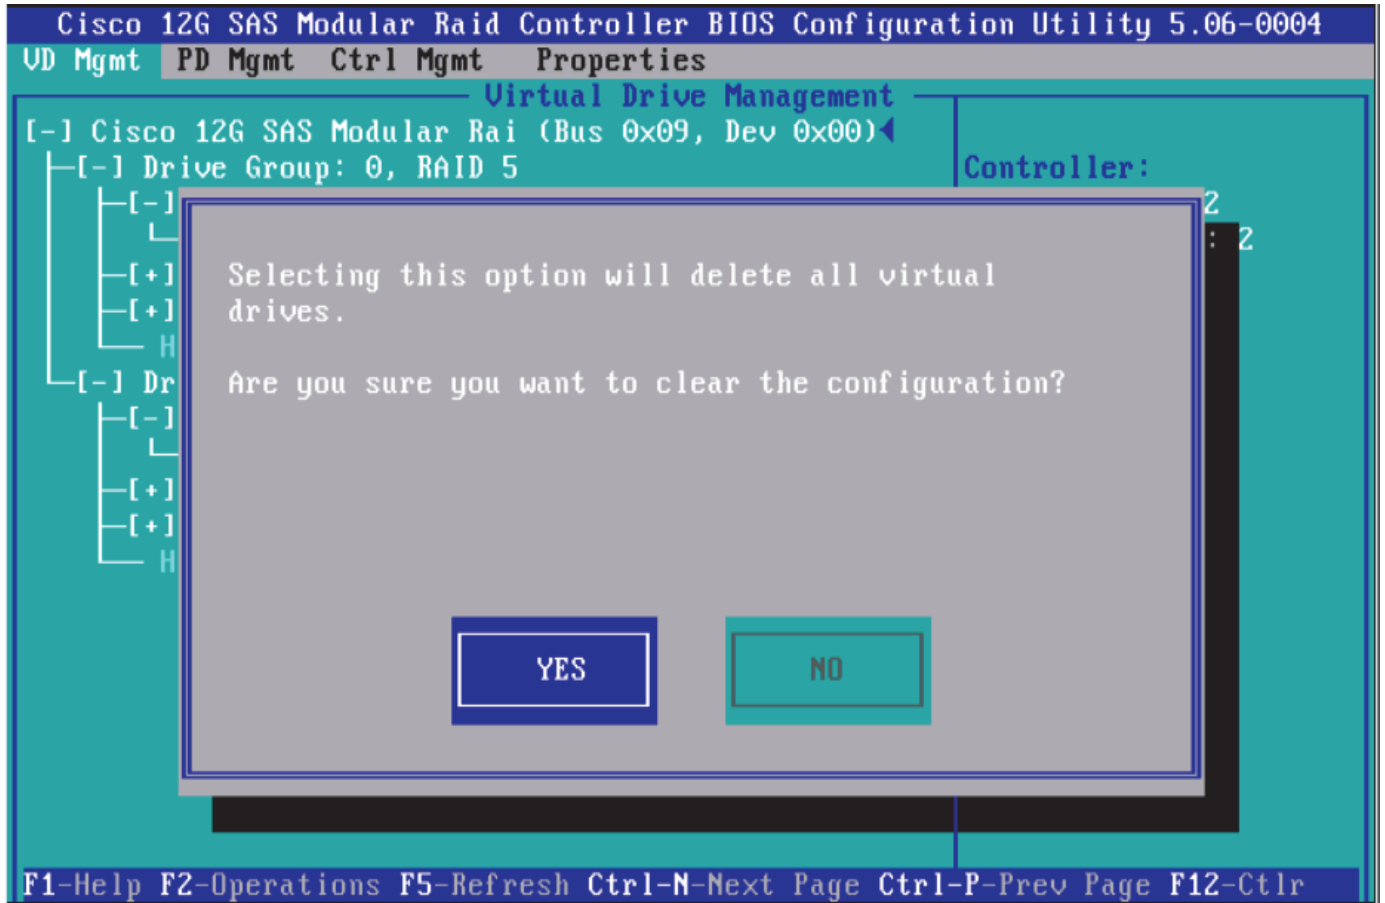
0 Virtual Drive(s) found on the host adapter.
0 Virtual Drive(s) handled by BIOS
Press <Ctrl><R> to Run MegaRAID Configuration Utility
10.29.160.60 admin 0.4 fps 0.001 KB/s

```

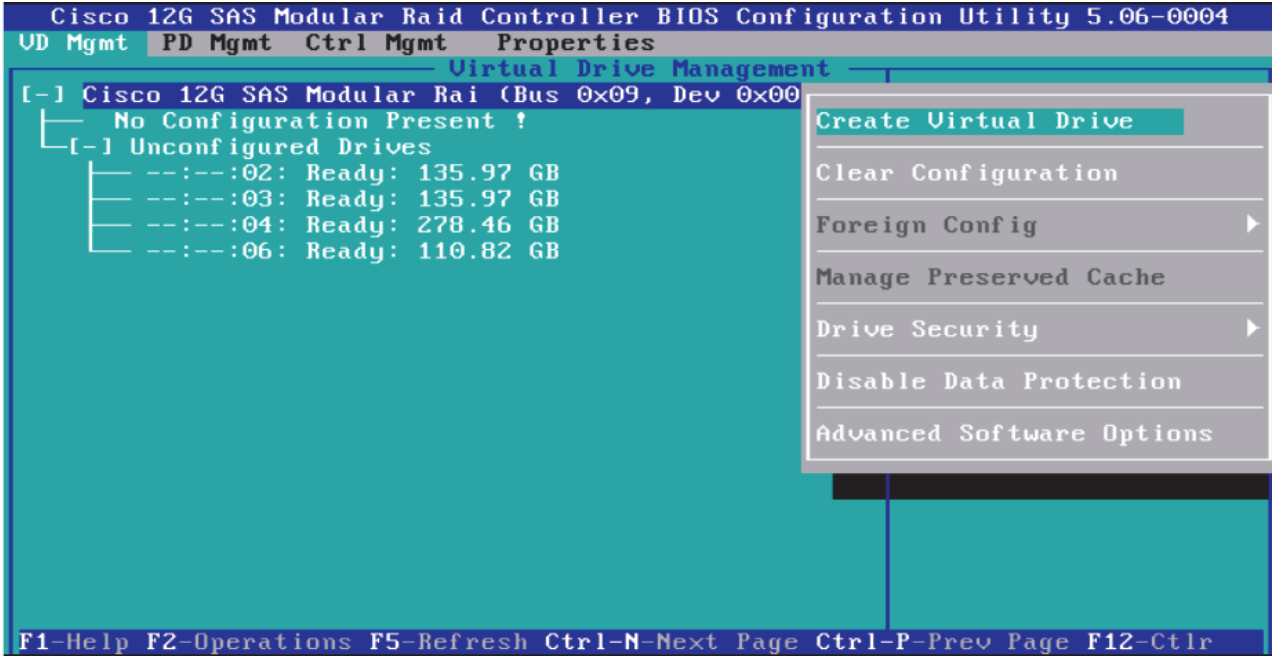
7. In the MegaRAID configuration utility, under VD Mgmt section, use the arrow keys to select the Cisco 12G SAS Modular RAID (Bus 0xNN, Dev 0xNN) line item.
8. Press the function key **F2**.
9. Select the option Clear Configuration, and press **ENTER**.



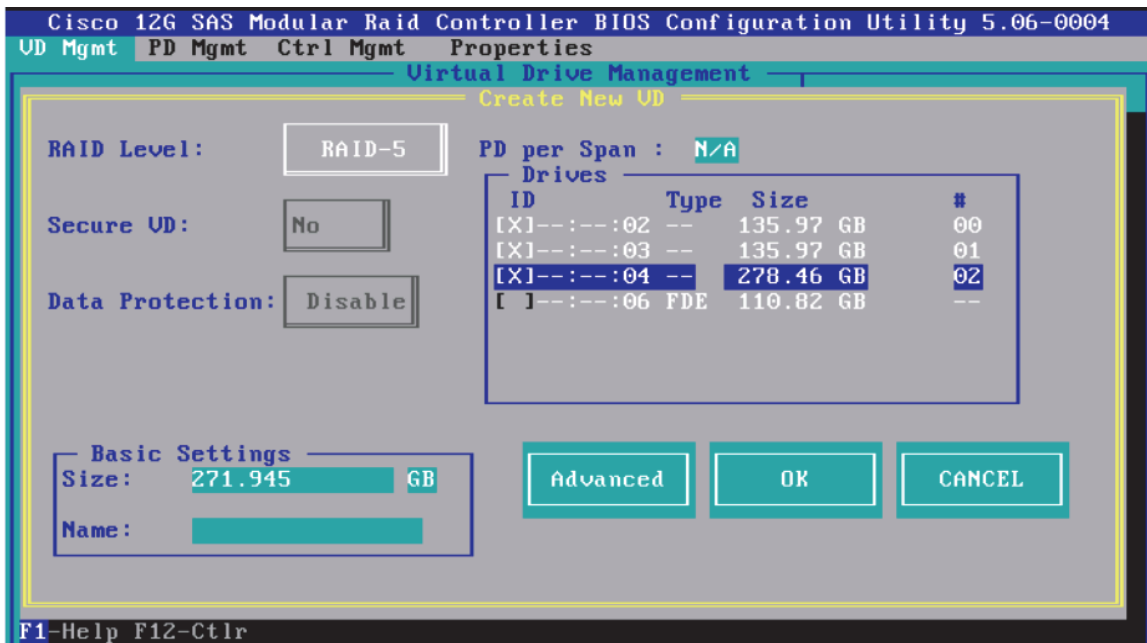
- To the question Are you sure you want to clear the configuration? click **YES** and press **ENTER** key.



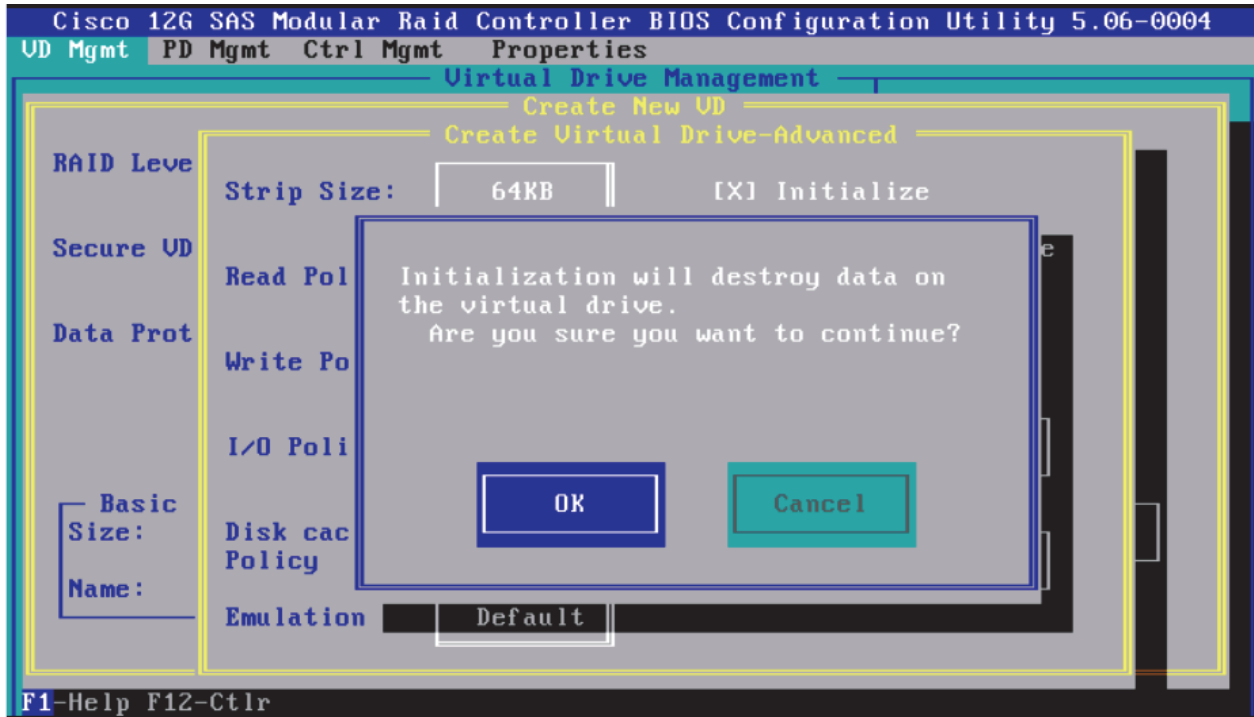
11. In the VD Mgmt section, use the arrow keys to select the Cisco 12G SAS Modular RAID (Bus 0xNN, Dev 0xNN) line item.
12. Press the function key **F2**, select Create Virtual Drive and press **ENTER**.



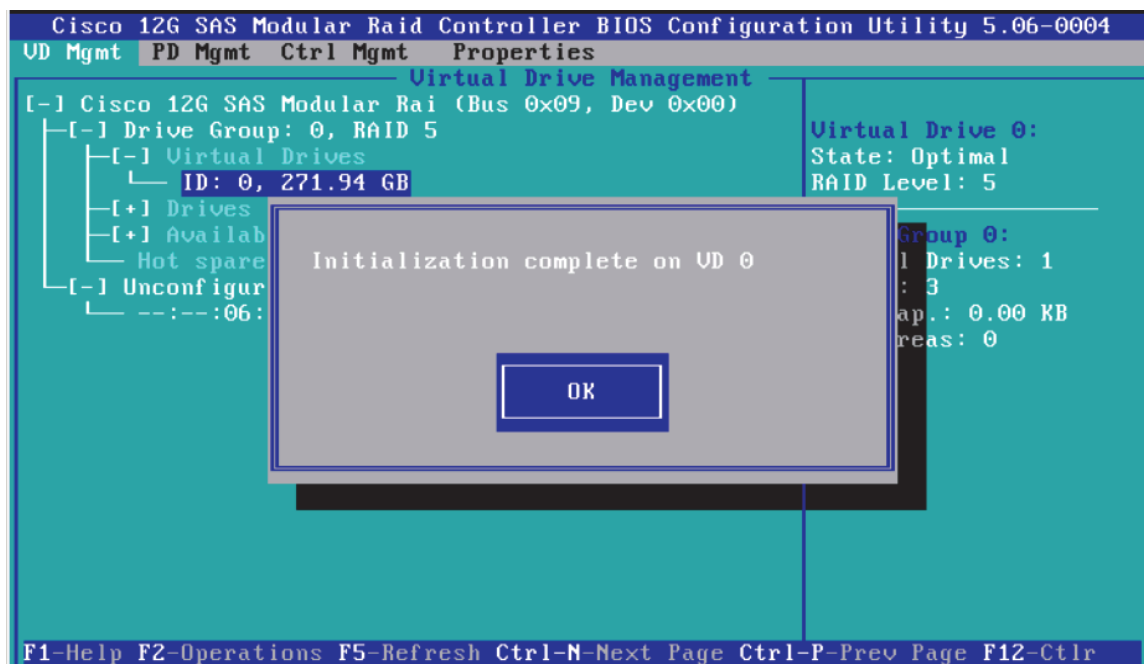
13. In the RAID Level: press **ENTER** and choose **RAID-5**.
14. In the Drives section, press **SPACE** on the desired number of drives to select them to be part of the RAID group. Use the Up and Down arrow keys to navigate.



15. Select the **Advanced** button, and Check the Initialize checkbox.
16. Press **OK** to continue with initialization.

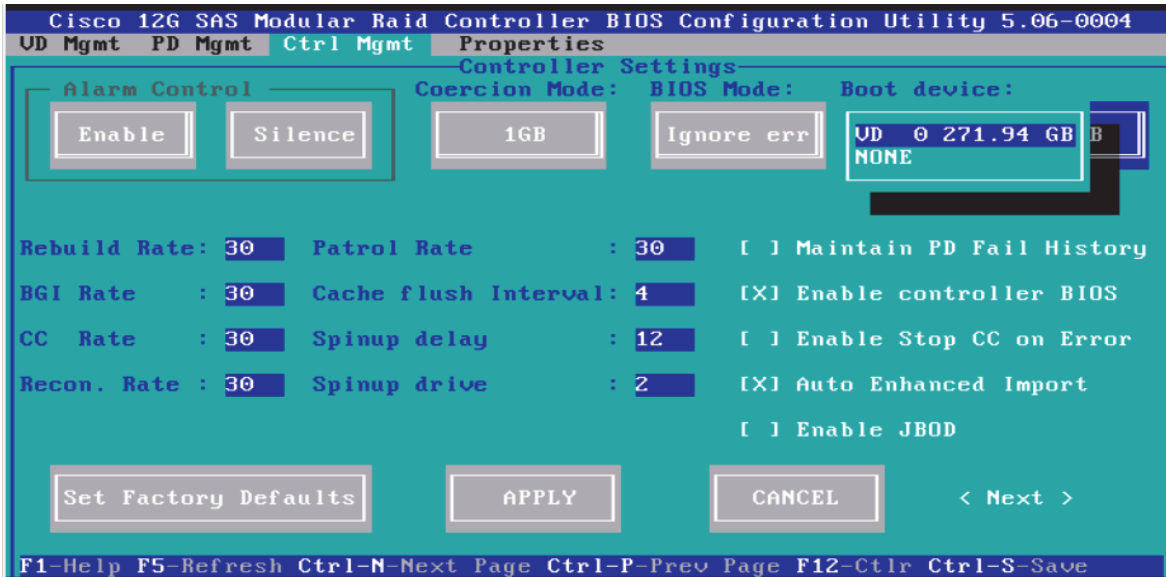


17. After the initialization is complete, the following message appears. Press **OK** to continue.

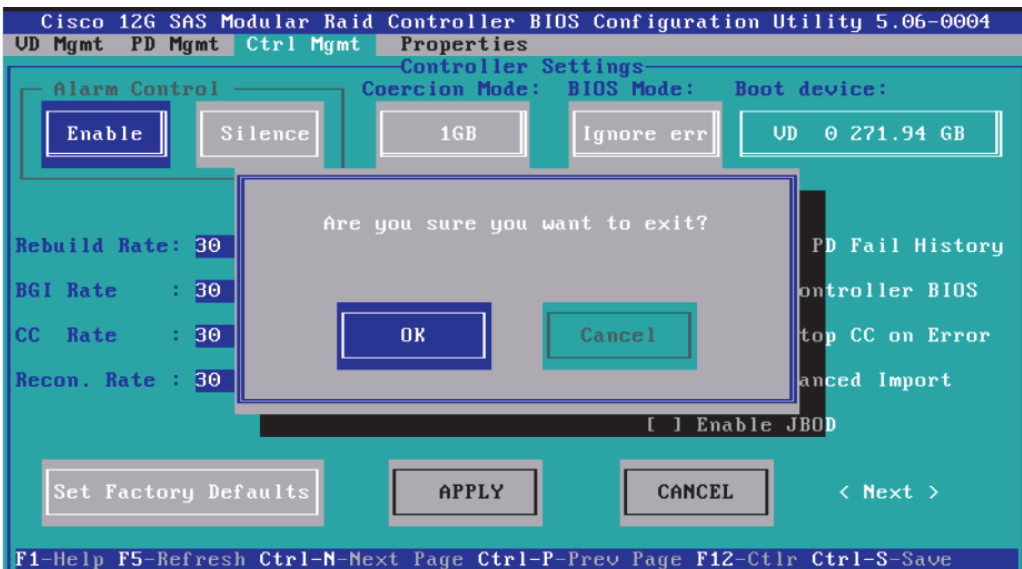


18. Press **Ctrl-N** twice to navigate to the Ctrl Mgmt screen.


19. Select Boot device field and press **ENTER**.



20. Select the **VD 0**, and press **ENTER** again.
21. Press **Ctrl+S** to save the configuration.
22. Press **ESC** to exit the MegaRAID configuration utility.

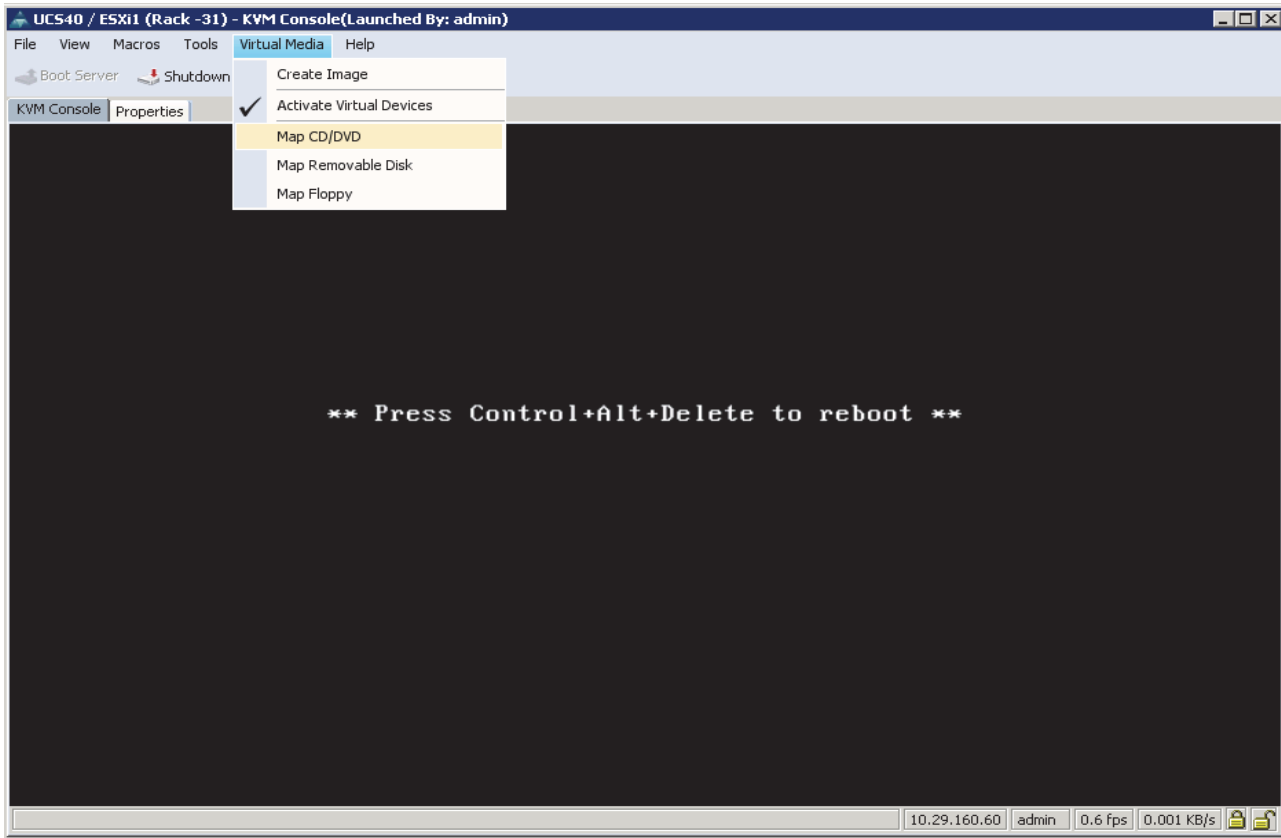


23. In the KVM window, select the Virtual Media menu.
24. Click the Activate Virtual Devices found in the right hand corner of the Virtual Media selection menu.



25. In the KVM window, select the Virtual Media menu and Select **Map CD/DVD**.

Figure 192 Mapping the CD/DVD Virtual Media



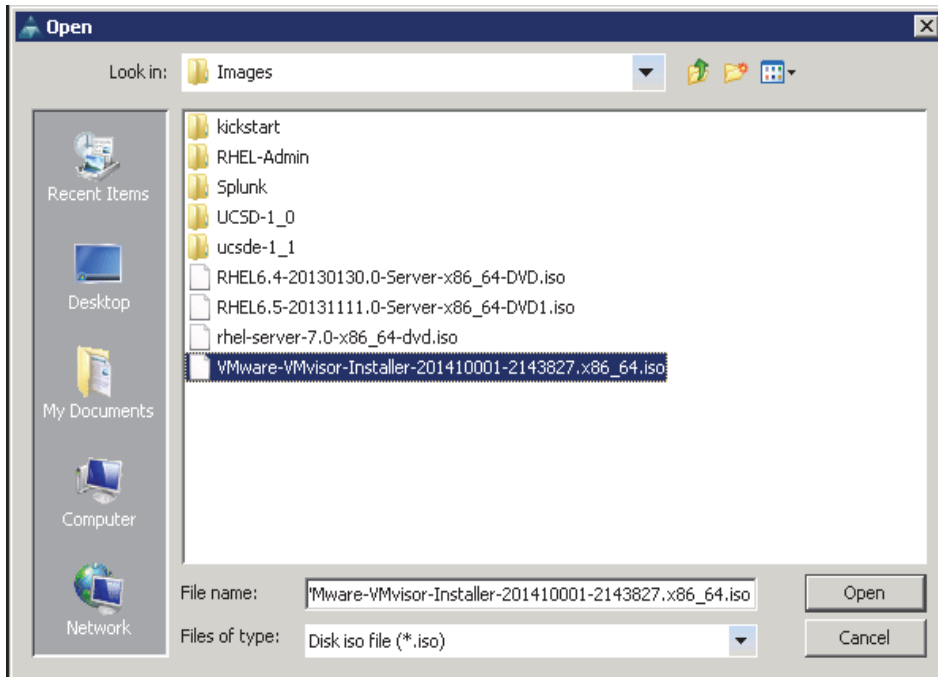
26. Browse to the VMware vSphere ESXi 5.5 installer ISO image file.



Note The VMware vSphere ESXi 5.5 installable ISO is assumed to be on the client machine.


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

Figure 193 Browse to VMWare ESXi Hypervisor ISO Image



28. In the KVM window, select the **KVM** tab to monitor during boot.
29. In the KVM window, select the **Macros > Static Macros > Ctrl-Alt-Del** button in the upper left corner.
30. Click **OK** to reboot the system.
31. On reboot, the machine detects the presence of the VMWare ESXi install media.

Figure 194 ESXi Standard Boot Menu



32. Select the ESXi-5.5.0-yyyyymmddnnnn-standard Installer. The installer begins automatically.

Figure 195 Loading the ESXi Installer

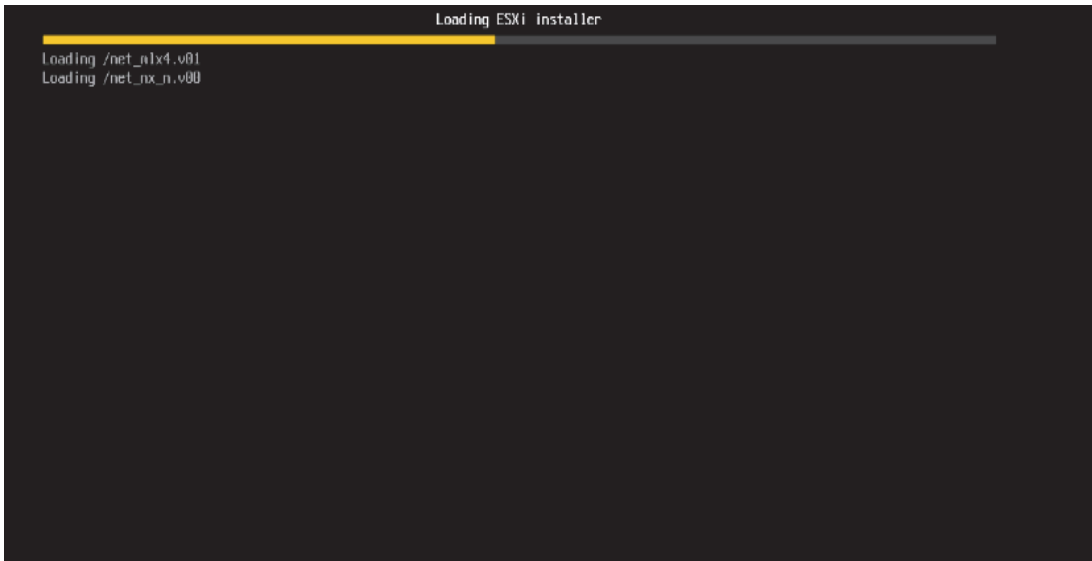


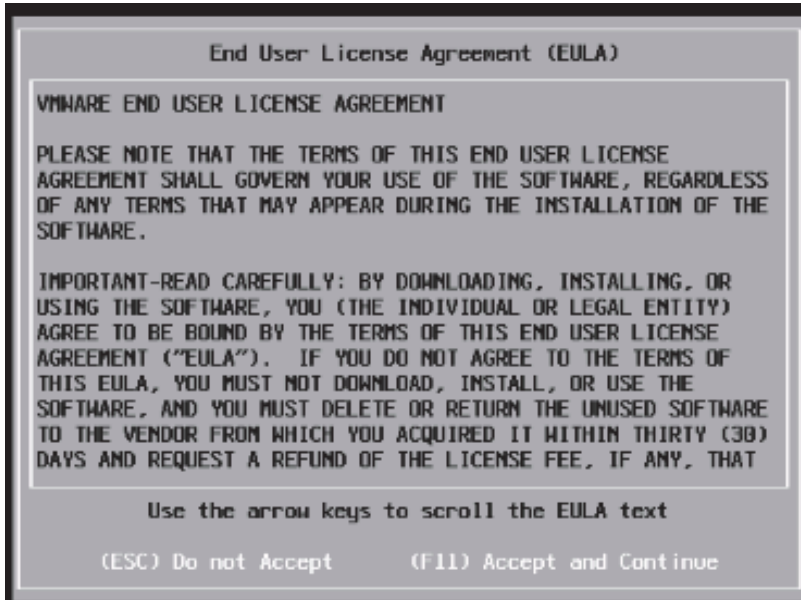
Figure 196 VMWare ESXi Installation screen



33. Press ENTER to continue.

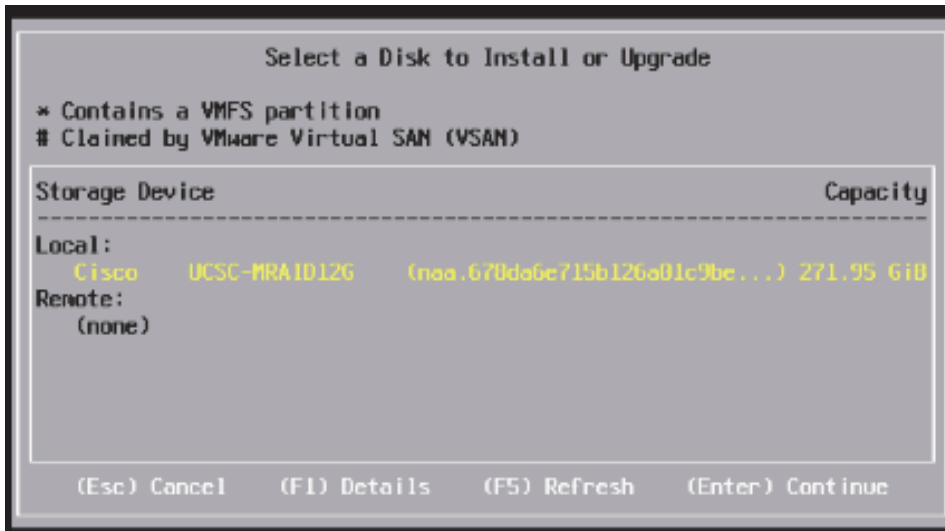
34. Press **F11** to accept End user License Agreement (EULA) and continue.

Figure 197 Accept End User License Agreement (EULA)



35. Select the storage device. Press **ENTER** to proceed with the installation.

Figure 198 Selecting the Storage Device for installing the ESXi operating system.



36. Select the Keyboard US Default. Press **ENTER** to continue.

Figure 199 *Choose the Keyboard layout*



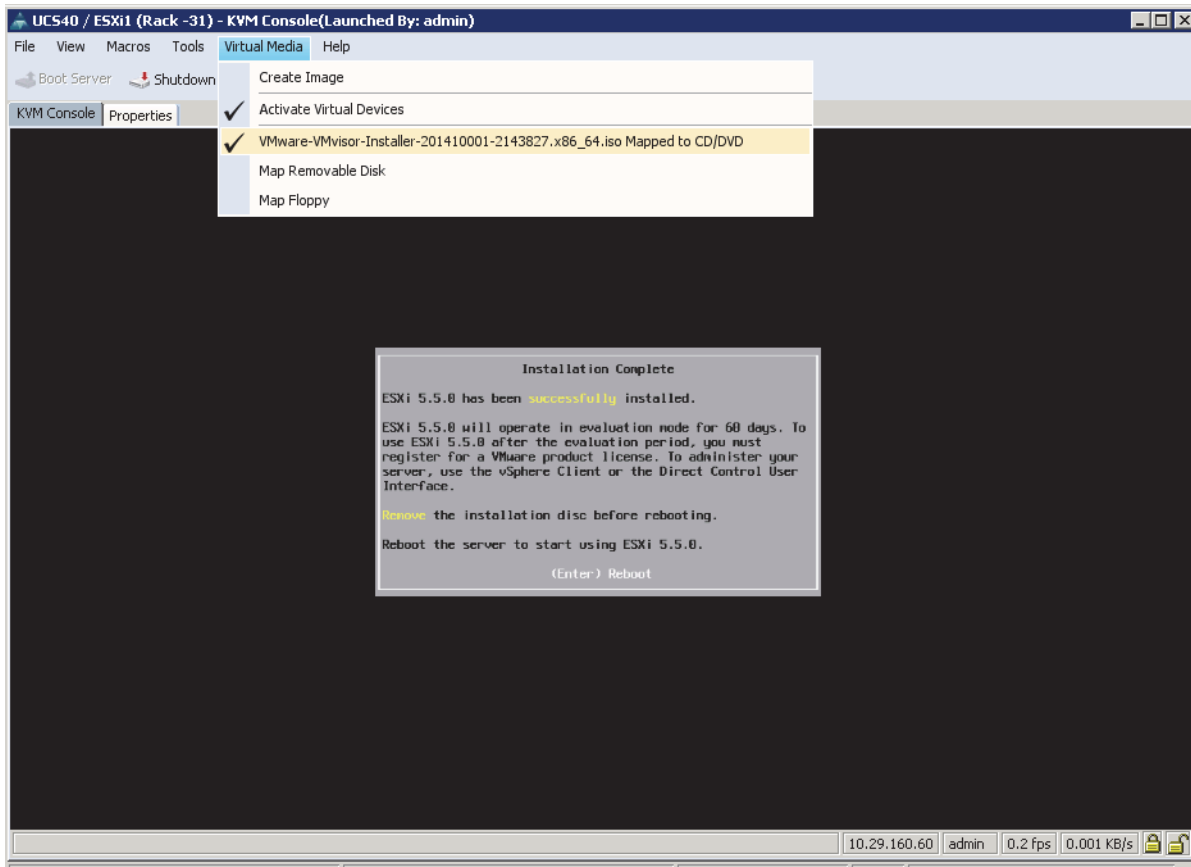
37. Choose the root password and confirm it. Press **ENTER** to continue.

Figure 200 *Choose the root password*



38. Press **F11** to confirm and begin installation.
39. Once the installation completes, the following message is displayed in the KVM.
40. Remove the VMWare vSphere Hypervisor's ISO from the Virtual Media menu, by selecting it as shown.

Figure 201 ESXi installation complete – Unmount the Virtual Media



41. Click **Yes** to proceed with un-mapping of the ISO.

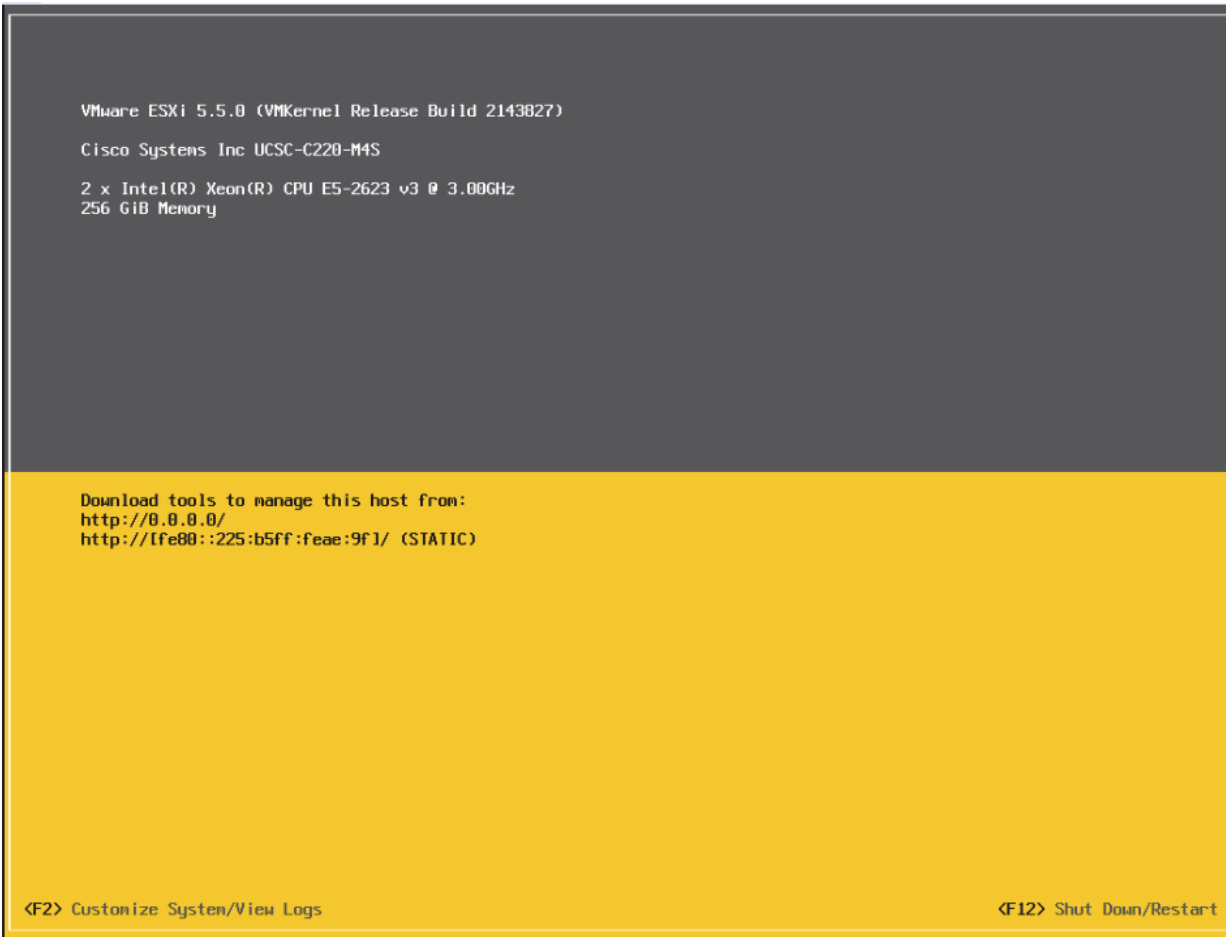
42. Press **ENTER** to reboot the server.

The VMWare vSphere ESXi installation is complete.

Configuring the Management Network

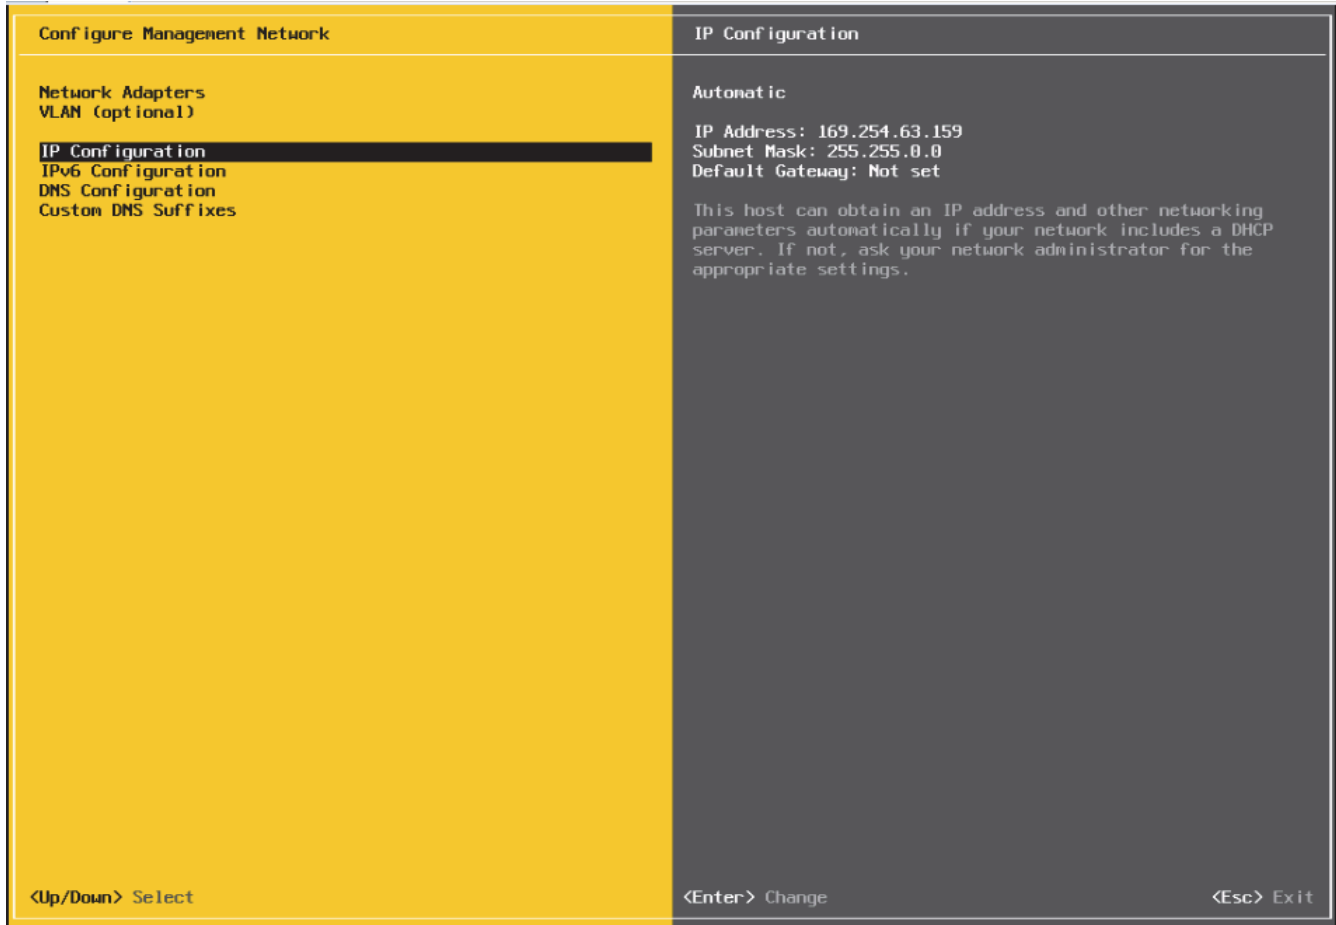
1. Once the server reboots, press **F2** to log on.
2. Enter username as root, and the password chosen above.

Figure 202 VMWare ESXi initial screen as seen via the KVM Console



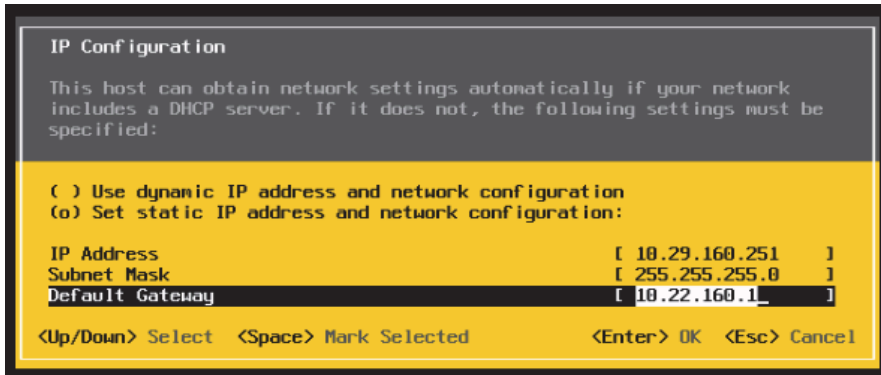
3. Press **F2** to continue
4. Select Configure Management Network, and press **ENTER**.
5. Select **IP Configuration** option.

Figure 203 Enter the IP configuration option of the Management Network



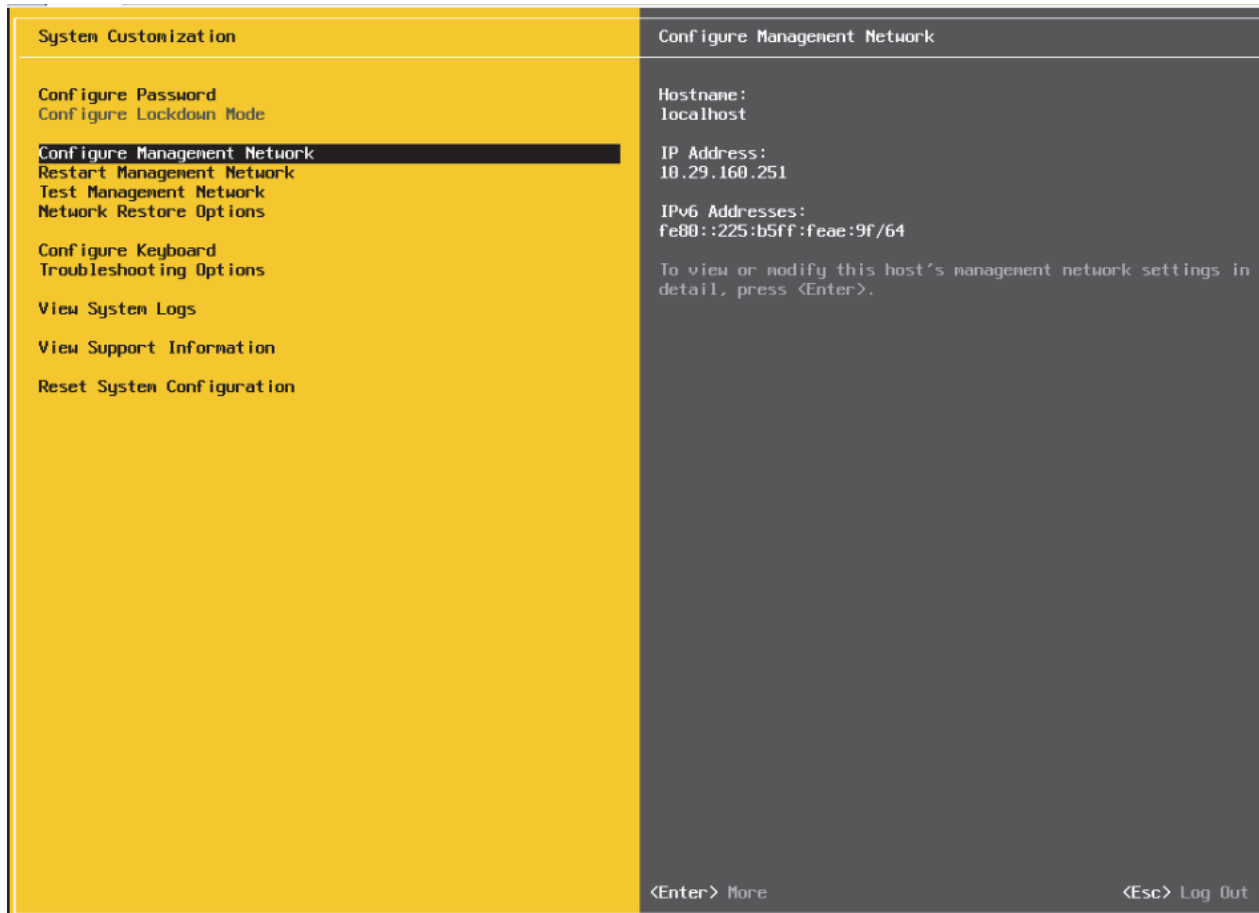
6. Press **ENTER** to continue.
7. Use the Up/Down arrow keys to highlight the Set Static IP address and network configuration option, and press **SPACE** key to select it.
8. Enter the static IP address, Subnet Mask and Default Gateway.

Figure 204 Enter the IP Address configuration details



9. Press **OK** to submit the changes.
10. Press **ESC** key exit the Management Network Screen.
11. In the Configure Management Network: Confirm dialog box, Press **Y** to restart the Management Network.
12. Verify the IP address settings in the System Customization screen.

Figure 205 Verify the IP address details in the System Customization screen



Installing the VMWare ESXi client software

1. Using a web browser, visit the url: <https://10.29.160.251/>
2. Click on Download vSphere Client.

Figure 206 Accessing the ESXi web interface

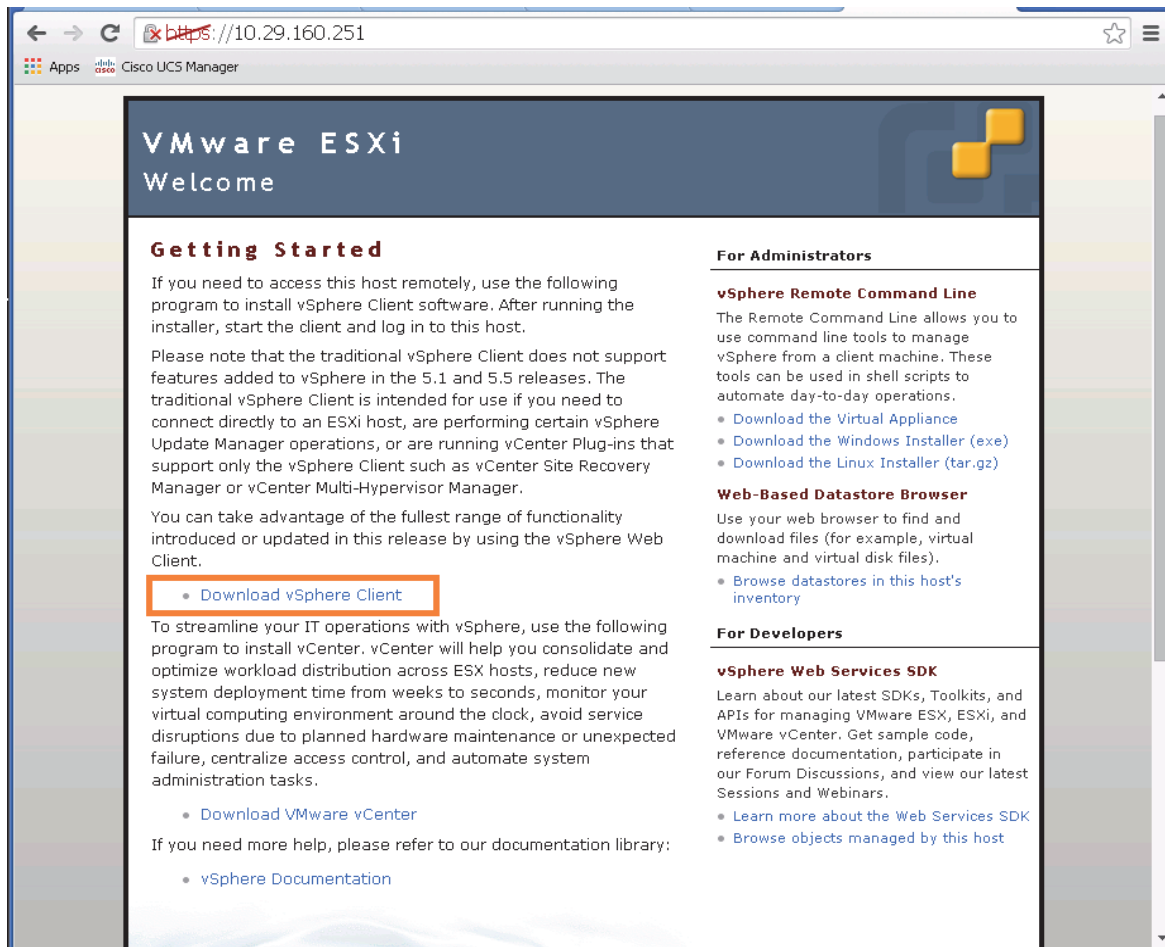


Figure 207 Download the VMWare vSphere ESXi client software




VMware-viclient-all-5.5.0-1993072.exe

<http://vsphereclient.vmware.com/vsphereclient/1/9/9/3/0/7/2/VMware-viclient-all-5.5.0-199...>

Show in folder Remove from list

3. Proceed to install the downloaded VMWare client software.

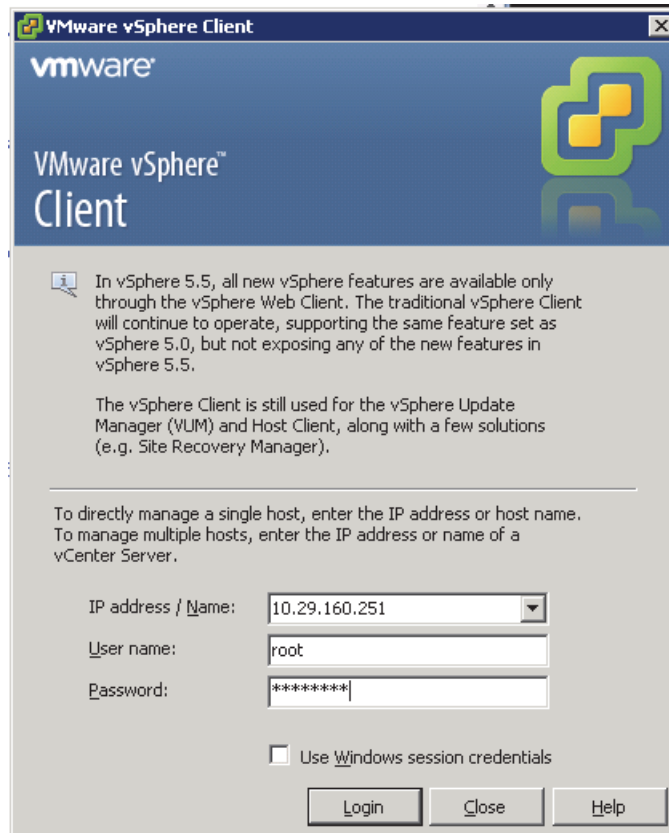
Figure 208 Installing the vSphere Client software



Configuring the vSphere ESXi hypervisor

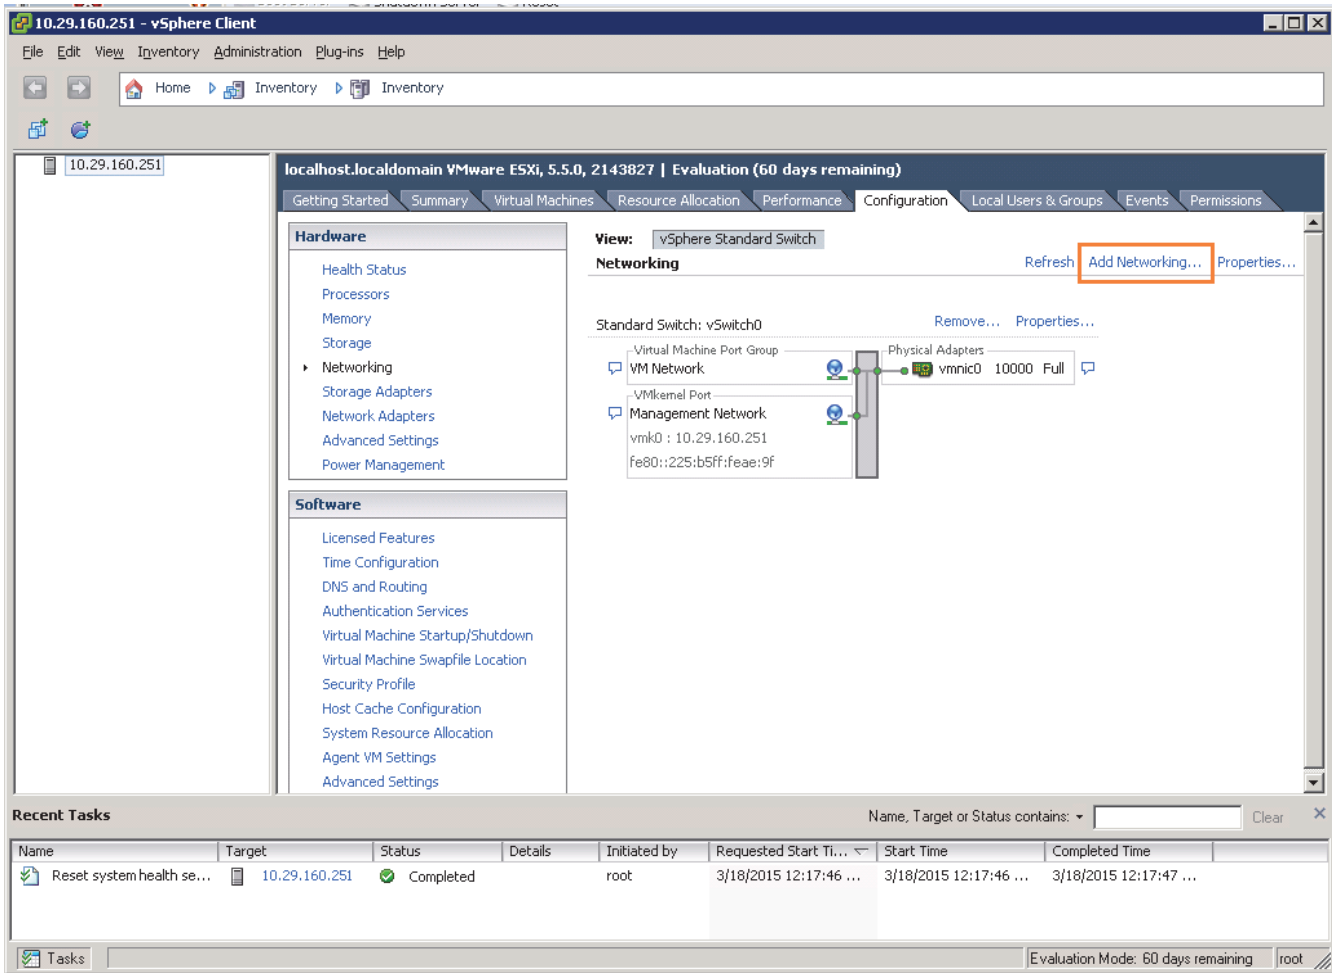
1. After the installation is complete, launch the VMWare vSphere client.
2. Enter the chosen IP address, the username as root, and the chosen password.
3. Click on **Login** to continue.

Figure 209 Logging into the ESXi using vSphere Client



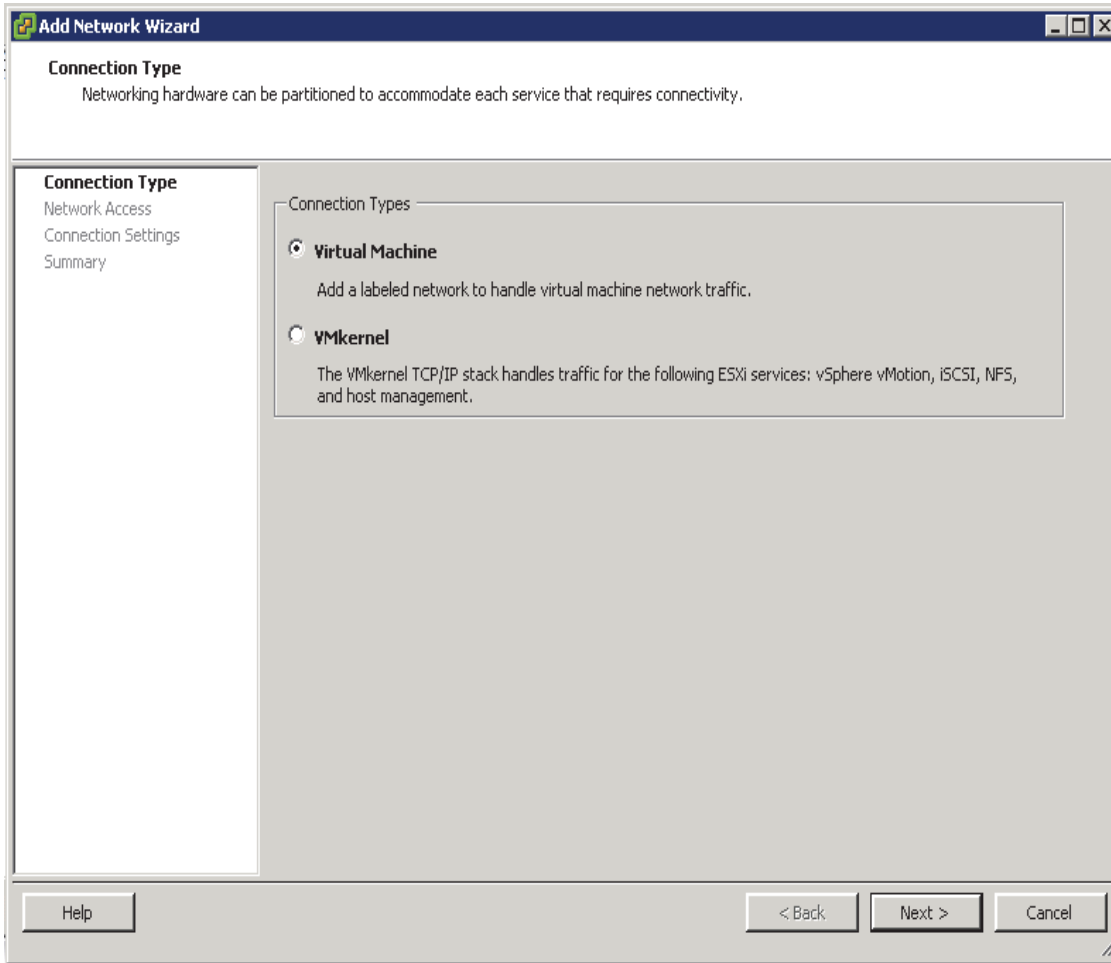
4. In the vSphere Client, click on the Configuration tab on the right, and within the Hardware section, click on Networking.
5. Click on Add Networking link on the upper right hand side.

Figure 210 vSphere Client Networking screen



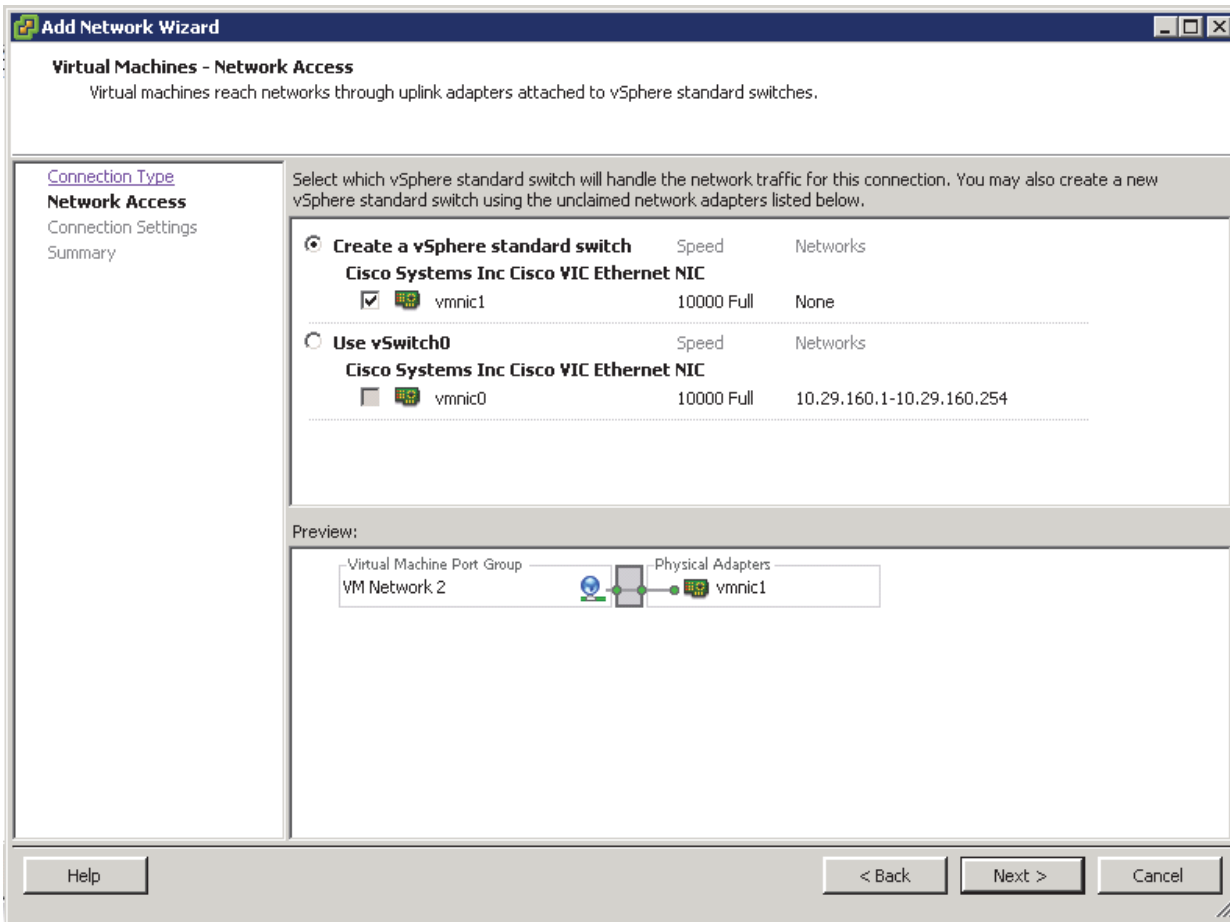
6. In the Add Networking dialog box, click the **Virtual Machine** radio button and click **Next**.

Figure 211 Adding a new Virtual Machine Network



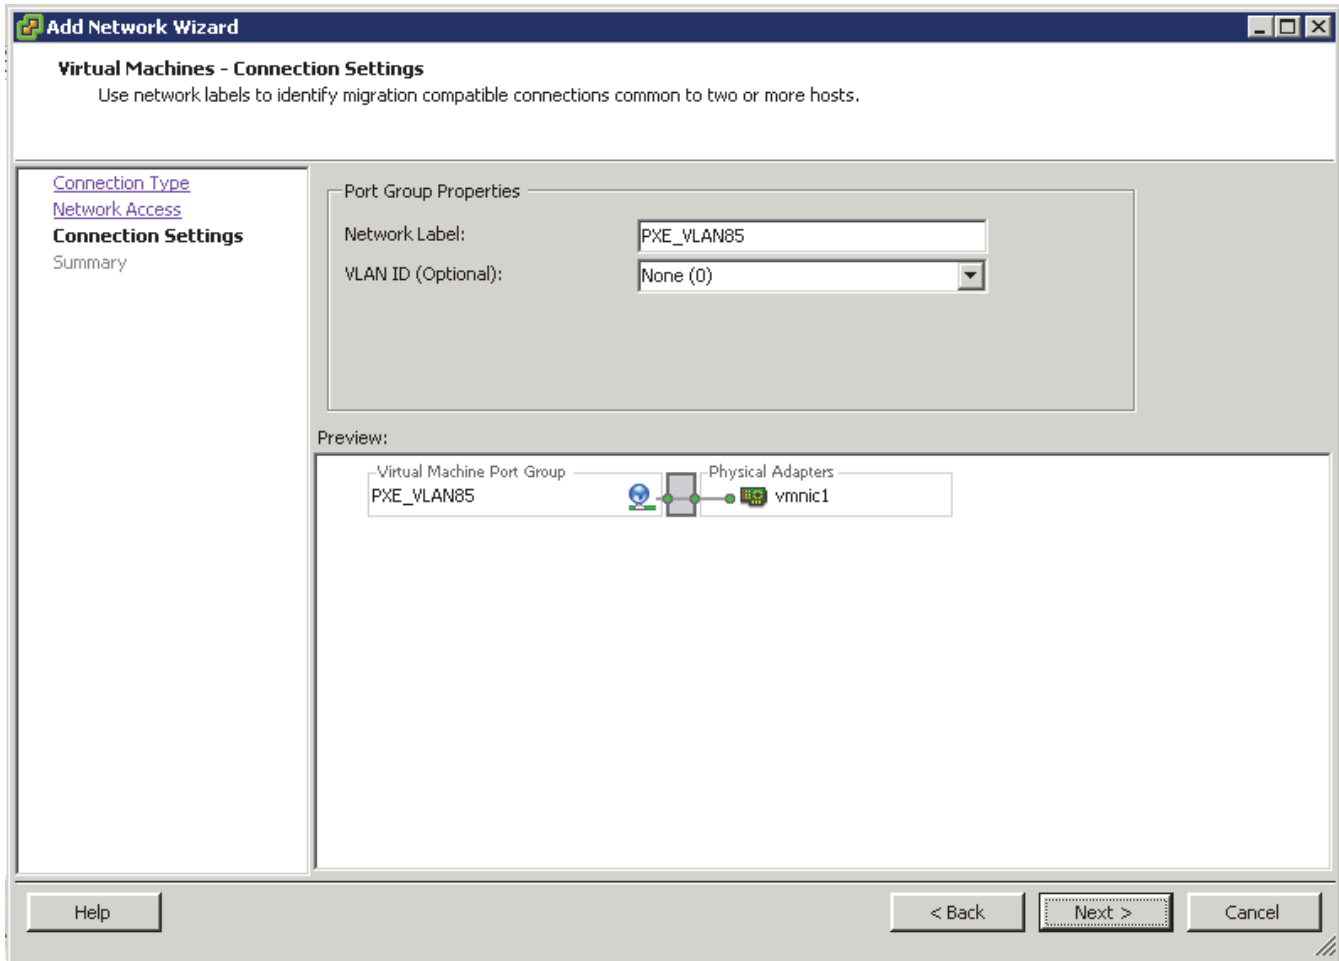
7. Click the **Create a vSphere standard switch** radio button and make sure that the checkbox next to `vmnic1` is checked.
8. Click **Next**.

Figure 212 Creating a new vSphere Standard Switch



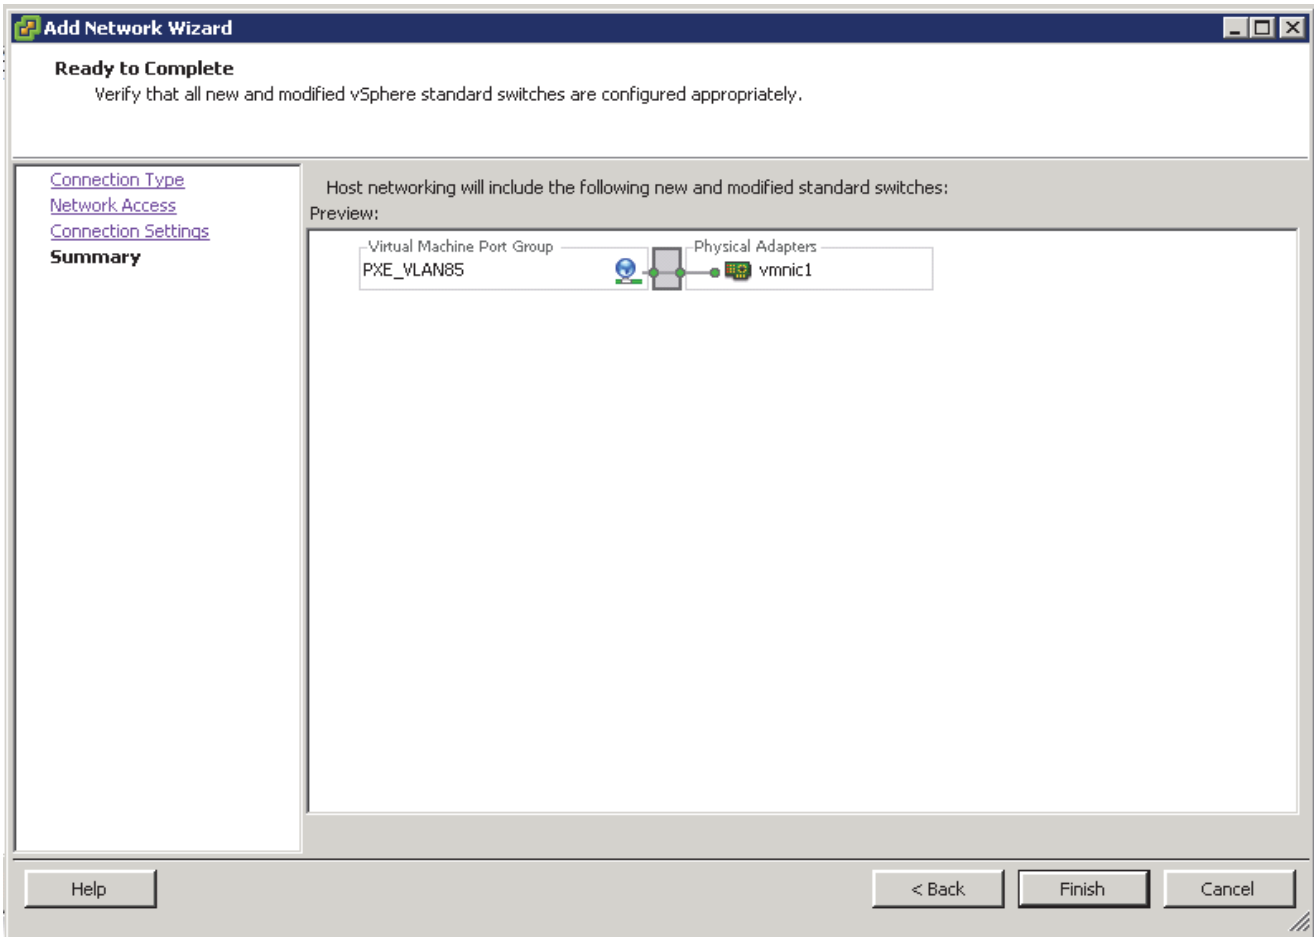
9. In the Port Group Properties, change the Network Label field to PXE_VLAN85.
10. Leave the VLAN ID(Optional) field as None(0).
11. Click **Next**.

Figure 213 Creating the Port Group for the PXE VLAN



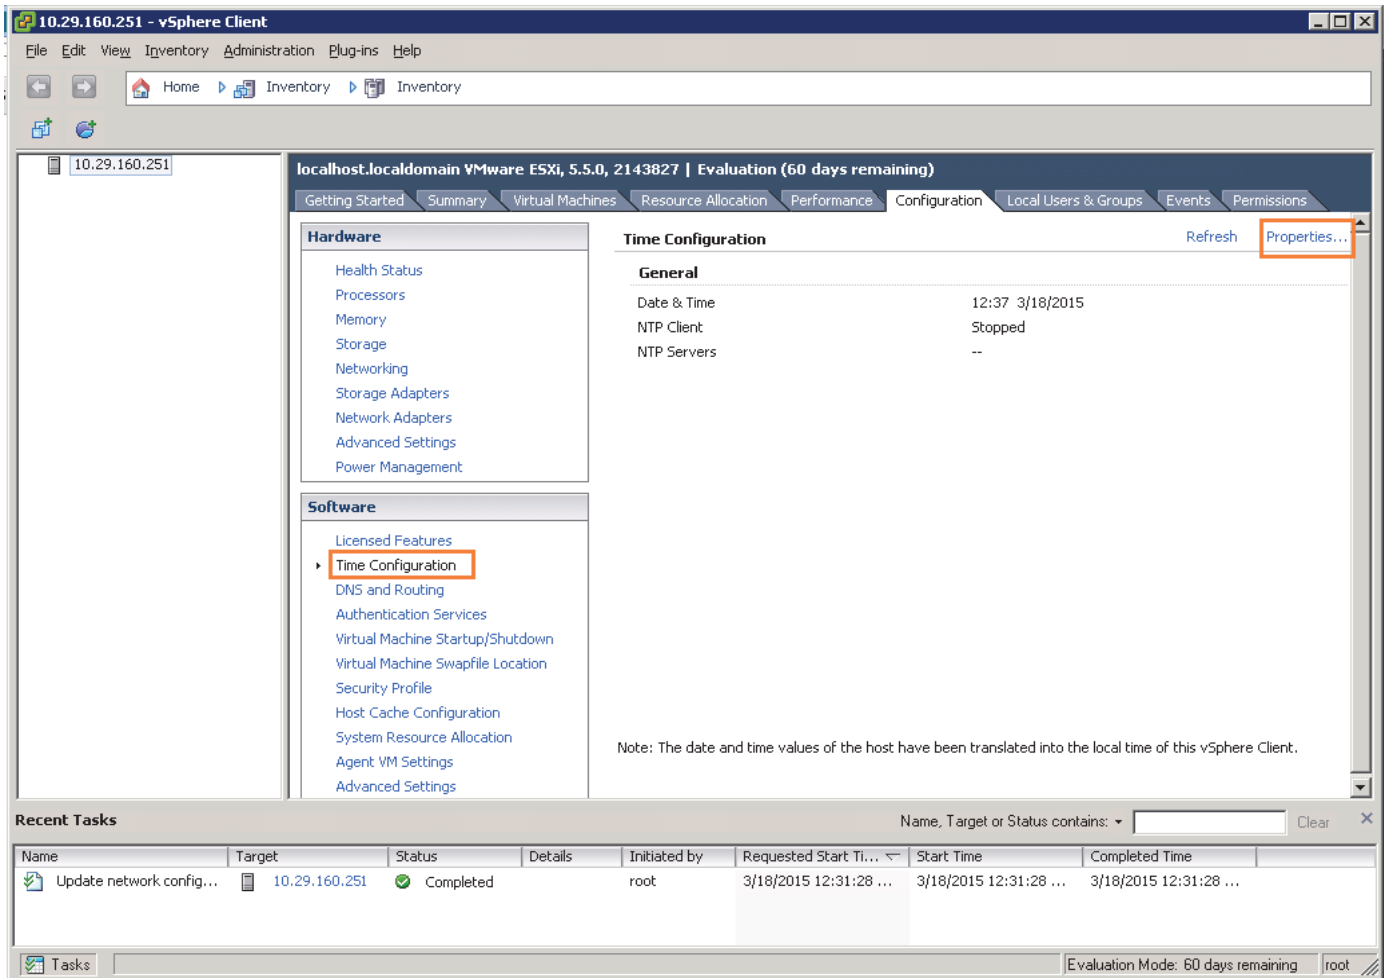
12. Click **Finish** to complete adding the Network.

Figure 214 Verify the Created vSphere Standard Switches



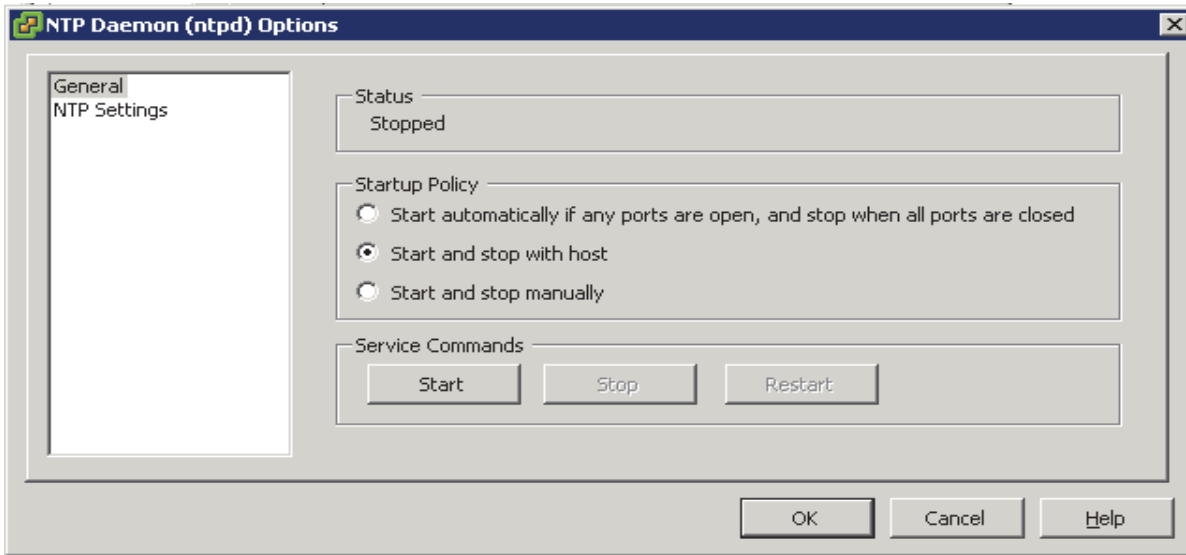
13. Click on the **Time Configuration** under the Software section.
14. Click on **Properties** at the upper right hand corner.

Figure 215 Enabling the NTP Client on the ESXi



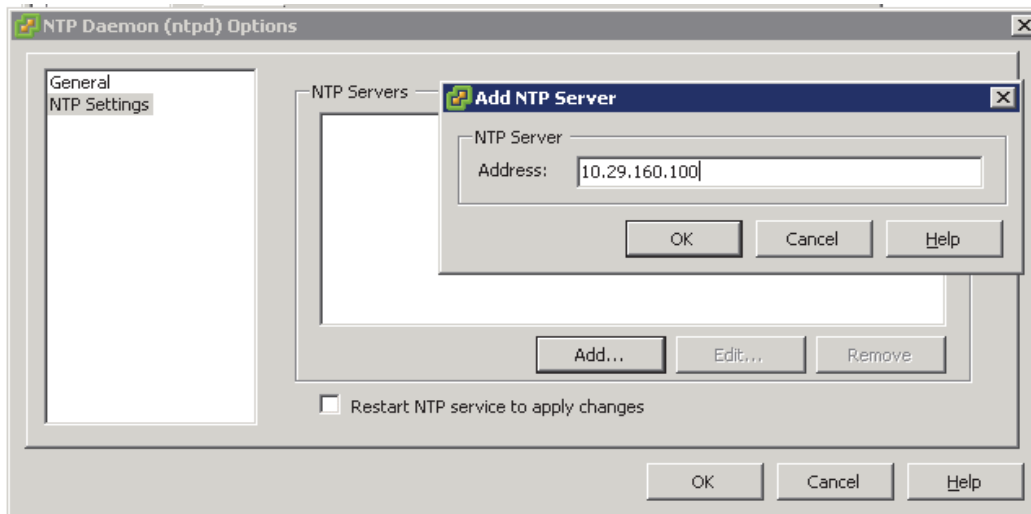
15. In the NTP Daemon (ntpd) Options dialog box, click **Options**.
16. Click on the **General** options.
17. Click to select the start and stop with **host** radio button.

Figure 216 NTP Daemon



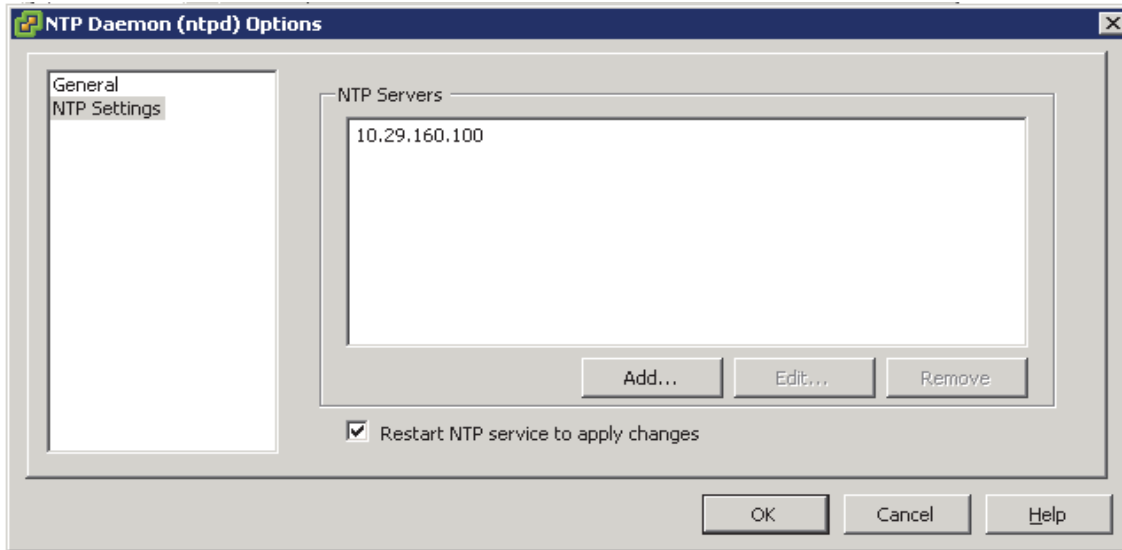
18. Click on NTP Settings option.
19. Click on **Add** button to add the NTP server's IP address.
20. Press **OK** to continue.

Figure 217 Adding a new NTP Server to the ESXi NTP Settings



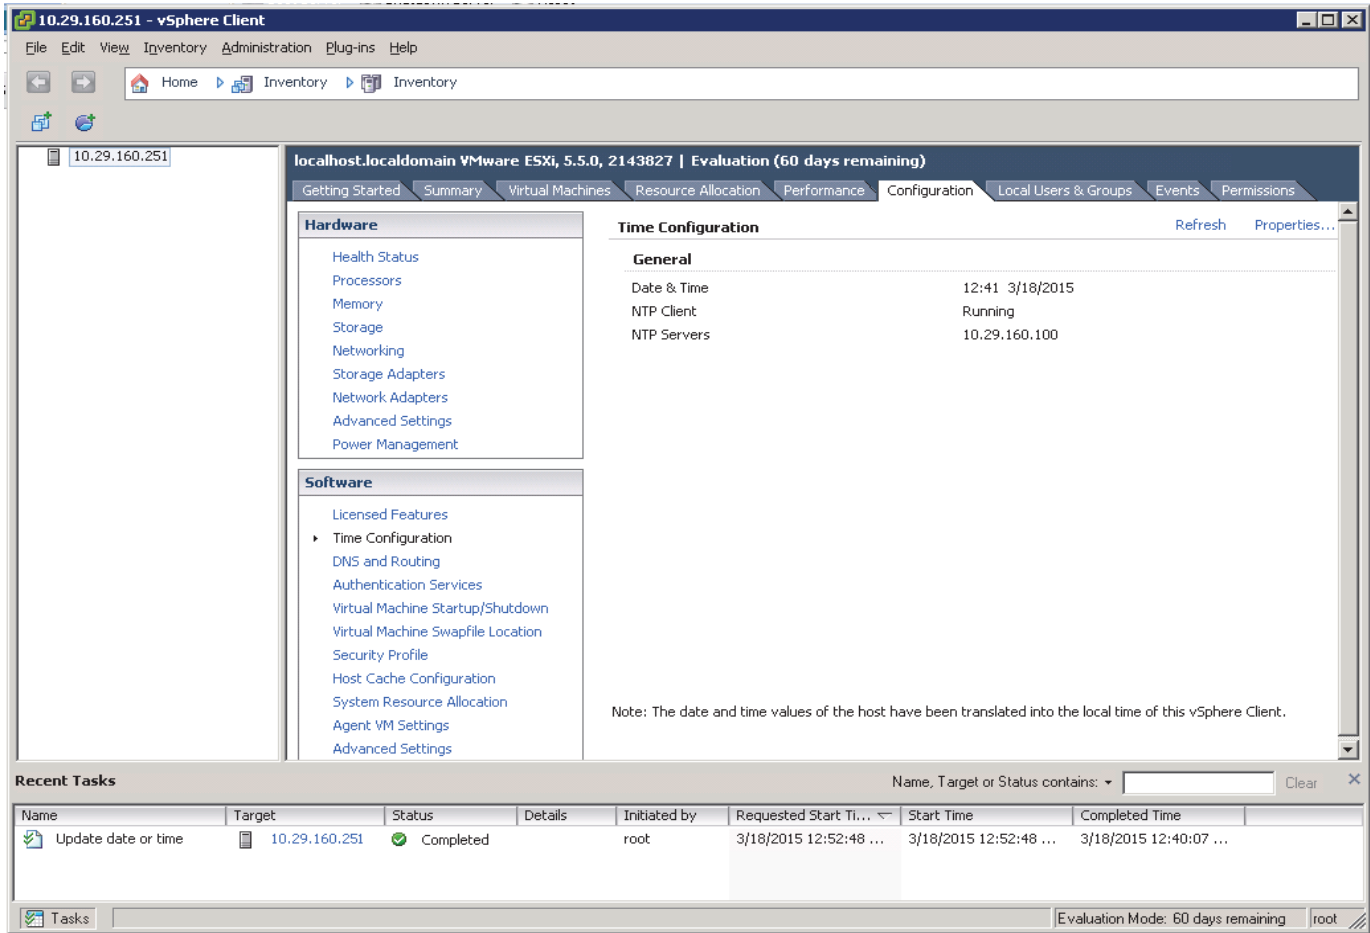
21. In the next screen, verify the IP-address in the NTP Servers list.
22. Click on the checkbox **Restart NTP service to apply changes**.
23. Press the button **OK** twice to complete the time configurations.

Figure 218 Restart NTP Service



24. Time configuration option would now show that the NTP client is running, along with the IP address of the NTP client.

Figure 219 Verifying the NTP Client



Downloading the UCS Director Express Software Components

The software components of UCS Director Express for Big Data need to be downloaded from three different locations.

Table 21 Cisco UCS Director Express Big Data 1.1 Software Components

Software component	File Names	Link to Download
Cisco UCS Director Express 1.0 OVF	CUCSD_Express_1_0_0_0_GA.zip	https://software.cisco.com/download/release.html?mdfid=286281255&flowid=71403&softwareid=285018084&release=1&relind=AVAILABLE&relicycle=&reltype=latest
Cisco UCS Director 5.2.0.1 patch	cucsd_patch_5_2_0_1.zip	https://software.cisco.com/download/release.html?mdfid=286283454&flowid=72903&softwareid=285018084&release=5&relind=AVAILABLE&relicycle=&reltype=latest
Cisco UCS Director Baremetal Agent 5.2 OVF	CUCSD_BMA_5_2_0_0_VMWARE_GA.zip	https://software.cisco.com/download/release.html?mdfid=286284995&flowid=73724&softwareid=285018084&release=1&relind=AVAILABLE&relicycle=&reltype=latest
Cisco UCS Director Express for Big Data 1.1 Upgrade Package	UCSDEExpress_Big_Data_1.1_Upgrade_Package.zip	https://software.cisco.com/download/release.html?mdfid=286284995&flowid=73724&softwareid=285018084&release=1&relind=AVAILABLE&relicycle=&reltype=latest
25. Cisco UCS Director Express for Big Data BMA Update Package	UCSDEExpress_BMA_Big_Data_1.1_Upgrade_Package.zip	https://software.cisco.com/download/release.html?mdfid=286284995&flowid=73724&softwareid=285018084&release=1&relind=AVAILABLE&relicycle=&reltype=latest

Download the Software Components

- Using the links provided Table 15 above, download the Cisco UCS Director Express for Big Data 1.1 OVF Appliance zip file.

Figure 220 Cisco UCS Director Express for Big Data 1.0 Download Page

The screenshot shows the Cisco website's download page for UCS Director Express for Big Data 1.0. The page is titled 'Download Software' and includes a navigation menu with options like 'Products & Services', 'Support', and 'How to Buy'. The main content area is titled 'UCS Director Express for Big Data 1.0' and features a 'Release 1' section. This section contains a table of software releases with columns for 'File Information', 'Release Date', and 'Size'. Each row includes a 'Download' button and an 'Add to cart' button. The row for 'Cisco UCS Director Express for Big Data 1.0 (OVF Appliance) MD5 Checksum' is highlighted with an orange border. Below the table, there is a 'Related Information' section with a 'Dashboard Information Sources' subsection.

File Information	Release Date	Size
Cisco UCS Director Bare Metal Agent Patch for Cisco UCS Director Express F or Big Data (Patch need to be applied on top Cisco UCS Director BMA 5.0. MD5 Checksum - 5b2a6c11950f07837e29bdcc52dca301)	19-NOV-2014	10.37 MB
Cisco UCS Director Express For Big Data Patch (Patch needs to be applied on 1.0. MD5 Checksum - ca44a9a25057af5072acafaf7fc7d933)	19-NOV-2014	1.76 MB
Cisco UCS Director Express Hotfix for Bash Code Injection Vulnerability (Bash Shellshock - CVE-2014-6271, CVE-2014-7169) Note: Patch has README that explains how to apply this patch	06-OCT-2014	1.82 MB
Cisco UCS Director Express for Big Data 1.0 (OVF Appliance) MD5 Checksum (8d6cb7dc36107ca5c1f93a9faf69d49c)	05-SEP-2014	2663.09 MB
Cisco UCS Director Express for Big Data BMA Update Package MD5 Checksum (517fa2a881b8cab6dff0c3ad17a1cc9b)	05-SEP-2014	343.95 MB

- Using the links provided Table 15 above; download the Cisco UCS Director 5.2.0.1 Patch zip file, and Cisco UCS Director Baremetal Agent 5.2 VMware vSphere OVF Appliance zip file.

Figure 221 Cisco UCS Director 5.2 Download Page

Downloads Home > Products > Servers - Unified Computing > UCS Director > UCS Director 5.2 > UCS Director Virtual Appliance Software-5

UCS Director 5.2

[Expand All](#) | [Collapse All](#)

- ▼ Latest 5
- ▼ All Releases
- ▶ 5

Release 5

CUCSD 5.2.0.1 Patch

File Information	Release Date	Size	
Cisco UCS Director 5.2.0.1 Patch (Patch need to be applied on top of 5.2 MD5 Checksum - 1ef745cd8bbd43a46aa1398247dbfc1c) cucsd_patch_5_2_0_1.zip	03-FEB-2015	1141.61 MB	<input type="button" value="Download"/> <input type="button" value="Add to cart"/>
Cisco UCS Director 5.2.0.0A HOTFIX Patch (PSIRT FIX FOR NTP - Patch need to be applied on top of 5.2.0.0 MD5 Checksum - 24f9a3c0c2c6aa1ab83fc0da70cf5ce7) cucsd_patch_5_2_0_0A.zip	15-JAN-2015	1.45 MB	<input type="button" value="Download"/> <input type="button" value="Add to cart"/>
Cisco UCS Director 5.2 (HyperV Appliance) MD5 Checksum - f04047c63e5c1422ff49fe575a77d143 CUCSD_5_2_0_0_HYPERV_GA.zip	20-DEC-2014	9344.73 MB	<input type="button" value="Download"/> <input type="button" value="Add to cart"/>
Cisco UCS Director 5.2 (VMWare vSphere OVF Appliance. MD5 Checksum - 06bfb6fe95aabef9c69555b535946363) CUCSD_5_2_0_0_VMWARE_GA.zip	20-DEC-2014	2869.15 MB	<input type="button" value="Download"/> <input type="button" value="Add to cart"/>
Cisco UCS Director Baremetal Agent 5.2 (HyperV Appliance MD5 Checksum - 0fd872b48f9f302416b6769a247cbbec) CUCSD_BMA_5_2_0_0_HYPERV_GA.zip	20-DEC-2014	8195.32 MB	<input type="button" value="Download"/> <input type="button" value="Add to cart"/>
Cisco UCS Director Baremetal Agent 5.2 (VMWare vSphere OVF Appliance MD5 Checksum - a0c34c4c924720dc9d2f9b099c5b9b5c) CUCSD_BMA_5_2_0_0_VMWARE_GA.zip	20-DEC-2014	1857.43 MB	<input type="button" value="Download"/> <input type="button" value="Add to cart"/>

Add Devices

Add Notification

3. Using the links provided Table 21 above; download the Cisco UCS Director 5.2.0.1 Patch zip file, and the Cisco UCS Director Baremetal Agent 5.2 VMWare vSphere OVF Appliance zip file.

Figure 222 Cisco UCS Director Express for Big Data 1.1 Download Page

[Downloads Home](#) > [Products](#) > [Servers - Unified Computing](#) > [UCS Director](#) > [UCS Director Express for Big Data 1.1](#) > [UCS Director Virtual Appliance Software-1](#)

UCS Director Express for Big Data 1.1

Search... [Add Devices](#)
[Expand All](#) | [Collapse All](#) [Add Notification](#)

▼ Latest **1**
 ▼ All Releases
 ▶ 0

Release 1

Cisco UCSD Express 1.1 (Upgrade Package and BMA Patch)

File Information	Release Date	Size	
Cisco UCS Director Express for Big Data 1.1 BMA Update Package (MD5 Checksum 25e434da9b06465cade4902e0e5b0d81) UCSDExpress_BMA_5_2_Big_Data_1.1_Upgrade_Package.zip	10-MAR-2015	353.13 MB	Download Add to cart
Cisco UCS Director Express for Big Data 1.1 Upgrade Package (MD5 Checksum 8748164497a2b42ee4ba079098a0a1e3) UCSDExpress_Big_Data_1.1_Upgrade_Package.zip	10-MAR-2015	2.05 MB	Download Add to cart

4. Please all the files in a directory in the client windows workstation.
5. Unzip the contents of the CUCSD_Express_1_0_0_0_GA.zip and CUCSD_BMA_5_2_0_0_VMWARE_GA.zip.

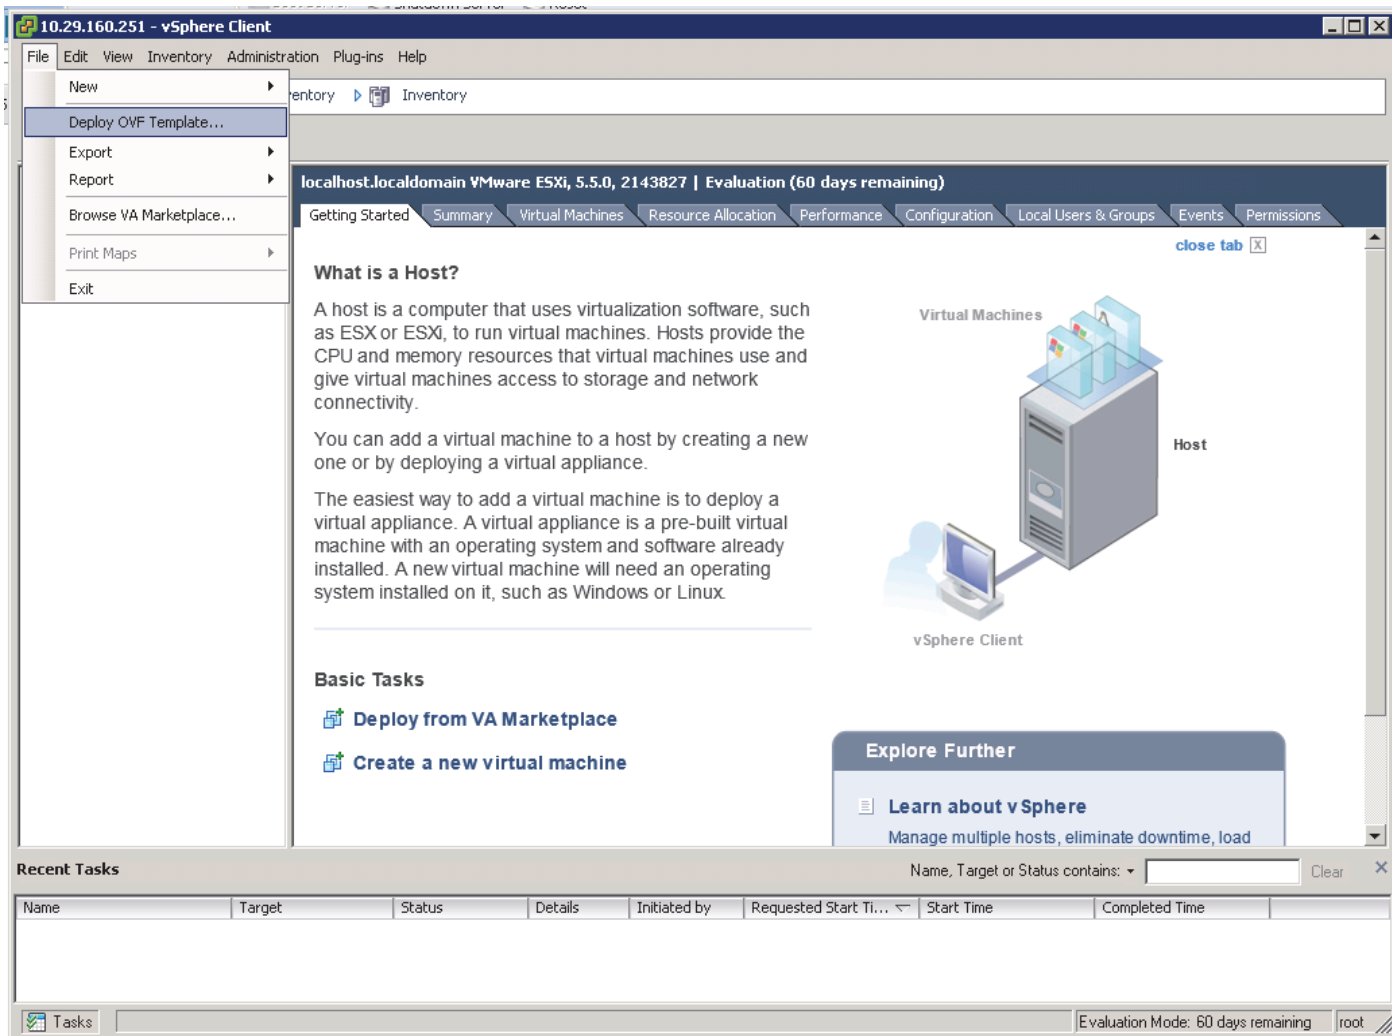
Installing Cisco UCS Director Express for Big Data

The Cisco UCS Director Express for Big Data shall be installed on the VMWare vSphere hypervisor using the vSphere Client software.

Deploying the Cisco UCS Director Baremetal Agent OVF

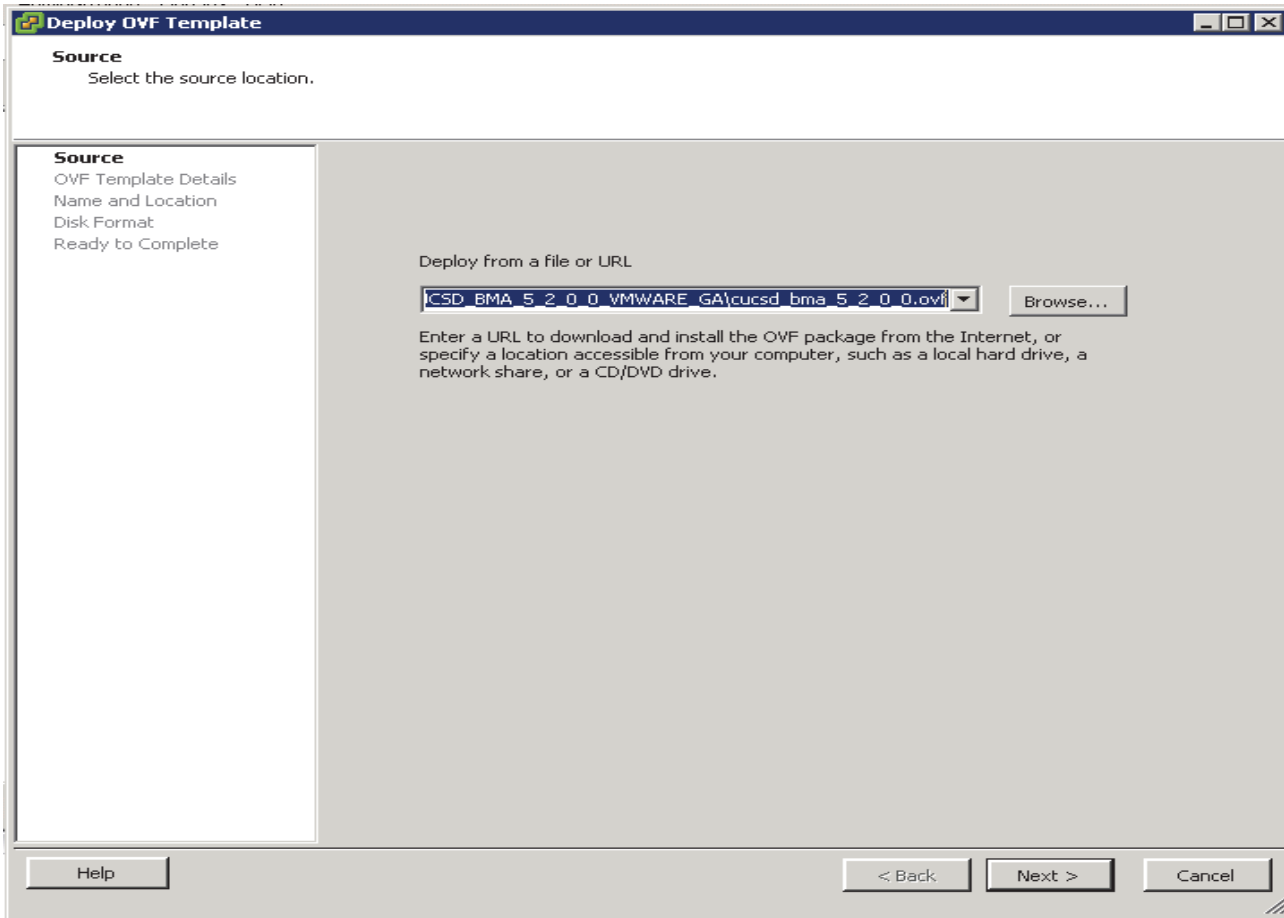
1. Launch the VMWare vSphere client software
2. Enter the chosen IP address, the username as root, and the chosen password.
3. Click on **Login** to continue.
4. From the **File** menu, Select **Deploy OVF Template**.

Figure 223 Deploy OVF in the vSphere Client



5. Choose the Cisco UCS Director Baremetal Agent 5.2.0.0 OVF template. Click **Open**.
6. Click **Next** to continue.

Figure 224 Select the Cisco UCS Director Baremetal Agent OVF file



7. Review the details of the OVF template, Click **Next**.
8. Accept the End User License Agreement. Click **Next** to continue.
9. In the **Name and Location** option, Enter the name of the VM. Click **Next** to continue.

Figure 225 Enter Cisco UCS Director Baremetal Agent VM Name

The screenshot shows a window titled "Deploy OVF Template" with a sub-header "Name and Location" and the instruction "Specify a name and location for the deployed template". On the left, a sidebar lists the wizard steps: "Source", "OVF Template Details", "End User License Agreement", "Name and Location" (highlighted), "Disk Format", "Network Mapping", and "Ready to Complete". The main area has a "Name:" label and a text input field containing "CUCSD-BM-5.2.0.0_36". Below the input field, a note states: "The name can contain up to 80 characters and it must be unique within the inventory folder." At the bottom, there are four buttons: "Help", "< Back", "Next >", and "Cancel".

10. In the Disk Format option, click the **Thick Provision Lazy Zeroed** radio button. Click **Next** to continue.

Figure 226 Select the Disk Format for the VM

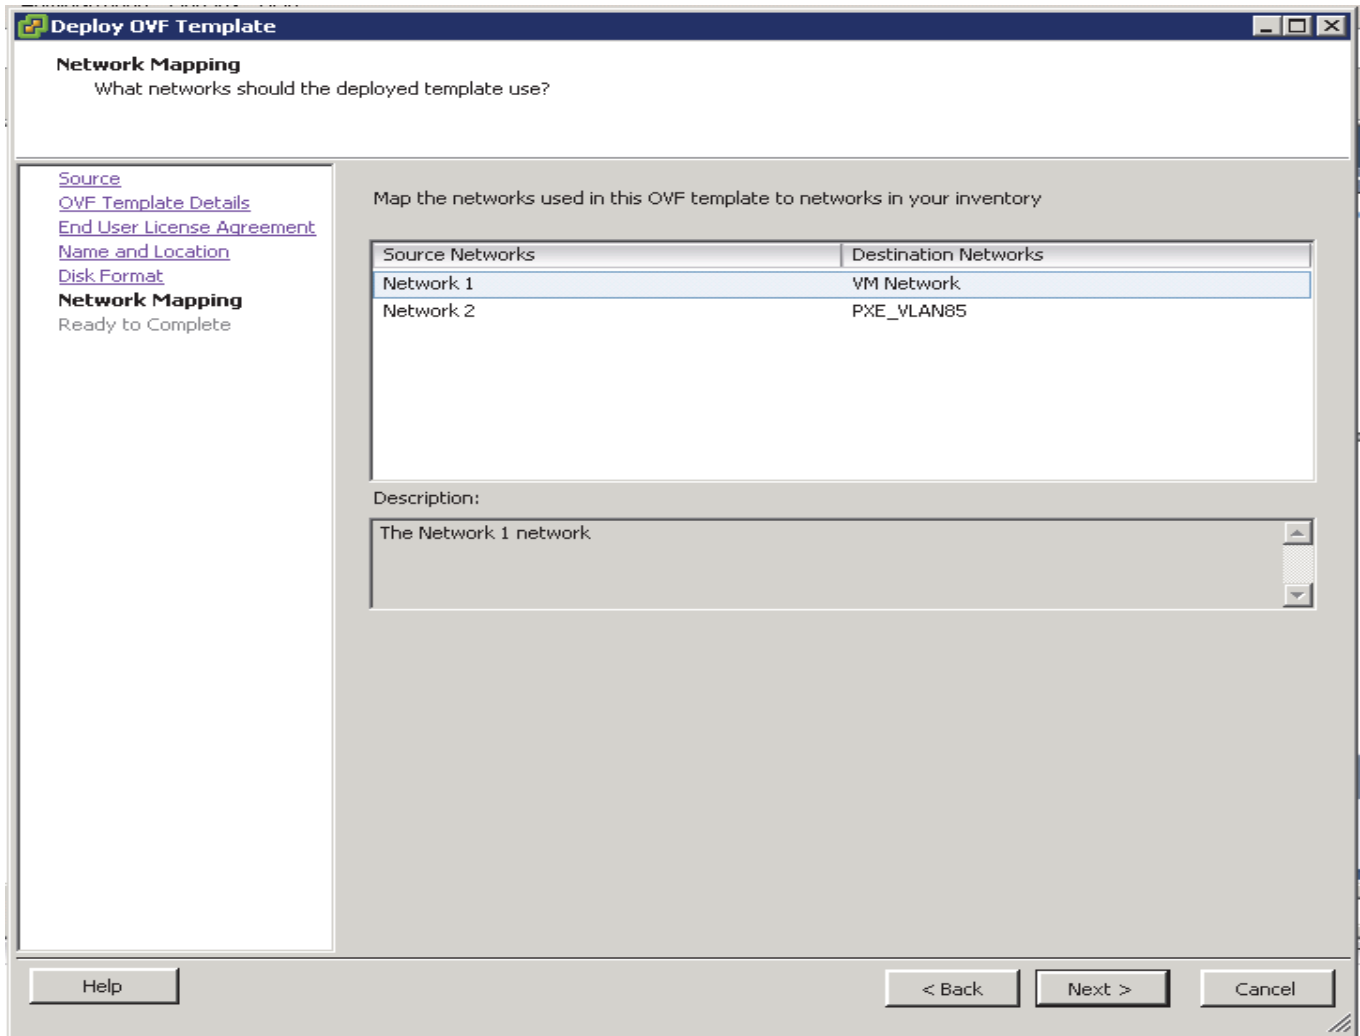
The screenshot shows the 'Deploy OVF Template' wizard in the 'Disk Format' step. The window title is 'Deploy OVF Template'. Below the title bar, the text reads 'Disk Format' and 'In which format do you want to store the virtual disks?'. On the left side, there is a navigation pane with links: 'Source', 'OVF Template Details', 'End User License Agreement', 'Name and Location', 'Disk Format' (which is highlighted), 'Network Mapping', and 'Ready to Complete'. The main area contains the following fields and options:

- Datstore:** A text box containing 'datastore1'.
- Available space (GB):** A text box containing '263,5'.
- Provisioning options:** Three radio buttons are present:
 - Thick Provision Lazy Zeroed
 - Thick Provision Eager Zeroed
 - Thin Provision

At the bottom of the window, there are three buttons: 'Help', '< Back', and 'Next >', and a 'Cancel' button on the far right.

11. In the Network Mapping option,
 - Choose **VM Network** as the destination network for source Network 1.
 - Choose **PXE_VLAN85** as the destination network for source Network 2.
12. Click **Next** to continue.

Figure 227 Network Mapping for Deployed Template



- Review the details of the VM, click the check box **Power on after deployment** and click **Finish** to proceed with the VM deployment.

Figure 228 Deploy the Cisco UCS Director Baremetal Agent VM

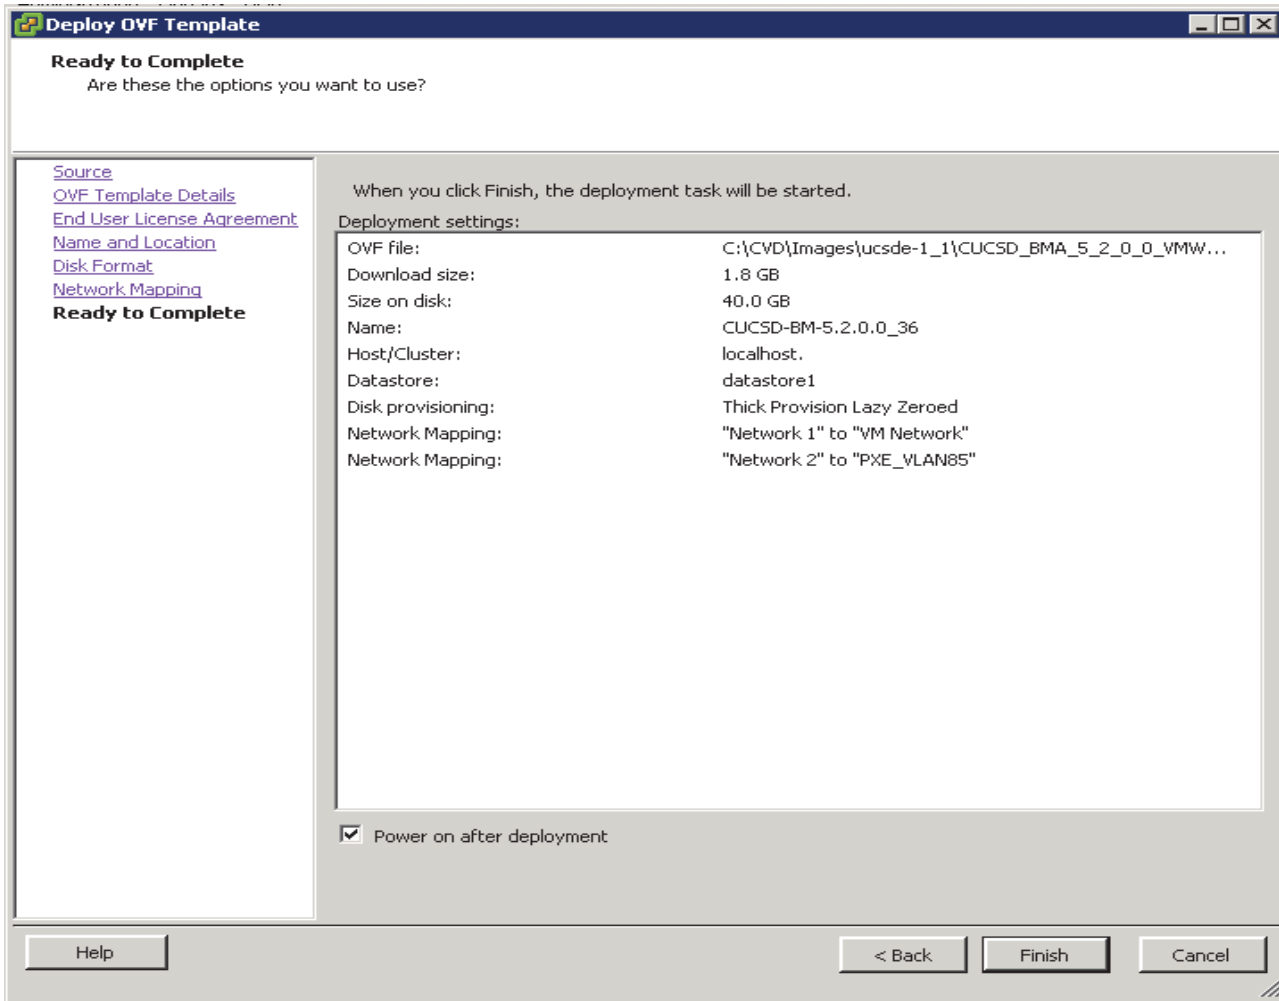
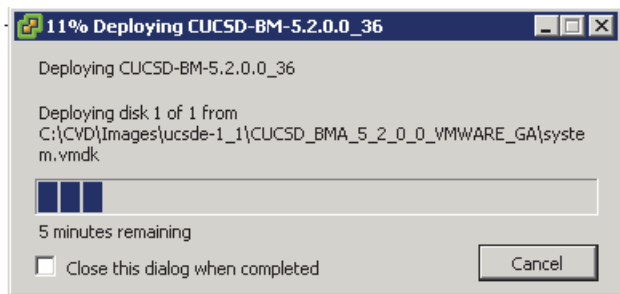


Figure 229 Cisco UCS Director Baremetal Agent VM Deployment in Progress

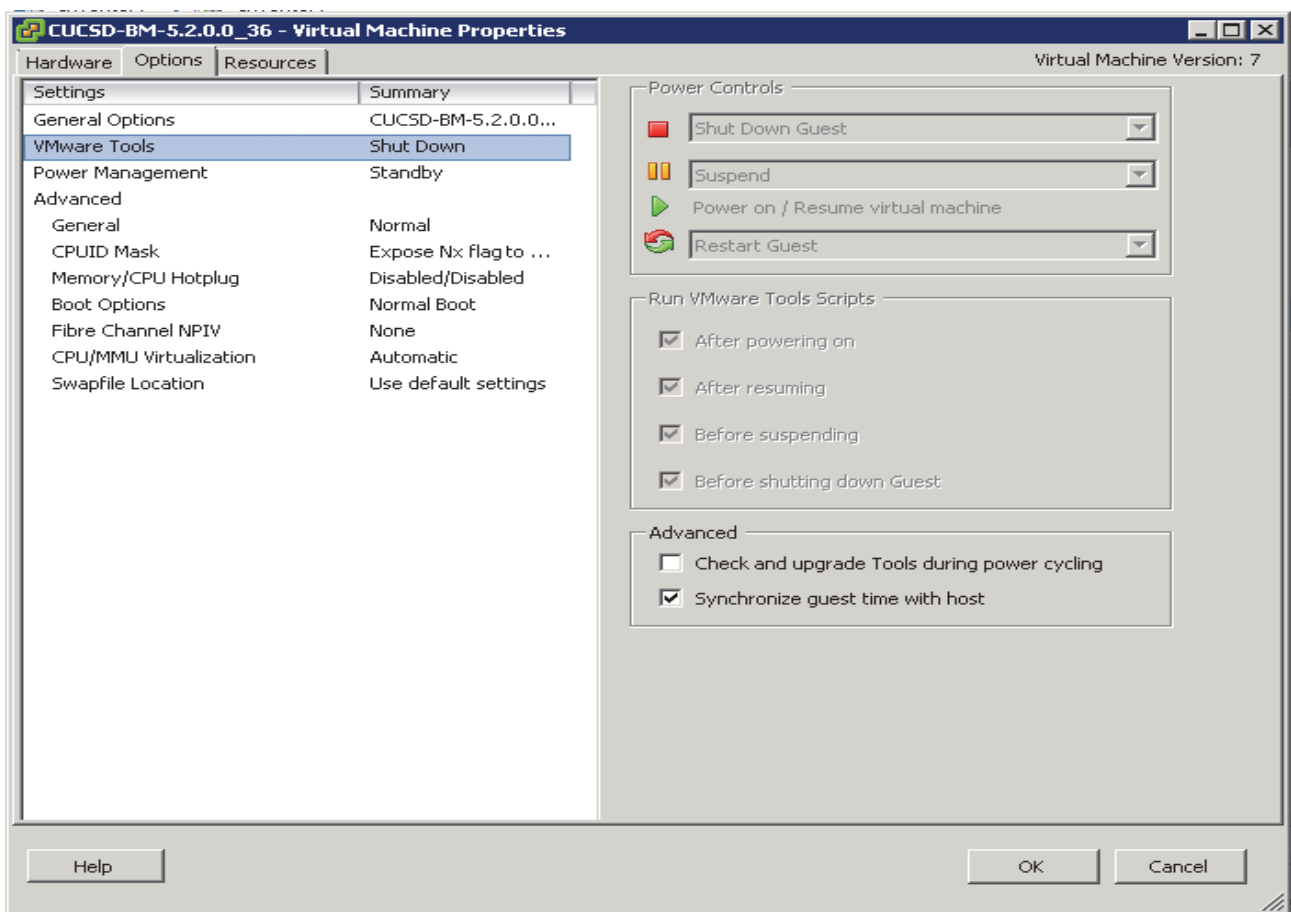


Configuring the Cisco UCS Director Baremetal Agent VM (BMA-VM)

The Cisco UCS Director Baremetal Agent VM named as CUCSD-BM-5.2.0.0_36 shall be known as BMA-VM here onwards.

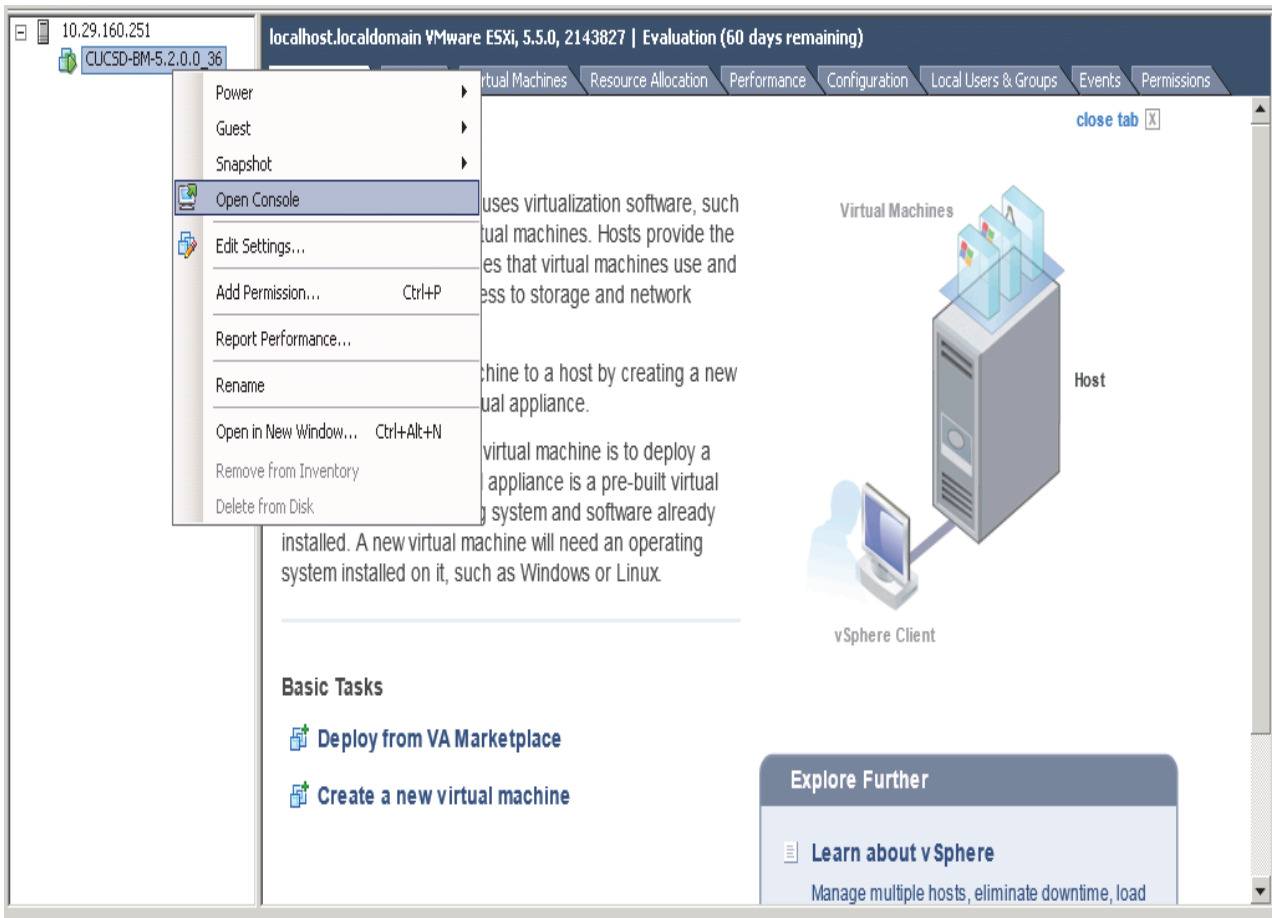
1. Right click on the BMA-VM, and select **Edit Settings**.
2. In the Virtual Machine Properties dialog box, click on the Options Tab.
3. Click on the VMWare **Tools**, Click on the **Synchronize guest time with host** option in the Advanced **section**.
4. Click on **OK** button to accept the changes.

Figure 230 Edit VM Settings to Synchronize the Guest Time with the ESXi Host



5. Right click on the BMA-VM, and select **Open Console**.

Figure 231 Access the VM Console of the BMA-VM



- In the console accept the End User License Agreement by typing **yes** and press **ENTER**.

Figure 232 Accept the EULA

```
Do you agree with the terms of the End User License Agreement?
yes/no [no]: yes_
```

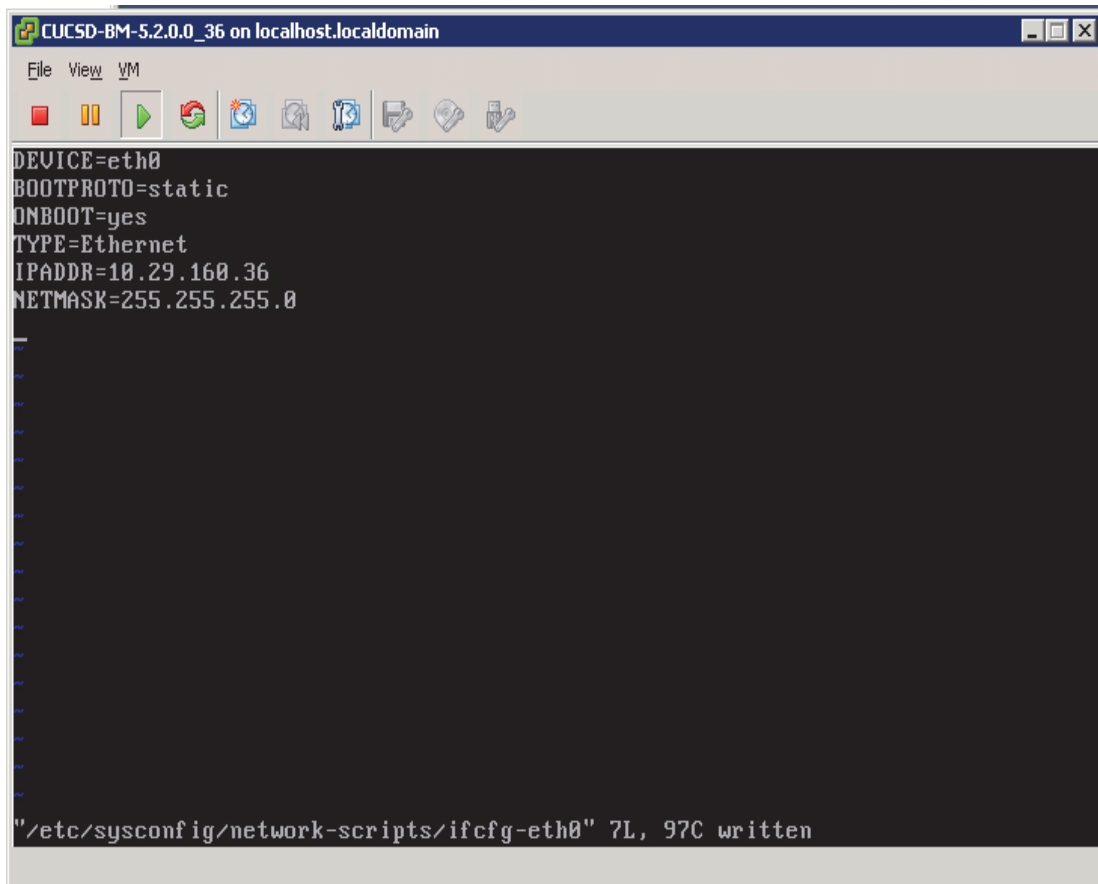
- Login as **root** user using the default password **pxeboot**.
- Configure the network interfaces by editing the `ifcfg-eth0` and `ifcfg-eth1` files located at `/etc/sysconfig/network-scripts/` directory, as follows:

Table 22 BMA-VM network configurations

Network Interface	Configuration
eth0	IP Address: 10.29.160.36, Subnet Mask: 255.255.255.0

eth1	IP Address: 192.168.85.36, Subnet Mask: 255.255.255.0
------	-------------------------------------------------------

Figure 233 Editing the BMA-VM NIC eth0



The screenshot shows a terminal window titled "CUCSD-BM-5.2.0.0_36 on localhost.localdomain". The terminal displays the following configuration for the eth0 interface:

```
DEVICE=eth0
BOOTPROTO=static
ONBOOT=yes
TYPE=Ethernet
IPADDR=10.29.160.36
NETMASK=255.255.255.0
```

At the bottom of the terminal, a message indicates that the configuration has been saved: `"/etc/sysconfig/network-scripts/ifcfg-eth0" 7L, 97C written`.

Figure 234 Editing the BMA-VM NIC eth1

The screenshot shows a terminal window titled "CUCSD-BM-5.2.0.0_36 on localhost.localdomain". The terminal displays the following configuration for the eth1 interface:

```

DEVICE=eth1
BOOTPROTO=static
ONBOOT=yes
TYPE=Ethernet
IPADDR=192.168.85.36
NETMASK=255.255.255.0

```

At the bottom of the terminal, a message indicates that the configuration file has been written:

```
"/etc/sysconfig/network-scripts/ifcfg-eth1" 7L, 98C written
```

- Restart the network service by using the service command.
service network restart

Figure 235 Restart the network

```

[root@localhost ~]# service network restart
Shutting down interface eth0: [ FAILED ]
Shutting down interface eth1: [ OK ]
Shutting down loopback interface: [ OK ]
Bringing up loopback interface: [ OK ]
Bringing up interface eth0: [ OK ]
Bringing up interface eth1: [ OK ]

```


Installing the Cisco UCS Director Express Big Data Upgrade Package

1. Copy over the `UCSDEExpress_BMA_5.2_Big_Data_1.1_Upgrade_Package.zip` that was downloaded from `cisco.com` to this VM, by using a secure shell FTP session.
2. Unzip the contents in a temporary staging directory.
3. Change directory into the `scripts/bin` directory.
4. Change the permissions to add execute permissions to the `copyfiles.sh` script file and execute it.

`chmod +x copyfiles.sh`

Figure 236 Install the Cisco UCS Director Express Big Data Upgrade Package

```
[root@localhost stage]# ls
CentOSLive      bd_bma_version.info  feature-bigdata-intel.jar
Hortonworks-2.1 cloudera-5.0.1       mapr_common_templates
Hortonworks-2.2 cloudera-5.2.0       ntp_server_config.sh
MapR-3.1.1      cloudera-5.2.1       run.sh.template
MapR-4.0.1      cloudera-5.3.0       scripts
bd-sw-rep      common_templates     templates
[root@localhost stage]# cd scripts/bin
[root@localhost bin]# chmod +x ./copyfiles.sh
```

5. Execute the `copyfiles.sh` script.

`./copyfiles.sh`

This script copies the number of software modules such as CentOSLive image into the BMA-VM and creates a new repository directory by name `bd-sw-rep` under the `/opt/cnsaroot` directory. This new directory acts as the repository of all the Big Data specific 3rd party hadoop distribution directories.

Configuring the Big Data software repositories

Copy the Contents of RHEL6.5 ISO into the BMA-VM

1. Copy over the contents of the RHEL6.5 ISO into the directory `/opt/cnsaroot/images/RHEL6.5` on the BMA-VM.
2. Copy the contents of the directory `/opt/cnsaroot/images/RHEL6.5/isolinux` into the directory `/opt/cnsaroot/RHEL6.5`.

Figure 237 Copy the Contents of RHEL6.5 ISO into the BMA-VM

```
[root@localhost ~]# cd /opt/cnsaroot/RHEL6.5
[root@localhost RHEL6.5]# cp /opt/cnsaroot/images/RHEL6.5/isolinux/* .
[root@localhost RHEL6.5]# ls
TRANS.TBL  boot.msg  initrd.img  isolinux.cfg  splash.jpg  vmlinuz
boot.cat   grub.conf  isolinux.bin  memtest      vesamenu.c32
```

Download and Place the Common Utility files in BMA-VM

- From a host connected to the Internet, download the Parallel-SSH and Cluster-Shell utility tools and copy them over to the `/opt/cnsaroot/bd-sw-rep` directory.
 - Download Parallel SSH archive from <https://pypi.python.org/packages/source/p/pssh/pssh-2.3.1.tar.gz>
 - Download Cluster-Shell RPM package from http://dl.fedoraproject.org/pub/epel/6/x86_64/clustershell-1.6-1.el6.noarch.rpm

Figure 238 Copy the Cluster-Shell and Passwordless-SSH Utilities

```
-rw-r--r-- 1 root root 250400 Feb 18 21:18 clustershell-1.6-1.el6.noarch.rpm
-rw-r--r-- 1 root root 23427 Feb 18 21:17 pssh-2.3.1.tar.gz
[root@localhost bd-sw-rep]# pwd
/opt/cnsaroot/bd-sw-rep
[root@localhost bd-sw-rep]#
```

- By following the instructions on this page of the BMA-Install guide, download and copy over the Hadoop Distro RPMs into their respective directories under `/opt/cnsaroot/bd-sw-rep`. http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/ucs-director-express/bma-install-config/1-1/b_ucsd_express_bma_install_config_guide_1-1/b_ucsd_express_bma_install_config_guide_chapter_0101.html#reference_F3FE769E6A114DAD8CD5E3296556B70E
- Upload the appropriate License files to the Hadoop distribution directories
 - Place the Cloudera License in a file called `ClouderaEnterpriseLicense.lic` and place it under the `/opt/cnsaroot/bd-sw-rep/cloudera05.x.y`.
 - Place the MapR license in a file called `license.txt` MapR License and place it under the directory `/opt/cnsaroot/bd-sw-rep/MapR-X.Y.Z`.



Note Hortonworks distribution does not require any license file.

Figure 239 Copy the RPM Packages for the Hadoop Distributions

```
[root@localhost ~]# cd /opt/cnsaroot/bd-sw-rep/
[root@localhost bd-sw-rep]# ls cloudera-5.3.0/
ClouderaEnterpriseLicense.lic  cm5.3.0-centos6.tar.gz
catalog.properties           mysql-connector-java-5.1.26.tar.gz
cdh5.3.0-centos6.tar.gz      userrpmlist.txt
[root@localhost bd-sw-rep]# ls Hortonworks-2.2/
HDP-2.2.0.0-centos6-rpm.tar.gz  catalog.properties
HDP-UTILS-1.1.0.20-centos6.tar.gz  openssl-1.0.1e-30.el6.x86_64.rpm
ambari-1.7.0-centos6.tar.gz      userrpmlist.txt
[root@localhost bd-sw-rep]# ls MapR-4.0.2
catalog.properties           mapr-v4.0.2GA.rpm.tgz
catalog.properties.txt      mapr-whirr-0.7.0.16780-1.noarch.rpm
ext-2.2.zip                 pdsh-2.27-1.el6.rf.x86_64.rpm
libgenders-1.14-2.el6.rf.x86_64.rpm  soci-3.2.1-1.el6.x86_64.rpm
libgenders-devel-1.14-2.el6.rf.x86_64.rpm  soci-mysql-3.2.1-1.el6.x86_64.rpm
license.txt                 sshpass-1.05-1.el6.x86_64.rpm
mapr-drill-0.7.0.29434-1.noarch.rpm  userrpmlist.txt
mapr-ecosystem-20150205.rpm.tgz
[root@localhost bd-sw-rep]#
```

Setup a UCSD Patch Directory in the BMA-VM

Cisco UCS Director Express for Big Data VM which will be installed in the next section, requires the patches to be kept in a web server. The BMA-VM comes pre-configured with a web-server used during PXE booting process. This section walks through the steps to create a directory to hold these patches in the BMA-VM.

1. In BMA-VM, create a directory by name patches under /var/www/html.

```
mkdir /var/www/html/patches
```

2. Copy over the Cisco UCS Director Express for Big Data 1.1 specific patch files (See Table 3) to this patch directory.

Figure 240 Setup a UCSD Patch Directory in the HTTP Root Path

```
[root@localhost ~]# ls -l /var/www/html/patches
total 1172256
-rw-r--r-- 1 root root 2139421 Feb 18 04:52 UCSDExpress_Big_Data_1.1_Upgrade_Package.zip
-rw-r--r-- 1 root root 1197064934 Feb 3 13:16 cucsd_patch_5_2_0_1.zip
```

3. Start the HTTPD server in the BMA-VM.

```
service httpd start
```

Figure 241 Start the HTTPD

```
[root@localhost bd-sw-rep]# service httpd start
Starting httpd: [ OK ]
```

4. Verify if these files are accessible by visiting the URL <http://<BMA-VM's >IP address/patches/>.

Figure 242 Verify the Accessibility of the Cisco UCS Director Express Patches

Name	Last modified	Size	Description
Parent Directory		-	
UCSDExpress_Big_Data_1.1_Upgrade_Package.zip	18-Feb-2015 04:52	2.0M	
cucsd_patch_5_2_0_1.zip	03-Feb-2015 13:16	1.1G	

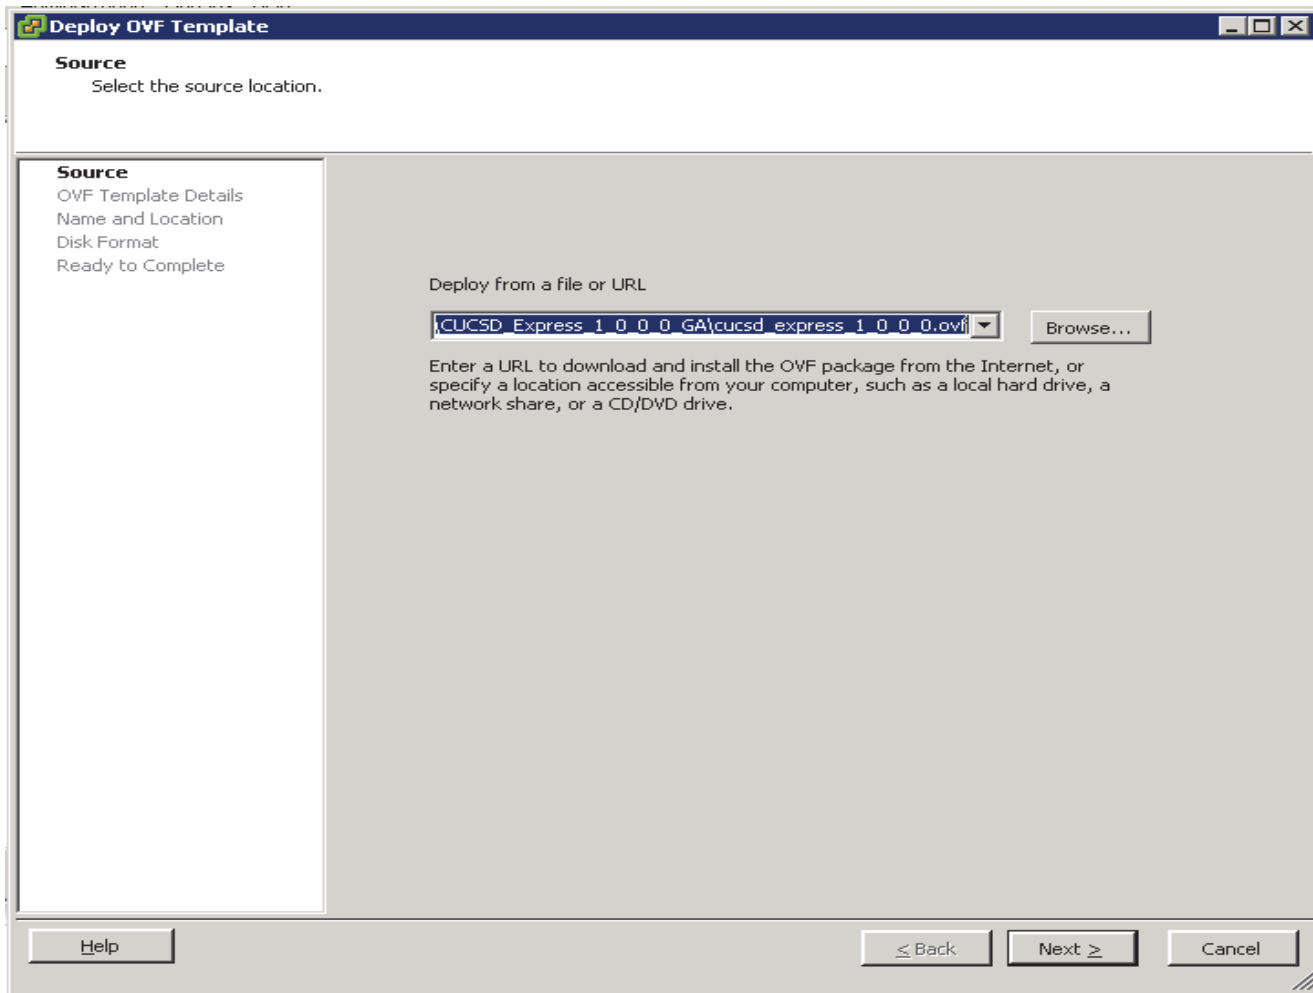
Apache/2.2.3 (CentOS) Server at 10.29.160.36 Port 80

BMA-VM configurations are complete.

Deploying the Cisco UCS Director Express OVF

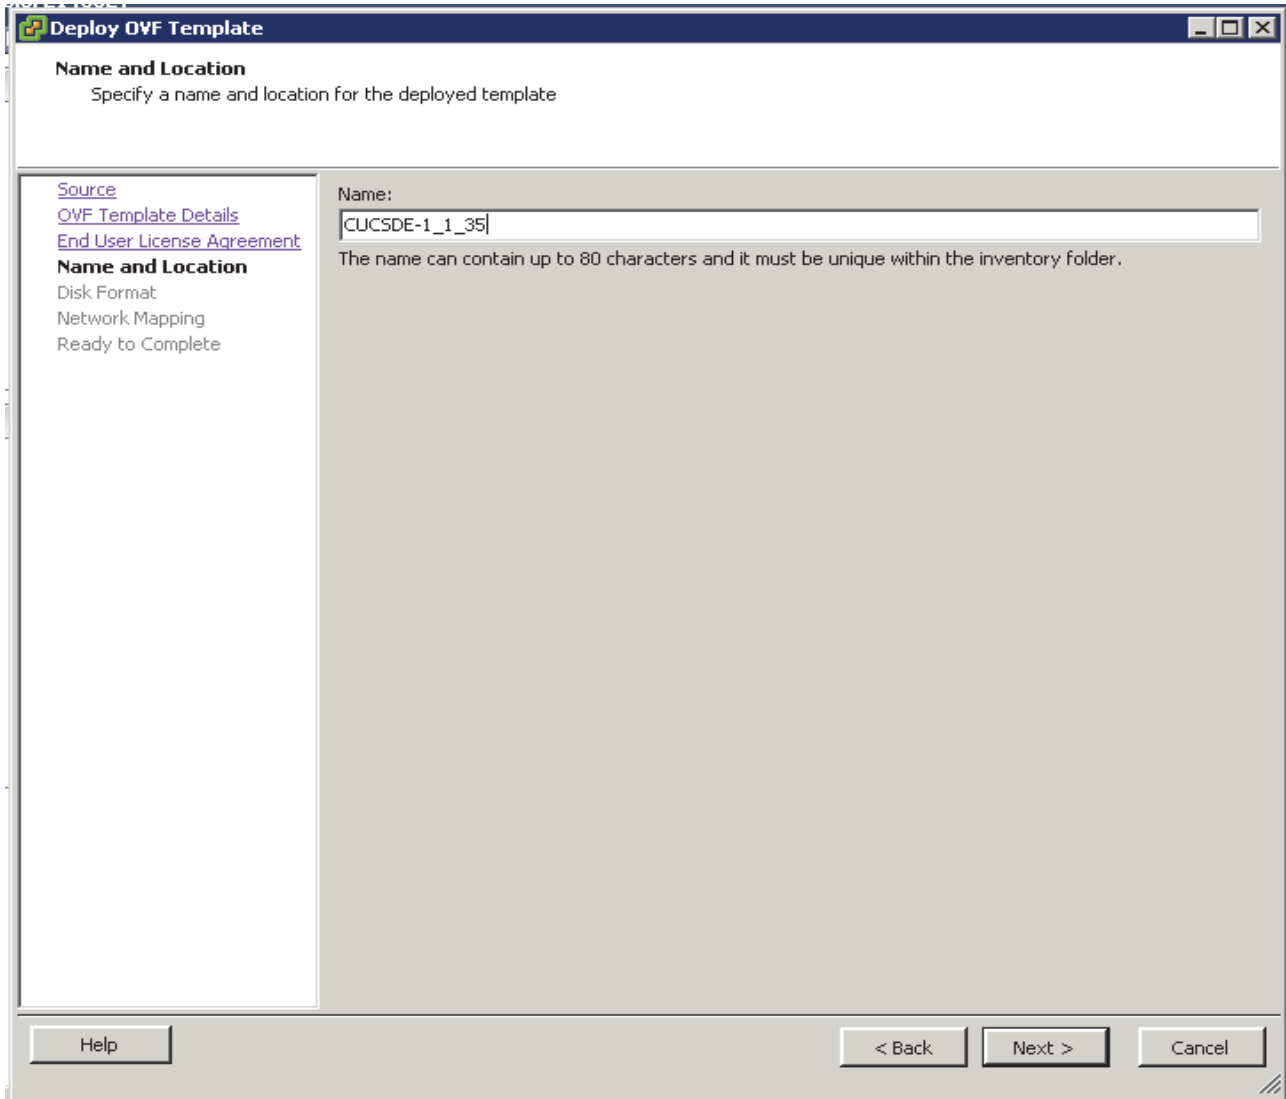
1. Launch the VMWare vSphere client software
2. Enter the chosen IP address, the username as root, and the chosen password.
3. Click **Login** to continue.
4. From the **File** menu, Select **Deploy OVF Template**.
5. Choose the Cisco UCS Director Express for Big Data 1.0 OVF template. Click **Open**.

Figure 243 Deploy the Cisco UCSD Express 1.0 OVF



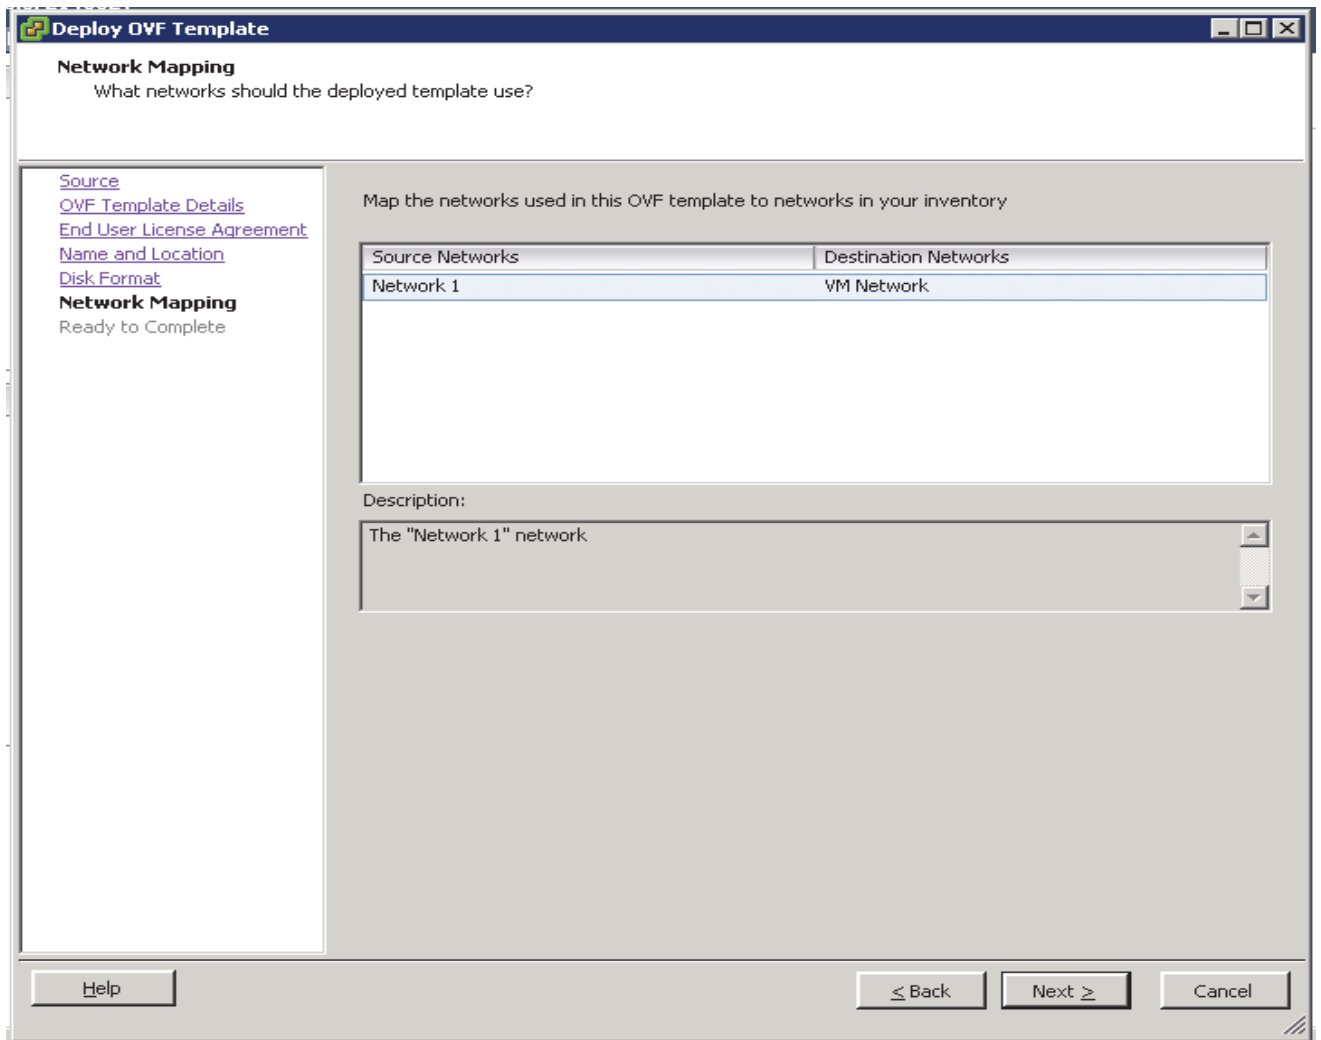
6. Review the details of the OVF, and Click **Next** to continue.
7. Accept the EULA, Click **Next** to continue.
8. Name the VM, Click **Next** to continue.

Figure 244 Name the Cisco UCS Director Express VM



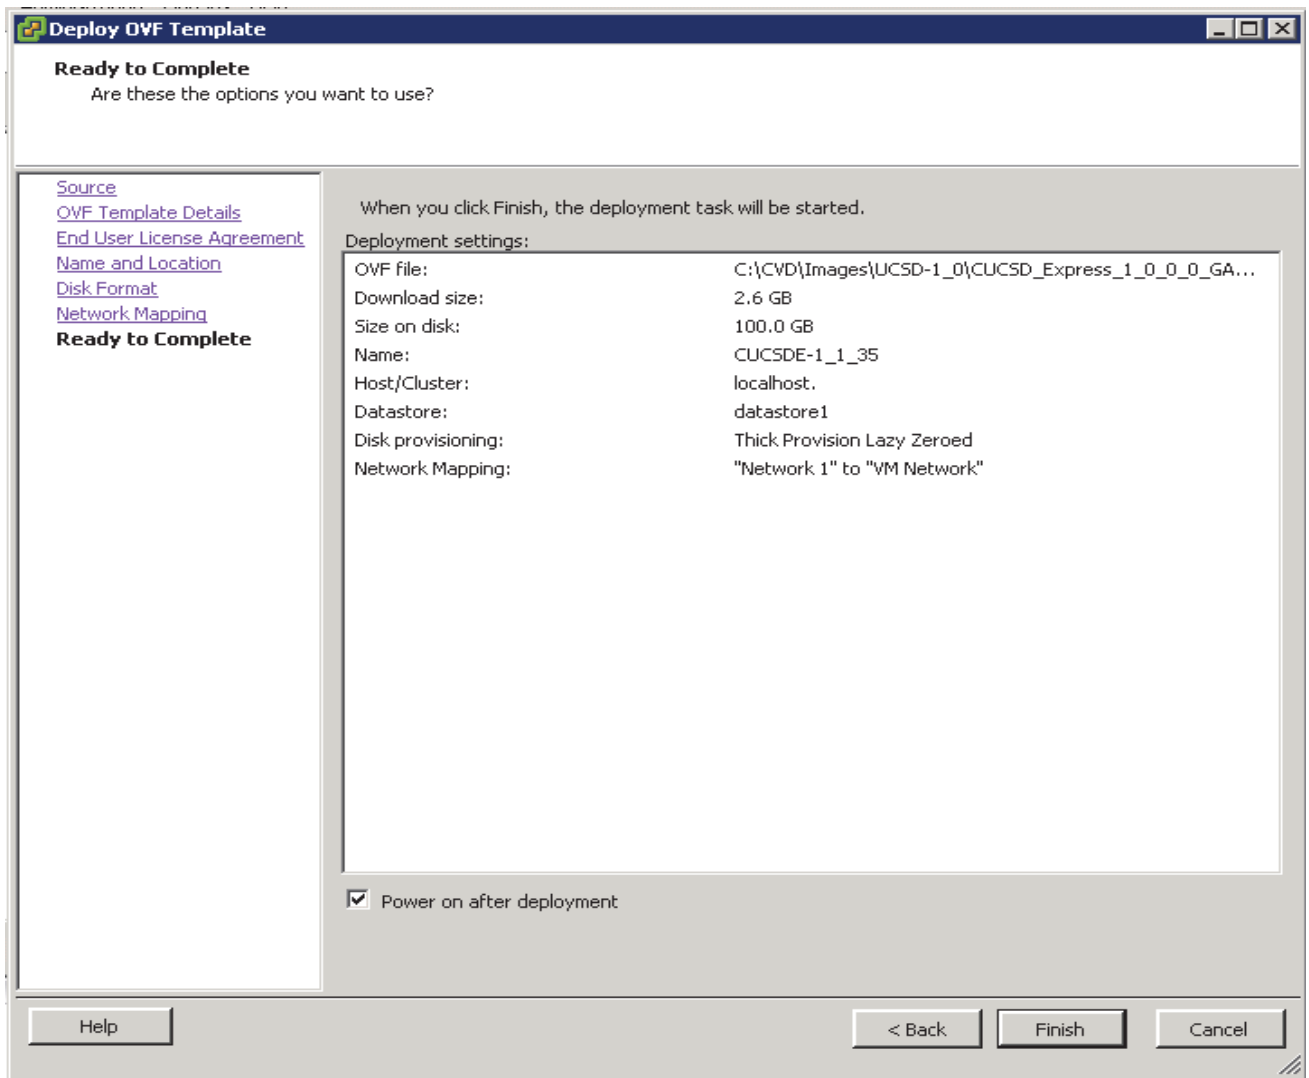
9. Choose the destination network **VM Network** for the source network **Network 1**. Click **Next** to continue.

Figure 245 Cisco UCS Director Express VM Network Configuration



10. In the Disk Format option, click the **Thick Provision Lazy Zeroed** radio button. Click **Next** to continue.
11. Review the details of the VM, Check the checkbox **Power On after deployment**.
12. Click **Finish** to proceed with deployment.

Figure 246 Deploy the Cisco UCS Director Express VM

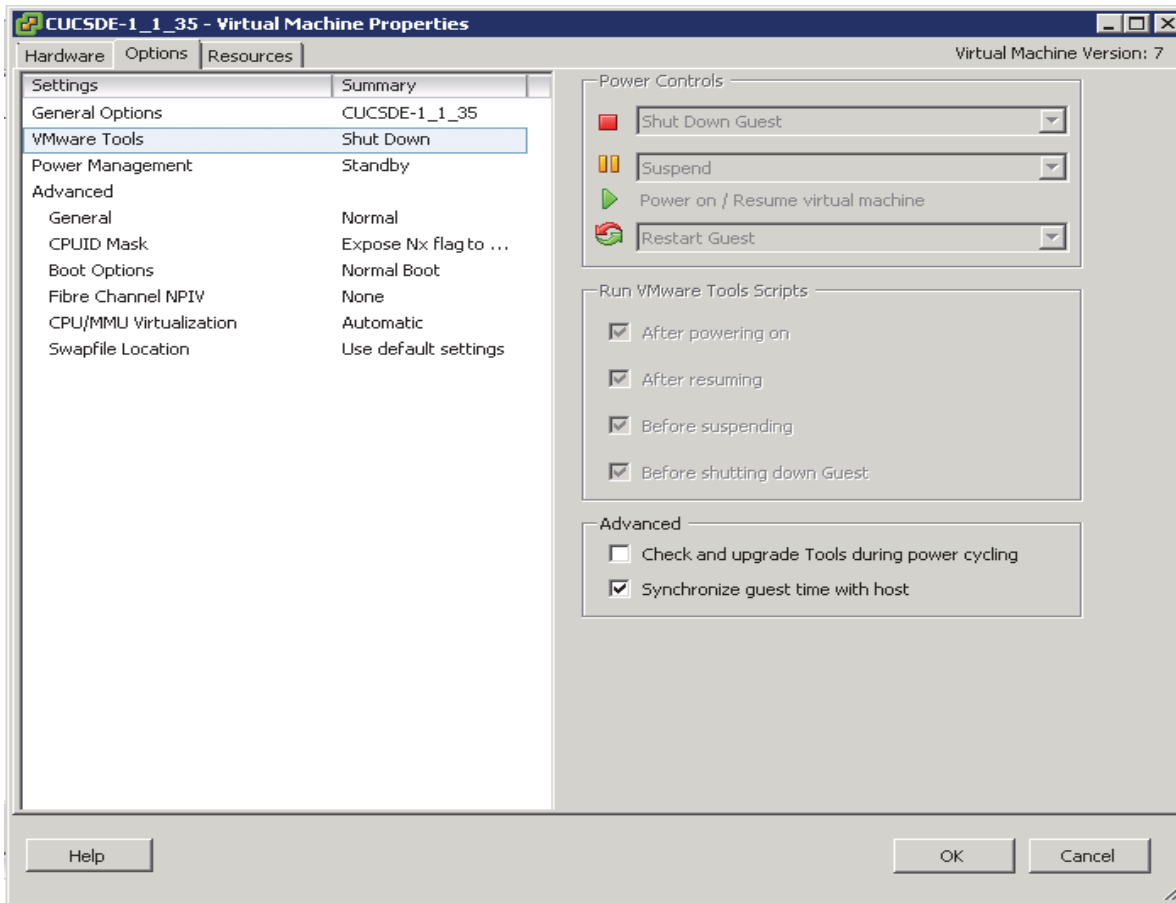


Configuring the Cisco UCS Director Express VM (UCSD-VM)

The Cisco UCS Director Express VM named as CUCSDE-1_1_35 shall be known as UCSD-VM here onwards.

1. Right click on the UCSD-VM, and select **Edit Settings**.
2. In the Virtual Machine Properties dialog box, click on the **Options** tab.
3. Click on the **VMware Tools**, Click on the **Synchronize guest time with host** option in the **Advanced** section.
4. Click on **OK** button to accept the changes.

Figure 247 Edit VM Settings to Synchronize the Guest Time with the ESXi Host



5. Right-click on the UCSD-VM and select **Open Console**.
6. Accept the End User License Agreement by typing **yes** and press the **ENTER**.
7. In the prompt to configure the static IP for the network interface, enter the IP address, Netmask and Gateway information.
8. Enter **y** to continue with the network configuration.

Figure 248 Assigning the Static IP Address to the UCSD-VM eth0

```

This script is executed on first boot only.
Configuring static IP configuration
Do you want to Configure static IP [y/n]? : y
Do you want to configure IPv4/IPv6 [v4/v6] ? : v4
Configuring static IP for appliance. Provide the necessary access credentials
IP Address: 10.29.160.35
Netmask: 255.255.255.0
Gateway: 10.29.160.1
Configuring Network with : IP(10.29.160.35), Netmask(255.255.255.0), Gateway(10.29.160.1)
Do you want to continue [y/n]? : y_
    
```


9. Configure the UCSD Express as the personality by entering the number 2.
10. At the prompt **Switching personality to UCSD Express, Are you sure to continue [y/n]?** Type **y** and hit **ENTER**.

Figure 249 Choose the UCSD Express Personality

```
Configuring Personality
  Select the personality

    1) Default - UCSD
    2) UCSD Express
    3) Cirrus

Personality : [1/2/3]? 2
Switching personality to UCSD Express. Are you sure to continue [y/n]? y_
```

11. The UCSD-VM goes through a personality change configuration as shown below.

Figure 250 UCSD-VM First-Boot Initializations

```
completed db privileges
copying my.cnf.template
Completed copying my.cnf.template
Forcing it to a login prompt
Completed forcing it to a login prompt
starting database
started database
sleep 1m
JRE Copy Start
JRE Copy End
Installing native files
Unzip of native files completed
Installing native (/usr/lib) files
Installed native (/usr/lib) files
Installing native (/usr/include) files
Installed native (/usr/include) files
Installing native (/usr/bin) files
Installed native (/usr/bin) files
Installing native (/etc) files
Installed native (/etc) files
Installing CUIC-vix files
Installed CUIC-vix files
JRE_HOME is
Wed Feb 18 09:31:47 UTC 2015 : Initializing CUIC Database schema
```



Note This step takes about 10-15 minutes to complete.

Applying the Upgrade Patches

1. Open a SSH/Putty session to the UCSD-VM.
2. Login as the user **shelladmin** with password **changeme**.

Figure 251 Logging onto the UCSD-VM Shell Administration Tool

```
login as: shelladmin
shelladmin@10.29.160.35's password: █
```

3. In the Shell Admin Menu, enter 3 to stop the services.
4. At the prompt, **Do you want to stop services [y/n]?** Type **y** to confirm and hit **ENTER** to continue.

Figure 252 Issuing the Command to Stop all the Services Via Shell Administration Tool.

```

Standalone Node

Select a number from the menu below

1) Change ShellAdmin Password
2) Display Services Status
3) Stop Services
4) Start Services
5) Stop Database
6) Start Database
7) Backup Database
8) Restore Database
9) Time Sync
10) Ping Hostname/IP Address
11) Show Version
12) Import CA Cert (JKS) File
13) Import CA Cert (PEM) File for VNC
14) Configure Network Interface
15) Display Network Details
16) Enable Database for Cisco UCS Director Baremetal Agent
17) Add Cisco UCS Director Baremetal Agent Hostname/IP
18) Tail Inframgr Logs
19) Apply Patch
20) Shutdown Appliance
21) Reboot Appliance
22) Manage Root Access
23) Login as Root
24) Configure Multi Node Setup (Advanced Deployment)
25) Clean-up Patch Files
26) Collect logs from a Node
27) Collect Diagnostics
28) Change Personality
29) Quit

SELECT> 3

Do you want to stop services [y/n]? : y █

```

5. In the Shell Admin menu, type 2 to view the status of the services. They all should be **NOT-RUNNING** as shown below.

Figure 253 Verifying the Status of the UCSD-VM Services

```

SELECT> 2

Service                Status                PID
-----                -
broker                 NOT-RUNNING          -
controller             NOT-RUNNING          -
eventmgr               NOT-RUNNING          -
client                 NOT-RUNNING          -
idaccessmgr           NOT-RUNNING          -
inframgr               NOT-RUNNING          -
TOMCAT                 NOT-RUNNING          -
websock                NOT-RUNNING          -

3467 ?                00:00:00 mysqld_safe
3888 ?                00:03:05 mysqld
Press return to continue ...

```

6. In the Shell Admin menu, type **19** and **ENTER** to start the patching process.
7. Type **n** to the prompt **Do you want to take database backup before applying patch[y/n]?**
8. At the prompt, Patch URL: enter **http://<BMA_IP>/patches/cucsd_patch_5_2_0_1.zip**
9. Hit **ENTER** to continue.

Figure 254 Cisco UCS Director 5.2.0.1 Patch Application Process

```

Select a number from the menu below

1) Change ShellAdmin Password
2) Display Services Status
3) Stop Services
4) Start Services
5) Stop Database
6) Start Database
7) Backup Database
8) Restore Database
9) Time Sync
10) Ping Hostname/IP Address
11) Show Version
12) Import CA Cert (JKS) File
13) Import CA Cert (PEM) File for VNC
14) Configure Network Interface
15) Display Network Details
16) Enable Database for Cisco UCS Director Baremetal Agent
17) Add Cisco UCS Director Baremetal Agent Hostname/IP
18) Tail Inframgr Logs
19) Apply Patch
20) Shutdown Appliance
21) Reboot Appliance
22) Manage Root Access
23) Login as Root
24) Configure Multi Node Setup (Advanced Deployment)
25) Clean-up Patch Files
26) Collect logs from a Node
27) Collect Diagnostics
28) Change Personality
29) Quit

SELECT> 19
Applying Patch...
Do you want to take database backup before applying patch[y/n]? n
User selected option not to take backup, proceeding with applying patch
Applying Patch:
Patch URL :http://10.29.160.36/patches/cucsd_patch_5_2_0_1.zip
Applying the Patch http://10.29.160.36/patches/cucsd_patch_5_2_0_1.zip [y/n]? y

```

This 5.2.0.1 patch that is being applied to the UCSD-VM's, upgrades all the core application software to the latest Cisco UCS Director's code base. After this step completes, the Big Data Upgrade package for release 1.1 needs to be applied.

10. In the Shell Admin menu, type **19** and **ENTER** to start the patching process.
11. Type **n** to the prompt **Do you want to take database backup before applying patch[y/n]?**
12. At the prompt, Patch URL:, enter **http://<BMA_IP>/patches/UCSDExpress_Big_Data_1.1_Upgrade_Package.zip**
13. Hit **ENTER** to continue.

Figure 255 Cisco UCS Director Express for Big Data 1.1 Upgrade Package Installation Process

```

1) Change ShellAdmin Password
2) Display Services Status
3) Stop Services
4) Start Services
5) Stop Database
6) Start Database
7) Backup Database
8) Restore Database
9) Time Sync
10) Ping Hostname/IP Address
11) Show Version
12) Import CA Cert (JKS) File
13) Import CA Cert(PEM) File for VNC
14) Configure Network Interface
15) Display Network Details
16) Enable Database for Cisco UCS Director Baremetal Agent
17) Add Cisco UCS Director Baremetal Agent Hostname/IP
18) Tail Inframgr Logs
19) Apply Patch
20) Shutdown Appliance
21) Reboot Appliance
22) Manage Root Access
23) Login as Root
24) Configure Multi Node Setup (Advanced Deployment)
25) Clean-up Patch Files
26) Collect logs from a Node
27) Collect Diagnostics
28) Change Personality
29) Quit

SELECT> 19
Applying Patch...
Do you want to take database backup before applying patch[y/n]? n
User selected option not to take backup, proceeding with applying patch
Applying Patch:
Patch URL :http://10.29.160.36/patches/UCSDExpress_Big_Data_1.1_Upgrade_Package.zip
ip

Applying the Patch http://10.29.160.36/patches/UCSDExpress_Big_Data_1.1_Upgrade_Pack
age.zip [y/n]? y

```

Figure 256 Cisco UCS Director Express for Big Data 1.1 Upgrade Package Application Complete

```

*****
Wed Jan 21 22:10:45 UTC 2015 : Copying ui.properties file
*****
Directory doesn't exit, continuing with installation process
*****
Wed Jan 21 22:10:45 UTC 2015 : Copying SSL File
*****
*****
Wed Jan 21 22:10:45 UTC 2015 : Copying VMWare Files & scalability folder
*****
Scalability folder exists, taking backup /opt/scalability-01-21-2015-22-10-45
Diagnostics folder exists, taking backup /opt/diagnostics-01-21-2015-22-10-45
*****
Wed Jan 21 22:10:45 UTC 2015 : Copying localization related files
*****
Japanese Directory exists.
TrueType folder is present
*****
Wed Jan 21 22:10:45 UTC 2015 : Copying sysmgr jar to T1 library locations if exist
*****
*****
Wed Jan 21 22:10:45 UTC 2015 : Personality specific changes for upgrade
*****
Personality details --> Product Name : UCSD Express for Big Data , Product Version :
1.0.0.0
Restored account-type-exclusion-list.properties for UCSD Express for Big Data
Restored DefaultRoleMenuMappings.properties for UCSD Express for Big Data
Restored RegularSet_menu.xml for UCSD Express for Big Data
Restored AdminSet_menu.xml for UCSD Express for Big Data
Restored feature-exclusion-list.properties for UCSD Express for Big Data
Restored reports.xml for UCSD Express for Big Data
Restored about.json for UCSD Express for Big Data
Restored signed-sku-mapping.xml for UCSD Express for Big Data
*****
Restart services and database for the changes to take effect
*****

INFO (FileUtil.java:958) *****
INFO (FileUtil.java:963)
INFO (FileUtil.java:967) 150121 22:10:45 [FileUtil] RunCommandThread: Completed thre
d:      Thread[Thread-1,5,main]

Completed installing package 0
*****
Press return to continue ...

```

14. After the successful application of the patch, type **4** and **ENTER** to start the services.



Note It takes about a few minutes for all the services to get started.

15. Type **2** to check on the services status. All the services should now be in **RUNNING** state.

Figure 257 Verify the Status of the Services in the UCSD-VM

```

SELECT> 2
Service          Status          PID
-----
broker           RUNNING        7756
controller       RUNNING        7888
eventmgr         RUNNING        7966
client           RUNNING        8025
idaccessmgr      RUNNING        8113
inframgr         RUNNING        8172
TOMCAT           RUNNING        8240
websock          RUNNING        8320

3467 ?           00:00:00 mysqld_safe
3888 ?           00:05:52 mysqld
Press return to continue ...

```



Note Even after all the services are in a RUNNING state, it would take an additional 3 to 5 minutes for the UCSD-VM client services to become available.

Configuring the Cisco UCS Director Express for Big Data (UCSD Express)

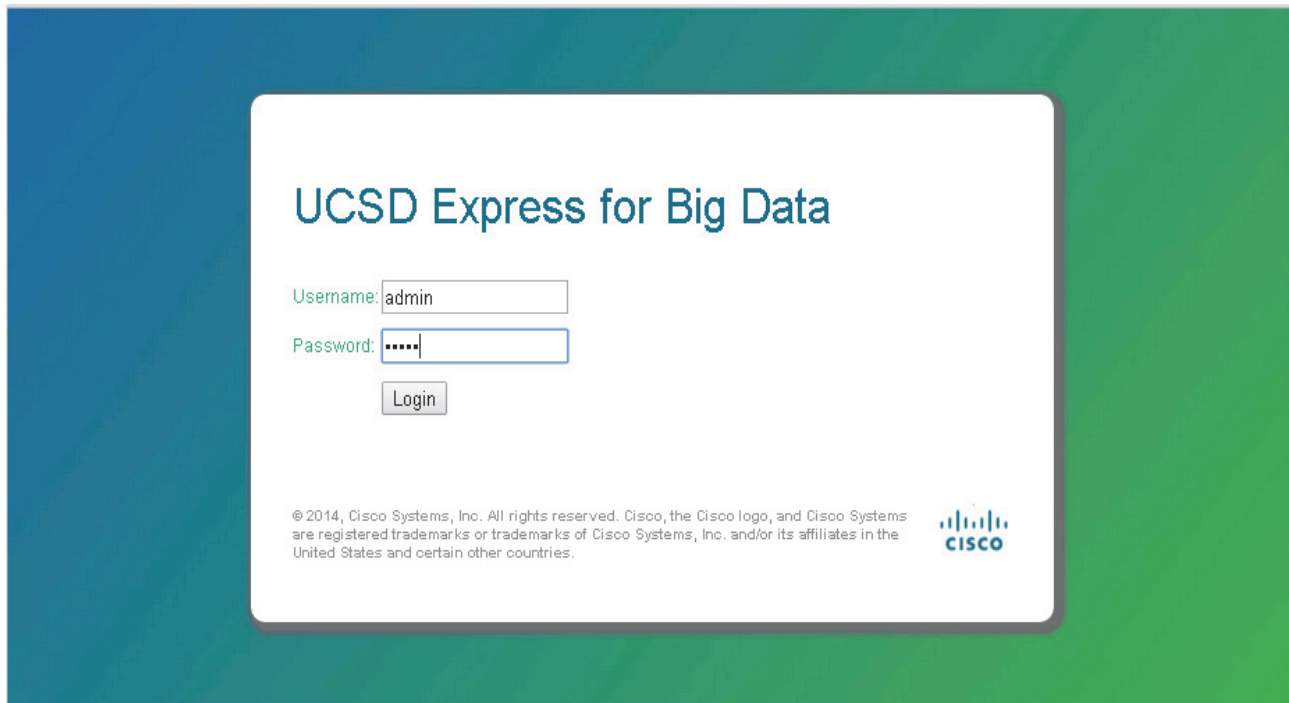
The Cisco UCS Director Express for Big Data, henceforth known as UCSD-Express, needs to be configured with the IP address to the UCS domain (i.e. UCS Manager's) physical account. This allows the UCSD-Express to query the UCS Manager and perform inventory collection.

The UCSD-Express will also need to be configured with the BMA's physical account and configure it's services such as DHCP.

Add the licenses to UCSD-Express

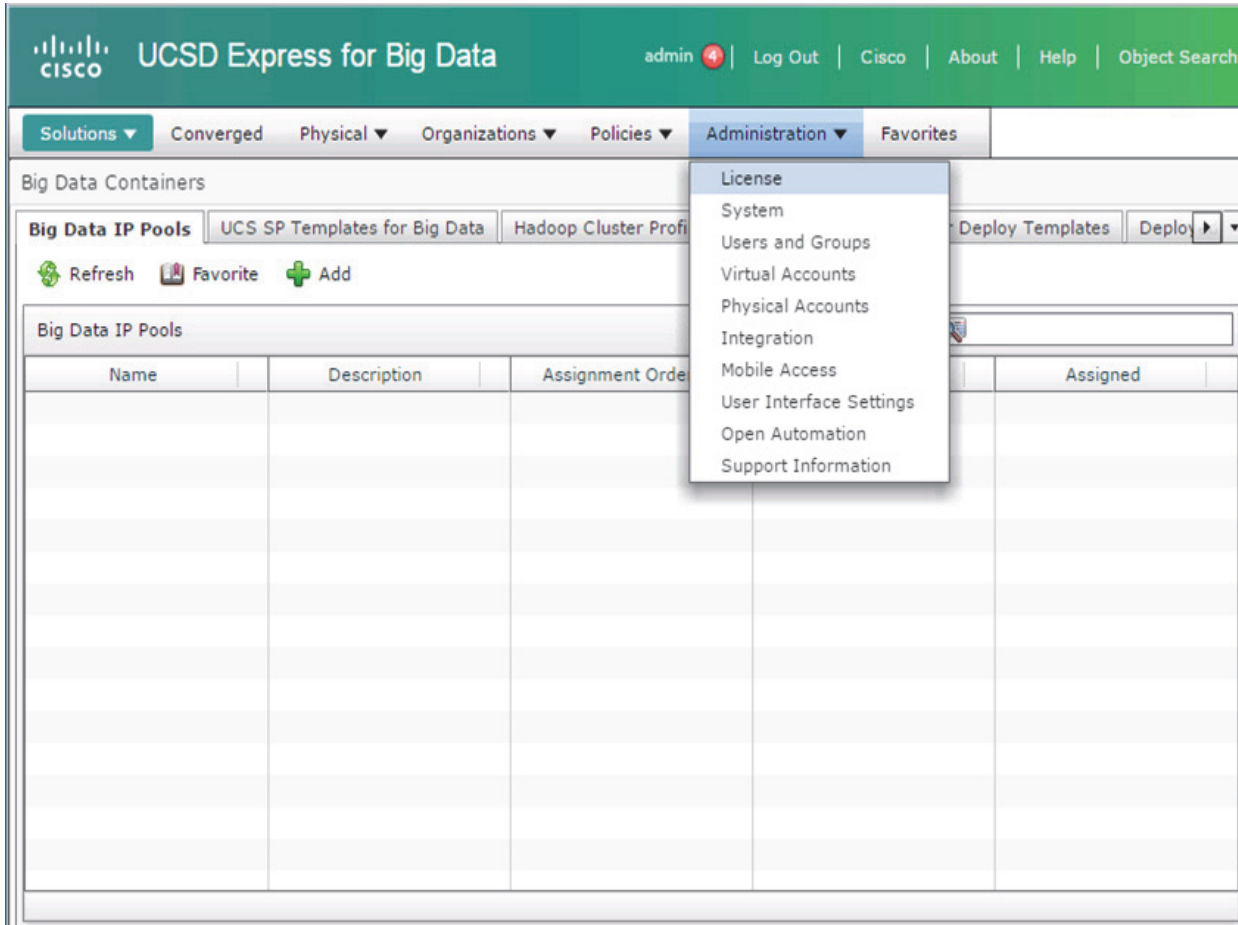
1. Using a web browser, visit the URL <http://<UCSD-VM's IP>/>.
2. Login as user **admin** with the default password **admin**.

Figure 258 Logging onto the Cisco UCS Director Express for Big Data



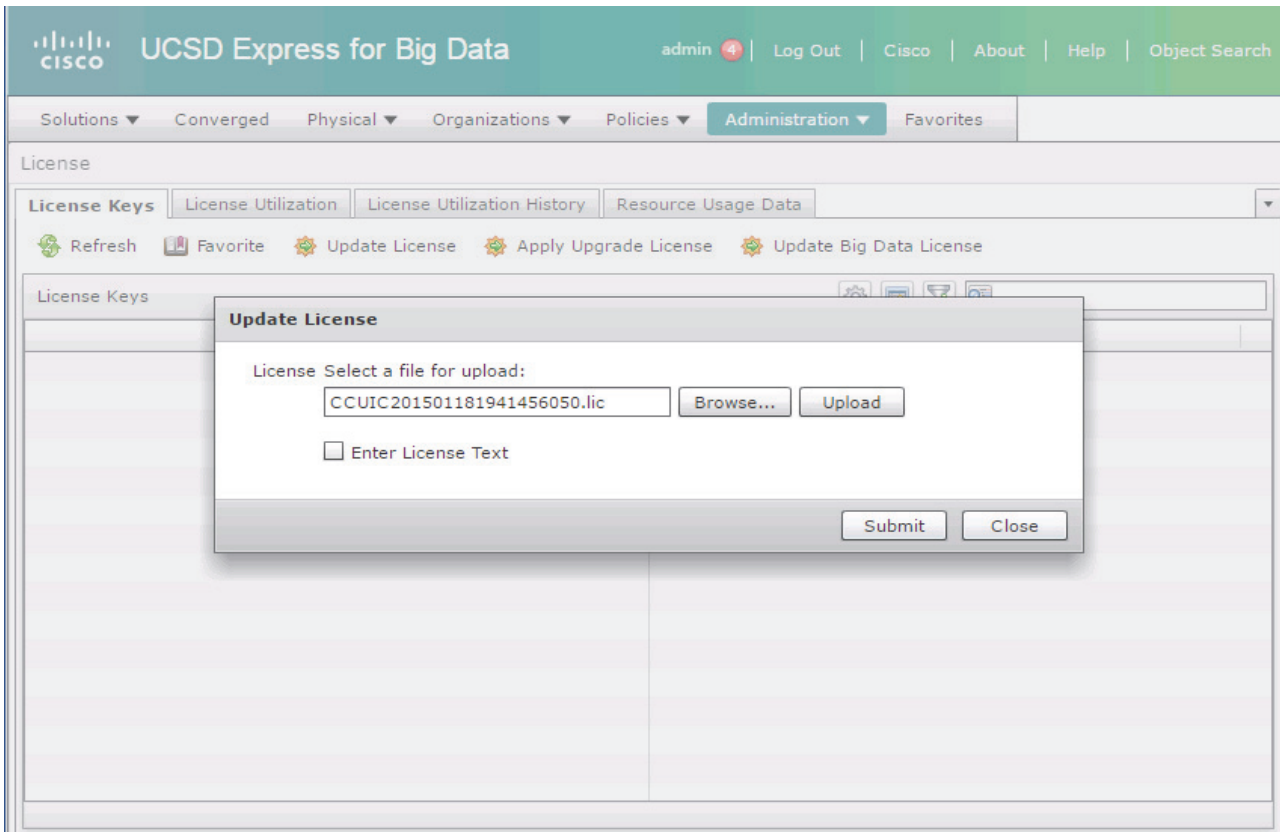
3. Navigate to **Administration > License** screen.

Figure 259 Accessing the License Administration Page



4. Click on **License Keys** tab.
5. Click on **Update License**.
6. In the **Update License** dialog box, click **Browse** to select the license file.
7. Click **Upload**.
8. After the license file gets uploaded, Click **Submit** to apply the license.

Figure 260 Applying the Base Cisco UCS Director License.



9. The license keys are displayed as shown below.

Figure 262 Applying the Cisco UCS Director Express Big Data Subscription License

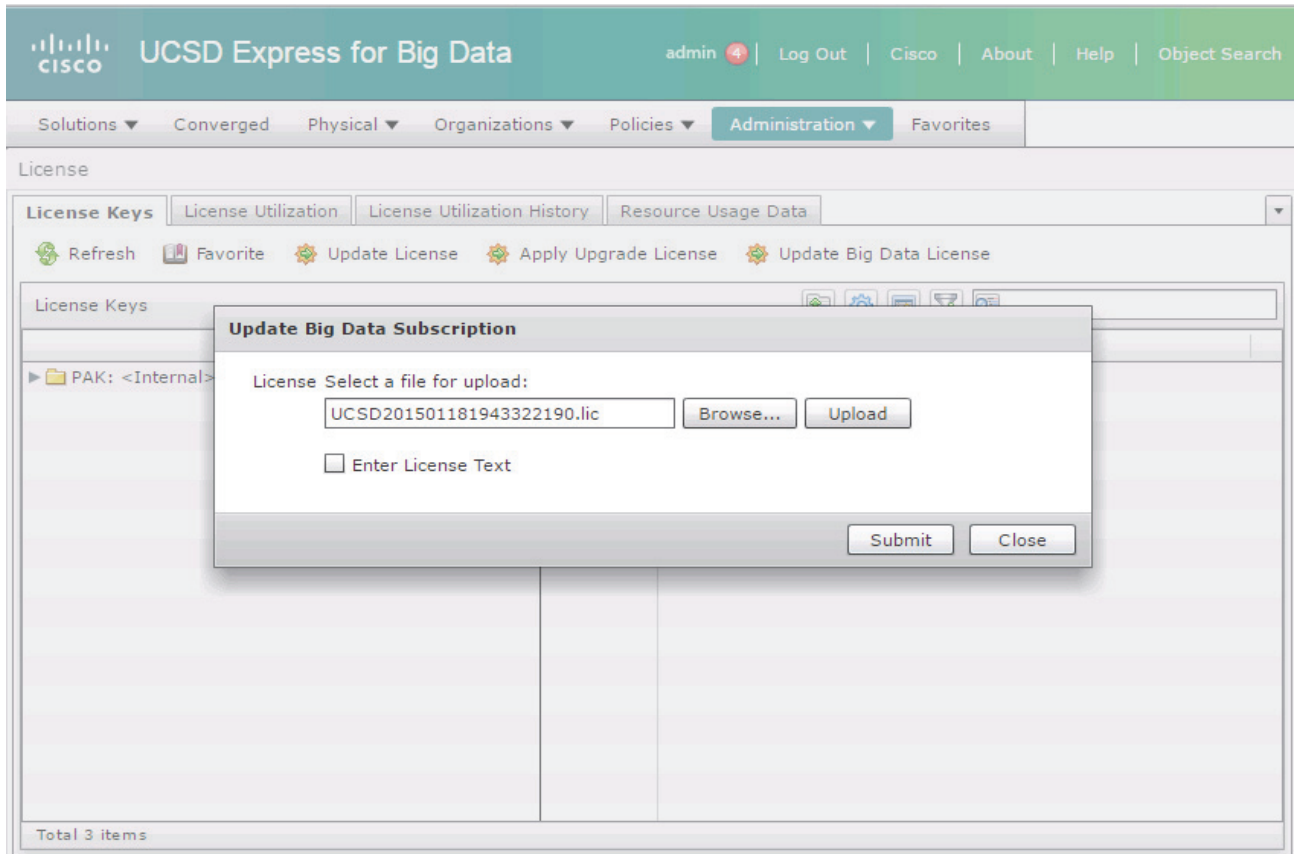


Figure 263 Completion of the License Application.

The screenshot shows the Cisco UCS Director Express for Big Data web console. The top navigation bar includes the Cisco logo, the product name, and user information (admin, Log Out, Cisco, About, Help, Object Search). Below the navigation bar, there are tabs for Solutions, Converged, Physical, Organizations, Policies, Administration (selected), and Favorites. The main content area is titled "License" and contains several sub-tabs: License Keys (selected), License Utilization, License Utilization History, and Resource Usage Data. Under the License Keys tab, there are action buttons: Refresh, Favorite, Update License, Apply Upgrade License, and Update Big Data License. Below these buttons is a table with the following data:

License Entry	License Value/Status
PAK: <Internal> (#20150118194332219 - 2)	
Expiration Date	March 18, 2015
License ID	PAK: <Internal> (#20150118194332219 - 2)
CUIIC-EBDS	1
CUIIC-EBDS	1
PAK: <Internal> (#20150118194145605 - 1)	
Expiration Date	March 19, 2015
License ID	PAK: <Internal> (#20150118194145605 - 1)
CUIIC-BASE-K9	1

Total 7 items

Add the UCS Manager physical account to the UCSD-Express

1. In the UCSD-Express web console, navigate to **Administration >Physical Accounts**.
2. Click + **ADD** button
 - a. Input the UCS Manager Account details as follows.
 - b. In the Account Name field, enter a name to this UCS Manager account.
 - c. In the Server Address field, enter the IP address of the UCS Manager.
 - d. In the User ID field, enter admin.
 - e. In the Password field, enter the password to the UCS Manager's admin user.
 - f. In the Transport Type field, choose https.
3. Click **Add**.

Figure 264 Adding the UCS Manager as a Physical Account in the UCSD-VM

The screenshot shows the 'Add Account' dialog box in the Cisco UCS Director Express for Big Data interface. The dialog is open over the 'Physical Accounts' section. The fields are: Pod (Default Pod), Category (Computing), Account Type (UCSM), Authentication Type (Locally Authenticated), Server Management (All Servers), Account Name (UCSM40), Server Address (10.29.160.40), Use Credential Policy (unchecked), User ID (admin), Password (masked with asterisks), Transport Type (https), Port (443), and Description (empty). The 'Add' and 'Close' buttons are at the bottom right of the dialog.

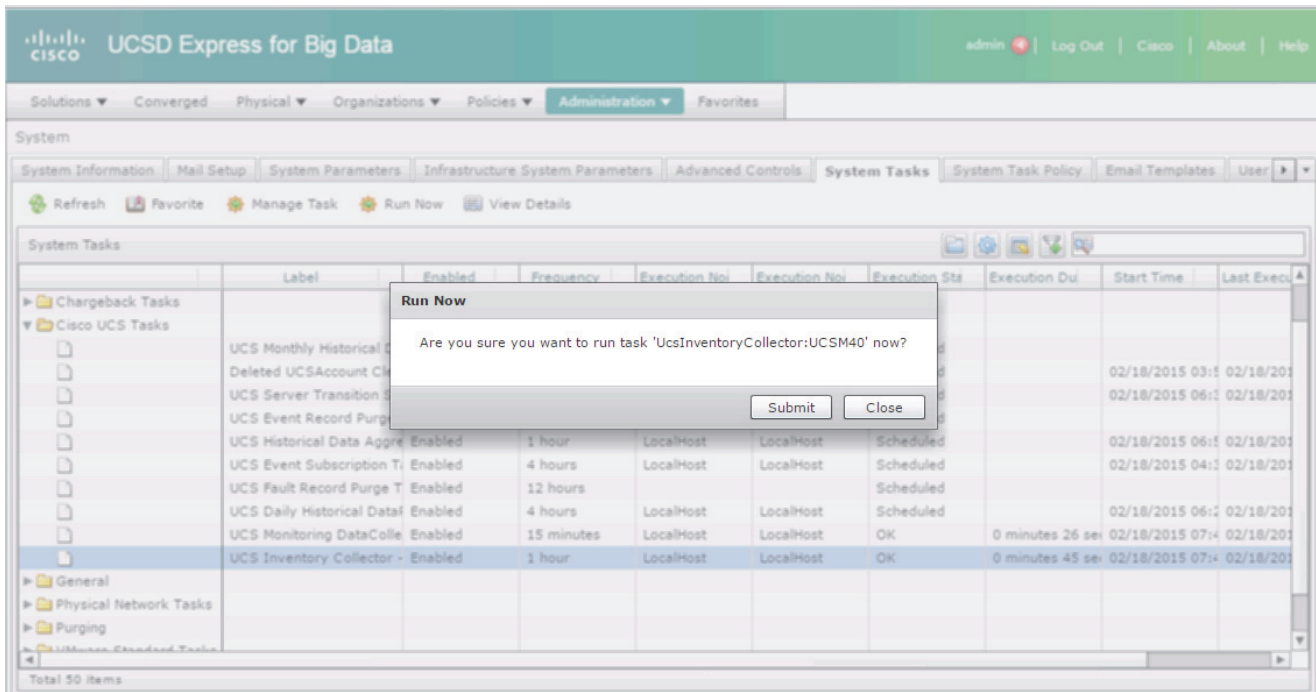


Note

After adding a physical account, the UCSD-Express will query the UCS Manager to perform the inventory collection. This process of inventory collection happens at scheduled intervals. Optionally, you may kick start the inventory collection process manually. These optional steps are described in the steps 4 to 8 below.

4. Goto **Administration > System**.
5. Click on **System Tasks** tab.
6. Open the folder Cisco UCS Tasks.
7. Click on UCS Inventory Collector Task.
8. Click **Run Now** button to execute the task.

Figure 265 Start the UCS Inventory Collection System Task



Add the Bare Metal Agent physical account to the UCSD-Express

1. In the UCSD-Express web console, navigate to **Administration > Physical Accounts**.
2. Click on **Bare Metal Agents** tab; Click **+ Add**.
3. Enter the BMA physical account information details as follows:
4. In **BMA Name** field, enter a name to this BMA physical account.
5. In the **BMA Management Address** field, enter the BMA-VM's IP address assigned to **NIC eth0**.
6. In the **Login ID** field, enter **root**.
7. In the **Password** field, enter the password. Default password is **pxeboot**.
8. Check the checkbox **BMA Uses Different Interfaces for Management and PXE Traffic**.
9. In the **BMA PXE Interface Address** field, enter PXE IP address i.e. BMA-VM's IP address assigned to **NIC eth1**.
10. Click **Submit**.

Figure 266 Adding the Bare Metal Agent Appliance Information

Add Bare Metal Agent Appliance

BMA Name *

BMA Management Address *

NOTE: This address must be reachable from the Cisco UCS Director appliance

Login ID *

Password *

BMA Uses Different Interfaces for Management and PXE Traffic

BMA PXE Interface Address *

Description

Location

UCSD Database Address *

Configure the Bare Metal Agent's DHCP services

1. Navigate to **Administration > Physical Accounts > Bare Metal Agents**.
2. Select the **BMA** entry.
3. On the menu items row, click on the downward facing arrow located at the far right.
4. Select **Configure DHCP**.

Figure 267 Configuring the DHCP

The screenshot shows the Cisco UCS Director Express for Big Data interface. The top navigation bar includes 'admin', 'Log Out', 'Cisco', 'About', 'Help', and 'Object Search'. The main menu has 'Administration' selected. The 'Physical Accounts' section is active, and the 'Bare Metal Agents' tab is selected. A table lists Bare Metal Agents, with one entry 'BMA36' visible. A context menu is open over this entry, showing options like 'Refresh', 'Favorite', 'Add', 'Edit', 'View Details', 'Delete', 'Start Services', 'Stop Services', 'Service Status', 'Set Default BMA', 'View DHCP Conf...', 'View BMA logs', and 'View DHCP Configuration' (highlighted).

BMA Name	BMA Management Address	PXE Server Address	Reachable	Location	Description	Default BMA	Status	Last Update
BMA36	10.29.160.36	192.168.100.36	YES			Yes	Active	03/21

Total 1 items

5. In the **Configure DHCP** dialog box, enter the following
6. In the **DHCP Subnet** field, enter the subnet that's associated with the BMA-VM's **eth1** NIC.
7. In the **DHCP Netmask**, enter the appropriate subnet mask value for this network.
8. In the **DHCP Start IP**, enter a starting IP address in the same subnet.
9. In the **DHCP End IP**, enter a starting IP address in the same subnet.
10. In the **Router IP Address**, enter the IP address of the gateway router in the network if available, if not may be left as blank or input the IP address of the BMA-VM's **eth1** NIC.
11. Click **Submit**.

Figure 268 Configuring the DHCP services on the BMA.

Configure DHCP

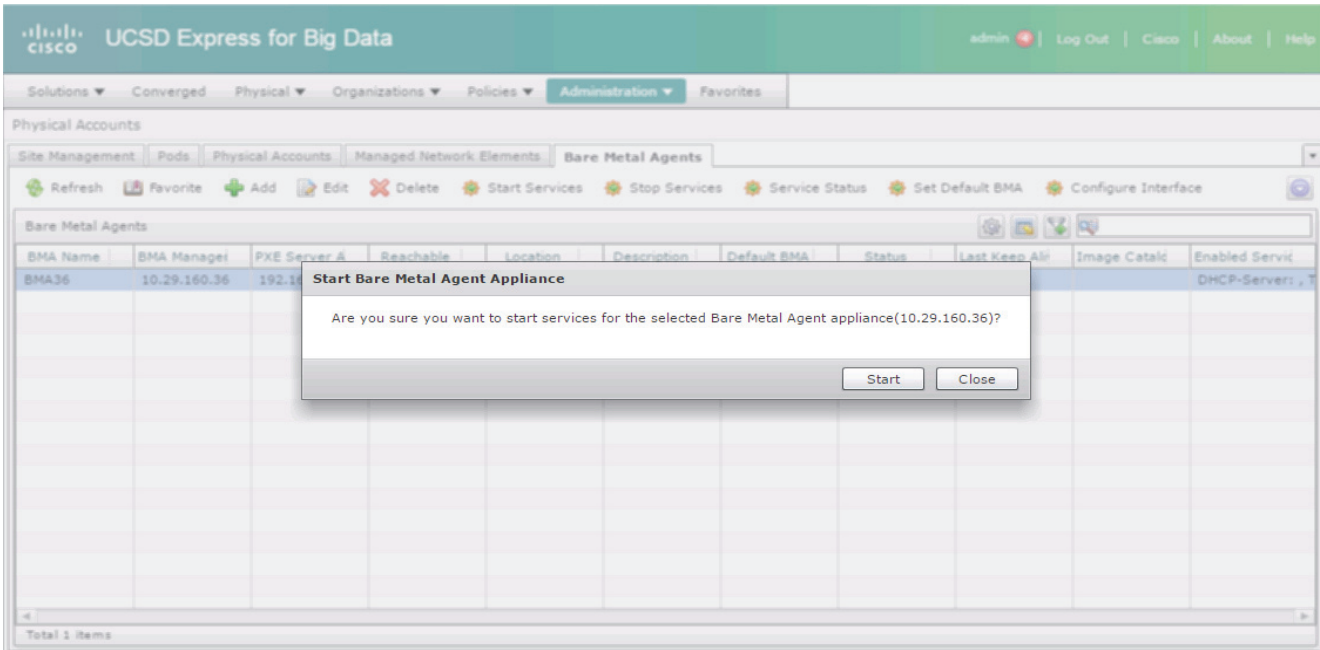
DHCP Subnet	<input type="text" value="192.168.85.0"/>	*
DHCP Netmask	<input type="text" value="255.255.255.0"/>	*
DHCP Start IP	<input type="text" value="192.168.85.160"/>	*
DHCP End IP	<input type="text" value="192.168.85.254"/>	*
Router IP Address	<input type="text" value="192.168.85.36"/>	

Submit Close

Start the BMA services

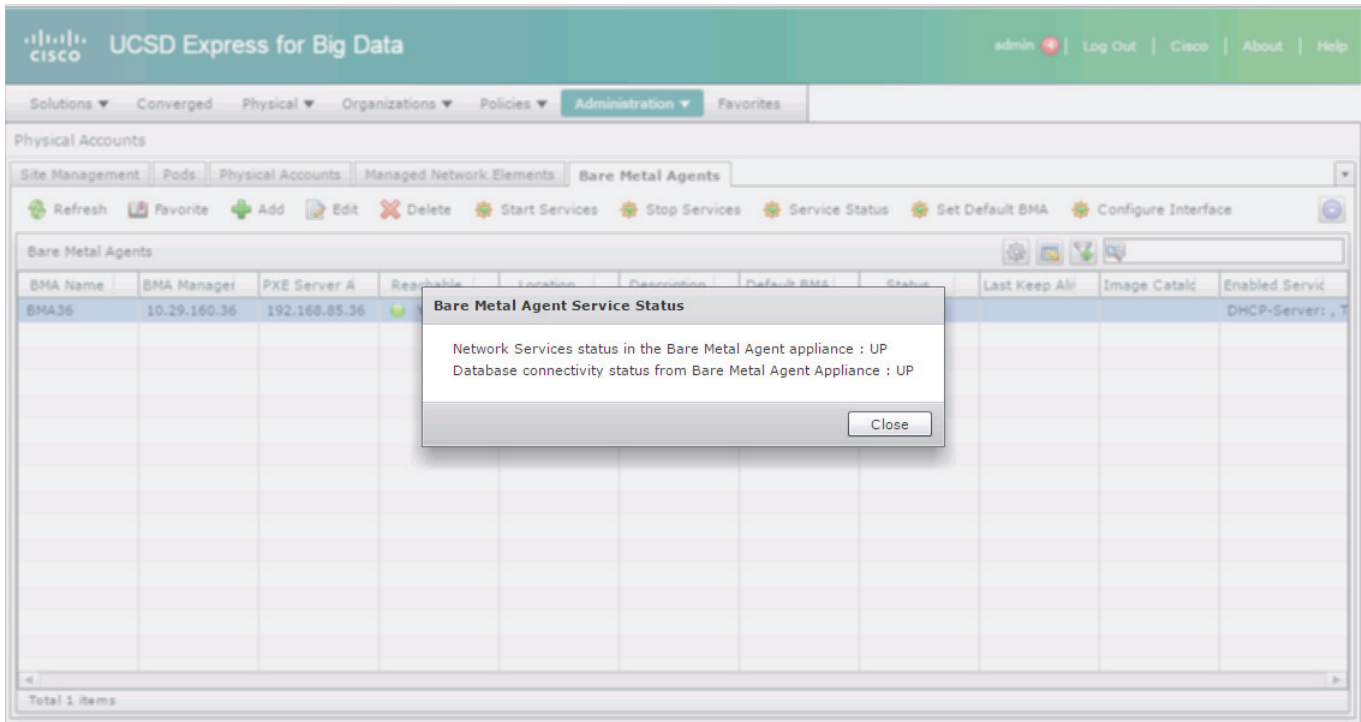
1. Navigate to **Administration >Physical Accounts >Bare Metal Agents**.
2. Select the BMA entry.
3. Click **Start Services**.
4. In the **Start Bare Metal Agent Appliance** dialog box, click **Start** to start the services.

Figure 269 Starting the BMA Services



5. Click on **Service Status**, to **check the status of the services**.
6. The Bare Metal Agent Service Status **message box should display both the Network Services status and Database connectivity status as UP**.

Figure 270 Verifying the Bare Metal Agent Services Status

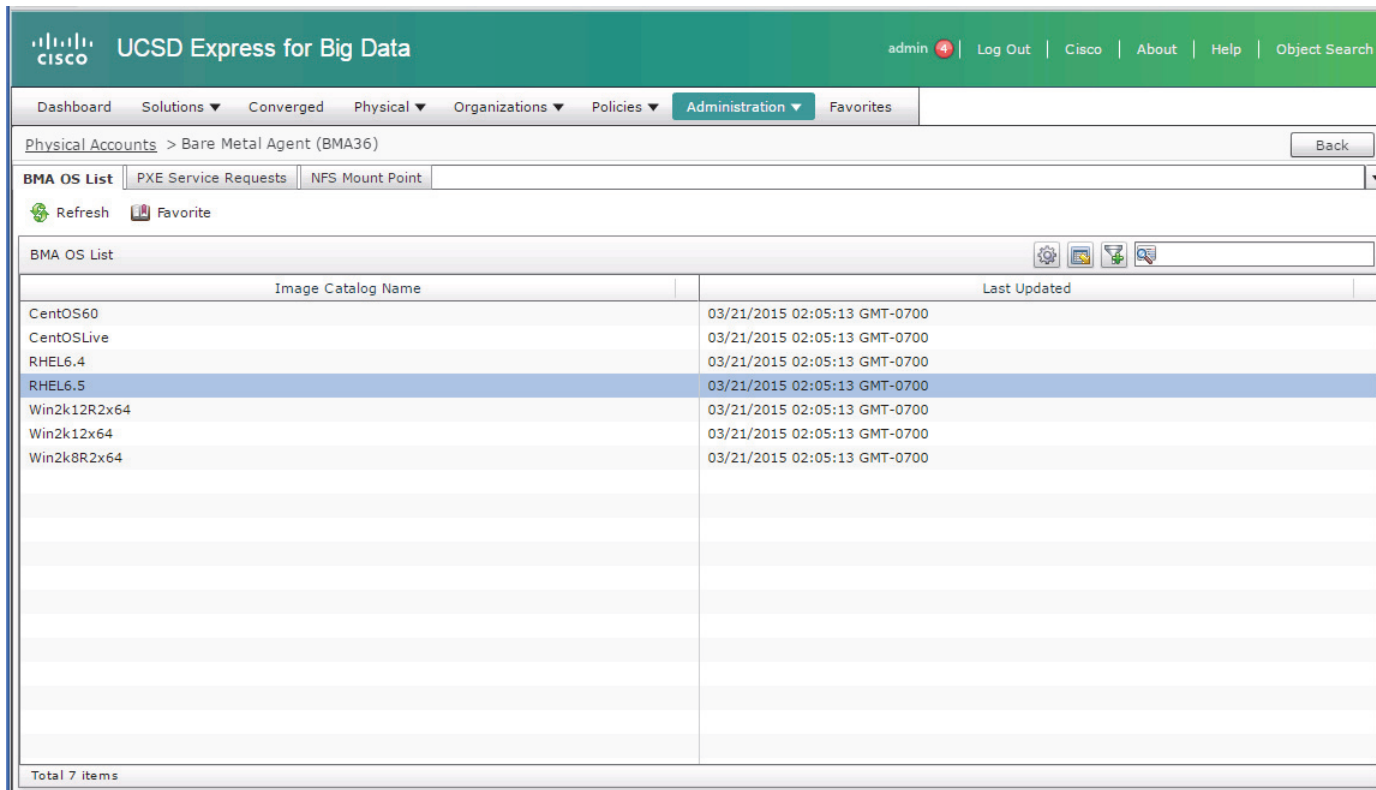


Note

It may take a little while for the service status and on the BMA entry to get updated. The UCSD-Express and the associated BMA parts are now ready.

7. Double click on the BMA entry to verify the RHEL operating system repository.

Figure 271 Verifying the RHEL Operating System Software



Note

BMA-VM software periodically scan the /opt/cnsaroot directory to update the available list of operating system software repositories.

Creating the Hadoop Cluster using UCSD-Express

For creating a Hadoop cluster of a desired distribution, the UCS Manager that’s managing the target servers must be pre-configured to meet the following requirements. For performing these configurations, refer to any Cisco UCS Integrated Infrastructure for Big Data Cisco Validated Designs found at http://www.cisco.com/go/bigdata_design

- a. The uplink ports fabric Interconnects must be reachable to that the UCSD-Express appliances management network (i.e. eth0).
- b. The UCS-Manager must be configured with a host firmware policy containing C-series rack mount server firmware packages.
- c. UCS Manager must be configured to discover the Rack Servers in its domain, and the respective ports are configured as server ports.
- d. The server pool must be configured with appropriate set of physical servers that are part of the UCS domain.
- e. The QOS System Classes Platinum and Best Effort must be configured and enabled.



Note C240/C220 M4 Rack Servers are supported from UCS firmware 2.2(3d) onwards.

Create the IP Address pools

1. Using a web browser, visit the URL **http://<UCSD-VM's IP>/**.
2. Login as user **admin** with the default password **admin**.
3. Navigate to **Solutions > Big Data Containers**.
4. Click on the **Big Data IP Pools** Tab.
5. Click on **+ Add**.

Figure 272 Creating the IP Address Pools

The screenshot shows the UCSD Express for Big Data web interface. The top navigation bar is green and contains the Cisco logo, the text "UCSD Express for Big Data", and user information: "admin" with a red lock icon, "Log Out", "Cisco", "About", "Help", and "Object Search". Below the navigation bar is a menu with "Dashboard", "Solutions" (selected), "Converged", "Physical", "Organizations", "Policies", "Administration", and "Favorites". The main content area is titled "Big Data Containers" and has several tabs: "Big Data IP Pools" (selected), "UCS SP Templates for Big Data", "Hadoop Cluster Profile Templates", "Hadoop Cluster Deploy Templates", and "Deployed Clusters". Below the tabs are "Refresh", "Favorite", and "+ Add" buttons. The "Big Data IP Pools" section features a table with the following columns: "Name", "Description", "Assignment C", "Size", and "Assigned". The table is currently empty. At the bottom left of the table area, it says "Total 0 item".

6. In the **Create an IP Pool** dialog box.
7. Enter the name **MGMT**. Click **Next** to continue.

Figure 273 *Creating the IP Address pool for MGMT VLAN*

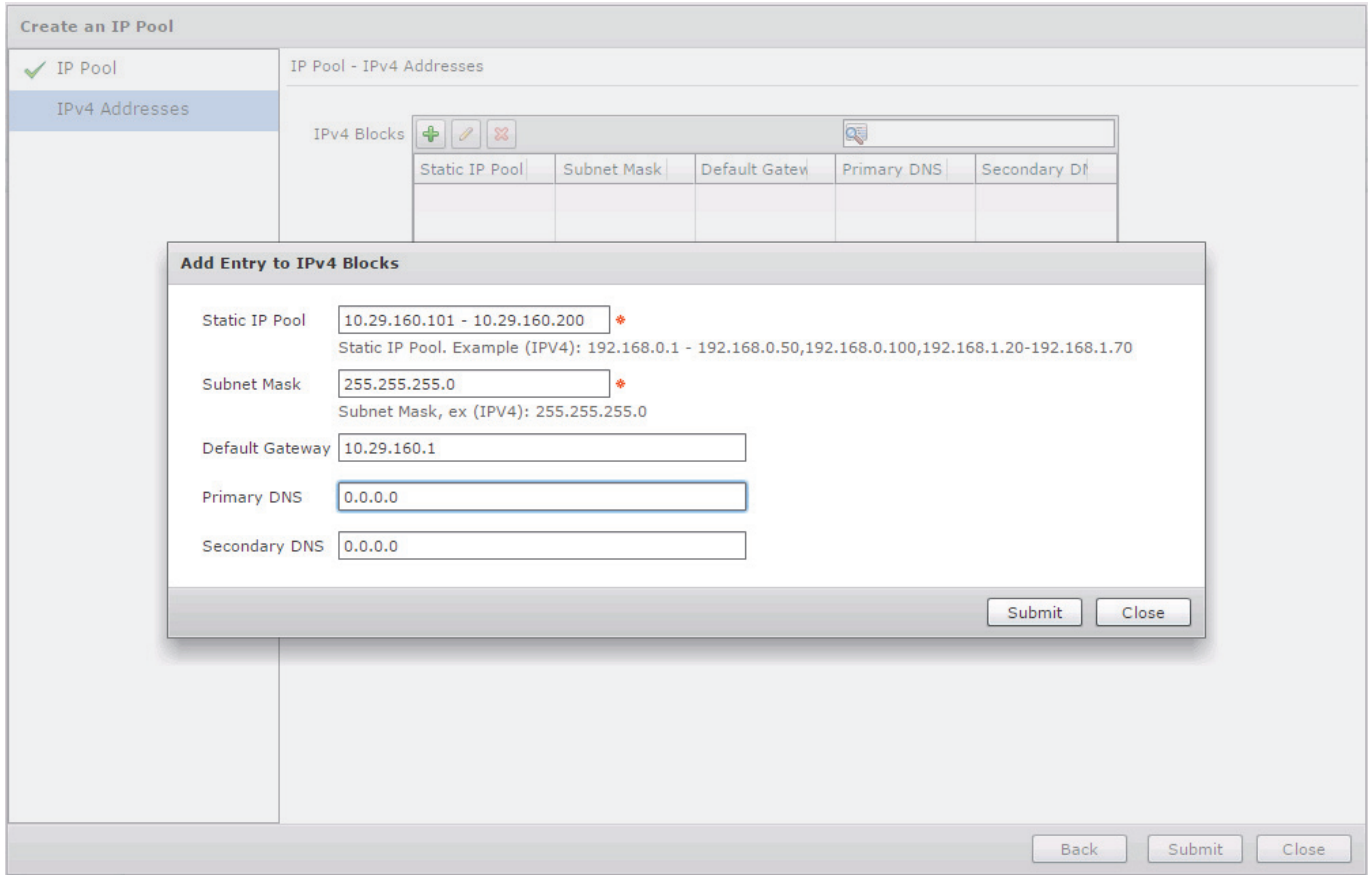
The screenshot shows a web-based configuration interface for creating an IP pool. The title bar reads "Create an IP Pool". On the left, a sidebar contains "IP Pool" (selected) and "IPv4 Addresses". The main content area is titled "IP Pool Management" and includes the following fields:

- IP Pool Name:** A text input field containing "MGMT" with a red asterisk indicating a required field.
- Description:** An empty text input field.
- Assignment Order:** A dropdown menu currently set to "Default".

At the bottom right of the interface, there are two buttons: "Next" and "Close".

8. In the IPv4 Blocks table, click on +.
9. In the Add Entry to IPv4 Blocks dialog box, enter the following.
 - In the Static IP Pool field, enter the Static IP Address pool range in the format A.B.C.X – A.B.C.Y.
 - In the Subnet Mask field, enter the appropriate subnet mask.
 - In the Default Gateway field, enter the IP address of the Gateway if present.
 - In the Primary DNS field, enter the IP address of the DNS server.
10. Click **Submit**.

Figure 274 Adding a Block of IP Address to the MGMT IP Address Pool

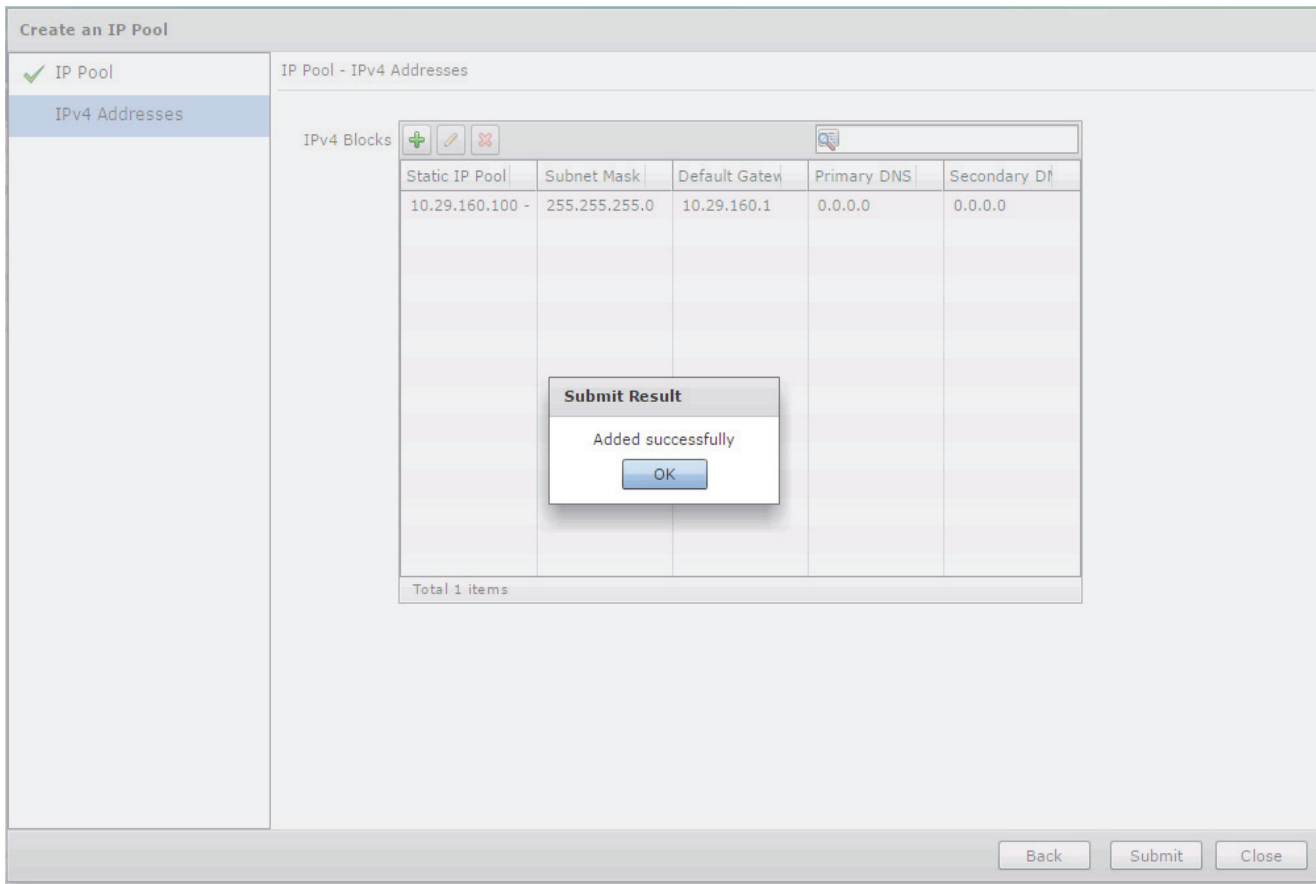


Note

The Default Gateway, Primary and Secondary DNS fields are optional.

11. Click **Submit** again to create the Big Data IP Pool.

Figure 275 *IP Address Pool Added Successfully*



Repeat this process for two more interfaces, by creating an IP address pool by name HDFS for Hadoop configurations to be associated with vNIC eth1, and an IP address pool by name DATA to be associated with vNIC eth2 in the service profiles. Please refer to “Configuring VLAN Section” above in Cisco UCS Integrated Infrastructure for Big Data CVDs.

The following figure shows the UCSD-Express that is fully provisioned all the necessary Big Data IP address Pools.

Figure 276 All the IP Address Pools have been Configured Successfully

The screenshot shows the UCSD Express for Big Data web interface. The main content area is titled "Big Data Containers" and contains a sub-section for "Big Data IP Pools". The table below shows the configuration for three IP pools: MGMT, HDFS, and DATA. Each pool has an assignment order of "default", a size of 100, and 0 nodes assigned.

Name	Description	Assignment Order	Size	Assigned
MGMT		default	100	0
HDFS		default	100	0
DATA		default	100	0

Total 3 items

Creating a Hadoop Cluster

1. Using a web browser, visit the URL **http://<UCSD-VM's IP>/**.
2. Login as user **admin** with the default password **admin**.
3. Navigate to **Solutions >Big Data Containers**.
4. Click on the **Hadoop Cluster Deploy Templates** Tab.
5. Click on **Create Instant Hadoop Cluster**.
6. In the Instant Hadoop Cluster Creation dialog box, enter the following.
7. In Big Data Account Name field, enter a preferred name.
8. In the UCS Manager Policy Name Prefix field, enter a prefix that is less than equal to 5 letters long.
9. In the Hadoop Cluster Name field, enter a preferred name of the cluster – this will be the name assigned to the Hadoop cluster within the context of selected Hadoop Manager.
10. In the Hadoop Node Count filed, enter the desired number of nodes.

The minimum number of nodes allowed for Cloudera and Hortonworks Hadoop cluster is 4 and for MapR cluster it is 3.


Note

There should be sufficient number of servers available in the server pool.

11. In the password fields, enter the preferred passwords and confirm them.
12. Choose the OS Version from the drop-down box. For C220 M4/C240 M4 rack servers, only OS supported is RHEL 6.5.


Note

At the time of this writing, RHEL6.5 is the only OS that is supported on C220 M4/C240 M4 rack servers.

13. In the Hadoop Distribution field, select **Hortonworks** from the drop-down list.
14. In the Hadoop Distribution Version field, select **Hortonworks-2.2** from the drop-down list.

Figure 277 *Selecting the Hadoop Distribution Version*

Cloudera	<p>Hadoop Distribution: cloudera *</p> <p>Hadoop Distribution Version: cloudera-5.2.0, cloudera-5.2.0, cloudera-5.3.0, cloudera-5.0.1, cloudera-5.2.1 *</p> <p>UCS Manager Account:</p>
Hortonworks	<p>Hadoop Distribution: Hortonworks *</p> <p>Hadoop Distribution Version: Hortonworks-2.1, Hortonworks-2.1, Hortonworks-2.2 *</p> <p>UCS Manager Account:</p>
MapR	<p>Hadoop Distribution: MapR *</p> <p>Hadoop Distribution Version: MapR-4.0.1, MapR-4.0.1, MapR-3.1.1, MapR-4.0.2 *</p> <p>UCS Manager Account:</p>

15. In the UCS Manager Account, select the appropriate UCS-Manager account.
16. Select the organization.
17. vNIC Template Entry
18. Double-click on row eth0 and select appropriate Mgmt IP-pool, MAC Address Pool and enter the MGMT VLAN id. Click Submit.

Figure 278 *Editing the vNIC Template to Provide the MGMT Network Configurations*

Edit Entry

vNIC Name: eth0 *

IP Pool: MGMT(10.29.160.101 - 10.29.160.200) *

MAC Address Pool: mac_pool1 (1978) *

VLAN ID: 1 *
[4048-4093],[1-3967]
(MGMT VLAN)

Submit Close

19. Double-click on **eth1** and select appropriate IP-pool, MAC Address Pool and enter the DATA1 VLAN ID. Click **Submit**.

Figure 279 *Editing the vNIC Template to Provide the DATA1 Network Configurations*

Edit Entry

vNIC Name: eth1 *

IP Pool: HDFS(192.168.11.101 - 192.168.11.200) *

MAC Address Pool: mac_pool1 (1978) *

VLAN ID: 11 *
[4048-4093],[1-3967]
(DATA1 VLAN)

Submit Close

20. Double-click on **eth2** and select appropriate IP-pool, MAC Address Pool and enter the DATA VLAN ID. Click **Submit**.

Figure 280 *Editing the vNIC Template to Provide the DATA2 Network Configurations*

The screenshot shows a dialog box titled "Edit Entry" with the following configuration fields:

- vNIC Name:** eth2
- IP Pool:** DATA(192.168.12.101 - 192.168.12.200)
- MAC Address Pool:** mac_pool1 (1978)
- VLAN ID:** 12

Below the VLAN ID field, there is additional information: [4048-4093],[1-3967] and (DATA2 VLAN). At the bottom right of the dialog box are "Submit" and "Close" buttons.



Note

The following figure show the expanded version of the Instant Hadoop Cluster Creation dialog box with all the fields filed in.

Figure 281 Creating an Instant Hortonworks Hadoop Cluster

Instant Hadoop Cluster Creation

Big Data Account Name: *

Account name can have atmost 10 alphanumeric characters

UCSM Policy Name Prefix: *

UCSM Policy Name Prefix can have atmost 5 characters

Hadoop Cluster Name: *

Hadoop Node Count: *

SSH (root) Password: *

Confirm SSH Password: *

Hadoop Manager Password: *

Confirm Hadoop Manager Password: *

Host Node Prefix: *

OS Version: *

Choose RHEL 6.5 for M4 Servers

Hadoop Distribution: *

Hadoop Distribution Version: *

UCS Manager Account: *

Organization: *

Server UUID pool: *

PXE VLAN ID: *

[4048-4093],[1-3967]

Server Pool:

	ID	Server Pool	Server Pool f	Assigned	Size
<input checked="" type="checkbox"/>	UCSM40;org-roo	M4_servers		8	13

Total 1 items

Host Firmware Package:

	Account Nam	Organization	Name	DN	Mode
<input type="checkbox"/>	UCSM40	root	default	org-root/fw-host	staged
<input checked="" type="checkbox"/>	UCSM40	root	C_series_FW	org-root/fw-host	staged
<input type="checkbox"/>	UCSM40	root	ESXi_FW_Packa	org-root/fw-host	staged

vNIC Template:

vNIC Name	IP Pool	First MAC Address	VLAN ID
eth0	MGMT:10.29.160.1	00:25:B5:00:00:00	1
eth1	HDFS:0.0.0.0	00:25:B5:00:00:00	11
eth2	DATA:0.0.0.0	00:25:B5:00:00:00	12

Total 3 items

21. Click **Submit**.

Monitoring the Hadoop Cluster Creation

1. In the UCSD-Express web console, navigate to Organization ? Service Requests.
2. Browse through the workflows. There are 3 types of workflows executed.
 - There would be one Master Workflows i.e. UCS CPA Multi-UCS Manager Hadoop cluster WF, per the Hadoop cluster creation request. Master workflow kick starts one or more UCS Manager-specific workflows. Besides that, this master workflow is responsible for Hadoop cluster provisioning.
 - UCS Manager specific workflows i.e. Single UCS Manager Server Configuration WF, would in turn kick start one or more UCS CPA Node Baremetal workflows.
 - UCS CPA Baremetal workflows provision the UCS service profiles and perform OS installation and custom configuration per node.

Figure 282 List of Workflows Recently Complete

Service Request ID	Request Type	Initiating User	Catalog/Workflow Name	Initiator C	Request Time	Request Status
348	Advanced	admin	UCS CPA Node BareMetal		03/17/2015 23:38:05 GMT-07	Complete
347	Advanced	admin	UCS CPA Node BareMetal		03/17/2015 23:38:05 GMT-07	Complete
346	Advanced	admin	UCS CPA Node BareMetal		03/17/2015 23:38:05 GMT-07	Complete
345	Advanced	admin	UCS CPA Node BareMetal		03/17/2015 23:38:04 GMT-07	Complete
344	Advanced	admin	Single UCSM Server Configuration WF		03/17/2015 23:36:19 GMT-07	Complete
343	Advanced	admin	UCS CPA Multi-UCSM Hadoop Cluster WF		03/17/2015 23:35:24 GMT-07	Complete
342	Advanced	admin	UCS CPA Node BareMetal		03/17/2015 14:31:27 GMT-07	Complete
341	Advanced	admin	UCS CPA Node BareMetal		03/17/2015 14:31:27 GMT-07	Complete
340	Advanced	admin	UCS CPA Node BareMetal		03/17/2015 14:31:27 GMT-07	Complete
339	Advanced	admin	Single UCSM Server Configuration WF		03/17/2015 14:29:38 GMT-07	Complete
338	Advanced	admin	UCS CPA Multi-UCSM Hadoop Cluster WF		03/17/2015 14:28:54 GMT-07	Complete
337	Advanced	admin	UCS CPA Node BareMetal		03/17/2015 11:24:20 GMT-07	Complete
336	Advanced	admin	UCS CPA Node BareMetal		03/17/2015 11:24:20 GMT-07	Complete
335	Advanced	admin	UCS CPA Node BareMetal		03/17/2015 11:24:20 GMT-07	Complete
334	Advanced	admin	UCS CPA Node BareMetal		03/17/2015 11:24:19 GMT-07	Complete
333	Advanced	admin	Single UCSM Server Configuration WF		03/17/2015 11:22:37 GMT-07	Complete
332	Advanced	admin	UCS CPA Multi-UCSM Hadoop Cluster WF		03/17/2015 11:21:44 GMT-07	Complete

3. Double-click on one of the master workflows i.e. UCS CPA Multi-UCS Manager Hadoop Cluster to view the various steps undertaken to provision a Hadoop cluster.

Figure 283 Viewing a Completed Master Workflow

The screenshot shows a 'Service Request' window with the following details:

Service Request		Current status for the service request.	
Status			Refresh
▼ Overview			
Request ID	343	1	Initiated by admin 03/17/2015 23:35:30
Request Type	Advanced	2	Multi-UCSM Hadoop Cluster Profile 03/17/2015 23:35:53
Workflow Name	UCS CPA Multi-UCSM Hadoop Cluster WF	3	Setup Hadoop Cluster Env 03/17/2015 23:36:13
Workflow Version Label	0	4	Multi UCSM Configuration WF 03/17/2015 23:36:20
Request Time	03/17/2015 23:35:24 GMT-0700	5	Multi BareMetal WF Monitor 03/18/2015 00:25:04
Request Status	Complete	6	Synchronized Command Execution 03/18/2015 00:25:27
Comments		7	Custom SSH Command 03/18/2015 00:26:02
▼ Ownership		8	Provision Hadoop Cluster Completed action 03/18/2015 00:41:06
Initiating User	admin	9	Complete Completed successfully. 03/18/2015 00:41:09



Note If necessary click on the Log tab to view the logs generated during the provisioning of the Hadoop Cluster.

4. Double-click on one of the child workflows: i.e. UCS CPA Node Baremetal.

Figure 284 A Completed UCS CPA Node Baremetal workflow.

Workflow Status | Log | Objects Created and Modified | Input/Output

Service Request

Status Refresh

▼ Overview		Current status for the service request.	
Request ID	345	1	Initiated by admin 03/17/2015 23:38:05
Request Type	Advanced	2	Modify Workflow Priority (High) 03/17/2015 23:38:08
Workflow Name	UCS CPA Node BareMetal	3	Assign BareMetal SR ID 03/17/2015 23:38:11
Workflow Version Label	0	4	Create UCS Service Profile from template 03/17/2015 23:38:17
Request Time	03/17/2015 23:38:04 GMT-0700	5	Service Profile unbind/rebind Action 03/17/2015 23:39:21
Request Status	Complete	6	Modify UCS Service Profile Boot Policy 03/17/2015 23:40:23
Comments		7	Associate UCS Service Profile 03/17/2015 23:45:59
▼ Ownership		8	Assign ServerIdentity 03/17/2015 23:46:00
Initiating User	admin	9	Bind/Unbind vNIC Template 03/17/2015 23:46:09
		10	Bind/Unbind vNIC Template 03/17/2015 23:46:13
		11	Setup PXE Boot (OS Type: CentOSLive) 03/17/2015 23:46:38
		12	Setup RAID Commands 03/17/2015 23:46:50
▼ Overview		Current status for the service request.	
Request ID	345	13	UCS Blade Power ON Action 03/17/2015 23:47:34
Request Type	Advanced	14	Monitor PXE Boot 03/17/2015 23:53:16
Workflow Name	UCS CPA Node BareMetal	15	Monitor RAID Configuration 03/17/2015 23:53:17
Workflow Version Label	0	16	UCS Blade Power OFF Action 03/17/2015 23:53:31
Request Time	03/17/2015 23:38:04 GMT-0700	17	Setup PXE Boot (OS Type: RHEL6.5) 03/17/2015 23:53:54
Request Status	Complete	18	Setup RAID Commands 03/17/2015 23:53:57
Comments		19	UCS Blade Power ON Action 03/17/2015 23:57:17
▼ Ownership		20	Monitor PXE Boot 03/18/2015 00:04:19
Initiating User	admin	21	Modify UCS Service Profile Boot Policy Server has Local Disks 03/18/2015 00:04:20
		22	Service Profile unbind/rebind Action 03/18/2015 00:05:23
		23	UCS Blade Power ON Action 03/18/2015 00:11:08
		24	Assign IP Status 03/18/2015 00:11:08
Request Status	Complete	24	Assign IP Status 03/18/2015 00:11:08
Comments		25	Custom SSH Command 03/18/2015 00:16:37
▼ Ownership		26	Custom SSH Command 03/18/2015 00:17:10
Initiating User	admin	27	Synchronized Command Execution 03/18/2015 00:18:14
		28	UCS Blade Power OFF Action 03/18/2015 00:18:27
		29	UCS Blade Power ON Action 03/18/2015 00:19:40
		30	Synchronized Command Execution Completed action 03/18/2015 00:24:29
		31	Complete Completed successfully. 03/18/2015 00:24:32

Close

Host and Cluster Performance Monitoring

1. In the UCSD-Express web console, navigate to **Solutions > Big Data Accounts** for viewing the Hadoop cluster accounts.

Figure 285 *Big Data Accounts Summary Screen*

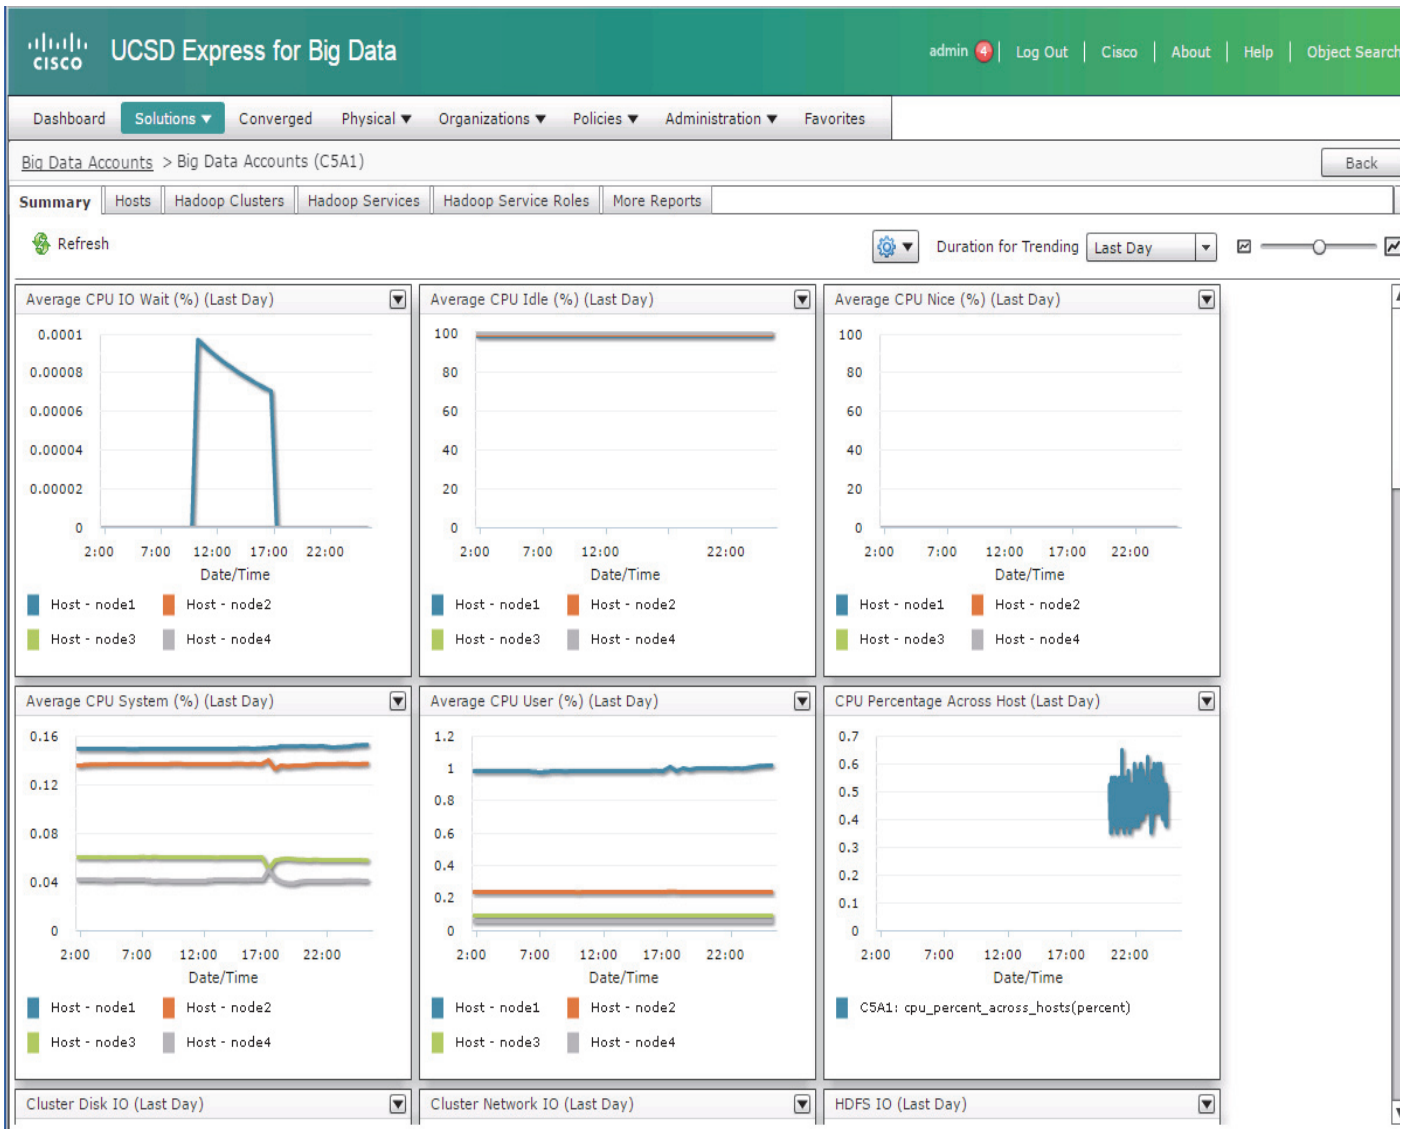
The screenshot shows the UCSD Express for Big Data web console. The main content area is titled 'Big Data Accounts' and contains a table with the following data:

Account Name	Account Type	Data Center	Management Console IP	Login
C5A1	Cloudera Derived Account	Default Pod	10.29.160.124	root
MapR1	MapR Derived Account	Default Pod	10.29.160.128	root
hw2	Hortonworks Derived Account	Default Pod	10.29.160.131	root

At the bottom of the table, it indicates 'Total 3 items'.

2. Double-click on one of the accounts to view the cluster-wide performance charts.

Figure 286 Hadoop Cluster Statistics



Cluster Management

1. In the UCSD-Express web console, navigate to **Solutions > Big Data Accounts** for viewing the Hadoop cluster accounts.
2. Double-Click on one of the accounts to drill into the cluster.
3. Click on the **Hosts** tab.

Figure 287 *Big Data Accounts – Viewing the List of Hosts of a Particular Hadoop Cluster*

The screenshot displays the Cisco UCS Express for Big Data interface. The top navigation bar includes the Cisco logo, the text "UCSD Express for Big Data", and user information "admin" with a "Log Out" link, along with "Cisco", "About", and "Help" links. Below the navigation bar is a menu with "Dashboard", "Solutions", "Converged", "Physical", "Organizations", "Policies", "Administration", and "Favorites". The main content area shows the "Big Data Accounts" section for "Big Data Accounts (C5A1)". A sub-menu includes "Summary", "Hosts", "Hadoop Clusters", "Hadoop Services", "Hadoop Service Roles", and "More Reports". A toolbar below the sub-menu contains icons for "Refresh", "Favorite", "Add Managed Node", "Add Live Node", "Add BareMetal Nodes", "View Details", "Delete Node", "Assign Rack", and "Recommission Node/Decommission Node". The "Hosts" section is active, displaying a table with the following data:

Host IP	Kernel Name	Host Name	Rack Name	Health	Server Identity	BareMetal Wf	Commission State
10.29.160.124	Linux	node1	/Default	Good	UCSM40;sys/rack-unit-5	334	Commissioned
10.29.160.125	Linux	node2	/Default	Good	UCSM40;sys/rack-unit-16	335	Commissioned
10.29.160.126	Linux	node3	/Default	Good	UCSM40;sys/rack-unit-10	336	Commissioned
10.29.160.127	Linux	node4	/Default	Good	UCSM40;sys/rack-unit-11	337	Commissioned

At the bottom of the table, it indicates "Total 4 items".

In this screen, the user can perform various management operations such as,

- Add one/more Baremetal nodes to the cluster.
 - Delete a node back to Baremetal
 - Decommission/Recommission
4. Click on the **Services** tab, where one could Start/Stop the Hadoop services.

Figure 288 Viewing the Services Provisioned in Specific Hadoop Cluster

The screenshot shows the Cisco UCS Express for Big Data interface. The top navigation bar includes the Cisco logo, the text 'UCSD Express for Big Data', and user information 'admin' with a red notification icon, along with links for 'Log Out', 'Cisco', 'About', 'Help', and 'Object Search'. Below the navigation bar is a menu with 'Dashboard', 'Solutions', 'Converged', 'Physical', 'Organizations', 'Policies', 'Administration', and 'Favorites'. The breadcrumb trail reads 'Big Data Accounts > Big Data Accounts (C5A1)' with a 'Back' button. The main content area is titled 'Hadoop Services' and includes a search bar and several action icons: Refresh, Favorite, Start All Services, and Stop All Services. A table displays the following data:

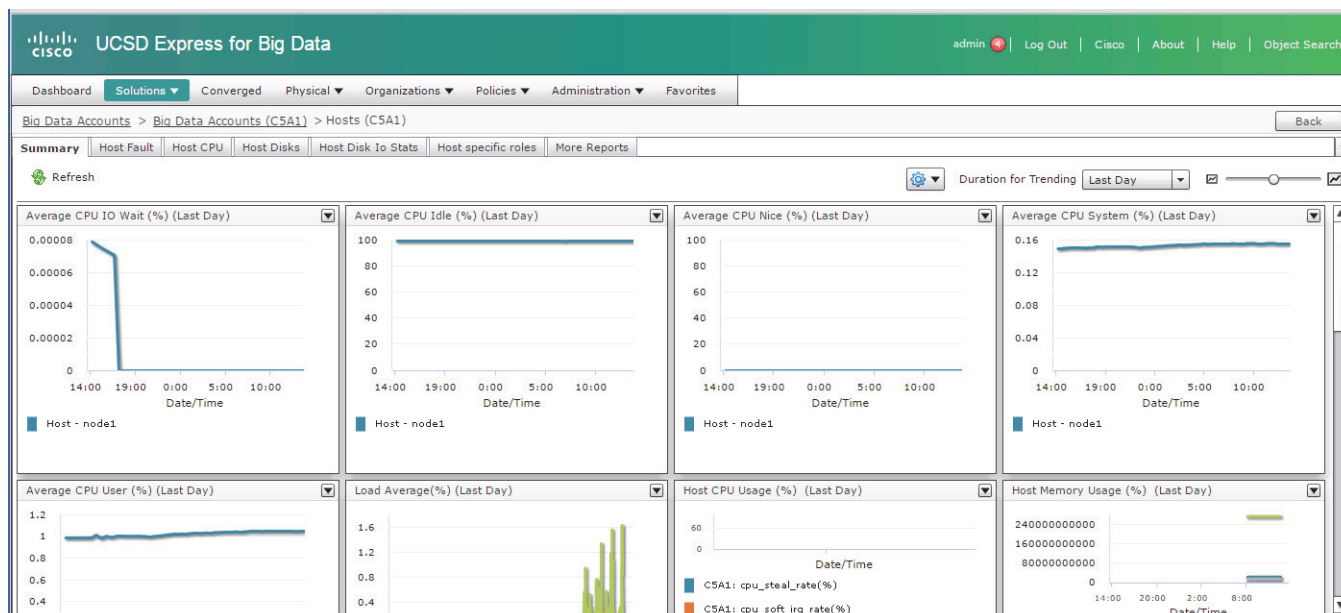
Status	Health	Service Type	Service Name
STARTED	GOOD	FLUME	flume
STARTED	GOOD	SQOOP	sqoop
STARTED	GOOD	KS_INDEXER	ks_indexer
STARTED	GOOD	HUE	hue
STARTED	GOOD	SENTRY	Sentry
STARTED	GOOD	ZOOKEEPER	zookeeper
STARTED	GOOD	OOZIE	oozie
STARTED	GOOD	IMPALA	impala
STARTED	GOOD	HDFS	hdfs
STARTED	GOOD	SOLR	solr
STARTED	GOOD	SPARK	spark
STARTED	GOOD	HBASE	hbase
STARTED	GOOD	YARN	yarn
STARTED	GOOD	HIVE	hive

At the bottom of the table, it indicates 'Total 14 items'.

Host level Monitoring

In the **Hosts** tab, double-click on one of the hosts to view the host's statistics.

Figure 289 Summary Statistics Screen of a Specific Host in a Hadoop Cluster



The user may monitor various resource utilization metrics of the particular host by clicking on the other tabs in this screen.

Reference

For details on managing the Hadoop clusters deployed on the Cisco UCS Integrated Infrastructure for Big Data, see the *Cisco UCS Director Express for Big Data Management Guide* at:

http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/ucs-director-express/management-guide/1-1/b_Management_Guide_for_Cisco_UCS_Director_Express_1_1.html

Bill of Materials

Table 23 provides the BOM for Cisco UCSD Big Data subscription licenses for up to 64 servers and Table 24 provides the BOM for the various Hadoop platforms.

Table 23 Bill of Material for UCSD for Big Data Subscription Licenses for up to 64 Servers

CUIC-SVR-OFFERS=	Cisco UCS Director Server Offerings	1
CON-SAU-SVROFFERS	Cisco UCS Director Server Offerings Software Application Sup	1
CUIC-BASE-K9	Cisco UCS Director Software License	1
CON-SAU-CUICBASE	SW APP SUPP + UPGR Cisco UCS Director Base Software	1
CUIC-TERM	Acceptance of Cisco UCS Director License Terms	1

Table 23 *Bill of Material for UCSD for Big Data Subscription Licenses for up to 64 Servers*

CUIC-EBDS-LIC=	UCSD Express for Big Data - Standard Edition (SE)	1
CUIC-EBDS-LIC	UCSD Express for Big Data - Standard Edition (SE)	64
CUIC-EBDS-S1-3YR	UCSD Express for Big Data - SE 3 year	64
CUIC-TERM	Acceptance of Cisco UCS Director License Terms	1

Table 24 *Bill of Material for Various Hadoop Platforms*

Part Number	Description
UCS-BD-CEBN=	CLOUDERA ENTERPRISE BASIC EDITION
UCS-BD-CEFNB=	CLOUDERA ENTERPRISE FLEX EDITION
UCS-BD-CEDN=	CLOUDERA ENTERPRISE DATA HUB EDITION
UCS-BD-HDP-ENT=	HORTONWORKS ENTERPRISE EDITION
UCS-BD-HDP-EPL=	HORTONWORKS ENTERPRISE PLUS EDITION
UCS-BD-M5-SL=	MapR M5 EDITION
UCS-BD-M7-SL=	MapR M7 EDITION